

Symbolic Execution Game Semantics

Yu-Yang Lin

Queen Mary University of London, UK

Nikos Tzevelekos

Queen Mary University of London, UK

Abstract

We present a framework for symbolically executing and model checking higher-order programs with external (open) methods. We focus on the client-library paradigm and in particular we aim to check libraries with respect to any definable client. We combine traditional symbolic execution techniques with operational game semantics to build a symbolic execution semantics that captures arbitrary external behaviour. We prove the symbolic semantics to be sound and complete. This yields a bounded technique by imposing bounds on the depth of recursion and callbacks. We provide an implementation of our technique in the \mathbb{K} framework and showcase its performance on a custom benchmark based on higher-order coding errors such as reentrancy bugs.

2012 ACM Subject Classification Theory of computation \rightarrow Semantics and reasoning

Keywords and phrases game semantics, symbolic execution, higher-order open programs

Digital Object Identifier 10.4230/LIPIcs.FSCD.2020.28

Funding Research funded by EPSRC (EP/P004172/1).

Acknowledgements We would like to thank members of the K framework for their consistent support with K, and anonymous reviewers for their insightful reviews.

1 Introduction

Two important challenges in program verification are state-space explosion and the environment problem. The former refers to the need to investigate infeasibly many states, while the latter concerns cases where the code depends on an environment that is not available for analysis. State-space explosion has been approached with a range of techniques, which have led to verification tools being nowadays routinely used on industrial-scale code (e.g. [10, 5, 7]). The environment problem, however, remains largely unanswered: verification techniques often require the whole code to be present for the analysis and, in particular, cannot analyse components like libraries where parts of the code are missing (e.g. the client using the library). This problem is particularly acute in higher-order programs, where the interaction between a program and its environment can be intricate and e.g. involve callbacks or reentrant calls. In this paper we address this latter problem by combining *game semantics*, a semantics theory for higher-order programs, with *symbolic execution*, a technique that uses *symbolic values* to explore multiple execution paths of a program.

To showcase the importance and challenges of the environment problem, following is a

```
1 import send:(int  $\rightarrow$  unit)
2 int balance := 100;
3
4 public withdraw (m:int) :(unit) =
5   if (not (!balance < m)) then
6     send(m);
7     balance := !balance - m;
8     assert(not(!balance < 0))
9   else ();
```

simple example of a library written in a sugared version of HOLi, the vehicle language of this paper. The example is a simplified implementation of “The DAO” smart contract, a failed decentralised autonomous organisation on the Ethereum blockchain platform [12]. As with libraries, the challenge in analysing smart contracts is that the client code is not available. We must



© Yu-Yang Lin, Nikos Tzevelekos;
licensed under Creative Commons License CC-BY

5th International Conference on Formal Structures for Computation and Deduction (FSCD 2020).

Editor: Zena M. Ariola; Article No. 28; pp. 28:1–28:24

Leibniz International Proceedings in Informatics



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

45 thus generate all possible contexts in which the contract can be called. In this case, the error
 46 is caused by a reentrant call from the `send()` method, which is provided by the environment.
 47 When this method is called, the environment takes control and is allowed to call any method
 48 in the library. If a client were to call `withdraw()` within its `send()` method, the recursive
 49 call would drain all the funds available, which is simulated in this example by a negative
 50 balance. This happens because the method is manipulating a global state, and is updating
 51 it after the external call. We can see that an analysis capturing this error would need to
 52 be able to predict an intricate environment behaviour. Moreover, such an analysis should
 53 ideally only predict realisable environment behaviours.

54 Symbolic execution [34, 13, 19] explores all paths of a program using symbolic values
 55 instead of concrete input values. Each symbolic path holds a path condition (a SAT formula)
 56 that is satisfiable if and only if the path can be concretely executed. While the resulting
 57 analysis is unbounded in general, by restricting our focus to bounded paths we can soundly
 58 catch errors, or affirm the absence thereof up to the used bound. Game semantics [2, 14],
 59 on the other hand, models higher-order program phrases in isolation as 2-player games:
 60 sequences of computational *moves* (method calls and returns) between the program and
 61 its hypothetical environment. The power of the technique lies in its use of combinatorial
 62 conditions to precisely allow those game plays that can be realised by including the program
 63 in an actual environment. Moreover, the theory can be formulated operationally in terms
 64 of a trace semantics for open terms [18, 21, 16] which, in turn, lends itself to a symbolic
 65 representation. The latter yields a symbolic execution technique that is *sound and complete*
 66 in the following sense: given an open program, its symbolic traces match its concrete traces,
 67 which match its realisable traces in some environment.

68 Returning to the DAO example, we can model the ensuing interaction as a sequence of
 69 moves, alternating between the environment and the library. Any finite sequence of moves
 70 (that leads to an assertion violation) is a trace defining a counterexample. Running the
 71 example in HOLiK, our implementation of the symbolic semantics in the \mathbb{K} Framework [33],
 72 the following minimal symbolic trace is automatically found:

```
74   call⟨withdraw, x1⟩ · call⟨send, x1⟩ · call⟨withdraw, x2⟩
75     · call⟨send, x2⟩ · ret⟨send, ()⟩ · ret⟨withdraw, ()⟩ · ret⟨send, ()⟩
```

76 where x_1 is the original call parameter, and x_2 is the parameter for the reentrant call,
 77 satisfiable with values $x_1 = 100$ and $x_2 = 1$. A fix would be to swap line 6 and 7, to update
 80 internal state before passing control.

81 In Appendix A we look at a few more examples of libraries that exhibit errors due to
 82 high-order behaviours. We provide three examples: a file lock example, a double deallocation
 83 example, and an unsafe implementation of flat-combining.

84 Overall, this paper contributes a novel symbolic execution technique based on game
 85 semantics to precisely model the behaviour of higher-order stateful programs. Specifically:
 86 – We present a symbolic trace semantics for higher-order libraries that captures the behaviour
 87 of an unknown environment, and prove it sound and complete: i.e. it produces no spurious
 88 error traces, and is able to produce the complete execution tree of any library. – By
 89 bounding the depth of nested calls and the *insistence* of the environment in calling library
 90 methods, we derive a sound and bounded-complete technique to check higher-order libraries
 91 for errors. – We implement the latter in the \mathbb{K} semantical framework [33] to produce a
 92 sound and bounded-complete tool for higher-order libraries as a proof of concept. We test
 93 our implementation with benchmarks adapted from the literature. Some material has been
 94 delegated to an Appendix. Full proofs can be found in an extended version of the paper [22].

<i>Libraries</i> $L ::= B \mid \mathbf{abstract} \ m; L$	<i>Terms</i> $M ::= m \mid i \mid () \mid x \mid \lambda x.M \mid r := M \mid !r$
<i>Blocks</i> $B ::= \varepsilon \mid \mathbf{public} \ m = \lambda x.M; B$	$\mid M \oplus M \mid \langle M, M \rangle \mid \pi_1 M \mid \pi_2 M$
$\mid m = \lambda x.M; B \mid \mathbf{global} \ r := i; B$	$\mid MM \mid \mathbf{if} \ M \ \mathbf{then} \ M \ \mathbf{else} \ M$
$\mid \mathbf{global} \ r := \lambda x.M; B$	$\mid \mathbf{letrec} \ x = \lambda x.M \ \mathbf{in} \ M$
<i>Clients</i> $C ::= L; \mathbf{main} = M$	$\mid \mathbf{let} \ x = M \ \mathbf{in} \ M \mid \mathbf{assert}(M)$

$() : \mathbf{unit}$	$i : \mathbf{int}$	$\frac{x \in \mathbf{Vars}_\theta}{x : \theta}$	$\frac{m \in \mathbf{Meths}_{\theta, \theta'}}{m : \theta \rightarrow \theta'}$	$\frac{M, M' : \mathbf{int}}{M \oplus M' : \mathbf{int}}$	$\frac{M : \mathbf{int} \quad M_1, M_0 : \theta}{\mathbf{if} \ M \ \mathbf{then} \ M_1 \ \mathbf{else} \ M_0 : \theta}$
$\frac{M : \theta_1 \quad M' : \theta_2}{\langle M, M' \rangle : \theta_1 \times \theta_2}$	$\frac{\langle M, M' \rangle : \theta_1 \times \theta_2}{\pi_i \langle M, M' \rangle : \theta_i}$	$\frac{r \in \mathbf{Refs}_\theta}{!r : \theta}$	$\frac{r \in \mathbf{Refs}_\theta}{r := M : \mathbf{unit}}$	$\frac{M : \theta \quad M' : \theta \rightarrow \theta' \quad M : \theta}{M' M : \theta'}$	
$\frac{M : \theta' \quad x : \theta}{\lambda x.M : \theta \rightarrow \theta'}$	$\frac{x, M : \theta \quad M' : \theta'}{\mathbf{let} \ x = M \ \mathbf{in} \ M' : \theta'}$	$\frac{x, \lambda y.M : \theta \rightarrow \theta'' \quad M' : \theta'}{\mathbf{letrec} \ x = \lambda y.M \ \mathbf{in} \ M' : \theta'}$	$\frac{M : \mathbf{int}}{\mathbf{assert}(M) : \mathbf{unit}}$		

■ **Figure 1** Syntax and typing rules of HOLi.

2 A Language for Higher-Order Libraries: HOLi

We introduce HOLi, a language for higher-order libraries which define methods to be used by an external client, and in turn require external methods (provided by the client). We give in HOLi an operational semantics for terms that integrates a counter for the depth of nested calls that a program phrase can make. We then extend this counting semantics to open terms by means of a trace semantics. We show that the trace semantics of libraries is sound and complete for reachability of errors under any external client.

2.1 Syntax and operational semantics

A library in HOLi is a collection of typed higher-order methods. A client is simply a library with a main body. Types are given by the grammar:

$$\theta ::= \mathbf{unit} \mid \mathbf{int} \mid \theta \times \theta \mid \theta \rightarrow \theta$$

We use countably infinite sets \mathbf{Meths} , \mathbf{Refs} and \mathbf{Vars} for method, global reference and variable names, ranged over by m , r and x respectively, and variants thereof; while i is for ranging over the integers. We use \oplus to range over a set of binary integer operations, which we leave unspecified. Each set of names is typed, that is, it can be expressed as a disjoint union as follows: $\mathbf{Meths} = \bigsqcup_{\theta, \theta'} \mathbf{Meths}_{\theta, \theta'}$, $\mathbf{Refs} = \bigsqcup_{\theta \neq \theta_1 \times \theta_2} \mathbf{Refs}_\theta$, $\mathbf{Vars} = \bigsqcup_\theta \mathbf{Vars}_\theta$.

The full syntax and typing rules are given in Figure 1. Thus, a library consists of abstract method declarations, followed by blocks of public and private method and reference definitions. A method is considered private unless it is declared **public**. Each public/private method and reference is defined once. Abstract methods are not given definitions: these methods are external to the library. Public, private and abstract methods are all disjoint.

Libraries are well typed if all their method and reference definitions are well typed (e.g. **public** $m = \lambda x.M$ is well typed if $m : \theta$ and $\lambda x.M : \theta$ are both valid for the same type θ) and only mention methods and references that are defined or abstract. A client $L; \mathbf{main} = M$ is well typed if $M : \mathbf{unit}$ is valid and $L; m = \lambda x.M$ is well typed for some fresh x, m . A library/client is *open* if it contains abstract methods. This is different to open/closed terms: we call a term *open* if it contains free variables.

