

Exact Separation Logic

Towards Bridging the Gap Between Verification and Bug-Finding

Petar Maksimović

Imperial College London, UK

Runtime Verification Inc., Urbana, IL, USA

Caroline Cronjäger

Ruhr-Universität Bochum, Germany

Andreas Löw

Imperial College London, UK

Julian Sutherland

Nethermind, London, UK

Philippa Gardner

Imperial College London, UK

Abstract

Over-approximating (OX) program logics, such as separation logic (SL), are used for *verifying* properties of heap-manipulating programs: all terminating behaviour is characterised, but established results and errors need not be reachable. OX function specifications are thus incompatible with true bug-finding supported by symbolic execution tools such as Pulse and Pulse-X. In contrast, under-approximating (UX) program logics, such as incorrectness separation logic, are used to *find* true results and bugs: established results and errors are reachable, but there is no mechanism for understanding if all terminating behaviour has been characterised.

We introduce exact separation logic (ESL), which provides fully-verified function specifications compatible with both OX verification and UX true bug-finding: all terminating behaviour is characterised and all established results and errors are reachable. We prove soundness for ESL with mutually recursive functions, demonstrating, for the first time, function compositionality for a UX logic. We show that UX program logics require subtle definitions of internal and external function specifications compared with the familiar definitions of OX logics. We investigate the expressivity of ESL and, for the first time, explore the role of abstraction in UX reasoning by verifying abstract ESL specifications of various data-structure algorithms. In doing so, we highlight the difference between *abstraction* (hiding information) and *over-approximation* (losing information). Our findings demonstrate that abstraction cannot be used as freely in UX logics as in OX logics, but also that it should be feasible to use ESL to provide tractable function specifications for self-contained, critical code, which would then be used for both verification and true bug-finding.

2012 ACM Subject Classification Theory of computation → Logic and verification; Theory of computation → Program reasoning; Theory of computation → Separation logic; Theory of computation → Hoare logic; Theory of computation → Abstraction

Keywords and phrases Separation logic, program correctness, program incorrectness, abstraction

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2023.19

Related Version *Extended Version*: <https://arxiv.org/abs/2208.07200>

Funding Maksimović, Löw and Gardner were partially supported by the EPSRC Fellowship “VetSpec: Verified Trustworthy Software Specification” (EP/R034567/1). Cronjäger was partially supported by the Erasmus Plus Student Mobility for Traineeships scheme.

Acknowledgements We would like to thank Sacha-Élie Ayoun and Daniele Nantes Sobrinho for the many discussions that have improved the quality of the paper. We would also like to thank the anonymous reviewers for their comments.



© Petar Maksimović, Caroline Cronjäger, Andreas Löw, Julian Sutherland, and Philippa Gardner; licensed under Creative Commons License CC-BY 4.0

37th European Conference on Object-Oriented Programming (ECOOP 2023).

Editors: Karim Ali and Guido Salvaneschi; Article No. 19; pp. 19:1–19:27

Leibniz International Proceedings in Informatics



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

1 Introduction

Over-approximating (OX) program logics were introduced to reason about program correctness, starting with Hoare logic [18] and evolving to separation logic (SL) [27, 30]. SL is used for *verification* and features function specifications of the form $\{P\} f(\vec{x}) \{Q\}$, the meaning of which is that all terminating executions of the function f that start from a state in the pre-condition P end in a state covered by the post-condition Q . SL has the standard rule of *forward consequence*, which allows one to *lose information* (for example, if we had a post-condition with $x = 42$, we could soundly weaken this precise information to the less precise $x > 0$ or even to the non-informative `true`). In essence, the philosophy underlying the OX approach in general can be stated as:

no paths can be cut, but information can be lost.

A key property of SL is that function specifications are *compositional*, enabling *scalable* reasoning about the heap. This is due to their *locality*, which allows the pre-condition to describe only the partial state sufficient for the function to execute, and the *frame* property, which allows the function to be called in any larger state. SL function specifications have been used for verification of complex, real-world code in tools such as VeriFast [19], Iris [20], and Gillian [10, 23]. However, given that their post-conditions may describe states that are not reachable from their pre-conditions, such OX specifications are not compatible with true bug-finding, as found, for example, in Meta’s Pulse [28] and Pulse-X [21] tools.

Under-approximating (UX) program logics were recently introduced, originating from reverse Hoare logic (RHL) [8] for reasoning about correctness of probabilistic programs, and coming to prominence with incorrectness logic [26] and incorrectness separation logic (ISL) [28], which identified their bug-finding potential. ISL function specifications are of the form $[P] f(\vec{x}) [ok : Q_{ok}]$ and $[P] f(\vec{x}) [err : Q_{err}]$, the meaning of which is that any state in the success post-condition Q_{ok} or the error post-condition Q_{err} is reachable from some state in the pre-condition P by executing the function f ; this guarantees that all results and bugs reported in the post-conditions will be true. In contrast to SL, ISL uses the rule of *backward consequence*, which allows one to *cut paths* (for example, if we had a post-condition with $x > 0$, we could soundly strengthen this information to consider only the path in which $x = 42$). Therefore, the philosophy underlying UX logics in general can be summarised as:

paths can be cut, but no information can be lost.

When it comes to the use of ISL function specifications, whilst this has been implemented in Pulse-X, as far as we are aware, ISL does not feature function-call rules, and function compositionality for ISL and UX logics has not been proven. Moreover, as it is not possible to determine if UX specifications cover all terminating behaviour, they remain incompatible with verification and cannot therefore be used in tools such as VeriFast, Iris, and Gillian.

Our challenge is to develop a program logic in which we can state and prove function specifications that are compatible with *both* verification and true bug-finding. Our motivation comes from the unique flexibility and expressivity that such specifications would provide, as they could be used by verification and bug-finding tools alike, closing the gap between these two contrasting paradigms. From our experience in program logics and associated tool-building, we believe that the main use case for exact specification should be self-contained, critical code, such as widely-used data-structure libraries.

We introduce exact separation logic (ESL), with *exact* (EX) function specifications of the form $(P) f(\vec{x}) (ok : Q_{ok}) (err : Q_{err})$, whose meaning combines that of SL and ISL specifications: all terminating executions of the function that start from a state in the

pre-condition P end in a state covered by the post-conditions; *and* all states in two post-conditions are reachable from a state in the pre-condition by executing the function. The exactness of ESL can be captured by the slogan:

no paths can be cut and no information can be lost.

The slogan is supported by the rule of *equivalence*, which combines the forward consequence of SL and the backward consequence of ISL. In fact, ESL proof rules form a common core of SL and ISL, and ESL should therefore be a familiar setting to those acquainted with either.

We prove soundness for ESL with mutually recursive functions, which we believe is the first proof of function compositionality for a UX logic, and which transfers immediately to ISL. In doing so, drawing inspiration from InsecSL [25], we provide formal definitions of *external* and *internal* function specifications, which describe, respectively, the interface a function exposes towards its clients and towards its implementation, and highlight the difference in complexity between these two types of specifications in OX and UX reasoning.

Using numerous examples, we demonstrate here and in the extended version [24] how ESL can be used to reason about data-structure libraries, language errors, mutual recursion, and non-termination. In doing so, we introduce, for the first time, abstract predicates to UX reasoning and provide abstract function specifications for a number of data-structure algorithms, focussing on singly-linked lists and binary trees. In doing so, we highlight an important difference between the concepts of abstraction and over-approximation: in particular, abstraction corresponds to *hiding* information whereas over-approximation corresponds to *losing* it. Our findings demonstrate that, while abstraction cannot be used as freely in UX logics as in OX logics, sometimes resulting in less abstract specifications and more complex proofs, it should be feasible to use ESL to provide tractable function specifications for self-contained, critical code that can then be used for both verification and true bug-finding.

2 Exact Separation Logic by Example

We guide the reader through what it means to write ESL specifications and proofs by intuition and example, contrasting our findings with those known from SL and ISL.

Illustrative Example. Consider the command $C \triangleq \text{if } (x > 0) \{y := 42\} \text{ else } \{y := 21\}$, which can be specified, starting from the pre-condition $x \in \mathbb{Z}$, in ESL, SL, and ISL as follows:

<pre>(x ∈ ℤ) if (x > 0) { (x > 0) y := 42 (Q₁ : x > 0 ∧ y = 42) } else { (x ≤ 0) y := 21 (Q₂ : x ≤ 0 ∧ y = 21) } (Q₁ ∨ Q₂)</pre>	<pre>{ x ∈ ℤ } if (x > 0) { ... // Same as ESL ... } { Q₁ ∨ Q₂ } // Losing information { x ∈ ℤ ∧ y > 0 }</pre>	<pre>[x ∈ ℤ] if (x > 0) { [x > 0] y := 42 [Q₁ : x > 0 ∧ y = 42] } else { y := 21 } // Path cutting [x > 0 ∧ y = 42]</pre>
---	--	--

As ESL specifications must neither cut paths nor lose information (in this example, about the values of x and y), the ESL post-condition of C must be equivalent to $(x > 0 \wedge y = 42) \vee (x \leq 0 \wedge y = 21)$. In SL, it is possible to use forward consequence to weaken this

information and obtain, for example, $x \in \mathbb{Z} \wedge y > 0$, or just $x \in \mathbb{Z}$, or even just true . In ISL, it is possible to cut, for example, the **else** branch of the **if** statement, but the values of x and y must be maintained in the post-condition of the **then** branch, $x > 0 \wedge y = 42$.

One question that we have been often asked is whether it is simpler to prove an exact specification $(P) \text{ C } (ok : Q_{ok}) (err : Q_{err})$ in ESL, or to prove it separately in SL and ISL. The answer is that it is simpler to prove the specification in ESL. If a specification is exact, then it does not cut paths and it does not lose information. Therefore, the tools that make SL and ISL proofs simpler than ESL proofs, namely forward consequence and backward consequence, can only be used in very limited ways, if at all. From our experience, the ISL proof of an exact specification will turn out to be almost identical to the ESL one, and an SL proof on top of that would duplicate a large part of the work. In fact, if one were to try to prove the exact specification $(x \in \mathbb{Z}) \text{ C } ((x > 0 \wedge y = 42) \vee (x \leq 0 \wedge y = 21))$ from the above example in either SL or ISL, they would obtain exactly the same proof as in ESL.

We also emphasise that ESL is not meant to replace either SL or ISL. If one is interested in only verification or only bug-finding, then one should use a formalism tailored to that type of analysis to exploit the available shortcuts. However, if one wanted to use the same codebase for both verification and bug-finding, then ESL offers a way of providing specifications useful for both. One example of such a codebase would be a widely-used data-structure library, where some of the users use it for verification and others for bug-finding.

List-length in ESL: Intuition. We consider a list-length function, $\text{LLen}(x)$, which takes a list at x , does not modify it, and returns its length, and the following ESL specification:

$$(x = x \star \text{list}(x, n)) \text{ LLen}(x) (ok : \text{list}(x, n) \star \text{ret} = n)$$

This specification uses a standard list-length predicate, $\text{list}(x, n)$, which states that the length of the list at x equals n and is defined as follows:

$$\text{list}(x, n) \triangleq (x = \text{null} \star n = 0) \vee (\exists v, x'. x \mapsto v, x' \star \text{list}(x', n - 1)),$$

hiding the information about the values and internal node addresses of the list. Before proving this specification, we establish some intuition about why it holds. Let us assume that it does not hold and try to find a counter-example: by the meaning of ESL specifications, it is either not OX-valid or it is not UX-valid. The former, however, is not possible, as the analogous SL specification holds. The latter means that it is possible to find a state in the post-condition not reachable by the execution of f from any state in the pre-condition, and may be unfamiliar to the reader as UX program logics have been introduced only recently.

We start looking for such a state in the post-condition (post-model) by choosing some values for x and n : say, $x = 0$ and $n = 2$. This also fixes ret to 2. Then, we fully unfold $\text{list}(0, 2)$ to obtain $\exists v_1, x_1, v_2. 0 \mapsto v_1, x_1 \star x_1 \mapsto v_2, \text{null}$, and instantiate the existentials v_1 , x_1 , and v_2 : say, with 1, 4, and 9, respectively. In this way, we obtain the state described by the assertion $0 \mapsto 1, 4 \star 4 \mapsto 9, \text{null}$. When it comes to the pre-condition, x and n (and also x) are fixed by the post-model choices, and when we unfold the list, the pre-condition becomes $x = 0 \star \exists v_1, x_1, v_2. 0 \mapsto v_1, x_1 \star x_1 \mapsto v_2, \text{null}$. As the algorithm does not modify the list, it becomes clear that if we choose v_1 , x_1 , and v_2 as for the post-model (that is, 1, 4, and 9, respectively), the algorithm will reach our post-model. Given that the same reasoning would apply for any choice of x and n , we realise that the given specification is, in fact, also UX-valid and hence exact. This reveals an important observation, which is that

*abstraction does not always equate to over-approximation, that is,
hiding information does not always mean losing information.*

For those used to OX reasoning, it might appear that the post-condition $\text{list}(x, n) \star \text{ret} = n$ loses information about the structure of the list, but the insight here is that this information was never known in the pre-condition in the first place, as we also only had $\text{list}(x, n)$ there.

