

1 Verifying wait-freedom for concurrent higher-order 2 programs

3 **Anonymous author**

4 Anonymous affiliation

5 **Anonymous author**

6 Anonymous affiliation

7 **Anonymous author**

8 Anonymous affiliation

9 — Abstract —

10 Wait-freedom is the strongest non-blocking progress guarantee for concurrent data structures,
11 ensuring that every operation completes in a finite number of steps regardless of interference from
12 other threads. While verification of wait-freedom has been studied for first-order languages, verifying
13 it for higher-order programming languages with general references remains an open challenge. In
14 such languages, operations may be used by arbitrary, unverified higher-order clients, making it
15 unclear how to even define wait-freedom formally in terms of programs' semantics, let alone prove it.

16 In this paper, we present the first framework for verifying wait-freedom of concurrent programs
17 written in a higher-order language with general references. Our approach is based on the Lawyer
18 concurrent separation logic which has been recently introduced for termination verification. We
19 identify a specification pattern in the Lawyer logic that captures wait-freedom. To establish this
20 connection formally, we extend Lawyer with a novel adequacy theorem that proves that programs
21 which are proven correct in the Lawyer logic against a specification in this aforementioned specification
22 pattern are wait-free. Proving wait-freedom requires us to show that all calls made to operations by
23 any arbitrary client terminate. Thus, as part of formally proving wait-freedom, *i.e.*, as part of the
24 proof of the adequacy theorem above, we need to prove that the behavior of the client of the data
25 structure is safe in the sense that it does not break the internal invariants of the data structure,
26 *e.g.*, by directly manipulating the data structure's internal state. To this end, we develop a logical
27 relations model that establishes safety for all clients once and for all.

28 We demonstrate the effectiveness of our approach by proving wait-freedom for several representa-
29 tive examples, including higher-order functions such as the list map function, and a memory-efficient
30 single-producer, single-consumer queue. For the latter, wait-freedom is conditional in that as the
31 name suggests there can be at most one enqueueer thread and one dequeuer thread. To capture
32 this formally we introduce the notion of restricted wait-freedom as a variant of wait-freedom that
33 restricts the number of concurrent threads, and show how our approach can support reasoning about
34 restricted wait-freedom. All our results have been mechanized on top of the Rocq Prover and using
35 the Iris separation logic framework that Lawyer is also based on.

36 **2012 ACM Subject Classification** Author: Please fill in 1 or more `\ccsdsc macro`

37 **Keywords and phrases** separation logic, refinement, higher-order logic, concurrency, formal verifica-
38 tion

39 **Digital Object Identifier** 10.4230/LIPIcs...

40 **1** Introduction

41 Any implementation of a concurrent data structure must account for possible interference
42 between multiple operations on that data structure running concurrently. Done wrongly, it
43 might lead to incorrect behavior of operations, or even worse, to undefined behavior of an
44 entire program. The simplest way to ensure concurrent operations' correctness is to resort
45 to *blocking concurrency*, *e.g.*, using a lock, allowing only one operation to progress at any



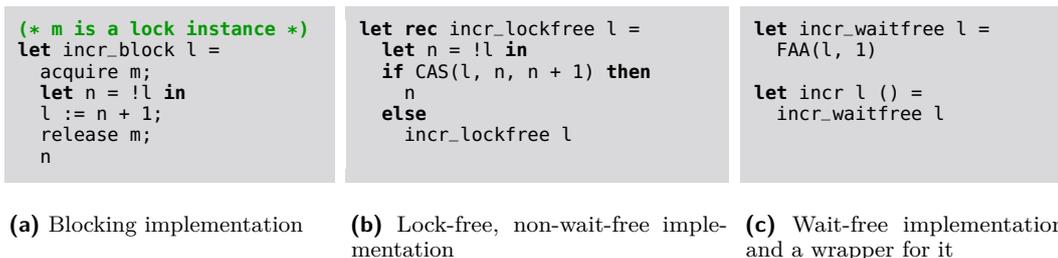
© Anonymous author(s);

licensed under Creative Commons License CC-BY 4.0

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

XX:2 Verifying wait-freedom for concurrent higher-order programs



■ **Figure 1** Comparison of counter increment implementations; we use an OCaml-like syntax. Location l stores an integer

46 given moment, and forcing others to wait for its completion. However, in many settings,
47 *e.g.*, in operating system kernels [5], blocking concurrency is discouraged as it might lead
48 to poor performance, and due to risk of deadlocking the entire system. Therefore, in such
49 cases a *lock-free* implementation is more desirable. Lock-free data structure implementations
50 guarantee that among the concurrently running operations acting on the database, at least
51 one of them must run to completion. Yet, with lock-freedom does not guarantee that there is
52 no starvation. That is, given a specific call to a lock-free data structure’s operation, we cannot
53 guarantee that that call must terminate. To avoid starvation, one must use so-called *wait-free*
54 [17] implementations. Wait-freedom is the strongest form of non-blocking concurrency. It
55 guarantees that *every* individual call to the data structure operations must run to completion
56 regardless of whether or not, or how many other operations of the same data structure are
57 run concurrently with it.

58 We illustrate the differences between blocking, lock-free, and wait-free concurrency using
59 and simple and contrived example: a concurrent counter implementation featuring an
60 increment operation, and a read operation (which we elide). Figure 1 shows multiple possible
61 implementations of the “increment” operation. The implementation in Figure 1a allows
62 only one thread at a time to access the counter, forcing all others to wait until the lock is
63 released. The starvation behavior of this implementation of the concurrent counter data
64 structure depends on whether the lock used is a fair lock or not. The lock-free implementation
65 (Figure 1b), implemented via an atomic Compare-And-Set (CAS) operation in a loop, does
66 not wait for other concurrent operations on the counter to complete. It attempts to increment
67 using the CAS operation, and if that fails, it retries. Note that if the CAS operation fails,
68 that means that the value of the counter (stored in memory location l) must have changed
69 since it was read on the line above the CAS operation. This in turn means that another
70 increment operation must have succeeded in another thread. Thus, this implementation is
71 lock-free. However, this lock-free implementation can starve threads, even under some fair
72 schedulers. (Even though unlikely, a pathological but fair scheduler could schedule the thread,
73 and subsequently preempt it right before the CAS operation, only to return to it, every time,
74 after another increment has succeeded.) Finally, we have the wait-free implementation in
75 Figure 1c. It uses the atomic Fetch-And-Add (FAA) operation which atomically increments
76 the counter. Thus, the increment always finishes in a single, atomic step. In our discussions
77 throughout the rest of the paper we will refer to this wait-free implementation of increment.
78 To this end, we introduce a wrapper function `incr`, so that we can treat `incr l` as a function
79 (of type `unit to unit`).

80 Of course, in this contrived example of a simple concurrent counter (Figure 1c) we can
81 relegate the problem of making the data structure wait-free to the hardware using the FAA

```

(* The function loop should not make any
   recursive calls. *)
let simple_op () =
  (* Could be allocated at address 0x42. *)
  let b = ref true in
  let rec loop () =
    if !b then
      ()
    else
      loop ()
  in loop ()

(* Parallel composition is written as || *)
let good_client =
  let r := ref simple_op in
  simple_op () || !r ()

(* Direct memory access is written as [.] *)
let bad_client =
  [0x42] := false || simple_op ()

```

■ **Figure 2** An example of a wait-free operation together with two possible clients

82 operation. In reality, as we will see below, for data structures ever so slightly more complex
 83 than this wait-free implementations are far from trivial to implement and to (formally) reason
 84 about. Indeed, to ensure wait-freedom, operations must be designed to execute concurrently
 85 in a way that neither do they have to wait for the completion of other operations, nor are
 86 they impeded in a way by the other operations that forces them to retry. Therefore, a
 87 wait-free implementation of a data structure is often much more complicated than a lock-free
 88 or blocking implementation of the same data structure (see *e.g.* wait-free queue of Kogan
 89 and Petrank [26] versus the lock-free queue of Michael and Scott [30]). Thus, it is often also
 90 more challenging to verify.

91 In this paper, we show how to use a higher-order concurrent program logic to verify
 92 wait-freedom for programs written in a concurrent *higher-order* programming language with
 93 general references. Such languages achieve greater modularity by supporting features such
 94 as closures and first-class modules. Our work goes significantly beyond the state of the
 95 art [9, 28] which has only considered wait-freedom of programs in first-order programming
 96 languages without general references. Working with higher-order programs, and with general
 97 references, introduces many new challenges that one needs to answer. In fact, it is not even
 98 clear a priori how to formally define wait-freedom in our setting. Informally, it is often
 99 explained that a wait-free operation should be able to make progress in any program context
 100 under any scheduler. However, it is not entirely clear here what one means by “any program
 101 context”, in particular in our higher-order setting. A “context” (or *client*) of an operation can
 102 be any program that calls that operation, either directly or indirectly (through a reference),
 103 possibly concurrently—see *e.g.*, the `good_client` in Figure 2. However, it does not make
 104 sense to consider literally any client. In particular, the client’s behavior should be ‘safe’ in
 105 that it should not be able to violate data encapsulation of the wait-free data structure, *i.e.*, it
 106 should not be able to directly accessing data structure’s internal state. For a simple example
 107 of what can go wrong see the `bad_client` in Figure 2 which directly writes to `0x42` while
 108 concurrently running `simple_op`. The example `bad_client` could potentially violate an
 109 invariant of `simple_op` (namely the invariant that the local state variable `b` is always true)
 110 and thus break the wait-freedom of `simple_op`.

111 The example `bad_client` in Figure 2 shows that the progress guarantee of a wait-free
 112 operation crucially relies on its client being safe. Indeed, prior works on wait-freedom
 113 verification for first-order programs enforce safety of clients one way or another, but their
 114 approaches do not apply in our setting. One way to enforce clients’ safety is by using a
 115 programming language that syntactically separates accesses to the operation’s and client’s
 116 states [28]. This approach, however, is not viable for the higher-order language we consider.
 117 (In general, the introduction of higher-order store enables behaviors not present in a first-order
 118 setting, including turning a terminating program into non-terminating, as demonstrated

XX:4 Verifying wait-freedom for concurrent higher-order programs

119 by [10]). Another approach to ensuring safety of clients is by requiring them to be verified
120 [9, 32]. In addition to not being applicable in our setting (explained below), this approach
121 also inherently limits the applicability of the wait-freedom result: now every single client
122 of potential interest must be verified in the corresponding framework. Rather, we take
123 an approach that does not require one to verify individual clients but instead establishes
124 operations' progress within arbitrary clients. This also includes those clients out of reach of
125 existing solutions, *e.g.*, higher-order clients.

126 In this work, we present the first solution to the verification of wait-freedom for a realistic,
127 higher-order language. Our solution is based on the recent Lawyer logic [31], an extension of
128 the Trillium [35] and Iris [25] frameworks. So far, Lawyer has only been used for establishing
129 (fair) termination of closed programs. The key insight of our work is that a specific form of
130 specifications given to a function forces that the function to terminate *in any context, without*
131 *either blocking or starving*. That is, such Lawyer specifications essentially, for all intents
132 and purposes codify wait-freedom in the Lawyer program logic. However, there are two
133 main technical issues that prevent the Lawyer logic from directly being applicable to proving
134 wait-freedom. Firstly, the so-called adequacy theorem of Lawyer is not applicable. The
135 adequacy applies to closed programs, *i.e.*, not a single data structure. Thus, the adequacy
136 theorem must be applied to the closed program, *i.e.*, the data structure and its client.
137 Moreover, the adequacy theorem states that when a program is proven correct in Lawyer,
138 then the entire program, meaning all threads must terminate. Hence, in this work we state
139 and prove a separate and novel adequacy theorem specifically for wait-freedom. Secondly, the
140 Lawyer logic allows only composition of two programs as long as they are both verified in the
141 Lawyer logic. Thus, if we were to only restrict ourselves to verified clients, we would also be
142 restricted to only considering terminating clients. On the other hand, as we discussed above
143 wait-freedom only makes sense when the client respects state encapsulation. Therefore,
144 in our approach to verification of wait-freedom we prove once and for all that *all programs*
145 *in our programming language respect state encapsulation*. For this purpose we construct a
146 so-called logical relations model for our programming language. This is similar to how prior
147 works on proving robust safety of programs [34, 13] have used logical relations models.

148 We use the approach outlined above for proving wait-freedom to prove wait-freedom
149 of a number of example programs written in our higher-order language, including those
150 that highlight these higher-order programming features. In particular, we modularly verify
151 wait-freedom of the higher order map function for lists. That is, we show that `list_map f` is
152 wait-free whenever it is applied to any wait-free argument `f`. We also show that our approach
153 supports practical wait-free algorithms by verifying a memory-efficient single-producer, single-
154 consumer queue [21]. An important point about this example, and in fact many practical
155 wait-free algorithms, is that these algorithms in fact assume fixed upper bounds on the
156 number of concurrent operations. Verification of this class of wait-free algorithms has been
157 largely ignored by the prior work on verification of wait-freedom. Supporting reasoning
158 about such examples formally requires us to introduce a notion of what we call *restricted*
159 wait-freedom, and to slightly adjust our approach to verification of wait-freedom.