► **Remark 1.** By typing variable, reference and method names, we do not need to provide a context in typing judgements. Note that the references we use are of non-product type and,

28:4 Symbolic Execution Game Semantics

$$\begin{array}{l}
(E[\mathbf{let} \ x = v \ \mathbf{in} \ M], R, S, k) \rightarrow (E[M\{v/x\}], R, S, k) \qquad (E[\pi_j \langle v_1, v_2 \rangle], R, S, k) \rightarrow (E[v_j], R, S, k) \\
(E[r := v], R, S, k) \rightarrow (E[()], R, S[r \mapsto v], k) \qquad (E[!r], R, S, k) \rightarrow (E[S(r)], R, S, k) \\
(E[\mathbf{if} \ i \ \mathbf{then} \ M_1 \ \mathbf{else} \ M_0], R, S, k) \rightarrow (E[M_j], R, S, k) \quad (1) \quad (E[i_1 \oplus i_2], R, S, k) \rightarrow (E[i], R, S, k) \quad (2) \\
(E[\lambda x.M], R, S, k) \rightarrow (E[m], R \uplus \{m \mapsto \lambda x.M\}, S, k) \qquad (E[\mathbf{assert}(i)], R, S, k) \rightarrow (E[()], R, S, k) \quad (3) \\
(E[mv], R, S, k) \rightarrow (E[(M\{v/x\})], R, S, k + 1) \quad (4) \qquad (E[(\!v)], R, S, k + 1) \rightarrow (E[v], R, S, k) \\
(E[\mathbf{letrec} \ f = \lambda x.M \ \mathbf{in} \ M'], R, S, k) \rightarrow (E[M'\{m/f\}], R \uplus \{m \mapsto \lambda x.M\{m/f\}\}, S, k) \\
\text{Conditions: } (1) : j = 1 \text{ iff } i \neq 0, \quad (2) : i = i_1 \oplus i_2, \quad (3) : i \neq 0, \quad (4) : R(m) = \lambda x.M.
\end{array}$$

$$\begin{array}{l}
\text{Values } v ::= m \mid i \mid () \mid \langle v, v \rangle \qquad \text{Terms (extended) } M ::= \dots \mid (\!M) \\
\text{Eval. Contexts } E ::= \bullet \mid \mathbf{assert}(E) \mid r := E \mid E \oplus M \mid v \oplus E \mid \langle E, M \rangle \mid \langle v, E \rangle \mid \pi_j E \\
EM \mid mE \mid \mathbf{let} \ x = E \ \mathbf{in} \ M \mid \mathbf{if} \ E \ \mathbf{then} \ M \ \mathbf{else} \ M \mid (E)
\end{array}$$

$$\begin{array}{l}
(\mathbf{abstract} \ m; L, R, S, \mathcal{P}, \mathcal{A}) \xrightarrow{bid} (L, R, S, \mathcal{P}, \mathcal{A} \uplus \{m\}) \\
(\mathbf{public} \ m = \lambda x.M; B, R, S, \mathcal{P}, \mathcal{A}) \xrightarrow{bid} (B, R \uplus \{m \mapsto \lambda x.M\}, S, \mathcal{P} \uplus \{m\}, \mathcal{A}) \\
(m = \lambda x.M; B, R, S, \mathcal{P}, \mathcal{A}) \xrightarrow{bid} (B, R \uplus \{m \mapsto \lambda x.M\}, S, \mathcal{P}, \mathcal{A}) \\
(\mathbf{global} \ r := i; B, R, S, \mathcal{P}, \mathcal{A}) \xrightarrow{bid} (B, R, S \uplus \{r \mapsto i\}, \mathcal{P}, \mathcal{A}) \\
(\mathbf{global} \ r := \lambda x.M; B, R, S, \mathcal{P}, \mathcal{A}) \xrightarrow{bid} (B, R \uplus \{m \mapsto \lambda x.M\}, S \uplus \{r \mapsto m\}, \mathcal{P}, \mathcal{A})
\end{array}$$

■ **Figure 2** Operational semantics (top); values and evaluation contexts (mid); library build (bottom).

124 more importantly, **global** to the library: a term can use references but not create them locally
125 or pass them as arguments (we discuss how to include such references in Appendix C).

126 ► **Example 2.** The DAO-attack example from the Introduction can be written in HOLi as:

```

127   abstract send; global bal := 100;
128   public wdraw =
129     λx. if !bal ≥ x then (send(x); bal := !bal - x; assert(!bal ≥ 0)) else ()
130

```

131 where $send, wdraw \in \mathbf{Meths}_{\text{int,unit}}$, $bal \in \mathbf{Refs}_{\text{int}}$, and $M; M'$ stands for $\mathbf{let} \ _ = M \ \mathbf{in} \ M'$.

132 A library contains public methods that can be called by a client. On the other hand,
133 a client contains a main body that can be executed. These two scenarios constitute the
134 operational semantics of HOLi. Both are based on evaluating (closed) terms, which we
135 define next. Term evaluation requires: the closed term being evaluated; method definitions,
136 provided by a method repository; reference values, provided by a store; and a call-depth
137 counter (a natural number). Since method application is the only source of infinite behaviour
138 in HOLi, bounding the depth of nested calls is enough to guarantee termination in program
139 analysis. Hence we provide a mechanism to keep track of call depth.

140 The operational semantics is given in Figure 2. The evaluation of terms (top part) involves
141 configurations of the form (M, R, S, k) , where:

- 142 ■ M is a closed term which may contain *evaluation boxes*, i.e. points inside a term where
143 a method call has been made and has not yet returned, and is taken from the syntax
144 extending the one of Figure 1 with the rule: $M ::= \dots \mid (\!M)$
- 145 ■ R is a *method repository*, i.e. a partial map from method names to their bodies
- 146 ■ S is a *store*, i.e. a partial map from reference names to their stored values
- 147 ■ k is a *counter*, i.e. a natural number.

148 Most of the rules are standard, but it is worth noting that lambdas are not values themselves
 149 but, rather, evaluate to method names that are freshly stored in the repository. Moreover,
 150 evaluation boxes interplay with the counter k in the semantics: they mark places where the
 151 depth has increased because of a nested call. The penultimate line of rules in the operational
 152 semantics keeps track of call depth, and illustrates the utility of evaluation boxes: making
 153 a call increases the counter and leaves behind an evaluation box; returning from the call
 154 removes the box and decreases the counter again.

155 A library L *builds* into a configuration of the form $(\varepsilon, R, S, \mathcal{P}, \mathcal{A})$, which includes its
 156 public methods according to the rules in Figure 2 (bottom). More precisely, R and S are as
 157 above, while $\mathcal{P}, \mathcal{A} \subseteq \mathbf{Meths}$ are (disjoint) sets of *public* and *abstract* method names. We say
 158 that (a well typed) L builds to $(\varepsilon, R, S, \mathcal{P}, \mathcal{A})$ if $(L, \emptyset, \emptyset, \emptyset) \xrightarrow{bld}^* (\varepsilon, R, S, \mathcal{P}, \mathcal{A})$. If L builds
 159 to $(\varepsilon, R, S, \mathcal{P}, \mathcal{A})$ then the client $L; \mathbf{main} = M$ builds to $(M, R, S, \mathcal{P}, \mathcal{A})$. Moreover, we can
 160 link libraries to clients and evaluate them, as in the following definition.

- 161 ► **Definition 3. 1.** Library L and client C are compatible if L builds to some $(\varepsilon, R, S, \mathcal{P}, \mathcal{A})$
 162 and C builds to some $(M, R', S', \mathcal{P}', \mathcal{A}')$ such that: $\mathcal{P} \supseteq \mathcal{A}'$ and $\mathcal{A} \supseteq \mathcal{P}'$ (complementation);
 163 $\text{dom}(S) \cap \text{dom}(S') = \emptyset$ (disjoint state); and $\text{dom}(R) \cap \text{dom}(R') = \emptyset$ (method ownership).
 164 2. For a library L , we let \hat{L} be L with all its abstract method declarations and **public**
 165 keywords removed; and similarly for \hat{C} . Given compatible library L and client C , we let
 166 their composition be the client: $L;C = \hat{L};\hat{C}$.
 167 3. Given compatible L, C , the semantics of $L;C$ is:

$$168 \quad \llbracket L;C \rrbracket = \{\rho \mid L;C \text{ builds to } (M, R, S, \emptyset, \emptyset) \wedge (M, R, S, 0) \rightarrow^* \rho\}$$

169 We say that $\llbracket L;C \rrbracket$ fails if it contains some $(E[\mathbf{assert}(0)], \dots)$.

170 ► **Example 4.** To illustrate how libraries and clients are used, consider the DAO example
 171 again as a library L_{DAO} . We can define a client C_{atk} :

```
172 abstract wdraw; global wlet := 0;
173 public send = λx.wlet := !wlet + x; if !wlet < 100 then wdraw(x) else ();
174 main = wdraw(1)
```

176 to produce the following linked client $L_{\text{DAO}};C_{\text{atk}}$ (modulo re-ordering):

```
177 global bal := 100; global wlet := 0;
178 wdraw = λx. if !bal ≥ x then (send(x); bal := !bal - x; assert(!bal > 0)) else ();
179 public send = λx.wlet := !wlet + x; if !wlet < 100 then wdraw(x) else ();
180 main = wdraw(1)
```

182 We can see how L_{DAO} is vulnerable to an attacker such as C_{atk} after linking them. The aim is
 183 thus to use bounded analysis to find counterexamples that define clients such as this one.

184 2.2 Trace Semantics

185 The semantics we defined only allows us to evaluate terms, and only so long as their method
 186 applications only involve methods that can be found in the repository R . We next extend
 187 this semantics to encompass libraries and terms that can also call abstract methods. The
 188 approach we follow is based on operational game semantics [18, 21, 16] and in particular the
 189 semantics is given by means of traces of method calls and returns (called *moves* in game

$$\begin{array}{l}
(\text{INT}) \quad \frac{(M, R, S, k) \rightarrow (M', R', S', k')}{(\mathcal{E}, M, R, S, \mathcal{P}, \mathcal{A}, k)_p \rightarrow (\mathcal{E}, M', R', S', \mathcal{P}, \mathcal{A}, k')_p} \\
(\text{PQ}) \quad (\mathcal{E}, E[mv], R, S, \mathcal{P}, \mathcal{A}, k)_p \xrightarrow{\text{call}(m,v)} ((m, E) :: \mathcal{E}, 0, R, S, \mathcal{P}', \mathcal{A}, k)_o \\
(\text{OQ}) \quad (\mathcal{E}, l, R, S, \mathcal{P}, \mathcal{A}, k)_o \xrightarrow{\text{call}(m,v)} ((m, l+1) :: \mathcal{E}, mv, R, S, \mathcal{P}, \mathcal{A}', k)_p \text{ if } R(m) = \lambda x.M \\
(\text{PA}) \quad ((m, l) :: \mathcal{E}, v, R, S, \mathcal{P}, \mathcal{A}, k)_p \xrightarrow{\text{ret}(m,v)} (\mathcal{E}, l, R, S, \mathcal{P}', \mathcal{A}, k)_o \\
(\text{OA}) \quad ((m, E) :: \mathcal{E}, l, R, S, \mathcal{P}, \mathcal{A}, k)_o \xrightarrow{\text{ret}(m,v)} (\mathcal{E}, E[v], R, \mathcal{P}, \mathcal{A}', k)_p \\
\hline
(\text{PC}) : m \in \mathcal{A} \wedge \mathcal{P}' = \mathcal{P} \cup (\text{Meths}(v) \cap \text{dom}(R)), \quad (\text{OC}) : m \in \mathcal{P} \wedge \mathcal{A}' = \mathcal{A} \cup (\text{Meths}(v) \setminus \text{dom}(R)).
\end{array}$$

■ **Figure 3** Trace semantics rules. Rules (PQ), (PA) assume the condition (PC), and similarly for (OQ), (OA) and (OC). $\text{Meths}(v)$ contains all method names appearing in v . INT stands for *internal* transition; PQ for *P-question* (i.e. call); PA for *P-answer* (i.e. return). Similarly for OQ and OA.

190 semantics jargon), between the library and its client. In between such moves, the semantics
 191 evolves as the operational semantics we already saw.

192 To maintain a terminating analysis, we need to keep track of an added source of infinite
 193 execution, namely endless consecutive calls from an external component: a library will never
 194 terminate if its client keeps calling its methods. This leads us to a semantics with two
 195 counters, k and l , where k keeps track of internal nested method calls and l records the
 196 number of consecutive calls made from the external component. This counter l is orthogonal
 197 to k and is refreshed at every call to the external context.

198 When computing the semantics of a library, the library and its methods are the *Player* (P)
 199 of the computation game, while the (intended) client is the *Opponent* (O). As the semantics
 200 is given in absence of an actual client, O actually represents every possible client. When
 201 computing the semantics of a client, the roles are reversed. In both cases, though, the same
 202 sets of rules is used and there is no need to specify who is P and O in the semantics.

203 The trace semantics uses *game configurations*, which are divided into *P-configurations*
 204 and *O-configurations* given respectively as:

$$205 \quad (\mathcal{E}, M, R, S, \mathcal{P}, \mathcal{A}, k)_p \quad \text{and} \quad (\mathcal{E}, l, R, S, \mathcal{P}, \mathcal{A}, k)_o.$$

206 In a P -configuration, a term M is being evaluated – this is P 's role. In an O -configuration,
 207 an external call has been made and the semantics waits for O to either return that call, or
 208 reply itself with another call. The components $M, R, S, \mathcal{P}, \mathcal{A}, k, l$ are as above, while \mathcal{E} is an
 209 *evaluation stack*:

$$210 \quad \mathcal{E} ::= \varepsilon \mid (m, E) :: \mathcal{E} \mid (m, l) :: \mathcal{E}$$

211 which keeps track of the computations that are on hold due to external calls. The trace
 212 semantics is generated by the rules given in Figure 3.

