

Completeness of Iris-Based Program Logics

JOHANNES HOSTERT, ETH Zurich, Switzerland

ZICHEN ZHANG, New York University, USA

PUMING LIU, NYU Shanghai, China

SIMON ODDERSHEDE GREGERSEN, CISPA Helmholtz Center for Information Security, Germany

RALF JUNG, ETH Zurich, Switzerland

JOSEPH TASSAROTTI*, New York University, USA

Traditionally, proof systems such as program logics come with two core theorems: *soundness* and *completeness*. The role of soundness is obvious: we want to be sure that arguments carried out inside the logic actually lead to correct conclusions. Completeness complements that by ensuring that the logic does not limit expressivity: in principle, any correct result can be obtained within the confines of the logic. This result typically has to be stated *relative* to the completeness of the assertion logic that is used to reason about pre- and postconditions.

Over the past decade, the Iris framework has emerged as a widely used foundation for building separation logics. While Iris-based logics typically come with a soundness proof, none of them come with a proof of completeness. In fact, it is not entirely clear how to translate the typical recipe for relative completeness to Iris-based logics: these logics are higher-order, which means there is no separation of assertion logic and specification logic. Furthermore, the languages these logics reason about are also typically higher-order, enabling forms of recursion that go beyond a simple while loop (e.g., Landin’s Knot). In this paper, we present the first approach for establishing completeness of Iris-based program logics, and we show the generality of our methodology by applying it to a range of different logics described in prior work, including partial and total concurrent separation logics for a higher-order ML-like language, two quantitative logics (for bounding execution time and probabilistic errors), and a relational logic for proving refinement. All our results have been mechanized in the Rocq prover.

Additional Key Words and Phrases: Iris, Separation Logic, Completeness, Probability, Data Races

1 Introduction

Iris [33, 35, 38] is a higher-order separation logic framework that has been used for numerous program logics and verification projects. It provides reusable building blocks for constructing different variants of separation logic, with applications ranging from verification tools for C [47, 58], Rust [16, 48], OCaml [2, 24, 59], and Go [9], to reasoning about crash recovery [8], distributed execution [41, 60, 64], weak memory [13, 36, 50], resource consumption [53, 57], and probabilistic programs [1, 21, 22, 25, 26, 44], and semantic models of challenging type system features [15, 18–20, 23, 28, 32, 52, 65, 67]. All of these projects come with foundational soundness proofs, mechanized in the Rocq prover, that build on or adapt the core soundness proofs provided by the framework itself. These soundness results show that proofs of specifications about programs in the logic actually imply the intended properties about programs’ executions.

Although the many applications of Iris-based logics provide some empirical evidence that these logics are highly expressive, there are no formal results that characterize the expressivity of these logics. In particular, it is unknown whether these logics are *complete*, that is, whether any program whose execution satisfies certain properties can be provably shown to do so by applying the proof

*Also affiliated with Amazon Web Services. This paper does not reflect the views of Amazon Web Services.

Authors’ Contact Information: Johannes Hostert, johannes.hostert@inf.ethz.ch, ETH Zurich, Department of Computer Science, Zurich, Switzerland; Zichen Zhang, zichenzhang@nyu.edu, New York University, New York, New York, USA; Puming Liu, pl2559@nyu.edu, NYU Shanghai, Shanghai, China; Simon Oddershede Gregersen, s.gregersen@nyu.edu, CISPA Helmholtz Center for Information Security, Germany; Ralf Jung, ralf.jung@inf.ethz.ch, ETH Zurich, Department of Computer Science, Zurich, Switzerland; Joseph Tassarotti, jt4767@nyu.edu, New York University, New York, New York, USA.

50 rules offered by these logics. Resolving this question is important because it tells us whether a logic
51 is “missing” any rules, or whether the rules suffice to in principle verify any program.

52 For classical Hoare logic, completeness was long ago established by Cook [12]. This result
53 is “relative” to the completeness of the assertion logic that is used for the rule of consequence.
54 Subsequent works developed similar completeness results for logics for concurrency, such as the
55 method of Owicki and Gries [56], shown complete by Owicki [55], and rely-guarantee reasoning [68].
56 Most recently, de Boer and Hiep [14] have shown how to adapt this style of approach to concurrent
57 separation logic. However, these works consider languages and logics that are quite different
58 from Iris. In prior work on completeness of concurrent program logics, the setting is a first order
59 language, and the only source of looping behavior is a while-like loop construct. The logics are
60 also first order with a strict stratification between assertion logic and specification logic. In contrast,
61 Iris is often used to reason about higher-order languages that can encode Landin’s knot [42], and
62 the logic makes no distinction between assertions and specifications by supporting Hoare triples in
63 pre- and postconditions, a crucial feature for reasoning about higher-order programs. As a result,
64 existing results on relative completeness shed little light on the situation for logics like Iris.

65 This paper introduces a technique for proving completeness of Iris-based program logics. In
66 doing so, we follow the usual approach of Iris: we provide reusable building blocks that can often
67 be directly applied, and that can be further customized when needed. Concretely, we establish a
68 core reusable lemma for obtaining completeness proofs for instances of Iris’s language-agnostic
69 default program logic. This nicely complements the existing language-agnostic soundness proof
70 for said logic. We also demonstrate that the proof pattern behind this lemma generalizes to other
71 Iris-based program logics that are not a direct instance of this framework.

72 Using our methodology, we prove, for the first time, that a number of Iris-based logics are
73 complete (or can be made so by adding one or two missing proof rules). Our case studies include
74 partial and total higher-order concurrent separation logics, two quantitative logics (for bounding
75 execution time and probabilistic errors), and a relational logic for proving refinement. As expected,
76 these proofs make use of the suite of reasoning principles that Iris provides (ghost state, invariants,
77 and ownership). In particular, *Löb induction* [45], the fundamental primitive for recursive reasoning
78 in step-indexed logics such as Iris, is powerful enough to let us show completeness even for
79 languages with higher-order state. We demonstrate this with a series of case studies. All our results
80 have been mechanized in the Rocq prover [63] using the Iris Proof Mode [37, 39].

81 We begin this paper in §2 by building an Iris-style program logic for a sequential language with
82 higher-order state and proving it complete. In §3, we explain how Iris is used to reason about
83 concurrent programs, and we extend the completeness proof to that setting. We continue with
84 a series of completeness case studies: the logic for Iris’ default language HeapLang (§4) and its
85 total variant (§5); λ_{Rust} , which has a non-standard memory model to rule out data races (§6); a logic
86 with time credits for reasoning about execution cost (§7); a probabilistic logic for reasoning about
87 error bounds (§8); and finally a relational logic for proving refinement of concurrent higher-order
88 programs (§9). We finish the technical contribution of the paper with a discussion of a general
89 semantic condition that all complete Iris logics have to satisfy (§10). We then survey prior proofs of
90 completeness for other program logics (§11) and conclude with a discussion about the significance
91 of completeness (§12).

92 2 Warm-up: Completeness for a Sequential Language

93 In this section, we slowly build up to our main completeness result by first considering an Iris-style
94 program logic for the simpler, sequential setting. Concretely, we define a language SeqLang, a call-
95 by-value lambda calculus with integers and mutable higher-order state. The syntax and semantics
96 can be found in Figure 1. We define a small-step operational semantics with the reduction relation
97
98

99	$v, w \in Val ::= z \mid \ell \mid \lambda x. e \mid ()$	$(z \in \mathbb{Z}, \ell \in Loc)$	
100	$e \in Expr ::= v \mid x \mid \lambda x. e \mid e_1 e_2 \mid e_1 \odot_2 e_2 \mid \mathbf{if} \ e \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2 \mid$	$(\odot_2 \in \{+, -, =, \dots\})$	
101	$\mathbf{ref}(e_1) \mid \mathbf{free}(e) \mid !e \mid e_1 \leftarrow e_2$		
102	$K \in Ctx ::= \bullet \mid e \ K \mid K \ v \mid e \odot_2 \ K \mid K \odot_2 \ v \mid \mathbf{if} \ K \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2 \mid$		
103	$\mathbf{ref}(K) \mid \mathbf{free}(K) \mid !K \mid e \leftarrow K \mid K \leftarrow v$		
104	$\sigma \in State \triangleq Loc \xrightarrow{\text{fin}} Val, \quad \rho \in Conf \triangleq Expr \times State$		
105			
106	$\frac{\text{PUREOP} \quad z_1 \odot_2 z_2 = z_3}{z_1 \odot_2 z_2 \rightarrow_{\text{pure}} z_3} \quad \dots \quad \frac{\text{BASEPURE} \quad e_1 \rightarrow_{\text{pure}} e_2}{(e_1, \sigma) \rightarrow_{\text{base}} (e_2, \sigma)}$		
107	$\frac{\text{PUREBETA} \quad (\lambda x. e) v \rightarrow_{\text{pure}} e[v/x]}{\dots} \quad \dots \quad \frac{\text{BASEALLOCC} \quad \ell \notin \text{dom}(\sigma)}{(\mathbf{ref}(v), \sigma) \rightarrow_{\text{base}} (\ell, \sigma[\ell \leftarrow v])}$	$\frac{\text{BASEFREE} \quad \ell \in \text{dom}(\sigma)}{(\mathbf{free}(\ell), \sigma) \rightarrow_{\text{base}} ((), \sigma \setminus \{\ell\})}$	
108	$\frac{\text{BASELOAD} \quad \sigma(\ell) = v}{(!\ell, \sigma) \rightarrow_{\text{base}} (v, \sigma)}$	$\frac{\text{BASESTORE} \quad \ell \in \text{dom}(\sigma)}{(\ell \leftarrow v, \sigma) \rightarrow_{\text{base}} ((), \sigma[\ell \leftarrow v])}$	$\frac{\text{STEPCTX} \quad (e_1, \sigma) \rightarrow_{\text{base}} (e_2, \sigma')}{(K[e_1], \sigma) \rightarrow (K[e_2], \sigma')}$
109			
110			
111			
112			
113			
114			
115			
116			
117			
118			

Fig. 1. Syntax and Semantics of SeqLang.

\rightarrow using evaluation contexts K and a *base* reduction relation $\rightarrow_{\text{base}}$. The definition of evaluation contexts induces a right-to-left evaluation order.

The step rules operate on configurations ρ , which are pairs consisting of an expression e and a heap σ . The heap is represented as a finite partial map from locations to values. The reduction of the pure, deterministic fragment of the language is described by the *pure* reduction relation $\rightarrow_{\text{pure}}$, which can (using **BASEPURE**) be turned into a base reduction step that leaves the heap untouched. Note that **Figure 1** omits some of the pure reductions for brevity. The rules for the load and store operations are standard. Deallocation removes the heap cell (which can then be re-used by a future allocation). For all of these operations, the reduction relation *gets stuck* if the location is not allocated. Allocation is the only non-deterministic operation in our language since it picks an arbitrary fresh location $\ell \notin \text{dom}(\sigma)$.

2.1 A Sequential Separation Logic

Figure 2 defines a separation logic for SeqLang. This separation logic can be used to derive judgments of the shape $P \vdash Q$, where P and Q are separation logic assertions. Separation logic enriches propositional logic with the separating conjunction $*$ and separating implication \multimap , sometimes called the “magic wand”. The separating conjunction $P * Q$ expresses ownership of P and Q *separately*, i.e., P and Q hold for disjoint parts of the heap. Thus, in general, $P \not\vdash P * P$. Most often, Iris-based separation logics are *affine* logics, meaning they allow the weakening rule $P * Q \vdash P$. The separating implication \multimap is a version of implication that interacts with separating conjunction the same way that regular implication interacts with conjunction: $P \vdash Q \multimap R$ iff $P * Q \vdash R$. We do not give a full introduction to separation logic here but refer to Chapter 3 of Birkedal and Bizjak [5] for an expository introduction.

148 The assertion $\ulcorner \varphi \urcorner$ embeds an arbitrary meta-level proposition φ (i.e., a Rocq Prop) into the syntax
 149 of the logic. In particular, if $\varphi \Rightarrow \varphi'$ holds at the meta-level then $\ulcorner \varphi \urcorner \vdash \ulcorner \varphi' \urcorner$. This means that the
 150 proof rules of the logic are not purely syntactic.

151 To turn our logic into a program logic, we first introduce the points-to connective $\ell \mapsto v$. This
 152 assertion expresses that the current heap stores the value v at location ℓ . Furthermore, this fact is
 153 *exclusive*, meaning it is not possible to own two points-to connectives for the same piece of state.
 154 This is also witnessed by the rule `POINTSToNE`.

155 We reason about program executions using $\text{wp } e \{ \Phi \}$. Like all separation logic assertions, this is
 156 an assertion about the current program state: wp holds in some program state if executing e from
 157 that state executes safely, and if furthermore whenever e terminates with a value v , then $\Phi(v)$ holds
 158 for the final state. This property expresses a form of partial correctness as it does not guarantee
 159 termination. We discuss the definition of wp in §10, but for now we treat it as an abstract predicate
 160 that satisfies the rules in Figure 2. The later modality \triangleright can be ignored for now.

161 The structural rule `WPWAND` can be used to prove both monotonicity of the postcondition as well
 162 as the *frame rule*, the hallmark of separation logic:¹

$$163 \quad \text{wp } e \{ \Phi \} * P \vdash \text{wp } e \{ v. \Phi(v) * P \}$$

164 This rule is the key ingredient to *modular* reasoning about programs: If we have proven a specifica-
 165 tion $\text{wp } e \{ \Phi \}$ about the program e , then any assertion P on some unrelated state (that is disjoint
 166 from the state used in the verification) will not affect the execution. As such, P can be passed along
 167 to the postcondition to be used unchanged once e has finished executing.

168 The rule `WPVALUE` is the “base case” of the wp , since we just plug the value into the postcondition.
 169 The rule `WPBIND` can be used to focus on a sub-expression in an evaluation context. All the
 170 remaining reasoning rules correspond to base reduction steps, with the rule `WPPURE` re-using the
 171 pure reduction relation from Figure 1. The rules for the heap-manipulating steps are written in a
 172 “predicate transformer style” with Φ serving as a form of continuation. For example, the rule `WPSTORE`
 173 says that to reason about a store of w to location ℓ , we must own $\ell \mapsto v$ for some value v . In addition,
 174 we must show that the post-condition Φ would follow from having $\ell \mapsto w$, the updated points-to
 175 that reflects the result of the store. From this rule, it is straightforward to derive an alternate “direct
 176 style” version that looks like a typical Hoare triple: $\ell \mapsto v \vdash \text{wp } \ell \leftarrow w \{ v. \ulcorner v = () \urcorner * \ell \mapsto w \}$.

177 There is no rule in Figure 2 for reasoning about loops as SeqLang does not have a primitive loop
 178 or recursion mechanism. However, since it is a higher-order language and supports higher-order
 179 state, recursion can be encoded in multiple ways. We postpone the discussion of how to reason
 180 about recursion until §2.3.

181 Before we consider the question of completeness, we verify two examples to demonstrate how
 182 the rules are to be used.

183 *Example 1.* We prove that the program $40 + (1 + 1)$ evaluates to the value 42. We start with the
 184 goal at the top and apply the rules one-by-one, noting what is left to show. This is how the rules are
 185 intended to be used, and the resulting order of proof steps follows the program’s execution order.

194
 195 ¹The notation $\text{wp } e \{ \Phi \}$ is sugar for $\text{wp } e \{ v. \Phi(v) \}$.

197	$P, Q \in iProp ::= \ulcorner \varphi \urcorner \mid P \wedge Q \mid P \vee Q \mid \forall x. P \mid \exists x. P \mid \dots$		(φ a meta-level proposition)
198	$\mid P * Q \mid P * \! * Q \mid \ell \mapsto v \mid \text{wp } e \{v. P\} \mid \triangleright P$		(Φ, Ψ predicates $Val \rightarrow iProp$)
199	<hr/>		
200	<hr/>		
201	POINTSToNE	WPWAND	
202	$\ell_1 \mapsto v_1 * \ell_2 \mapsto v_2 \vdash \ulcorner \ell_1 \neq \ell_2 \urcorner$	$\text{wp } e \{\Phi\} * (\forall x. \Phi(x) * \Psi(x)) \vdash \text{wp } e \{\Psi\}$	
203	<hr/>		
204	WPVALUE		
205	$\Phi(v) \dashv\vdash \text{wp } v \{\Phi\}$	WPBIND	WPPURE
206	$\text{wp } e \{v. \text{wp } K[v] \{\Phi\}\} \dashv\vdash \text{wp } K[e] \{\Phi\}$	$\frac{e_1 \rightarrow \text{pure } e_2}{\triangleright \text{wp } e_2 \{\Phi\} \vdash \text{wp } e_1 \{\Phi\}}$	
207	<hr/>		
208	WPFREE	WPALLOC	
209	$\ell \mapsto v * \triangleright \Phi(()) \vdash \text{wp } \mathbf{free}(\ell) \{\Phi\}$	$\triangleright \forall \ell. \ell \mapsto v * \Phi(\ell) \vdash \text{wp } \mathbf{ref}(v) \{\Phi\}$	
210	<hr/>		
211	WPLoad	WPSTORE	
212	$\ell \mapsto v * \triangleright (\ell \mapsto v * \Phi(v)) \vdash \text{wp } !\ell \{\Phi\}$	$\ell \mapsto v * \triangleright (\ell \mapsto w * \Phi(())) \vdash \text{wp } \ell \leftarrow w \{\Phi\}$	

Fig. 2. Separation Logic for SeqLang.