160 In summary, we make the following contributions in this paper:

- 161 ■ A formal definition of wait-freedom for higher-order languages together with two variations
162 needed to account for realistic examples
- 163 ■ A Lawyer-based specification that internalizes the notion of wait-freedom in the Lawyer
164 logic
- 165 ■ A solution for verifying reasoning about arbitrary client programs based on combining a
166 logical relations model for safety together with Lawyer logic's approach to liveness

- 167 ■ Verification of wait-freedom of a number of case studies including those not supported by
 168 prior work
 169 ■ All results have been mechanized in the Rocq Prover

170 The rest of the paper is organized as follows. We formally define wait-freedom in Section 2
 171 where we also state the adequacy theorem that reduces proving wait-freedom of an operation
 172 to proving its wait-freedom specification in Lawyer. Then, in Section 3 we prove such
 173 specification for a simple example program to illustrate our approach. In Section 4 we
 174 continue by verifying more complicated programs which require adapting both the definition
 175 of wait-freedom and the Lawyer specification as we explained earlier. We explain the key
 176 ideas for proving our adequacy theorem in Section 5. Finally, we discuss a limitations of our
 177 solution in Section 6, compare it with related work in Section 7, and conclude in Section 8.

178 2 Overview

179 In this section, we formally define the notion of wait-freedom as a property of execution
 180 traces of the programming language under consideration. To avoid tedious direct reasoning
 181 about execution traces, our approach reduces proving this property to proving a specification
 182 with Lawyer program logic. We therefore briefly show the specification format of Lawyer and
 183 state the adequacy theorem for wait-freedom that establishes the aforementioned reduction.

184 Throughout the paper, we will be using the notation $m[k]$ to denote a key k lookup in
 185 a map-like structure m . That includes m being a list (and k an index in it), or a partial
 186 function $m : A \xrightarrow{\text{fin}} B$ for some types A and B (and k an element of type A). Moreover, we
 187 will use this notation for the trace configuration lookup which we define below. When we
 188 write $m[k]$, we implicitly assume that k belongs to the domain of m and thus the lookup is
 189 well-defined. Then, by $m[k \mapsto v]$ we denote the structure m with the key k updated to the
 190 value v .

191 2.1 Definition of wait-freedom for our language

192 We consider the programs written in the language $\lambda_{\text{ref}}^{\text{conc}}$ — an ML-like higher-order language
 193 with general references and concurrency. An excerpt of its semantics shown in Figure 3,
 194 whereas the full semantics is provided in the Rocq development. A less common feature of
 195 $\lambda_{\text{ref}}^{\text{conc}}$ — **to_int** primitive — is explained further in Section 4.2. The *fill* ($K : ECtx$) ($e : Expr$)
 196 operation fills the hole \bullet in the context K with the expression e ; we avoid the more common
 197 $K[e]$ notation, as it is used for lookup operations in this paper.

198 We define wait-freedom as a property of the execution traces of $\lambda_{\text{ref}}^{\text{conc}}$. Intuitively, an
 199 operation $\text{op} : Val$ is wait-free if, for any execution trace, any “call” to op is followed by a
 200 “return” later in this trace. A “call” in $\lambda_{\text{ref}}^{\text{conc}}$ is simply an application $\text{op } a$ for some ($a : Val$).
 201 To identify a call in a thread pool \mathcal{E} , we need to provide a thread identifier $\tau : \text{Thread}$ and
 202 an evaluation context $E : ECtx$. Then, a return in our approach is simply a value, identified
 203 with a thread identifier and an evaluation context.

► **Definition 1** (Call and return in a thread pool).

$$204 \quad \text{Call}(\mathcal{E}, \tau, E, \text{op}) \triangleq \exists a \in Val. \mathcal{E}[\tau] = \text{fill } E (\text{op } a)$$

$$205 \quad \text{Ret}(\mathcal{E}, \tau, E) \triangleq \exists r \in Val. \mathcal{E}[\tau] = \text{fill } E r$$

206 To define wait-freedom, we consider execution traces of $\lambda_{\text{ref}}^{\text{conc}}$. Before proceeding, we intro-
 207 duce a number of general notations for traces. Below, $tr[i]$ and $tr\langle i \rangle$ denote correspondingly

XX:6 Verifying wait-freedom for concurrent higher-order programs

Syntax

$x, y, f \in \text{Var}$	$l \in \text{Loc}$	$n \in \mathbb{Z}$	$b \in \mathbb{B}$
	$\odot ::= + \mid - \mid * \mid = \mid < \mid \dots$		
Val	$v ::= () \mid b \mid n \mid l \mid (v, v) \mid \mathbf{inl} v \mid \mathbf{inr} v \mid \mathbf{rec} f(x) := e$		
	$\mid \mathbf{match} e \mathbf{with} \mathbf{inl} x \Rightarrow e \mid \mathbf{inr} y \Rightarrow e \mathbf{end} \mid \dots$		
Expr	$e ::= x \mid v \mid e \odot e \mid \mathbf{ref}(e) \mid \mathbf{if} e \mathbf{then} e \mathbf{else} e \mid \mathbf{fork}(e) \mid !e \mid ee$		
	$\mid \mathbf{to_int} e \mid \mathbf{free} e \mid \mathbf{fst} e \mid \mathbf{snd} e \mid \dots$		
ECtx	$E ::= \bullet \mid e \odot E \mid E \odot v \mid \mathbf{if} E \mathbf{then} e \mathbf{else} e \mid (e, E) \mid (E, v)$		
	$\mid eE \mid Ev \mid \dots$		
Heap	$h \in \text{Loc} \xrightarrow{\text{fin}} \text{Val}$		
ThreadPool	$\mathcal{E} \in \text{Thread} \xrightarrow{\text{fin}} \text{Expr}$		
Config	$c ::= (\mathcal{E}, h)$		

Pure reductions

$$\begin{array}{l}
 (\mathbf{rec} f(x) := e) v \xrightarrow{\text{pure}} e[\mathbf{rec} f(x) := e/f][v/x] \quad (\mathbf{if} \text{ true then } e_1 \mathbf{else} e_2 \xrightarrow{\text{pure}} e_1 \\
 \vdots
 \end{array}$$

Per-thread reductions

$$\begin{array}{l}
 (e, h) \rightsquigarrow (e', h) \quad \text{if } e \xrightarrow{\text{pure}} e' \\
 (\mathbf{ref}(v), h) \rightsquigarrow (\ell, h[\ell \mapsto v]) \quad \text{if } \ell \notin \text{dom}(h) \\
 \vdots
 \end{array}$$

Configuration reductions

$$\begin{array}{c}
 \frac{(e, h) \rightsquigarrow (e', h')}{(\mathcal{E}[\tau \mapsto \text{fill } E e], h) \xrightarrow{\tau} (\mathcal{E}[\tau \mapsto \text{fill } E e'], h')} \\
 \frac{\tau' \notin \text{dom}(\mathcal{E}) \cup \{\tau\} \quad \mathcal{E}' \triangleq \mathcal{E}[\tau \mapsto \text{fill } E ()][\tau' \mapsto e]}{(\mathcal{E}[\tau \mapsto \text{fill } E(\mathbf{fork}(e))], h) \xrightarrow{\tau} (\mathcal{E}', h)}
 \end{array}$$

■ **Figure 3** An excerpt of $\lambda_{\text{ref}}^{\text{conc}}$ semantics

208 the i th state and label of a trace tr . In particular, for an execution trace etr of $\lambda_{\text{ref}}^{\text{conc}}$ and
 209 any index i of it, $tr[i]$ is the configuration (a pair of a thread pool and a heap, according to
 210 Figure 3) before the i th step, and $etr\langle i \rangle$ is the identifier of the thread taking the i th step.
 211 Then, $\text{dom}(tr)$ denotes the set of indices i for which $tr[i]$ is defined, *i.e.*, natural numbers
 212 between 0 and the (potentially infinite) length of tr .

213 We'd like to state that every call must return. Of course, it only makes sense when the
 214 calling thread is kept being scheduled until the return. This is captured by the following
 215 definition.

216 ► **Definition 2** (Eventual return of calls). *For an operation $\text{op} : \text{Val}$, execution trace etr and*
 217 *thread identifier τ , we say that all τ 's calls to op in etr eventually return iff*

$$\begin{array}{l}
 \text{alwaysReturnsInTrace}(\text{op}, etr, \tau) \triangleq \forall (E : \text{ECtx}), i. \\
 \text{Call}(etr[i].1, \tau, E, \text{op}) \implies \text{schedUntilRet}(etr, \tau, E, i) \implies \exists j \geq i. \text{Ret}(etr[j].1, \tau, E)
 \end{array}$$

220 where

$$\begin{array}{l}
 \text{schedUntilRet}(etr, \tau, E, i) \triangleq \forall j \geq i. j \in \text{dom } etr \implies \\
 \neg(\exists k \in [i..j]. \text{Ret}(etr[k].1, \tau, E)) \implies \exists p \geq j. etr\langle p \rangle = \tau,
 \end{array}$$

223 Finally, we restrict the starting configuration of traces under consideration:

224 ■ The program of every thread in the starting configuration must be a substitution of the
 225 wait-free operation's code into some expression with a free variable x . Moreover, this
 226 expression must not contain hard-coded locations and pointer arithmetic (denoted by
 227 $noLocs$). An example from Figure 2 illustrates why it is necessary to restrict random
 228 memory access by the client program.

229 ■ The data structure accessed by the operation must be initialized. What exactly “initial-
 230 ization” means depends on the operation (see an example below), therefore our definition
 231 of wait-freedom is parameterized with the set of starting configurations.

232 As a result, we define wait-freedom as follows.

233 ► **Definition 3** (Wait-freedom). *Given a set of “initialized” configurations $C \subseteq Config$, we*
 234 *define wait-freedom of an operation op as*

$$\begin{aligned}
 235 \quad waitFree_C(op) \triangleq & \forall etr. \bigwedge_{\tau \mapsto e \in etr[0].1} (\exists e'. e = e'[op/x] \wedge noLocs(e')) \implies \\
 236 \quad & etr[0] \in C \implies \\
 237 \quad & \forall \tau. alwaysReturnsInTrace(op, etr, \tau)
 \end{aligned}$$

238 For example, in Section 4 we show that a FAA-based `incr` operation of counter (Figure 1c)
 239 is wait-free if the location l of the counter is initialized with an integer. In terms of Definition 3,
 240 it is denoted as

$$241 \quad waitFree_{hasInt(l)}(\mathbf{incr} \ l), \quad \text{where } hasInt(l) \triangleq \{c \mid c.2[l] \in \mathbb{Z}\}$$

242 We note that, while Definition 3 captures the behavior of simple wait-free operations,
 243 more complicated algorithms only satisfy weaker forms of that property. We will see examples
 244 of it in Section 4.1 and Section 4.2.

245 2.2 Proving wait-freedom with Lawyer logic

246 The notion of wait-freedom as defined in Definition 3 is formulated at the level of execution
 247 traces, which are notoriously hard to reason about directly. To avoid that, our approach
 248 reduces establishing wait-freedom of a given operation to verifying this operation with
 249 Lawyer[31] — a concurrent separation logic built on top of Iris [25] and Trillium [35] frameworks.
 250 We explain the details of Lawyer logic in the subsequent sections; here, we present the high-
 251 level form of the specification that a wait-free operation must satisfy.

252 ► **Definition 4** (Lawyer specification of wait-freedom). *For a set of allowed starting configura-*
 253 *tions $C \subseteq Config$, we define the Lawyer specification of op 's wait-freedom as*

$$254 \quad WaitFreeSpec_C(op) \triangleq \forall c \in C. heap_repr(c) \implies NoInfExec(op) \wedge PresModInv(op)$$

255 The above specification can be read as follows. For any allowed starting configuration
 256 c , we assume the ownership of c 's entire heap, which is represented as an Iris proposition
 257 $heap_repr(c)$. Given that (\implies can be read as an implication), we can establish two speci-
 258 fications for the operation op . First, $NoInfExec(op)$ states that for any argument $a : Val$,
 259 execution of $op \ a$ cannot be infinite. Second, $PresModInv(op)$ states that, again, for any a ,
 260 execution of $op \ a$ preserves the internal invariants that op relies on (*e.g.* that the location l
 261 of counter in Figure 1c always stores a number).

262 Our adequacy theorem states that the definition above indeed implies wait-freedom.

XX:8 Verifying wait-freedom for concurrent higher-order programs

263 ► **Theorem 5 (Adequacy).** *If $\text{WaitFreeSpec}_C(\text{op})$ is provable in Lawyer, then $\text{waitFree}_C(\text{op})$*
 264 *holds in the meta-logic.*

265 As we note in Section 2.1, the notion of “wait-freedom” for more complicated operations
 266 is more elaborate. Therefore, for such operations we prove slightly different versions of
 267 Definition 4.