213 The formulation follows closely the operational game semantics technique. For example,
 214 from a P -configuration $(\mathcal{E}, M, R, S, \mathcal{P}, \mathcal{A}, k)_p$, there are 3 options:

- 215 1. If M can make an internal reduction, i.e. in the operational semantics in context (R, S, k) ,
 216 then $(\mathcal{E}, M, R, S, \mathcal{P}, \mathcal{A}, k)_p$ performs this reduction (via (INT)).
- 217 2. If M is stuck at a method application for a method that is not in the repository R , then
 218 that method must be abstract (i.e. external) and needs to be called externally. This is
 219 achieved by issuing a call move and moving to an O -configuration (via (PQ)). The current
 220 evaluation context and the called method name are stored, in order to resume once the
 221 call is returned (via (OA)).

222 3. If M is a value and the evaluation stack is non-empty, then P has completed a method
223 call that was issued by O (via (OQ)) and can now return (via (PA)).

224 On the other hand, from an O -configuration $(\mathcal{E}, l, R, S, \mathcal{P}, \mathcal{A}, k)_o$, there are 2 options:

- 225 1. either return the last open method call (made by P) via (OA), or
- 226 2. call one of the public methods (from \mathcal{P}) using (OQ).

227 The role of conditions (PC) and (OC) is to ensure that each player calls the methods
228 owned by the other, or returns their own, and update the sets of public and abstract names
229 according to the method names passed inside v .

230 ► **Remark 5.** The novelty of Figure 3 with respect to previous work on trace semantics for
231 open libraries (e.g. [26]) lies in the use of l in order to bound the ability of O to ask repeated
232 questions for finite analysis. The way rules (OQ) and (PA) are designed is such that any
233 sequence of consecutive O -calls and P -returns has maximum length $2n$ if we bound l to n
234 (i.e. $l \leq n$), as each such pair of moves increases l by 1. On the other hand, each P -call
235 supplies to O a fresh counter ($l = 0$) to be used in contiguous (OQ)-(PA)'s. Thus, l can be
236 seen as keeping track of the *insistence* of O in calling.

237 Finally, we can define the trace semantics of libraries.

238 ► **Definition 6.** Let L be a library. The semantics of L is :

$$239 \quad \llbracket L \rrbracket = \{(\tau, \rho) \mid (L, \emptyset, \emptyset, \emptyset) \xrightarrow{bld}^* (\varepsilon, R, S, \mathcal{P}, \mathcal{A}) \wedge (\varepsilon, 0, R, S, \mathcal{P}, \mathcal{A}, 0)_o \xrightarrow{\tau} \rho\}$$

241 We say that $\llbracket L \rrbracket$ fails if it contains some $(\tau, (\mathcal{E}, E[\mathbf{assert}(0)], \dots))$.

242 ► **Example 7.** Consider the DAO example as library L_{DAO} once again. Evaluating the game
243 semantics we know the following sequence is in $\llbracket L_{\text{DAO}} \rrbracket$. For economy, we hide $R, \mathcal{P}, \mathcal{A}$ and
244 show only the top of the stack in the configurations. We also use $m(v)?$ and $m(v)!$ for calls
245 and returns. We write S_i for the store $[bal \mapsto i]$.

$$246 \quad (\varepsilon, 0, S_{100}, 0)_o \xrightarrow{wdraw(42)?} ((wdraw, 1), wdraw(42), S_{100}, 0)_p$$

$$247 \quad \rightarrow^* ((wdraw, 1), E[send(42)], S_{100}, 1)_p \xrightarrow{send(42)?} ((send, E), 2, S_{100}, 1)_o$$

$$248 \quad \xrightarrow{wdraw(100)?} ((wdraw, 1), wdraw(100), S_{100}, 1)_p$$

$$249 \quad \rightarrow^* ((wdraw, 1), E'[send(100)], S_{100}, 2)_p \xrightarrow{send(100)?} ((send, E), 2, S_{100}, 2)_o$$

$$250 \quad \xrightarrow{send()!} ((wdraw, 1), E'[], S_{100}, 2)_p \rightarrow^* ((wdraw, 1), (), S_0, 2)_p$$

$$251 \quad \xrightarrow{wdraw()!} ((send, E), 1, S_0, 2)_o \xrightarrow{send()!} ((wdraw, 1), E[], S_0, 1)_p$$

$$252 \quad \rightarrow^* ((wdraw, 1), E[\mathbf{assert}(-42 \geq 0)], S_{-42}, 1)_p$$

254 This transition sequence is an instance of the symbolic trace provided in the Introduction.
255 Here, a call is made with parameter 42, and a reentrant call with 100, which leads to the
256 assertion violation $\mathbf{assert}(-42 \geq 0)$. Note that a bound of $k \leq 2$ is sufficient to find this
257 assertion violation.

258 We next establish two focal properties of the trace semantics: bounding k and l ensures
259 termination (Theorem 8), and that it is sound and complete with respect to library errors
260 (Theorem 9). Notice Theorem 9 captures both soundness and completeness as it states that
261 the game semantics eventually reaches every error that is concretely reachable for any client
262 while finding only errors that can be reached concretely by a definable client.

263 ► **Theorem 8** (Boundedness). *For any game configuration ρ , provided an upper bound k_0*
 264 *and l_0 for call counters k and l , the labelled transition system starting from ρ is strongly*
 265 *normalising.*

266 **Proof.** For any transition sequence $\rho = \rho_0 \rightarrow \dots \rightarrow \rho_i \rightarrow \dots$ and each $i > 0$, we set the
 267 following two classes of configurations:

$$268 \quad (A) = \{\rho_i \mid |\rho_i| < |\rho_{i-1}|\} \quad (B) = \{\rho_i \mid \exists j < i - 1. |\rho_i| < |\rho_j|\}$$

269 where $|\rho| = (k_0 - k, |M|, l_0 - l)$ is the *size* of ρ , and $|\rho| < |\rho'|$ is defined by the lexicographic
 270 ordering of the triple $(k_0 - k, |M|, l_0 - l)$, with bounds k_0 and l_0 such that $k \leq k_0$ and $l \leq l_0$
 271 for semantic transitions to be applicable. If not present in the configuration, we look at
 272 the evaluation stack \mathcal{E} to find the top-most missing component. In other words, opponent
 273 configurations will have size $(k_0 - k, |E|, l_0 - l)$ where E is the top-most one in \mathcal{E} , whereas
 274 proponent configurations will have size $(k_0 - k, |M|, l_0 - l)$ where l is the top-most one in \mathcal{E} .

275 We approach the proof in two steps: (1) we show that, for any transition sequence out of
 276 ρ , each reachable configuration belongs to (at least) one of the above classes; and (2) prove
 277 that the classes form a terminating sequence. For (1), considering all moves available to ρ ,
 278 we have the following cases.

279 1. If $\rho \rightarrow \rho'$ is an (INT) move, we have two possibilities.

280 a. For a transition $(E[\llbracket v \rrbracket], R, S, k) \rightarrow (E[v], R, S, k+1)$, where $k+1 \leq k_0$, we have a class
 281 (B) configuration since there must be a $(E[mv], R, S, k)$ such that $(E[mv], R, S, k) \rightarrow^*$
 282 $(E[v], R, S, k)$ which is lexicographically ordered since $|v| < |mv|$.

283 b. Every other transition sequence is class (A) since they reduce the size of the term.

284 2. If $\rho \rightarrow \rho'$ is a (PQ) move, we have that ρ' is a class (A) configuration since $(k, |E|, l_0) <$
 285 $(k, |E[mv]|, l_0 - l)$ by lexicographic ordering.

286 3. If $\rho \rightarrow \rho'$ is an (OA) move, we have a transition

$$287 \quad ((m, E) :: \mathcal{E}, l, \dots, k)_o \xrightarrow{ret(m,v)} (\mathcal{E}, E[v], \dots, k)_p$$

288 which must be a result of the prior proponent question, meaning \mathcal{E} holds an l' on top.
 289 We thus have the following sequence

$$290 \quad (\mathcal{E}, E[mv], \dots, k)_p \rightarrow^* (\mathcal{E}, E[v], \dots, k)_o$$

291 where $(k, |E[v]|, l) < (k, |E[mv]|, l')$, so ρ' is a class (B) configuration.

292 4. If $\rho \rightarrow \rho'$ is an (OQ) move, we have the transition

$$293 \quad (\mathcal{E}, l, \dots, k)_o \xrightarrow{call(m,v)} ((m, l+1) :: \mathcal{E}, mv, \dots, k)_p$$

$$294 \quad \rightarrow ((m, l+1) :: \mathcal{E}, \llbracket M\{v/x\} \rrbracket, \dots, k+1)$$

296 Simplifying the transition, we remove the configuration in between and take

$$297 \quad (\mathcal{E}, l, R, S, \mathcal{P}, \mathcal{A}, k)_o \xrightarrow{call(m,v)} ((m, l+1) :: \mathcal{E}, \llbracket M\{v/x\} \rrbracket, R, S, \mathcal{P}, \mathcal{A}, k+1)_p$$

298 to be our new equivalent transition. We thus have that ρ' is a class (A) configuration since
 299 $(k_0 - (k+1), |\llbracket M\{v/x\} \rrbracket|, l_0 - (l+1)) < (k_0 - k, |E|, l_0 - l)$ by lexicographic ordering.

300 5. If $\rho \rightarrow \rho'$ is a (PA) move, we have the transition

$$301 \quad ((m, l) :: \mathcal{E}, v, \dots, k)_p \xrightarrow{ret(m,v)} (\mathcal{E}, l, \dots, k)_o$$

302 which must be the result of a prior opponent question

$$303 \quad (\mathcal{E}, l + 1, \dots, k)_o \xrightarrow{call(m,v)} ((m, l) :: \mathcal{E}, \langle M\{v/x\} \rangle, \dots, k + 1)_p$$

$$304 \quad \rightarrow^* ((m, l) :: \mathcal{E}, \langle v \rangle, \dots, k + 1)_p \rightarrow ((m, l) :: \mathcal{E}, v, \dots, k)_p \xrightarrow{ret(m,v)} (\mathcal{E}, l, \dots, k)_o$$

306 where E' is the topmost evaluation context in \mathcal{E} . We thus have that $(k_0 - k, E', l_0 - l) <$
 307 $(k_0 - k, E', l_0 - (l + 1))$, so ρ' is a class (B) configuration.

308 For (2), let us assume there is an infinite sequence

$$309 \quad \rho_0 \rightarrow \dots \rightarrow \rho_j \rightarrow \dots \rightarrow \rho_i \rightarrow \dots$$

310 Since all reachable configurations fall into either (A) or (B) class, we know that the sequence
 311 must comprise only (A) and (B) configurations. In this infinite sequence, we know that all
 312 sequences of (A) configurations are in descending size, so (A) sequences cannot be infinite.
 313 We also observe that (B) configurations are padded with (A) sequences. For instance, if
 314 ρ_i is a (B) configuration, and ρ_j is its matching configuration, there may exist nested (B)
 315 configurations between ρ_j and ρ_i , as well as (A) sequences padding these.

316 Additionally, these (B) configurations can only occur as a return to a call, so we know
 317 they only occur together with the introduction of evaluation boxes ($\langle \bullet \rangle$). Since these brackets
 318 occur in pairs and are introduced in a nested fashion, we know \mathcal{E} can only contain evaluation
 319 contexts with well-bracketed evaluation boxes, meaning that there cannot be interleaved
 320 sequences of (B) configurations where their target configurations intersect. More specifically,
 321 the sequence

$$322 \quad \rho_0 \rightarrow \dots \rightarrow \rho_j \rightarrow \dots \rightarrow \rho'_j \rightarrow \dots \rightarrow \rho_i \rightarrow \dots \rightarrow \rho'_i \rightarrow \dots$$

323 where ρ'_i matches ρ'_j and ρ_i matches ρ_j is not possible.

