



1 Compositional Correctness and Completeness for 2 Symbolic Partial Order Reduction

3 Åsmund Aqissiaq Arild Kløvstad ✉ 

4 University of Oslo, Norway

5 Eduard Kamburjan ✉

6 University of Oslo, Norway

7 Einar Broch Johnsen ✉ 

8 University of Oslo, Norway


9 — Abstract —

10 Partial Order Reduction (POR) and Symbolic Execution (SE) are two fundamental abstraction
11 techniques in program analysis. SE is particularly useful as a state abstraction technique for sequential
12 programs, while POR addresses equivalent interleavings in the execution of concurrent programs.
13 Recently, several promising connections between these two approaches have been investigated, which
14 result in *symbolic partial order reduction*: partial order reduction of symbolically executed programs.
15 In this work, we provide *compositional* notions of completeness and correctness for symbolic partial
16 order reduction. We formalize completeness and correctness for (1) abstraction over program states
17 and (2) trace equivalence, such that the abstraction gives rise to a complete and correct SE, the trace
18 equivalence gives rise to a complete and correct POR, and their combination results in complete
19 and correct symbolic partial order reduction. We develop our results for a core parallel imperative
20 programming language and mechanize the proofs in Coq.

21 **2012 ACM Subject Classification** Theory of computation → Parallel computing models

22 **Keywords and phrases** Symbolic Execution, Coq, Trace Semantics, Partial Order Reduction

23 **Digital Object Identifier** 10.4230/LIPIcs.CONCUR.2023.23

24 **Supplementary Material** The Coq logo  links directly to the Coq formalization [20] available at
25 <https://github.com/Aqissiaq/symex-formally-formalized>

27 **Acknowledgements** The first author would like to thank Yannick Zakowski for help with Coq
28 formatting and Erik Voogd for valuable insights on symbolic semantics.

29 **1** Introduction

30 Program analyses rely on representing the possible reachable states and traces of a program
31 run efficiently and are commonly accompanied by a correctness theorem (*all representable*
32 *states and traces are reachable*) and possibly a completeness theorem (*all reachable states*
33 *and traces are represented*). Explicitly listing all states or traces leads to the “state space
34 explosion”, as even for simple programs, the number of possible program states may grow so
35 fast that examining them all explicitly becomes infeasible.

36 One source of this growth is the domain of data — the number of possible values is very
37 large, even for a single integer. Symbolic execution [7, 18, 19] (SE) mitigates this problem by
38 representing values symbolically, thus covering many possible concrete states at once. SE is
39 utilized to great effect in program analysis [3]. Another source of growth is *concurrency*, as
40 the number of possible interleavings grows exponentially. Partial Order Reduction (POR)
41 is a technique for tackling this explosion by taking advantage of the fact that independent
42 events can be reordered without affecting the final result [16].



© Åsmund Aqissiaq Arild Kløvstad, Eduard Kamburjan, Einar Broch Johnsen;
licensed under Creative Commons License CC-BY 4.0

34th International Conference on Concurrency Theory (CONCUR 2023).


Editors: John Q. Open and Joan R. Access, Article No. 23; pp. 23:1–23:18

Leibniz International Proceedings in Informatics



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

43 The combined use of both POR and SE has recently begun to be investigated [6, 29],
 44 called *symbolic partial order reduction (SPOR)*. Notions of correctness and completeness are
 45 available for both SE and POR, but how these notions can be composed to obtain correctness
 46 and completeness of SPOR remains an open challenge. In this paper, we tackle this challenge
 47 and give a compositional notion of correctness and completeness for SPOR, based on the
 48 abstraction and equivalence notions that define SE and POR. To formulate such a theory we
 49 use *trace*-based semantics. Trace semantics is both expressive [24, 32] and compositional [11],
 50 and allows a natural formulation of partial order reduction [6].



■ **Figure 1** State of the art and our contribution

51 **State of the Art.** Figure 1 shows the available correctness and completeness results for SE
 52 and for POR. Each corner denotes a program semantics, and the arrows denote correctness
 53 and completeness. First, let us examine the left side of the square, which is concerned with
 54 SE.

55 The left edge of Figure 1, labeled (1), is provided by de Boer and Bonsangue [5], who
 56 define symbolic and concrete semantics for several minimal imperative languages to formulate
 57 and prove notions of correctness and completeness for SE. However, their work is limited to
 58 a *sequential setting*. The proof is based on using a suitable abstraction between concrete and
 59 symbolic states, that defines the SE.