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

233

234

235

236

237

238

239

240

241

242

243

244

245

217	$\vdash \text{wp } 40 + (1 + 1) \{v. \ulcorner v = 42 \urcorner\}$		
218	$\Leftarrow \vdash \text{wp } 1 + 1 \{w. \text{wp } 40 + w \{v. \ulcorner v = 42 \urcorner\}\}$	by WPBIND with $K := 40 + \bullet$	
219	$\Leftarrow \vdash \text{wp } 2 \{w. \text{wp } 40 + w \{v. \ulcorner v = 42 \urcorner\}\}$	by WPPURE	
220	$\Leftarrow \vdash \text{wp } 40 + 2 \{v. \ulcorner v = 42 \urcorner\}$	by WPVALUE	
221	$\Leftarrow \vdash \text{wp } 42 \{v. \ulcorner v = 42 \urcorner\}$	by WPPURE	
222	$\Leftarrow \vdash \ulcorner 42 = 42 \urcorner$	by WPVALUE	

Example 2. We prove that the program $(\lambda x. !x) (\mathbf{ref}(42))$ evaluates to the value 42 as well. This program involves mutable state, so a formal proof involves manipulating separation logic assertions; we omit those details in the proof outline shown below.

229	$\vdash \text{wp } (\lambda x. !x) (\mathbf{ref}(42)) \{v. \ulcorner v = 42 \urcorner\}$		
230	$\Leftarrow \vdash \text{wp } \mathbf{ref}(42) \{w. \text{wp } (\lambda x. !x) w \{v. \ulcorner v = 42 \urcorner\}\}$	by WPBIND	
231	$\Leftarrow \ell \mapsto 42 \vdash \text{wp } \ell \{w. \text{wp } (\lambda x. !x) w \{v. \ulcorner v = 42 \urcorner\}\}$	by WPALLOC	
232	$\Leftarrow \ell \mapsto 42 \vdash \text{wp } (\lambda x. !x) \ell \{v. \ulcorner v = 42 \urcorner\}$	by WPVALUE	
233	$\Leftarrow \ell \mapsto 42 \vdash \text{wp } !\ell \{v. \ulcorner v = 42 \urcorner\}$	by WPPURE	
234	$\Leftarrow \ell \mapsto 42 \vdash \ulcorner 42 = 42 \urcorner$	by WPLoad	

Soundness. We have seen how the logic can be used to verify simple programs, and we have discussed that the wp encodes a partial correctness criterion. However, we still need to show that the logic is sound: proving a wp about a program should establish partial correctness of the program at the meta-level. To state this formally, we first define a meta-level predicate $\text{safe}_\varphi(e, \sigma)$:

$$\begin{aligned} \text{red}(e, \sigma) &\triangleq \exists e', \sigma'. (e, \sigma) \rightarrow (e', \sigma') \\ \text{safe}_\varphi(e, \sigma) &\triangleq \forall e', \sigma'. (e, \sigma) \rightarrow^* (e', \sigma') \implies (\exists v. e' = v \wedge \varphi(v)) \vee \text{red}(e', \sigma') \end{aligned}$$

The predicate $\text{safe}_\varphi(e, \sigma)$ says that executing e never gets stuck, and if it terminates in a value v then the value satisfies the predicate φ .

The following *soundness*² theorem connects the wp to *safe*.

Theorem 3 (Soundness). *If $\vdash \text{wp } e \{v. \ulcorner \varphi(v) \urcorner\}$, then $\text{safe}_\varphi(e, \sigma)$ for all σ .*

We do not prove soundness here and instead continue with completeness. Completeness is the converse of soundness, where we go from a safety proof at the meta-level to a proof of a corresponding wp:

Theorem 4 (Completeness). *If $\text{safe}_\varphi(e, \sigma)$ for all σ , then $\vdash \text{wp } e \{v. \ulcorner \varphi(v) \urcorner\}$.*

The goal for the remainder of this section is to prove this theorem.

2.2 Sequential Completeness for Terminating Programs

Before proving [Theorem 4](#), we will warm up by proving a weaker version that makes the additional assumption that e is strongly normalizing under \rightarrow (written SN_\rightarrow), *i.e.*, terminating when run from any state. Additionally, as our proof is going to proceed by induction, it will be helpful to generalize the statement to get a better induction hypothesis. Thus, we will prove the following

Theorem 5 (Completeness from Strong Normalization). *If $\text{safe}_\varphi(e, \sigma)$ and $\text{SN}_\rightarrow(e, \sigma)$, then $\bigstar_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v \vdash \text{wp } e \{v. \ulcorner \varphi(v) \urcorner\}$.*

In this theorem, we use an iterated separating conjunction to get points-to assertions for all the locations in σ and have to derive the wp. The proof will need two lemmas about *safe*.

Lemma 6. *If $\text{safe}_\varphi(e, \sigma)$, then either e is a value satisfying φ , or (e, σ) is reducible.*

Lemma 7. *If $\text{safe}_\varphi(e, \sigma)$ and $(e, \sigma) \rightarrow (e_2, \sigma_2)$, then $\text{safe}_\varphi(e_2, \sigma_2)$.*

PROOF OF THEOREM 5. The proof is by induction on the assumption that (e, σ) is strongly normalizing. This means that we get to assume the induction hypothesis for all immediate successors (e', σ') of (e, σ) . Formally, we have the following:

$$\forall e', \sigma'. (e, \sigma) \rightarrow (e', \sigma') \Rightarrow \text{safe}_\varphi(e', \sigma') \Rightarrow \bigstar_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v \vdash \text{wp } e' \{w. \ulcorner \varphi(w) \urcorner\} \quad (\text{IH})$$

and we are left to prove $\text{safe}_\varphi(e, \sigma) \Rightarrow \bigstar_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v \vdash \text{wp } e \{w. \ulcorner \varphi(w) \urcorner\}$ under this induction hypothesis.³ Now, we just need to reason about (at least) one reduction step of e and then apply the induction hypothesis.

We first apply [Lemma 6](#). If e is a value, we apply [WPVALUE](#) which concludes the proof. In the remaining case, since we know our expression is reducible, we know that $e = K[e_1]$ for some K, e_1 and that $(e_1, \sigma) \rightarrow_{\text{base}} (e_2, \sigma')$ for some e_2, σ' . Thus, using [WPBIND](#), we turn our goal into $\text{wp } e_1 \{v. \text{wp } K[v] \{w. \ulcorner \varphi(w) \urcorner\}\}$. We continue with a case distinction on the base reduction.

Consider the case [BASEPURE](#) where $e_1 \rightarrow_{\text{pure}} e_2$. We apply [WPPURE](#) to turn our goal into $\text{wp } e_2 \{v. \text{wp } K[v] \{w. \ulcorner \varphi(w) \urcorner\}\}$. Applying [WPBIND](#) “backwards” turns our goal into $\text{wp } K[e_2] \{w. \ulcorner \varphi(w) \urcorner\}$. Since $(e, \sigma) \rightarrow (K[e_2], \sigma)$, we instantiate our induction hypothesis with $(K[e_2], \sigma)$. By [Lemma 7](#), $(K[e_2], \sigma)$ is safe. Also, the state has not changed so we have the needed points-tos. Thus all premises of the induction hypothesis hold, finishing the case.

The four heap-dependent cases remain. In the [BASELOAD](#) case, we get that $e_1 = !\ell$, and that $\sigma(\ell) = v$. First, we remove the points-to for ℓ from the iterated separating conjunction over σ . Next, we use

²In Iris, this theorem is usually called *adequacy* for reasons we discuss later in [§10](#).

³As usual with this kind of induction on a well-founded reduction relation, there is no separate base case to consider.

295 it to instantiate and apply **Wpload**; the points-to for ℓ is returned in the postcondition. Finally, we
 296 re-establish the iterated separating conjunction for σ and finish the case by applying the induction
 297 hypothesis. For **BASESTORE**, we do the same but with **WpSTORE**. Of course, the points-to now stores a
 298 new value, but σ is altered in lockstep, so we can still re-establish the big separating conjunction
 299 and apply the induction hypothesis. For **BASEFREE**, we also remove the points-to for ℓ from the
 300 iterated separating conjunction, but **WpFREE** does not return the points-to in the postcondition.
 301 Since it is also deleted from σ , we can still establish the iterated separating conjunction for the new
 302 state before applying the induction hypothesis.

303 Only the case **BASEALLOC** remains. Here, $e_1 = \mathbf{ref}(v)$ and $e_2 = \ell_1$ for some $\ell_1 \notin \text{dom}(\sigma)$. When
 304 applying **WpALLOC**, we get a (fresh) non-deterministic location ℓ_2 along with a points-to, which
 305 might be different from ℓ_1 . Since we own a points-to for all locations in σ , we use **POINTSTONE**
 306 to establish that $\ell_2 \notin \text{dom}(\sigma)$. Since $\ell_2 \notin \text{dom}(\sigma)$ and our allocator is non-deterministic, we have that
 307 $(e_1, \sigma) \rightarrow_{\mathbf{base}} (\ell_2, \sigma[\ell_2 \leftarrow v])$ is a valid reduction step. Thus, we apply the induction hypothesis to
 308 $(K[\ell_2], \sigma[\ell_2 \leftarrow v])$ and finish the proof. \square
 309

310 Taking a step back, the structure of this proof exploits the fact that either e is already a value, and
 311 hence satisfies φ by assumption, or else e is reducible, in which case we perform a case distinction on
 312 the reduction rule, using **WpBIND** along the way. After applying a proof rule for the corresponding
 313 reduction rule, we use **WpBIND** in the opposite direction to return to something that the induction
 314 hypothesis can be applied to. In the following section, we introduce an alternative induction
 315 principle that allows us to conduct a similar proof but *without* the strong normalization assumption.
 316

317 2.3 Sequential Completeness for Partial Programs

318 Before we attempt to eliminate the strong normalization assumption from our proof, we explain
 319 how Iris-based logics typically allow us to reason about non-terminating programs. Consider the
 320 program $\Omega \triangleq (\lambda x. x x) (\lambda x. x x)$. It is easy to see that $\Omega \rightarrow_{\mathbf{pure}} \Omega$. Since our wp is intended to cover
 321 partial correctness, we should be able to prove the specification $\text{wp } \Omega \{v. \text{False}\}$. But how?

322 The core reasoning principle for recursion in typical Iris-based logics is *Löb induction*, which
 323 surfaces the underlying step-indexed nature of Iris [62, 66]. In the logic, this is exposed by means
 324 of the so-called *later modality* [3, 6, 54], written $\triangleright P$. Intuitively, the assertion $\triangleright P$ says that P is only
 325 “approximately” true now, but will be actually true after one step of computation.

326 The connection between program steps and the \triangleright modality can be seen in the rules for the
 327 wp in Figure 2. For example, in the rule **WpSTORE** the assumption to the left of the turnstile is
 328 $\ell \mapsto v * \triangleright(\ell \mapsto w * \Phi())$. The left conjunct requires us to own the points-to for ℓ containing
 329 some value v before the store occurs. The right conjunct says that—one step later—after the store has
 330 happened, we get back $\ell \mapsto w$ (reflecting the updated value) and have to establish the postcondition.

331 The later modality commutes with most logical con-
 332 nectives, except for the magic wand and the existential
 333 quantifier.⁴ There are three key rules for working with
 334 the later modality: **LATERINTRO**, **LATERMONO**, and **LÖB**.

335 Intuitively, **LATERINTRO** tells us that if something is true now, it remains true one step later. (In
 336 particular, this means the wp rules with \triangleright imply the ones without \triangleright .) The rule **LATERMONO** is the
 337 main elimination rule for the later modality: we can eliminate a later modality in our assumptions
 338 if we can remove a later modality from the conclusion as well.

339 The critical rule for reasoning about loops is **LÖB** induction, which reifies induction on the step
 340 index: To prove P , it suffices to prove P under the induction hypothesis $\triangleright P$. This allows a coinductive
 341

$$\begin{array}{c}
 \text{LATERINTRO} \\
 \frac{P \vdash \triangleright P}{\triangleright P \vdash P} \\
 \\
 \text{LATERMONO} \quad \text{LÖB} \\
 \frac{P \vdash Q \quad \triangleright P \vdash P}{\triangleright P \vdash \triangleright Q} \quad \frac{\triangleright P \vdash P}{\vdash P}
 \end{array}$$

342 ⁴It commutes with the existential quantifier if the domain of quantification is non-empty.
 343

style of reasoning: we can use the induction hypothesis $\triangleright P$ only after we made “progress” by taking a step in the program, which lets us apply **LATERMONO** to “unlock” the induction hypothesis.

Example 8. We demonstrate how Löb induction allows us to verify the program Ω :

$$\begin{array}{lll}
 & \vdash \text{wp } \Omega \{v. \text{False}\} & \\
 \Leftarrow & \triangleright \text{wp } \Omega \{v. \text{False}\} \vdash \text{wp } \Omega \{v. \text{False}\} & \text{by } \mathbf{LÖB} \\
 \Leftarrow & \Omega \rightarrow_{\text{pure}} \Omega & \text{by } \mathbf{WPPURE}
 \end{array}$$

Here, applying **LÖB** allows us to immediately conclude using **WPPURE**. Standard rules for partial-correctness reasoning about recursive functions or while loops can be derived via Löb induction. The power of Löb induction is that it is not tied to a particular program construct but works for all propositions, and the induction hypothesis can have an arbitrary shape. This is useful for proving our completeness theorem.

Theorem 9 (Completeness of SeqLang). *If $\text{safe}_\varphi(e, \sigma)$, then $\bigstar_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v \vdash \text{wp } e \{v. \ulcorner \varphi(v) \urcorner\}$.*

The proof proceeds similarly to the proof of **Theorem 5**, but instead of doing induction on the fact that (e, σ) is strongly normalizing, we do Löb induction.⁵ Instead of the requirement to exhibit an explicit step with the \rightarrow relation, the induction hypothesis is now guarded by a later. But this is equivalent, since the later modality is the logic’s way of requiring us to take a step! Indeed, the old proof continues to work. In all cases, we apply one of the wp rules which introduces a later modality in the goal. By applying **LATERMONO**, we “unlock” the induction hypothesis and use it to conclude the proof.

3 Completeness for a Concurrent Language

We have seen how to prove completeness for an Iris-based program logic for a sequential language. In this section, we show how that approach generalizes to a concurrent language, **ConcLang**.

3.1 A Concurrent Separation Logic

We begin by extending the grammar, reduction rules, and program logic with support for concurrency. These additions are outlined in **Figure 3**. On the language side, we add two new primitives. The first is **fork**, which implements unstructured concurrency. The second is an atomic compare-and-swap operation **CAS**, which allows us to write interesting concurrent programs. Threads are modeled using a thread-pool semantics, defined by the relation \rightarrow_{tp} , which works on configurations ρ now storing a list of threads (expressions) instead of just a single one. We will treat this list as isomorphic to a finite partial function $\mathbb{N} \xrightarrow{\text{fin}} \text{Expr}$ mapping n to the n -th element in the list. To create new threads, the relations \rightarrow and $\rightarrow_{\text{base}}$ have a new parameter \vec{e}_f , which is a list of forked-off expressions. This list is empty in all rules except **TPFORK**, where it is a singleton list since **fork** creates one new thread.

On the program logic side, the important new rule is **WPFORK**, which allows us to spawn a new thread. Applying this rule lets us reason modularly about concurrent threads by splitting ownership of resources between the current thread and the new thread. For example, if we are trying to establish $\text{wp}(\mathbf{fork } e_1; e_2) \{ \Phi \}$, then by applying **WPBIND** and **WPFORK**, we are left with a goal that requires us to prove $\text{wp } e_2 \{v. \text{True}\} * \text{wp } e_1 \{ \Phi \}$ (ignoring the \models and the mask \mathcal{E} for now). That is, we must prove two *separate* wps, one for each thread. This is the key power of concurrent separation logic! We can verify threads in isolation, as long as we can separate the assumptions needed to verify each.

⁵This requires that we move the assumption $\text{safe}_\varphi(e, \sigma)$ into the logic as the assertion $\ulcorner \text{safe}_\varphi(e, \sigma) \urcorner$.

