

Staged Specification Logic for Verifying Higher-Order Imperative Programs

Darius Foo(✉)^[0000-0002-3279-5827], Yahui Song^[0000-0002-9760-5895], and
Wei-Ngan Chin^[0000-0002-9660-5682]

School of Computing, National University of Singapore, Singapore
{dariusf,yahuis,chinwn}@comp.nus.edu.sg

Abstract. Higher-order functions and imperative states are language features supported by many mainstream languages. Their combination is expressive and useful, but complicates specification and reasoning, due to the use of yet-to-be-instantiated function parameters. One inherent limitation of existing specification mechanisms is its reliance on *only two stages*: an initial stage to denote the precondition at the start of the method and a final stage to capture the postcondition. Such two-stage specifications force *abstract properties* to be imposed on unknown function parameters, leading to less precise specifications for higher-order methods. To overcome this limitation, we introduce a novel extension to Hoare logic that supports *multiple stages* for a call-by-value higher-order language with ML-like local references. Multiple stages allow the behavior of unknown function-type parameters to be captured abstractly as uninterpreted relations; and can also model the repetitive behavior of each recursion as a separate stage. In this paper, we define our staged logic with its semantics, prove its soundness and develop a new automated higher-order verifier, called Heifer, for a core ML-like language.

1 Introduction

Programs written in modern languages today are rife with higher-order functions [3, 33], but specifying and verifying them remains challenging, especially if they contain imperative effects. Consider the *foldr* function from OCaml. Here is a good specification for it in Iris [17], a state-of-the-art framework for higher-order concurrent separation logic that is built using Coq proof assistant.

$$\forall P, Inv, f, xs, l. \left\{ \begin{array}{l} (\forall x, a', ys. \{P \ x * Inv \ ys \ a'\} \ f(x, a') \ \{r. \ Inv \ (x::ys) \ r\}) \\ * \ isList \ l \ xs * all \ P \ xs * Inv \ [] \ a \end{array} \right\} \\ \text{foldr } f \ a \ l \\ \{r. \ isList \ l \ xs * Inv \ xs \ r\}$$

While this specification is conventional in weakest-precondition calculi like Iris, one might argue that that this specification is not the best possible specification for *foldr*, since it requires two *abstract properties* *Inv* and *P* to summarize the behaviour of *f*. Moreover, the input list *l* is also immutable, through the

same *isList* predicate in both its pre- and postcondition. (If mutation of list is allowed, a more complex *Inv* with an extra mutated list parameter is required.)

These abstract properties must be correspondingly instantiated for each instance of f , but unfortunately some usage scenarios (to be highlighted later in Sec 2.2) of *foldr* cannot be captured by this particular pre/post specification of Iris, despite how well-designed it was. Thus, the conventional pre/post approach to specifying higher-order functions currently suffers from possible *loss in precision* in its specifications since the presence of these abstract properties implicitly *strengthens* the preconditions for higher-order imperative methods.

This paper proposes a new logic, *Higher-Order Staged Specification Logic (HSSL)*, for specifying and verifying higher-order imperative methods. It is designed for automated verification via SMT and uses separation logic as its core stateful logic, aiming at more precise specifications of heap-based changes. While we have adopted separation logic to support heap-based mutations, *HSSL* may also be used with other base logics, such as those using dynamic frames [23]. We next provide an overview of our methodology by examples before providing formal details and an experimental evaluation of our proposal.

2 Illustrative Examples

We provide three examples to highlight the key features of our methodology.

2.1 A Simple Example

We introduce the specification logic using a simple example (Fig. 1), to highlight a key challenge we hope to solve, namely how should we specify the behavior of *hello* without pre-committing to some abstract property on f ? To do that, we can model f using an *uninterpreted relation*. We use uninterpreted relation rather than a function here in order to model both over-approximation and possible

```

1 let hello f x y =
2   x := !x + 1;
3   let r = f y in
4   let r2 = !x + r in
5   y := r2;
6   r2

```

Fig. 1: A Simple Example

side-effect. Since f is effectful and may modify arbitrary state, including the references x and y , a modular specification of *hello* must express the ordering of the call to f with respect to the other statements in it so that the caller of *hello* may reason precisely about its effects. Therefore, a first approximation is the following specification. We adopt standard separation logic pre/post assertions and extend them with *sequential composition* and *uninterpreted relations*. A final parameter (named as *res* here) is added to denote the result of each staged specification's relation (*hello* here), a convention we follow henceforth.

```

hello(f, x, y, res) =
  ∃ a · req x ↦ a;           // Stage 1: requiring x be pre-allocated
  ens[_] x ↦ a+1;         // Stage 2: ensuring x is updated
  ∃ r · f(y, r);           // Stage 3: unknown higher-order f call
  ∃ b · req x ↦ b * y ↦ -; // Stage 4: requiring x, y be pre-allocated
  ens[res] x ↦ b * y ↦ res ∧ res = b+r // Stage 5: y is updated, and x is unchanged

```

We can summarize the imperative behavior of *hello* before the call to *f* with a read from *x*, followed by a write to *x*, as captured by Stages 1-2. The same applies to the portion after the call to *f* (lines 4-6), but here we only consider the scenario when *x* and *y* are disjoint¹. Stages 4 and 5 state that memory location *x* is being read while *y* will be correspondingly updated.

The ordering of the unknown *f* call with respect to the parts before and after does matter, so the call can be seen as *stratifying* the temporal behavior of the function into *stages*. Should a specification for *f* become known, usually at a call site, its *instantiation* may lead to a staged formula with only **req/ens** stages; which can always be *compacted* into a single **req/ens** pair. We detail a *normalization* procedure to do this in Section 3.2.

As mentioned before, *f* can modify *x* despite not having direct access to it via an argument, as it could capture *x* from the environment of the caller of *hello*. To model this, we make worst-case assumptions on the footprints of unknown functions, resulting in the precondition $x \mapsto b$ in stage 4.

2.2 Pre/Post vs Staged Specifications via *foldr*

We now specify *foldr* and compare it with the Iris specification from Section 1.

<pre> 1 let rec foldr f a l = 2 match l with 3 [] => a 4 h :: t => 5 f h (foldr f a t) </pre>	<pre> foldr(f, a, l, rr) = ens[rr] l = [] ∧ rr = a ∨ ∃ x, r, l₁ · ens[_] l = x :: l₁; foldr(f, a, l₁, r); f(x, r, rr) </pre>
---	--

We model *foldr* as a recursive predicate whose body is a staged formula. The top-level disjunction represents the two possible paths that result from pattern matching. In the base case, when *l* is the empty list, and the result of *foldr* is *a*. In the recursive case, when *l* is nonempty, the specification expresses that the behavior of *foldr* is given by a recursive call to *foldr* on the tail of *l* to produce a result *r*, followed by a call to *f* with *r* to produce a value for *rr*. Crucially, we are able to represent the call to the unknown function *f* directly in the specification, without being forced to impose a stronger precondition on *f*.

foldr's specification's is actually very precise, to the point of mirroring the *foldr* program. Nevertheless, abstraction may readily be recovered by proving that this predicate entails a weaker formula, and a convenient point for this would be when the unknown function-typed parameter is instantiated at each of *foldr*'s call sites; we discuss an example of this shortly. The point of specifying *foldr* this way is that the precision of stages enables us not to have to commit to an abstraction prematurely. We should, of course, summarize as early as is appropriate to keep our proving process tractable.