List-length in ESL: Proof Sketch. Reasoning about function specifications in the UX/EX setting has not been studied previously and requires subtle definitions of *external function specifications*, which provide the interface that the function exposes to the client, and *internal function specifications*, which provide the interface to the function's implementation. With OX logics, these are well-understood and the gap between them is small. For UX/EX logics, this gap is larger. We illustrate these concepts informally using the list-length example, and give the corresponding formal definitions in §4.

The proof sketch of the ESL external specification of the list-length algorithm is given in Figure 1. It is more complex than its SL counterpart (cf. [24]), but is manageable and comes with the benefit that this ESL specification can be used for both verification and bug-finding.

First, as the function is recursive, we have to provide a measure and prove the specification extended with this measure: in this case, the measure is $\alpha = n$, given by the length of the list. This measure is necessary to ensure the finite reachability property for mutually recursive functions in UX logics, and is a known technique from the work on total correctness specifications for OX logics [7, 9]. Recursive function calls are then allowed only if they use specifications of a strictly smaller measure, represented in the proof sketch by the function specification context $\Gamma(\alpha)$, which contains the specification of $\text{LLen}(x)$ for all $\beta < \alpha$.

The move from the external to the internal pre-condition initialises the local function variables to `null`. The ESL rule for the `if` statement, just like in SL, adds the condition to the then-branch, its negation to the else-branch, and collects the branch post-conditions using disjunction. The rules for the basic commands (here, the assignments $r := 0$ and $r := r + 1$ and the lookup $x := [x + 1]$) are also the same as in SL, as these are already exact. The unfolding of the list is also done in the same way, as unfolding always preserves equivalence; note how the condition of the `if` statement determines the appropriate disjunct for the list predicate. The recursive function call is allowed to go through as it is used with measure $n - 1$ (with the parts of the assertion representing the pre- and the post-condition highlighted).

The major difference between ISL/ESL and SL proofs is that we cannot lose information about the function parameters and local variables in the middle of the former. Therefore, we cannot simplify the assertions Q'_1 and Q'_2 further and cannot fold back the list predicate within the internal specification, as we would do in SL (cf. corresponding proof in [24]).

The most complex part of the proof sketch is the transition from the internal to the external post-condition, in which we have to somehow forget the local variables of the function, given that they must not spill out into the calling context. This is done by replacing them with fresh, existentially quantified logical variables, which in this case also allows us to use equivalence to fold back the list predicate and reach the target post-condition. The details of this transition, in which we denote $\text{ret} = n \star \alpha = n$ by R , are as follows:

$$\begin{aligned}
& \exists x_q, r_q. Q'[x_q/x][r_q/r] \star \text{ret} = r[x_q/x][r_q/r] \\
& \Leftrightarrow ((x = \text{null} \star n = 0) \vee (\exists x_q, r_q, v, x'. x_q = x' \star x \mapsto v, x' \star \text{list}(x', n - 1) \star r_q = n)) \star R \\
& \Leftrightarrow ((x = \text{null} \star n = 0) \vee (\exists v, x'. x \mapsto v, x' \star \text{list}(x', n - 1))) \star R \text{ // can fold now} \\
& \Leftrightarrow \text{list}(x, n) \star (n = 0 \vee n > 0) \star R \\
& \Leftrightarrow \text{list}(x, n) \star \text{ret} = n \star \alpha = n
\end{aligned}$$

Observe that, since we are proving an EX specification, we are not allowed to cut paths. This means that the ISL proof of the analogous ISL specification of $\text{LLen}(x)$ would be identical, noting that the use of equivalence would technically be replaced by backward consequence.

```

// Function is recursive and requires a measure:  $\alpha = n$ 
 $\Gamma(\alpha) \vdash (x = x \star \text{list}(x, n) \star \alpha = n)$ 
  LLen( $x$ ) {
    // Transition from external to internal pre-condition: initialise locals to null
    ( $x = x \star \text{list}(x, n) \star \alpha = n \star r = \text{null}$ )
    if ( $x = \text{null}$ ) {
      ( $x = x \star \text{list}(x, n) \star \alpha = n \star r = \text{null} \star x = \text{null}$ )
       $r := 0$ 
      ( $Q'_1 : x = x \star \text{list}(x, n) \star \alpha = n \star r = 0 \star x = \text{null}$ )
    } else {
      ( $x = x \star \text{list}(x, n) \star \alpha = n \star r = \text{null} \star x \neq \text{null}$ )
      // Unfold  $\text{list}(x, n)$  using the equivalence
      //  $\models \text{list}(x, n) \star x \neq \text{null} \Leftrightarrow \exists v, x'. x \mapsto v, x' \star \text{list}(x', n - 1)$ 
      ( $\exists v, x'. x = x \star x \mapsto v, x' \star \text{list}(x', n - 1) \star \alpha = n \star r = \text{null}$ )
       $x := [x + 1]$ ;
      ( $\exists v, x'. x = x' \star x \mapsto v, x' \star \text{list}(x', n - 1) \star \alpha = n \star r = \text{null}$ )
      // As  $\alpha - 1 < \alpha$ , we can use the specification of LLen( $x$ ) with measure  $\alpha - 1$ 
      ( $\exists v, x'. x = x' \star x \mapsto v, x' \star \text{list}(x', n - 1) \star \alpha - 1 = n - 1 \star r = \text{null}$ )
       $r := \text{LLen}(x)$ ;
      ( $\exists v, x'. x = x' \star x \mapsto v, x' \star \text{list}(x', n - 1) \star \alpha - 1 = n - 1 \star r = n - 1$ )
       $r := r + 1$ 
      ( $Q'_2 : \exists v, x'. x = x' \star x \mapsto v, x' \star \text{list}(x', n - 1) \star \alpha - 1 = n - 1 \star r = n$ )
    };
    ( $Q' : Q'_1 \vee Q'_2$ )
    return  $r$ 
    ( $Q' \star \text{ret} = r$ )
    // Transition from internal to external post-condition given in text
  }
  ( $\text{list}(x, n) \star \text{ret} = n \star \alpha = n$ )

```

■ **Figure 1** ESL proof sketch: $(x = x \star \text{list}(x, n)) \text{ LLen}(x) \text{ (ok : list}(x, n) \star \text{ret} = n)$.

List-insert in ESL: Intuition. The list-length function, $\text{LLen}(x)$, is an example of an algorithm where the EX specification is analogous to the traditional OX specification. At times, however, ESL specifications have to be more complex. Consider, for example, the list-insert algorithm $\text{LInsertFirst}(x, v)$, which inserts the element v at the beginning of the list x . Its traditional OX specification is:

$$\{x = x \star v = v \star \text{list}(x, vs)\} \text{ LInsertFirst}(x, v) \{ \text{list}(\text{ret}, v : vs) \}$$

where $\text{list}(x, vs)$ is the standard list predicate that exposes the values of the list:

$$\text{list}(x, vs) \triangleq (x = \text{null} \star vs = []) \vee (\exists v, x', vs'. x \mapsto v, x' \star \text{list}(x', vs') \star vs = v : vs')$$

Using the counter-example approach to check if this specification is EX-valid, we easily see that it loses information: in particular, no end-state where x is not the *second* pointer in the returned list ret is reachable from the given pre-condition. Consequently, for EX validity, we are required to use the following, less abstract, ESL specification for LInsertFirst :

$$(x = x \star v = v \star \text{list}(x, xs, vs)) \text{ LInsertFirst}(x, v) (\text{list}(\text{ret}, \text{ret} : xs, v : vs) \star \text{listHead}(x, xs))$$

where $\text{list}(x, xs, vs)$ is a predicate that exposes the internal pointers of a given list in addition to the values, and $\text{listHead}(x, xs)$ states that the list xs starts with x .

Further Examples. In §5 and [24], we give many additional examples of ESL specifications and proofs to illustrate reasoning about list algorithms and binary trees, as well as language errors, mutual recursion, non-termination, and client programs.

3 The Programming Language

We introduce ESL using a simple programming language, the syntax of which is given below.

Language Syntax

$$\begin{array}{l}
 v \in \text{Val} ::= n \in \text{Nat} \mid b \in \text{Bool} \mid s \in \text{Str} \mid \text{null} \mid \bar{v} \quad x \in \text{PVar} \\
 E \in \text{PExp} ::= v \mid x \mid E + E \mid E - E \mid \dots \mid E = E \mid E < E \mid \neg E \mid E \wedge E \mid \dots \mid E : E \mid E @ E \mid \dots \\
 C \in \text{Cmd} ::= \text{skip} \mid x := E \mid x := \text{nondet} \mid \text{error}(E) \mid \text{if } (E) C \text{ else } C \mid \text{while } (E) C \mid C ; C \mid \\
 \quad y := f(\bar{E}) \mid x := [E] \mid [E] := E \mid x := \text{new}(E) \mid \text{free}(E)
 \end{array}$$

Values, $v \in \text{Val}$, include: natural numbers, $n \in \text{Nat}$; Booleans, $b \in \text{Bool} \triangleq \{\text{true}, \text{false}\}$; strings, $s \in \text{Str}$; a dedicated value `null`; and lists of values, $\bar{v} \in \text{List}$. Expressions, $E \in \text{PExp}$, comprise values, program variables, $x \in \text{PVar}$, and various unary and binary operators (e.g., addition, equality, negation, conjunction, list prepending, and list concatenation). Commands comprise: the variable assignment; non-deterministic number generation; error raising; the `if` statement; the `while` loop; command sequencing; function call; and memory management commands, that is, lookup, mutation, allocation, and deallocation. The sets of program variables for expressions and commands, denoted by $\text{pv}(E)$ and $\text{pv}(C)$ respectively, and the sets of modified variables for commands, denoted by $\text{mod}(C)$, are defined in the standard way.

► **Definition 1 (Functions).** A function, denoted by $f(\bar{x}) \{ C; \text{return } E \}$, comprises: a function identifier, $f \in \text{Str}$; the function parameters, \bar{x} , given by a list of distinct program variables; a function body, $C \in \text{Cmd}$; and a return expression, $E \in \text{PExp}$, with $\text{pv}(E) \subseteq \{\bar{x}\} \cup \text{pv}(C)$.

Program variables in function bodies that are not the function parameters are treated as local variables initialised to `null`, with their scope not extending beyond the function.

► **Definition 2 (Function Implementation Contexts).** A function implementation context, $\gamma : \text{Str} \rightarrow_{\text{fin}} \text{PVar List} \times \text{Cmd} \times \text{PExp}$, is a finite partial function from function identifiers to their implementations. For $\gamma(f) = (\bar{x}, C, E)$, we also write $f(\bar{x}) \{ C; \text{return } E \} \in \gamma$.

We next define an operational semantics that gives a complete account of the behaviour of commands and does not get stuck on any input, as we explicitly account for language errors and missing resource errors.

► **Definition 3 (Stores, Heaps, States).** Variable stores, $s : \text{PVar} \rightarrow_{\text{fin}} \text{Val}$, are partial finite functions from program variables to values. Heaps, $h : \text{Nat} \rightarrow_{\text{fin}} (\text{Val} \uplus \emptyset)$, are partial finite functions from natural numbers to values extended with a dedicated symbol $\emptyset \notin \text{Val}$. Program states, $\sigma = (s, h)$, consist of a store and a heap.

Heaps are used to model the memory, and the dedicated symbol $\emptyset \notin \text{Val}$ is required for UX frame preservation¹ to hold (cf. Definition 10). In particular, $h(n) = v$ means that an allocated heap cell with address n contains the value v ; and $h(n) = \emptyset$ means that a heap

¹ UX frame preservation means that if a program runs with a non-missing outcome to a given final state, then it also runs with the same outcome to an extended final state, with the extension (the *frame*) unaffected by the execution. From ISL [28], it is known that losing deallocation information breaks UX frame preservation; the solution is to keep track of deallocated cells, which we achieve by using \emptyset .

cell with address n has been deallocated [11–14, 28]. This linear memory model is used in much of the SL literature, including ISL [28]. Onward, \emptyset denotes the empty heap, $h_1 \uplus h_2$ denotes heap disjoint union, and $h_1 \# h_2$ denotes that h_1 and h_2 are disjoint.

► **Definition 4** (Expression Evaluation). *The evaluation of an expression E with respect to a store s , denoted $\llbracket E \rrbracket_s$, results in either a value or a dedicated symbol denoting an evaluation error, $\dagger \notin \text{Val}$. Some illustrative cases are:*

$$\llbracket v \rrbracket_s = v \quad \llbracket x \rrbracket_s = \begin{cases} s(x), & x \in \text{dom}(s) \\ \dagger, & \text{otherwise} \end{cases} \quad \llbracket E_1 + E_2 \rrbracket_s = \begin{cases} \llbracket E_1 \rrbracket_s + \llbracket E_2 \rrbracket_s, & \llbracket E_1 \rrbracket_s, \llbracket E_2 \rrbracket_s \in \text{Nat} \\ \dagger, & \text{otherwise} \end{cases}$$

The big-step operational semantics uses judgements of the form $\sigma, C \Downarrow_\gamma o : \sigma'$, read: given implementation context γ and starting from state σ , the execution of command C results in outcome $o \in O = \{ok, err, miss\}$ and state σ' . The outcome can either equal: *ok* (elided where possible), denoting a successful execution; *err*, denoting an execution faulting with a language error, or *miss*, denoting an execution faulting with a missing resource error.