268 We show how the specifications from Definition 4 are proven for simple programs in
 269 Section 3. Then, in Section 4 we show how our approach can be adapted to verify programs
 270 that satisfy weaker versions of Definition 3. Finally, we discuss the key insights behind the
 271 proof of Theorem 5 in Section 5.

3 Verifying wait-freedom in Lawyer logic

273 In this section, we show how an operation is verified against the wait-freedom specification of
 274 Definition 4. As an example, we consider the wrapper `incr l` over the wait-free implementation
 275 of counter increment from Figure 1c. Specifically, we verify $\text{WaitFreeSpec}_{\text{hasInt}(l)}(\text{incr } l)$
 276 in the Lawyer logic for any location l . Doing so allows us to use Theorem 5 to conclude
 277 $\text{waitFree}_{\text{hasInt}(l)}(\text{incr } l)$ holds at the meta-level. We will introduce the details of Iris and the
 278 Lawyer logic as they are needed.

279 Consider some location ℓ . By definition of WaitFreeSpec , we need to prove

$$280 \quad \forall c \in \text{hasInt}(l). \text{heap_repr}(c) \Rightarrow \text{NoInfExec}(\text{incr } l) \wedge \text{PresModInv}(\text{incr } l)$$

To do so, we take an arbitrary configuration $c : \text{Config}$ for which we have $c.2[l] = k$ for some
 $k \in \mathbb{Z}$, and further assume ownership of each heap location of c :

$$\text{heap_repr}(c) \triangleq \bigstar_{\{(\ell, v) \mid c.2[\ell] = v\}} \ell \mapsto v$$

Iris is a separation logic and thus Iris propositions assert ownership of *resources* representing both the state of the program, and so-called ghost resources. For example, ownership of a single heap location is expressed with a *points-to* proposition, written $\ell \mapsto v$, with the usual intuitive meaning: it asserts both exclusive ownership of the memory location ℓ , and the fact that it currently stores the value v . To assume ownership of the entire heap, we take the *separating conjunction* of points-to propositions for all locations in the heap. The separating conjunction $P * Q$ asserts ownership of resources which can be split into two *disjoint* parts, one satisfying P and the other Q . The notion of being disjoint depends on particular resource kind. For points-to propositions, $l_1 \mapsto v_1 * l_2 \mapsto v_2$ means that $l_1 \neq l_2$; hence the exclusivity of ownership of the memory location. Thus, given the ownership of the entire heap, we have to prove the specifications $\text{NoInfExec}(\text{incr } l)$ and $\text{PresModInv}(\text{incr } l)$. However, proving both will rely on the fact that l always stores some number. Such a fact can be expressed through Iris *invariants* as follows:

$$\text{cntInv} \triangleq \boxed{\exists k \in \mathbb{Z}. l \mapsto k}$$

281 An invariant proposition \boxed{P} means that P always holds throughout program’s execution.
 282 While verifying a single program step, one can assume that P holds by *opening* the invariant,
 283 but it must be *closed* again right after taking that step of computation by reestablishing
 284 that P holds again. Importantly, invariants (but not their content in general) are duplicable.
 285 That is, $\boxed{P} \vdash \boxed{P} * \boxed{P}$. For concurrent operations such as `incr l` that we consider, it means
 286 that every thread running an operation can be given a knowledge about the invariant.

287 Note that the definition of `WaitFreeSpec` features a so-called *viewshift*, written \Rightarrow . View-
 288 shifts are similar to logical implication, but additionally allow give access invariants and
 289 resources. That is, to prove $P \Rightarrow Q$, in addition to assuming P , we can open/close invariants
 290 and update resources so as to establish Q . Thus, given ownership of the entire heap, and the
 291 fact that l currently points to some integer, we establish `cntlnv` by providing $l \mapsto k$ (ignoring
 292 the rest of the heap). From that point on we can assume that `cntlnv` holds while proving the
 293 specifications.

The first specification we need to prove, `NoInfExec(incr l)`, informally captures that execution of our operation applied to any argument cannot run forever. Formally, it is defined as a dependent pair of a natural number and a specification that mentions it¹:

$$\text{NoInfExec}(\text{op}) \triangleq$$

$$\{\mathbf{F} : \mathbb{N} \mid \forall \tau : \text{Thread}, \pi : \text{Phase}, a : \text{Val}. \{\mathbf{F} \cdot \text{bar}_\circ \pi * \text{ph}_\circ \tau \pi\} \text{op } a \{ _ . \text{ph}_\circ \tau \pi \}^\tau \}$$

294 The number \mathbf{F} is (an over-approximation of) the upper bound on the execution time of
 295 `op`. The actual specification for `op` is defined as a Lawyer *Hoare triple*. In general, a triple
 296 $\{P\} e \{\Phi\}^\tau$ means that assuming that the precondition P holds, the program e without
 297 getting stuck (*i.e.*, it does not crash) when run under the thread τ , and that whenever e
 298 reduces to a value v , the postcondition $\Phi(v)$ holds. The Hoare triple also enforces that
 299 throughout the execution of the program all invariants are preserved. Moreover, Lawyer
 300 Hoare triples additionally enforce that every single step of the program e is refined by a
 301 transition in the so-called Obligations Model of Lawyer—a transition system that Lawyer
 302 instantiates Trillium with. As a result, given a Lawyer Hoare triple for a program e implies
 303 that any execution trace of e refines some trace of the Obligations Model. As all traces of
 304 the Obligation Model are finite, this refinement allows Lawyer to establish termination of
 305 verified programs.

306 The precondition and postcondition of the triple `NoInfExec(op)` above consist of Lawyer
 307 logic resources. First, the precondition can be read as an assignment of \mathbf{F} barrels of *fuel*
 308 to the thread τ , with an indirection via the *phase* π that marks the barrels, and is assigned
 309 to τ . The fuel is represented by the $\mathbf{F} \cdot \text{bar}_\circ \pi$ resource (a separating conjunction of \mathbf{F}
 310 individual barrels $\text{bar}_\circ \pi$). The phase assignment is written in the Lawyer logic as $\text{ph}_\circ \tau \pi$.
 311 The fuel resource is a logical representation of the number of remaining steps the program
 312 is allowed to take. As we will see later, verifying every single execution step consumes one
 313 barrel of fuel. As an aside we mention that in Lawyer [31], both the fuel and the phase of
 314 resources have one extra parameter each. Here, we assign these parameters the values that
 315 are largely irrelevant for this presentation; instead of showing these values, we just mark
 316 both resources with a \circ subscript.

317 Later on in the paper, in Section 6, we will argue why specifically the specific specification
 318 pattern of `NoInfExec(op)`, as well as certain restrictions on the operation's invariants, captures
 319 the nature of wait-freedom in the Lawyer logic.

To prove `NoInfExec(incr l)`, we take \mathbf{F} to be 3, this requires us to show that for any τ , π ,
 and a we have the following:

$$\text{cntlnv} \vdash \{3 \cdot \text{bar}_\circ \pi * \text{ph}_\circ \tau \pi\} \text{incr } l a \{ _ . \text{ph}_\circ \tau \pi \}^\tau$$

The triple above essentially states that we can spend up to 3 barrels of fuel to complete the execution of `op a`. To prove this triple, we use the rules of the Lawyer logic [31]. The Lawyer

¹ The reason why \mathbf{F} cannot be existentially quantified inside the triple is related to the issues of step-indexing which are described in detail by Spies *et al.* [33].

XX:10 Verifying wait-freedom for concurrent higher-order programs

logic provides rules that cover updates to both the physical state of execution (*i.e.*, the heap) and the state of the Obligations Model. For simplicity, we only show a number of derived rules that are necessary for our specific example.² An execution of `incr l a` consists of two beta-reductions and one FAA step. For the former, we use the following rule:

$$\begin{array}{c} \text{LAWYER-BETA-FUEL} \\ \{\text{ph}_o\tau\pi * P\} e[v/x] \{\text{ph}_o\tau\pi\}^\tau \vdash \{\text{bar}_o\pi * \text{ph}_o\tau\pi * P\} (\lambda x. e) v \{\text{ph}_o\tau\pi\}^\tau \end{array}$$

The rule above says that to verify a beta-reduction, user must spend a barrel of fuel and preserve the phase resource. After the beta-reduction step the user can continue verifying the reduct with the remaining resources. Thus, we apply this rule twice by choosing P to be $2 \cdot \text{bar}_o\pi$ and $\text{bar}_o\pi$, respectively. After these two beta-reduction steps, we are left to prove the following:

$$\text{cntInv} \vdash \{\text{bar}_o\pi * \text{ph}_o\tau\pi\} \mathbf{FAA}(l, 1) \{_ . \text{ph}_o\tau\pi\}^\tau$$

The fetch-and-add primitive can only proceed when the location it accesses contains an integer, which is the case for our example as evidenced by our invariant. Thus, to verify the FAA step we need to open the invariant, get the current value under location l , increment it, and close the invariant again. Moreover, as with the previous steps, we need to spend fuel for this step as well. All of that is captured by the following rule:

$$\begin{array}{c} \text{LAWYER-FAA-FUEL-INV} \\ P \multimap \exists m : \mathbb{Z}. \ell \mapsto m * (\ell \mapsto (m + n) \multimap P) \\ \hline \boxed{P} \vdash \{\text{bar}_o\pi * \text{ph}_o\tau\pi\} \mathbf{FAA}(\ell, n) \{_ . \text{ph}_o\tau\pi\}^\tau \end{array}$$

320 The (*magic*) *wand* connective of separation logic, written \multimap , is a separating implication:
 321 $P \multimap Q$ holds for a resource such that when added to any resource satisfying P , once can
 322 satisfy Q . Hence, the premise of the rule `LAWYER-FAA-FUEL-INV` says that the content of
 323 invariant should contain a points-to resource for the location, and that the invariant could be
 324 reestablished after the location has been updated. This is trivially satisfied by `cntInv`. The
 325 expression $\mathbf{FAA}(l, 1)$ reduces to the value stored under l in the heap before the update takes
 326 place. This value is ignored by our postcondition which only requires the phase resource.
 327 Thus, after using the rule `LAWYER-FAA-FUEL-INV` the proof of `NoInExec(incr l)` is finished.

The next specification that we need to prove, *i.e.*, `PresModInv(incr l)`, is defined as a different Hoare triple as follows:

$$\text{PresModInv}(\text{op}) \triangleq \forall \tau : \text{Thread}, a : \text{Val}. \{\text{RS}_V(a)\}_{\text{noMod}}^{\frac{1}{2}} \text{op } a \{v. \text{RS}_V(v)\}^\tau$$

328 The overall intuitive meaning of this triple is similar to one described above: if the
 329 precondition holds, then an execution of `op a` preserves invariants, and if it reduces to a
 330 value, then the postcondition holds for that value. However, there are two major differences.
 331 First, the execution is allowed to get stuck (denoted by $\frac{1}{2}$ superscript). Second, this triple
 332 does not establish a refinement between the program and the Obligation model and thus
 333 its precondition and postcondition, as well as invariants it relies on, cannot mention any
 334 Lawyer resources.³ We use the `noMod` subscript to denote that the triple does not deal with

² The rules that we present below are derived from the rules of Lawyer logic which are more general for more flexibility. However, that level of generality and flexibility is not necessary for the results presented in this paper, and thus presenting them would only muddle the presentation.

³ Technically, this Hoare triple is defined by directly instantiating the Trillium's program logic with a trivial, always looping transition system.

```

let rec list_map f l =
  match l with
  | None => None
  | Some p => Some (f (fst p), list_map f (snd p))
end

```

(a) Implementation

$$\text{isList}(v : \text{Val}) \triangleq v = \text{None} \vee \exists h, v'. v = \text{Some}(h, v') \wedge \text{isList}(v')$$
(b) Representation of lists in $\lambda_{\text{ref}}^{\text{conc}}$

■ **Figure 4** Higher-order list mapping function

335 refinement to a model. As we will explain in Section 5.2, such Hoare triples ensure that
 336 `op` can be used by an arbitrary client program, and that it would not violate the internal
 337 invariants of `op`. Technically, a client program *uses* `op` by providing it a *safe* argument,
 338 captured by the predicate $\text{RS}_{\mathbf{V}}(a)$, and expects `op` to return a safe value back. We will
 339 explain the details of what it means to be safe value and the motivation as to why it is
 340 necessary to restrict only to safe values in Section 5.2.