393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441

$$\begin{array}{l}
e \in \text{Expr} ::= \dots \mid \mathbf{fork} \ e \mid \mathbf{CAS}(e_1, e_2, e_3), \quad \vec{e} \in \mathcal{L}(\text{Expr}), \quad \rho \in \text{Conf} \triangleq \mathcal{L}(\text{Expr}) \times \text{State} \\
P, Q \in \text{iProp} ::= \dots \mid \text{wp}_{\mathcal{E}} \ e \{v. \Phi(v)\} \mid \Vdash_{\mathcal{E}} P \mid \boxed{P}^{\mathcal{N}} \mid \bullet^{\mathcal{Y}} m \mid k \hookrightarrow^{\mathcal{Y}} v \\
\text{TPSTEP} \\
\frac{(T(n), \sigma) \rightarrow (e_2, \sigma, \vec{e}_f)}{(T, \sigma) \rightarrow_{\text{tp}} (T[n \leftarrow e_2] \uparrow \vec{e}_f, \sigma)} \qquad \text{TPFORK} \\
(\mathbf{fork} \ e, \sigma) \rightarrow_{\text{base}} ((\), \sigma, [e]) \\
\text{INVALLOC} \qquad \text{UPDATEINTRO} \qquad \text{UPDATETRANS} \qquad \text{UPDATEFRAME} \\
P \vdash \Vdash_{\mathcal{E}} \boxed{P}^{\mathcal{N}} \qquad P \vdash \Vdash_{\mathcal{E}} P \qquad \Vdash_{\mathcal{E}} \Vdash_{\mathcal{E}} P \vdash \Vdash_{\mathcal{E}} P \qquad (\Vdash_{\mathcal{E}} P) * Q \vdash \Vdash_{\mathcal{E}} (P * Q) \\
\text{UPDATEINV} \qquad \text{WPATOMICINV} \\
\frac{\mathcal{N} \subseteq \mathcal{E} \quad P * Q \vdash \Vdash_{\mathcal{E} \setminus \mathcal{N}} P * R}{\boxed{P}^{\mathcal{N}} * Q \vdash \Vdash_{\mathcal{E}} R} \qquad \frac{\mathcal{N} \subseteq \mathcal{E} \quad \text{Atomic } e \quad P \vdash \text{wp}_{\mathcal{E} \setminus \mathcal{N}} \ e \{v. P * \Phi(v)\}}{\boxed{P}^{\mathcal{N}} \vdash \text{wp}_{\mathcal{E}} \ e \{\Phi\}} \\
\text{WPUPDATEELIM} \qquad \text{WPWAND} \\
\frac{P \vdash \text{wp}_{\mathcal{E}} \ e \{\Phi\}}{\Vdash_{\mathcal{E}} P \vdash \text{wp}_{\mathcal{E}} \ e \{\Phi\}} \qquad \text{wp}_{\mathcal{E}} \ e \{\Phi\} * (\forall x. \Phi(x) \multimap \Vdash_{\mathcal{E}} \Psi(x)) \vdash \text{wp}_{\mathcal{E}} \ e \{\Psi\} \\
\text{WPVALUE} \qquad \text{WPFORK} \\
\frac{}{\Vdash_{\mathcal{E}} \Phi(v) \uparrow \text{wp}_{\mathcal{E}} \ v \{\Phi\}} \qquad \triangleright \Vdash_{\mathcal{E}} (\Phi((\)) * \text{wp}_{\top} \ e \{v. \text{True}\}) \vdash \text{wp}_{\mathcal{E}} \ \mathbf{fork} \ e \{\Phi\}
\end{array}$$

Fig. 3. Syntax, Semantics, and Program Logic for ConcLang.

But this separation means that if we have a points-to assertion $\ell \mapsto v$, we can only give it to one of the two threads. That suffices if the two threads access disjoint pieces of state, but if they need to access shared state, then we need other features in the logic: *invariants* and *ghost state*. Readers familiar with these features of Iris can skip to §3.2.

Invariants and Masks. If we want two threads to operate on the same heap cell, we need to somehow give both of them access to it. Since the executions of both threads interleave, we can no longer reason locally in each thread, and instead need to describe the *protocol* under which both threads cooperate. In Iris, this is done with the invariant assertions $\boxed{P}^{\mathcal{N}}$, which says that P is an invariant of program execution (it holds before and after each step), and this invariant has been allocated in namespace \mathcal{N} . Invariants are freely duplicable and thus shareable among several threads. There is no exclusive ownership attached to an invariant assertion; all it conveys is *knowledge* that the invariant is maintained. For the purposes of this paper,⁶ a thread can use $\boxed{P}^{\mathcal{N}}$ to get access to P for *one atomic step* so long as it shows that P holds again after that step is completed. The *mask* \mathcal{E} on the wp is a set of names of invariants that are currently *enabled* and can be opened; this mechanism is used to prevent a thread from opening the same invariant twice. The process of opening invariants is captured by the rule **WPATOMICINV**. After opening an invariant, the invariant’s namespace is removed from the mask. Additionally, we have the side condition *Atomic* e which says that e is *atomic* in the sense that it reduces to a value in a single reduction step.

⁶For soundness reasons, invariants only give access to $\triangleright P$. In this paper, our invariants are always *timeless*, so we can eliminate this later modality. See Jung et al. [33] for the details and a definition of timelessness.

442 *Example 10.* Let us now verify the following specification about a program which shares a
 443 location between two threads.⁷

$$444 \quad \text{wp}_\top (\mathbf{let } x = \mathbf{ref}(0) \mathbf{in fork } (x \leftarrow 1); !x) \{v. \ulcorner v \in \{0, 1\} \urcorner\}$$

445
 446 The proof proceeds by first allocating the location ℓ_x . Next, we can allocate an invariant using
 447 `INVALLOC` (keep ignoring the \Rightarrow for now). The invariant we allocate is $\boxed{\exists v \in \{0, 1\}. \ell_x \mapsto v}^{\mathcal{N}}$
 448 for some namespace \mathcal{N} we choose. The choice does not matter in this example. Since invariant
 449 assertions are duplicable, we can pass it to both threads when we apply `WPFORK`. Both threads will
 450 now proceed to access the invariant. For the forked-off thread, we use `WPATOMICINV` and since a store
 451 operation is atomic, we get access to the points-to for the duration of the operation. Afterwards, we
 452 restore the invariant. This is possible since we store 1, and $1 \in \{0, 1\}$. Similarly, when performing
 453 the load in the primary thread, we open the invariant to learn $v \in \{0, 1\}$ which establishes the
 454 postcondition.
 455

456 What we have just done is created an invariant, shared it between two threads, and opened it
 457 around atomic operations in each thread. The choice of invariant assertion was up to us, but we
 458 had to ensure that all threads restored the invariant after each atomic step. This is the typical base
 459 pattern of concurrent program verification.
 460

461 **Ghost State and Updates.** While the invariants we just presented are sufficient for reasoning
 462 about simple concurrent programs, it has long been known [56] that for general reasoning about
 463 concurrent programs, we need some way to augment the physical state of the program with
 464 additional *auxiliary* or *ghost* state. Iris’s solution is a general form of *user-defined ghost state*, but
 465 for this paper we will only need *ghost map* assertions governed by the following rules.
 466

$$\begin{array}{c}
 \text{GHOSTMAPLOOKUP} \\
 \bullet^y m * k \hookrightarrow^y v \vdash \ulcorner m(k) = v \urcorner \\
 \\
 \text{GHOSTMAPUPDATE} \\
 \bullet^y m * k \hookrightarrow^y v \vdash \Rightarrow_{\mathcal{E}} \bullet^y m[k \leftarrow v'] * k \hookrightarrow^y v' \\
 \\
 \text{GHOSTMAPALLOC} \\
 \vdash \Rightarrow_{\mathcal{E}} \exists \gamma. \bullet^y \emptyset \\
 \\
 \text{GHOSTMAPINSERT} \\
 \frac{k \notin \text{dom}(m)}{\bullet^y m \vdash \Rightarrow_{\mathcal{E}} \bullet^y m[k \leftarrow v] * k \hookrightarrow^y v} \\
 \\
 \text{GHOSTMAPDELETE} \\
 \bullet^y m * k \hookrightarrow^y v \vdash \Rightarrow_{\mathcal{E}} \bullet^y m \setminus \{k\}
 \end{array}$$

470
 471 Ghost maps work with an arbitrary set of values and a countable set of keys. The ghost points-to
 472 $k \hookrightarrow^y v$ is very similar to the regular points-to, except that it is tied to a ghost heap instance named y
 473 instead of the regular heap. To avoid confusing ghost points-tos with the (physical) points-tos,
 474 we sometimes call the former “ghosts-tos”. The *authoritative* state $\bullet^y m$ describes the entire ghost
 475 heap map m . Updating the ghost heap at key k requires ownership of the authoritative state and
 476 of the ghosts-to for k . This update is entirely a logical operation (no physical program operation
 477 is involved), which is captured by the *update modality* $\Rightarrow_{\mathcal{E}}$. Intuitively, $\Rightarrow_{\mathcal{E}} P$ means that we can
 478 update the current ghost state to make P hold. Such updates can be executed in the postcondition
 479 of a `wp` (`WPWAND`), and whenever the goal is itself a `wp` (`WPUPDATEELIM`). The structure of the
 480 modality is given by `UPDATEINTRO`, `UPDATETRANS`, and `UPDATEFRAME`. Together, those rules show that
 481 this modality is a strong monad with respect to separating conjunction.
 482

483
 484 The update also shows up in the invariant allocation rule `INVALLOC`, since allocating a new
 485 invariant also updates ghost state. Additionally, we can use such an update to *peek* into an invariant
 486 with `UPDATEINV`. This opens the invariant “for 0 steps,” which is occasionally useful when one wants
 487 to perform ghost updates using parts of the invariant, or derive pure facts from it.
 488

489 ⁷The program uses let-bindings and sequencing, which are easily encoded using lambdas.
 490

3.2 Completeness for Concurrent Programs

Now that we have extended the logic with concurrency primitives, we revisit the question of soundness and completeness. First, our soundness result changes due to the switch to thread-pool semantics, which requires us to state soundness on the level of configurations:

$$\begin{aligned} \text{safe-tp}_\varphi(\vec{e}, \sigma) &\triangleq \forall \vec{e}', \sigma'. (\vec{e}, \sigma) \rightarrow_{\text{tp}}^* (\vec{e}', \sigma') \implies \\ &\quad \forall n, e''. \vec{e}'[n] = e'' \implies (\exists v. e'' = v \wedge (n = 0 \implies \varphi(v))) \vee \text{red}(e'', \sigma') \\ \text{safe}_\varphi(e, \sigma) &\triangleq \text{safe-tp}_\varphi([e], \sigma) \end{aligned}$$

Intuitively, threads are never allowed to be stuck, but only the “main” thread, the first in \vec{e}' , has to satisfy the postcondition φ on termination.⁸ For this modified definition, we have the following analogues of [Lemma 6](#) and [Lemma 7](#).

Lemma 11. *If $\text{safe-tp}_\varphi(\vec{e}, \sigma)$, then $\vec{e}[n]$ is either a value (satisfying φ if $n = 0$), or $(\vec{e}[n], \sigma)$ is reducible.*

Lemma 12. *If $\text{safe-tp}_\varphi(\vec{e}, \sigma)$ and $(\vec{e}, \sigma) \rightarrow_{\text{tp}} (\vec{e}_2, \sigma_2)$, then $\text{safe-tp}_\varphi(\vec{e}_2, \sigma_2)$.*

The soundness theorem establishes safety for a thread pool with e as the initial main thread.

Theorem 13 (Soundness). *If $\vdash \text{wp}_\top e \{v. \ulcorner \varphi(v) \urcorner\}$, then $\text{safe}_\varphi(e, \sigma)$ for all σ .*

For the completeness theorem, recall the two main facts that were maintained inductively in our previous sequential proof: (1) the current expression e and state σ satisfies *safe*, and (2) we have points-to for all of the locations in σ . As we move to the concurrent setting, safety is a property of the entire thread pool, and the points-to for the locations potentially need to be shared across all threads. To deal with these changes, we first introduce ghost state to track the status of the concurrently running threads. Specifically, we use a ghost map of type $\mathbb{N} \xrightarrow{\text{fin}} \text{Expr}$ to track the thread pool. This way, even though we reason about one thread at a time when proving a wp, we maintain that the overall system satisfies *safe*. Each thread will own the ghost points-to for its thread identifier. Second, in order to allow threads to share access to the state, we put all the heaps points-to in a shared invariant, along with the authoritative copy of the thread pool ghost state. This invariant looks as follows.

$$I_{\text{compl}} \triangleq \boxed{\exists \vec{e}, \sigma. \ulcorner \text{safe-tp}_\varphi(\vec{e}, \sigma) \urcorner * \bullet^Y \vec{e} * \bigstar_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v}^N}$$

We can now state the core lemma we are going to prove by Löb induction.

Lemma 14 (Per-thread Completeness). $I_{\text{compl}} * n \hookrightarrow^Y e \vdash \text{wp}_\top e \{v. n \hookrightarrow^Y v\}$.

PROOF. As in the sequential setting, our proof will eliminate the safety assumption to deduce that e is either a value or reducible; in the latter case, we perform a case analysis on which step it takes. Since the safety assumption is maintained by the invariant, we will need to first open the invariant.

This immediately presents us with an issue: We can only keep the invariant open around the next step if e is atomic, but without opening the invariant we do not know anything about e . Therefore, we do not know whether to apply [WPATOMICINV](#) or [UPDATEINV](#). It seems like the existing rules do not allow us to open the invariant and only then make the decision of whether to keep it open for the next step or not.

However, it turns out that it is possible to derive a rule where this choice is delayed until after the invariant has been opened. In particular, with this rule we can prove Atomic e after opening,

⁸We think of the overall return value of the concurrent execution as being the value that the first thread returns. This also matches the way forked-off threads only have a postcondition of True in [WPFORK](#).

540 using the assumptions previously kept in the invariant. The statement and proof of this rule is
 541 somewhat technical (and it uses the law of the excluded middle in the meta-logic); see [Lemma 31](#) in
 542 [Appendix A](#) for more details. With this lemma, we can open the invariant and use [GHOSTMAPLOOKUP](#)
 543 and [Lemma 11](#) to learn that e is a value or reducible. If it is a value, we immediately conclude with
 544 [WPVALUE](#). In the remaining case, we derive $e = K[e_1]$ where e_1 is base-reducible. More specifically,
 545 we get that $(e_1, \sigma) \rightarrow_{\text{base}} (e', \sigma', \vec{e}_f)$ for some σ, e', σ' .

546 For *pure steps* $e_1 \rightarrow_{\text{pure}} e'$, we have to immediately close the invariant again, since pure steps are
 547 in general not atomic. While this forces us to give up all assumptions about the state, we do not
 548 actually need them, since pure steps do not care about the state. Using [WPBIND](#) and [WPPURE](#), we
 549 are left with proving $\text{wp}_{\top} K[e'] \{ \cdot \cdot \cdot \}$. We also unlocked the Löb induction hypothesis since the
 550 pure step allowed us to eliminate a later. But to apply the induction hypothesis, we would need
 551 $n \hookrightarrow^{\gamma} K[e']$ while we currently only have $n \hookrightarrow^{\gamma} K[e_1]$. We can update our ghosts-to by opening
 552 the invariant again (for 0 steps) and using [GHOSTMAPUPDATE](#). To then close the invariant again, we
 553 use [Lemma 12](#) to prove that the new configuration is safe.

554 For [TPFORK](#), we have $e_1 = \text{fork } e_f$. Here, we also close the invariant and proceed similarly to
 555 the previous case. But now we need to prove two wp assertions, one for each thread. To do so, we
 556 can again open the invariant to update $n \hookrightarrow^{\gamma} K[\text{fork } e_f]$ to $K[()]$, but also insert a new entry n'
 557 (using [GHOSTMAPINSERT](#)) to create a new ghosts-to $n' \hookrightarrow^{\gamma} e_f$ for the newly created thread. We close
 558 the invariant again. We now use our Löb induction hypothesis *twice*, once for the parent thread
 559 and once for the new thread. This is possible because our Löb induction hypothesis holds without
 560 any resources and is thus duplicable.⁹

561 For *atomic expressions* modifying memory, we can keep the invariant open. This lets us obtain
 562 ownership of the entire heap (the iterated separating conjunction with all points-tos), so we can
 563 proceed like in the sequential case. Handling [CAS](#) is also straightforward since it is very similar to a
 564 load or a store. Once this is done, we close the invariant, which once again requires updating the
 565 ghosts-tos like in the other cases. \square

566
 567 This lemma, while having a rather technical precondition, contains the “meat” of the completeness
 568 proof. Proving a lemma that is the obvious inverse of [soundness](#) is now easy:

569 **Theorem 15** (Completeness of [ConcLang](#)). *If $\text{safe}_{\varphi}(e, \sigma)$ for all σ , then $\vdash \text{wp}_{\top} e \{v. \ulcorner \varphi(v) \urcorner\}$.*

571 **PROOF.** We instantiate the precondition with the empty heap \emptyset to get $\text{safe-tp}_{\varphi}([e], \emptyset)$. Next, we
 572 use [GHOSTMAPALLOC](#) to get an empty ghost map at γ , and insert $0 \leftarrow e$ into it giving us $0 \hookrightarrow^{\gamma} e$. We
 573 allocate the invariant I_{compl} since we have all required resources, including the iterated separating
 574 conjunction over the empty set. From [Lemma 14](#), we get $\text{wp}_{\top} e \{v. 0 \hookrightarrow^{\gamma} v\}$, while we must derive
 575 $\text{wp}_{\top} e \{v. \ulcorner \varphi(v) \urcorner\}$. For this, we use [WPWAND](#) which also allows us to open invariants with [UPDATEINV](#).
 576 We assume $0 \hookrightarrow^{\gamma} v$ and must show $\models_{\top} \ulcorner \varphi(v) \urcorner$, for which we once again open I_{compl} , and then use
 577 [GHOSTMAPLOOKUP](#) and [Lemma 11](#). Since v is a value and thus not reducible, and since $n = 0$, we get
 578 $\varphi(v)$ which finishes the proof. \square

580 3.3 The General Recipe

581 So far, we considered proving completeness for a *concrete* language. However, Iris provides a default
 582 program logic based on a *generic* wp construction, suitable for any language Λ that has a concept of
 583 expressions e , values v , state σ and evaluation contexts K . The language Λ must also have contextual
 584 operational semantics derived from a base reduction rule $\rightarrow_{\text{base}}$, with concurrency handled via
 585 thread-pool semantics. This wp comes with a language-independent soundness theorem that factors
 586

587 ⁹Using Iris terminology, the induction hypothesis is *persistent*, which allows us to use it several times.

most of the work of proving soundness for a particular language into a reusable building block. This raises the question: is there a similar language-independent building block for completeness?