► **Definition 5** (Operational Semantics). *The representative cases of the big-step operational semantics are given in Figure 2. The complete semantics is given in [24].*

The successful transitions are straightforward: for example, the `nondet` command generates an arbitrary natural number; the function call executes the function body in a store where the function parameters are given the values of the function arguments and the function locals are initialised to `null`; and the control flow statements behave as expected. Allocation requires the specified amount of contiguous cells (always available as heaps are finite), and lookup, mutation, and deallocation require the targeted cell not to have been freed.

The semantics stores error information in a dedicated program variable `err`, not available to the programmer. For simplicity of error messages, we assume to have a function `str : PExp → Str`, which serialises program expressions into strings. The faulting semantic transitions are split into *language errors*, which can be captured by program-logic reasoning, and *missing resource errors*, which cannot, as such errors break the frame property. Language errors arise due to, for example, expressions being incorrectly typed (e.g. `null + 1`) or an attempt to access deallocated cells (that is, the use-after-free error). On the other hand, missing resource errors arise from accessing cells that are not present in memory.

4 Exact Separation Logic

We introduce an exact separation logic for our programming language, giving the assertion language in §4.1, specifications in §4.2, and the program logic rules in §4.3.

4.1 Assertion Language

To define assertions and their meaning, we introduce *logical variables*, $x, y, z, \in \text{LVar}$, distinct from program variables, and define the set of *logical expressions* as follows:

$$E \in \text{LExp} \triangleq v \mid x \mid \times \mid E + E \mid E - E \mid \dots \mid E = E \mid \neg E \mid E \wedge E \mid \dots \mid E \cdot E \mid E : E \mid \dots$$

Note that we can use program expressions in assertions (for example, $E \in \text{Val}$), as they form a proper subset of logical expressions.

$$\begin{array}{c}
\frac{\llbracket E \rrbracket_s = v \quad s' = s[x \rightarrow v]}{(s, h), x := E \Downarrow_\gamma (s', h)} \quad \frac{n \in \mathbb{N} \quad s' = s[x \rightarrow n]}{(s, h), x := \text{nondet} \Downarrow_\gamma (s', h)} \quad \frac{\llbracket E \rrbracket_s = \text{false}}{(s, h), \text{while } (E) \text{ C} \Downarrow_\gamma (s, h)} \\
\frac{\llbracket E \rrbracket_s = \text{true} \quad (s, h), \text{C} \Downarrow_\gamma \sigma'' \quad \sigma'', \text{while } (E) \text{ C} \Downarrow_\gamma o : \sigma'}{(s, h), \text{while } (E) \text{ C} \Downarrow_\gamma o : \sigma'} \quad \frac{f(\vec{x}) \{ \text{C}; \text{return } E' \} \in \gamma \quad \llbracket \vec{E} \rrbracket_s = \vec{v} \quad \text{pv}(\text{C}) \setminus \{ \vec{x} \} = \{ \vec{z} \} \quad s_p = \emptyset [\vec{x} \rightarrow \vec{v}] [\vec{z} \rightarrow \text{null}] \quad (s_p, h), \text{C} \Downarrow_\gamma (s_q, h') \quad \llbracket E' \rrbracket_{s_q} = v'}{(s, h), y := f(\vec{E}) \Downarrow_\gamma (s[y \rightarrow v'], h')} \\
\frac{\llbracket E \rrbracket_s = n \quad h(n) = v}{(s, h), x := [E] \Downarrow_\gamma (s[x \rightarrow v], h)} \quad \frac{\llbracket E_1 \rrbracket_s = n \quad h(n) \in \text{Val} \quad \llbracket E_2 \rrbracket_s = v \quad h' = h[n \mapsto v]}{(s, h), [E_1] := E_2 \Downarrow_\gamma (s, h')} \\
\frac{(n' + i \notin \text{dom}(h)) \Big|_{i=1}^{\llbracket E \rrbracket_s - 1} \quad h' = h \uplus \{ (n' + i \mapsto \text{null}) \Big|_{i=1}^{\llbracket E \rrbracket_s - 1} \}}{(s, h), x := \text{new}(E) \Downarrow_\gamma (s[x \rightarrow n'], h')} \quad \frac{\llbracket E \rrbracket_s = n \quad h(n) \in \text{Val}}{(s, h), \text{free}(E) \Downarrow_\gamma (s, h[n \mapsto \emptyset])} \\
\frac{\llbracket E \rrbracket_s = \downarrow \quad v_{\text{err}} = [\text{"ExprEval"}, \text{str}(E)]}{(s, h), x := [E] \Downarrow_\gamma \text{err} : (s_{\text{err}}, h)} \quad \frac{\llbracket E \rrbracket_s = n \notin \text{dom}(h) \quad v_{\text{err}} = [\text{"MissingCell"}, \text{str}(E), n]}{(s, h), x := [E] \Downarrow_\gamma \text{miss} : (s_{\text{err}}, h)} \\
\frac{h(\llbracket E \rrbracket_s) = \emptyset \quad v_{\text{err}} = [\text{"UseAfterFree"}, \text{str}(E_1), \llbracket E \rrbracket_s]}{(s, h), x := [E] \Downarrow_\gamma \text{err} : (s_{\text{err}}, h)} \quad \frac{\llbracket E \rrbracket_s = v \quad v_{\text{err}} = [\text{"Error"}, v]}{(s, h), \text{error}(E) \Downarrow_\gamma \text{err} : (s_{\text{err}}, h)}
\end{array}$$

■ **Figure 2** Operational semantics (excerpt), with $s_{\text{err}} \triangleq s[\text{err} \rightarrow v_{\text{err}}]$ and $\text{str} : \text{PExp} \rightarrow \text{Str}$.

► **Definition 6** (Assertion Language). *The assertion language is defined as follows:*

$$\begin{aligned}
\pi \in \text{BArt} &\triangleq E_1 = E_2 \mid E_1 < E_2 \mid E \in X \mid \dots \mid \neg \pi \mid \pi_1 \Rightarrow \pi_2 \\
P \in \text{Asrt} &\triangleq \pi \mid \text{False} \mid P_1 \Rightarrow P_2 \mid \exists x. P \mid \text{emp} \mid E_1 \mapsto E_2 \mid E \mapsto \emptyset \mid P_1 \star P_2 \mid \otimes_{E_1 \leq x < E_2} P
\end{aligned}$$

where $E, E_1, E_2 \in \text{LExp}$, $X \subseteq \text{Val}$, and $x \in \text{LVar}$.

Boolean assertions, $\pi \in \text{BArt}$, lift Boolean logical expressions to assertions. Assertions, $P \in \text{Asrt}$, contain Boolean assertions, standard first-order connectives and quantifiers, and spatial assertions. Spatial assertions include: the empty memory assertion, emp ; the positive cell assertion, $E_1 \mapsto E_2$; the negative cell assertion, $E \mapsto \emptyset$ (as in [11–14] and denoted in ISL by $E \not\mapsto$ [28]), the separating conjunction (star); and its iteration (iterated star).

To define assertion satisfiability, we introduce *substitutions*, $\theta : \text{LVar} \rightarrow_{\text{fin}} \text{Val}$, which are partial finite mappings from logical variables to values, extending expression evaluation of Definition 4 to $\llbracket E \rrbracket_{\theta, s}$ straightforwardly, with a new base case for logical variables:

$$\llbracket x \rrbracket_{\theta, s} = \theta(x), \text{ if } x \in \text{dom}(\theta) \quad \llbracket x \rrbracket_{\theta, s} = \downarrow, \text{ if } x \notin \text{dom}(\theta)$$

► **Definition 7** (Satisfiability). *The assertion satisfiability relation, denoted by $\theta, \sigma \models P$, is defined as follows:*

$$\begin{aligned}
\theta, (s, h) \models \pi &\Leftrightarrow \llbracket \pi \rrbracket_{\theta, s} = \text{true} \wedge h = \emptyset \\
\theta, (s, h) \models \text{False} &\Leftrightarrow \text{never} \\
\theta, (s, h) \models P_1 \Rightarrow P_2 &\Leftrightarrow \theta, (s, h) \models P_1 \Rightarrow \theta, (s, h) \models P_2 \\
\theta, (s, h) \models \exists x. P &\Leftrightarrow \exists v \in \text{Val}. \theta[x \mapsto v], (s, h) \models P \\
\theta, (s, h) \models \text{emp} &\Leftrightarrow h = \emptyset \\
\theta, (s, h) \models E_1 \mapsto E_2 &\Leftrightarrow h = \{ \llbracket E_1 \rrbracket_{\theta, s} \mapsto \llbracket E_2 \rrbracket_{\theta, s} \} \\
\theta, (s, h) \models E_1 \mapsto \emptyset &\Leftrightarrow h = \{ \llbracket E_1 \rrbracket_{\theta, s} \mapsto \emptyset \} \\
\theta, (s, h) \models P_1 \star P_2 &\Leftrightarrow \exists h_1, h_2. h = h_1 \uplus h_2 \wedge \theta, (s, h_1) \models P_1 \wedge \theta, (s, h_2) \models P_2 \\
\theta, (s, h) \models \otimes_{E_1 \leq x < E_2} P &\Leftrightarrow \exists h_i, \dots, h_{k-1}. h = \uplus_{j=i}^{k-1} h_j \wedge \forall j. i \leq j < k \Rightarrow \theta, (s, h_j) \models P[j/x] \\
&\text{where } i = \llbracket E_1 \rrbracket_{\theta, s}, k = \llbracket E_2 \rrbracket_{\theta, s}, \text{ and } x \text{ is not free in } E_1 \text{ or } E_2.
\end{aligned}$$

19:10 Exact Separation Logic

Assertion satisfiability is defined in the standard way. For convenience, we choose Boolean assertions to be satisfiable only in the empty heap.

► **Definition 8** (Validity). *An assertion P is valid, denoted by $\models P$, iff $\forall \theta, \sigma. \theta, \sigma \models P$.*

4.2 Specifications

We define specifications for commands and functions, focussing in particular on external and internal function specifications and the relationship between them.

► **Definition 9.** Specifications, $t = (P) (ok : Q_{ok}) (err : Q_{err}) \in \mathcal{Spec}$, comprise a pre-condition, P , a success post-condition, Q_{ok} , and a faulting post-condition, Q_{err} .

We denote that command C has specification t by $C : t$, or by $(P) C (ok : Q_{ok}) (err : Q_{err})$ in quadruple form. Additionally, we use the following shorthand:

$$\begin{aligned} (P) C (Q) &\triangleq (P) C (ok : Q) (err : \text{False}) \\ (P) C (err : Q) &\triangleq (P) C (ok : \text{False}) (err : Q) \\ (P) C (Q) &\triangleq (P) C (ok : -) (err : -) \end{aligned}$$

noting the use of Q for cases in which the post-condition details are not relevant. We use quadruples rather than triples since, even though the post-condition could be expressed as a disjunction of *ok*- and *err*-labelled assertions, we find the quadruple distinction helpful as compound commands (e.g. sequence) treat the two differently (cf. Figure 3).

The EX-validity of a specification t for a command C in an implementation context γ requires both OX and UX frame-preserving validity.

► **Definition 10** (γ -Valid Specifications). *Given implementation context γ , command C , and specification $t = (P) (ok : Q_{ok}) (err : Q_{err})$, t is γ -valid for C , denoted by $\gamma \models C : t$ or $\gamma \models (P) C (ok : Q_{ok}) (err : Q_{err})$, if and only if:*

$$\begin{aligned} & // \text{Frame-preserving over-approximating validity} \\ & (\forall \theta, s, h, h_f, o, s', h''. \theta, (s, h) \models P \implies \\ & \quad (s, h \uplus h_f), C \Downarrow_{\gamma} o : (s', h'') \implies (o \neq \text{miss} \wedge \exists h'. h'' = h' \uplus h_f \wedge \theta, (s', h') \models Q_o)) \wedge \\ & // \text{Frame-preserving under-approximating validity} \\ & (\forall \theta, s', h', h_f, o. \theta, (s', h') \models Q_o \implies h_f \# h' \implies \\ & \quad (\exists s, h. \theta, (s, h) \models P \wedge (s, h \uplus h_f), C \Downarrow_{\gamma} o : (s', h' \uplus h_f))) \end{aligned}$$

Observe that the outcome o can either be success or a language error; it cannot be a missing resource error as this would break UX frame preservation. As our operational semantics is complete, we can also use ESL to characterise non-termination. In particular, if a command satisfies a specification with both post-conditions **False**, then the execution of the command is guaranteed to not terminate if executed from a state satisfying the pre-condition. Were the semantics incomplete (for example, if it did not reason about errors), then such a specification might also indicate the absence of a semantic transition.

Compared to traditional OX reasoning, UX reasoning brings additional complexity to proofs of function specifications. To handle this complexity, we introduce two types of function specifications: *external* specifications, which provide the interface the function exposes to the client, and the related *internal* specifications, which provide the interface to the function implementation. This terminology is also used informally in InsecSL [25]. We use these in subsequent sections to show that ESL exhibits function compositionality.