341 The proof of $\text{PresModInv}(\text{incr } l)$ proceeds similarly to one above. Since the input argument
 342 a is not used, we ignore the $\text{RS}_{\mathbf{V}}(a)$ resource. The rules used to verify beta-reductions and
 343 FAA can be obtained by simply removing Lawyer resources from the corresponding rules
 344 above and adjusting the kind of Hoare triple. Similarly to the proof above, we use `cntInv` to
 345 verify the FAA step. The main difference is that the return value v is not ignored—it needs
 346 to be shown to be safe, *i.e.* $\text{RS}_{\mathbf{V}}(v)$. However, as we will see in Section 5.2 all integers are
 347 safe, *i.e.*, $\forall n : \mathbb{Z}. \vdash \text{RS}_{\mathbf{V}}(n)$.

348 4 Case studies

349 As we discussed in the Introduction, while the simple definition of wait-freedom that we
 350 worked with in Sections 2 and 3 works for many simple programs, for more involved programs,
 351 this definition needs to be adjusted. In this section, we motivate and explain two independent
 352 extensions to our approach to verification of wait-freedom. One where we allow to verify
 353 (higher-order) functions that potentially get stuck, and one where the data structure restricts
 354 the number of operations that can run concurrently at a given time. Throughout this section
 355 we will use the following helper definitions:

356 $\text{None} \triangleq \text{inl } ()$ $\text{Some } v \triangleq \text{inr } v$

357 and similarly define a corresponding notation for in pattern matching where we simply write
 358 `match e with None => e1 | Some v => e2 end`.

359 4.1 Modular verification of higher-order wait-free functions

360 Consider the higher-order `list_map` function in Figure 4a. It applies the given function f
 361 to all elements of the list l . In our language, lists are values satisfying the recursive predicate
 362 `isList` defined in Figure 4b.

363 Note that if the input function f does not run forever for any argument, the partially
 364 applied `list_map f` does not run forever for any argument either. However, it might get
 365 stuck even if f never does: if its second argument l does not satisfy `isList`, then `list_map` will
 366 get stuck, *e.g.*, on taking projections. Thus, the behavior of calls to `list_map f` is captured
 367 by the following variant of Definition 2, with the changes [highlighted](#):

XX:12 Verifying wait-freedom for concurrent higher-order programs

368 ► **Definition 6** (Eventual return or stuckness of calls to `op`). *For an operation $op : Val$, execution*
369 *trace etr and thread identifier τ , we say that all τ 's calls to op in etr eventually return or*
370 *get stuck, written $alwaysReturnsInTrace^{\sharp}(op, etr, \tau)$, iff*

$$\begin{aligned} 371 \quad & alwaysReturnsInTrace^{\sharp}(op, etr, \tau) \triangleq \\ 372 \quad & \forall \mathcal{E}, i. Call(etr[i].1, \tau, E, op) \wedge schedUntilRet(etr, \tau, E, i) \implies \\ 373 \quad & \exists j \geq i. Ret(etr[j].1, \tau, E) \vee stuck(etr[j], \tau) \end{aligned}$$

374 *where stuckness is defined as follows:*

$$375 \quad stuck(c, \tau) \triangleq \neg \exists c'. c \xrightarrow{\tau} c'$$

376 Similarly, we define possibly-stuck wait-freedom $waitFree^{\sharp}$ as a variant of Definition 3
377 that uses $alwaysReturnsInTrace^{\sharp}$ instead of $alwaysReturnsInTrace$.

378 Our approach to proving wait-freedom supports proving verifying possibly-stuck wait-
379 freedom with minimal changes compared to the ordinary wait-freedom we had discussed
380 in earlier section. That is, we prove a variant of Theorem 5 that establishes possibly-stuck
381 wait-freedom of an operation, given that it satisfies the Lawyer specification $WaitFreeSpec^{\sharp}$,
382 which in turn, similarly to definitions above, is a weakening of $WaitFreeSpec$ that allows the
383 operation to possibly get stuck.

384 Now, we can apply the approach outlined in Sections 2 and 3 to prove that `list_map f`
385 is possibly-stuck wait-free. Crucially, we do so modularly by verifying both `list_map` and `f`
386 independently of one another. Formally, we prove the following higher-order possibly-stuck
387 wait-freedom specification for `list_map`:

$$388 \quad \forall f, C. WaitFreeSpec_C^{\sharp}(f) \implies WaitFreeSpec_C^{\sharp}(list_map\ f) \quad (\text{list-map-HO-spec})$$

389 As in Section 3, we prove the two specifications for `list_map f`. For each of them, we
390 make use of the corresponding specification for `f`. The logical inference rules used in the
391 process are the same used in Section 3, with the addition of rules to verify stuck computations,
392 *e.g.*, the following Hoare triple for projecting of the unit value: $\{\text{True}\}^{\sharp} \mathbf{fst}() \{P\}^{\tau}$.

393 Given the (list-map-HO-spec) spec above, together with the fact that ordinary Hoare
394 triples imply possibly-stuck Hoare triples, we can reuse already the proven (possibly-stuck)
395 wait-freedom specifications to derive possibly-stuck wait-freedom of `list_map` partially
396 applied to specific function argument. For example, by weakening the Hoare triple that we
397 proved in Section 3, we obtain a proof of $WaitFreeSpec_{hasInt(l)}^{\sharp}(list_map(incr\ l))$ for any
398 location l , thus establishing $waitFree_{hasInt(l)}^{\sharp}(list_map(incr\ l))$.

399 4.2 Restricted wait-freedom

400 Many practical implementations of wait-free algorithms [21, 27, 36, 14] assume an extra
401 condition not accounted for by Definition 3: that the number of concurrent operations has
402 a fixed upper bound. An example of such implementation is shown in Figure 5a. This
403 implementation of a concurrent wait-free queue has two uncommon features. First, besides
404 adding and removing elements to the queue, it allows reading the head element without
405 removing it. Second, it optimizes the memory usage by deallocating the head node when it
406 is removed. Supporting both of these features, and simultaneously safety and Linearizability
407 is highly non-trivial. Therefore, the implementation in Figure 5a limits the concurrency to
408 just two threads: an enqueueer thread that can both enqueue elements and read the queue
409 head, and a dequeuer thread that can both dequeue elements and read the queue head.

```

(* Node = {val: Value; next: pointer to Node} *)
(* Head, Tail, BeingRead, FreeLater : pointer to Node *)
(* OldHeadVal: Value *)

(* methods used by "enqueueer" thread: *)

let read_head_e () =
  let cur_head = !Head in
  if (cur_head == !Tail) then None
  else
    BeingRead := cur_head;
    let v =
      if (cur_head == !Head) then
        cur_head.val
      else
        !OldHeadVal
    in
    Some v

let enqueue v =
  let dummy = Node {val: 0, next: null} in
  let new_tail = ref dummy in
  let cur_tail = !Tail in
  cur_tail.val := v;
  cur_tail.next := new_tail;
  Tail := new_tail

(* methods used by "dequeuer" thread: *)

let deque () =
  let cur_head = !Head in
  if (cur_head == !Tail) then None
  else
    let v = cur_head.val in
    OldHeadVal := v;
    Head := cur_head.next;
    let to_free =
      if (cur_head == !BeingRead) then
        let old_read = !FreeLater in
        FreeLater := cur_head;
        old_read
      else
        cur_head
    in
    free to_free; Some v

let read_head_d () =
  let cur_head = !Head in
  if (cur_head == !Tail) then
    None
  else
    Some cur_head.val

```

(a) Original implementation by Jayanti and Petrovic [21] expressed in λ_{ref}^{conc}

```

let enqueueer v =
  if (v == ()) then
    Some (read_head_e ())
  else
    match (to_int v) with
    | Some n => Some (enqueue n)
    | None => None
  end

let dequeuer v =
  if (v == true) then
    Some (read_head_d ())
  else if (v == false) then
    Some (deque ())
  else
    None

```

(b) Wrappers to be verified

■ **Figure 5** Single-enqueueer, single-dequeueer queue

410 To support verifying such algorithms, which we call *restricted wait-free* requires multiple
 411 changes to our approach to verification of wait-freedom as we have described so far. First of
 412 all, we define two wrapper functions `enqueueer` and `dequeuer` (see Figure 5b) which wrap
 413 the possible operations of the queue for the enqueueer and dequeuer threads explained above
 414 into a single operation respectively—each function checks its argument and dispatches the
 415 appropriate queue operation accordingly. Then, for technical reasons that we explain in
 416 Section 6, we restrict the queue to only store integers. To that end, the `enqueueer` wrapper
 417 includes a call to the `to_int` primitive (which we have added to λ_{ref}^{conc} for this example), the
 418 semantics of which are as follows:

$$419 \quad \text{to_int}(n \in \mathbb{Z}) \stackrel{\text{pure}}{\rightsquigarrow} \text{Some } n \quad \text{to_int}(v \notin \mathbb{Z}) \stackrel{\text{pure}}{\rightsquigarrow} \text{None}$$

420 Restricted wait-freedom can be stated as a variant of Definition 3 as follows, with changes
 421 highlighted:⁴

⁴ Note that below we assume that $\text{Thread} \triangleq \mathbb{N}$ and that a thread pool can be treated as a list of expressions.

XX:14 Verifying wait-freedom for concurrent higher-order programs

422 ► **Definition 7** (Wait-freedom). *Given a set of “initialized” configurations $C \subseteq \text{Config}$, we*
 423 *define restricted wait-freedom of list of operations $\text{ops} : \text{List}(\text{Val})$ as*

$$\begin{aligned}
 424 \quad & \text{waitFreeRestr}_C(\text{ops}) \triangleq \\
 425 \quad & \forall \text{etr}. \text{length}(\text{etr}[0].1) = \text{length}(\text{ops}) \wedge \\
 426 \quad & \left(\bigwedge_{\tau \mapsto e \in \text{etr}[0].1} \exists e'. \text{etr}[0].1[\tau] = e'[\text{ops}[\tau]/x] \wedge \text{noLocs}(e') \wedge \text{noForks}(e') \right) \wedge \\
 427 \quad & \text{etr}[0] \in C \implies \forall \tau. \text{alwaysReturnsInTrace}(\text{ops}[\tau], \text{etr}, \tau)
 \end{aligned}$$

428 This definition establishes a property of a *list* of methods—one operation may appear
 429 multiple times in this list. Initially (at the start of etr), we have $\text{length}(\text{ops})$ many threads,
 430 and the i^{th} thread is a client of the i^{th} method in ops . Importantly, none of threads
 431 threads should fork any other threads—otherwise the restriction on the number of concurrent
 432 operations might be violated. The *alwaysReturnsInTrace* conclusion states that every call to
 433 one of the operations in any of the threads must always return.

434 For example, restricted wait-freedom of the queue algorithm in Figure 5b is formally
 435 stated using Definition 7 by taking $\text{ops} \triangleq [\text{enqueueer}, \text{dequeueer}]$. For more complicated
 436 n -enqueueer, single-dequeueer algorithm (also presented in [21]) one would need to use a list
 437 consisting of n enqueueer methods and a single dequeueer method, *i.e.*, $\text{ops} \triangleq \text{dequeueer} ::$
 438 $\text{repeat}(n, \text{enqueueer})$.

439 Verifying restricted wait-freedom requires *changes* to the specification we had given in
 440 Definition 4. We will first give the formal definition and explain the details afterwards.