¹ For simplicity, the intersection of specifications \wedge_{sp} that arises from disjoint preconditions is omitted (with some loss in precision) in the main paper, but its core mechanism is briefly described in Appendix D [13].

Recursive staged formulae are needed mainly to specify higher-order functions with unknown function-typed parameters. Otherwise, our preference is to apply summarization to obtain non-recursive staged formulae whenever unknown function-type parameters have been suitably instantiated. Under this scenario, we may still use recursive pure predicates or recursive shape predicates in order to obtain best possible modular specifications for our program code.

Now, we show how the staged specification for *foldr* can be used by proving that we can sum a list by folding it. *sum* can be specified in a similar way to *foldr*, but since this is a pure function that can be additionally checked for termination, we can automatically convert it into a *pure predicate* (without any stages or imperative side effects) to be used in (the pure fragment of) our specification logic. Termination of pure predicates is required for them to be safely used in specifications. (Techniques to check for purity and termination are well-known and thus omitted.) Also, each pure predicate may be used as either a staged predicate or a pure predicate. In case a pure predicate $p(v^*, res)$ is used as a staged predicate; its staged definition is:

$$p(v^*, res) = \mathbf{req} \text{ emp} \wedge \text{pre}(v^*); \mathbf{ens}[_] \text{ emp} \wedge p(v^*, res)$$

where $\text{pre}(v^*)$ denotes the precondition to guarantee termination and avoids exceptions. Note that $p(v^*, res)$ is overloaded to be used as either a staged predicate or a pure predicate. This is unambiguous from the context of its use.

```

6  let rec sum li =                               sum(li, res) =
7    match li with                               l=[] ^ res=0
8    | [] -> 0
9    | x :: xs -> x + sum xs                     ∨ ∃ r, l1 · l=x::l1 ^ sum(l1, r) ^ res=x+r

```

We can now re-summarize an imperative use of *foldr* with the help of *sum*.

```

10 let foldr_sum_state x xs init
11 foldr_sum_state(x, xs, init, res) =
    ∃ i, r · req x ↦ i; ens[res] x ↦ i+r ^ res=r+init ^ sum(xs, r)
12 = let g c t = x := !x + c; c + t in foldr g xs init

```

This summarization gives rise to the following entailment:

$$\begin{aligned} & \forall m, xs, init, res. \text{foldr}(g, xs, init, res) \\ \sqsubseteq & \exists i, r. \mathbf{req} x \mapsto i; \mathbf{ens}[res] x \mapsto i+r \wedge res=r+init \wedge \text{sum}(xs, r) \end{aligned}$$

We have implemented a proof system for subsumption (denoted by \sqsubseteq) between staged formulae in our verifier, called Heifer. This particular entailment can be proved automatically by induction on *xs*. While Iris's earlier pre/post specification for *foldr* can handle this example through a suitable instantiation of (*Inv - .*), it is unable to handle the following three other call instances.

```

13 let foldr_ex1 l = foldr (fun x r -> let v = !x
14                               in x := v+1; v+r) l 0
15 let foldr_ex2 l = foldr (fun x r -> assert(x+r>=0);x+r) l 0
16 let foldr_ex3 l = foldr (fun x r -> if x>=0 then x+r
17                               else raise Exc()) l 0
    
```

The first example cannot be handled since Iris's current specification for *foldr* expects its input list l to be immutable. The second example fails since the precondition required cannot be expressed using just the abstract property ($P\ x$). The last example fails because the abstract property ($Inv\ (x :: ys)\ r$) used in the postcondition of f expects its method calls to return normally. In contrast, using our approach via staged specification, we can re-summarize the above three call instances to use the following subsumed specifications.

$$\begin{aligned}
 foldr_ex1(l, res) &\sqsubseteq \exists xs \cdot \mathbf{req}\ List(l, xs) ; \exists ys \cdot \\
 &\quad \mathbf{ens}[res]\ List(l, ys) \wedge mapinc(xs, ys) \wedge sum(xs, res) \\
 foldr_ex2(l, res) &\sqsubseteq \mathbf{req}\ allSPos(l) ; \mathbf{ens}[res]\ sum(l, res) \\
 foldr_ex3(l, res) &\sqsubseteq \mathbf{ens}[res]\ allPos(l) \wedge sum(l, res) \vee (\mathbf{ens}[_]\ \neg allPos(l); Exc())
 \end{aligned}$$

Note that the first example utilizes a recursive spatial $List(l, xs)$ predicate, while the last example used $Exc()$ as a relation to model exception as a stage in our specification. The three pure predicates and one spatial predicate used in the above can be formally defined, as shown below.

$$\begin{aligned}
 mapinc(xs, ys) &= (xs = [] \wedge ys = []) \vee (\exists x, xs_1, ys_1 \cdot xs = x :: xs_1 \wedge ys = (x+1) :: ys_1 \\
 &\quad \wedge mapinc(xs_1, ys_1)) \\
 allPos(l) &= (l = []) \vee (\exists x, l_1 \cdot l = x :: l_1 \wedge allPos(l_1) \wedge x \geq 0) \\
 allSPos(l) &= (l = []) \vee (\exists x, r, l_1 \cdot l = x :: l_1 \wedge allSPos(l_1) \wedge sum(l, r) \wedge r \geq 0) \\
 List(l, rs) &= (emp \wedge l = []) \vee (\exists x, rs_1, l_1 \cdot x \mapsto r * List(l_1, rs_1) \wedge l = x :: l_1 \wedge rs = r :: rs_1)
 \end{aligned}$$

We emphasize that our proposal for staged logics is strictly more *expressive* than traditional two-stage pre/post specifications, since the latter can be viewed as an instance of staged logics. As an example, the earlier two-stage specification for *foldr* can be modelled non-recursively in our staged logics as:

$$\begin{aligned}
 foldr(f, a, l, res) &= \\
 &\exists P, Inv, xs \cdot \mathbf{req}\ List(l, xs) * Inv([], a) \wedge all(P, xs) \\
 &\quad \wedge f(x, a', r) \sqsubseteq (\exists ys \cdot \mathbf{req}\ Inv(ys, a') \wedge P(x); \mathbf{ens}[r]\ Inv(x :: ys, r)); \\
 &\mathbf{ens}[res]\ List(l, xs) * Inv(xs, res)
 \end{aligned}$$

2.3 Inferrable vs User-provided Specifications via *map*

Our methodology for higher-order functions is further explicated by the *map* method, shown in Fig. 2. Specifications typeset in lavender must be user-supplied, whereas those shown in red (with the small circle) may be automated or inferred (using the rules of Section 4). Like *sum* before, *length* and *incrg* may be viewed as *ghost* functions, written only for their specifications to be used to describe behavior. These specifications are also routine and can be mechanically derived; we elide them here and provide them in Appendix A [13]. The method

```

1  let rec length xs =
2    ◦length(xs, res) = ...
3    match xs with
4    | [] -> 0
5    | x :: xs1 ->
6      1 + length xs1
7
8  let rec incrg init li =
9    ◦incrg(init, li, res) = ...
10   match li with
11   | [] -> []
12   | x :: xs -> init ::
13     incrg (init + 1) xs
14
15 let rec map f xs =
16   ◦map(f, xs, res) = ...
17   match xs with
18   | [] -> []
19   | x :: xs1 ->
20     f x :: map f xs1
21
22 let map_incr xs x
23   = let f a = x := !x+1; !x
24     in map f xs

```

Fig. 2: Implementation of *map_incr* with a Summarized Specification from *map*

map_incr describes the scenario we are interested in, where *the state of the closure affects the result of map*. Its specification states that the pointer x must have its value incremented by the length of xs . Moreover, the *contents* of the resulting list is captured by another pure function *incrg*, which builds a list of as many increasing values as there are elements in its input list.