324 Now, ignoring all (A) and nested (B) sequences, we are left with an infinite stream of
 325 top-level (B) sequences which are also in descending order. Since starting size is finite, we
 326 cannot have an infinite stream of (B) sequences. Thus, the assumption that the sequence is
 327 infinite does not hold, meaning our semantics is strongly normalising. \blacktriangleleft

328 **► Theorem 9 (S and C).** *We call a client good if it contains no assertions. For any library*
 329 *L , the following are equivalent:*

- 330 1. $\llbracket L \rrbracket$ fails (reaches an assertion violation)
- 331 2. there exists a good client C such that $\llbracket L; C \rrbracket$ fails

332 **Proof.** 1 to 2: Suppose now that $(\tau, \rho) \in \llbracket L \rrbracket$ for some trace τ and failed ρ . By Theorem 11,
 333 we have that there is a good client C realising the trace τ . So then, by Lemma 10, we have
 334 that $\llbracket L; C \rrbracket$ fails.

335 2 to 1: Suppose $\llbracket L; C \rrbracket$ fails for some good client C . Then, by Lemma 10, there are τ, ρ, ρ'
 336 such that $(\tau, \rho) \in \llbracket L \rrbracket$, $(\tau, \rho') \in \llbracket C \rrbracket$, and ρ is failed (i.e. is of the shape $(\mathcal{E}, E[\text{assert}(0)], \dots)$).
 337 \blacktriangleleft

338 The latter relies on an auxiliary lemma (well-composing of libraries and clients, see [22]),
 339 and a definability result akin to game semantics definability arguments (see Appendix D).

340 ▶ **Lemma 10** (L-C Compositionality). *For any library L and compatible good client C , $\llbracket L;C \rrbracket$*
 341 *fails if and only if there exist $(\tau_1, \rho_1) \in \llbracket L \rrbracket$ and $(\tau_2, \rho_2) \in \llbracket C \rrbracket$ such that $\tau_1 = \tau_2$ and*
 342 *$\rho_1 = (\mathcal{E}, E[\text{assert}(0)], \dots)$.*

343 ▶ **Theorem 11** (Definability). *Let L be a library and $(\tau, \rho) \in \llbracket L \rrbracket$. There is a good client C*
 344 *compatible with L such that $(\tau, \rho') \in \llbracket C \rrbracket$ for some ρ' .*

3 Symbolic Semantics

346 Checking libraries for errors using the semantics of the previous section is infeasible, even when
 347 the traces are bounded in length, as ground values are concretely represented. In particular,
 348 integer values provided by O as arguments to calls or return values range over all integers.
 349 The typical way to mitigate this limitation is to execute the semantics symbolically, using
 350 symbolic variables for integers and path conditions to bind these variables to plausible values.
 351 We use this technique to devise a symbolic version of the trace semantics, corresponding to a
 352 symbolic execution which will enable us in the next sections to introduce a practical method
 353 and implementation for checking libraries for errors. The symbolic semantics is fully formal,
 354 closely following the developments of the previous section, and allows us to prove a strong
 355 form of correspondence between concrete and symbolic semantics (a bisimulation).

356 Apart from integers, another class of concrete values provided by O are method names.
 357 For them, the semantics we defined is symbolic by design: all method names played by O are
 358 going to be fresh and therefore picking just one of those fresh choices is sufficient (formally
 359 speaking, the semantics lives in nominal sets [32]). The reason why using fresh names for
 360 methods played by O is sound is that the effect of O calling a higher-order public method
 361 with an argument m (where m is another public method), and with $\lambda x.mx$, is equivalent as
 362 far as reachability of an error is concerned. In the latter case, the client semantics would
 363 create a fresh name m' , bind it to $\lambda x.mx$, and pass m' as an argument. We therefore just
 364 focus on this latter case.

365 The symbolic semantics involves terms that may contain symbolic values for integers. We
 366 therefore extend the syntax for values and terms to include such values, and abuse notation
 367 by continuing to use M to range over them. We let \mathbf{SInts} be a set of symbolic integers
 368 ranged over by κ and variants, and define:

$$369 \quad \text{Sym. Values } \tilde{v} ::= m \mid i \mid () \mid \kappa \mid \tilde{v} \oplus \tilde{v} \mid \langle \tilde{v}, \tilde{v} \rangle$$

$$370 \quad \text{Sym. Terms } M ::= \dots \mid \kappa$$

372 where, in $\tilde{v} \oplus \tilde{v}$, not both \tilde{v} can be integers. We moreover use a symbolic environment to
 373 store symbolic values for references, but also to keep track of arithmetic performed with
 374 symbolic integers. More precisely, we let σ be a finite partial map from the set $\mathbf{SInts} \cup \mathbf{Refs}$
 375 to symbolic values. Finally, we use pc to range over program conditions, which will be
 376 quantifier-free first-order formulas with variables taken from \mathbf{SInts} , and with \top, \perp denoting
 377 true and false respectively.

378 The semantics for closed symbolic terms involves configurations of the form (M, R, σ, pc, k) .
 379 Its rules include copies of those from Figure 2 (top) where the pc and σ are simply carried
 380 over. For example:

$$381 \quad (E[\lambda x.M], R, \sigma, pc, k) \rightarrow_s (E[m], R \uplus \{m \mapsto \lambda x.M\}, \sigma, pc, k)$$

382 where m is fresh. On the other hand, the following rules directly involve symbolic reasoning:

$$383 \quad (E[\text{assert}(\kappa)], R, \sigma, pc, k) \rightarrow_s (E[\text{assert}(0)], \sigma, pc \wedge (\kappa = 0), k)$$

$$\begin{array}{l}
(\widetilde{\text{INT}}) \quad \frac{(M, R, \sigma, pc, k) \rightarrow_s (M', R', \sigma, pc', k')}{(\mathcal{E}, M, R, \mathcal{P}, \mathcal{A}, \sigma, pc, k)_p \rightarrow_s (\mathcal{E}, M', R', \mathcal{P}, \mathcal{A}, \sigma', pc', k')_p} \\
(\widetilde{\text{PQ}}) \quad (\mathcal{E}, E[m\tilde{v}], R, \mathcal{P}, \mathcal{A}, \sigma, pc, k)_p \xrightarrow{\text{call}(m, \tilde{v})}_s ((m, E) :: \mathcal{E}, 0, R, \mathcal{P}', \mathcal{A}, \sigma, k)_o \\
(\widetilde{\text{OQ}}) \quad (\mathcal{E}, l, R, \mathcal{P}, \mathcal{A}, \sigma, pc, k)_o \xrightarrow{\text{call}(m, \tilde{v})}_s ((m, l+1) :: \mathcal{E}, m\tilde{v}, R, \mathcal{P}, \mathcal{A}', \sigma, pc, k)_p \\
(\widetilde{\text{PA}}) \quad ((m, l) :: \mathcal{E}, \tilde{v}, R, \mathcal{P}, \mathcal{A}, \sigma, pc, k)_p \xrightarrow{\text{ret}(m, \tilde{v})}_s (\mathcal{E}, l, R, \mathcal{P}', \mathcal{A}, \sigma, pc, k)_o \\
(\widetilde{\text{OA}}) \quad ((m, E) :: \mathcal{E}, l, R, \mathcal{P}, \mathcal{A}, \sigma, pc, k)_o \xrightarrow{\text{ret}(m, \tilde{v})}_s (\mathcal{E}, E[\tilde{v}], R, \mathcal{P}, \mathcal{A}', \sigma, pc, k)_p \\
\hline
(\widetilde{\text{PC}}) \quad m \in \mathcal{A} \text{ and } \mathcal{P}' = \mathcal{P} \cup (\text{Meths}(\tilde{v}) \cap \text{dom}(R)). \\
(\widetilde{\text{OC}}) \quad m \in \mathcal{P} \text{ and } (\tilde{v}', \mathcal{A}') \in \text{symval}(\theta, \mathcal{A}) \text{ where } \theta \text{ is the expected type of } \tilde{v}. \text{ Moreover:} \\
\text{symval}(\theta, \mathcal{A}) = \begin{cases} \{(\cdot, \mathcal{A})\} & \text{if } \theta = \text{unit} \\ \{(\kappa, \mathcal{A} \uplus \{\kappa\}) \mid \kappa \text{ is fresh in } \text{dom}(\sigma) \uplus \mathcal{A}\} & \text{if } \theta = \text{int} \\ \{(m, \mathcal{A} \uplus \{m\}) \mid m \text{ is fresh in } \text{dom}(R) \uplus \mathcal{A}\} & \text{if } \theta = \theta_1 \rightarrow \theta_2 \\ \{((\tilde{v}_1, \tilde{v}_2), \mathcal{A}_2) \mid (\tilde{v}_1, \mathcal{A}_1) \in \text{symval}(\theta_1, \mathcal{A}) \\ \quad (\tilde{v}_2, \mathcal{A}_2) \in \text{symval}(\theta_2, \mathcal{A}_1)\} & \text{if } \theta = \theta_1 \times \theta_2 \end{cases}
\end{array}$$

■ **Figure 4** Symbolic trace semantics rules. Rules $(\widetilde{\text{PQ}})$, $(\widetilde{\text{PA}})$ assume the condition $(\widetilde{\text{PC}})$, and similarly for $(\widetilde{\text{OQ}})$, $(\widetilde{\text{OA}})$ and $(\widetilde{\text{OC}})$. Note that $(\widetilde{\text{OQ}})$, $(\widetilde{\text{OA}})$ are non-deterministic as they introduce \tilde{v} .

$$\begin{array}{l}
384 \quad (E[\text{assert}(\kappa)], R, \sigma, pc, k) \rightarrow_s (E[()], R, \sigma, pc \wedge (\kappa \neq 0), k) \\
385 \quad (E[r], R, \sigma, pc, k) \rightarrow_s (E[\sigma(r)], R, \sigma, pc, k) \\
386 \quad (E[r := \tilde{v}], R, \sigma, pc, k) \rightarrow_s (E[()], R, \sigma[r \mapsto \tilde{v}], pc, k) \\
387 \quad (E[\tilde{v}_1 \oplus \tilde{v}_2], R, \sigma, pc, k) \rightarrow_s (E[\kappa], R, \sigma \uplus \{\kappa \mapsto \tilde{v}_1 \oplus \tilde{v}_2\}, pc, k) \quad \text{where } \kappa \text{ is fresh} \\
388 \quad (E[\text{if } \kappa \text{ then } M_1 \text{ else } M_0], R, \sigma, pc, k) \rightarrow_s (E[M_0], R, \sigma, pc \wedge (\kappa = 0), k) \\
389 \quad (E[\text{if } \kappa \text{ then } M_1 \text{ else } M_0], R, \sigma, pc, k) \rightarrow_s (E[M_1], R, \sigma, pc \wedge (\kappa \neq 0), k) \\
390
\end{array}$$

391 and where $\tilde{v}_1 \oplus \tilde{v}_2$ is a symbolic value (for $i_1 \oplus i_2$ the rule from Figure 1 applies).

392 We now extend the symbolic setting to the trace semantics. We define symbolic configurations for P and O respectively as:

$$394 \quad (\mathcal{E}, M, R, \mathcal{P}, \mathcal{A}, \sigma, pc, k)_p \qquad (\mathcal{E}, l, R, \mathcal{P}, \mathcal{A}, \sigma, pc, k)_o$$

395 with evaluation stack \mathcal{E} , proponent term M , counters $k, l \in \mathbb{N}$, method repository R , public
396 method name set \mathcal{P} , σ and pc as previously. The abstract name set \mathcal{A} is now a finite subset
397 of $\text{Meths} \cup \text{SInts}$, as we also need to keep track of the symbolic integers introduced by
398 O (in order to be able to introduce fresh such names). The rules for the symbolic trace
399 semantics are given in Figure 4. Note that O always refreshes names it passes. This is a
400 sound overapproximation of all names passed for the sake of analysis.

401 Similarly to Definition 6, we can define the symbolic semantics of libraries.

402 ► **Definition 12.** Given library L , the symbolic semantics of L is:

$$\begin{aligned}
403 \quad \llbracket L \rrbracket_s &= \{(\tau, \rho) \mid (L, \emptyset, \emptyset, \emptyset, \emptyset) \xrightarrow{\text{bld}}^* (\varepsilon, R, S, \mathcal{P}, \mathcal{A}) \\
404 \quad &\quad \wedge (\varepsilon, 0, R, \mathcal{P}, \mathcal{A}, S, \top, 0)_o \xrightarrow{\tau}_s \rho \wedge \exists \mathcal{M}. \mathcal{M} \models \rho(\sigma)^\circ \wedge \rho(pc)\} \\
405
\end{aligned}$$

406 where $\rho(\chi)$ is component χ in configuration ρ , and \mathcal{M} is a model as defined in the next
407 section. We say that $\llbracket L \rrbracket_s$ fails if it contains some $(\tau, (\mathcal{E}, E[\text{assert}(0)], \dots))$.