441 ► **Definition 8** (Lawyer specification of restricted wait-freedom). *For a set of allowed starting*
 442 *configurations $C \subseteq \text{Config}$, we define the Lawyer specification of $\text{ops} : \text{List}(\text{Val})$ restricted*
 443 *wait-freedom as follows, for some $\mathbf{F} : \mathbb{N}$:*

$$\begin{aligned}
 444 \quad & \text{WaitFreeRestrSpec}_C(\text{ops}) \triangleq \forall c \in C. \text{heap_repr}(c) \Rightarrow \text{mtoks}(\text{ops}) * \\
 445 \quad & \bigstar_{\text{op} \in \text{ops}} \left(\forall \tau, \pi, a. \{ \mathbf{F} \cdot \text{bar}_\circ \pi * \text{ph}_\circ \tau \pi * \text{mtok}(\text{op}) \} \text{op } a \{ _ . \text{ph}_\circ \tau \pi * \text{mtok}(\text{op}) \}^\tau \right) * \\
 446 \quad & \bigstar_{\text{op} \in \text{ops}} \left(\forall \tau, a. \{ \text{mtok}(\text{op}) \}_{\text{noMod}}^\dagger \text{op } a \{ v. \text{isGroundVal}(v) * \text{mtok}(\text{op}) \}^\tau \right)
 \end{aligned}$$

where

$$447 \quad \text{mtoks}(l : \text{List}(\text{Val})) \triangleq \text{toks}_{\text{ops}}(l) \quad \text{and} \quad \text{mtok}(m : \text{Val}) \triangleq \text{mtoks}([m])$$

In the definition above, $\text{isGroundVal}(v : \text{Val})$ states that the value v does not contain any
 locations or lambda functions. This is related to the restriction where we require the queue
 to only store integers which we will discuss in Section 6. The proposition $\text{mtoks}(l)$ (an
 abbreviation for more general $\text{toks}_{\text{ops}}(l)$) is a proposition which asserts ownership of the list l
 of *method tokens*. The list of tokens can be split into individual tokens, and most importantly,
 the total amount of tokens is determined upon creation and cannot be increased. Thus,
 requiring ownership of $\text{mtok}(\text{op})$ for calling an operation op allows us to restrict the number
 of calls to op in the program logic. In logic, this is expressed by the following laws (where all
 parameters of toks are lists of values and $\text{count}(m, l)$ counts occurrences of m in list l):

$$\text{toks}_L(l_1) * \text{toks}_L(l_2) \dashv\vdash \text{toks}_L(l_1 ++ l_2) \quad \text{toks}_L(l) \vdash \forall m \in l. \text{count}(m, l) \leq \text{count}(m, L)$$

448 Finally, the adequacy theorem for restricted wait-freedom is similar to Theorem 5:

449 ► **Theorem 9** (Adequacy for restricted wait-freedom). *If $\text{WaitFreeRestrSpec}_C(\text{ops})$ is provable*
 450 *in Lawyer, then $\text{waitFreeRestr}_C(\text{ops})$ holds in the meta-logic.*

451 We use Theorem 9 to establish $\text{waitFreeRestr}_{C_Q}([\text{enqueuer}, \text{dequeuer}])$ (where C_Q is some
 452 set of allowed initial states for the queue in Figure 5a) by verifying the wrappers in Figure 5b.
 453 We note that despite the fact that termination of queue methods follows almost immediately—
 454 they all terminate in constant time—proving their safety is highly non-trivial: reading of the
 455 head node might execute concurrently with dequeuing that involves freeing the head. We
 456 refer the reader to our Rocq development for the details of the proof.

457 **5 Proving the wait-freedom adequacy theorem**

458 As the example in Section 3 demonstrates, proving the specification in Definition 4 for a
 459 given operation is relatively straightforward. Indeed, a wait-free operation cannot block, or
 460 be delayed by other operations running concurrently, thus it must terminate in finite number
 461 of steps. To account for this, the rules of Lawyer logic simply consume one fuel resource
 462 upon every step; essentially, the user is mainly responsible for providing enough fuel initially.
 463 This applies both to the original specification in Definition 4, and its variants described in
 464 Section 4. Thus, the non-trivial parts of the proof might only concern safety of the operation
 465 (as we mentioned in Section 4.2).

466 However, showing that the specification in Definition 4 implies wait-freedom, *i.e.* proving
 467 Theorem 5, is quite challenging. Ideally, we would just reuse the existing work on termination
 468 verification in Lawyer for proving wait-freedom. That is, we would ideally reuse Lawyer’s
 469 proof that verifying a program in Lawyer establishes a refinement between the program
 470 and an always terminating transition system, which then guarantees that the program
 471 terminates. We might hope that this would help us derive termination of all calls to our
 472 wait-free operation. However, we face the following challenges:

- 473 C1 The property established by Lawyer’s adequacy theorem is termination of closed programs,
 474 *i.e.*, in our case the client program together with the wait-free data structure. In our
 475 case, we aim to prove termination of each individual call to the wait-free operation. The
 476 latter property would follow from the former, if the client program is terminating, but
 477 the client program does not have to be terminating in general.
- 478 C2 Lawyer (and Trillium in general) requires the entire closed program to be verified for
 479 the adequacy theorem to apply, and to allow us to derive a property of its execution
 480 traces. For wait-freedom, we are given an arbitrary client program (that uses the wait-free
 481 operation) and cannot verify it directly.
- 482 C3 The Lawyer logic does not contain rules that support verifying non-terminating programs.
 483 For example, there are no rules that allow generating fuel indefinitely long to be spent in
 484 an infinite execution.

485 We address these challenges as follows. First, in Section 5.1 we reduce the problem of
 486 proving wait-freedom to showing termination of the closed program. With that, we are
 487 able to reuse the Lawyer’s adequacy theorem, thus addressing the Challenge C1. Then, in
 488 Section 5.2 we employ the well-known method of logical relations to establish “robust safety”
 489 of the wait-free operation, *i.e.*, we prove that an arbitrary client does not violate operation’s
 490 internal invariants. This solves the Challenge C2. Finally, in Section 5.3 we show that the
 491 aforementioned reduction in fact allows to apply the Lawyer logic to generate fuel consumed
 492 by an infinitely running program. Together with the robust safety result, it allows to verify
 493 a non-terminating program, which solves Challenge C3.

XX:16 Verifying wait-freedom for concurrent higher-order programs

494 Finally, we note that Theorem 5 and its variants for possibly-stuck wait-freedom from
495 Section 4.1 are proven in the same way with minimal variations. However, proving Theorem 9
496 requires non-trivial changes to the logical relation and the way it is used. For space reasons,
497 we refer the reader to technical Rocq development for that. Therefore, the rest of the section
498 only explains how the original Theorem 5 is proven.

499 5.1 Reducing wait-freedom to termination

500 The first step in proving the wait-freedom adequacy theorem is to reduce proving wait-freedom
501 to showing termination of the closed program. For that, the proof of Theorem 5 proceeds via
502 a proof by contradiction. Namely, for the operation op and some set of starting configurations
503 $C \subseteq \text{Config}$ we assume that an execution trace etr , evaluation context E , thread τ and trace
504 index i such that:

- 505 ■ $\text{WaitFreeSpec}_C(\text{op})$ is provable in Lawyer.
- 506 ■ The premises of both $\text{waitFree}_C(\text{op})$ and $\text{alwaysReturnsInTrace}(\text{op}, \text{etr}, \tau)$ hold. In partic-
507 ular, a call to op occurs in etr at index i , thread τ under evaluation context E .
- 508 ■ Yet, the conclusion of $\text{alwaysReturnsInTrace}(\text{op}, \text{etr}, \tau)$ (i.e. $\exists j \geq i. \text{Ret}(\text{etr}[j].1, \tau, E)$)
509 does not hold.

510 Note that the only way this is possible is in the case where etr is infinite. Indeed, if the
511 call to op never returns, then by the schedUntilRet assumption of $\text{alwaysReturnsInTrace}$ the
512 calling thread τ is always eventually scheduled. We will prove the contradiction by showing
513 that the same set of assumptions simultaneously implies that etr must terminate.

514 To prove that etr terminates, we refine it to a trace of the Obligations Model of Lawyer;
515 the adequacy theorem of Lawyer implies the finiteness of both traces. For establishing the
516 refinement, we need to verify the program (i.e., each element of the expressions in the starting
517 thread pool, $\text{etr}[0].1$). As mentioned in Challenge C2 and Challenge C3, the Lawyer logic
518 can only support terminating programs, whereas we already concluded above that etr is
519 infinite, and moreover we cannot apply the rules of Lawyer directly, as the program to verify
520 (a client of the wait-free operation) is universally quantified over.

521 The key to establishing the refinement is exploiting the assumptions above. Indeed, in
522 Section 5.3 we show how to reuse Lawyer logic and construct a refinement *assuming* the
523 premises above and given the specification for the wait-free operation. In other words, we
524 establish the refinement *only* for execution traces that satisfy the premises above. This is in
525 contrast with the adequacy theorem of Trillium (and Lawyer) which establish the refinement
526 for *all* execution traces. Therefore, we need to generalize the adequacy theorem of Trillium
527 to establish a refinement only when some condition on execution holds. In return, the fact
528 that this condition holds is propagated to the refinement construction via user-provided
529 “progress resource” which we describe in more detail in Section 5.3. Below, we provide the
530 generalized adequacy theorem of Trillium and [highlight](#) the main changes compared to the
531 original theorem.

532 ► **Theorem 10** (Conditional adequacy theorem of Trillium). *Consider a transition system \mathcal{M}
533 and a relative image-finite relation ξ on finite traces of $\lambda_{\text{ref}}^{\text{conc}}$ and the model. Then, consider
534 $\text{TI} : \text{FinTrace}(\lambda_{\text{ref}}^{\text{conc}}) \rightarrow \text{FinTrace}(\mathcal{M}) \rightarrow \text{iProp}$ — a representation of the aforementioned
535 finite traces in the Iris logic. Then, consider a predicate F on finite execution traces. Finally,
536 consider a “progress resource” $\text{PR}_F : \text{FinTrace}(\lambda_{\text{ref}}^{\text{conc}}) \rightarrow \text{iProp}$ that satisfies a number of
537 laws described in [appendix\[2\]](#).*

538 *Let etr be an execution trace and m a state of \mathcal{M} . If $\xi(etr[0], m)$ holds (treating its*
 539 *arguments as singleton traces), and furthermore we have*

$$540 \quad \vdash \text{TI}(etr[0], m) * \text{AlwaysHolds}(\xi, etr[0], m) * \\ 541 \quad \text{PR}_F(etr[0]) \text{ (* treating } etr[0] \text{ as a singleton trace *)}$$

542 *and $F(etr')$ holds for any finite prefix etr' of etr , then there exists such model trace mtr*
 543 *that $mtr[0] = m$ and $\hat{\xi}(etr, mtr)$ holds.*

544 We refer the reader to Trillium [35] for definitions of `AlwaysHolds` and $\hat{\xi}$; the latter is
 545 essentially the refinement we're looking for. Intuitively, the three resources in the Iris premise
 546 of Theorem 10 correspondingly describe *what* is the current state of the execution, *why* this
 547 state representation implies refinement at the meta-level, and *how* this representation is
 548 preserved on every step of the execution.

549 We prove the finiteness of our *etr* by applying Theorem 10 and establishing the refine-
 550 ment to the Obligations Model. For that, we choose the trace interpretation, the starting
 551 model state m , and the refinement relation that the Lawyer's adequacy theorem does when
 552 it uses Trillium's adequacy theorem (the latter being slightly adjusted — see our Rocq
 553 implementation for details). We defer the choice of the predicate F and the progress
 554 resource to Section 5.3; we only note that the chosen F indeed holds for all prefixes of
 555 *etr*. Thus, it remains to prove the Iris premise of Theorem 10. There, $\text{TI}(etr[0], m)$ and
 556 $\text{AlwaysHolds}(\xi, etr[0], m)$ are proved similarly to how it is done in Lawyer. We will explain
 557 how to establish the progress resource in Section 5.3. Thus, using the Trillium's adequacy
 558 theorem we prove that *etr* is finite, and thus derive a contradiction.

559 5.2 Logical relation for ensuring “robust safety” of a wait-free operation

560 As the Challenge C2 states, using a Lawyer-based approach requires us to verify the entire
 561 program under consideration, including the arbitrary and thus unknown client of the operation.
 562 Moreover, as the Challenge C3 states, the Lawyer logic in general does not support verification
 563 of non-terminating programs which we would like to consider as clients of wait-free operations.

564 However, it turns out that verifying the operation's client can be split into two smaller
 565 steps. The first step is proving the *robust safety* of the operation, *i.e.* the fact that any
 566 operation's client preserves operation's internal invariants. The next step, described further
 567 in Section 5.3, is showing that any execution of a client can be represented in the Obligations
 568 Model of Lawyer.