To determine what that could look like, we return to the proof of [Theorem 15](#) and split it into a generic and a language-specific part. On the language-specific side, we have the individual cases for each base reduction case, and the fact that we need an iterated separating conjunction over the entire state σ (which can look very different for other languages). To abstract over the iterated separating conjunction, we introduce a predicate $S_\circ : \Lambda.State \rightarrow iProp$ collecting all language-specific resources required to represent the state. To abstract over the base reduction case distinction, we formulate a general condition on base reductions that is sufficient to obtain completeness. Our general completeness theorem then only requires that Λ satisfies this condition.

Theorem 16 (Generic Completeness). *Let Λ be a language satisfying [Condition 17](#) given below. If $safe_\varphi(e, \sigma)$, then $S_\circ(\sigma) \vdash wp_\top e \{v. \ulcorner \varphi(v) \urcorner\}$.*

To define this general condition, we introduce two auxiliary definitions:

$$\begin{aligned} AtomicPre(\mathcal{E}, e, \sigma, \Phi) &\triangleq \forall v', \sigma', \vec{e}_f. \ulcorner (e, \sigma) \rightarrow (v', \sigma', \vec{e}_f) \urcorner * S_\circ(\sigma') \multimap_{\mathcal{E}} \\ &\quad \Phi(v') * \multimap_{e_f \in \vec{e}_f} wp_\top e_f \{v. True\} \\ NonAtomicPre(\mathcal{E}, e, \Phi) &\triangleq \forall e', \vec{e}_f. (\forall \sigma. S_\circ(\sigma) \multimap_{\mathcal{E}} \exists \sigma'. \ulcorner (e, \sigma) \rightarrow^+ (e', \sigma', \vec{e}_f) \urcorner * S_\circ(\sigma')) \multimap_{\top} \\ &\quad wp_\top e' \{\Phi\} * \multimap_{e_f \in \vec{e}_f} wp_\top e_f \{v. True\} \end{aligned}$$

Note that we use $P \multimap_{\mathcal{E}} Q$ as an abbreviation for $P \multimap \multimap_{\mathcal{E}} Q$.

Condition 17 (Completeness).

$$\begin{aligned} \ulcorner base-red(e, \sigma) \urcorner * S_\circ(\sigma) \vdash \multimap_{\mathcal{E}} (\ulcorner Atomic e \urcorner * \forall \Phi. \triangleright AtomicPre(\mathcal{E}, e, \sigma, \Phi) \multimap wp_{\mathcal{E}} e \{\Phi\}) \vee \\ (S_\circ(\sigma) * \forall \Phi. \triangleright NonAtomicPre(\mathcal{E}, e, \Phi) \multimap wp_{\top} e \{\Phi\}) \end{aligned}$$

To understand this condition, imagine we are trying to prove it for some language Λ , and ignore the grayed-out parts for now (they are only relevant for reduction steps that fork new threads). We get to assume that e is base-reducible, as well as $S_\circ(\sigma)$. Our proof can now go ahead by case distinction on base-reducibility. The condition offers us two choices for proceeding.

If e is atomic, we proceed by proving $wp_{\mathcal{E}} e \{\Phi\}$ for some arbitrary Φ , under the precondition that $\triangleright AtomicPre(\mathcal{E}, e, \sigma, \Phi)$. Proving this wp will require a language-specific rule and since e is atomic, this rule should just leave us with $\Phi(v')$ for some v' . The only way to prove $\Phi(v')$ is to use *AtomicPre*. That means we have to prove that (e, σ) actually steps to (v', σ') and that we can update $S_\circ(\sigma)$ to $S_\circ(\sigma')$.

If e is not atomic, then we only use the second case.¹⁰ In this case, we must give up $S_\circ(\sigma)$ since the invariant can not be kept open around a non-atomic expression; this is why our goal is now $wp_\top e \{\Phi\}$ with the mask \top . Again, the next step is to apply some language-specific proof rule. The idea is that this rule makes progress on the proof by reducing e to e' (think of [WPPURE](#)). As long as we take at least one step in this reduction, we can apply *NonAtomicPre* to obtain a wp for e' . More specifically, applying *NonAtomicPre* forces us to prove that we can reduce e to e' starting in *any* state σ , potentially changing it to σ' .¹¹ As part of this, we are given $S_\circ(\sigma)$ and have to update it to $S_\circ(\sigma')$. If we can prove such a state-independent reduction from e to e' , then in exchange we obtain a wp for e' .

¹⁰Note that the second case is also usable for atomic expressions, which comes in handy since many technically atomic expressions do not affect the state at all.

¹¹In the example we considered above, the state will never change here, but in [§6](#) we will see an example of a language where the non-atomic case is useful for state-changing reductions. Similarly, our examples have so far only taken one step whereas the condition allows taking multiple; this is a straightforward generalization of [Lemma 12](#) and also used in [§6](#).

638 The grayed-out parts handle fork-based concurrency: If any steps we take fork off new threads \vec{e}_f ,
 639 then we need wp assertions not just for e' (or the postcondition for v') but also the new threads.

640 **Theorem 16** is powerful enough to verify many of our case studies. A similar theorem can be
 641 proven for sequential languages; this gives rise to a slightly stronger conclusion (see [Appendix B](#)).
 642 Some case studies come with a custom definition of wp which makes **Theorem 16** formally inappli-
 643 cable; for these we are able to re-prove a very similar theorem by closely following the pattern
 644 outlined so far.

645 4 Case Study: Completeness of HeapLang

646 HeapLang is the built-in language of Iris and is often used as a starting point for other Iris-based
 647 program logics. HeapLang is very similar to ConLang, with two main differences. First, freeing a
 648 heap location leaves behind a tombstone that prevents this location from being re-allocated in the
 649 future. Second, HeapLang has support for prophecy variables [34].

650 The first point is relevant for completeness because it means a program can assert that two
 651 separately allocated locations are different even if the first location is freed before the second one is
 652 allocated. HeapLang did not come with a proof rule intended to cover this case, so it may seem like
 653 we cannot prove the correctness of such a program and hence the logic is not complete. However,
 654 it turns out that the existing rules for HeapLang’s “metadata” mechanism (which allows the proof
 655 to associate arbitrary ghost data with a physical location) are sufficient to derive correctness of
 656 that program. By adding “metadata” tokens in I_{compl} , we are able to finish the completeness proof.

657 Prophecy variables are technically challenging because of how they are embedded in the opera-
 658 tional semantics. For space reasons, we do not explain the full details here. The main take-away is
 659 that the general recipe from §3.3 is strong enough to also cover this case. It is worth noting that we
 660 slightly adjust the operational semantics to obtain completeness: the operation for “resolving” a
 661 prophecy variable does not actually require that the variable exists, but no sensible proof rule can
 662 be given for the case where the variable does not exist.

663 Overall, this case study shows that our approach to proving completeness has the expected effect
 664 of identifying examples of programs that are correct with respect to the operational semantics, but
 665 cannot be verified with the existing proof rules. With only minor tweaks to the language and the
 666 logic, we establish completeness of HeapLang.

667 5 Case Study: Completeness of a Total Correctness Logic

668 Apart from the wp $e \{ \Phi \}$ assertion for partial correctness, Iris also comes with a variant wp $e [\Phi]$
 669 for total correctness [40]. This new connective satisfies the following soundness theorem.

670 **Theorem 18** (Total Soundness). *If $\vdash \text{wp}_{\top} e [v. \ulcorner \varphi(v) \urcorner]$, then $\text{safe}_{\varphi}(e, \sigma) \wedge \text{SN}_{\rightarrow_{\text{tp}}}([e], \sigma)$ for all σ .*

671 This total wp connective is concurrent and proves *scheduler-independent* termination, requiring
 672 the program to terminate even under a maximally demonic scheduler.¹² The rules for the total wp
 673 are very similar to the ones discussed so far. The only difference is that none of the rules contain any
 674 later modalities, so that we are not allowed to eliminate a later modality when taking a step. This
 675 makes Löb induction impossible, forcing us to prove that programs actually terminate, typically by
 676 induction in the meta-level logic.

677 This also means that we cannot use Löb induction in our completeness proof either. Instead,
 678 we use the strong normalization assumption, echoing §2.2. Our proof mostly follows the general
 679 recipe of §3.3, specifically **Theorem 16**, and can be split into a language-specific total completeness
 680

681 ¹²See [40, p. 11] for a discussion on the limitations of this definition.

condition and a generic part. In fact, the total completeness condition is identical to [Condition 17](#), except that the later modalities are removed.¹³

Theorem 19 (Generic Total Completeness). *Let Λ be a language that satisfies the total completeness condition. If $\text{safe}_\varphi(e, \sigma) \wedge \text{SN}_{\rightarrow_{\text{tp}}}([e], \sigma)$, then $\mathcal{S}_\sigma(\sigma) \vdash \text{wp}_\top e \left[v. \ulcorner \varphi(v) \urcorner \right]$.*

The heart of the proof is a thread-local lemma akin to [Lemma 14](#). The challenge in this proof is that it has to be done by induction on the strong normalization of some configuration $\rho_l = (\vec{e}_l, \sigma_l)$. But since all our knowledge about the “current” configuration is held inside the invariant, we do not have this “current” configuration available to do our induction. It seems like we need to open the invariant *before* starting the induction, but this does not lead to a sufficient induction hypothesis.

The solution is to do induction on a *lower bound* of the current configuration. Since the current configuration is not accessible at the start of the proof, we use additional ghost state to track past configurations that can reach the current configuration, one of which will be ρ_l . The proof is then by strong induction on this lower bound ρ_l being strongly normalizing.

As such, our induction hypothesis holds for all (transitive) successors of ρ_l . Once we open the invariant, we learn (from the additional ghost state) that $\rho_l \rightarrow_{\text{tp}}^* (\vec{e}, \sigma)$, where (\vec{e}, σ) is the “current” configuration. As usual in our completeness proofs, we need to take a step $(\vec{e}, \sigma) \rightarrow_{\text{tp}} (\vec{e}', \sigma')$ to a new configuration (\vec{e}', σ') . But by transitivity, this is also reachable from the lower bound ρ_l , so the induction hypothesis is applicable and we conclude the proof.

We instantiate this general recipe with the total correctness logic for `HeapLang`.

6 Case Study: Completeness of λ_{Rust}

`RustBelt` [32] develops a logical relation for a Rust-like type system. As part of this, they define a core calculus called λ_{Rust} that captures the most interesting aspects of Rust, along with a program logic that is used to define the logical relation. We have applied our methodology to their calculus, uncovering some missing proof rules and, more interestingly, showing that the existing proof rules for memory accesses are complete for the non-standard memory model of this language.

λ_{Rust} is superficially similar to `ConcLang`. The main difference is the memory model, which distinguishes between atomic and non-atomic accesses, and is defined in a way that causes programs with data races to be stuck. We say a data race occurs when there are two concurrent accesses to the same location, in which at least one is a write and at least one is non-atomic. By defining the semantics to make data races get stuck, if we then show that a program is safe (meaning it never gets stuck), then we know it is data race free. To detect data races in the semantics, λ_{Rust} essentially equips every memory location with a reader-writer lock. More precisely, a non-atomic access takes two steps: the first step acquires the lock, getting stuck if this is not possible; this produces an administrative redex. The second step performs the desired operation and releases the lock. Atomic accesses just ensure the lock is not held in a conflicting way; they only take a single step.

The program logic of λ_{Rust} entirely hides the existence of these reader-writer locks. The points-to assertion $\ell \mapsto v$ ensures that the lock is currently not taken. λ_{Rust} also supports fractional permissions [7], with $\ell \mapsto_q v$ for $q < 1$ allowing for the existence of concurrent readers.¹⁴ With this setup, the proof rules for loads and stores are entirely standard. In particular, the rules for atomic and non-atomic loads and stores are *the same!* However, only atomic accesses are Atomic, so `WPAtomicInv` cannot be used with non-atomic accesses.

¹³See [Condition 34](#) in [Appendix C](#) for its exact form.

¹⁴In fact, all our ghosts-tos and points-tos (in `ConcLang` and `HeapLang`) are fractional, but we have so far omitted this since it was not relevant for completeness. In λ_{Rust} , the completeness proof actually has to make use of fractions.

Proving completeness of the data-race-free program logic. The completeness proof uses the general recipe from §3.3. However, we have to adjust I_{compl} : previously, that invariant had exclusive ownership of *the entire heap*. This poses an issue since we cannot open the invariant around non-atomic accesses, preventing access to the points-to we need. But observe that in the specific case of a non-atomic store, there is no reason why the points-to *must* be in the invariant: No other thread will access this location, because there would otherwise be a data race which would contradict the assumption that the program is safe. Similarly, for reads, we know that other threads will not perform writes, so that we should be able to take a fraction of the points-to out of the invariant.

The task is now to construct an invariant that—for each location—tracks which (if any) accesses are currently happening on that location. If an access is ongoing, we temporarily remove the points-to from the invariant and replace it with a proof that no other thread can possibly be performing a concurrent access. (This is inspired by the construction used to reason about optimizations exploiting the absence of data races in Simuliris [17, §3].) Ignoring concurrent reads for now, a simplified version of this invariant for each heap location $\sigma(\ell) = (v, d)$ with value v and reader-writer lock state d is:

$$(\ell \mapsto_1 v \vee (\exists n, e. n \xrightarrow{Y}_{1/2} e * \lceil \text{write-about-to-happen}(e) \rceil)) * \lceil d = \text{Unlocked} \rceil$$

That is, we either have the full points-to, or we know that thread n is about to perform a non-atomic write. The invariant also asserts that the lock is not held, *i.e.*, there are no memory accesses *currently* going on. This may seem surprising—since the invariant always holds, does that not mean that there can never be an ongoing memory access? Let us take a look at what happens in the proof.

The relevant case in our completeness proof (*i.e.*, in the proof of [Condition 17](#)) is the case where a thread is about to perform a non-atomic write to some location ℓ . We open the invariant and prove that only the left disjunct of the per-location invariant can be satisfied: if the right disjunct were active, this would mean another thread were performing a write concurrently with the current thread, which contradicts our assumption that the program is safe. We can therefore switch the invariant to the right disjunct, taking the (physical) points-to out of the invariant while trading in half the thread ghosts-to. Next, we use the *non-atomic* case in [Condition 17](#) (the one that we only used for pure steps so far). This means we have to close the invariant—but crucially, we still own the points-to! We can now take *two* steps in the current thread using the usual proof rule for non-atomic writes. Afterwards, we apply *NonAtomicPre*, which requires us to prove that this reduction is indeed state-independent. This relies on us holding exclusive ownership of ℓ , which implies that no other thread can interfere with us. We know that the per-location invariant is in the right disjunct, so we can switch it to the left disjunct by giving back the points-to and obtaining the second half of our thread ghosts-to. Since we took two steps at once, we skipped over the intermediate state where the reader-writer lock is held, which means that the $d = \text{Unlocked}$ part of the invariant is also maintained. In other words, thanks to the exclusive ownership of ℓ , we do not actually have to show the invariant for every intermediate state. That is how $d = \text{Unlocked}$ can be “always” true.

To also handle non-atomic reads, the real invariant is more complex, requiring more detailed accounting of the fractional parts loaned out to each thread. We do not spell this out here. The overall completeness theorem has the additional assumption that there are no non-atomic accesses *currently happening* in the starting state σ (*i.e.*, all locks are unlocked), which is a reasonable restriction.

Identified completeness gaps. While performing our completeness proofs, we noticed that the existing rules in λ_{Rust} were not complete. The missing rules were unrelated to data races, but

785 instead about the semantics of comparisons, which can be non-deterministic when comparing
 786 pointers to already deallocated parts of memory. We presume that the absence of these rules was
 787 an oversight, since RustBelt never verified any programs that depended on this non-determinism.

788 Since λ_{Rust} used a block-based memory model [43], the program logic also exposes *block tokens*
 789 tracking the size of each block, which we put into the state invariant \mathbb{S}_b . In doing so, we realized
 790 that these block tokens lacked an exclusivity law, which we therefore added.

791 7 Case Study: Completeness of a Logic for Execution Time Bounds

792 Mével et al. [51] develop a logic for bounding the running time of programs by embedding time
 793 credits [4] into Iris. Here, we consider a completeness proof for such a logic. Similar to their
 794 approach, we extend ConLang with a **tick** operation,¹⁵ where **tick**(e) encodes the idea that
 795 executing e incurs a cost. For instance, if we want to bound the number of comparisons that a sorting
 796 function performs, we would replace each comparison $e_1 \leq e_2$ in the function with **tick**($e_1 \leq e_2$).