408 The symbolic rules follow those of the concrete semantics, the biggest change being the
409 treatment of symbolic values played by O . Condition $(\widetilde{\text{OC}})$ stipulates that O plays distinct

28:12 Symbolic Execution Game Semantics

410 fresh symbolic integers as well as fresh method names, in each appropriate position in \tilde{v} , and
 411 all these names are included in the set \mathcal{A} .

412 ► **Example 13.** As with Example 7, we consider the DAO attack. Running the symbolic
 413 semantics, we find the following minimal class of errors. We write $\sigma_{\tilde{v}}$ for a symbolic
 414 environment $[bal \mapsto \tilde{v}]$.

$$\begin{aligned}
 & (\varepsilon, 2, \sigma_{100}, k_0)_o \xrightarrow{wdraw(\kappa_1)?} ((wdraw, 1), wdraw(\kappa_1), \sigma_{100}, 2)_p \\
 & \quad \rightarrow^* ((wdraw, 1), E[send(\kappa_1)], \sigma_{100}, 1)_p \xrightarrow{send(\kappa_1)?} ((send, E), 2, \sigma_{100}, 1)_o \\
 & \quad \xrightarrow{wdraw(\kappa_2)?} ((wdraw, 1), wdraw(\kappa_2), \sigma_{100}, 1)_p \\
 & \quad \rightarrow^* ((wdraw, 1), E'[send(\kappa_2)], \sigma_{100}, 0)_p \xrightarrow{send(\kappa_2)?} ((send, E), 2, \sigma_{100}, 0)_o \\
 & \quad \xrightarrow{send()!} ((wdraw, 1), E'[], \sigma_{100}, 0)_p \\
 & \quad \rightarrow^* ((wdraw, 1), (), \sigma_{100-\kappa_2}, 0)_p \xrightarrow{wdraw()!} ((send, E), 1, \sigma_{100-\kappa_2}, 0)_o \\
 & \quad \xrightarrow{send()!} ((wdraw, 1), E[()], \sigma_{100-\kappa_2}, 1)_p \\
 & \quad \rightarrow^* ((wdraw, 1), E[assert(!bal \geq 0)], \sigma_{100-\kappa_2-\kappa_1}, 1)_p
 \end{aligned}$$

424 For this to be a valid error, we require $(\kappa_1, \kappa_2 \leq 100) \wedge (100 - \kappa_2 - \kappa_1 < 0)$ to be satisfiable.
 425 Taking assignment $\{\kappa_1 \mapsto 100, \kappa_2 \mapsto 1\}$, we show the path is valid.

426 3.1 Soundness

427 The main result of this section is establishing the soundness of the symbolic semantics: a
 428 trace and a specific configuration can be achieved symbolically iff they can be achieved
 429 concretely as well. In fact, we will need to quantify this statement as, by construction, the
 430 symbolic semantics requires O to always place fresh method names, whereas in the concrete
 431 semantics O is given the freedom to play old names as well. What we show is that the
 432 symbolic semantics corresponds (via *bisimilarity*) to a restriction of the concrete semantics
 433 where O plays fresh names only. This restriction is sound, in the sense that it is sufficient for
 434 identifying when a configuration can fail. We make this precise below.

435 A **model** \mathcal{M} is a finite partial map from symbolic integers to concrete integers. Given
 436 such an \mathcal{M} and a formula ϕ , we define $\mathcal{M} \models \phi$ using a standard first-order logic interpretation
 437 with integers and arithmetic operators (in particular, we require that all symbolic integers in
 438 ϕ are in the domain of \mathcal{M}). Moreover, for any symbolic term M (or trace, move, etc.), we
 439 denote by $M\{\mathcal{M}\}$ the concrete term we obtain by substituting any symbolic integer κ of M
 440 with its corresponding concrete integer $\mathcal{M}(\kappa)$. Finally, given a symbolic environment σ , we
 441 define its formula representation σ° recursively by:

$$442 \quad \emptyset^\circ = \top, \quad (\sigma \uplus \{r \mapsto v\})^\circ = \sigma^\circ, \quad (\sigma \uplus \{\kappa \mapsto v\})^\circ = \sigma^\circ \wedge (\kappa = v).$$

443 We now define notions for equivalence between symbolic and concrete configurations.
 444 Let \mathcal{M} be a model. For any concrete configuration $\rho = (\mathcal{E}, \chi, R, S, \mathcal{P}, \mathcal{A}, k)$ and symbolic
 445 configuration $\rho_s = (\mathcal{E}', \chi', R', \mathcal{P}', \mathcal{A}', \sigma, pc, k')$, we say they are *equivalent in \mathcal{M}* , written
 446 $\rho =_{\mathcal{M}} \rho_s$, if:

- 447 ■ $(\mathcal{E}, \chi, R) = (\mathcal{E}', \chi', R')\{\mathcal{M}\}, \mathcal{P} = \mathcal{P}', \mathcal{A} = \mathcal{A}' \cap \mathbf{Meths}$ and $S = (\sigma \upharpoonright \mathbf{Refs})\{\mathcal{M}\};$
- 448 ■ $\text{dom}(\mathcal{M}) = (\mathcal{A}' \cup \text{dom}(\sigma)) \cap \mathbf{SInts}$ and $\mathcal{M} \models pc \wedge \sigma^\circ.$

449 The notion of equivalence we require between concrete configurations and their symbolic
450 counterparts is behavioural equivalence, modulo O playing fresh names.

451 More precisely, a transition $\rho \xrightarrow{\chi} \rho'$ is called *O-refreshing* if, when ρ is an O -configuration
452 and $\chi = \text{call}/\text{ret}(m, v)$ then all names in v are fresh and distinct. A set \mathcal{R} with elements
453 of the form $(\rho, \mathcal{M}, \rho_s)$ is a **bisimulation** if, whenever $(\rho, \mathcal{M}, \rho_s) \in \mathcal{R}$, written $\rho \mathcal{R}_{\mathcal{M}} \rho_s$ then
454 $\rho =_{\mathcal{M}} \rho_s$ and, using χ to range over moves and ε (i.e. no move):

- 455 ■ if $\rho \xrightarrow{\chi} \rho'$ is O -refreshing then there exists $\mathcal{M}' \supseteq \mathcal{M}$ such that $\rho_s \xrightarrow{\chi_s} \rho'_s$, with $\chi =$
456 $\chi_s\{\mathcal{M}'\}$, and $\rho' \mathcal{R}_{\mathcal{M}'} \rho'_s$;
- 457 ■ if $\rho_s \xrightarrow{\chi_s} \rho'_s$ then there exists $\mathcal{M}' \supseteq \mathcal{M}$ such that $\rho \xrightarrow{\chi\{\mathcal{M}'\}} \rho'$ and $\rho' \mathcal{R}_{\mathcal{M}'} \rho'_s$.

458 We let \sim be the largest bisimulation relation: $\rho \sim_{\mathcal{M}} \rho_s$ iff there is bisimulation \mathcal{R} such that
459 $\rho \mathcal{R}_{\mathcal{M}} \rho_s$.

460 We can show that concrete and symbolic configurations are bisimilar.

461 ► **Lemma 14.** *Given ρ, ρ_s a concrete and symbolic configuration respectively, and \mathcal{M} a model*
462 *such that $\rho =_{\mathcal{M}} (\rho')$, we have $\rho \sim_{\mathcal{M}} \rho_s$.*

463 **Proof (sketch).** We show that $\{(\rho, \mathcal{M}, \rho') \mid \rho =_{\mathcal{M}} \rho'\}$ is a bisimulation. ◀

464 Next, we argue that O -refreshing transitions suffice for examining failure of concrete
465 configurations. Indeed, suppose τ is a trace leading to fail, and where O plays an old name
466 m in argument position in a given move. Then, τ can be simulated by a trace τ' that uses
467 a fresh m' in place of m . If m is an O -name, we obtain τ' from τ by following exactly the
468 same transitions, only that some P -calls to m are replaced by calls to m' (and accordingly
469 for returns). If, on the other hand, m is a P -name, then the simulation performed by τ'
470 is somewhat more elaborate: some internal calls to m will be replaced by P -calls to m' ,
471 immediately followed by the required calls to m (and dually for returns).

472 ► **Lemma 15 (O-Refreshing).** *Let ρ be a concrete configuration. Then, ρ fails iff it fails using*
473 *only O -refreshing transitions.*

474 With the above, we can prove soundness.

475 ► **Theorem 16 (Soundness).** *For any L , $\llbracket L \rrbracket$ fails iff $\llbracket L \rrbracket_s$ fails.*

476 **Proof.** Lemma 14 implies that $\llbracket L \rrbracket_s$ fails iff $\llbracket L \rrbracket$ fails with O -refreshing transitions, which in
477 turns occurs iff $\llbracket L \rrbracket$ fails, by Lemma 15. ◀

478 3.2 Bounded Analysis for Libraries

479 Definition 12 states how the symbolic trace semantics can be used to independently check
480 libraries for errors. As with the trace semantics in Definition 6, this is strongly normalising
481 when given an upper limit to the call counters. As such, $\llbracket L \rrbracket_s$ with counter bounds $k_0, l_0 \in \mathbb{N}$,
482 for k, l respectively, defines a finite set (modulo selecting of fresh names) of reachable valid
483 configurations within $k \leq k_0, l \leq l_0$, where validity is defined by the satisfiability of the
484 symbolic environment σ and the path condition pc of the configuration reached. By virtue of
485 Theorems 9 and 16, every valid reachable configuration that is failed (evaluates an invalid
486 assertion) is realisable by some client. And viceversa.

487 Given a library L , taking $\mathcal{F}\llbracket L \rrbracket_s$ to be all reachable final configurations, we have the
488 exhaustive set of paths L can reach. In $\mathcal{F}\llbracket L \rrbracket_s$, every failed configuration (τ, ρ) , i.e. such
489 that ρ holds a term $E[\text{assert}(0)]$, defines a reachable assertion violation, where τ is a true

	$l \leq 1$	$l \leq 2$	$l \leq 3$
$k \leq 2$	226/70/45 (555s)	5708/60/44 (4710s)	9656/3/23 (12471s)
$k \leq 3$	1254/67/51 (1475s)	4092/27/18 (13482s)	4187/17/12 (16649s)
$k \leq 4$	3392/63/48 (3180s)	3069/19/14 (15903s)	1335/12/10 (17765s)
$k \leq 5$	3659/57/45 (4787s)	895/15/10 (16757s)	215/11/9 (17796s)

$a/b/c$ (d) for a traces found in b successful runs taking d seconds in total where c out of 59 unsafe files were found to have bugs, per bound.
59 of 59 unsafe files found to have bugs over the various bounds checked

■ **Table 1** Table recording performance of HOLiK on our benchmarks

490 counterexample. Hence, to check L for assertion violations it suffices to produce a finite
491 representation of the set $\mathcal{F}[[L]]_s$. One approach is to bound the depth of analysis by setting an
492 upper bound to the call counters, using a name generator to make deterministic the creation
493 of fresh names, and then exhaustively search all final configurations for failed elements. In
494 the following section we implement this routine and test it.

495 4 Implementation and Experiments

496 We implemented the syntax and symbolic trace semantics (symbolic games) for HOLi in
497 the \mathbb{K} semantic framework [33] as a proof of concept, and tested it on 70 sample libraries.¹
498 Using \mathbb{K} 's option to exhaustively expand all transitions, \mathbb{K} is able to build a closure of all
499 applicable rules. By providing a bound on the call counters, we produce a finite set of all
500 reachable valid symbolic configurations up to the given depth (equivalent to finding every
501 valid $\rho \in \mathcal{F}[[L]]_s$) which thus implements our bounded symbolic execution.

502 We wrote and adapted examples of coding errors into a set of 70 sample libraries written
503 in HOLi, totalling 6,510 lines of code (LoC). Examples adapted from literature include:
504 reentrancy bugs from smart contracts [3, 24]; variations of the “awkward example” [31];
505 various programs from the MoChi benchmark [36]; and simple implementations related to
506 concurrent programming (e.g. flat combining and race conditions) where errors may occur
507 in a single thread due to higher-order behaviour. We also combined several libraries, by
508 concatenating refactored method and reference definitions, to generate larger libraries that
509 are harder to solve. Combined files range from 150 to 520 LoC.