569 To prove robust safety, we employ the well-known idea of logical relations [34, 13]. This
 570 approach is based on proving, for a given expression e , an Iris proposition $\text{RS}_{\mathbf{E}}(e)$; the
 571 predicate $\text{RS}_{\mathbf{E}}(\cdot)$ is called the *expression relation*. The expression relation intuitively captures
 572 that any execution of e preserves all invariants in the system (including those of the wait-free
 573 operation), but it may get stuck. Crucially, the expression relation is proven to be closed
 574 under program compositions, *e.g.*, if both e_1 and e_2 are in the expression relation, so is the
 575 expression $e_1 \ e_2$, *i.e.*, calling the function e_1 with argument e_2 —intuitively, the worst that
 576 can happen here is e_1 not a function in which case $e_1 \ e_2$ gets stuck, which is allowed. This
 577 compositionality of the expression relation is the reason why to establish robust safety of
 578 the operation it is sufficient to prove $\text{RS}_{\mathbf{E}}(e)$ for all clients e of the operation. In Section 3
 579 we described a `PresModInv` Hoare triple that states a property similar to $\text{RS}_{\mathbf{E}}(\cdot)$. Indeed,
 580 expression relation is defined similarly to `PresModInv` (we explain $\text{RS}_{\mathbf{V}}(v)$ below):

$$581 \quad \text{RS}_{\mathbf{E}}(e) \triangleq \forall \tau. \{ \text{True} \}_{\text{noMod}}^{\xi} e \{ v. \text{RS}_{\mathbf{V}}(v) \}^{\tau}$$

XX:18 Verifying wait-freedom for concurrent higher-order programs

582 To understand how the expression relation is proven, first consider some e that does not
 583 actually call the wait-free operation. The potentially unsafe steps of e execution are writing
 584 to memory location (or freeing it) or evaluating an arbitrary, potentially unsafe lambda
 585 expression. For these steps to be safe, we need to ensure that the corresponding location
 586 or lambda expression are also “safe to use”. More generally, we introduce a *value relation*
 587 $\text{RS}_{\mathbf{V}}(v)$ defined as follows:⁵

► **Definition 11** (Value relation).

$$\begin{aligned}
 588 \quad & \text{RS}_{\mathbf{V}}(()) \triangleq \text{RS}_{\mathbf{V}}(b : \mathbb{B}) \triangleq \text{RS}_{\mathbf{V}}(n : \mathbb{Z}) \triangleq \text{True} \\
 589 \quad & \text{RS}_{\mathbf{V}}(\mathbf{inl} \ v) \triangleq \text{RS}_{\mathbf{V}}(\mathbf{inr} \ v) \triangleq \text{RS}_{\mathbf{V}}(v) \\
 590 \quad & \text{RS}_{\mathbf{V}}((v_1, v_2)) \triangleq \text{RS}_{\mathbf{V}}(v_1) * \text{RS}_{\mathbf{V}}(v_2) \\
 591 \quad & \text{RS}_{\mathbf{V}}(\ell : \text{Loc}) \triangleq \boxed{(\exists v : \text{Val}. \ell \mapsto v * \text{RS}_{\mathbf{V}}(v)) \vee \text{freed}(\ell)} \\
 592 \quad & \text{RS}_{\mathbf{V}}(\mathbf{rec} \ f \ x := e) \triangleq \forall \tau, v. \{\text{RS}_{\mathbf{V}}(v)\}_{\text{noMod}}^{\ddagger} (\mathbf{rec} \ f \ x := e)v \{r. \text{RS}_{\mathbf{V}}(r)\}^{\tau}
 \end{aligned}$$

593 Note that the last two cases of this definition reference $\text{RS}_{\mathbf{V}}(v')$ for some value v' which
 594 might be structurally larger than one being considered. The soundness of such definitions is
 595 guaranteed by the guarded recursion of Iris; we refer the reader to [23] for more details.

596 Thus, proving $\text{RS}_{\mathbf{E}}(e)$ for an expression e that does not call the wait-free operation
 597 involves keeping track of the values produced by executing e so far and making sure they
 598 all satisfy the value relation. With that, if e does not contain hard-coded locations initially
 599 and does not use pointer arithmetic, we can indeed prove $\text{RS}_{\mathbf{E}}(e)$ by structural induction⁶,
 600 showing that evaluating its sub-expressions produce values that can be safely used afterwards.

601 However, for an e that actually calls a wait-free operation \mathbf{op} , it requires some extra
 602 work. Indeed, values produced during \mathbf{op} execution might not satisfy the value relation (*e.g.*
 603 locations used by \mathbf{op} internally are covered by \mathbf{op} ’s own invariants), and also \mathbf{op} might use
 604 pointer arithmetic (*e.g.* queue in Figure 5a uses it for accessing \mathbf{Node} structure fields). Thus,
 605 we cannot automatically derive expression relation for calls to \mathbf{op} —instead, the user must
 606 show it. And indeed, the $\text{PresModInv}(\mathbf{op})$ part of \mathbf{op} wait-free specification is exactly the
 607 lambda-expression case of Definition 11, stating that it is always safe to call \mathbf{op} with any
 608 argument. Another way to think about this is that showing $\text{PresModInv}(\mathbf{op})$ is also required
 609 because it ensures that the operation does not expose its own internal state, which would of
 610 course not be safe.

611 A more general fact is stated by the “fundamental theorem” of logical relations. Namely,
 612 that the expression relation holds for any expressions whose free variables are substituted
 613 with values that are themselves in the value relation:

614 ► **Theorem 12** (Fundamental theorem). *Let e be an expression that does not contain any Loc*
 615 *literals, or any pointer arithmetic, and whose free variables are $\vec{x}\$$. Given any list of values*
 616 *$\vec{v}\$$ such that $|\vec{v}\$| = |\vec{x}\$|$. Then the following holds:*

$$617 \quad \left(\bigstar_{v \in \vec{v}\$} \text{RS}_{\mathbf{V}}(v) \right) \vdash \text{RS}_{\mathbf{E}}(e[\vec{v}\$/\vec{x}\$])$$

618 A client of the operation \mathbf{op} can be represented as a thread pool made of expressions into
 619 which \mathbf{op} is substituted (see the first premise of Definition 3). To prove robust safety of \mathbf{op} ,

⁵ The resource $\text{freed}(\ell)$ is obtained after executing $\mathbf{free} \ \ell$ and is irrelevant for the rest of the paper.

⁶ For lambda expressions, it also involves so-called Löb induction available in step-indexed logics such as Iris.

620 we apply Theorem 12 and provide $\text{PresModInv}(\text{op})$ to satisfy its premise. As a result, we
 621 establish the following:

622 ► **Theorem 13** (Robust safety of the wait-free operation). *Let \mathcal{E} be a thread pool such that its*
 623 *every element e has a single free variable x , and does not contain any Loc literals, or any*
 624 *pointer arithmetic. Assume that $\text{WaitFreeSpec}_C(\text{op})$ holds for some C . Then the following*
 625 *holds:*

$$626 \quad \vdash \bigstar_{e \in \mathcal{E}} \text{RS}_{\mathbf{E}}(e[\text{op}/x])$$

627 5.3 Refining an infinite execution to a trace of Obligations Model

628 In Section 5.1, we have reduced proving wait-freedom to that of termination, which we
 629 establish by applying Theorem 10 and refining a given execution to the trace of Obligations
 630 Model of Lawyer. However, that reduction, by contradiction, considers an *infinite* execution,
 631 whereas the main property of Obligations Model is the finiteness of its traces. As the
 632 Challenge C3 states, the Lawyer logic that establishes refinement of the program by the
 633 Obligations Model was designed to only verify terminating programs. Thus, in this section
 634 we briefly describe the solution that allows us to repurpose the Lawyer logic for verifying
 635 non-terminating programs. Due to lack of space we refer the reader to our Rocq formalization
 636 [1] for the details.

637 The central part in establishing refinement by verifying a program is the progress resource
 638 used by Theorem 10. Intuitively, the resource $\text{PR}_F(\text{etr})$ means that if the finite execution
 639 trace etr is currently refined by some trace of the Obligations Model, and if it takes a step
 640 into a new trace etr' such that $F(\text{etr}')$ holds, then etr' will also refine some Obligations
 641 Model trace, all the invariants are preserved, and $\text{PR}_F(\text{etr}')$ is reestablished. By providing
 642 the progress resource for a singleton trace $\text{PR}_F(c)$, one essentially proves that any execution
 643 starting from the configuration c is refined by a trace of the Obligations Model, as long as it
 644 does not violate the predicate F .

645 Crucially, the user is free to choose an arbitrary progress resource, as long as it satisfies a
 646 number of laws listed in the appendix [2]. The original adequacy theorem of Lawyer can
 647 be stated as a corollary of Theorem 10 where the progress resource ignores the predicate
 648 F , and for the starting configuration it consists of Hoare triples for every thread present in
 649 that configuration. In our case, we cannot prove such triples, as we consider non-terminating
 650 client programs, and the Lawyer logic is designed to establish termination.

651 Thus, the progress resource for wait-freedom consists of two parts: one that guarantees
 652 preservation of invariants and another that ensures the refinement of the Obligations Model.
 653 The former is given by the robust safety result of Theorem 13. The construction of the latter
 654 is more elaborate.

655 The main part for establishing a refinement of the Obligations Model is providing a barrel
 656 of fuel on every execution step. For that, we exploit the assumptions on the execution that
 657 were obtained during the reduction in Section 5.1. Namely, that our progress resource is
 658 defined as $F(\text{etr}) \triangleq \text{infCallPrefix}(E, \tau, i, \text{etr})$, where E, τ and i are fixed during the reduction.
 659 We refer the reader to Rocq development [1] for its formal definition. Intuitively, this predicate
 660 means that etr is a possible finite prefix of some infinite trace in which thread τ has an
 661 infinite call to op under evaluation context E at index i . Now, knowing i , we can pre-allocate
 662 i barrels of fuel in the initial progress resource, so that the refinement can be preserved for
 663 the first i steps before the infinite call.

664 To provide fuel for the rest of execution, we first note that the thread τ that runs that call
 665 does not need infinite fuel: the specification $\text{NoInfExec}(\text{op})$ states that any call to op refines
 666 a trace of the Obligations Model, as long as a fixed amount of fuel is provided to satisfy the
 667 precondition of that specification. Therefore, we add this fuel to the initial progress resource.

668 To provide fuel to all other threads, we employ a trick that allows to *generate fuel*
 669 *indefinitely*. That is, we use the Lawyer logic’s obligations mechanism originally designed for
 670 verifying blocking concurrency. We refer the reader to Lawyer [31] for more details; in short,
 671 threads can wait for other threads’ actions by generating fuel, as long as this waiting is not
 672 circular. In our case, we can assign the calling thread τ an obligation that is never fulfilled
 673 which does not correspond to any action, but merely enables the fuel generation mechanism.
 674 This lets all other threads to wait for this obligation to be fulfilled. By assumption, τ never
 675 terminates, and thus the other threads can always receive fuel. Of course, this trick is
 676 only possible because we started the proof in Section 5.1 by contradiction—under normal
 677 circumstances, a thread with an obligation will eventually fulfill it and thus disable fuel
 678 generation.

679 **6 Discussion and limitations**

680 **6.1 Lawyer specifications of non-blocking properties**

681 In Section 3, we claimed that the specification $\text{NoInfExec}(\text{op})$ captures that all calls to op
 682 must terminate. We, furthermore claimed that $\text{PresModInv}(\text{op})$ implies that op does not
 683 rely on any invariants that involve Lawyer resources. Thus, according to Liang and Feng
 684 [28], these conditions together imply that op is wait-free as it must never block itself, nor
 685 can it delay other threads. Below we will explain how these specifications prohibit blocking
 686 and delaying, and discuss the changes needed in our approach in order to support other
 687 non-blocking properties.

688 To understand why $\text{NoInfExec}(\text{op})$ prohibits blocking, we need to explain what this
 689 specification *does not* include. Remember that it is defined as follows, for some fixed $\mathbf{F} : \mathbb{N}$:

$$690 \quad \forall \tau : \text{Thread}, \pi : \text{Phase}, a : \text{Val}. \{ \mathbf{F} \cdot \text{bar}_o \pi * \text{ph}_o \tau \pi \} \text{op } a \{ _ . \text{ph}_o \tau \pi \}^\tau$$

691 Neither its precondition- nor postcondition mention any obligation resource $\text{obls}_\tau \mathcal{O}$ of
 692 Lawyer. This resource indicates that the thread τ holds a set \mathcal{O} of *obligations* that it must
 693 fulfill, *e.g.*, setting a shared flag that other threads can wait for. In Lawyer, a thread having
 694 an obligation allows the other threads to busy-wait for this obligation to be fulfilled. In
 695 practice, waiting of another thread’s obligation means that the waiting thread can *generate*
 696 fuel to be spent during its busy-waiting. To prevent circular waiting, a thread can only
 697 generate fuel by presenting its own set of obligations, and crucially proving that all their
 698 obligations are “ordered after” the obligation (of the other thread) that it waits for.⁷ Thus,
 699 $\text{NoInfExec}(\text{op})$ not including obligations resources means that it is non-blocking, because it
 700 cannot generate fuel by waiting for other threads’ obligations, and thus it can only spend
 701 the fuel that is explicitly provided to it in $\text{NoInfExec}(\text{op})$.

702 A delay in a lock-free algorithm is a situation when the successful completion of one
 703 operation makes another operation running concurrently with it to fail and have to retry.
 704 An example of it can be seen in Figure 1b where a successful **CAS** in one thread can cause

⁷ For more details of what this “ordered after” relation is see the Lawyer paper. We only emphasize here that this order is well-founded.

705 another thread to fail its own **CAS** which forces it attempt to increment again by making
 706 a recursive call. To prohibit delays, our approach disables the fuel sharing mechanism of
 707 Lawyer. This mechanism, introduced by [19], requires a thread that delays others to provide
 708 them with extra fuel: since the delayed thread will retry its operation, it will need more fuel.
 709 In Lawyer (see [31, Delaying example]), the fuel is shared via an invariant that is accessible
 710 by both the delaying thread, and the delayed thread. To disable fuel sharing, we simply
 711 require that invariants of the wait-free operation do not contain any Lawyer resources.⁸

712 Besides wait-freedom, there are two other non-blocking properties: obstruction-freedom
 713 [16] and lock-freedom. The former only requires an operation to progress if it is delayed by
 714 other operations a finite number of times. When this condition is violated, an obstruction-
 715 free operation might block. Thus, we believe that adapting our approach to verifying
 716 obstruction-freedom could be achieved by adding obligation resources to the specification.
 717 For lock-freedom, one could follow the approach of Lawyer and enable fuel sharing by
 718 allowing fuel sharing, *i.e.* by lifting the aforementioned restriction on invariants. However,
 719 the soundness of fuel sharing in Lawyer is relies on the “phases” mechanism [19], which only
 720 allows to share fuel among the threads that have the same ancestor in the fork tree. The
 721 definition of lock-freedom does not place any restrictions on threads’ ancestry, and therefore
 722 this restriction too must be lifted from the fuel sharing mechanism.