60 The bottom edge of Figure 1, labeled (2), is studied by de Boer et al. [6], who formulate
 61 partial order reduction for symbolic execution with explicit threads using a syntactic notion
 62 of interference freedom and implement this approach in the rewriting logic framework
 63 Maude [10]. Their results are not connected to the concrete semantics. The result is based
 64 on an equivalence relation between symbolic traces, that defines the SPOR, but does not
 65 use an explicit abstraction between states. We discuss further, related results on symbolic
 66 execution in Sec. 6.

67 The top of Figure 1 concerns POR [1, 13, 16, 26] for concrete executions, where numerous
 68 implementations are available. The correctness of such a reduction corresponds to the top
 69 edge of Figure 1, though it is not usually presented in terms of an equivalence relation as
 70 proposed by de Boer et al. Results directly of SPOR are given by Schemmel et al. [29], who
 71 apply (dynamic) partial order reduction to symbolic execution using “unfolding” to explore
 72 paths. This shows that POR is applicable directly to SE, but does not discuss a generic
 73 notion of state abstraction and trace equivalence.

74 While all four corners of Figure 1 are well established, and several edges have been
 75 explored, there exists no general formalization of the properties for state abstraction and
 76 trace equivalence needed for a uniform and compositional treatment of different POR
 77 algorithms and SE techniques. Hence, the present work unifies notions of correctness and
 78 completeness for symbolic execution and partial order reduction, and fills in the remaining
 79 (black) edges of Figure 1. By compositional completeness and correctness, we mean that the
 80 diagonal follows automatically from the other edges of the figure.

81 **Approach.** To fill the gap we formulate concrete and symbolic trace semantics for a

82 small imperative language with parallel composition and show that these semantics enjoy
83 a bisimulation relationship. We then formulate partial order reduction in terms of an
84 equivalence relation on traces, and show that this also leads to a bisimulation of reduced
85 and non-reduced semantics. These bisimulations extend to correctness and completeness
86 results, and compose naturally to semantics with *both* symbolic execution and partial order
87 reduction.

88 The results are obtained in a framework extending the work of de Boer et al. and are
89 centered around the notions of state abstraction and trace equivalence. Following de Boer et
90 al., state abstraction is given by transforming concrete states according to symbolic states,
91 and a concrete state is abstracted if it can be obtained by some symbolic transformation.
92 Trace equivalence defines an equivalence relation on sequences of events which allows for
93 partial order reduction. In particular, it suffices to explore one trace per equivalence class.

94 Both symbolic and concrete states are implemented by total functions of variable names
95 with generic properties. To reduce the number of rules and allow for elegant parallel
96 composition the semantics are given by a reduction system in the style of Felleisen and
97 Hieb [12] with contexts formalized as functions on statements and an inductive relation [21].
98 The full semantics are obtained by stepwise transitive closure, which allows for proofs by
99 induction and case analysis of the final step.

100 **Contributions.** Our contribution is threefold.

- 101 1. We unify and fill in the remaining edges in the above diagram. In particular we give
102 correctness and completeness relations for concrete partial order reduction, directly relate
103 partial order reduction in the symbolic and concrete case, and compose the results to
104 relate concrete semantics to reduced symbolic semantics.
- 105 2. Correctness and completeness for both symbolic execution and partial order reduction
106 are formulated in a parametric fashion, allowing for different implementations of both,
107 providing they fulfill certain conditions.
- 108 3. Finally, the entire development is mechanized in Coq [4,33]. This lends credence to the
109 results and allows for extensions and further work in a systematic manner.

110 **Structure.** Section 2 introduces basic notions for symbolic execution with trace semantics
111 for a basic imperative language with parallel composition. Then both concrete and symbolic
112 semantics are given as reduction systems with contexts to handle both sequential and parallel
113 composition. Finally we formulate and prove correctness and completeness of the symbolic
114 semantics with respect to the concrete semantics. Section 3 introduces a notion of trace
115 equivalence that connects correctness and completeness to partial order reduction, which is
116 used in Section 4 to define independence of events in a semantic manner. We then define new
117 PO-reduced semantics for both symbolic and concrete cases, and show that they bisimulate
118 their non-reduced counterparts. Finally, Section 5 connects previous results and shows that
119 bisimulation carries through POR to fill in the upper right half of the diagram. Section 6
120 and 7 give further related work and concludes.