510 We ran HOLiK on all sample libraries, lexicographically increasing the bounds from
511 $k \leq 2, l \leq 1$ to $k \leq 5, l \leq 3$ (totalling 78,120 LoC checked), with a timeout set to five minutes
512 per library. We start from $k \leq 2$ because it provides the minimum nesting needed to observe
513 higher-order semantics. All experiments ran on an Ubuntu 19.04 machine with 16GB RAM,
514 Intel Core i7 3.40GHz CPU, with intermediate calls to Z3 to prune invalid configurations. Per
515 bound, the number of counterexamples found, the time taken in seconds, and the execution
516 status, i.e. whether it terminated or not, are recorded in Table 1.

517 We can observe that independently increasing the bounds for k and l causes exponential
518 growth in the total time taken, which is expected from symbolic execution. Note that the
519 time tends towards 21000 seconds because of the timeout set to 5 minutes for 70 programs. In
520 particular, while the number of errors found grows exponentially with respect to the increase

¹ The tool and its benchmarks can be found at: <https://github.com/LaifsV1/HOLiK>.

521 in bounds – which is due to the exponential growth in paths – this trend does not continue
522 indefinitely because programs start timing out without reporting any errors as the bounds
523 grow. With bounds $k \leq 2$ and $l \leq 1$, all 70 programs in our benchmark were successfully
524 analysed, though not all minimal errors were found until the bounds were increased further.
525 Cumulatively, all unsafe programs in our benchmark were correctly identified.

526 While the table may suggest that increasing bound for l is more beneficial than that
527 for k , the number of errors reported does not imply every trace is useful. For instance,
528 increasing the bound for l can lead to errors re-merging in a higher-order version, which
529 suggests potential gain from a partial order reduction. Overall, the k and l counters are
530 incomparable as they keep track of different behaviours. Finally, since HOLiK was able
531 to handle every file and correctly identified all unsafe files in the benchmark, we conclude
532 that HOLiK, as a proof of concept, captures the full range of behaviours in higher-order
533 libraries. Results suggest that the tool scales up to at least medium-sized programs (<1000
534 LoC), which is promising because real-world medium-size higher-order programs have been
535 proven infeasible to check with standard techniques (e.g. the DAO withdraw contract was
536 approximately 100 LoC).

537 **5** Related Work

538 Game semantics techniques have been applied to program equivalence verification by reducing
539 program equivalence to language equivalence in a decidable automata class [15, 1]. Equivalence
540 tools can be used for reachability but, as they perform full verification, they can only cover
541 lower-order recursion-free language fragments. For example, the Coneqct [25] tool can verify
542 the simplified DAO attack, but cannot check higher-order or recursive functions (e.g. the
543 “file lock” and “flat combiner” examples), and operates on integers concretely. Close to our
544 approach is also Symbolic GameChecker [11], which performs symbolic model checking by
545 using a representation of games based on symbolic finite-state automata. The tool works
546 on recursion-free Idealized Algol with first-order functions, which supports only integer
547 references. On the other hand, it is complete (not bounded) on the fragment that it covers.

548 Besides games techniques, a recent line of work on verification of contracts in Racket
549 [28, 27] is the work closest to ours. Racket contracts exist in a higher-order setting similar
550 to ours, and generalise higher-order pre and post conditions, and thus specify safety. To
551 verify these, [28] defines a symbolic execution based on what they call “demonic context” in
552 prior work [39]. This either returns a symbolic value to a call, or performs a call to a known
553 method within some unknown context, thus approximating all the possible higher-order
554 behaviours, and is equivalent to the role the opponent plays in our games. In [27], the
555 technique is extended to handle state, and finitised for total verification. The approaches
556 are notionally similar to ours, since both amount to Symbolic Execution for an unknown
557 environment. In substance, the techniques are very different and in particular ours is based
558 on a semantics theory which allows us to obtain compositionality and definability results,
559 which are not proven for [27] and proven for [28] only in a stateless setting. On the other
560 hand, Racket contracts can be used for richer verification questions than assertion violations.
561 In terms of tool performance, we provide a comparison of the techniques in Appendix B.

562 Another relevant line of work is that of verifying programs in the Ethereum Platform.
563 Smart contracts call for techniques that handle the environment, with a focus on reentrancy.
564 Tools like Oyente [24] and Majan [29] use pre-defined patterns to find bugs in the transaction
565 order, but are not sound or complete. ReGuard [23] finds sound reentrancy bugs using a
566 fuzzing engine to generate random transactions to check with a reentrancy automaton. In

567 principle, it may detect reentrancy faster than symbolic execution (native execution is faster
 568 [41]), but, is incomplete even in a bounded setting. More closely related to our approach,
 569 [17] considers the possibility of an *unknown* contract $c?$ calling a *known* contract c^* at each
 570 higher call level. This can be generalised in our game semantics as *abstract* and *public* names
 571 calling each other, but their focus is on modelling reentrancy, while we handle the full range
 572 of higher-order behaviours.

573 Like KLEE [4] and jCUTE [37], our implementation is a symbolic execution tool. These
 574 are generally able to find first-order counterexamples, but are unable to produce higher-order
 575 traces involving unknown code. Particularly, KLEE and jCUTE only handle symbolic calls
 576 provided these can be concretised. This partially models the environment, but calls are often
 577 impossible to concretise with libraries. The CBMC [6, 20] bounded model checking approach,
 578 which also bounds function application to a fixed depth, partially handle calls to unknown
 579 code by returning a non-deterministic value to such calls. This is equivalent to a game where
 580 only move available to the opponent is to answer questions. This restriction allows CBMC
 581 to find some bugs caused by interaction with the environment, but misses errors that arise
 582 from transferring flow of control (e.g. reentrancy). The typical BMC approach also misses
 583 bugs involving disclosure of names.

584 Higher-order model checking tools like MoCHi [36] are also related. MoCHi model checks
 585 a pure subset of OCaml and is based on predicate abstraction and CEGAR and higher-order
 586 recursion scheme model checkers. The modular approach [35] further extends this idea
 587 with modular analysis that guesses refinement intersection types for each top-level function.
 588 Although generally incomparable, HOLiK covers program features that MoCHi does not:
 589 MoCHi does not handle references and support for open code is limited (from experiments,
 590 and private communication with the authors).

591 **6 Future Directions**

592 Observing errors resurface deeper in the trace suggests the possibility of defining a partial
 593 order for our semantics to obtain equivalence classes for configurations and thus eliminate
 594 paths that involve known errors [30, 40]. Additionally, while k and l successfully bound
 595 infinite behaviour, a notion of bounding can be arbitrarily chosen. In fact, while we chose to
 596 directly bound the sources of infinite behaviour in method calls for simplicity of proofs and
 597 implementation, the theory does not prevent the generalisation of k and l as a monotonic
 598 cost function that bounds the semantics. It may also be worth considering the elimination of
 599 bounds entirely for the sake of unbounded verification. For this, one direction is abstract
 600 interpretation [9, 8], which amounts to defining overapproximations for values in our language
 601 to then attempt to compute a fixpoint for the range of values that assertions may take.
 602 However, defining and using abstract domains that maintain enough precision to check higher-
 603 order behaviours, such as reentrancy, is not a simple extension of the theory. Another direction,
 604 similar to Coneqct [25], is to define a push-down system for our semantics. Particularly,
 605 the approach in [25] is based on the decidability of reachability in fresh-register pushdown
 606 automata, and would require overapproximations for methods and integers. As with abstract
 607 interpretation, this would require defining abstract domains for methods and integers. While
 608 methods could be approximated using a finite set of names, as with k -CFA [38], an extension
 609 using integer abstract domains would need refinement to tackle reentrancy attacks. Finally,
 610 MoCHi [36] shows that it is possible to use CEGAR and higher-order recursion schemes
 611 for unbounded verification of higher-order programs. However, an extension of the MoCHi
 612 approach to include references and open code is not obvious.

613 ——— **References** ———

- 614 1 S. Abramsky, D. R. Ghica, L. Ong, and A. Murawski. Algorithmic game semantics and
615 component-based verification. In *Proceedings of SAVBCS 2003: Specification and Verification*
616 *of Component-Based Systems, Workshop at ESEC/FASE 2003*, pages 66–74, 2003. published
617 as Technical Report 03-11, Department of Computer Science, Iowa State University. URL:
618 <http://www.cs.iastate.edu/~leavens/SAVBCS/2003/papers/SAVBCS03.pdf>.
- 619 2 Samson Abramsky and Guy McCusker. Game semantics. In Ulrich Berger and Helmut
620 Schwichtenberg, editors, *Computational Logic*, pages 1–55, Berlin, Heidelberg, 1999. Springer
621 Berlin Heidelberg.
- 622 3 Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. A survey of attacks on ethereum smart
623 contracts sok. In *Proceedings of the 6th International Conference on Principles of Security*
624 *and Trust - Volume 10204*, pages 164–186, New York, NY, USA, 2017. Springer-Verlag New
625 York, Inc. doi:10.1007/978-3-662-54455-6_8.
- 626 4 Cristian Cadar, Daniel Dunbar, and Dawson Engler. Klee: Unassisted and automatic generation
627 of high-coverage tests for complex systems programs. In *Proceedings of the 8th USENIX*
628 *Conference on Operating Systems Design and Implementation, OSDI'08*, pages 209–224,
629 Berkeley, CA, USA, 2008. USENIX Association. URL: [http://dl.acm.org/citation.cfm?](http://dl.acm.org/citation.cfm?id=1855741.1855756)
630 [id=1855741.1855756](http://dl.acm.org/citation.cfm?id=1855741.1855756).
- 631 5 Cristiano Calcagno, Dino Distefano, J r my Dubreil, Dominik Gabi, Pieter Hooimeijer, Martino
632 Luca, Peter W. O'Hearn, Irene Papakonstantinou, Jim Purbrick, and Dulma Rodriguez. Moving
633 fast with software verification. In Klaus Havelund, Gerard J. Holzmann, and Rajeev Joshi,
634 editors, *NASA Formal Methods - 7th International Symposium, NFM 2015, Pasadena, CA,*
635 *USA, April 27-29, 2015, Proceedings*, volume 9058 of *Lecture Notes in Computer Science*,
636 pages 3–11. Springer, 2015.
- 637 6 Edmund M. Clarke, Daniel Kroening, and Flavio Lerda. A tool for checking ANSI-C programs.
638 In Kurt Jensen and Andreas Podelski, editors, *Tools and Algorithms for the Construction*
639 *and Analysis of Systems, 10th International Conference, TACAS 2004, Proceedings*, volume
640 2988 of *Lecture Notes in Computer Science*, pages 168–176. Springer, 2004. doi:10.1007/
641 978-3-540-24730-2_15.
- 642 7 Byron Cook, Kareem Khazem, Daniel Kroening, Serdar Tasiran, Michael Tautschnig, and
643 Mark R. Tuttle. Model checking boot code from AWS data centers. In Hana Chockler and
644 Georg Weissenbacher, editors, *Computer Aided Verification - 30th International Conference,*
645 *CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July*
646 *14-17, 2018, Proceedings, Part II*, volume 10982 of *Lecture Notes in Computer Science*, pages
647 467–486. Springer, 2018.
- 648 8 Agostino Cortesi and Matteo Zanioli. Widening and narrowing operators for abstract
649 interpretation. *Computer Languages, Systems & Structures*, 37(1):24–42, 2011. doi:
650 10.1016/j.cl.2010.09.001.
- 651 9 Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for
652 static analysis of programs by construction or approximation of fixpoints. In Robert M.
653 Graham, Michael A. Harrison, and Ravi Sethi, editors, *Conference Record of the Fourth ACM*
654 *Symposium on Principles of Programming Languages, Los Angeles, California, USA, January*
655 *1977*, pages 238–252. ACM, 1977. doi:10.1145/512950.512973.
- 656 10 Patrick Cousot, Radhia Cousot, J r me Feret, Laurent Mauborgne, Antoine Min , and Xavier
657 Rival. Why does astr e scale up? *Formal Methods in System Design*, 35(3):229–264, 2009.
- 658 11 Aleksandar S. Dimovski. Program verification using symbolic game semantics. *Theor. Comput.*
659 *Sci.*, 560:364–379, 2014. doi:10.1016/j.tcs.2014.01.016.
- 660 12 Quinn Dupont. *Experiments in Algorithmic Governance: A history and ethnography of " The*
661 *DAO, " a failed Decentralized Autonomous Organization*, chapter 8. Routledge, 01 2017.
- 662 13 William E. Howden. Symbolic testing and the dissect symbolic evaluation system. *Software*
663 *Engineering, IEEE Transactions on*, SE-3:266–278, 08 1977. doi:10.1109/TSE.1977.231144.