797 For the semantics of **tick**, the state σ of the program is extended with a natural number counter
 798 field $\sigma.tc$. Evaluating **tick**(e) first evaluates e until it reaches a value v , then decrements the $\sigma.tc$
 799 counter by one, and returns v . The **tick** operation gets stuck if the $\sigma.tc$ counter is zero. Thus, if
 800 $\sigma.tc = m$ and (e, σ) never gets stuck, then e performs no more than m **tick** operations.¹⁶

801 To reason about **tick**, the logic adds an assertion $\$(n)$ that represents permission to perform up
 802 to n **tick** operations. These assertions are called *time credits* and have the following rules:

$$803 \begin{array}{ll} \text{TIMECREDITSPLIT} & \text{WFTICK} \\ 804 \$(m + n) \dashv\vdash \$(m) * \$(n) & \$(1) * \triangleright \Phi(v) \vdash \text{wp}_{\top} \mathbf{tick}(v) \{\Phi\} \end{array}$$

805 **TIMECREDITSPLIT** allows us to split and join time credits with the separating conjunction. **WFTICK**
 806 lets us reason about a **tick** operation by spending 1 credit.

807 To prove that a program does at most m **tick** operations, we prove a wp starting with an initial
 808 budget of $\$(m)$, as expressed in the logic's soundness theorem.

809 **Theorem 20** (Soundness of Time Credits).

810 If $\$(m) \vdash \text{wp}_{\top} e \{v. \ulcorner \varphi(v) \urcorner\}$, then $\text{safe}_{\varphi}(e, \sigma)$ for all σ with $\sigma.tc = m$.

811 Because this theorem implies e never gets stuck and $\sigma.tc = m$, we know that e performs at most
 812 m **tick** operations. Using the techniques of this paper, it is straightforward to prove a completeness
 813 result showing that the above two rules for time credits are all we need:

814 **Theorem 21** (Completeness of Time Credits).

815 If $\text{safe}_{\varphi}(e, \sigma)$ for all σ with $\sigma.tc = m$, then $\$(m) \vdash \text{wp}_{\top} e \{v. \ulcorner \varphi(v) \urcorner\}$.

816 **PROOF SKETCH.** Apply the generic completeness recipe (**Theorem 16**), and instantiate $\mathbb{S}_b(\sigma)$ as
 817 $(\ast_{(\ell \leftarrow v) \in \sigma.h} \ell \mapsto v) * \$(\sigma.tc)$. The proof largely follows the proof of **Lemma 14**, with an additional
 818 case for **tick**. For that case, from the safety assumption we know that $\sigma.tc$ is at least 1, so by using
 819 **TIMECREDITSPLIT**, we can pull out $\$(1)$, which we spend using **WFTICK** to justify the **tick** step. \square

824 8 Case Study: Completeness of a Logic for Probabilistic Error Bounds

825 In this section, we turn to Eris [1], an Iris-based separation logic for proving bounds on the
 826 probabilities of errors in higher-order sequential probabilistic programs. The idea is that many
 827 randomized programs are designed to satisfy a specification with very high probability, and only
 828

829 ¹⁵Instead of extending ConLang, Mével et al. [51] actually implement **tick** as a library in ConLang. However, such an
 830 approach can only be complete for programs that use the library “correctly,” which is hard to state, let alone reason about.

831 ¹⁶It is essential to add **tick** around every operation we actually wish to count as incurring a cost. Mével et al. [51] discuss
 832 a translation that adds a **tick** to all steps of the language, while other work has instead parameterized the operational
 833 semantics with a *cost model* for each transition, but the basic idea is the same.

fail when some rare event occurs. Eris allows one to prove a bound on this probability of failure. Because Eris’s specifications cover a probabilistic property, it does not use Iris’s default definition of wp , so we cannot directly apply [Theorem 16](#) or [Theorem 19](#) to establish completeness. Nevertheless, as we will see, the general proof approach developed in the previous sections can be used to obtain a completeness result for Eris.

Eris targets a language that extends SeqLang with a probabilistic choice operator $\mathbf{rand} N$, which evaluates to a random integer drawn uniformly from the set $\{0, \dots, N\}$. The language removes \mathbf{free} and makes allocation deterministic, so that the only non-determinism is the probabilistic choice in \mathbf{rand} . To define a semantics for this language, the authors first define the function $\text{step} : \text{Expr} \times \text{State} \rightarrow \text{Expr} \times \text{State} \rightarrow [0, 1]$, where $\text{step}(e, \sigma)(e', \sigma')$ gives the probability that (e, σ) steps to (e', σ') in a single step. This plays an analogous role to the \rightarrow relation in SeqLang. They then inductively define the function $\text{exec}_n : \text{Expr} \times \text{State} \rightarrow \text{Val} \rightarrow [0, 1]$, where $\text{exec}_n(e, \sigma)v$ gives the probability that (e, σ) will terminate and return v in at most n steps. The partially applied function $\text{exec}_n(e, \sigma)$ can be seen as a *sub-distribution* on values, meaning that the sum of the probabilities it assigns to all values is a number $p \leq 1$. The gap between p and 1 represents the probability of non-termination after n steps. Finally, $\text{exec}(e, \sigma)v$ is defined as the limit of exec_n as $n \rightarrow \infty$.

Eris comes in two variants, a *partial* version and a *total* version. In this section, we will focus on the total version, though our Rocq development proves a completeness result for the partial version as well. The key feature of Eris is the *error credit* assertion written $\mathcal{E}(\varepsilon)$ with $\varepsilon \in [0, 1]$, which represents a “budget” or “permission” to err with probability ε . Just as time credits represent permission to incur a cost and must be spent every time a **tick** occurs, error credits are spent when a specification fails to hold with some probability. By constraining the initial budget, the probability of failure is bounded from above, as formalized in the soundness theorem for total Eris:

Theorem 22 (Soundness of Total Eris).

If $\mathcal{E}(\varepsilon) \vdash \text{wp } e [v. \ulcorner \varphi(v) \urcorner]$, then $\text{Pr}_{\text{exec}(e, \sigma)}[\varphi] \geq 1 - \varepsilon$ for all σ .

The conclusion of this statement says that with probability at least $1 - \varepsilon$, evaluating (e, σ) will terminate with a value satisfying φ . In other words, the probability that it either diverges, gets stuck, or returns a value that fails to satisfy φ is at most ε , the initial error budget.

We reason about error credits using the following five rules:

$$\begin{array}{c}
 \text{ERRCOMBINE} \qquad \qquad \qquad \text{ERRWEAKEN} \qquad \qquad \qquad \text{ERR1} \\
 \mathcal{E}(\varepsilon_1) * \mathcal{E}(\varepsilon_2) \dashv\vdash \mathcal{E}(\varepsilon_1 + \varepsilon_2) \qquad \mathcal{E}(\varepsilon_1) * \ulcorner \varepsilon_2 < \varepsilon_1 \urcorner \vdash \mathcal{E}(\varepsilon_2) \qquad \mathcal{E}(1) \vdash \text{False} \\
 \\
 \text{TWP RANDEXP} \qquad \qquad \qquad \text{TWP THINAIR} \\
 \frac{\sum_{i=0}^N \frac{\mathcal{E}(i)}{N+1} = \varepsilon}{\mathcal{E}(\varepsilon) \vdash \text{wp } \mathbf{rand} N [n. \mathcal{E}(n)]} \qquad \frac{\forall \varepsilon'. \ulcorner \varepsilon' > \varepsilon \urcorner * \mathcal{E}(\varepsilon') * P \vdash \text{wp } e [\Phi]}{\mathcal{E}(\varepsilon) * P \vdash \text{wp } e [\Phi]}
 \end{array}$$

ERRCOMBINE allows us to split and join error credits, much like time credits. Using **ERRWEAKEN**, we can weaken our credits to a smaller amount. **ERR1** says that if we have 1 error credit, then we can derive False. The intuition here is that if we look at the statement of soundness, then owning $\mathcal{E}(1)$ means we only have to show our specification holds with probability at least $1 - 1 = 0$, which is trivial, so there is nothing to prove. The rule **TWP RANDEXP** allows us to re-distribute our credits across the different branches of $\mathbf{rand} N$, as long as the resulting error credits have the same *expected value* as the amount we started with. The user instantiates the rule with a function $\mathcal{E} : \{0, \dots, N\} \rightarrow [0, 1]$, where $\mathcal{E}(i)$ represents how many error credits we will have in the postcondition when the random sample returns i . The idea is that, if an error will occur on some branch based on the outcome of $\mathbf{rand} N$, we shift credits from the non-erroring branches to the erroring branches, so that they

883 have $\not\leq(1)$ and can apply **ERR1**. Finally, **TWP_{THIN}AIR** says that when proving a wp, if we start with ε
 884 credits, then it suffices to show that the proof can be carried out with any ε' that is strictly greater
 885 than ε . In other words, this rule allows us to conjure up additional error credits “out of thin air,” but
 886 the additional error might be arbitrarily small. The justification for this rule relies on a continuity
 887 property of wp, which we use to take the limit as ε' converges to ε .

888 Using these rules, it is possible to verify almost-sure termination of programs, *i.e.*, show that
 889 they terminate with probability 1. Almost-surely terminating programs are not necessarily strongly
 890 normalizing, but the probability of a diverging execution is 0. In general, reasoning about almost-
 891 sure termination is challenging, and the research literature has a number of variants of program
 892 logics and proof rules for showing almost-sure termination of probabilistic while loops, with recent
 893 interest in characterizing the completeness of these proof rules [46, 49].

894 We will show that Total Eris is *nearly* complete. The failure of completeness has nothing to do
 895 with probabilistic reasoning, but rather comes from the fact that the language uses a deterministic
 896 allocator, yet the proof rule for allocation is the same as the one we saw in §2, where we non-
 897 deterministically get a points-to for the returned location. Thus, instead we prove completeness for
 898 programs that do not allocate (but may make use of existing allocated locations):

899 **Theorem 23** (Completeness of Total Eris). *For all configurations (e, σ) not containing any **ref**
 900 expressions, if $\Pr_{\text{exec}}(e, \sigma)[\varphi] \geq 1 - \varepsilon$, then $\not\leq(\varepsilon) * \mathbf{*}_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v \vdash \text{wp } e [v. \ulcorner \varphi(v) \urcorner]$.*

901 To approach the completeness theorem, we start from executions up to a finite number of steps.

902 **Lemma 24.** *For all configurations (e, σ) not containing any **ref** expressions, if $\Pr_{\text{exec}_n}(e, \sigma)[\varphi] \geq 1 - \varepsilon$,
 903 then $\not\leq(\varepsilon) * \mathbf{*}_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v \vdash \text{wp } e [v. \ulcorner \varphi(v) \urcorner]$.*

904 **PROOF.** We first weaken the $\not\leq(\varepsilon)$ to $\not\leq(1 - \Pr_{\text{exec}_n}(e, \sigma)[\varphi])$ via **ERRWEAKEN**, and then perform
 905 induction on n , with e and σ quantified. The structure mirrors the completeness proof in §2, but
 906 with one key difference: whereas that proof performed induction on strong normalization (or, for
 907 partial programs, used Löb induction), here we induct on the index n of exec_n .

908 When $n = 0$, either $e = v$ for some $v \in \varphi$, or $\Pr_{\text{exec}_0}(e, \sigma)[\varphi] = 0$. In both cases, the wp can be
 909 easily derived. In the inductive step, we need to show:

$$910 \quad \text{If } \forall e', \sigma'. \not\leq(1 - \Pr_{\text{exec}_n}(e', \sigma')[\varphi]) * \mathbf{*}_{(\ell \leftarrow v) \in \sigma'} \ell \mapsto v \vdash \text{wp } e' [v. \ulcorner \varphi(v) \urcorner] \text{ (the IH),}$$

$$911 \quad \text{then } \not\leq(1 - \Pr_{\text{exec}_{n+1}}(e, \sigma)[\varphi]) * \mathbf{*}_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v \vdash \text{wp } e [v. \ulcorner \varphi(v) \urcorner].$$

912 Here, the nontrivial case happens when e is a reducible expression (and hence not a value). As
 913 in §2, we decompose e into $K[e_1]$, where e_1 is base-reducible. We then apply **WP_{BIND}** and proceed
 914 by case distinction on the base step e_1 takes. For pure reductions and heap-dependent operations
 915 (only loads and stores are possible), the proof follows exactly the same pattern as before.

916 The genuinely new case arises when e_1 is **rand** N . In that case, $e = K[\text{rand } N]$ and the configu-
 917 ration (e, σ) steps to $(K[i], \sigma)$ for every $i = 0, \dots, N$ with equal probability. The failure probability
 918 of this expression is given by:

$$919 \quad 1 - \Pr_{\text{exec}_{n+1}}(e, \sigma)[\varphi] = \sum_{i=0}^N \frac{1 - \Pr_{\text{exec}_n}(K[i], \sigma)[\varphi]}{N + 1}$$

920 We can therefore apply **TWP_{RAND}EXP** with $\mathcal{E}(i) \triangleq 1 - \Pr_{\text{exec}_n}(K[i], \sigma)[\varphi]$. This yields $\not\leq(\mathcal{E}(i))$ for
 921 each outcome i , with which we can obtain wp $K[i] [v. \ulcorner \varphi(v) \urcorner]$ by the induction hypothesis. \square

922 With this lemma, we can prove **Theorem 23**. We start by applying **TWP_{THIN}AIR**, giving us $\not\leq(\varepsilon')$
 923 for some $\varepsilon' > \varepsilon$. From this inequality, we have that $\Pr_{\text{exec}}(e, \sigma)[\varphi] > 1 - \varepsilon'$. By monotonicity of

exec, there exists n such that $\Pr_{\text{exec},n}(e,\sigma)[\varphi] \geq 1 - \varepsilon'$. It follows by [Lemma 24](#) that $\ast_{(\ell \leftarrow v) \in \sigma} \ell \mapsto v \ast \not\vdash (\varepsilon') \vdash \text{wp } e [v. \varphi(v)]$.

9 Case Study: Completeness of a Relational Logic

So far, we have looked at instances of Iris used for *unary* reasoning, where a specification describes a single program's behavior. However, Iris has also been used for *relational* reasoning, in which specifications relate the behavior of two programs. This kind of reasoning is useful because, among other applications, it can be used to show that a program *refines* a specification program. A common approach is to embed relational reasoning on top of an existing unary program logic by representing the specification program as a form of ghost state, as in the CaReSL logic [66]. With this approach, one introduces assertions to describe the state of this specification program. In this section, we will look at how this approach works for the case of the ConLang language.

The logic comes with three new assertions. First, we have an assertion $j \Rightarrow e$ which states that the thread at index j in the spec program's thread pool is executing expression e , while $\ell \mapsto_s v$ says that in the spec program's heap, location ℓ points to v . Under the hood, both of these are just Iris ghost state, similar to $k \hookrightarrow v$. Finally, we have an assertion specCtx^N which represents the invariant that stores the authoritative copies of the spec program's state. This invariant ensures that the spec state can only be updated in a way that reflects the reduction semantics of the spec program. Then one derives rules for executing steps of the spec program by performing Iris ghost state updates. For example, for any \mathcal{E} such that $N \subseteq \mathcal{E}$, we have:

$$\begin{aligned} \text{specCtx}^N \ast j \Rightarrow K[\ell \leftarrow w] \ast \ell \mapsto_s v \vdash \Vdash_{\mathcal{E}} (j \Rightarrow K[()] \ast \ell \mapsto_s w) \\ \text{specCtx}^N \ast j \Rightarrow K[!\ell] \ast \ell \mapsto_s v \vdash \Vdash_{\mathcal{E}} (j \Rightarrow K[v] \ast \ell \mapsto_s v) \\ \text{specCtx}^N \ast j \Rightarrow K[\text{ref } v] \vdash \Vdash_{\mathcal{E}} (\exists \ell. j \Rightarrow K[\ell] \ast \ell \mapsto_s v) \\ \text{specCtx}^N \ast j \Rightarrow K[\text{fork } e] \vdash \Vdash_{\mathcal{E}} (\exists j'. j \Rightarrow K[()] \ast j' \Rightarrow e) \end{aligned}$$

The first two rules perform stores and loads, respectively, using the spec program's points-to assertions for the corresponding location. The third rule allocates a new location, creating a new spec points-to for the new location. The fourth rule forks a thread and creates a fresh $j' \Rightarrow e$ that represents the forked-off spec thread. Similar rules can be shown for all of the other primitive commands and pure steps of the language.

Then, to prove that an implementation e refines a specification e' , one proves a particular wp for e , as captured by the following soundness theorem.

Theorem 25 (Relational Soundness). *If $\text{specCtx}^N \ast 0 \Rightarrow e' \vdash \text{wp } e \{v. \exists v'. 0 \Rightarrow v' \ast \ulcorner \varphi(v, v') \urcorner\}$, then*

- (1) (e, \emptyset) is safe, and
- (2) for all v, \vec{e} , and σ such that $([e], \emptyset) \rightarrow_{\text{tp}}^{\ast} (v :: \vec{e}, \sigma)$ there exists v', \vec{e}' and σ' such that $([e'], \emptyset) \rightarrow_{\text{tp}}^{\ast} (v' :: \vec{e}', \sigma')$ and $\varphi(v, v')$.