These examples illustrate the methodology involved with staged specifications. They inherit the modular verification and biabduction-based[4] specification inference of separation logic, adding the ability to describe imperative behavior using function stages to the mix; biabduction then doubles as a means to normalize and compact stages. There is emphasis on the inference of specifications and proof automation, and proofs are built out of simple lemmas, which help summarize behavior and the shapes of data, and either remove recursion or move it into a pure ghost function where it is easier to comprehend.

In summary, staged logic for specifying imperative higher-order functions represents a fundamentally new approach that is *more general* and yet can be *more precise* than what is currently possible via state-of-the-art pre/post specification logics for imperative higher-order methods. Our main technical contributions to support this new approach include:

1. **Higher-Order Staged Specification Logic (*HSSL*):** we design a novel program logic to specify the behaviors of imperative higher-order methods and give its formal semantics.
2. **Biabduction-based Normalization:** we propose a normalization procedure for *HSSL* that serves two purposes: (i) it allows us to produce succinct staged formulae for programs automatically, and (ii) it helps structure entailment proof obligations, allowing them to be discharged via SMT.
3. **Entailment:** we develop a proof system to solve subsumption entailments between normalized *HSSL* formulae, prove its soundness, and implement an automated prover based on it.
4. **Evaluation:** we report on initial experimental results, and present various case studies highlighting *HSSL*'s capabilities.

3 Language and Specification Logic

We target a minimal OCaml-like imperative language with higher-order functions and state. The syntax is given in Fig. 3. Expressions are in ANF (A-normal form); sequencing and control over evaluation order may be achieved using let-bindings, which define immutable variables. Mutation may occur through heap-allocated *refs*. Functions are defined by naming lambda expressions, which may be annotated with a specification Φ (covered below). For simplicity, they are always in tupled form and their calls are always fully applied. Pattern matching is encoded using recognizer functions (e.g., *is_cons*) and *if* statements. *assert* allows proofs of program properties to be carried out at arbitrary points.

(Expressions)	$e ::= v \mid x \mid \text{let } x=e_1 \text{ in } e_2 \mid f(x^*) \mid \text{ref } x \mid x_1 := x_2 \mid !x \mid x_1 :: x_2 \mid \text{assert } D \mid \text{if } x \text{ then } e_1 \text{ else } e_2$	
(Values)	$v ::= c \mid \text{nil} \mid x_1 :: x_2 \mid \text{fun } (x^*) \Phi[r] \rightarrow e$	
(Staged)	$\Phi ::= E \mid \Phi_1 \vee \Phi_2 \mid \Phi_1 ; \Phi_2 \mid \exists x^* . \Phi$	
(Stage)	$E ::= \mathbf{req} D \mid \mathbf{ens}[r] D \mid f(x^*, r)$	(State) $D ::= \sigma \wedge \pi$
(Heap)	$\sigma ::= \text{emp} \mid x_1 \mapsto x_2 \mid \sigma_1 * \sigma_2$	
(Pure)	$\pi ::= \text{true} \mid \pi_1 \vee \pi_2 \mid \neg \pi \mid \exists x . \pi \mid t_1 = t_2 \mid a_1 < a_2 \mid \Phi_1 \sqsubseteq \Phi_2$	
(A-Terms)	$a ::= i \mid x \mid a_1 + a_2 \mid -a$	
(Terms)	$t ::= \text{nil} \mid t_1 :: t_2 \mid c \mid a \mid f \mid \lambda (x^*, r) \rightarrow \Phi$	
	$c \in \mathbb{B} \cup \mathbb{Z} \cup \mathbf{unit}$	$i \in \mathbb{Z}$
	<hr style="width: 100%; border: 0.5px solid black;"/>	
		$x, f, r \in \text{var}$

Fig. 3: Syntax of the Core Language and Staged Logics

Program behavior is specified using *staged formulae* Φ , which are disjunctions and/or sequences of *stages* E . A stage is an assertion about program state *at a specific point*. Each stage takes one of three forms: a precondition $\mathbf{req} D$, a postcondition $\mathbf{ens}[r] D$ with a named result r , or a *function stage* $f(v^*, r)$, representing the specification of a (possibly-unknown) function call. For brevity, we use a context notation $\Phi[r]$ where r explicitly identifies the final result of specification Φ . Program states D are described using separation logic formulae from the *symbolic heap* fragment [4], without recursive spatial predicates (for simplicity of presentation). Most values of the core language are as usual also terms of the (pure) logic; a notable exception is the lambda expression $\text{fun } (x^*) \Phi[r] \rightarrow e$, which occurs in the logic as $\lambda (x^*, r) \rightarrow \Phi[r]$, without its body. Subsumption assertions between two staged formulae (Sec 5) are denoted by $\Phi_1 \sqsubseteq \Phi_2$.

3.1 Semantics of Staged Formulae

From Triples to Stages Staged formulae generalize standard Hoare triples. The standard partial-correctness interpretation of the separation logic Hoare triple $\{ P(v^*, x^*) \} e \{ \exists y^* . Q(v^*, x^*, y^*, \text{res}) \}$ where v^* denote valid program variables and x^* denote specification variables (e.g., ghost variables) is that for all states st satisfying $P(v^*, x^*)$, given a reduction $e, st \rightsquigarrow^* v, st'$, if

$e, st \not\sim^* \text{fault}$, then st' satisfies $\exists y^*. Q(v^*, x^*, y^*, res)$. The staged equivalent is $\{ \Phi \} e \{ \Phi; \exists x^* \cdot \mathbf{req} P(v^*, x^*); \exists y^* \cdot \mathbf{ens}[_] Q(v^*, x^*, y^*, res) \}$. Apart from mentioning the *history* Φ , which remains unchanged, its meaning is identical. Consider, then, $\{ \Phi \} e \{ \Phi; \mathbf{req} P_1; \mathbf{ens}[_] Q_1; \mathbf{req} P_2; \mathbf{ens}[_] Q_2 \}$ – an intuitive extension of the semantics of triples is that given $e, st \rightsquigarrow^* e_1, st_1$, where st_1 satisfies Q_1 , the extended judgment holds if st_1 *further* satisfies P_2 , and reduction from there, $e_1, st_1 \rightsquigarrow^* e_2, st_2$, results in a state st_2 that satisfies Q_2 .