► **Definition 11** (External Specifications). *A specification $(P) (ok : Q_{ok}) (err : Q_{err})$ is an external function specification if and only if:*

- $P = (\vec{x} = \vec{x} \star P')$, for some distinct program variables \vec{x} , distinct logical variables \vec{x} , and assertion P' , with $\text{pv}(P') = \emptyset$; and
- either $\text{pv}(Q_{ok}) = \{\text{ret}\}$ or $Q_{ok} = \text{False}$, and either $\text{pv}(Q_{err}) = \{\text{err}\}$ or $Q_{err} = \text{False}$.

The set of external specifications is denoted by \mathcal{ESpec} .

► **Definition 12** (Function Specification Contexts). *A function specification context, $\Gamma : \text{Fid} \rightarrow_{\text{fin}} \mathcal{P}(\mathcal{ESpec})$, is a finite partial function from function identifiers to a set of external specifications, with the more familiar notation $(\vec{x} = \vec{x} \star P) f(\vec{x}) (ok : Q_{ok}) (err : Q_{err}) \in \Gamma$ at times used in place of $(\vec{x} = \vec{x} \star P) (ok : Q_{ok}) (err : Q_{err}) \in \Gamma(f)$.*

The constraints on external specifications are well-known from OX logics and follow the usual scoping of function parameters and local variables, which are limited to the function body: the pre-conditions only contain the function parameters, \vec{x} ; and the post-conditions may only have the (dedicated) program variables ret or err , which hold, respectively, the return value on successful termination or the error value on faulting termination.

Internal function specifications are more interesting for exact and UX than for OX reasoning. The internal pre-condition is straightforward, extending the external pre-condition by instantiating the local variables to null . The internal post-condition must therefore include information about the parameters and local variables, as the internal specification cannot lose information. This means that the connection between internal and external post-conditions is subtle, given the constraints on the latter. To address this, we define an internalisation function, relating an external function specification with a set of possible internal specifications. In particular, the external post-condition has to be equivalent to an internal one in which the parameters and local variables of the internal post-condition have been replaced by fresh existentially quantified logical variables.

► **Definition 13** (Internalisation). *Given implementation context γ and function $f \in \text{dom}(\gamma)$, a function specification internalisation, $\text{Int}_{\gamma, f} : \mathcal{ESpec} \rightarrow \mathcal{P}(\text{Spec})$, is defined as follows:*

$$\begin{aligned} \text{Int}_{\gamma, f}((P) (ok : Q_{ok}) (err : Q_{err})) = \\ \{(P \star \vec{z} = \text{null}) (ok : Q'_{ok}) (err : Q'_{err}) \mid & \models Q'_{ok} \Rightarrow E \in \text{Val} \star \text{True} \text{ and} \\ & \models Q_{ok} \Leftrightarrow \exists \vec{p}. Q'_{ok}[\vec{p}/\vec{p}] \star \text{ret} = E[\vec{p}/\vec{p}] \text{ and} \\ & \models Q_{err} \Leftrightarrow \exists \vec{p}. Q'_{err}[\vec{p}/\vec{p}]\}, \end{aligned}$$

where $f(\vec{x})\{C; \text{return } E\} \in \gamma$, $\vec{z} = \text{pv}(C) \setminus \text{pv}(P)$, $\vec{p} = \text{pv}(P) \uplus \{\vec{z}\}$, and the logical variables \vec{p} are fresh with respect to Q_{ok} and Q_{err} .

This approach also works for SL and ISL as well (with \Leftarrow instead of \Leftrightarrow for the post-conditions for SL, and \Rightarrow instead of \Leftrightarrow for ISL). It is not strictly necessary for SL, however, as information about program variables can be forgotten in the internal post-conditions before the transition to the external post-condition.

► **Definition 14** (Environments). *An environment, (γ, Γ) , is a pair consisting of an implementation context γ and a specification context Γ .*

An environment (γ, Γ) is *valid* if and only if every function specified in Γ has an implementation in γ and every specification in Γ has a γ -valid internal specification.

► **Definition 15** (Valid Environments). *Given an implementation context γ and a specification context Γ , the environment (γ, Γ) is valid, written $\models (\gamma, \Gamma)$, if and only if*

$$\text{dom}(\Gamma) \subseteq \text{dom}(\gamma) \wedge (\forall f, \vec{x}, C, E. f(\vec{x})\{C; \text{return } E\} \in \gamma \implies (\forall t. t \in \Gamma(f) \implies \exists t' \in \text{Int}_{\gamma, f}(t). \gamma \models C : t'))$$

Finally, a specification t is valid for a command C in a specification context Γ if and only if t is γ -valid for all implementation contexts γ that validate Γ .

► **Definition 16** (Γ -Valid Specifications). *Given a specification context Γ , a command C , and a specification $t = (P) (Q)$, the specification t is Γ -valid for command C , written $\Gamma \models C : t$ or $\Gamma \models (P) C (Q)$, if and only if $\forall \gamma. \models (\gamma, \Gamma) \implies \gamma \models (P) C (Q)$.*

4.3 Program Logic

We give the representative ESL proof rules in Figure 3 and all in [24]. We introduce and discuss in detail the function-related rules, given for the first time in a UX setting. We denote the repetition of the pre-condition in the post-condition by *pre*. When reading the rules, it is important to remember that we *must not drop paths and must not lose information*. The judgement $\Gamma \vdash (P) C (ok : Q_{ok}) (err : Q_{err})$ means that the specification t is *derivable* for a command C given the specifications recorded in Γ .

The basic command rules are fairly straightforward. The [NONDET] rule existentially quantifies the generated value via $x \in \mathbb{N}$ to capture all paths, in contrast with the RHL [8] and ISL [28] rules, which explicitly choose one value to describe one path. The $E' \in \text{Val}$ in the post-condition is necessary as we know that E' evaluates to a value from the pre-condition and cannot lose information; the same principle applies to many other rules. The [ASSIGN] rule requires that the evaluation of E does not fault in the pre-condition via $E \in \text{Val}$. Strictly speaking, we should have an additional case in which the assigned variable is not in the store. To avoid this clutter, we instead assume that program variables are always in the store as we are analysing function bodies and, in our programming language, all local variables are initialised on function entry. The error-related rules capture cases in which expression evaluation faults (e.g. [LOOKUP-ERR-VAL] rule, using $E \notin \text{Val}$), expressions are of the incorrect type, or memory is accessed after it has been freed (e.g. [LOOKUP-ERR-USE-AFTER-FREE] rule, using $E \mapsto \emptyset$). Note that missing resource errors cannot be captured without breaking frame preservation, as the added-on frame could contain the missing resource.

When it comes to composite commands, we opt for two *if*-rules, covering the branches separately. The sequencing rule shows how exact quadruples of successive commands can be joined together, highlighting, in particular, how errors are collected using disjunction. One interesting aspect of this rule is what happens when C_1 only throws an error or does not terminate, meaning that $R = \text{False}$. In both those cases, given the exactness of the rules, it has to be that $Q_{ok} = Q_{err}^2 = \text{False}$, and the post-condition of the sequence becomes $(ok : \text{False}) (err : Q_{err}^1 \vee \text{False})$, meaning that, if C_1 only throws an error (that is, $Q_{err}^1 \neq \text{False}$) then that is the only error that can come out of the sequence, and if C_1 does not terminate (that is, $Q_{err}^1 = \text{False}$) then the sequence does not terminate either.

The while rule is an adaptation of the RHL while rule [8], generalising the invariant of the SL while rule with two natural-number-indexed families of variants, P_i and Q_i , which explicitly maintain the iteration index. Note how the i in the premise is a meta-variable representing a natural number, which in the conclusion gets substituted for an existentially quantified logical variable; a similar principle will be applied later when dealing with environment extension. Interestingly, this rule does not require adjustment to reason about non-termination.

$$\begin{array}{c}
\text{SKIP} \\
\Gamma \vdash (\text{emp}) \text{ skip } (\text{emp}) \\
\text{NONDET} \\
\frac{x \notin \text{pv}(E') \quad Q \triangleq E' \in \text{Val} \star x \in \mathbb{N}}{\Gamma \vdash (x = E') x := \text{nondet } (Q)} \\
\text{ASSIGN} \\
\frac{x \notin \text{pv}(E') \quad Q \triangleq E' \in \text{Val} \star x = E[E'/x]}{\Gamma \vdash (x = E' \star E \in \text{Val}) x := E (Q)} \\
\text{LOOKUP} \\
\frac{Q \triangleq E' \in \text{Val} \star x = E_1[E'/x] \star E[E'/x] \mapsto E_1[E'/x]}{\Gamma \vdash (x = E' \star E \mapsto E_1) x := [E] (Q)} \\
\text{MUTATE} \\
\frac{Q \triangleq E_1 \mapsto E_2 \star E \in \text{Val}}{\Gamma \vdash (E_1 \mapsto E \star E_2 \in \text{Val}) [E_1] := E_2 (Q)} \\
\text{NEW} \\
\frac{Q \triangleq E' \in \text{Val} \star \bigotimes_{0 \leq i < E[E'/x]} ((x+i) \mapsto \text{null})}{\Gamma \vdash (x = E' \star E \in \mathbb{N}) x := \text{new}(E) (ok : Q)} \\
\text{ERROR} \\
\frac{E_{\text{err}} \triangleq [\text{"Error"}, E]}{\Gamma \vdash (E \in \text{Val}) \text{ error}(E) (err : err = E_{\text{err}})} \\
\text{FREE} \\
\frac{Q \triangleq E' \in \text{Val} \star E \mapsto \emptyset}{\Gamma \vdash (E \mapsto E') \text{ free}(E) (ok : Q)} \\
\text{LOOKUP-ERR-VAL} \\
\frac{P \triangleq x = E' \star E \notin \text{Val} \quad E_{\text{err}} \triangleq [\text{"ExprEval"}, \text{str}(E)]}{\Gamma \vdash (P) x := [E] (err : Q_{\text{err}}^*)} \\
\text{LOOKUP-ERR-USE-AFTER-FREE} \\
\frac{P \triangleq x = E' \star E \mapsto \emptyset \quad E_{\text{err}} \triangleq [\text{"UseAfterFree"}, \text{str}(E), E]}{\Gamma \vdash (P) x := [E] (err : Q_{\text{err}}^*)} \\
\text{IF-THEN} \\
\frac{C \triangleq \text{if } (E) C_1 \text{ else } C_2 \quad \Gamma \vdash (P \star E) C_1 (Q)}{\Gamma \vdash (P \star E) C (Q)} \\
\text{IF-ELSE} \\
\frac{C \triangleq \text{if } (E) C_1 \text{ else } C_2 \quad \Gamma \vdash (P \star \neg E) C_2 (Q)}{\Gamma \vdash (P \star \neg E) C (Q)} \\
\text{IF-ERR-VAL} \\
\frac{C \triangleq \text{if } (E) C_1 \text{ else } C_2 \quad E_{\text{err}} \triangleq [\text{"ExprEval"}, \text{str}(E)]}{\Gamma \vdash (P \star E \notin \text{Val}) C (err : Q_{\text{err}}^*)} \\
\text{SEQ} \\
\frac{\Gamma \vdash (P) C_1 (ok : R) (err : Q_{\text{err}}^1) \quad \Gamma \vdash (R) C_2 (ok : Q_{ok}) (err : Q_{\text{err}}^2)}{\Gamma \vdash (P) C_1; C_2 (ok : Q_{ok}) (err : Q_{\text{err}}^1 \vee Q_{\text{err}}^2)} \\
\text{WHILE} \\
\frac{\forall i \in \mathbb{N}. \models P_i \Rightarrow E \in \mathbb{B} \star \text{True} \quad \forall i \in \mathbb{N}. \Gamma \vdash (P_i \star E) C (ok : P_{i+1}) (err : Q_i)}{\Gamma \vdash (P_0) \text{ while } (E) C (ok : \neg E \star \exists i. P_i) (err : \exists i. Q_i)} \\
\text{EQUIV} \\
\frac{\Gamma \vdash (P') C (ok : Q'_{ok}) (err : Q'_{err}) \quad \models P', Q'_{ok}, Q'_{err} \Leftrightarrow P, Q_{ok}, Q_{err}}{\Gamma \vdash (P) C (ok : Q_{ok}) (err : Q_{err})} \\
\text{FRAME} \\
\frac{\text{mod}(C) \cap \text{fv}(R) = \emptyset \quad \Gamma \vdash (P) C (ok : Q_{ok}) (err : Q_{err})}{\Gamma \vdash (P \star R) C (ok : Q_{ok} \star R) (err : Q_{err} \star R)} \\
\text{DISJ} \\
\frac{\Gamma \vdash (P_1) C (ok : Q_{ok}^1) (err : Q_{err}^1) \quad \Gamma \vdash (P_2) C (ok : Q_{ok}^2) (err : Q_{err}^2)}{\Gamma \vdash (P_1 \vee P_2) C (ok : Q_{ok}^1 \vee Q_{ok}^2) (err : Q_{err}^1 \vee Q_{err}^2)} \\
\text{EXISTS} \\
\frac{\Gamma \vdash (P) C (ok : Q_{ok}) (err : Q_{err})}{\Gamma \vdash (\exists x. P) C (ok : \exists x. Q_{ok}) (err : \exists x. Q_{err})}
\end{array}$$

■ **Figure 3** ESL proof rules (excerpt), with $Q_{\text{err}}^* = (\text{pre} \star \text{err} = E_{\text{err}})$.