- 664 14 Dan R. Ghica. Applications of game semantics: From program analysis to hardware synthesis.
665 In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS*
666 *2009, 11-14 August 2009, Los Angeles, CA, USA*, pages 17–26. IEEE Computer Society, 2009.
667 doi:10.1109/LICS.2009.26.
- 668 15 Dan R. Ghica and Guy McCusker. Reasoning about idealized ALGOL using regular languages.
669 In Ugo Montanari, José D. P. Rolim, and Emo Welzl, editors, *Automata, Languages and*
670 *Programming, 27th International Colloquium, ICALP 2000, Geneva, Switzerland, July 9-15,*
671 *2000, Proceedings*, volume 1853 of *Lecture Notes in Computer Science*, pages 103–115. Springer,
672 2000. doi:10.1007/3-540-45022-X_10.
- 673 16 Dan R. Ghica and Nikos Tzevelekos. A system-level game semantics. *Electr. Notes Theor.*
674 *Comput. Sci.*, 286:191–211, 2012. doi:10.1016/j.entcs.2012.08.013.
- 675 17 Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. Foundations and tools for the
676 static analysis of ethereum smart contracts. In Hana Chockler and Georg Weissenbacher,
677 editors, *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part*
678 *of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings,*
679 *Part I*, volume 10981 of *Lecture Notes in Computer Science*, pages 51–78. Springer, 2018.
680 doi:10.1007/978-3-319-96145-3_4.
- 681 18 A. Jeffrey and J. Rathke. A fully abstract may testing semantics for concurrent objects. In
682 *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*, pages 101–112, July
683 2002. doi:10.1109/LICS.2002.1029820.
- 684 19 James C. King. A new approach to program testing. *SIGPLAN Not.*, 10(6):228–233, April
685 1975. URL: <http://doi.acm.org/10.1145/390016.808444>, doi:10.1145/390016.808444.
- 686 20 Daniel Kroening. The CBMC homepage. <http://www.cprover.org/cbmc/>, 2017. [Online;
687 accessed 13-Jun-2017].
- 688 21 James Laird. A fully abstract trace semantics for general references. In Lars Arge, Christian
689 Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *Automata, Languages and Pro-*
690 *gramming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland, July 9-13, 2007,*
691 *Proceedings*, volume 4596 of *Lecture Notes in Computer Science*, pages 667–679. Springer, 2007.
692 doi:10.1007/978-3-540-73420-8_58.
- 693 22 Yu-Yang Lin and Nikos Tzevelekos. Symbolic execution game semantics. Extended version
694 with full proofs, Feb 2020. URL: [https://github.com/LaifsV1/HOLiK/raw/master/paper/](https://github.com/LaifsV1/HOLiK/raw/master/paper/full-paper.pdf)
695 [full-paper.pdf](https://github.com/LaifsV1/HOLiK/raw/master/paper/full-paper.pdf).
- 696 23 C. Liu, H. Liu, Z. Cao, Z. Chen, B. Chen, and B. Roscoe. Reguard: Finding reentrancy bugs in
697 smart contracts. In *2018 IEEE/ACM 40th International Conference on Software Engineering:*
698 *Companion (ICSE-Companion)*, pages 65–68, May 2018.
- 699 24 Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. Making smart
700 contracts smarter. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer*
701 *and Communications Security, CCS '16*, pages 254–269, New York, NY, USA, 2016. ACM.
702 doi:10.1145/2976749.2978309.
- 703 25 Andrzej S. Murawski, Steven J. Ramsay, and Nikos Tzevelekos. A contextual equivalence
704 checker for IMJ*. In Bernd Finkbeiner, Geguang Pu, and Lijun Zhang, editors, *Automated*
705 *Technology for Verification and Analysis*, pages 234–240, Cham, 2015. Springer International
706 Publishing.
- 707 26 Andrzej S. Murawski and Nikos Tzevelekos. Higher-order linearisability. In *28th International*
708 *Conference on Concurrency Theory, CONCUR 2017, September 5-8, 2017, Berlin, Germany,*
709 *pages 34:1–34:18*, 2017. doi:10.4230/LIPIcs.CONCUR.2017.34.
- 710 27 Phuc C. Nguyen, Thomas Gilray, Sam Tobin-Hochstadt, and David Van Horn. Soft contract
711 verification for higher-order stateful programs. *PACMPL*, 2(POPL):51:1–51:30, 2018. doi:
712 10.1145/3158139.
- 713 28 Phuc C. Nguyen and David Van Horn. Relatively complete counterexamples for higher-order
714 programs. In David Grove and Steve Blackburn, editors, *Proceedings of the 36th ACM*

- 715 *SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR,*
716 *USA, June 15-17, 2015*, pages 446–456. ACM, 2015. doi:10.1145/2737924.2737971.
- 717 29 Ivica Nikolić, Aashish Kolluri, Ilya Sergey, Prateek Saxena, and Aquinas Hobor. Finding the
718 greedy, prodigal, and suicidal contracts at scale. In *Proceedings of the 34th Annual Computer*
719 *Security Applications Conference, ACSAC '18*, pages 653–663, New York, NY, USA, 2018.
720 ACM. doi:10.1145/3274694.3274743.
- 721 30 Doron A. Peled. All from one, one for all: on model checking using representatives. In Costas
722 Courcoubetis, editor, *Computer Aided Verification, 5th International Conference, CAV '93,*
723 *Elounda, Greece, June 28 - July 1, 1993, Proceedings*, volume 697 of *Lecture Notes in Computer*
724 *Science*, pages 409–423. Springer, 1993. doi:10.1007/3-540-56922-7_34.
- 725 31 Andrew Pitts and Ian Stark. Operational reasoning for functions with local state. In *Higher*
726 *Order Operational Techniques in Semantics*, pages 227–273. Cambridge University Press, 1998.
- 727 32 Andrew M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge
728 University Press, New York, NY, USA, 2013.
- 729 33 Grigore Roşu and Traian Şerbănuţă. An overview of the k semantic framework. *The Journal of*
730 *Logic and Algebraic Programming*, 79:397–434, 08 2010. doi:10.1016/j.jlap.2010.03.012.
- 731 34 Robert S. Boyer, Bernard Elspas, and Karl Levitt. Select—a formal system for testing and
732 debugging programs by symbolic execution. *ACM SIGPLAN Notices*, 10:234–245, 06 1975.
733 doi:10.1145/390016.808445.
- 734 35 Ryosuke Sato and Naoki Kobayashi. Modular verification of higher-order functional programs.
735 In Hongseok Yang, editor, *Programming Languages and Systems - 26th European Symposium*
736 *on Programming, ESOP 2017*, volume 10201 of *Lecture Notes in Computer Science*, pages
737 831–854. Springer, 2017. doi:10.1007/978-3-662-54434-1_31.
- 738 36 Ryosuke Sato, Hiroshi Unno, and Naoki Kobayashi. Towards a scalable software model checker
739 for higher-order programs. In Elvira Albert and Shin-Cheng Mu, editors, *Proceedings of the*
740 *ACM SIGPLAN 2013 Workshop on Partial Evaluation and Program Manipulation, PEPM 2013,*
741 *Rome, Italy, January 21-22, 2013*, pages 53–62. ACM, 2013. doi:10.1145/2426890.2426900.
- 742 37 Koushik Sen and Gul Agha. Cute and jcute: Concolic unit testing and explicit path model-
743 checking tools. In Thomas Ball and Robert B. Jones, editors, *Computer Aided Verification,*
744 pages 419–423, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- 745 38 Olin Shivers. *Control-Flow Analysis of Higher-Order Languages*. PhD thesis, Carnegie-Mellon
746 University, 1991.
- 747 39 Sam Tobin-Hochstadt and David Van Horn. Higher-order symbolic execution via contracts.
748 In Gary T. Leavens and Matthew B. Dwyer, editors, *Proceedings of the 27th Annual ACM*
749 *SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications,*
750 *OOPSLA 2012, part of SPLASH 2012, Tucson, AZ, USA, October 21-25, 2012*, pages 537–554.
751 ACM, 2012. doi:10.1145/2384616.2384655.
- 752 40 Antti Valmari. Stubborn sets for reduced state space generation. In Grzegorz Rozenberg,
753 editor, *Advances in Petri Nets 1990 [10th International Conference on Applications and*
754 *Theory of Petri Nets, Bonn, Germany, June 1989, Proceedings]*, volume 483 of *Lecture Notes*
755 *in Computer Science*, pages 491–515. Springer, 1989. doi:10.1007/3-540-53863-1_36.
- 756 41 Insu Yun, Sangho Lee, Meng Xu, Yeongjin Jang, and Taesoo Kim. Qsym: A practical concolic
757 execution engine tailored for hybrid fuzzing. In *Proceedings of the 27th USENIX Conference on*
758 *Security Symposium, SEC'18*, pages 745–761, Berkeley, CA, USA, 2018. USENIX Association.
759 URL: <http://dl.acm.org/citation.cfm?id=3277203.3277260>.

760 **A** Motivating examples

761 Our file lock example provides a scenario where the library makes it possible for the client to
 762 update a file without first reacquiring the lock for it. The library contains an empty private
 763 method `updateFile` that simulates file access. The library also provides a public method
 764 `openFile`, which locks the file, allows the user to update the file indirectly, and then releases
 765 the lock.

```
766
767 1 import userExec :((unit → unit) → unit)
768 2 int lock := 0;
769 3 private updateFile(x:unit) :(unit) = { () };
770 4 public openFile (u:unit) :(unit) = {
771 5   if (!lock) then ()
772 6   else (lock := 1;
773 7         let write = fun(x:unit):(unit) → (assert(!lock);updateFile());
774 8         in userExec(write); lock := 0) };
775
```

776 The bug here is that `openFile` creates a `write` method, which it then passes to the client,
 777 via `userExec(write)`, to use whenever they want. This provides the client indirect access to
 778 the private method `updateFile`, which it can call without first acquiring the lock. Running
 779 this example in HOLiK we obtain the following minimal trace:

```
780   call⟨openFile, ()⟩ · call⟨userExec, m2⟩ · ret⟨userExec, ()⟩
781   · ret⟨openFile, ()⟩ · call⟨m2, ()⟩
782
```

783 where m_2 is the method *name* generated by the library and bound to the variable `write`.
 784 This example serves as a representative of a class of bugs caused by revealing methods to
 785 the environment, a higher-order problem, in this case involving the second-order method
 786 `userExec` revealing m_2 .

787 Next, we simulate double deallocation using a global reference `addr` as the memory
 788 address. The library defines private methods `alloc` and `free` to simulate allocation and
 789 freeing. The empty private method `doSomething` serves as a placeholder for internal computation
 790 that does not free memory.

```
791
792 1 import getInput :(unit → int)
793 2 int addr := 0; // 0 means address is free
794 3 private alloc (u:unit) :(unit) = {
795 4   if not(!addr) then addr := 1 else () };
796 5 private free (u:unit) :(unit) = {
797 6   assert(!addr); addr := 0 };
798 7 private doSomething (i:int) :(unit) = { () };
799 8 public run (u:unit) :(unit) = {
800 9   alloc(); doSomething(getInput ()); free() };
801
```

802 The error occurs in line 9, which calls the client method `getInput`. This passes control to
 803 the client, who can now call `run` again, thus causing `free` to be called twice. Executing the
 804 example on HOLiK, we obtain the following trace:

```
805   call⟨run, ()⟩ · call⟨getInput, ()⟩ · call⟨run, ()⟩ · call⟨getInput, ()⟩
806   · ret⟨getInput, x1⟩ · ret⟨run, ()⟩ · ret⟨getInput, x2⟩
807
```

808 As with the DAO attack, this is a reentrancy bug.

809 Finally, we have an unsafe implementation of a flat combiner. The library defines two
 810 public methods: `enlist`, which allows the client to add procedures to be executed by the

library, and `run`, which lets the client run all procedures added so far. The higher-order global reference `list` implements a list of methods.

```

813
814 1 private empty(x:int) : (unit) = { () };
815 2 fun list := empty;
816 3 int cnt := 0; int running := 0;
817 4 public enlist(f:(unit → unit)) :(unit) = {
818 5   if (!running) then ()
819 6   else
820 7     cnt := !cnt + 1;
821 8     (let c = !cnt in let l = !list in
822 9       list := (fun(z:int):(unit) → if (z == c) then f() else l(z)));
823 10 public run(x:unit) :(unit) = {
824 11   running := 1;
825 12   if (0 < !cnt) then
826 13     (!list)(!cnt);
827 14     cnt := !cnt - 1; assert(not (!cnt < 0)); run()
828 15   else (list := empty; running := 0) };
829

```

The bug here is also due to a reentrant call in line 13. However, this is a much tougher example as it involves a higher-order reference `list`, a recursive method `run`, and a second-order method `enlist` that reveals client names to the library. With HOLiK, we obtain the following minimal counterexample:

```

834   call⟨enlist, m1⟩ · ret⟨enlist, ()⟩ · call⟨run, ()⟩ · call⟨m1, ()⟩
835   · call⟨run, ()⟩ · call⟨m1, ()⟩ · ret⟨m1, ()⟩ · ret⟨run, ()⟩ · ret⟨m1, ()⟩
836

```

where m_1 is a client name revealed to the library. In the trace above, `enlist` reveals the method m_1 to the library. This name is then added to the list of procedures to execute. In `run`, the library passes control to the client by calling m_1 . At this point, the client is allowed to call `run` again before the list is updated.