While heap formulae are satisfied by program states, staged formulae (like triples), are satisfied by traces which begin and end at particular states. Uninterpreted function stages further allow stages to describe the *intermediate states* of programs in specifications – a useful ability in the presence of unknown higher-order imperative functions, as we illustrate in Section 2 and Appendix C [13]. To formalize all this, we give a semantics for staged formulae next.

Formal Semantics We first recall the standard semantics for separation logic formulae in Fig. 4, which provides a useful starting point.

$$\begin{aligned}
S, h \models \sigma \wedge \pi & \quad \text{iff} \quad \llbracket \pi \rrbracket_S \text{ and } S, h \models \sigma \\
S, h \models \text{emp} & \quad \text{iff} \quad \text{dom}(h) = \{\} \\
S, h \models x \mapsto y & \quad \text{iff} \quad \text{dom}(h) = \{S(x)\} \text{ and } h(S(x)) = \llbracket y \rrbracket_S \\
S, h \models \sigma_1 * \sigma_2 & \quad \text{iff} \quad \exists h_1 h_2. h_1 \circ h_2 = h \text{ such that } S, h_1 \models \sigma_1 \text{ and } S, h_2 \models \sigma_2
\end{aligned}$$

Fig. 4: Semantics of Separation Logic Formulae

Let var be the set of program variables, val the set of primitive values, and $loc \subset val$ the set of heap locations; ℓ is a metavariable ranging over locations. The models are program states, comprising a *store* of variables S , a partial mapping from a finite set of variables to values $var \rightarrow val$, and the heap h , a partial mapping from locations to values $loc \rightarrow val$. $\llbracket \pi \rrbracket_S$ denotes the valuation of pure formula π under store S . $\text{dom}(h)$ denotes the domain of heap h . $h_1 \circ h_2 = h$ denotes disjoint union of heaps; if $\text{dom}(h_1) \cap \text{dom}(h_2) = \{\}$, $h_1 \cup h_2 = h$. We write $h_1 \subseteq h_2$ to denote that h_1 is a subheap of h_2 , i.e., $\exists h_3 \cdot h_1 \circ h_3 = h_2$. $s[x:=v]$ and $s[x;\neq]$ stand for store/heap updates and removal of keys.

We define the semantics of *HSSL* formulae in Fig. 5. Let $S, h \rightsquigarrow S_1, h_1, R \models \Phi$ denote the *models* relation, i.e., starting from the program state with store S and heap h , the formula Φ transforms the state into S_1, h_1 , with an intermediate result R . R is either $Norm(r)$ for partial correctness, Err for precondition failure, or \top for *possible* precondition failure in one of its execution paths.

When Φ is of the form $\mathbf{req} \sigma \wedge \pi$, the heap h is split into a heaplet h_1 satisfying $\sigma \wedge \pi$, which is consumed, and a frame h_2 , which is left as the new heap. *Read-only heap assertions* $(\sigma \wedge \pi) @ R$ under \mathbf{req} check but do not change the heap.

When Φ is of the form $\mathbf{ens}[_] \sigma \wedge \pi$, σ describes locations which are to be added to the current heap. The semantics allows some concrete heaplet h_1 that satisfies $\sigma \wedge \pi$ (containing new or updated locations) be (re-)added to heap h .

When Φ is a function stage $f(x^*, r)$, its semantics depends on the specification of f . A staged existential causes the store to be extended with a binding from x

$S, h \rightsquigarrow S, h_1, Norm(_) \models \mathbf{req} \sigma \wedge \pi$	<i>iff</i> $h_1 \subseteq h$ and $S, h_1 \models \sigma \wedge \pi$
$S, h \rightsquigarrow S, h, Err \models \mathbf{req} \sigma \wedge \pi$	<i>iff</i> $\forall h_1 \cdot h_1 \subseteq h \Rightarrow S, h_1 \not\models \sigma \wedge \pi$
$S, h \rightsquigarrow S, h, R \models \mathbf{req} (\sigma \wedge \pi) @ R$	<i>iff</i> $S, h \rightsquigarrow S, h_1, R \models \mathbf{req} (\sigma \wedge \pi)$
$S, h \rightsquigarrow S, h \circ h_1, Norm(r) \models \mathbf{ens}[r] \sigma \wedge \pi$	<i>iff</i> $S, h_1 \models \sigma \wedge \pi$ and $dom(h_1) \cap dom(h) = \{\}$
$S, h \rightsquigarrow S_1, h_1, R \models f(x^*, r)$	<i>iff</i> $S(f) = fun(y^*) \Phi[r'] \rightarrow e,$ $S, h \rightsquigarrow S_1, h_1, R \models [r' := r][y^* := x^*] \Phi$
$S, h \rightsquigarrow S_1, h_1, R \models \exists x \cdot \Phi$	<i>iff</i> $\exists v \cdot S[x := v], h \rightsquigarrow S_1, h_1, R \models \Phi$
$S, h \rightsquigarrow S_2, h_2, R \models \Phi_1; \Phi_2$	<i>iff</i> $S, h \rightsquigarrow S_1, h_1, Norm(r) \models \Phi_1,$ $S_1, h_1 \rightsquigarrow S_2, h_2, R \models \Phi_2$
$S, h \rightsquigarrow S_1, h_1, \top \models \Phi_1; \Phi_2$	<i>iff</i> $S, h \rightsquigarrow S_1, h_1, \top \models \Phi_1$
$S, h \rightsquigarrow S_3, h_3, Norm(r_3) \models \Phi_1 \vee \Phi_2$	<i>iff</i> $\exists h_1, h_2, r_1, r_2 \cdot S, h \rightsquigarrow S_1, h_1, Norm(r_1) \models \Phi_1$ and $S, h \rightsquigarrow S_2, h_2, Norm(r_2) \models \Phi_2,$ and $(S_3, h_3, r_3) \in \{(S_1, h_1, r_1), (S_2, h_2, r_2)\}$
$S, h \rightsquigarrow S_1, h_1, \top \models \Phi_1 \vee \Phi_2$	<i>iff</i> $S, h \rightsquigarrow S_1, h_1, \top \models \Phi_1$ or $S, h \rightsquigarrow S_1, h_1, \top \models \Phi_2$

Fig. 5: Semantics of Staged Formulae

to an existential value v . Sequential composition $\Phi_1; \Phi_2$ results in a failure \top if Φ_1 does, while disjunction $\Phi_1 \vee \Phi_2$ requires both branches not to fail.

3.2 Compaction

Staged formulae subsume separation logic triples, but triples suffice for many verification tasks, particularly those without calls to unknown functions, and we would like to recover their succinctness in cases where intermediate states are not required. This motivates a *compaction* or *normalization* procedure for staged formulae, written $\Phi \Longrightarrow \Phi$ (Fig. 6). Compaction is also useful for aligning staged formulae, allowing entailment proofs to be carried out stage by stage; we elaborate on this use in Section 5.