723 6.2 Restricting the content of wait-free data structures

724 When verifying the queue implementation in Section 4.2, we introduced a seemingly unnec-
 725 essary constraint. Namely, we defined **enqueueer** wrapper in Figure 5b in a way that only
 726 allows the integer values to be enqueued. Note that the original implementation in Figure 5a
 727 never inspects the enqueued values and thus does not rely on this constraint. However, in
 728 short, lifting this constraint would imply adding an extra premise to the precondition of
 729 **NoInExec** which must hold for the argument — but since we show termination of calls to an
 730 operation with *any* argument, we would not be able to prove it when using **NoInExec** in the
 731 adequacy proof.

732 Before proceeding to details of this constraint, we explain its scope. Despite arising when
 733 verifying restricted wait-freedom in Section 4.2, the issue is not related to tokens, or the
 734 restricted wait-freedom in general. However, it arises with all data structures that contain
 735 methods that return values previously stored by other methods (queues, stacks, linked lists
 736 *etc.*). In our context, such structures happen to be restricted wait-free.

737 Recall that establishing robust safety of **dequeuer** in Figure 5a using Theorem 13 requires
 738 proving the $\text{PresModInv}(\text{dequeuer})$ specification, which is the lambda-expression case of
 739 Definition 11. (For technical reasons, Definition 8 requires a stronger specification from
 740 which $\text{PresModInv}(\text{dequeuer})$ is derivable, but the issue would still hold without this stronger
 741 specification.) Note that when **dequeuer** returns a dequeued value, it cannot, by itself, prove
 742 that it is safe-to-use. For this to hold, the queue invariant must enforce that all values stored
 743 are safe-to-use by storing the $\text{RS}_{\mathbf{V}}(v)$ resources for every stored value v . To maintain this
 744 invariant, **enqueueer** must only call **enqueue** with safe-to-use values — and it indeed does, as
 745 only integers are being enqueued.

746 If we remove the only-integers constraint, the specifications for **enqueueer** would have
 747 to include the $\text{RS}_{\mathbf{V}}(v)$ premise in the precondition. In particular, it would appear in the
 748 precondition of $\text{NoInExec}(\text{enqueueer})$. When proving wait-freedom, we consider a potentially

⁸ More formally, we require that the invariant can be defined in non-Lawyer instantiations of Trillium.

XX:22 Verifying wait-freedom for concurrent higher-order programs

749 non-terminating call to `enqueueer` with some argument a . Internally, the adequacy theorem
750 uses (the token-based version of) `NoInfExec(enqueueer)` to verify the call `enqueueer a`. Since
751 $RS_{\mathbf{V}}(a)$ is not provable in general for arbitrary a , we would not be able to use this specification.

752 To lift this limitation, one has to exploit the fact that a client program satisfying the
753 conditions of Theorem 13 can only produce safe-to-use values as it executes (except when
754 the wait-free operation is in progress). We leave establishing this formally to future work.

755 6.3 Proving linearizability of wait-free data structures

756 An important criteria for the correctness of a concurrent data structure is *linearizability* [18].
757 A data structure is linearizable if whenever operations on that structure are run concurrently,
758 each of them appear to take effect atomically, even if they actually take multiple steps to
759 finish. Linearizability makes it easier to reason about the data structure as one does not have
760 to consider all possible interleavings of the operations—one can simply reason as though
761 these operations happen one after another. It is thus not surprising that many wait-free data
762 structures are designed to be linearizable [21, 27, 36, 14]. Therefore, it is desirable to verify
763 both wait-freedom and linearizability for wait-free data structure.

764 In fact, proving linearizability with program logics is well-studied [20, 8, 25, 4]. Birkedal *et*
765 *al.* [4] have formally shown that linearizability is internalized by specifications given in terms
766 of the so-called *logically atomic Hoare triples* [25]. Moreover, the Lawyer approach [31] that
767 we build upon provides a liveness-aware version of the logically atomic specification. However,
768 even simple wait-free implementations contain what is known as *future-dependent linearization*
769 *points*[37] where the point at which linearization takes place within an operation can depend
770 on how the other operations started afterwards execute. Moreover, the linearization point
771 of an operation might happen inside another operation that runs concurrently with it in
772 a different thread. It is the case for the queue in Figure 5a. Consider the scenario where
773 the dequeuer thread runs a number of `dequeue` operations one by one, while the enqueueer
774 thread is running the single `read_head_e` operation. In that case, the linearization point of
775 `read_head_e`, if the comparison after the second read of `Head` fails, will occur *when the last*
776 *dequeue operation writes to OldHeadVal*.

777 Iris supports verifying future-dependent linearizability with using so-called prophecy
778 variables [24]. However, as is, Lawyer does not support prophecy variables because Trillium
779 [35] does not. Thus, verifying both wait-freedom and linearizability in our framework would
780 require extending the Lawyer logic and the Trillium logic first.

781 6.4 Wait-free operations of a lock-free data structure

782 Some data structures provide both lock-free and wait-free methods. For example, an
783 implementation of counter can include a lock-free increment from Figure 1b and a wait-free
784 `read` method that simply reads the current counter value in a single step. Thus, in general,
785 wait-freedom is a property of a single method, not an entire data structure. Our approach
786 could also be used for verifying wait-freedom in such a setting as well. First, the meta-level
787 definition of wait-freedom should mention both a wait-free operation, and a number of
788 lock-free operations that may run concurrently with it (similarly to how Definition 7 applies
789 to the list of wait-free operations). As for the specifications, the approach suggested by Total
790 TaDA [9] applies: the operations running concurrently with the wait-free one only need to
791 preserve invariants. With that, we can prove a generalization of Theorem 5 that establishes
792 wait-freedom of an operation `op` given `WaitFreeSpec(op)` and proofs of `PresModInv(op')` for
793 every lock-free operation `op'` running concurrently with `op`.

794 **7** Related work

795 All prior solutions verify wait-freedom for first-order languages. These solutions all guarantee
796 wait-freedom under the assumption, in one way or another, that the client of a wait-free
797 operation is safe, *i.e.* that the client does not violate the data structure’s internal invariants.
798 In our case, we prove once and for all that all possible clients in our higher-order programming
799 language satisfy this property via a logical relations argument as explained in Section 5.2.
800 Moreover, most of the prior solutions only support unrestricted wait-freedom, lack case
801 studies, and are not mechanized.

802 **7.1 Program logics for wait-freedom**

803 **Total TaDA** The program logic of [9] supports verification of non-blocking program prop-
804 erties, including unrestricted wait-freedom. Proving a program specification in that logic
805 establishes a meta-level judgement about program execution. This judgement is defined as
806 a least fixed point and thus guarantees the program’s termination. The non-trivial (from
807 liveness perspective) rules of the Total TaDA logic allow verification of both while loops
808 and recursive functions by choosing a decreasing measure of termination. Moreover, their
809 specifications are also capable of expressing the linearizability, which we do not consider.

810 Da Rocha Pinto *et al.* [9] demonstrate their approach by verifying wait-freedom of `read`
811 operation of a counter.

812 Total TaDA assumes that the operation’s client is safe by employing a rely relation [22]
813 into the aforementioned meta-level judgement established for the operation. The only way
814 to ensure that the client respects the rely relation is to verify it; thus their definition of
815 wait-freedom does not quantify over arbitrary client as ours does. However, they support
816 clients verified with the non-total TaDA [8] logic. Contrary to that, our approach supports
817 arbitrary client programs (up to assumptions of Theorem 12). Finally, Total Tada is not
818 mechanized.

819 **Lili** The approach of [28] verifies liveness of both blocking and non-blocking operations of
820 concurrent objects, as well as their linearizability. In that work, (unrestricted) wait-freedom
821 is treated as starvation-freedom (a property usually applied to blocking algorithms) without
822 the fair scheduling assumption. The Lili logic is parametric, and by instantiating it in a
823 specific way one can exclude rules that justify threads being delayed and/or blocked by each
824 other. This is similar to how Lawyer proves *e.g.* termination under unfair scheduler with a
825 more restrictive instantiation of Obligations Model.

826 The adequacy statement of Lili is similar to ours. It states that for an arbitrary execution
827 trace starting from an arbitrary client program configuration, and an “appropriate” state,
828 all calls to the object’s methods terminate. However, the major difference compared to
829 our approach is that the physical state of the concurrent object and its client are explicitly
830 separated. That way, in Lili the invariants of the concurrent object are trivially preserved by
831 its clients who do not have access to them at all. In our work, our higher-order language
832 does not allow us easily state and use such a separation, and thus we use our logical relations
833 model to prove safety of arbitrary client.

834 The Lili logic is neither higher-order nor mechanized. Moreover, neither Lili, nor the
835 lock-free fragment of its predecessor [29] have any case studies on wait-freedom verification.

836 **Automatic tool for verifying non-blocking algorithms** The work of [15] verifies progress
837 properties of non-blocking algorithms. They notice that progress of these algorithms can

838 be established by identifying a set of ordered abstract actions. An action at a higher level
 839 cannot be executed infinitely often unless there is an action of lower level which is also
 840 executed infinitely often. In this setting, (unrestricted) wait-freedom corresponds to discrete
 841 orders, *i.e.* no action can ever be executed infinitely often. In the absence of non-trivial
 842 orders between actions and unbounded `while` loops, Gotsman *et al.* [15] reduce verifying
 843 wait-freedom to verifying safety.

844 This solution relies on an automatic tool that establishes safety of a given operation.
 845 Moreover, it establishes the rely relation that the operation’s client is expected to preserve.
 846 Similarly to Total TaDA, this means that wait-freedom is only established for verified clients.

847 Gotsman *et al.* [15] present no wait-freedom case studies. They explain that their
 848 approach might not apply to non-trivial wait-free operations due to the limitations of their
 849 tool. As an example, they mention the wait-free `contains` operation of the concurrent set
 850 algorithm by [38]. The `while` loop in that operation traverses a sorted list of integers up to
 851 a given number. As the list can be concurrently modified during the traversal, termination
 852 of such loop relies on the list being sorted. The tool by [15] cannot exploit this fact, whereas
 853 in our approach the termination of such loops can be shown by providing the amount of fuel
 854 proportional to the given number.

855 Moreover, their tool has limited safety reasoning capabilities and thus cannot verify
 856 examples like Figure 5a where the safety argument is highly non-trivial. The tool is automated,
 857 but the presented theory is not mechanized in a proof assistant.

858 **Wait-freedom of “fast-path-slow-path” in VST** The work of [32] considers non-blocking
 859 algorithms that are instances of a specific variation of the “fast-path-slow-path” approach
 860 [27]. This approach aims to design efficient wait-free data structures by combining a lock-free
 861 implementation (which is often faster, but does not have wait-free progress guarantee) and a
 862 wait-free implementation (slower, but has guaranteed progress).

863 The authors of [32] establish a pen-and-paper proof that such algorithms, if safe, are
 864 indeed wait-free, and provide an extension of the VST logic [3] to verify their safety, and thus
 865 also wait-freedom. This approach limits the number of concurrent threads, and therefore it
 866 can only be used to establish restricted wait-freedom. Similar to Total TaDA and Gotsman
 867 *et al.* [15], such approach requires the user to verify the entire program.

868 As a case study, Peterson *et al.* [32] verify wait-freedom of the queue by Kogan and
 869 Petrank [27], which is in turn an optimization of the queue by Kogan and Petrank [26].
 870 In the implementation by [27], a thread helps up to N other concurrent operations before
 871 performing its own operation (where N is the maximal number of threads). We believe
 872 that our approach can also support verifying this case study by requiring an amount of fuel
 873 proportional to N .

874 7.2 Universal constructions

875 Designing a wait-free implementation for every possible data structure is tedious, and therefore
 876 a question arises as to whether it is possible to design a universal wait-free construction.
 877 And indeed, that is possible. These are algorithms that turn any *sequential* data structure
 878 implementation into a (restricted) wait-free one. The early universal constructions such as
 879 [17] are considered impractical due to their high time and space complexity of the resulting
 880 wait-free implementations. Luckily, this is less of an issue for more recent constructions
 881 mentioned below. However, none of the universal constructions has a mechanized proof of
 882 wait-freedom.