The structural rules are not surprising, with equivalence replacing the forward/backward consequence of OX/UX reasoning and with frame, existential introduction, and disjunction affecting both post-conditions. Disjunction allows us to derive the standard SL if rule, which captures both branches at the same time. Note that there, however, is no sound conjunction rule, as the conjunction rules of SL and ISL cannot be combined in ESL, since conjunction does not distribute over the star in both directions, breaking frame preservation.

Function Call. We discuss the ESL function-call rule in detail, creating it starting from the standard OX-sound SL rule, adapted for quadruples:

$$\frac{\{\vec{x} = \vec{x} \star P\} f(\vec{x}) \{ok : Q_{ok}\} \{err : Q_{err}\} \in \Gamma \quad y \notin \text{pv}(E_y)}{\Gamma \vdash \{y = E_y \star \vec{E} = \vec{x} \star P\} y := f(\vec{E}) \{ok : Q_{ok}[y/\text{ret}]\} \{err : y = E_y \star Q_{err}\}}$$

In order to make this rule UX-sound, we only have to ensure that no information from the pre-condition is lost in the post-conditions. In particular, we have to remember:

19:14 Exact Separation Logic

- that the evaluation of E_y does not fault, captured by $E_y \in \text{Val}$ and needed in the success post-condition only, as it is already implied by the error post-condition; and
- that $\vec{E} = \vec{x}$ holds (with the substitution $[E_y/y]$ needed in the success post-condition as the value of y may change if the function call succeeds);

bringing us to the ESL function call rule:

$$\frac{(\vec{x} = \vec{x} \star P) \quad f(\vec{x}) \quad (ok : Q_{ok}) \quad (err : Q_{err}) \in \Gamma \quad y \notin \text{pv}(E_y) \quad Q'_{ok} \triangleq E_y \in \text{Val} \star \vec{E}[E_y/y] = \vec{x} \star Q_{ok}[y/\text{ret}] \quad Q'_{err} \triangleq y = E_y \star \vec{E} = \vec{x} \star Q_{err}}{\Gamma \vdash (y = E_y \star \vec{E} = \vec{x} \star P) \quad y := f(\vec{E}) \quad (ok : Q'_{ok}) \quad (err : Q'_{err})}$$

Environment Formation. Whereas the ESL function-call rule does not deviate substantially from its OX counterpart, the environment formation rules illustrate the difference in complexity between OX and UX function compositionality. These rules use the judgement $\vdash (\gamma, \Gamma)$ to state that the environment (γ, Γ) is *well-formed*. The base case, $\vdash (\emptyset, \emptyset)$, is trivial and the same as for SL, stating that the environment consisting of an empty implementation context and an empty specification context is well-formed. For illustrative purposes, we give a simplified version of the extension rule, extending the environment with a single, possibly recursive, function. The full rule, which extends the environment with a group of mutually recursive functions, is given in [24]. We start from the corresponding OX-sound rule from SL:

$$\frac{\text{ENV-EXTEND-SL} \quad \vdash (\gamma, \Gamma) \quad f \notin \text{dom}(\gamma) \quad \gamma' = \gamma[f \mapsto (\vec{x}, C, E)] \quad \Gamma' = \Gamma[f \mapsto \{t\}] \quad \exists t' \in \text{Int}_{\gamma', f}(t). \Gamma' \vdash C : t'}{\vdash (\gamma', \Gamma')}$$

which states that a well-formed environment (γ, Γ) can be extended with a given function f and its external specification t to (γ', Γ') if some corresponding internal specification of f can be proven for the body of f under the extended specification context Γ' .² Note that using Γ' means that a specification can be used to prove itself, which is sound in SL but unsound in ISL: specifically, it would allow us to prove UX-invalid specifications of non-terminating functions. For example, we would be able to prove that the function $\mathbf{f}() \{ r := \mathbf{f}(); \text{return } r \}$ satisfies the EX-valid specification **(emp)** $\mathbf{f}()$ **(False)**, but also the EX-invalid specification **(emp)** $\mathbf{f}()$ **(ret = 42)**. The latter is vacuously OX-valid as there are no terminating executions, but when considered from the UX viewpoint, it implies the existence of an execution path from the pre- to the post-condition, contradicting the non-termination of \mathbf{f} .

Therefore, to be soundly usable in UX reasoning, specifications with satisfiable post-conditions (onward: terminating specifications) must come with a mechanism that disallows the above counter-example. We achieve this by following a standard approach for reasoning about termination [7, 9], based on decreasing measures on well-ordered sets. In particular, we require the terminating specification to be proven, $t \triangleq (P) \quad (ok : Q_{ok}) \quad (err : Q_{err})$ to be extended with a *measure* $\alpha \in \mathbb{N}$, denoting this extension by $t(\alpha)$:³

$$t(\alpha) \triangleq (P \star \alpha = E_\mu) \quad (ok : Q_{ok} \star \alpha = E_\mu) \quad (err : Q_{err} \star \alpha = E_\mu)$$

where E_μ is a logical expression describing how the measure is computed. Then, we prove that $t(\alpha)$ holds for every specific α , assuming that recursive calls to f can only use the terminating

² Internalisation is normally omitted in SL as forward consequence allows information about program variables to be lost, making internal and external post-conditions of SL specifications almost the same.

³ We can extend the measure beyond natural numbers to computable ordinals, $\mathcal{O} \triangleq \omega_1^{\text{CK}}$, allowing us to reason about a broader set of functions, such as those with non-deterministic nested recursion (cf. [24]).

specifications $t(\beta)$ of a measure β *strictly smaller* than α . This restriction is standard and, if the proof succeeds, ensures that f has at least one terminating execution. Also, it disallows the above-mentioned counter-example, as no measure given in the pre-condition would be able to decrease before the recursive call. Importantly, the measure is only a tool required for proving specification validity and once this proof has been completed, the specification without the measure is added to Γ and can be used in proofs of client code.

In addition, we incorporate reasoning about non-terminating specifications (NT-specifications). This is relevant in situations in which the operational semantics of the analysed language is complete, which allows non-termination to be captured using the post-condition **False**. As NT-specifications are vacuously UX-sound, our focus is on ensuring their OX-soundness, which we do by again imposing a measure α , but allowing recursive calls for a measure β *smaller or equal* than α , that is, for an NT-specification to be used to prove itself. This, for example, allows for a proof of the specification **(emp)** $f() \{ r := f(); \text{return } r \}$ **(False)** by choosing a constant measure α . The ESL environment extension rule, therefore, is as follows:

```

ENV-EXTEND
// Assume valid environment, extend implementation context with new function f
⊢ (γ, Γ)    f ∉ dom(γ)    γ' = γ[f ↦ (x̄, C, E)]

// Extend the specifications of f with a measure α
t ≜ (P) (ok : Qok) (err : Qerr)    t∞ ≜ (P∞) (False)    t∞(α) ≜ (P∞ * α = Eμ) (False)
t(α) ≜ (P * α = Eμ) (ok : Qok * α = Eμ) (err : Qerr * α = Eμ)

// Construct Γ(α): assume t for measure β < α and t∞ for measure β ≤ α
Γ(α) = Γ[f ↦ {t(β) | β < α} ∪ {t∞ | β ≤ α}]

// For every α, prove internal specifications of f corresponding to t and t∞
∀α. ∃t' ∈ Intγ',f(t(α)). Γ(α) ⊢ C : t'    ∀α. ∃t' ∈ Intγ',f(t∞(α)). Γ(α) ⊢ C : t'

// Extend Γ with t and t∞
Γ' := Γ[f ↦ {t, t∞}]

```

$$\vdash (\gamma', \Gamma')$$

4.4 Soundness

We state the soundness results for ESL and give intuition about the proofs; the full proofs can be found in [24].

► **Theorem 17.** *Any derivable specification is valid:* $\Gamma \vdash (P) \text{ C } (Q) \implies \Gamma \models (P) \text{ C } (Q)$.

Proof. By induction on $\Gamma \vdash (P) \text{ C } (Q)$. Most cases are straightforward; the [FUN-CALL] rule obtains a valid specification for the function body from the validity of the environment. ◀

► **Theorem 18.** *Any well-formed environment is valid:* $\vdash (\gamma, \Gamma) \implies \models (\gamma, \Gamma)$.

Proof. At the core of the proof is a lemma stating that $\models (\gamma, \Gamma) \implies (\forall \alpha. \models (\gamma', \Gamma(\alpha)))$, where γ' and $\Gamma(\alpha)$ have been obtained from γ and Γ as per [ENV-EXTEND]. Using this lemma, we derive the desired $\models (\gamma', \Gamma')$, where Γ' is obtained from Γ and $\Gamma(\alpha)$ as per [ENV-EXTEND]. The proof of this lemma is done by transfinite induction on α and has the standard zero, successor, and limit ordinal cases. We outline the proof for the case in which a single, possibly recursive, function f with body C_f is added; the generalisation to n mutually recursive functions is straightforward and can be seen in [24].

In all three cases, the soundness of all specifications except the NT-specification with the highest considered ordinal follows from the inductive hypothesis. This remaining NT-specification is vacuously UX-valid, meaning that we only need to prove its OX-validity. For this, we use a form of fixpoint induction called Scott induction (see, e.g., Winskel [31]), required when specifications can be used to prove themselves (e.g. any SL specification).

We set up the Scott induction by extending the set of commands with two pseudo-commands, `scope` and `choice`, with the former modelling the function call but allowing arbitrary commands to be executed in place of the function body, and the latter denoting non-deterministic choice. We then construct the greatest-fixpoint closure of these extended commands, denoted by \mathbb{C} , whose elements may contain infinite applications of the command constructors. We define a behavioural equivalence relation $\simeq_{\gamma'}$ on \mathbb{C} and denote by $\mathbb{C}_{\gamma'}$ the obtained quotient space. This relation induces a partial order $\sqsubseteq_{\gamma'}$, and a join operator that coincides with `choice`, and we show that $(\mathbb{C}_{\gamma'}, \sqsubseteq_{\gamma'})$ is a domain.

We next define S^α as the set of all equivalence classes that hold an element that, for every specification in $(\Gamma(\alpha))(f)$, OX-satisfies at least one of its internal specifications, and show that S^α is an admissible subset of $\mathbb{C}_{\gamma'}$, that is, that it contains the least element of $\mathbb{C}_{\gamma'}$ (represented, for example, by the infinite loop `while (true) { skip }`) and is chain-closed.

We then define the function $h(\mathbb{C}) \triangleq C_f[\mathbb{C}, \gamma', f]$, which replaces all function calls to f in C_f with \mathbb{C} using the `scope` command, and the function g as the lifting of h to $\mathbb{C}_{\gamma'}$: $g([\mathbb{C}]) := [h(\mathbb{C})]$. We next prove that g is continuous (that is, monotonic and supremum-preserving) and that $g(S^\alpha) \subseteq S^\alpha$, from which we can apply the Scott induction principle, together with a well-known identity of the least-fixpoint, which implies that $C_f \in \text{lfp}(g)$, to obtain that $[C_f] \in S^\alpha$. From there, we are finally able to prove that $\models (\gamma', \Gamma(\alpha))$. ◀

These two theorems, to the best of our knowledge, are the first to demonstrate sound function compositionality for UX logics. Previous work on UX logics [25, 28, 29] used function specifications in examples but did not include rules in the logic for calling functions and managing a function specification environment. Program logics that do not include function rules and function specification environments in effect delegate soundness responsibilities to the meta-logic within which they are embedded. This might be appropriate in some contexts, such as in interactive theorem provers, whose meta-logic is reliable. Charguéraud’s clean-slate tutorial SL implementation in Coq [6], for example, does not provide either function call rules or program-logic-level infrastructure for a function specification environment; instead, it relies on Coq’s induction mechanism and definitional mechanism to use and store function specifications. However, when implementing an SL/ISL/ESL-based tool in a mainstream non-ITP language, such as C++ or OCaml, no reliable meta-logic that can act as a safety net is available. This is particularly concerning for UX logics, which require complex rules for handling functions, including forgetting information about function-local mutable variables. In ESL, the handling of functions is fully internalised into the logic, no meta-logic facilities are required to handle function calls, and the program-logic-level facilities of ESL for handling functions are validated by the soundness proof.

The proof of Theorem 18 can be adjusted for ISL: the function call rule would remain the same, and [ENV-EXTEND] would not include NT-specifications, removing the need for Scott induction. On the other hand, the Scott induction itself could be easily adapted for SL.

5 Examples: ESL in Practice

We demonstrate how to use ESL to specify and verify correctness and incorrectness properties of data-structure algorithms, focussing on singly-linked lists and binary search trees.

We investigate, for the first time, the use of abstract predicates in a UX program logic, decoupling abstraction from over-approximation. Our findings show that UX/EX specifications can soundly incorporate abstraction, but also that it, ultimately, cannot be used as freely as in the OX setting. Firstly, since UX reasoning cannot lose information, not all algorithms can be UX-specified at all levels of abstraction, and hence sometimes specifications have to be less abstract than in OX reasoning. Secondly, because specifications are only composable when expressed at the same level of abstraction, specifications of a library client have to be written at the “least common level of abstraction” of the specifications of all of the library functions that the client calls.