$\mathbf{ens}[-] false; \Phi \Longrightarrow \mathbf{ens}[-] false$	$\mathbf{req} D_1; \mathbf{req} D_2 \Longrightarrow \mathbf{req} (D_1 * D_2)$
$emp; \Phi \Longrightarrow \Phi$	$\mathbf{ens}[-] D_1; \mathbf{ens}[r] D_2 \Longrightarrow \mathbf{ens}[r] (D_1 * D_2)$
$\Phi; emp \Longrightarrow \Phi$	$\frac{D_A * D_1 \vdash D_2 * D_F}{\mathbf{ens}[r] D_1; \mathbf{req} D_2 \Longrightarrow \mathbf{req} D_A; \mathbf{ens}[r] D_F}$

Fig. 6: Select compaction rules

The three rules on the left simplify flows. A false postcondition ($\mathbf{ens} \sigma \wedge false$) models an unreachable or nonterminating program state, so the rest of a flow may be safely ignored. emp in the next two rules is *either* ($\mathbf{req} emp \wedge true$) or ($\mathbf{ens} emp \wedge true$); either may serve as an identity for flows. The first two rules on the right merge consecutive pre- and postconditions. Intuitively, they are sound because symbolic heaps separated by sequential composition must be disjoint to be meaningful – this follows from the use of disjoint union in Fig. 5. The last rule

allows a precondition **req** D_2 to be transposed with a preceding postcondition **ens** D_1 . This is done using biabduction [4], which computes a pair of antiframe D_A and frame D_F such that the antiframe is the new precondition required, and frame is what remains after proving the known precondition. The given rule assumes that D_1 and D_2 are disjoint². A read-only $@R$ heap assertion under **req** would be handled by matching but not removing from D_F (see [7]).

Thus staged formulae can always be compacted into the following form, consisting of a disjunction of *flows* θ (a disjunction-free staged formula)³, each consisting of a prefix of function stages (preceded by a description of the intermediate state at that point), followed by a final pre- and postcondition, capturing any behavior remaining after calling unknown functions.

$$\begin{aligned} \Phi &::= \theta \mid \Phi \vee \Phi \\ \theta &::= (\exists x^* \cdot \mathbf{req} D; \exists x^* \cdot \mathbf{ens}[-] D; f(v^*, r) ;)^* \exists x^* \cdot \mathbf{req} D; \exists x^* \cdot \mathbf{ens}[-] D \end{aligned}$$

An example of compaction is given below (Fig. 7, left). We start at the first two stages of the flow and solve a biabduction problem (shown on the right, with solution immediately below) to infer a precondition for the whole flow, or, more operationally, to “push” the **req** to the left. We will later be able to rely on the new precondition to know that $a = 1$ when proving properties of the rest of the flow. Finally, we may combine the two **ens** stages because sequential composition guarantees disjointness. Normalization is *sound* in the sense that it transforms staged formulae without changing their meaning.

$$\begin{array}{ll} \mathbf{ens} \ x \mapsto 1 * y \mapsto 2; \mathbf{req} \ x \mapsto a; \mathbf{ens} \ x \mapsto a+1 & \\ \iff \mathbf{req} \ a=1; \mathbf{ens} \ y \mapsto 2; \mathbf{ens} \ x \mapsto a+1 & D_A * x \mapsto 1 * y \mapsto 2 \vdash x \mapsto a * D_F \\ \iff \mathbf{req} \ a=1; \mathbf{ens} \ y \mapsto 2 * x \mapsto a+1 & D_A=(a=1), D_F=(y \mapsto 2) \end{array}$$

Fig. 7: An example of compaction

Theorem 1 (Soundness of Normalization). *Given $\Phi_1 \implies \Phi_2$, if $S, H \rightsquigarrow S_1, H_1, R_1 \models \Phi_1$, then $S, H \rightsquigarrow S_1, H_1, R_1 \models \Phi_2$.*

Proof. By case analysis on the derivation of $\Phi_1 \implies \Phi_2$. See Appendix I.2 [13].

4 Forward Rules for Staged Logics

To verify that a program satisfies a given specification Φ_s , we utilize a set of rules (presented in Fig. 8) to compute an abstraction or summary of the program Φ_p , then discharge the proof obligation $\Phi_p \sqsubseteq \Phi_s$ (covered in Section 5), in a manner similar to strongest postcondition calculations.

We make use of the following notations. $_$ denotes an anonymous existentially quantified variable. $[x:=v]\Phi$ denotes the substitution of x with v in Φ , giving

² More exhaustive aliasing scenarios are considered in Appendix D [13].

³ Using further normalization rules such as $(\Phi_1 \vee \Phi_2); \Phi_3 \implies (\Phi_1; \Phi_3) \vee (\Phi_2; \Phi_3)$

$$\begin{array}{c}
 \frac{\Phi_1 \sqsubseteq \Phi_3 \quad \{ \Phi_3 \} e \{ \Phi_4 \} \quad \Phi_4 \sqsubseteq \Phi_2}{\{ \Phi_1 \} e \{ \Phi_2 \}} \text{Conseq} \quad \frac{\{ \Phi_1 \} e \{ \Phi_2 \}}{\{ \Phi; \Phi_1 \} e \{ \Phi; \Phi_2 \}} \text{Frame} \\
 \\
 \frac{}{\{ \Phi \} x \{ \Phi; \mathbf{ens}[x] \mathit{emp} \}} \text{Var} \quad \frac{\text{fresh } r \quad v ::= c \mid \mathit{nil} \mid x_1 :: x_2}{\{ \Phi \} v \{ \Phi; \exists r. \mathbf{ens}[r] r = v \}} \text{Val} \\
 \\
 \frac{\text{fresh } r}{\{ \Phi \} \mathit{ref } x \{ \Phi; \exists r. \mathbf{ens}[r] r \mapsto x \}} \text{Ref} \\
 \\
 \frac{\text{fresh } a, \mathit{res}}{\{ \Phi \} !x \{ \Phi; \exists a, \mathit{res}. \mathbf{req } x \mapsto a; \mathbf{ens}[\mathit{res}] x \mapsto a \wedge \mathit{res} = a \}} \text{Deref} \\
 \\
 \frac{}{\{ \Phi \} x_1 := x_2 \{ \Phi; \mathbf{req } x_1 \mapsto -; \mathbf{ens}[_] x_1 \mapsto x_2 \}} \text{Assign} \\
 \\
 \frac{\{ \Phi; \mathbf{ens}[_] x \} e_1 \{ \Phi_1 \} \quad \{ \Phi; \mathbf{ens}[_] \neg x \} e_2 \{ \Phi_2 \}}{\{ \Phi \} \mathit{if } x \mathit{then } e_1 \mathit{else } e_2 \{ \Phi_1 \vee \Phi_2 \}} \text{If} \\
 \\
 \frac{\text{fresh } x \quad \{ \Phi \} e_1 \{ \exists r. \Phi_1[r] \} \quad \{ [r := x] \Phi_1 \} e_2 \{ \Phi_2 \}}{\{ \Phi \} \mathit{let } x = e_1 \mathit{in } e_2 \{ \exists x. \Phi_2 \}} \text{Let} \\
 \\
 \frac{\text{fresh } \mathit{res} \quad \{ \mathbf{ens}[_] \mathit{Pure}(\Phi) \} e \{ \exists r'. \Phi_p[r'] \} \quad ([r' := r] \Phi_p) \sqsubseteq \Phi_s}{\{ \Phi \} \mathit{fun } (x^*) \exists r. \Phi_s[r] \rightarrow e \{ \Phi; \exists \mathit{res}. \mathbf{ens}[\mathit{res}] \mathit{res} = \lambda (x^*, r) \rightarrow \Phi_s \}} \text{Lambda} \\
 \\
 \frac{\text{fresh } r}{\{ \Phi \} f(x^*) \{ \Phi; \exists r. f(x^*, r) \}} \text{Call} \quad \frac{}{\{ \Phi \} \mathit{assert } D \{ \Phi; \mathbf{req } D @ R \}} \text{Assert}
 \end{array}$$

Fig. 8: Forward Reasoning Hoare Rules with Staged Logics

priority to recently bound variables. We lift sequencing from flows to disjunctive staged formulae in the natural way: $\Phi_1; \Phi_2 \triangleq \bigvee \{ \theta_1; \theta_2 \mid \theta_1 \in \Phi_2, \theta_2 \in \Phi_2 \}$.