883 Therefore, it is interesting to consider whether our solution can be used to verify such
884 constructions. To do so, we would need to prove a specification for the sequential implemen-
885 tation, and use it together with a general proof of the universal construction in question; a
886 process that is similar to how we verified possibly-stuck wait-freedom of `list_map(incr l)` in
887 Section 4.1. Therefore, we believe that our approach can be applied to universal constructions
888 as well.

889 Examples of modern wait-free universal constructions include P-Sim [11] (along with its
890 extensions such as [12]) and CX [7]. The time complexity of these approaches is proportional
891 to the number of concurrent operations, and the execution time of the provided sequential
892 implementation. Given that these parameters are known, we believe that our approach can
893 verify restricted wait-freedom of the resulting implementations, as the required amount of
894 fuel is known in advance.

895 **8 Conclusion and future work**

896 In this work, we have presented the first solution to verifying wait-freedom in a higher-
897 order language. The key idea is internalizing the notion of wait-freedom in the Lawyer
898 separation logic with a particular specification pattern. To establish this connection formally
899 we presented a new adequacy theorem for Lawyer specifically for programs proven against a
900 spec in the aforementioned pattern. We also used a logical relations model to prove that a
901 wait-free operation can be safely used by an arbitrary, higher-order client. We demonstrated
902 our approach by verifying a number of examples, with the more complicated of them requiring
903 a slightly different notion of wait-freedom which we call restricted wait-freedom, and to
904 accordingly adjust the wait-freedom specification pattern. Our solution and case studies are
905 all mechanized in the Rocq Prover.

906 There are multiple notable directions of future work. First, we aim to extend our approach
907 to verify related properties, such as linearizability and bounded wait-freedom [6]. It would
908 also be interesting to verify one of the state-of-the-art universal constructions for wait-freedom.
909 Finally, we hope to be able to improve the usability of our approach by deriving the two
910 specifications needed by Definition 4 from a single, stronger specification that implies both.

911 **References**

- 912 **1** Anonymous. Supplementary rocq development for the wait-freedom paper, 2026.
- 913 **2** Anonymous. Technical appendix for the wait-freedom paper, 2026.
- 914 **3** Andrew W. Appel. Verified software toolchain. In Gilles Barthe, editor, *Programming*
915 *Languages and Systems*, pages 1–17, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- 916 **4** Lars Birkedal, Thomas Dinsdale-Young, Armaël Guéneau, Guilhem Jaber, Kasper Svendsen,
917 and Nikos Tzevelekos. Theorems for free from separation logic specifications. *Proc. ACM*
918 *Program. Lang.*, 5(ICFP), aug 2021. doi:10.1145/3473586.
- 919 **5** Fabian Bläse. Non-blocking synchronization in operating systems, 2021. URL: https://www4.cs.fau.de/Lehre/WS20/MS_AKSS/arbeiten/paper_blaese_final.pdf.
- 920 **6** Hagit Brit and Shlomo Moran. Wait-freedom vs. bounded wait-freedom in public data
921 structures (extended abstract). In *Proceedings of the Thirteenth Annual ACM Symposium*
922 *on Principles of Distributed Computing*, PODC '94, page 52–60, New York, NY, USA, 1994.
923 Association for Computing Machinery. doi:10.1145/197917.197950.
- 924 **7** Andreia Correia, Pedro Ramalhete, and Pascal Felber. A wait-free universal construction
925 for large objects. In *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and*
926 *Practice of Parallel Programming*, PPoPP '20, page 102–116, New York, NY, USA, 2020.
927 Association for Computing Machinery. doi:10.1145/3332466.3374523.
- 928

- 929 **8** Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. Tada: A logic for time
930 and data abstraction. In Richard Jones, editor, *ECOOP 2014 – Object-Oriented Programming*,
931 pages 207–231, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- 932 **9** Pedro da Rocha Pinto, Thomas Dinsdale-Young, Philippa Gardner, and Julian Suther-
933 land. Modular termination verification for non-blocking concurrency. In Peter Thie-
934 mann, editor, *Programming Languages and Systems - 25th European Symposium on Pro-
935 gramming, ESOP 2016, Held as Part of the European Joint Conferences on Theory and
936 Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceed-
937 ings*, volume 9632 of *Lecture Notes in Computer Science*, pages 176–201. Springer, 2016.
938 doi:10.1007/978-3-662-49498-1_8.
- 939 **10** Derek Dreyer, Georg Neis, and Lars Birkedal. The impact of higher-order state and control
940 effects on local relational reasoning. *SIGPLAN Not.*, 45(9):143–156, September 2010. doi:
941 10.1145/1932681.1863566.
- 942 **11** Panagiota Fatourou and Nikolaos D. Kallimanis. A highly-efficient wait-free universal con-
943 struction. In *Proceedings of the Twenty-Third Annual ACM Symposium on Parallelism in
944 Algorithms and Architectures, SPAA '11*, page 325–334, New York, NY, USA, 2011. Association
945 for Computing Machinery. doi:10.1145/1989493.1989549.
- 946 **12** Panagiota Fatourou, Nikolaos D. Kallimanis, and Eleni Kanellou. An Efficient Univer-
947 sal Construction for Large Objects. In Pascal Felber, Roy Friedman, Seth Gilbert, and
948 Avery Miller, editors, *23rd International Conference on Principles of Distributed Systems
949 (OPODIS 2019)*, volume 153 of *Leibniz International Proceedings in Informatics (LIPIcs)*,
950 pages 18:1–18:15, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Infor-
951 matik. URL: [https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.
952 OPODIS.2019.18](https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.OPODIS.2019.18), doi:10.4230/LIPIcs.OPODIS.2019.18.
- 953 **13** {Aina Linn} Georges. Designing and proving robust safety of efficient capability machine
954 programs, July 2023.
- 955 **14** Seep Goel, Pooja Aggarwal, and Smruti R. Sarangi. A wait-free stack. In Nikolaj Bjørner,
956 Sanjiva Prasad, and Laxmi Parida, editors, *Distributed Computing and Internet Technology*,
957 pages 43–55, Cham, 2016. Springer International Publishing.
- 958 **15** Alexey Gotsman, Byron Cook, Matthew Parkinson, and Viktor Vafeiadis. Proving that
959 non-blocking algorithms don’t block. *SIGPLAN Not.*, 44(1):16–28, January 2009. doi:
960 10.1145/1594834.1480886.
- 961 **16** M. Herlihy, V. Luchangco, and M. Moir. Obstruction-free synchronization: double-ended
962 queues as an example. In *23rd International Conference on Distributed Computing Systems,
963 2003. Proceedings.*, pages 522–529, 2003. doi:10.1109/ICDCS.2003.1203503.
- 964 **17** Maurice Herlihy. Wait-free synchronization. *ACM Trans. Program. Lang. Syst.*, 13(1):124–149,
965 January 1991. doi:10.1145/114005.102808.
- 966 **18** Maurice P. Herlihy and Jeannette M. Wing. Linearizability: a correctness condition for
967 concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, July 1990. doi:
968 10.1145/78969.78972.
- 969 **19** Bart Jacobs, Dragan Bosnacki, and Ruurd Kuiper. Modular termination verification of single-
970 threaded and multithreaded programs. *ACM Trans. Program. Lang. Syst.*, 40(3):12:1–12:59,
971 2018. doi:10.1145/3210258.
- 972 **20** Bart Jacobs and Frank Piessens. Expressive modular fine-grained concurrency specification.
973 In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of
974 Programming Languages, POPL '11*, page 271–282, New York, NY, USA, 2011. Association
975 for Computing Machinery. doi:10.1145/1926385.1926417.
- 976 **21** Prasad Jayanti and Srdjan Petrovic. Logarithmic-time single deleter, multiple inserter wait-free
977 queues and stacks. In Sundar Sarukkai and Sandeep Sen, editors, *FSTTCS 2005: Foundations
978 of Software Technology and Theoretical Computer Science*, pages 408–419, Berlin, Heidelberg,
979 2005. Springer Berlin Heidelberg.

- 980 22 C. B. Jones. Tentative steps toward a development method for interfering programs. *ACM*
981 *Trans. Program. Lang. Syst.*, 5(4):596–619, October 1983. doi:10.1145/69575.69577.
- 982 23 Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek
983 Dreyer. Iris from the ground up: A modular foundation for higher-order concurrent separation
984 logic. *J. Funct. Program.*, 28:e20, 2018. doi:10.1017/S0956796818000151.
- 985 24 Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy, Marianna Rapoport, Amin Timany,
986 Derek Dreyer, and Bart Jacobs. The future is ours: prophecy variables in separation logic.
987 *Proc. ACM Program. Lang.*, 4(POPL), dec 2019. doi:10.1145/3371113.
- 988 25 Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal,
989 and Derek Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent
990 reasoning. In Sriram K. Rajamani and David Walker, editors, *Proceedings of the 42nd Annual*
991 *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015,*
992 *Mumbai, India, January 15-17, 2015*, pages 637–650. ACM, 2015. doi:10.1145/2676726.
993 2676980.
- 994 26 Alex Kogan and Erez Petrank. Wait-free queues with multiple enqueueers and dequeuers.
995 *SIGPLAN Not.*, 46(8):223–234, February 2011. doi:10.1145/2038037.1941585.
- 996 27 Alex Kogan and Erez Petrank. A methodology for creating fast wait-free data structures.
997 *SIGPLAN Not.*, 47(8):141–150, February 2012. doi:10.1145/2370036.2145835.
- 998 28 Hongjin Liang and Xinyu Feng. A program logic for concurrent objects under fair scheduling.
999 In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of*
1000 *Programming Languages, POPL ’16*, page 385–399, New York, NY, USA, 2016. Association
1001 for Computing Machinery. doi:10.1145/2837614.2837635.
- 1002 29 Hongjin Liang, Xinyu Feng, and Zhong Shao. Compositional verification of termination-
1003 preserving refinement of concurrent programs. In *Proceedings of the Joint Meeting of the*
1004 *Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-*
1005 *Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS ’14,*
1006 *New York, NY, USA, 2014*. Association for Computing Machinery. doi:10.1145/2603088.
1007 2603123.
- 1008 30 Maged M. Michael and Michael L. Scott. Simple, fast, and practical non-blocking and blocking
1009 concurrent queue algorithms. In *Proceedings of the Fifteenth Annual ACM Symposium on*
1010 *Principles of Distributed Computing, PODC ’96*, page 267–275, New York, NY, USA, 1996.
1011 Association for Computing Machinery. doi:10.1145/248052.248106.
- 1012 31 Egor Namakonov, Justus Fasse, Bart Jacobs, Lars Birkedal, and Amin Timany. Lawyer: Mod-
1013 ular obligations-based liveness reasoning in higher-order impredicative concurrent separation
1014 logic, 2026.
- 1015 32 Christina Peterson, Victor Cook, and Damian Dechev. Practical progress verification of
1016 descriptor-based non-blocking data structures. In *2019 IEEE 27th International Symposium on*
1017 *Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS),*
1018 *pages 83–93, 2019*. doi:10.1109/MASCOTS.2019.00019.
- 1019 33 Simon Spies, Lennard Gäher, Daniel Gratzer, Joseph Tassarotti, Robbert Krebbers, Derek
1020 Dreyer, and Lars Birkedal. Transfinite iris: resolving an existential dilemma of step-indexed
1021 separation logic. In Stephen N. Freund and Eran Yahav, editors, *PLDI ’21: 42nd ACM*
1022 *SIGPLAN International Conference on Programming Language Design and Implementation,*
1023 *Virtual Event, Canada, June 20-25, 2021*, pages 80–95. ACM, 2021. doi:10.1145/3453483.
1024 3454031.
- 1025 34 David Swasey, Deepak Garg, and Derek Dreyer. Robust and compositional verification of
1026 object capability patterns. *Proc. ACM Program. Lang.*, 1(OOPSLA), October 2017. doi:
1027 10.1145/3133913.
- 1028 35 Amin Timany, Simon Oddershede Gregersen, Léo Stefanescu, Jonas Kastberg Hinrichsen, Léon
1029 Gondelman, Abel Nieto, and Lars Birkedal. Trillium: Higher-order concurrent and distributed
1030 separation logic for intensional refinement. *Proc. ACM Program. Lang.*, 8(POPL), jan 2024.
1031 doi:10.1145/3632851.

XX:28 Verifying wait-freedom for concurrent higher-order programs

- 1032 **36** Shahar Timnat, Anastasia Braginsky, Alex Kogan, and Erez Petrank. Wait-free linked-lists.
1033 *SIGPLAN Not.*, 47(8):309–310, February 2012. doi:10.1145/2370036.2145869.
- 1034 **37** Viktor Vafeiadis. *Modular fine-grained concurrency verification*. PhD thesis, University of
1035 Cambridge, 2008.
- 1036 **38** Martin Vechev and Eran Yahav. Deriving linearizable fine-grained concurrent objects. *SIG-*
1037 *PLAN Not.*, 43(6):125–135, June 2008. doi:10.1145/1379022.1375598.