Building on §2, we give further intuition on how to think informally about UX/EX specifications using a number of list algorithms and predicates describing lists with various degrees of abstraction (§5.1, §5.2), more detail on how to write formal ESL proofs (§5.3), and examples of ESL reasoning for binary-search-tree algorithms (§5.4).

5.1 List Predicates

We implement singly-linked lists in the standard way: every list consists of two contiguous cells in the heap (denoted $x \mapsto a, b$, meaning $x \mapsto a \star x + 1 \mapsto b$), with the first holding the value of the node, the second holding a pointer to the next node in the list, and the list terminating with a `null` pointer. To capture lists in ESL, we use a number of list predicates:

$$\begin{aligned} \text{list}(x) &\triangleq (x = \text{null}) \vee (\exists v, x'. x \mapsto v, x' \star \text{list}(x')) \\ \text{list}(x, n) &\triangleq (x = \text{null} \star n = 0) \vee (\exists v, x'. x \mapsto v, x' \star \text{list}(x', n - 1)) \\ \text{list}(x, vs) &\triangleq (x = \text{null} \star vs = []) \vee (\exists v, x', vs'. x \mapsto v, x' \star \text{list}(x', vs') \star vs = v : vs') \\ \text{list}(x, xs) &\triangleq (x = \text{null} \star xs = []) \vee (\exists v, x', xs'. x \mapsto v, x' \star \text{list}(x', xs') \star xs = x : xs') \\ \text{list}(x, xs, vs) &\triangleq (x = \text{null} \star xs = [] \star vs = []) \vee \\ &\quad (\exists v, x', xs', vs'. x \mapsto v, x' \star \text{list}(x', xs', vs') \star xs = x : xs' \star vs = v : vs') \end{aligned}$$

These predicates expose different parts of the list structure in their parameters, *hiding* the rest via existential quantification: the $\text{list}(x)$ predicate hides all information about the represented mathematical list, just declaring that there is a singly-linked list at address x ; the $\text{list}(x, n)$ predicate hides the internal node addresses and values, exposing the list length via the parameter n ; the $\text{list}(x, xs)$ predicate hides information about the values of the mathematical list, exposing the internal addresses of the list via the parameter xs ; the $\text{list}(x, vs)$ predicate hides information about the internal addresses, exposing the list’s values via the parameter vs ; and the list predicate $\text{list}(x, xs, vs)$ hides nothing, exposing the entire node-value structure via the parameters xs and vs . These predicates are related to each other via logical equivalence; for example, it holds that:

$$\begin{aligned} \models \text{list}(x) &\Leftrightarrow \exists n. \text{list}(x, n) & \models \text{list}(x) &\Leftrightarrow \exists vs. \text{list}(x, vs) \\ \models \text{list}(x) &\Leftrightarrow \exists xs, vs. \text{list}(x, xs, vs) & \models \text{list}(x, n) &\Leftrightarrow \exists vs. \text{list}(x, vs) \star |vs| = n \end{aligned}$$

5.2 Writing UX/EX Abstract Specifications

We consider a number of list algorithms, described in words, and guide the reader on how to write correct UX/EX specifications for these algorithms using the list abstractions given in the previous section (§5.1), comparing how the UX/EX approach and specifications differ from their OX counterparts. We provide detailed proofs and implementations for each type of algorithm (iterative/recursive, allocating/deallocating, pure/mutative, etc.) in [24].

An important point is to understand how to look for counter-examples to a given specification: from the definition of OX validity, it follows that breaking OX reasoning amounts to “finding a state in the pre-condition (pre-model) for which the execution of f terminates and does not end in a state in the post-condition”; from the definition of UX validity, breaking UX reasoning means “finding a model of the post-condition (post-model) not reachable by execution of f from any state in the pre-condition”; and from the definition of EX validity, it follows that breaking EX reasoning means breaking either OX or UX reasoning. In addition, it is useful to remember that, for breaking UX validity, it is sufficient to find information known in the pre-condition but lost in the post-condition.

Length. We first revisit the list-length function $\text{LLen}(x)$, which takes a list at address x , does not modify it, and returns its length. In §2, we have shown that it satisfies the exact specification $(x = x \star \text{list}(x, n)) \text{LLen}(x) (\text{list}(x, n) \star \text{ret} = n)$ and observed that

(O1) abstraction does not always equal over-approximation.

Using similar reasoning, we can come to the conclusion that the following, less abstract specifications for $\text{LLen}(x)$ are also EX-valid:

$$\begin{aligned} (x = x \star \text{list}(x, vs)) \text{LLen}(x) (\text{list}(x, vs) \star \text{ret} = |vs|) \\ (x = x \star \text{list}(x, xs, vs)) \text{LLen}(x) (\text{list}(x, xs, vs) \star \text{ret} = |xs|) \end{aligned}$$

On the other hand, if we consider the following OX-valid specification:

$$\{x = x \star \text{list}(x)\} \text{LLen}(x) \{\exists n \in \mathbb{N}. \text{list}(x) \star \text{ret} = n\}$$

we see that it is not UX-valid as the post-condition does not connect the return value to the list. In particular, if we choose a post-model for $\text{list}(x)$ that has length 2, but then choose, for example, $\text{ret} = 42$, we run into a problem: as the algorithm does not modify the list, we have to choose the same model of $\text{list}(x)$ for the pre-model to have a chance of reaching the post-model, but then the algorithm will return 2, not 42, meaning that this specification is indeed not UX/EX-valid. From this discussion, we observe that:

(O2) in valid UX/EX specifications, data-structure abstractions used in a post-condition must expose enough information to capture the behaviour of the function being specified with respect to the information given in the pre-condition.

Note that, given a specification less abstract than strictly needed, one can obtain more abstract ones by using validity-preserving transformations on specifications that correspond to the structural rules of the logic. We refer to these as *admissible* transformations, give the ones for existential introduction and equivalence below, and the rest in [24]:

$$\text{ADM-EXISTS} \frac{\Gamma \models (\vec{x} = \vec{x} \star P) f(\vec{x}) (ok : Q_{ok}) (err : Q_{err}) \quad y \notin \vec{x}}{\Gamma \models (\vec{x} = \vec{x} \star \exists y. P) f(\vec{x}) (ok : \exists y. Q_{ok}) (err : err : \exists y. Q_{err})}$$

$$\text{ADM-EQUIV} \frac{\Gamma \models (\vec{x} = \vec{x} \star P') f(\vec{x}) (ok : Q'_{ok}) (err : Q'_{err}) \quad \models P', Q'_{ok}, Q'_{err} \Leftrightarrow P, Q_{ok}, Q_{err}}{\Gamma \models (\vec{x} = \vec{x} \star P) f(\vec{x}) (ok : Q_{ok}) (err : Q_{err})}$$

For list-length, starting from the least abstract specification using $\text{list}(x, xs, vs)$, we can derive, for example, the specification using $\text{list}(x, n)$, as follows:

$$\begin{aligned} & (x = x \star \text{list}(x, xs, vs)) \text{LLen}(x) (\text{list}(x, xs, vs) \star \text{ret} = |xs|) \\ [\text{ADM-EXISTS}] & (\exists vs. x = x \star \text{list}(x, xs, vs)) \text{LLen}(x) (\exists vs. \text{list}(x, xs, vs) \star \text{ret} = |xs|) \\ [\text{ADM-EXISTS}] & (\exists xs, vs. x = x \star \text{list}(x, xs, vs)) \text{LLen}(x) (\exists xs, vs. \text{list}(x, xs, vs) \star \text{ret} = |xs|) \\ [\text{ADM-EQUIV}] & (x = x \star \text{list}(x, n)) \text{LLen}(x) (\text{list}(x, n) \star \text{ret} = n) \end{aligned}$$

Interestingly, from one more application of [ADM-EXISTS] and [ADM-EQUIV], we can derive

$$(x = x \star \text{list}(x)) \text{ LLen}(x) (\exists n. \text{list}(x, n) \star \text{ret} = n)$$

which further illustrates observation (O2), in that even though the pre-condition does not talk about the length of the list, the post-condition has to expose it because the function output depends on it, and hence the post-condition must connect up the return value to the length of the list, here by an existentially quantified variable.

This approach of deriving abstract specifications can be used in general for working with ESL: for a given algorithm, first prove the least abstract specification, which exposes all details, and then adjust the degree of abstraction to fit the needs of the client code. We discuss this further in the upcoming paragraph on reasoning about client programs.

Membership. Next, we consider the list-membership function $\text{LMem}(x, v)$, which takes a list at address x , does not modify it, and returns `true` if v is in the list, and `false` otherwise. Given (O2) and the fact that the function output depends on the values in the list, we understand that, for its UX/EX specification, we should be using a list abstraction that exposes at least the values, that is, $\text{list}(x, vs)$ or $\text{list}(x, xs, vs)$. The corresponding specifications are:

$$\begin{aligned} (x = x \star v = v \star \text{list}(x, vs)) \text{ LMem}(x, v) (\text{list}(x, vs) \star \text{ret} = (v \in vs)) \\ (x = x \star v = v \star \text{list}(x, xs, vs)) \text{ LMem}(x, v) (\text{list}(x, xs, vs) \star \text{ret} = (v \in vs)) \end{aligned}$$

and are proven similarly to list-length. We can check that a more abstract specification, say:

$$\{x = x \star v = v \star \text{list}(x)\} \text{ LMem}(x, v) \{\exists b \in \mathbb{B}. \text{list}(x) \star \text{ret} = b\}$$

is not UX-valid, by choosing, as the post-model, b to be `false` and the list at x to contain v . As for list-length, since list-membership does not modify the list, we have to choose the same model of $\text{list}(x)$ for the pre-model to have a chance of reaching the post-model, but then the algorithm will return `true`, not `false`, so this post-model is not reachable.

Swap-First-Two. Next, we consider the list-swap-first-two function $\text{LSwapFirstTwo}(x)$, which takes a list at address x , swaps its first two values if the list is of sufficient length, returning `null`, and throws an error otherwise without modifying the list. Given (O2), to specify this function we need an abstraction that captures list length and, apparently, also the list values; for example, $\text{list}(x, vs)$. As this function can throw errors, its full EX specification has to use the ESL quadruple, in which the two post-conditions are constrained with the corresponding, shaded, branching conditions:

$$\begin{aligned} (x = x \star \text{list}(x, vs)) \\ \text{LSwapFirstTwo}(x, v) \\ (ok : \exists v_1, v_2, vs'. \text{vs} = v_1 : v_2 : vs' \star \text{list}(x, v_2 : v_1 : vs') \star \text{ret} = \text{null}) \\ (err : \text{list}(x, vs) \star |vs| < 2 \star err = \text{“List too short!”}) \end{aligned}$$

observing that the success post-condition, given the used abstraction, has to not only state that the length of the list is not less than two, but also how the values are manipulated, and also that the error message is chosen for illustrative purposes.

However, note that the swapped values *are not featured in the function output*, but instead remain contained within the predicate. This indicates that a more abstract specification:

$$\begin{aligned} (x = x \star \text{list}(x, n)) \\ \text{LSwapFirstTwo}(x, v) \\ (ok : \text{list}(x, n) \star n \geq 2 \star \text{ret} = \text{null}) (err : \text{list}(x, n) \star n < 2 \star err = \text{“List too short!”}) \end{aligned}$$

which only reveals the list length, might be EX-valid, and indeed it is. Any list we choose in the error post-model will have length less than two, and can then be used in the pre-model to reach the post-model. On the other hand, whichever list we choose in the success post-model will have length at least two, that is, its values will be of the form $v_1 : v_2 : vs$ and it will have some addresses, and then we can choose a list with the same addresses and values $v_2 : v_1 : vs$ in the pre-model and we will reach the post-model by executing the function.

Pointer-Reverse. Let us now examine the list-pointer-reverse function, $\text{LPRev}(x)$, which takes a list at address x and reverses it by reversing the direction of the next-pointers, returning the head of the reversed list. Given (O2) and the fact that the algorithm manipulates pointers and returns an address, but the actual values in the list are not exposed, we will try to use the address-only $\text{list}(x, xs)$ predicate to specify this function as in the following OX triple, where xs^\dagger denotes the reverse of the mathematical list xs :

$$\{x = x \star \text{list}(x, xs)\} \text{LPRev}(x) \{\text{list}(\text{ret}, xs^\dagger)\}$$

which would seem to be UX-valid given our OX experience and previous examples, but is not. In particular, it has no information about the logical variable x , which exists only in the pre-condition. This is not an issue in OX reasoning, but in UX reasoning it would mean that there exists a logical environment that interprets the post-condition but not the pre-condition, and such a specification, by the definition, could never be UX-valid.