The first two rules in Fig. 8 are structural. The **Conseq** rule uses *specification subsumption* (detailed in Section 5) in place of implication – a form of behavioral subtyping. The **Frame** rule has both a *temporal* interpretation, which is that the reasoning rules are compositional with respect to the *history* of the current flow, and a *spatial* interpretation, consistent with the usual one from separation logic, if one uses the normalization rules (Section 3.2) to move untouched p from the final states of Φ_1 and Φ_2 into the frame Φ .

The **Var** and **Val** rules illustrate how the results of pure expressions are tracked via named **ens** results. The **Ref** rule results in a new, existentially-quantified location being added to the current state. The **Deref** and **Assign** rules are similar, both requiring proof that a named location exists with a value, then respectively either returning the value of the location and leaving it unchanged, or changing the location and returning the unit value. **Assert** checks the current heap state without modifying it using the $@R$ read-only annotation. **If** introduces disjunction. **Let** sequences expressions, renaming the intermediate result of e_1 accordingly; the scope of x in e_2 is represented by the scope of the introduced existential in the conclusion of the rule.

The **Lambda** rule handles function definition annotated with a given specification Φ_s . The body of the lambda is summarized into Φ_p starting from pure

information $Pure(\Phi)$ from its program context. Its behavior must be subsumed by the given specification. The result is then the lambda expression itself.

The **Call** rule is completely trivial, yet perhaps the most illuminating as to the design of *HSSL*. A standard modular verifier would utilize this rule to look up the specification associated with f , prove its precondition, then assume its postcondition. In our setting, however, there is the possibility that f is higher-order, unknown, and/or unspecified. Moreover, there is no need to prove the precondition of f immediately, due to the use of flows for describing program behaviors. Both of these point to the simple use of a function stage, which stands for a *possibly-unknown* function call. Utilizing the specification of f , if it is provided, is deferred to the unfolding done in the entailment procedure.

We prove soundness of these rules, which is to say that derived specifications faithfully overapproximate the programs they are derived from. In the following theorem, $e, h, S \rightsquigarrow h_1, S_1$ is a standard big-step reduction relation whose definition we leave to Appendix I.1 [13]. Termination is also considered in Appendix I.5 [13]. However, completeness is yet to be established.

Theorem 1 (Soundness of Forward Rules) *Given $\{ emp \} e \{ \Phi \}$, then $\forall S, h, S_2, h_1 \cdot (S, h \rightsquigarrow S_2, h_1, Norm(r) \models \Phi) \Rightarrow \exists S_1 \cdot e, h, S \rightsquigarrow Norm(v), h_1, S_1$ and $S_1 \subseteq S_2$ and $S_1(r) = v$.*

Proof. By induction on the derivation of $e, h, S_1 \rightsquigarrow R_1, h_1, S_1$. See Appendix I.3 [13].

5 Staged Entailment Checking and its Soundness

In this section, we outline how entailments of the form $F \vdash \Phi_p \sqsubseteq \Phi_s$ may be automatically checked. F denotes heap and pure frames that are propagated by our staged logics entailment rules. Our entailment is always conducted over the compacted form where non-recursive staged predicate definitions are unfolded, while unknown predicates are matched exactly. Lemmas are also used to try re-summarize each instantiation of recursive staged predicates to simpler forms, where feasible. As staged entailment ensures that all execution traces that satisfy Φ_p must also satisfy Φ_s , we rely on theory of *behavioral subtyping* [18] to relate them. Specifically, we check that *contravariance holds* for pre-condition entailment, while *covariance holds* for post-condition entailment, as follows:

$$\frac{\text{fresh } y^* \quad F_0 * D_2 \vdash (\exists x^*. D_1) * F \quad F \vdash \theta_a \sqsubseteq \theta_c}{F_0 \vdash (\exists x^* \cdot \mathbf{req} D_1; \theta_a) \sqsubseteq (\exists y^* \cdot \mathbf{req} D_2; \theta_c)} \mathbf{EntReq}$$

$$\frac{\text{fresh } x^* \quad F_0 * D_1 \vdash (\exists y^*. D_2) * F \quad F \vdash \theta_a \sqsubseteq \theta_c}{F_0 \vdash (\exists x^* \cdot \mathbf{ens}[r] D_1; \theta_a) \sqsubseteq (\exists y^* \cdot \mathbf{ens}[r] D_2; \theta_c)} \mathbf{EntEns}$$

More details of staged entailment rules are given in Appendix G [13]. Note that we use another entailment over separation logic $D_1 \vdash D_2 * F_r$ that can propagate residual frame, F_r . Lastly, we outline the soundness of staged entailment against the semantics of staged formulae, ensuring that all derivations are valid.

Theorem 2 (Soundness of Entailment). *Given $\Phi_1 \sqsubseteq \Phi_2$ and $S, h \rightsquigarrow \text{Norm}(r_1), S_1, h_1 \models \Phi_1$, then there exists h_2 such that $S, h \rightsquigarrow \text{Norm}(r_1), S_2, h_2 \models \Phi_2$ where $h_2 \subseteq h_1$. (Here, $h_1 \subseteq h_2$ denotes that $\exists h_3. h_1 \circ h_3 = h_2$.)*

Proof. By induction on the derivation of $\Phi_1 \sqsubseteq \Phi_2$. See Appendix I.4 [13].

6 Implementation and Initial Results

We prototyped our verification methodology in a tool named Heifer. Our tool takes input programs written in a subset of OCaml annotated with user-provided specifications. It analyzes input programs to produce normalized staged formulae (Section 3.2, Section 4), which it then translates to first-order verification conditions (Section 5) suitable for an off-the-shelf SMT solver. Here, our prototype targets SMT encodings via Why3 [11]. As an optimization, it uses Z3 [8] directly for queries which do not require Why3’s added features.

Table 1: A Comparison with Cameleer and Prusti. (Programs that are natively inexpressible are marked with “ \times ”. Programs that cannot be reproduced from Prusti’s artifact [1] are marked with “-” denoting incomparable. We use T to denote the total verification time (in seconds) and T_P to record the time spent on the external provers.)