This theorem requires a proof of the implementation e where in the precondition, spec thread 0 executes e' , and in the postcondition, the spec thread has reached some value v' that is related to the result v of the implementation. In the conclusion of the theorem, the first part states the usual safety of e that we expect from the wp. The second part captures the relational nature of the logic: it shows that for every execution of e in which the first thread terminates in some value v , there is a corresponding execution of the spec program e' in which its first thread terminates in a value v' that is related to v under φ . In other words, this shows that e *refines* e' . Intuitively, this soundness theorem follows from the fact that, from the way that specCtx and the spec program

are defined, the only way we could have gotten $0 \Rightarrow v'$ in the postcondition is by constructing such an execution of e' . Because the relational reasoning is constructed by embedding on top of the unary wp, we get to re-use the soundness proof of the underlying wp. As a result, the proof of [Theorem 25](#) is an easy consequence of the underlying soundness of the wp and the definition of the invariant in specCtx.

Now, given that we know the underlying wp is also complete, can we also derive completeness of the relational logic, *i.e.*, a converse to [Theorem 25](#)? To show such a converse, we would get as assumptions that (e, σ) is safe and that e refines e' under relation φ , and we would need to derive the entailment. Applying the completeness property of the underlying wp ([Theorem 15](#)), we derive a wp about e of the following form:

$$\text{wp } e \{v. \lceil \exists v', \vec{e}, \vec{e}', \sigma, \sigma'. ([e], \emptyset) \rightarrow_{\text{tp}}^* (v :: \vec{e}, \sigma) \wedge ([e'], \emptyset) \rightarrow_{\text{tp}}^* (v' :: \vec{e}', \sigma') \wedge \varphi(v, v') \rceil \}$$

The postcondition here has the pure property about the existence of a corresponding execution of e' . To finish the proof, we now “just” need to show from this pure fact that it is possible to update $0 \Rightarrow e'$ to $0 \Rightarrow v'$. An obvious idea would be to do so by induction on $([e'], \emptyset) \rightarrow_{\text{tp}}^* (v' :: \vec{e}', \sigma')$, showing that for each step in this trace, we can do an update on the ghost state.

However, perhaps surprisingly, this approach is blocked by one fundamental issue: the ghost update rule for allocation shown above. The conclusion of this allocation rule says that we just learn that there exists some location at which the allocation has occurred. But because allocation is non-deterministic, this location may not be the same location as the one allocation happens at in our assumed execution about e' !

In some sense, the notion of refinement we are working with is too brittle: we would like to say that the refinement does not depend on the specific locations selected by the allocation. In fact, by changing the definition of the ghost state and the invariant in specCtx, we can derive the following *stronger* version of the soundness theorem:

Theorem 26 (Strong Relational Soundness).

If $\text{specCtx}^N * 0 \Rightarrow e' \vdash \text{wp } e \{v. \exists v'. 0 \Rightarrow v' * \lceil \varphi(v, v') \rceil \}$, then

- (1) (e, \emptyset) is safe, and
- (2) For any infinite set S of locations, for all v, \vec{e} , and σ such that $([e], \emptyset) \rightarrow_{\text{tp}}^* (v :: \vec{e}, \sigma)$ there exists v', \vec{e}' and σ' such that $([e'], \emptyset) \rightarrow_{\text{tp}}^* (v' :: \vec{e}', \sigma')$ where $\varphi(v, v')$ and $\text{dom}(\sigma') \subseteq S$.

In other words, with this version of the theorem, we get to pick an infinite set S of locations,¹⁷ and we get that *all* of the resulting executions of e' only allocate to locations in S . This ensures that the executions of e' cannot depend too much on some particular choice of allocation. With this stronger definition of specCtx, we introduce an additional ghost state assertion $\text{reserve}(S)$ that allows us to *reserve* a set of locations. This assertion is used in the following two rules for $N \subseteq \mathcal{E}$:

$$\text{specCtx}^N \vdash \models_{\mathcal{E}} \exists S. \lceil \text{infinite}(S) \rceil * \text{reserve}(S)$$

$$\text{specCtx}^N * j \Rightarrow K[\mathbf{ref } v] * \text{reserve}(S) * \lceil \ell \in S \rceil \vdash \models_{\mathcal{E}} (j \Rightarrow K[\ell] * \ell \mapsto_s v * \text{reserve}(S \setminus \{\ell\}))$$

The first rule allows us to obtain, at any point, some infinite set S of reservations. The second rule says that if $\ell \in S$, where S is a set we have reserved, then we can force the allocation to be ℓ , which removes ℓ from the reservation set. With these stronger rules, we can prove a form of completeness that is converse to [Theorem 26](#).

Theorem 27 (Strong Relational Completeness). *If e is an expression such that*

¹⁷More precisely, in our Rocq development, we use a form of constructive infinite sets that allows for splitting an infinite set into two infinite sets without using a choice principle.

1030 (1) (e, \emptyset) is safe, and

1031 (2) For any infinite set S of locations, for all v, \vec{e} , and σ such that $(e, \emptyset) \rightarrow_{\text{tp}}^* (v :: \vec{e}, \sigma)$ there exists
 1032 v', \vec{e}' and σ' such that $([e'], \emptyset) \rightarrow_{\text{tp}}^* (v' :: \vec{e}', \sigma')$ where $\varphi(v, v')$ and $\text{dom}(\sigma') \subseteq S$.

1033 then $\text{specCtx}^N * 0 \Rightarrow e' \vdash \text{wp } e \{v. \exists v'. 0 \Rightarrow v' * \lceil \varphi(v, v') \rceil\}$.

1035 The proof follows the sketch we started with above, except that we start by getting $\text{reserve}(S)$
 1036 for some infinite set S . From our assumptions, we can then get that there exists an execution of
 1037 e' that yields a related value where all allocations are in S . Now, because we have $\text{reserve}(S)$, we
 1038 can perform ghost updates that match this execution. By using the stronger ghost update rule for
 1039 allocations, we ensure that all of the allocated locations in our ghost state match the ones in the
 1040 execution we are given.

1042 10 Semantic Conditions for Completeness

1043 In earlier sections, we examined several existing instantiations of Iris and established the complete-
 1044 ness of their proof rules. In this section, we turn to the question of how, when coming up with a
 1045 new instantiation of Iris, one can try to ensure that the resulting logic will be complete. For this
 1046 purpose, we zoom out from any concrete case study and consider again the default Iris program
 1047 logic with its language-independent wp.

1048 As mentioned in §3.3, this logic is first parameterized by a notion of a language Λ , which
 1049 must satisfy some basic structural properties. The next parameter of the framework is the *state*
 1050 *interpretation predicate*, \mathbf{S}_\bullet , which is a predicate over program states.¹⁸ In Iris, assertions such as
 1051 $\ell \mapsto v$ do not directly talk about the physical state; they are just ghost state. The role of \mathbf{S}_\bullet is to
 1052 define a relationship between that ghost state and the program's physical state. Under the hood,
 1053 $\ell \mapsto v$ is defined using ghosts-tos, and the corresponding authoritative state is held by the state
 1054 interpretation. This allows the user to define ghost state assertions that track the physical state.
 1055 Since \mathbf{S}_\bullet often needs to refer to some ghost state name γ , we typically write $\mathbf{S}_\bullet^\gamma$ to indicate this
 1056 dependency on γ .

1057 The \mathbf{S}_\bullet is used by Iris in the *definition* of the wp. Whereas many program logics *define* the wp so
 1058 that it is the weakest precondition, in Iris, the wp is instead defined recursively inside the logic.

1059 **Definition 28** (wp). The assertion $\text{wp}_\mathcal{E} e \{\Phi\}$ is the fixed point of the following guarded recursive
 1060 definition [33]

$$\begin{aligned}
 1062 \quad (v \in \text{Val}) \quad \text{wp}_\mathcal{E} v \{\Phi\} &\triangleq \mathbb{H}_\mathcal{E} \Phi(v) \\
 1063 \quad (e \notin \text{Val}) \quad \text{wp}_\mathcal{E} e \{\Phi\} &\triangleq \forall \sigma. \mathbf{S}_\bullet(\sigma) * \mathbb{E}_\mathcal{E} \mathbb{H}_\emptyset \lceil \text{red}(e, \sigma) \rceil * \\
 1064 &\quad \left(\forall e', \sigma', \vec{e}_f. \lceil (e, \sigma) \rightarrow (e', \sigma', \vec{e}_f) \rceil * \mathbb{H}_\emptyset \triangleright_\emptyset \mathbb{H}_\mathcal{E} \right. \\
 1065 &\quad \left. \mathbf{S}_\bullet(\sigma') * \text{wp}_\mathcal{E} e' \{\Phi\} * \bigstar_{e_f \in \vec{e}_f} \text{wp}_\top e_f \{_ \cdot \text{True}\} \right)
 \end{aligned}$$

1066 The base case of this definition says that if e is already a value v , then the postcondition $\Phi(v)$
 1067 must hold immediately (under an update modality $\mathbb{H}_\mathcal{E}$). The inductive step consists of two parts.
 1068 First, given the state interpretation for the current state σ , we must show $\text{red}(e, \sigma)$, *i.e.*, that the
 1069 current configuration is reducible. Then, for any e' that e may reduce to under σ , we must show
 1070 the state interpretation holds for the new state σ' , and that the wp holds for e' and any forked
 1071 threads \vec{e}_f . This definition uses a mask-changing variant of the update modality, where the update
 1072 can access some invariants and leave them open (or close some previously open invariants). The
 1073

1074 ¹⁸This is somewhat simplified; we focus on what is relevant for this discussion.

various update modalities along the way allow us to open invariants and modify ghost state in the course of proving these facts, while the later modality \triangleright in the definition guards the recursive occurrence of the wp to ensure that the fixed point exists [11, 33]. The later modality here is also the reason why we can strip a later in hypotheses at each step.

On top of this definition, the framework derives various language-generic rules about the wp, such as WPVALUE , WPWAND , and WPATOMICINV . In turn, the user instantiating the framework is responsible for deriving wp rules for each of the language-specific primitives. These proofs require unfolding the definition of the wp and exploiting the user's selected definition of \mathbf{S}_\bullet .

The framework provides a generic soundness proof for this definition of wp. More precisely, in Iris this property is called *adequacy*, by analogy to the use of that term in denotational semantics, as it shows that the model of the wp actually says something about a program's behavior.

Theorem 29 (Generic Adequacy). *For all states σ , let \mathbf{S}_\bullet^Y be an Iris predicate on states such that*

$$\vdash \exists \gamma. \mathbf{S}_\bullet^Y(\sigma) * \mathbf{S}_\circ^Y(\sigma) \quad (\text{STATEALLOC})$$

then for any meta-level predicate φ , $(\forall \gamma. \mathbf{S}_\bullet^Y(\sigma) \vdash \text{wp}_\top e \{v. \ulcorner \varphi(v) \urcorner\})$ implies $\text{safe}_\varphi(e, \sigma)$.

The condition of this theorem requires us to show the state interpretation and state invariant hold for some initial state σ . In turn, if we then derive a wp about e with $\mathbf{S}_\bullet^Y(\sigma)$ as an assumption, we get that (e, σ) satisfies safe_φ . All of the specific soundness theorems we have looked at for partial correctness have been consequences of this generic adequacy theorem.

A natural question is whether we can show a kind of converse to this theorem under some assumptions about \mathbf{S}_\bullet and \mathbf{S}_\circ . That is, can we show that whenever an expression satisfies *safe*, the wp holds? Such a converse would not exactly be a completeness theorem, because it would not tell us whether a particular collection of *proof rules* suffice for deriving the wp. But such a property is a necessary condition for completeness, since any complete logic would satisfy it. Thus, if our logic fails to satisfy this property, it has no hope of being complete. Moreover, establishing this would show that the wp is in fact the *weakest* precondition, *i.e.*, it is both a necessary and sufficient precondition. That is because any other sufficient precondition would imply *safe*, and therefore imply the wp.

We say that an instance of wp with this property is *requisite*. The following theorem gives sufficient conditions for requisiteness.

Theorem 30 (Generic Requisiteness). *For all γ , let \mathbf{S}_\bullet^Y be an Iris predicate on states such that*

$$\begin{aligned} \forall e, \sigma, \sigma_l. \ulcorner \text{red}(e, \sigma_l) \urcorner * \mathbf{S}_\bullet^Y(\sigma) * \mathbf{S}_\circ^Y(\sigma_l) &\Rightarrow \ulcorner \text{red}(e, \sigma) \urcorner * & (\text{STATEUPD}) \\ (\forall e', \sigma', \vec{e}_f. \ulcorner (e, \sigma) \rightarrow_{\text{tp}} (e', \sigma', \vec{e}_f) \urcorner &\Rightarrow \exists \sigma_l'. \ulcorner (e, \sigma_l) \rightarrow_{\text{tp}} (e', \sigma_l', \vec{e}_f) \urcorner * \mathbf{S}_\bullet^Y(\sigma') * \mathbf{S}_\circ^Y(\sigma_l') \end{aligned}$$

then for any meta-level predicate φ , $\text{safe}_\varphi(e, \sigma)$ implies $\mathbf{S}_\bullet^Y(\sigma) \vdash \text{wp}_\top e \{v. \ulcorner \varphi(v) \urcorner\}$.

To get some intuition for the condition STATEUPD , we think of \mathbf{S}_\bullet as the “authoritative” full description of the physical state σ , whereas \mathbf{S}_\circ describes some “fragment” or subset of the state σ_l . Then, if e is reducible in the subset of state σ_l for which we have \mathbf{S}_\circ , and we are given \mathbf{S}_\bullet for the full state σ , then we must first be able to show that e is also reducible in σ . Furthermore, for every e' that e can step to from σ , we must update the \mathbf{S}_\bullet accordingly, and moreover, we must show that there is a corresponding step to e' from σ_l that forks the same threads. For this step from σ_l , we have to be able to update \mathbf{S}_\circ .

This captures a sort of locality of the program semantics: ownership of σ_l is enough to characterize how e steps in any larger state that is compatible with ownership of $\mathbf{S}_\bullet^Y(\sigma_l)$.

PROOF SKETCH OF **THEOREM 30**. As in the proof of **Theorem 15**, we first allocate ghost state that will be used to track the status of all concurrent threads that are created. Exactly like before, we create an invariant that stores S_0 , the ghost map for the threads, and the pure fact that the configuration satisfies safe-tp_φ .

$$I_{\text{compl}} \triangleq \boxed{\exists \vec{e}, \sigma_1. \ulcorner \text{safe-tp}_\varphi(\vec{e}, \sigma_1) \urcorner * \bullet^Y \vec{e} * S_0(\sigma_1)}^N$$

Still using Löb induction, we prove $I_{\text{compl}} * n \hookrightarrow^Y e \vdash \text{wp}_\top e \{v. n \hookrightarrow^Y v\}$, which has the same shape as the entailment in **Lemma 14**. However, rather than inverting on the base step as we did in that lemma (which we cannot do here as this proof is generic with respect to the language Λ), we instead unfold the definition of the wp. The proof is trivial if e is a value. Otherwise, we open the invariant to show that e can execute for one step and update the thread pool accordingly. We can do these updates by applying **STATEUPD**. We then close the invariant with the updated state and thread pool, and use the induction hypothesis to conclude. \square

11 Related Work

Cook's Original Completeness Result. Cook [12] gave the first form of a completeness proof for a variant of Hoare logic. He noted that if one considers a fully syntactic proof system for both the assertion logic and the Hoare triple rules, then when the programming language under consideration is sufficiently expressive, any such proof system must be incomplete, since otherwise it could be used to decide the halting problem. His key insight was that one can side-step this problem and focus on whether there is any incompleteness that can be attributed to a deficiency in the Hoare logic rules, as opposed to the assertion logic rules. To do so, he examines completeness of the Hoare triple rules under the assumption that one has a complete proof system for the assertion logic, and calls this property *relative* completeness. In the case of Iris, we effectively do something similar, since Iris embeds pure assertions from the meta-logic, and adds a proof rule that reflects all implications from the meta-logic as entailments between pure assertions.

Cook next shows that for a sufficiently expressive assertion language, one can define the strongest postcondition sp (the dual of the weakest precondition) as an assertion in the logic. The main challenge lies in the rule for while loops, where one has to express a loop invariant. Once that is done, completeness follows by observing that if $\{P\} e \{Q\}$ is semantically valid, then $\vdash \{P\} e \{\text{sp}(e, P)\}$ and $\text{sp}(e, P) \Rightarrow Q$ by construction. Hence, $\text{sp}(e, P) \vdash Q$ by the assumed completeness of the assertion logic, and we have that $\vdash \{P\} e \{Q\}$ by the rule of consequence. In contrast, for our proofs, rather than *defining* a weakest precondition or strongest postcondition, we effectively end up showing that the wp provided by Iris is in fact the *weakest* precondition, in some sense.

Completeness for Concurrency Logics. Owicki and Gries [56] introduced a proof system for reasoning about concurrent programs. With their approach, to apply the parallel composition rule to $e_1 \parallel e_2$, we need not just Hoare *triples* about e_1 and e_2 , but rather full Hoare proof outlines that include each intermediate assertion in the Hoare proof between program statements. Then one checks a non-interference property that requires showing each intermediate assertion used in e_1 is *stable* under each of the possible steps of e_2 and vice-versa. Owicki [55] proved the relative completeness of this system. The idea behind her proof is to start by annotating a program with a collection of *auxiliary variables*. These variables are defined so that by inspecting their values during execution, one can determine the full history of the program's execution and interleaving order up to that point. Using these variables, she defines intermediate assertions that are stable by construction, because they can characterize the reachable states of the program. In our proof, the ghosts-tos that track the status of each thread play a similar role to Owicki's auxiliary variables.