To understand which specific information about x is required, we first add the general $x \in \text{Val}$, making the post-condition $\text{list}(\text{ret}, xs^\dagger) \star x \in \text{Val}$, and then try to choose a post-model by picking values for ret , xs , and x . Note that, given the definition of $\text{list}(x, xs)$, we cannot just pick any non-correlated values for ret and xs : in particular, either xs is an empty list and ret is `null`, or xs is non-empty and ret is its last element. This observation, in fact, reveals the information needed about x : either x is `null` and xs is empty, or xs is non-empty and x is its first element. We capture this information using the $\text{listHead}(x, xs)$ predicate:

$$\text{listHead}(x, xs) \triangleq (xs = [] \star x = \text{null}) \vee (\exists xs'. xs = x : xs')$$

and arrive at the desired EX specification of the list-pointer-reverse algorithm:

$$(x = x \star \text{list}(x, xs)) \text{LPRev}(x) (\text{list}(\text{ret}, xs^\dagger) \star \text{listHead}(x, xs))$$

Let us make sure that this specification is UX-valid. If we pick a post-model with $xs = []$, then $x = \text{ret} = \text{null}$ and the pre-model with the same x and xs will work, as the list holds no values. For a post-model with non-empty xs , x must equal the head of xs , ret must equal the tail of xs , and we also have to pick some arbitrary values vs , with $|vs| = |xs|$. Then, given the described behaviour of the algorithm, we know that this post-model is reachable from a pre-model which has the list at x with addresses xs and values vs^\dagger .

Free. Next, we take a look at the $\text{LFree}(x)$ function, which frees a given list at address x . Its OX specification is $\{x = x \star \text{list}(x)\} \text{LFree}(x) \{\text{ret} = \text{null}\}$, but it does not transfer to UX contexts because no resource from the pre-condition can be forgotten in the post-condition as that would break the UX frame property [28]. Instead, we have to keep track of the addresses to be freed, which we can do using the $\text{list}(x, xs)$ predicate (or $\text{list}(x, xs, vs)$), and we also have to explicitly state in the post-condition that these addresses have been freed:

$$(x = x \star \text{list}(x, xs)) \text{LFree}(x) (\text{freed}(xs) \star \text{listHead}(x, xs) \star \text{ret} = \text{null})$$

using the $\text{freed}(xs)$ predicate, which is defined as follows:

$$\text{freed}(xs) \triangleq (xs = []) \vee (\exists x, xs'. xs = x : xs' \star x \mapsto \emptyset, \emptyset \star \text{freed}(xs'))$$

Client Programs and Specification Composition. We discuss the usability of ESL specifications in general and abstraction in particular in the context of client programs that call multiple library functions. Consider the following (slightly contrived) client program, which takes a list and: pointer-reverses it if its length is between 5 and 10; frees it and then throws an error if its length is smaller than 5; and does not terminate otherwise:

```

LClient(x) {
  l := LLen(x);
  if (l < 5)
    { r := LFree(x); error("List too short!") } else
    { if (l > 10) { while (true) { skip } } else { r := LPRev(l) } };
  return r
}

```

Our first goal is to understand which is the most abstract list predicate that could be used for reasoning about this client, since we want to minimise the amount of details we need to carry along in the proof, noting that the least abstract one, $\text{list}(x, xs, vs)$, will always work. Observe that, importantly, only specifications expressed at the same abstraction level are composable with each other, because they must be composed using equivalence. We explore this in more detail in the subsequent formal discussion (see, in particular, observation (O5)).

When it comes to $\text{LClient}(x)$, for list-length, we need information about the list length, meaning that we can use either $\text{list}(x, n)$, $\text{list}(x, xs)$, or $\text{list}(x, vs)$, but not $\text{list}(x)$. For list-free, we must have information about the addresses, meaning that $\text{list}(x, n)$ and $\text{list}(x, vs)$ will not work, leaving us with $\text{list}(x, xs)$, which is also usable for list-pointer-reverse. Therefore, we can write the specification of this client using the $\text{list}(x, xs)$ predicate, as follows:

$$\begin{aligned}
& (x = x \star \text{list}(x, xs)) \\
& \quad \text{LClient}(x) \\
& \quad (ok : 5 \leq |xs| \leq 10 \star \text{list}(\text{ret}, xs^\dagger) \star \text{listHead}(x, xs)) \\
& \quad (err : |xs| < 5 \star \text{freed}(xs) \star \text{listHead}(x, xs) \star err = \text{"List too short!"})
\end{aligned}$$

In general, however, it is sufficient for a client to call one function that works with addresses and another that works with values for the only applicable predicate to be $\text{list}(x, xs, vs)$, which is still abstract in the sense that it allows for unbounded reasoning about lists, but does not hide any of its internal information. This leads us to the following observation:

(O3) specifications that use predicates which hide data-structure information, albeit provable, may have limited use in UX client reasoning.

As a final remark on abstraction, note that we have only considered predicates that expose the data-structure sub-parts (for lists, these sub-parts are values vs and addresses xs) either entirely or not at all. It would be also possible to expose *some* of this structure for some of the algorithms, but because of (O3), specifications using such abstractions are only composable with specifications exposing the same partial structure, and hence likely to be of limited use.

Non-termination. We conclude our discussion on specifications with two remarks on EX reasoning about non-terminating behaviour. First, consider the non-terminating branch of the LClient function, which is triggered when $|xs| > 10$. Observe that this branch is implicit in the client specification, in that it is subsumed by the success post-condition (since $\models P \vee (|xs| > 10 \star \text{False}) \Leftrightarrow P$). However, to demonstrate that it exists, we can constrain the pre-condition appropriately to prove the specification $(x = x \star \text{list}(x, xs) \star |xs| > 10) \text{LClient}(x) (\text{False})$. This implicit loss of non-terminating branches can be characterised informally as follows:

(O4) if the post-conditions do not cover all paths allowed by the pre-condition, then the “gap” is non-terminating.

In this case, the pre-condition implies $|xs| \in \mathbb{N}$ and the post-conditions cover the cases where $|xs| \leq 10$, leaving the gap when $|xs| > 10$, for which we provably have client non-termination.

Second, we observe that, in contrast to terminating behaviour, for non-terminating behaviour EX is as expressive as OX; that is, the EX triple $(P) C (\text{False})$ is equivalent to the OX triple $\{P\} C \{\text{False}\}$ as the UX triple $[P] C [\text{False}]$ is vacuously true. This is not to say that all non-terminating behaviour can be captured by ESL specifications. For example, as in OX, if the code branches on a value that does not come from the pre-condition, and if one of the resulting branches does not terminate, and if the code can also terminate successfully, then the non-terminating branch will be implicit in the pre-condition, but no gap in the sense of (O4) will be present. This is illustrated by the code and specification below, where the pre- and the post-condition are the same, but a non-terminating path still exists:

```
(x = 0) x := nondet; if (x > 42) { while (true) { skip } } else { x := 0 } (x = 0)
```

5.3 More ESL Proofs: Iterative list-length

In §2, we have shown a proof sketch for a recursive implementation of the list-length algorithm, demonstrating how to handle the measure for recursive function calls; how the folding of predicates works in the presence of equivalence; and how to move between external and internal specifications. We highlight again the UX-specific issue that we raised and that is related to predicate folding, which can be formulated generally as follows:

(O5) if the code accesses data-structure information that the used predicate hides, then that predicate might not be foldable in a UX-proof in all of the places in which it would be foldable in the corresponding OX-proof.

Here, we show how to write ESL proofs for looping code, using as example an iterative implementation of the list-length algorithm. Proofs for the majority of the other algorithms mentioned in §5.2 can be found in [24]; the rest are similar.

Iterative list-length in ESL: Proof Sketch. In Figure 4, we give an iterative implementation of the list-length algorithm and show that it satisfies the same ESL specification as its recursive counterpart, $(x = x \star \text{list}(x, n)) \text{LLen}(x) (ok : \text{list}(x, n) \star \text{ret} = n)$. Since there is no recursion, we elide the (trivial) measure. To state the loop variant, we use the list-segment predicate, defined as follows:

$$\text{lseg}(x, y, n) \triangleq (x = y \star n = 0) \vee (\exists v, x'. x \mapsto v, x' \star \text{lseg}(x', y, n - 1))$$

and to apply the [WHILE] rule, we define:

$$P_i \triangleq \exists j. \text{lseg}(x, x, i) \star \text{list}(x, j) \star n = i + j \star r = i$$

Note that we could have chosen to elide i from the body of P_i in this simple example, but since this is not necessarily possible or evident in general as well as for instructive purposes, we chose to keep it in the proof. Note how, on exiting the loop, the negation of the loop condition collapses the existentials i and j . This allows us to obtain the given internal post-condition, from which we then easily move to the desired external post-condition. For this proof, we also use three equivalence lemmas, which state that a non-empty list segment can be separated into its last element and the rest, that the length of an empty list equals zero, and that a null-terminated list-segment is a list.

```

 $\Gamma \vdash (x = x \star \text{list}(x, n))$ 
  LLen(x) {
    (  $x = x \star \text{list}(x, n) \star r = \text{null}$  )
    r := 0
    (  $x = x \star \text{list}(x, n) \star r = 0$  )
    (  $P_0$  )
    while (x  $\neq$  null) {
      (  $P_i \star x \neq \text{null}$  )
      (  $\exists j, v, x'. \text{lseg}(x, x, i) \star x \mapsto v, x' \star \text{list}(x', j - 1) \star n = i + j \star r = i$  )
      x := [x + 1];
      (  $\exists j, v, x'. \text{lseg}(x, x', i) \star x' \mapsto v, x \star \text{list}(x, j) \star n = i + (j + 1) \star r = i$  )
      // equivalence:  $\models \text{lseg}(x, y, n + 1) \Leftrightarrow \exists x', v. \text{lseg}(x, x', n) \star x' \mapsto v, y]$ 
      (  $\exists j. \text{lseg}(x, x, i + 1) \star \text{list}(x, j) \star n = (i + 1) + j \star r = i$  )
      r := r + 1
      (  $P_{i+1}$  )
    }
    (  $x = \text{null} \star \exists i. P_i$  )
    (  $\exists i, j. \text{lseg}(x, x, i) \star \text{list}(x, j) \star n = i + j \star r = i \star x = \text{null}$  )
    (  $\text{lseg}(x, \text{null}, n) \star r = n \star x = \text{null}$  ) // equivalence:  $\models \text{list}(\text{null}, j) \Leftrightarrow j = 0$ 
    (  $\text{list}(x, n) \star r = n \star x = \text{null}$  ) // equivalence:  $\models \text{lseg}(x, \text{null}, n) \Leftrightarrow \text{list}(x, n)$ 
    return r
    (  $\text{list}(x, n) \star r = n \star x = \text{null} \star \text{ret} = r$  )
    (  $\exists x_q, r_q. \text{list}(x, n) \star r_q = n \star x_q = \text{null} \star \text{ret} = r_q$  )
    (  $\text{list}(x, n) \star \text{ret} = n$  )
  }
  (  $\text{list}(x, n) \star \text{ret} = n$  )

```

■ **Figure 4** ESL proof sketch: iterative list-length.

5.4 Beyond List Examples: Binary Search Trees

While list algorithms illustrate many aspects of exact reasoning, it is also important to understand how ESL specification and verification works with other data structures. For this reason, we discuss two algorithms operating over binary search trees (BSTs) that are intended to represent sets of natural numbers. We use two abstractions for BSTs, one in which only their values are considered as a mathematical set:

$$\text{BST}(x, K) \triangleq (x = \text{null} \star K = \emptyset) \vee (\exists k, l, r, K_l, K_r. x \mapsto k, l, r \star \text{BST}(r, K_r) \star \text{BST}(l, K_l) \star K = K_l \uplus \{k\} \uplus K_r \star K_l < k \star k < K_r)$$

and another that fully exposes the BST structure:

$$\text{BST}(x, \tau) \triangleq (x = \text{null} \star \tau = \tau_\emptyset) \vee (\exists k, l, r, \tau_l, \tau_r. E \mapsto k, l, r \star \text{BST}(r, \tau_r) \star \text{BST}(l, \tau_l) \star \tau = ((x, k), \tau_l, \tau_r) \star \tau_l < k \star k < \tau_r)$$

where τ is a mathematical tree, that is, an algebraic data type with two constructors representing, respectively, an empty tree and a root node with two child trees: $\tau \in \text{Tree} \triangleq \tau_\emptyset \mid ((x, n), \tau_l, \tau_r)$, where the notation (x, n) represents a BST node with address x and value n . Note the overloaded $<$ notation, where one of the operands can be a set or a tree, which carry the intuitive meaning.