Benchmark	Heifer				Cameleer [21]			Prusti [30]		
	LoC	LoS	T	T_P	LoC	LoS	T	LoC	LoS	T
map	13	11	0.66	0.58	10	45	1.25	-	-	-
map_closure	18	7	1.06	0.77		\times		-	-	-
fold	23	12	1.06	0.87	21	48	8.08	-	-	-
fold_closure	23	12	1.25	0.89		\times		-	-	-
iter	11	4	0.40	0.32		\times		-	-	-
compose	3	1	0.11	0.09	2	6	0.05	-	-	-
compose_closure	23	4	0.44	0.32		\times			\times	
closure [26]	27	5	0.37	0.27		\times		13	11	6.75
closure_list	7	1	0.15	0.09		\times		-	-	-
applyN	6	1	0.19	0.17	12	13	0.37	-	-	-
blameassgn [12]	14	6	0.31	0.28		\times		13	9	6.24
counter	16	4	0.24	0.18		\times		11	7	6.37
lambda	13	5	0.25	0.22		\times		-	-	-
	197	73			45	112		37	27	

We have verified a suite of programs (Table 1) involving higher-order functions and closures. As the focus of our work is to explore a new program logic and subsumption-based verification methodology (rather than to verify existing programs), the benchmarks are small in size, and are meant to illustrate the style of specification and give a flavor of the potential for automation.

Table 1 provides an overview of the benchmark suite. The first two sub-columns show the size of each program (LoC) and the number of lines of user-provided specifications (LoS) required. The next two give the total wall-clock time taken (in seconds) to verify all functions in each program against the provided specifications, and the amount of time spent in external provers.

The next column shows the same programs verified using Cameleer [21, 24], a state-of-the-art deductive verifier. Cameleer serves as a good baseline for several reasons: it is representative of the dominant paradigm of pre/post specifications and, like Heifer, targets (a subset of) OCaml. It supports higher-order functions in both programs and specifications [25]. The most significant differences between Cameleer and Heifer are that Cameleer does not support *effectful* higher-order functions and is intended to be used via the Why3 IDE in a semi-interactive way (allowing tactic-like *proof transformations*, used in the above programs).

The last column shows results for Prusti [30]. Despite Rust’s ownership type system, we compare it against Prusti because of its state-of-the-art support for mutable closures, highlighting differences below. While we were able to reproduce the claims made in Prusti’s OOPSLA 2021 artifact [1], we were not able to verify many of our own benchmark programs due to two technical reasons, namely lacking support for Rust’s `impl Trait` (to return closures) and ML-like cons lists (which caused timeouts and crashes). Support for closures is also not yet in mainline Prusti [2]. Nevertheless, we verified the programs we could use for the artifact, the results of which are shown in Table 1. All experiments were performed on macOS using a 2.3 GHz Quad-Core Intel Core i7 CPU with 16 GB of RAM. Why3 1.7.0 was used, with SMT solvers Z3 4.12.2, CVC4 1.8, and Alt-Ergo 2.5.2. The Prusti artifact, a Docker image, was run using Moby 25.0.1.

User annotations required. Significantly less specification than code is required in Heifer, with an average LoS/LoC ratio of 0.37. This is helped by two things: the use of function stages in specifications, and the use of biabduction-based normalization, which allows the specifications of functions to be mostly automated, requiring only properties and auxiliary lemmas to be provided. In contrast, Cameleer’s ratio is 2.49, due to the need to adequately summarize the behaviors of the function arguments and accompany these summaries with invariants and auxiliary lemmas. Two examples illustrating this are detailed in Appendix F [13]. Prusti’s ratio is 0.73, but a caveat is that in the programs for it, only closure reasoning was used, without lemmas or summarization.

Expressiveness. Heifer is able to express many programs that Cameleer cannot, particularly closure-manipulating ones. This accounts for the **X** rows in Table 1. While some of these can be verified with Prusti, unlike stages, Prusti’s call descriptions do not capture ordering [10, 1]; an explicit limitation as shown by the **X** rows in Prusti’s column. Prusti is able to use history invariants and the ownership of the Rust type system, but this difference is more than mitigated in Heifer with the adoption of an expressive staged logic with spatial heap state; more appropriate for the weaker (but more general) type system of OCaml.

7 Related Work

The use of sequential composition in specifications goes back to classic theories of program refinement, such as Morgan’s refinement calculus [19] and Hoare and He’s Unifying Theories [15], as well session types [9] and logics [6]. It has also been

used to structure verification conditions and give users control over the order in which they are given to provers [14], allowing more reliable proof automation. We extend both lines of work, developing the use of sequential composition as a precise specification mechanism for higher-order imperative functions, and using it to guide entailment proofs of staged formulae.

Higher-order imperative functions were classically specified in program logics using *evaluation formulae* [16] and *reference-reachability predicates* [32]. The advent of separation logic has allowed for simpler specifications using *invariants* and *nested triples* (Section 1). These techniques are common in higher-order separation logics, such as HTT [20], CFML [5], Iris [17] and Steel/Pulse [28], which are encoded in proof assistants (e.g. Coq, F \star [27]) which do not natively support closures or heap reasoning. While the resulting object logics are highly expressive, they are much more complex (owing to highly nontrivial encodings) and consequently less automated than systems that discharge obligations via SMT. We push the boundaries in this area by proposing stages as a new, precise specification mechanism which is compatible with automated verification.

The guarantees of an expressive type system can significantly simplify how higher-order state is specified and managed. Prusti [30] exploits this with *call descriptions* (an alternative to function stages, as pure assertions saying that a call has taken place with a given pre/post) and *history invariants*, which rely on the ownership of mutable locations that closures have in Rust. Creusot [10] uses a *prophetic mutable value semantics* to achieve a similar goal with pre/post specifications of closures. Our solution is not dependent on an ownership type system, applying more generally to languages with unrestricted mutation.

Defunctionalization [22] is another promising means of reasoning about higher-order effectful programs [25], pioneered by the Why3-based Cameleer [21]. This approach currently does not support closures.

Our approach to automated verification is currently based on strict evaluation. It would be interesting to see how staged specifications can be extended to support verification of lazy programs, as had been explored in [31] and [29].

8 Conclusion

We have explored how best to *modularly specify and verify higher-order imperative programs*. Our contributions are manifold: we propose a new staged specification logic, rules for deriving staged formulae from programs and normalizing them using biabduction, and an entailment proof system. This forms the basis of a new verification methodology, which we validate with our prototype Heifer.

To the best of the authors' knowledge, this work is the first to introduce a *fundamental* staged specification mechanism for verifying higher-order imperative programs *without* any presumptions; being *more concise* (without the need for specifying abstract properties) and *more precise* (without imposing preconditions on function-typed parameters) when compared to existing solutions.

Acknowledgments This research is supported by the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2023) (Project Title: Automated Verification for Imperative Higher-Order Programs).