1177 This custom ghost state makes it simpler to track the status of all of the concurrent threads, as
 1178 compared to annotating the program with variables that encode execution ordering.

1179 One drawback of the Owicki-Gries method is that it is non-modular, because the parallel compo-
 1180 sition rule requires re-inspecting the proofs of the threads and checking them against each other.
 1181 Rely-guarantee [31] addresses this by using predicates to abstract over the behaviors of threads:
 1182 we verify each thread under a *rely* assertion that describes what interference could be caused by
 1183 its environment, while showing that the thread’s behavior upholds a *guarantee* assertion. Then,
 1184 for parallel composition, we simply check for the compatibility of each thread’s rely assertion
 1185 against the other thread’s guarantee. Stirling [61] observes that rely-guarantee is complete relative
 1186 to Owicki-Gries, and since the latter is complete, so too is rely-guarantee. Xu et al. [68] give a direct
 1187 proof of relative completeness for rely-guarantee, adapting and simplifying the auxiliary variable
 1188 construction of Owicki. In Iris, rely-guarantee style reasoning is subsumed by the use of invariants
 1189 and appropriate ghost state.

1190 Most recently, de Boer and Hiep [14] establish relative completeness of a concurrent separation
 1191 logic by adapting Owicki’s technique for proving relative completeness. They define an abstract
 1192 generalized version of the idea behind Owicki’s proof, which can then be instantiated to derive
 1193 completeness of different logics. They apply this in particular to a CSL that supports dynamic thread
 1194 creation. As compared to our work, they restrict attention to a first-order language, whereas we
 1195 have considered higher-order languages with higher-order state. On the other hand, their general
 1196 theorem would apply to any logic that satisfies their abstract conditions, whereas we have focused
 1197 on logics built on the Iris framework and its logical mechanisms.

1198 **Sequential Separation Logics.** Ishtiaq and O’Hearn [30] establish the completeness of a se-
 1199 quential separation logic that uses Hoare triples by showing that the weakest precondition can
 1200 be formulated as an assertion in the logic. Yang [69] and Yang and O’Hearn [70] also establish
 1201 completeness results for the original formulation of sequential separation logic.

1202 Haslbeck and Nipkow [27] prove the relative completeness of three program logics for reasoning
 1203 about time bounds, including a separation logic with time credits. They restrict attention to a
 1204 language without a dynamic heap or allocation, so as to be able to compare to the other two
 1205 non-separation logics under consideration. Our results in §7 consider time credits in a language
 1206 with allocation.

1207 All of the prior completeness results mentioned so far restrict attention to first-order programs.
 1208 Honda et al. [29] developed a complete program logic for imperative higher-order programs using
 1209 so-called characteristic formulae that completely capture the behavior of a program. They note
 1210 that in the first-order case, such formulae are easily constructed, but that in the higher-order
 1211 case, the challenge is that the description of a higher-order function’s behavior must be somehow
 1212 parameterized by the behavior of its input. Charguéraud [10] adapted characteristic formulae to
 1213 the setting of higher-order separation logic and established relative completeness in that setting.
 1214 In our approach, we side-step the question of finding a way to characterize a program’s behavior
 1215 as a formula in the logic by instead representing the threads of the program as ghost state and
 1216 embedding pure assertions about safety in our invariants.

1218 12 Conclusion

1219 We have introduced a general pattern for completeness proofs of Iris-based program logics. For
 1220 the family of logics that are based on Iris’ shared program logic, we provide reusable lemmas that
 1221 simplify such proofs for future work. As part of this, we have identified several small gaps in existing
 1222 logics, indicating that indeed these proofs do the job that one would expect from completeness:
 1223 ensuring that the proof rules cover everything relevant about every construct in the language.

1224
 1225

1226 Interestingly, we did not have to use most of the features that make Iris so expressive: there is no
 1227 higher-order ghost state, and while we did use invariants, all our invariants are timeless, meaning
 1228 we did not have a need for general impredicative invariants. The notable exception here is Löb
 1229 induction, which *does* play an important role in our proofs. Arguably, this is not very surprising.
 1230 Features like impredicative invariants and higher-order ghost state are motivated by the desire
 1231 to give powerful, modular specifications to interesting *libraries*, while completeness talks about
 1232 proving correctness of a whole program in a single, monolithic proof. Library case studies and
 1233 proofs of completeness are orthogonal means of assessing the expressivity of a logic. Iris has been
 1234 extensively evaluated for the former; by adding the latter, our work closes a long-standing gap in
 1235 the work on Iris.

1237 References

- 1238 [1] Alejandro Aguirre, Philipp G. Haselwarter, Markus de Medeiros, Kwing Hei Li, Simon Oddershede Gregersen, Joseph
 1239 Tassarotti, and Lars Birkedal. 2024. Error Credits: Resourceful Reasoning about Error Bounds for Higher-Order
 1240 Probabilistic Programs. *Proc. ACM Program. Lang.* 8, ICFP, Article 246 (Aug. 2024), 33 pages. doi:10.1145/3674635
- 1241 [2] Clément Allain and Gabriel Scherer. 2026. Zoo: A Framework for the Verification of Concurrent OCaml 5 Programs
 1242 using Separation Logic. *Proc. ACM Program. Lang.* 10, POPL (2026), 1702–1729. doi:10.1145/3776701
- 1243 [3] Andrew W. Appel, Paul-André Mellies, Christopher D. Richards, and Jérôme Vouillon. 2007. A very modal model of a
 1244 modern, major, general type system. In *Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of
 1245 Programming Languages, POPL 2007, Nice, France, January 17–19, 2007*. 109–122. doi:10.1145/1190216.1190235
- 1246 [4] Robert Atkey. 2010. Amortised Resource Analysis with Separation Logic. In *ESOP (LNCS, Vol. 6012)*. Springer, 85–103.
 1247 doi:10.1007/978-3-642-11957-6_6
- 1248 [5] Lars Birkedal and Aleš Bizjak. 2023. Lecture Notes on Iris: Higher-Order Concurrent Separation Logic. [https://iris-
 1249 project.org/tutorial-material.html](https://iris-project.org/tutorial-material.html). <https://iris-project.org/tutorial-pdfs/iris-lecture-notes.pdf> Aarhus University.
- 1250 [6] Lars Birkedal, Rasmus Ejlers Møgelberg, Jan Schwinghammer, and Kristian Støvring. 2012. First steps in synthetic
 1251 guarded domain theory: step-indexing in the topos of trees. *Log. Methods Comput. Sci.* 8, 4 (2012). doi:10.2168/LMCS-
 1252 8(4:1)2012
- 1253 [7] John Boyland. 2003. Checking Interference with Fractional Permissions. In *SAS (LNCS, Vol. 2694)*. Springer, 55–72.
 1254 https://doi.org/10.1007/3-540-44898-5_4
- 1255 [8] Tej Chajed, Joseph Tassarotti, M. Frans Kaashoek, and Nikolai Zeldovich. 2019. Verifying concurrent, crash-safe
 1256 systems with Perennial. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP 2019,
 1257 Huntsville, ON, Canada, October 27–30, 2019*. 243–258. doi:10.1145/3341301.3359632
- 1258 [9] Tej Chajed, Joseph Tassarotti, Mark Theng, Ralf Jung, M. Frans Kaashoek, and Nikolai Zeldovich. 2021. GoJournal:
 1259 a verified, concurrent, crash-safe journaling system. In *15th USENIX Symposium on Operating Systems Design and
 1260 Implementation, OSDI 2021, July 14–16, 2021*. 423–439. <https://www.usenix.org/conference/osdi21/presentation/chajed>
- 1261 [10] Arthur Charguéraud. 2011. Characteristic formulae for the verification of imperative programs. *SIGPLAN Not.* 46, 9
 1262 (Sept. 2011), 418–430. doi:10.1145/2034574.2034828
- 1263 [11] Krzysztof Ciesielski. 2007. On Stefan Banach and some of his results. *Banach Journal of Mathematical Analysis* 1, 1
 1264 (2007), 1–10. doi:10.15352/bjma/1240321550
- 1265 [12] Stephen A. Cook. 1978. Soundness and Completeness of an Axiom System for Program Verification. *SIAM J. Comput.*
 1266 7, 1 (1978), 70–90. doi:10.1137/0207005
- 1267 [13] Hoang-Hai Dang, Jacques-Henri Jourdan, Jan-Oliver Kaiser, and Derek Dreyer. 2020. RustBelt meets relaxed memory.
 1268 *Proc. ACM Program. Lang.* 4, POPL (2020), 34:1–34:29. doi:10.1145/3371102
- 1269 [14] Frank S. de Boer and Hans-Dieter A. Hiep. 2025. Beyond Concurrent Separation Logic: Who is Afraid of Completeness
 1270 Proofs?. In *Principles of Formal Quantitative Analysis - Essays Dedicated to Christel Baier on the Occasion of Her 60th
 1271 Birthday (LNCS, Vol. 15760)*, Nathalie Bertrand, Clemens Dubslaff, and Sascha Klüppelholz (Eds.). Springer, 301–319.
 1272 doi:10.1007/978-3-031-97439-7_15
- 1273 [15] Paulo Emilio de Vilhena and François Pottier. 2023. A Type System for Effect Handlers and Dynamic Labels. In
 1274 *Programming Languages and Systems - 32nd European Symposium on Programming, ESOP 2023, Held as Part of the
 European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22–27, 2023, Proceedings*.
 225–252. doi:10.1007/978-3-031-30044-8_9
- [16] Lennard Gäher, Michael Sammler, Ralf Jung, Robbert Krebbers, and Derek Dreyer. 2024. RefinedRust: A Type
 System for High-Assurance Verification of Rust Programs. *Proc. ACM Program. Lang.* 8, PLDI (2024), 1115–1139.
 doi:10.1145/3656422

- 1275 [17] Lennard Gäher, Michael Sammler, Simon Spies, Ralf Jung, Hoang-Hai Dang, Robbert Krebbers, Jeehoon Kang, and
1276 Derek Dreyer. 2022. Simuliris: a Separation Logic Framework for Verifying Concurrent Program Optimizations. *Proceedings of the ACM on Programming Languages* 6, POPL (2022), 1–31.
- 1277 [18] Aina Linn Georges, Benjamin Peters, Laila Elbeheiry, Leo White, Stephen Dolan, Richard A. Eisenberg, Chris Casinghino,
1278 François Pottier, and Derek Dreyer. 2025. Data Race Freedom à la Mode. *Proc. ACM Program. Lang.* 9, POPL (2025),
1279 656–686. <https://doi.org/10.1145/3704859>
- 1280 [19] Paolo G. Giarrusso, Léo Stefanescu, Amin Timany, Lars Birkedal, and Robbert Krebbers. 2020. Scala step-by-step:
1281 soundness for DOT with step-indexed logical relations in Iris. *Proc. ACM Program. Lang.* 4, ICFP (2020), 114:1–114:29.
1282 [doi:10.1145/3408996](https://doi.org/10.1145/3408996)
- 1283 [20] Simon Oddershede Gregersen, Chaitanya Agarwal, and Joseph Tassarotti. 2025. Logical Relations for Formally Verified
1284 Authenticated Data Structures. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications
1285 Security, CCS 2025, Taipei, Taiwan, October 13–17, 2025*. 1394–1408. [doi:10.1145/3719027.3744801](https://doi.org/10.1145/3719027.3744801)
- 1286 [21] Simon Oddershede Gregersen, Alejandro Aguirre, Philipp G. Haselwarter, Joseph Tassarotti, and Lars Birkedal. 2024.
1287 Almost-Sure Termination by Guarded Refinement. *Proc. ACM Program. Lang.* 8, ICFP (2024), 203–233. [doi:10.1145/
3674632](https://doi.org/10.1145/3674632)
- 1288 [22] Simon Oddershede Gregersen, Alejandro Aguirre, Philipp G. Haselwarter, Joseph Tassarotti, and Lars Birkedal. 2024.
1289 Asynchronous Probabilistic Couplings in Higher-Order Separation Logic. *Proc. ACM Program. Lang.* 8, POPL (2024),
753–784. [doi:10.1145/3632868](https://doi.org/10.1145/3632868)
- 1290 [23] Simon Oddershede Gregersen, Johan Bay, Amin Timany, and Lars Birkedal. 2021. Mechanized logical relations for
1291 termination-insensitive noninterference. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–29. [doi:10.1145/3434291](https://doi.org/10.1145/3434291)
- 1292 [24] Armaël Guéneau, Johannes Hostert, Simon Spies, Michael Sammler, Lars Birkedal, and Derek Dreyer. 2023. Melocoton:
1293 A Program Logic for Verified Interoperability Between OCaml and C. *Proc. ACM Program. Lang.* 7, OOPSLA2 (2023),
1294 716–744. <https://doi.org/10.1145/3622823>
- 1295 [25] Philipp G. Haselwarter, Kwing Hei Li, Alejandro Aguirre, Simon Oddershede Gregersen, Joseph Tassarotti, and Lars
1296 Birkedal. 2025. Approximate Relational Reasoning for Higher-Order Probabilistic Programs. *Proc. ACM Program. Lang.*
9, POPL (2025), 1196–1226. [doi:10.1145/3704877](https://doi.org/10.1145/3704877)
- 1297 [26] Philipp G. Haselwarter, Kwing Hei Li, Markus de Medeiros, Simon Oddershede Gregersen, Alejandro Aguirre, Joseph
1298 Tassarotti, and Lars Birkedal. 2024. Tachis: Higher-Order Separation Logic with Credits for Expected Costs. *Proc.
1299 ACM Program. Lang.* 8, OOPSLA2, Article 313 (Oct. 2024), 30 pages. [doi:10.1145/3689753](https://doi.org/10.1145/3689753)
- 1300 [27] Maximilian P. L. Haslbeck and Tobias Nipkow. 2018. Hoare Logics for Time Bounds - A Study in Meta Theory. 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 (LNCS, Vol. 10805)*, Dirk Beyer and Marieke Huisman (Eds.). Springer, 155–171. [doi:10.1007/978-3-319-89960-2_9](https://doi.org/10.1007/978-3-319-89960-2_9)
- 1301 [28] Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2020. Actris: session-type based reasoning in
1302 separation logic. *Proc. ACM Program. Lang.* 4, POPL (2020), 6:1–6:30. [doi:10.1145/3371074](https://doi.org/10.1145/3371074)
- 1303 [29] Kohei Honda, Nobuko Yoshida, and Martin Berger. 2005. An Observationally Complete Program Logic for Imperative
1304 Higher-Order Functions. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26–29 June 2005, Chicago, IL, USA, Proceedings*. IEEE Computer Society, 270–279. [doi:10.1109/LICS.2005.5](https://doi.org/10.1109/LICS.2005.5)
- 1305 [30] Samin S. Ishtiaq and Peter W. O’Hearn. 2001. BI as an Assertion Language for Mutable Data Structures. In *Conference
1306 Record of POPL 2001: The 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, London, UK, January 17–19, 2001*. 14–26. [doi:10.1145/360204.375719](https://doi.org/10.1145/360204.375719)
- 1307 [31] Cliff B. Jones. 1983. Tentative Steps Toward a Development Method for Interfering Programs. *ACM Trans. Program.
1308 Lang. Syst.* 5, 4 (1983), 596–619. [doi:10.1145/69575.69577](https://doi.org/10.1145/69575.69577)
- 1309 [32] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2017. RustBelt: Securing the Foundations of
1310 the Rust Programming Language. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 1–34.
- 1311 [33] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from
1312 the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20.
1313 <https://doi.org/10.1017/S0956796818000151>
- 1314 [34] Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy, Marianna Rapoport, Amin Timany, Derek Dreyer, and Bart Jacobs.
1315 2019. The future is ours: prophecy variables in separation logic. *Proc. ACM Program. Lang.* 4, POPL, Article 45 (dec
1316 2019), 32 pages. [doi:10.1145/3371113](https://doi.org/10.1145/3371113)
- 1317 [35] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015.
1318 Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. ACM, 637–650. <https://doi.org/10.1145/2676726.2676980>
- 1319
1320
1321
1322
1323