BST algorithms. We first consider the BST-find-minimum algorithm, $\text{BSTFindMin}(x)$, which takes a tree with root at x , does not modify it, and returns its minimum element or throws an empty-tree error. Since that algorithm operates only on the values in the tree, we are able to state its ESL specification using the $\text{BST}(x, K)$ predicate as follows:

$$(x = x \star \text{BST}(x, K)) \text{BSTFindMin}(x) \begin{array}{l} (\text{ok}: x \neq \text{null} \star \text{BST}(x, K) \star \text{ret} = \min(K)) \\ (\text{err}: x = \text{null} \star \text{BST}(x, K) \star \text{err} = \text{“Empty tree!”}) \end{array}$$

We have also considered the BST-insert algorithm, $\text{BSTInsert}(x, v)$, which takes a tree with root at x and inserts a new node with value v into it as a leaf if v is not already in the tree, or leaves the tree unmodified if it is. As this algorithm interacts both with values and addresses in the tree, the appropriate abstraction for it is $\text{BST}(x, \tau)$, and its ESL specification is:

$$\begin{array}{c} (x = x \star v = v \star \text{BST}(x, \tau)) \\ \text{BSTInsert}(x, v) \\ (\exists x'. \text{BST}(\text{ret}, \text{BSTInsert}(\tau, (x', v))) \star \text{BSTRoot}(x, \tau)) \end{array}$$

where $\text{BSTInsert}(\tau, \nu)$ is the mathematical algorithm that inserts the node ν into the tree τ :

$$\text{BSTInsert}(\tau_\emptyset, (x', v)) \triangleq \text{BSTInsert}(((x, k), \tau_l, \tau_r), (x', v)) \triangleq \begin{array}{l} \text{if } v < k \text{ then } ((x, k), \text{BSTInsert}(\tau_l, (x', v)), \tau_r) \\ \text{else if } k < v \text{ then } ((x, k), \tau_l, \text{BSTInsert}(\tau_r, (x', v))) \\ \text{else } ((x, k), \tau_l, \tau_r) \end{array}$$

and the predicate $\text{BSTRoot}(x, \tau)$ is defined analogously to $\text{listHead}(x, xs)$:

$$\text{BSTRoot}(x, \tau) \triangleq (\tau = \tau_\emptyset \star x = \text{null}) \vee (\exists k, \tau_l, \tau_r. \tau = ((x, k), \tau_l, \tau_r))$$

This example shows how in EX verification, just as in OX verification, we end up relating an imperative heap-manipulating algorithm to its mathematical/functional counterpart (cf. Appel [3] for a recent reiteration of this idea). The additional work required is that EX mathematical models must be more detailed: we are, yet again, not allowed to lose information. In particular, in OX verification we could relate $\text{BSTInsert}(x, v)$ to mathematical sets, but in EX verification we must relate our imperative implementation to tree models, including both values and pointers. Moreover, our mathematical model of the algorithm, $\text{BSTInsert}(\tau, (x', v))$, must insert elements in the same way as the imperative implementation, that is, in this case at the leaves of the tree. The proofs for both algorithms are given in [24].

6 Related Work

In the previous sections, we have placed ESL in the context of related work on OX and UX logics and associated tools. Here, we discuss formalisms capable of reasoning both about program correctness and program incorrectness, as well as existing approaches to the use of function specifications (summaries) and abstraction in symbolic execution.

Program Logics for Both Correctness and Incorrectness. Developed in parallel but independently of ESL, Outcome Logic (OL) [33], much like ESL, brings together reasoning about correctness and incorrectness into one logic. Both OL and ESL rely on the traditional meaning of correctness, but OL introduces a new approach to incorrectness, based on reachability of sets of states. It has not yet been shown that this approach has the same bug-finding potential as that of ISL: in particular, bi-abduction has not yet been demonstrated to be compatible with OL. In addition, the OL work, in contrast to ESL, does not discuss function compositionality or the interaction between abstraction, reachability, and incorrectness.

LCL_A [4,5] is a non-function-compositional, first-order logic that combines UX and OX reasoning using abstract interpretation. It is parametric on an abstract domain A , and proves UX triples of the form $\vdash_A [P] C [Q]$ where, under certain conditions, the triple also guarantees verification. These conditions, however, normally mean that only a limited number of pre-conditions can be handled. The conditions also have to be checked per-command and if they fail to hold (due to, e.g., issues with Boolean guards, which are known to be a major source of incompleteness), then the abstract domain has to be incrementally adjusted; the complexity of this adjustment and the expressivity of the resulting formalism is unclear.

Compositional Symbolic Execution. There exists a substantial body of work on symbolic execution with function summaries (e.g. [1, 15–17, 22, 32]), which is primarily based on first-order logic. We highlight the work of Godefroid et al., which initially used exact summaries of bounded program behaviour to drive the compositional dynamic test generation of SMART [15], and later distinguished between may (OX) and must (UX) summaries, leveraging the interaction between them to design the SMASH tool for compositional property checking and test generation [16]. SMASH, however, is limited in its ability to reason about heap-manipulating programs because, for example, it lacks support for pointer arithmetic. Nevertheless, it shows that interactions between OX and UX summaries can be exploited for automation, which is an important consideration for any automation of ESL. For example, SMASH is able to use not-may summaries (which amount to non-reachability) when constructing must-summaries (which amount to reachability), using the former to restrict the latter. When it comes to abstraction, for example, Anand et al. [2] implement linked-list and array abstractions for true bug-finding in non-compositional symbolic execution, in the context of the Java PathFinder, and use it to find bugs in list and array partitioning algorithms. True bug-finding is maintained by checking for state subsumption, which requires code modification rather than annotation and a record of all previously visited states.

7 Conclusions

We have introduced ESL, a program logic for exact reasoning about heap-manipulating programs. ESL specifications provide a sweet spot between verification and true bug-finding: as SL specifications, they capture all terminating behaviour, and, as ISL specifications, they describe only results and errors that are reachable. ESL specifications are therefore compatible with tools that target OX verification, such as VeriFast [19] and Iris [20], tools that target UX true bug-finding, such as Pulse [26,28], and tools capable of targeting both, such as Gillian [10,23]. ESL supports reasoning about mutually recursive functions and comes with a soundness result that immediately transfers to SL and ISL, thus demonstrating, for the first time, scalable functional compositionality for UX logics.

We have verified exact specifications for a number of illustrative examples, showing that ESL can reason about data-structure libraries, language errors, mutual recursion, and non-termination. In doing so, we emphasise the distinction between the often-conflated concepts of abstraction and over-approximation. We have demonstrated that abstract predicates can be soundly used in EX and UX reasoning, albeit not as freely as in OX reasoning.

We believe that ESL reasoning, in its intended context of semi-automatic verification of functional correctness properties, is useful for the verification of self-contained, critical code that underpins a larger codebase. To demonstrate this, we will in the future incorporate UX and EX verification inside Gillian [10,23], which already has support for function compositionality and semi-automatic predicate management as part of its OX verification.

References

- 1 Saswat Anand, Patrice Godefroid, and Nikolai Tillmann. Demand-driven compositional symbolic execution. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 2008. doi:10.1007/978-3-540-78800-3_28.
- 2 Saswat Anand, Corina S. Pasareanu, and Willem Visser. Symbolic execution with abstraction. *International Journal on Software Tools for Technology Transfer*, 11(1), 2009. doi:10.1007/s10009-008-0090-1.
- 3 Andrew W. Appel. Coq’s vibrant ecosystem for verification engineering. In *Conference on Certified Programs and Proofs (CPP)*, 2022. doi:10.1145/3497775.3503951.
- 4 Roberto Bruni, Roberto Giacobazzi, Roberta Gori, and Francesco Ranzato. A logic for locally complete abstract interpretations. In *Symposium on Logic in Computer Science (LICS)*, 2021. doi:10.1109/LICS52264.2021.9470608.
- 5 Roberto Bruni, Roberto Giacobazzi, Roberta Gori, and Francesco Ranzato. A correctness and incorrectness program logic. *Journal of the ACM*, 70(2), 2023. doi:10.1145/3582267.
- 6 Arthur Charguéraud. Separation logic for sequential programs (functional pearl). *Proceedings of the ACM on Programming Languages*, 4(ICFP), 2020. doi:10.1145/3408998.
- 7 Pedro da Rocha Pinto, Thomas Dinsdale-Young, Philippa Gardner, and Julian Sutherland. Modular termination verification for non-blocking concurrency. In *European Symposium on Programming (ESOP)*, 2016. doi:10.1007/978-3-662-49498-1_8.
- 8 Edsko de Vries and Vasileios Koutavas. Reverse Hoare logic. In *Software Engineering and Formal Methods (SEFM)*, 2011. doi:10.1007/978-3-642-24690-6_12.
- 9 Robert W. Floyd. Assigning meanings to programs. *Proceedings of Symposium on Applied Mathematics*, 19, 1967.
- 10 José Fragoso Santos, Petar Maksimović, Sacha-Élie Ayoun, and Philippa Gardner. Gillian, part I: A multi-language platform for symbolic execution. In *Programming Language Design and Implementation (PLDI)*, 2020. doi:10.1145/3385412.3386014.
- 11 José Fragoso Santos, Petar Maksimović, Théotime Grohens, Julian Dolby, and Philippa Gardner. Symbolic execution for JavaScript. In *Principles and Practice of Declarative Programming (PPDP)*, 2018. doi:10.1145/3236950.3236956.
- 12 José Fragoso Santos, Petar Maksimović, Daiva Naudžiūnienė, Thomas Wood, and Philippa Gardner. JaVerT: JavaScript verification toolchain. *Proceedings of the ACM on Programming Languages*, 2(POPL), 2018. doi:10.1145/3158138.
- 13 José Fragoso Santos, Petar Maksimović, Gabriela Sampaio, and Philippa Gardner. JaVerT 2.0: Compositional symbolic execution for JavaScript. *Proceedings of the ACM on Programming Languages*, 3(POPL), 2019. doi:10.1145/3290379.
- 14 Philippa Gardner, Sergio Maffeis, and Gareth David Smith. Towards a program logic for JavaScript. In *Principles of Programming Languages (POPL)*, 2012. doi:10.1145/2103656.2103663.
- 15 Patrice Godefroid. Compositional dynamic test generation. In *Principles of Programming Languages (POPL)*, 2007. doi:10.1145/1190216.1190226.
- 16 Patrice Godefroid, Aditya V. Nori, Sriram K. Rajamani, and Sai Deep Tetali. Compositional may-must program analysis: Unleashing the power of alternation. In *Principles of Programming Languages (POPL)*, 2010. doi:10.1145/1706299.1706307.
- 17 Benjamin Hillery, Eric Mercer, Neha Rungta, and Suzette Person. Exact heap summaries for symbolic execution. In *Verification, Model Checking, and Abstract Interpretation (VMCAI)*, 2016. doi:10.1007/978-3-662-49122-5_10.
- 18 C. A. R. Hoare. An Axiomatic Basis for Computer Programming. *Communications of the ACM (CACM)*, 12(10), 1969. doi:10.1145/363235.363259.
- 19 Bart Jacobs, Jan Smans, Pieter Philippaerts, Frédéric Vogels, Willem Penninckx, and Frank Piessens. VeriFast: A powerful, sound, predictable, fast verifier for C and Java. In *NASA Formal Methods Symposium (NFM)*, 2011. doi:10.1007/978-3-642-20398-5_4.

- 20 Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *Principles of Programming Languages (POPL)*, 2015. doi:10.1145/2676726.2676980.
- 21 Quang Loc Le, Azalea Raad, Jules Villard, Josh Berdine, Derek Dreyer, and Peter W. O'Hearn. Finding real bugs in big programs with incorrectness logic. *Proceedings of the ACM on Programming Languages*, 6(OOPSLA1), 2022. doi:10.1145/3527325.
- 22 Yude Lin, Tim Miller, and Harald Sondergaard. Compositional symbolic execution using fine-grained summaries. In *Australasian Software Engineering Conference*, 2015. doi:10.1109/ASWEC.2015.32.
- 23 Petar Maksimović, Sacha-Élie Ayoun, José Fragoso Santos, and Philippa Gardner. Gillian, part II: Real-world verification for JavaScript and C. In *Computer Aided Verification (CAV)*, 2021. doi:10.1007/978-3-030-81688-9_38.
- 24 Petar Maksimović, Caroline Cronjäger, Andreas Löow, Julian Sutherland, and Philippa Gardner. Exact separation logic (extended version), 2023. arXiv:2208.07200.
- 25 Toby Murray, Pengbo Yan, and Gidon Ernst. Incremental vulnerability detection with insecurity separation logic, 2021. arXiv:2107.05225.
- 26 Peter W. O'Hearn. Incorrectness logic. *Proceedings of the ACM on Programming Languages*, 4(POPL), 2019. doi:10.1145/3371078.
- 27 Peter W. O'Hearn, John C. Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *Computer Science Logic*, 2001. doi:10.1007/3-540-44802-0_1.
- 28 Azalea Raad, Josh Berdine, Hoang-Hai Dang, Derek Dreyer, Peter O'Hearn, and Jules Villard. Local reasoning about the presence of bugs: Incorrectness separation logic. In *Computer Aided Verification (CAV)*, 2020. doi:10.1007/978-3-030-53291-8_14.
- 29 Azalea Raad, Josh Berdine, Derek Dreyer, and Peter W. O'Hearn. Concurrent incorrectness separation logic. *Proceedings of the ACM on Programming Languages*, 6(POPL), 2022. doi:10.1145/3498695.
- 30 John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Logic in Computer Science (LICS)*, 2002. doi:10.1109/LICS.2002.1029817.
- 31 Glynn Winskel. *The Formal Semantics of Programming Languages: An Introduction, Chapter 10*. MIT Press, 1993.
- 32 Greta Yorsh, Eran Yahav, and Satish Chandra. Generating precise and concise procedure summaries. In *Principles of Programming Languages (POPL)*, 2008. doi:10.1145/1328438.1328467.
- 33 Noam Zilberstein, Derek Dreyer, and Alexandra Silva. Outcome logic: A unifying foundation for correctness and incorrectness reasoning. *Proceedings of the ACM on Programming Languages*, 7(OOPSLA1), 2023. doi:10.1145/3586045.