Bibliography

- [1] Modular specification and verification of closures in Rust (artefact), 2021.
- [2] Documentation of closures. <https://github.com/viperproject/prusti-dev/issues/1431>, 2024.
- [3] Fernando Alves, Delano Oliveira, Fernanda Madeiral, and Fernando Castor. On the bug-proneness of structures inspired by functional programming in JavaScript projects. *CoRR*, abs/2206.08849, 2022.
- [4] Cristiano Calcagno, Dino Distefano, Peter W. O’Hearn, and Hongseok Yang. Compositional shape analysis by means of bi-abduction. In Zhong Shao and Benjamin C. Pierce, editors, *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, pages 289–300. ACM, 2009.
- [5] Arthur Charguéraud. Characteristic formulae for the verification of imperative programs. In Manuel M. T. Chakravarty, Zhenjiang Hu, and Olivier Danvy, editors, *Proceeding of the 16th ACM SIGPLAN international conference on Functional Programming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*, pages 418–430. ACM, 2011.
- [6] Andreea Costea, Wei-Ngan Chin, Shengchao Qin, and Florin Craciun. Automated modular verification for relaxed communication protocols. In Sukyoung Ryu, editor, *Programming Languages and Systems - 16th Asian Symposium, APLAS 2018, Wellington, New Zealand, December 2-6, 2018, Proceedings*, volume 11275 of *Lecture Notes in Computer Science*, pages 284–305. Springer, 2018.
- [7] Cristina David and Wei-Ngan Chin. Immutable specifications for more concise and precise verification. In Cristina Videira Lopes and Kathleen Fisher, editors, *Proceedings of the 26th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2011, part of SPLASH 2011, Portland, OR, USA, October 22 - 27, 2011*, pages 359–374. ACM, 2011.
- [8] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. Z3: an efficient SMT solver. In C. R. Ramakrishnan and Jakob Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.
- [9] Pierre-Malo Deniérou, Nobuko Yoshida, Andi Bejleri, and Raymond Hu. Parameterised multiparty session types. *Log. Methods Comput. Sci.*, 8(4), 2012.
- [10] Xavier Denis and Jacques-Henri Jourdan. Specifying and verifying higher-order Rust iterators. In Sriram Sankaranarayanan and Natasha Sharygina, editors, *Tools and Algorithms for the Construction and Analysis of Systems*

- *29th International Conference, TACAS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Paris, France, April 22-27, 2023, Proceedings, Part II*, volume 13994 of *Lecture Notes in Computer Science*, pages 93–110. Springer, 2023.
- [11] Jean-Christophe Filliâtre and Andrei Paskevich. Why3 - where programs meet provers. In *European Symposium on Programming*, 2013.
- [12] Robert Bruce Findler and Matthias Felleisen. Contracts for Higher-order Functions. In Mitchell Wand and Simon L. Peyton Jones, editors, *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP '02), Pittsburgh, Pennsylvania, USA, October 4-6, 2002*, pages 48–59. ACM, 2002.
- [13] Darius Foo, Yahui Song, and Wei-Ngan Chin. Staged specification logic for verifying higher-order imperative programs (technical report). https://github.com/hipsleek/Heifer/blob/StagedSL/docs/FM2024_TR.pdf, 2024.
- [14] Cristian Gherghina, Cristina David, Shengchao Qin, and Wei-Ngan Chin. Structured specifications for better verification of heap-manipulating programs. In Michael J. Butler and Wolfram Schulte, editors, *FM 2011: Formal Methods - 17th International Symposium on Formal Methods, Limerick, Ireland, June 20-24, 2011. Proceedings*, volume 6664 of *Lecture Notes in Computer Science*, pages 386–401. Springer, 2011.
- [15] Jifeng He and Charles Antony Richard Hoare. Unifying theories of programming. In *RelMiCS*, 1998.
- [16] Kohei Honda, Nobuko Yoshida, and Martin Berger. An observationally complete program logic for imperative higher-order frame rules. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings*, pages 270–279. IEEE Computer Society, 2005.
- [17] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.*, 28:e20, 2018.
- [18] Gary T. Leavens and David A. Naumann. Behavioral subtyping, specification inheritance, and modular reasoning. *ACM Trans. Program. Lang. Syst.*, 37(4):13:1–13:88, 2015.
- [19] Carroll Morgan. The refinement calculus. In *NATO ASI PDC*, 1994.
- [20] Aleksandar Nanevski, J. Gregory Morrisett, and Lars Birkedal. Hoare type theory, polymorphism and separation. *J. Funct. Program.*, 18(5-6):865–911, 2008.
- [21] Mário Pereira and António Ravara. Cameleer: A deductive verification tool for OCaml. In Alexandra Silva and K. Rustan M. Leino, editors, *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part II*, volume 12760 of *Lecture Notes in Computer Science*, pages 677–689. Springer, 2021.
- [22] John C. Reynolds. Definitional interpreters for higher-order programming languages. In John J. Donovan and Rosemary Shields, editors, *Proceedings*

- of the ACM annual conference, ACM 1972, 1972, Volume 2, pages 717–740. ACM, 1972.
- [23] Jan Smans, Bart Jacobs, and Frank Piessens. Implicit dynamic frames. *ACM Trans. Program. Lang. Syst.*, 34(1):2:1–2:58, 2012.
 - [24] Tiago Soares and Mário Pereira. A framework for the automated verification of algebraic effects and handlers (extended version). *ArXiv*, abs/2302.01265, 2023.
 - [25] Tiago Lopes Soares. A deductive verification framework for higher order programs. *CoRR*, abs/2011.14044, 2020.
 - [26] Kasper Svendsen. *Modular specification and verification for higher-order languages with state*. IT-Universitetet i København, 2013.
 - [27] Nikhil Swamy, Catalin Hritcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean Karim Zinzindohoue, and Santiago Zanella Béguelin. Dependent types and multi-monadic effects in F. In Rastislav Bodík and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 256–270. ACM, 2016.
 - [28] Nikhil Swamy, Aseem Rastogi, Aymeric Fromherz, Denis Merigoux, Danel Ahman, and Guido Martínez. Steelcore: an extensible concurrent separation logic for effectful dependently typed programs. *Proc. ACM Program. Lang.*, 4(ICFP):121:1–121:30, 2020.
 - [29] Niki Vazou, Eric L. Seidel, and Ranjit Jhala. LiquidHaskell: experience with refinement types in the real world. In Wouter Swierstra, editor, *Proceedings of the 2014 ACM SIGPLAN symposium on Haskell, Gothenburg, Sweden, September 4-5, 2014*, pages 39–51. ACM, 2014.
 - [30] Fabian Wolff, Aurel Bilý, Christoph Matheja, Peter Müller, and Alexander J. Summers. Modular specification and verification of closures in Rust. *Proc. ACM Program. Lang.*, 5(OOPSLA):1–29, 2021.
 - [31] Dana N. Xu, Simon L. Peyton Jones, and Koen Claessen. Static contract checking for Haskell. In Zhong Shao and Benjamin C. Pierce, editors, *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, pages 41–52. ACM, 2009.
 - [32] Nobuko Yoshida, Kohei Honda, and Martin Berger. Logical reasoning for higher-order functions with local state. In Helmut Seidl, editor, *Foundations of Software Science and Computational Structures, 10th International Conference, FOSSACS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007, Braga, Portugal, March 24-April 1, 2007, Proceedings*, volume 4423 of *Lecture Notes in Computer Science*, pages 361–377. Springer, 2007.
 - [33] Fiorella Zampetti, Francois Belias, Cyrine Zid, and Giuliano Antonioland Massimiliano Di Penta. An empirical study on the fault-inducing effect of functional constructs in Python. *2022 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 47–58, 2022.