- 1324 [36] Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. 2017. Strong Logic for Weak
 1325 Memory: Reasoning About Release-Acquire Consistency in Iris. In *31st European Conference on Object-Oriented*
 1326 *Programming, ECOOP 2017, Barcelona, Spain, June 19-23, 2017*. 17:1–17:29. doi:10.4230/LIPICS.ECOOP.2017.17
- 1327 [37] Robbert Krebbers, Jacques-Henri Jourdan, Ralf Jung, Joseph Tassarotti, Jan-Oliver Kaiser, Amin Timany, Arthur
 1328 Charguéraud, and Derek Dreyer. 2018. MoSeL: A general, extensible modal framework for interactive proofs in
 1329 separation logic. *PACMPL* 2, ICFP (2018), 77:1–77:30. <https://doi.org/10.1145/3236772>
- 1330 [38] Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017. The Essence of
 1331 Higher-Order Concurrent Separation Logic. In *ESOP (LNCS, Vol. 10201)*. Springer, 696–723. https://doi.org/10.1007/978-3-662-54434-1_26
- 1332 [39] Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017. Interactive proofs in higher-order concurrent separation
 1333 logic. In *POPL*. 205–217. doi:10.1145/3009837.3009855
- 1334 [40] Robbert Krebbers, Luko van der Maas, and Enrico Tassi. 2025. Inductive Predicates via Least Fixpoints in Higher-Order
 1335 Separation Logic. In *16th International Conference on Interactive Theorem Proving (ITP 2025) (Leibniz International*
 1336 *Proceedings in Informatics (LIPIcs), Vol. 352)*, Yannick Forster and Chantal Keller (Eds.). Schloss Dagstuhl – Leibniz-
 1337 Zentrum für Informatik, Dagstuhl, Germany, 27:1–27:21. doi:10.4230/LIPICS.ITP.2025.27
- 1338 [41] Morten Krogh-Jespersen, Amin Timany, Marit Edna Ohlenbusch, Simon Oddershede Gregersen, and Lars Birkedal.
 1339 2020. Aneris: A Mechanised Logic for Modular Reasoning about Distributed Systems. In *Programming Languages and*
 1340 *Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the European Joint Conferences on*
 1341 *Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings*. 336–365. doi:10.1007/978-3-030-44914-8_13
- 1342 [42] Peter John Landin. 1964. The Mechanical Evaluation of Expressions. *Comput. J.* (1964).
- 1343 [43] Xavier Leroy, Andrew Appel, Sandrine Blazy, and Gordon Stewart. 2012. *The CompCert memory model, version 2*.
 1344 Technical Report RR-7987. Inria. <https://hal.inria.fr/hal-00703441>
- 1345 [44] Kwing Hei Li, Alejandro Aguirre, Simon Oddershede Gregersen, Philipp G. Haselwarter, Joseph Tassarotti, and Lars
 1346 Birkedal. 2025. Modular Reasoning about Error Bounds for Concurrent Probabilistic Programs. *Proc. ACM Program.*
 1347 *Lang.* 9, ICFP, Article 245 (Aug. 2025), 30 pages. doi:10.1145/3747514
- 1348 [45] Martin Hugo Löb. 1955. Solution of a problem of Leon Henkin. *Journal of Symbolic Logic* 20, 2 (1955), 115–118.
 1349 doi:10.2307/2266895
- 1350 [46] Rupak Majumdar and V. R. Sathiyarayanan. 2025. Sound and Complete Proof Rules for Probabilistic Termination.
 1351 *Proc. ACM Program. Lang.* 9, POPL (2025), 1871–1902. doi:10.1145/3704899
- 1352 [47] William Mansky and Ke Du. 2024. An Iris Instance for Verifying CompCert C Programs. *Proc. ACM Program. Lang.* 8,
 1353 POPL (2024), 148–174. doi:10.1145/3632848
- 1354 [48] Yusuke Matsushita, Xavier Denis, Jacques-Henri Jourdan, and Derek Dreyer. 2022. RustHornBelt: a semantic foundation
 1355 for functional verification of Rust programs with unsafe code. In *PLDI '22: 43rd ACM SIGPLAN International Conference*
 1356 *on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*. 841–856. doi:10.1145/3519939.3523704
- 1357 [49] Annabelle McIver, Carroll Morgan, Benjamin Lucien Kaminski, and Joost-Pieter Katoen. 2018. A new proof rule for
 1358 almost-sure termination. *Proc. ACM Program. Lang.* 2, POPL (2018), 33:1–33:28. doi:10.1145/3158121
- 1359 [50] Glen Mével and Jacques-Henri Jourdan. 2021. Formal verification of a concurrent bounded queue in a weak memory
 1360 model. *Proc. ACM Program. Lang.* 5, ICFP (2021), 1–29. doi:10.1145/3473571
- 1361 [51] Glen Mével, Jacques-Henri Jourdan, and François Pottier. 2019. Time Credits and Time Receipts in Iris. In *Programming*
 1362 *Languages and Systems, Luís Caires (Ed.)*. Springer International Publishing, Cham, 3–29. doi:10.1007/978-3-030-17184-1_1
- 1363 [52] Alexandre Moine, Stephanie Balzer, Alex Xu, and Sam Westrick. 2026. TypeDis: A Type System for Disentanglement.
 1364 *Proc. ACM Program. Lang.* 10, POPL (2026), 354–383. doi:10.1145/3776655
- 1365 [53] Alexandre Moine, Arthur Charguéraud, and François Pottier. 2023. A High-Level Separation Logic for Heap Space
 1366 under Garbage Collection. *Proc. ACM Program. Lang.* 7, POPL (2023), 718–747. doi:10.1145/3571218
- 1367 [54] Hiroshi Nakano. 2000. A Modality for Recursion. In *15th Annual IEEE Symposium on Logic in Computer Science, Santa*
 1368 *Barbara, California, USA, June 26-29, 2000*. 255–266. doi:10.1109/LICS.2000.855774
- 1369 [55] Susan S. Owicki. 1975. *Axiomatic Proof Techniques for Parallel Programs*. Technical Report. USA.
- 1370 [56] Susan S. Owicki and David Gries. 1976. An Axiomatic Proof Technique for Parallel Programs I. *Acta Informatica* 6
 1371 (1976), 319–340. doi:10.1007/BF00268134
- 1372 [57] François Pottier, Armaël Guéneau, Jacques-Henri Jourdan, and Glen Mével. 2024. Thunks and Debits in Separation
 Logic with Time Credits. *Proc. ACM Program. Lang.* 8, POPL (2024), 1482–1508. doi:10.1145/3632892
- [58] Michael Sammler, Rodolphe Lepigre, Robbert Krebbers, Kayvan Memarian, Derek Dreyer, and Deepak Garg. 2021. RefinedC: Automating the Foundational Verification of C Code with Refined Ownership Types. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 158–174.

- 1373 [59] Remy Seassau, Irene Yoon, Jean-Marie Madiot, and François Pottier. 2025. Formal Semantics and Program Logics for a
1374 Fragment of OCaml. *Proc. ACM Program. Lang.* 9, ICFP (2025), 128–159. doi:10.1145/3747509
- 1375 [60] Upamanyu Sharma, Ralf Jung, Joseph Tassarotti, M. Frans Kaashoek, and Nikolai Zeldovich. 2023. Grove: a Separation-
1376 Logic Library for Verifying Distributed Systems. In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP 2023, Koblenz, Germany, October 23-26, 2023*. 113–129. doi:10.1145/3600006.3613172
- 1377 [61] Colin Stirling. 1988. A Generalization of Owicki-Gries’s Hoare Logic for a Concurrent while Language. *Theor. Comput.*
1378 *Sci.* 58 (1988), 347–359. doi:10.1016/0304-3975(88)90033-3
- 1379 [62] Kasper Svendsen and Lars Birkedal. 2014. Impredicative Concurrent Abstract Predicates. In *ESOP (LNCS, Vol. 8410)*.
1380 Springer, 149–168. https://doi.org/10.1007/978-3-642-54833-8_9
- 1381 [63] The Rocq Team. 2026. The Rocq Prover. <https://rocq-prover.org/>.
- 1382 [64] Amin Timany, Simon Oddershede Gregersen, Léo Stefanescu, Jonas Kastberg Hinrichsen, Léon Gondelman, Abel
1383 Nieto, and Lars Birkedal. 2024. Trillium: Higher-Order Concurrent and Distributed Separation Logic for Intensional
1384 Refinement. *Proc. ACM Program. Lang.* 8, POPL (2024), 241–272. doi:10.1145/3632851
- 1385 [65] Amin Timany, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2024. A Logical Approach to Type Soundness. *J.*
1386 *ACM* 71, 6 (2024), 40:1–40:75. doi:10.1145/3676954
- 1387 [66] Aaron Turon, Derek Dreyer, and Lars Birkedal. 2013. Unifying refinement and hoare-style reasoning in a logic for
1388 higher-order concurrency. In *ICFP*. ACM, 377–390. doi:10.1145/2500365.2500600
- 1389 [67] Orpheas van Rooij and Robbert Krebbers. 2025. Affect: An Affine Type and Effect System. *Proc. ACM Program. Lang.* 9,
1390 POPL (2025), 126–154. doi:10.1145/3704841
- 1391 [68] Qiwen Xu, Willem P. de Roeper, and Jifeng He. 1997. The Rely-Guarantee Method for Verifying Shared Variable
1392 Concurrent Programs. *Formal Aspects Comput.* 9, 2 (1997), 149–174. doi:10.1007/BF01211617
- 1393 [69] Hongseok Yang. 2001. *Local reasoning for stateful programs*. Ph.D. Dissertation. USA. Advisor(s) Reddy, Uday S.
1394 AAI3023240.
- 1395 [70] Hongseok Yang and Peter W. O’Hearn. 2002. A Semantic Basis for Local Reasoning. In *Proceedings of the 5th*
1396 *International Conference on Foundations of Software Science and Computation Structures (FoSSaCS ’02)*. Springer-Verlag,
1397 Berlin, Heidelberg, 402–416.

1398 Received 2026-02-19

1399

1400

1401

1402

1403

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422 A The Strong Invariant Opening Lemma

1423 We provide the following lemma which allows us to open invariants and only then decide if we
 1424 want to keep them open or not, in particular it allows delaying the proof that some expression is
 1425 atomic until the invariant has been opened.

1427 **Lemma 31.** *If $\mathcal{N} \subseteq \mathcal{E}$, and if e is reducible under some σ , then it holds that*

$$\begin{aligned}
 & \boxed{P}^{\mathcal{N}*} \\
 & (P \Rightarrow_{\mathcal{E} \setminus \mathcal{N}}^* \\
 & \quad (\exists K, e'. \ulcorner e = K[e'] \urcorner * \ulcorner \text{Atomic } e' \urcorner * (\exists \sigma'. \ulcorner \text{base-red } (e', \sigma') \urcorner)) * \\
 & \quad \text{wp}_{\mathcal{E} \setminus \mathcal{N}} e' \{v. P * \text{wp}_{\mathcal{E}} K[v] \{\Psi\}\}) \vee \\
 & \quad P * \text{wp}_{\mathcal{E}} e \{\Psi\}) \\
 & \vdash \text{wp}_{\mathcal{E}} e \{\Psi\}
 \end{aligned}$$

1437 **PROOF.** Since we know that e is reducible, we have $e = K[e']$ and that e' is base-reducible. Case
 1438 distinction on if e' is atomic:

- 1439 • If it is, then we apply **WPBIND** followed by **WPATOMICINV**. We can assume P , specialize the
 1440 main assumption with it, and also eliminate the update there to expose the big disjunction.
 1441 We now do a case distinction on it.
 - 1442 – The first case is easy, as we get K', e'' such that $e = K'[e'']$ and also e'' is base-reducible.
 1443 Since the base redex is unique¹⁹, we have $K = K'$ and $e'' = e'$. Our goal is then precisely
 1444 equal to our assumption.
 - 1445 – In the second case, we frame P into the postcondition then and apply **WPBIND** backwards
 1446 to conclude.
- 1447 • If it is not, we use **WPUPDATEELIM** and also do a case distinction on our assumption.
 - 1448 – In the first case, we again get that $e' = e''$ for an atomic e'' , which is a contradiction.
 - 1449 – In the second case, we are done by assumption. □

1451 Note that this lemma is proven using classical reasoning, to decide if e' is atomic. This is usually
 1452 decidable for a specific language, but we did not want to add this as an extra assumption to our
 1453 recipe. The lemma can also be proven using the definition of wp (see **Definition 28** in §10) without
 1454 using classical logic, but since our proofs treat wp axiomatically, this would mean that this lemma
 1455 becomes an additional axiom.

1456 Further, when using this lemma in the proof of **Lemma 14** or **Theorem 16**, we did not have the
 1457 assumption that e is reducible yet. But in these proofs, it is easily obtained by an additional opening
 1458 and closing of the invariant at the start.

1460 B Stronger Sequential Versions of Completeness and Requisitesness Theorems

1461 The wp derived by **Theorem 16** takes the state invariant $S_{\circ}(\sigma)$ in the precondition but only gives
 1462 back a meta-level postcondition $\varphi(v)$. It turns out that if we restrict the expression e to be sequential,
 1463 we can derive a stronger version of completeness that gives back a S_{\circ} for the final state.

1464 To keep the result language-agnostic, rather than restrict the syntax of the language (like
 1465 **Theorem 23**), we define sequentiality as a meta-level semantic property saying that the thread pool
 1466 only ever contains 1 thread:

$$\text{seq}(e, \sigma) \triangleq \forall \vec{e}', \sigma'. ([e], \sigma) \rightarrow_{\text{tp}}^* (\vec{e}', \sigma') \Rightarrow |\vec{e}'| = 1$$

1469 ¹⁹This base redex uniqueness follows from the general laws that must hold for any language used to instantiate Iris's wp .

We also extend the notion of *safe* to account for the final state.

$$\begin{aligned} \text{ssafe-tp}_\phi(\vec{e}, \sigma) &\triangleq \forall \vec{e}', \sigma'. (\vec{e}, \sigma) \rightarrow_{\text{tp}}^* (\vec{e}', \sigma') \implies \\ &\quad \forall n, e''. \vec{e}'[n] = e'' \implies (\exists v. e'' = v \wedge (n = 0 \implies \phi(v, \sigma'))) \vee \text{red}(e', \sigma') \\ \text{ssafe}_\phi(e, \sigma) &\triangleq \text{ssafe-tp}_\phi([e], \sigma) \end{aligned}$$

The sequential version of completeness is then stated as following.

Theorem 32 (Generic Sequential Completeness). *Let Λ be a language satisfying the [Condition 17](#). If $\text{ssafe}_\phi(e, \sigma) \wedge \text{seq}(e, \sigma)$, then $\mathbb{S}_\circ(\sigma) \vdash \text{wp}_\top e \{v. \exists \sigma'. \mathbb{S}_\circ(\sigma') * \ulcorner \phi(v, \sigma') \urcorner\}$.*

PROOF SKETCH. Similar to [§2.2](#), but directly do Löb induction on the theorem statement. \square

Similarly, we also have a sequential version of requisiteness.

Theorem 33 (Generic Sequential Requisiteness). *For all γ , let \mathbb{S}_\circ^γ be an Iris predicate on states satisfying [STATEUPD](#), then for any meta-level predicate ϕ , $\text{ssafe}_\phi(e, \sigma) \wedge \text{seq}(e, \sigma)$ implies $\mathbb{S}_\circ^\gamma(\sigma) \vdash \text{wp}_\top e \{v. \exists \sigma'. \mathbb{S}_\circ(\sigma') * \ulcorner \phi(v, \sigma') \urcorner\}$.*

The significance of [Theorem 32](#) and [Theorem 33](#) is that they unlock the reuse of proofs across program logics. With sequential completeness/requisiteness, we can now verify subprograms using external tools and formally embed the resulting safety theorems into Iris.

Note that the sequentiality here is used in an essential way. When the main thread terminates, it must frame a fragment of the final state, σ' , into $\mathbb{S}_\circ(\sigma')$ to conclude its postcondition. This prevents other threads from accessing σ' since this point. If we allowed e to fork child threads, then we would have to show child threads can still make progress after the main thread terminates, without accessing σ' , which is impossible because safety permits a child thread to rely on σ' to make progress.

C The Exact Form of Generic Total Completeness Theorem

Condition 34 (Total Completeness).

$$\begin{aligned} \ulcorner \text{base-red}(e, \sigma) \urcorner * \mathbb{S}_\circ(\sigma) \vdash \text{E} \left(\ulcorner \text{Atomic } e \urcorner * \forall \Phi. \text{TAtomicPre}(\mathcal{E}, e, \sigma, \Phi) \multimap \text{wp}_\mathcal{E} e [\Phi] \right) \vee \\ (\mathbb{S}_\circ(\sigma) * \forall \Phi. \text{TNonAtomicPre}(\mathcal{E}, e, \Phi) \multimap \text{wp}_\top e [\Phi]) \end{aligned}$$

where

$$\text{TAtomicPre}(\mathcal{E}, e, \sigma, \Phi) \triangleq \forall v', \sigma', \vec{e}_f. \ulcorner (e, \sigma) \rightarrow (v', \sigma', \vec{e}_f) \urcorner * \mathbb{S}_\circ(\sigma') \multimap_{\mathcal{E}}$$

$$\Phi(v') * \multimap_{e_f \in \vec{e}_f} \text{wp}_\top e_f [v. \text{True}]$$

$$\text{TNonAtomicPre}(\mathcal{E}, e, \Phi) \triangleq \forall e', \vec{e}_f. (\forall \sigma. \mathbb{S}_\circ(\sigma) \multimap_{\mathcal{E}} \exists \sigma'. \ulcorner (e, \sigma) \rightarrow^+ (e', \sigma', \vec{e}_f) \urcorner * \mathbb{S}_\circ(\sigma')) \multimap_{\top}$$

$$\text{wp}_\top e' [\Phi] * \multimap_{e_f \in \vec{e}_f} \text{wp}_\top e_f [v. \text{True}]$$

Theorem 19 (Generic Total Completeness). *Let Λ be a language that satisfies [Condition 34](#). If $\text{safe}_\phi(e, \sigma) \wedge \text{SN}_{\rightarrow_{\text{tp}}}([e], \sigma)$, then $\mathbb{S}_\circ(\sigma) \vdash \text{wp}_\top e [v. \ulcorner \phi(v) \urcorner]$.*