841 **B** Comparison with Racket Contract Verification

842 We shall consider the latest version of the tool [27] since it handles state, which we refer to as
843 SCV (Software Contract Verifier). A small benchmark (19 programs) based on HOLiK and
844 SCV benchmarks was used for testing. Programs were manually translated between HOLi and
845 Racket. Care was taken to translate programs whilst maintaining their semantics: contracts
846 enforcing an input-output relation were translated into HOLi using wrapper functions that
847 define the relation through an if statement. In the other direction, since contracts do not
848 directly access references inside a term, stateful functions were translated from HOLi to
849 return any references we wish to reason about.

850 Table 2 records the comparison. On one hand, HOLiK only found real errors, whereas
851 SCV reported several spurious errors—a third of all errors were spurious. On the other
852 hand, SCV was able to prove total correctness of 3 of the 7 safe files present. SCV also
853 scales much better than HOLiK with respect to program size, which is in exchange of
854 precision. The difference in time for small programs is mainly due to initialisation time.
855 Subtle differences in the nature of each tool can also be observed. e.g., HOLiK reports 1 real
856 error for `ack-simple-e`, whereas SCV reports 2 errors. The difference is because SCV takes
857 into account constraints for integers (e.g. > 0 and $= 0$). More interestingly, for `various`,
858 HOLiK reports 19 ways to reach assertion violations, whereas SCV reports only 6 real ways
859 to violate contracts. The difference is because HOLiK reports paths through the execution

Program	LoC	Traces	Time (s)	LoC	Errors	Time (s)	False Errors
ack	17	0	6.0	9	N/A	2.4	N/A
ack-simple	13	0	6.5	9	0	2.4	0
ack-simple-e	13	1	6.5	9	2	2.5	0
dao	10	0	5.0	15	1	2.6	1
dao-e	16	1	5.5	15	1	2.7	0
dao-various	85	5	22.5	122	10	3.0	5
dao2-e	85	10	23.5	122	10	2.9	0
escape	9	0	5.0	9	0	2.6	0
escape-e	9	2	5.0	10	1	2.7	0
escape2-e	10	14	6.0	10	1	2.7	0
factorial	10	0	5.0	9	0	2.2	0
mc91	12	0	5.0	9	1	2.2	1
mc91-e	12	1	5.0	8	1	2.4	0
mult	14	0	5.0	11	2	2.7	2
mult-e	14	1	5.0	11	2	2.4	0
succ	7	0	5.0	7	1	2.5	1
succ-e	7	1	5.0	7	1	2.8	0
various	116	19	14.0	108	11	6.2	5
total	459	55	140.5	500	45	49.8	15

■ **Table 2** Comparison of HOLiK (left) and SCV (right). N/A is recorded for `ack` as in our attempts SCV crashed due to unknown reasons.

860 tree that reach errors, whereas SCV reports a set of terms that may violate the contracts. For
 861 instance, independently safe methods A and B that may call an unsafe method C would be,
 862 from testing, reported as three valid traces ($call\langle A \rangle \cdot call\langle C \rangle$, $call\langle B \rangle \cdot call\langle C \rangle$ and $call\langle C \rangle$)
 863 by HOLiK. In contrast, SCV reports a single contract violation blaming C . Finally, `ack`
 864 failed to run on SCV due to unknown errors; Racket reported an error internal to the tool.
 865 Further testing proved the file is a valid Racket program that can be executed manually.

866 C ML-like References

867 HOLi has global higher-order references. These are enough for coding all of our examples
 868 and, moreover, allow us to prove completeness (every error has a realising client). We here
 869 present a sketch of how games can be extended with (locally created, scope extruding)
 870 ML-like references, following e.g. [21, 16]. First, the following extension to types and terms
 871 are required.

$$872 \quad \theta ::= \dots \mid \mathbf{ref} \theta \quad M ::= \dots \mid !M \mid \mathbf{ref} M \mid M = M \quad v ::= \dots \mid r$$

874 The term $!M$ allows dereferencing terms M which evaluate to references, while $\mathbf{ref} v$ creates
 875 dynamically a fresh name $r \in \mathbf{Refs}_\theta$ (if $v : \theta$), and the semantic purpose is to update the
 876 store $S \uplus \{r \mapsto v\}$ when evaluating $\mathbf{ref} v$. Note that this allows us to store references to
 877 references, etc. Finally, the construct $M = M$ is for comparing references for name equality.

878 With terms handling general references concretely and symbolically, we extend game
 879 configurations with sets $\mathcal{L}_p, \mathcal{L}_o \subseteq \mathbf{Refs}$ that keep track of reference names disclosed by the
 880 proponent and opponent respectively. References being passed as values means that the
 881 client can update the references belonging to the client, and viceversa. When making a move,
 882 for each reference r they own that is passed, the proponent adds r to \mathcal{L}_p . Passing of names in
 883 a move can be done either by method argument and return value, but also via the common

```

1  global cnt := 0
2  global meth := 0
3  global refi := mi           # for each mi ∈ P
4  global refi := defval       # for each mi ∈ P'
5  global valθ := defval      # for each θ ∈ Θv
6  public mi = λx.           # for each mi ∈ A
7    cnt++; meth:=i; valθ1:=x; oracle ()
8  mi = λx.                 # for each mi ∈ A'
9    cnt++; meth:=i; valθ1:=x; oracle ()
10 oracle = λ().
11   match (!cnt) with      # number of P-moves played so far (max |τ|/2)
12     | i →
13       # if i > 0 and i-th P-move of τ is cr mj(v), with mj : θ1 → θ2, then
14       # - if cr = ret then d = 0 and θ = θ2
15       # - if cr = call then d = j and θ = θ1
16       # diverge if the last P-move played is different from cr mj(v)
17       if not (!meth = d and !valθ  $\hat{=}_{\theta}$  v) then diverge
18       else for mi in fresh(!valθ) do refi := mi
19       # if (i + 1)-th O-move of τ is cr' mk(u), with mk : θ1 → θ2, then
20       # - if cr' = ret then c = 0
21       # - if cr' = call then c = k
22       if c then let x = (!refk)u in # call mk(u)
23         cnt++; meth:=0; valθ2:=x; oracle (); !valθ2
24       else valθ2:=u # return u
25 main = oracle ()

```

■ **Figure 5** The client $C_{\tau, \mathcal{P}, \mathcal{A}}$.

884 part of the store (i.e. via the references known to both players). Similarly, opponent passes
885 names in their moves, which are added to \mathcal{L}_o . Concretely, when the opponent passes control,
886 all references in \mathcal{L}_p are updated with opponent values. Symbolically, the references r are
887 updated with distinct fresh symbolic integers κ if $r \in \mathbf{Refs}_{\text{Int}}$, distinct fresh method names
888 if $r \in \mathbf{Refs}_{\theta_1 \rightarrow \theta_2}$, or to arbitrary reference names if $r \in \mathbf{Refs}_{\mathbf{Refs}_{\theta}}$.

889 D Definability

890 In this section we show that every trace τ in the semantics of a library L has a corresponding
891 good client that realises the same trace in its semantics.

892 Let L be a library with public names \mathcal{P} and abstract names \mathcal{A} . Given a trace τ produced
893 by L , with \mathcal{P}' and \mathcal{A}' respectively the public and abstract names introduced in τ , we set:

$$\begin{aligned}
894 \quad \mathcal{N} &= \mathcal{P} \cup \mathcal{P}' \cup \mathcal{A} \cup \mathcal{A}' \\
895 \quad \Theta_v &= \{\theta \mid \exists m \in \mathcal{N}. m : \theta' \wedge \theta \text{ a syntactic subtype of } \theta'\} \\
896 \quad \Theta_m &= \{\theta \in \Theta \mid \theta \text{ a method type}\}
\end{aligned}$$

898 Note that the above sets are finite, since $\tau, \mathcal{P}, \mathcal{A}$ are finite. We assume a fixed enumeration of
899 $\mathcal{N} = \{m_1, m_2, \dots, m_n\}$. Moreover, for each type θ , we let **defval** _{θ} be a default value, and
900 **diverge** _{θ} a term that on evaluation diverges by infinite recursion. We then construct a client
901 $C_{\tau, \mathcal{P}, \mathcal{A}}$ as in Figure 5.

902 The code is structured as follows.

- 903 1. We start off by defining global references:
- 904 – *cnt* counts the number of P (Library) moves played so far;
 - 905 – *meth* stores an index that records the move made by P: if the move was a return then
 - 906 *meth* stores 0; if it was call to m_i then *meth* stores i ;
 - 907 – each ref_i will store the method $m_i \in \mathcal{P} \cup \mathcal{P}'$, either since the beginning (if $m_i \in \mathcal{P}$),
 - 908 or once P plays it (if $m_i \in \mathcal{P}'$);
 - 909 – each val_θ will be used for storing the value played by P in their last move.

910 In the latter case above, there is a light abuse of syntax as θ can be a product type, of
 911 which HOLi does not have references. But we can in fact simulate references of arbitrary
 912 type by several HOLi references.

- 913 2. For each $m_i : \theta_1 \rightarrow \theta_2 \in \mathcal{A}$, we define a public method m_i that simulates the behaviour
 914 of O whenever m_i is called in τ :

- 915 – it starts by increasing *cnt*, as a call to m_i corresponds to a P-move being played;
- 916 – it continues by storing i and x in *meth* and val_{θ_1} respectively;
- 917 – it calls the private method *oracle*, which is tasked with simulating the rest of τ and
- 918 storing the value that m_i will return in val_{θ_2} ;
- 919 – it returns the value in val_{θ_2} .

- 920 3. For each $m_i : \theta_1 \rightarrow \theta_2 \in \mathcal{A}'$ we produce a method just like above, but keep it private (for
 921 the time being).

- 922 4. The method *oracle* performs the bulk of the computations, by checking that the last
 923 move played by P was the expected one and selecting the next move to play (and playing
 924 it if is a call).

- 925 – The oracle is called after each P-move is played, so it starts with increasing *cnt*.
- 926 – It then performs a case analysis on the value of *cnt*, which above we denote collectively
 927 by assuming the value is i – this notation hides the fact that we have one case for each
 928 of the finitely many values of i .

929 For each such i , the oracle first checks if the previous P-move (if there was one), was
 930 the expected one. If the move was a call, it checks whether the called method was
 931 the expected one (via an appropriate value of d), and also whether the value was the
 932 expected one. Value comparisons ($\stackrel{\Delta}{=}_\theta$) only compare the integer components of θ , since
 933 we cannot compare method names. If this check is successful, the oracle extracts from
 934 u any method names played fresh by P and stores them in the corresponding ref_i .

935 Next, the oracle prepares the next move. If, for the given i , the next move is a call,
 936 then the oracle issues the call, stores the return value of that call, increases *cnt* and
 937 recurs to itself – when the issued call returns, it would be through a P-move. If, on the
 938 other hand, the next move is a return, the oracle simply stores the value to be returned
 939 in the respective *val* reference – this would allow to the respective m_i to return that
 940 value.

- 941 5. The **main** method simply calls the oracle.

942 We can then show the following (proof provided in full version [22]). For any library L
 943 and $(\tau, \rho) \in \llbracket L \rrbracket$, C_τ is such that $(\tau, \rho') \in \llbracket C_\tau \rrbracket$ for some ρ' .