121 **2 Symbolic Trace Semantics**

122 In this section we introduce the basic notions of our framework. In particular, we define a
123 small imperative language with parallel composition and formulate symbolic and concrete
124 trace semantics for it. We relate the two semantics by a bisimulation defining both trace
125 completeness and trace correctness.

23:4 Compositional Symbolic POR

$$\begin{array}{ll}
 e ::= n \mid x \mid e_1 + e_2 & \text{arith. expr.} \\
 b ::= \text{true} \mid \text{false} \mid \neg b \mid b_1 \wedge b_2 \mid e_1 \leq e_2 & \text{bool. expr.} \\
 s ::= x := e \mid s_1 ; s_2 \mid s_1 \parallel s_2 \mid \text{if } b \{s_1\}\{s_2\} \mid \text{while } b \{s\} \mid \text{skip} & \text{statements}
 \end{array}$$

■ **Figure 2** Grammar for expressions 🍌 and statements 🍌

126 2.1 Basic Notions

127 For the basic setup we assume a set of program variables Var , a set of arithmetic expressions
 128 $Aexpr$ and a set of Boolean expressions $Bexpr$. Our basic programming language is an
 129 imperative language with (side effect free) assignment, conditional branching, iteration and
 130 both sequential and parallel composition.

131 ► **Definition 2.1** (Syntax). *The sets of arithmetic expressions $Aexpr$, Boolean expressions*
 132 *$Bexpr$, and statements $Stmt$ are defined by the grammar in Figure 2, where we let x range*
 133 *over Var , n over \mathbb{N} , b over $Bexpr$, e over $Aexpr$ and s over statements.*

134 Before we define the semantics, we require a notion of store to express program state. We
 135 distinguish between symbolic stores, for symbolic execution, and concrete stores, for concrete
 136 execution.

137 ► **Definition 2.2** (Symbolic Store). *A symbolic store σ is a substitution, i.e., a map from*
 138 *Var to $Aexpr$ denoted by σ .*

139 We take equality of substitutions to be extensional, that is $\sigma = \sigma'$ if $\sigma(x) = \sigma'(x)$ for all
 140 x . An *update* to a substitution is denoted by $\sigma[x := e]$. A substitution can be recursively
 141 applied to a Boolean or arithmetic expression, resulting in a new expression. We denote such
 142 an application by $e\sigma$.

143 ► **Definition 2.3** (Concrete Store). *A concrete store V is a valuation, i.e., a map from Var*
 144 *to \mathbb{N} denoted by V .*

145 Like substitutions, valuations can be updated (denoted $V[x := n]$) and a valuation can be
 146 used to evaluate an expression. This evaluation is denoted $V(e)$ and results in a natural
 147 number for arithmetic expressions and a Boolean for Boolean expressions. For a Boolean
 148 expression b , we say V is a model of b if $V(b) = \text{true}$ and denote this by $V \models b$. The
 149 definitions of substitution and evaluation are standard and given in the auxiliary material.

150 2.2 Trace Semantics

151 Based on the notion of symbolic and concrete stores, we now give the symbolic and concrete
 152 semantics. Both semantics are based on traces, i.e., sequences of *events*. Events are
 153 assignments or guards in the symbolic case, or just assignments in the concrete case.

154 ► **Definition 2.4** (Symbolic Trace). *A symbolic trace is a sequence of conditions or symbolic*
 155 *assignments defined by the grammar*

156
$$\tau_S ::= [] \mid \tau_S :: (x := e) \mid \tau_S :: b$$

157 ► **Definition 2.5** (Concrete Trace). *A concrete trace is a sequence of concrete assignments*
 158 *defined by the grammar*

159
$$\tau_C ::= [] \mid \tau_C :: (x := e)$$

160 In both cases $[]$ denotes the empty trace and we write the trace $[] :: x :: y :: z \dots$ simply as
 161 $[x, y, z \dots]$. The concatenation of τ and τ' is denoted by $\tau \cdot \tau'$. The trace syntax is shared
 162 between symbolic and concrete traces, but the difference will be clear from context.

163 We represent the current program state as a pair of a statement (the program remaining
 164 to be executed) and the trace generated so far. Evaluating expressions requires to evaluate
 165 the expression in the last substitution or valuation of the trace. To do so, we extract this
 166 *final* substitution or valuation from a trace and an initial substitution or valuation by folding
 167 over the trace. In the case of a symbolic trace, the result is a symbolic substitution, while a
 168 concrete trace results in a concrete valuation.

169 ► **Definition 2.6** (Final Substitution 🍷). *Given an initial substitution σ , the final substitution*
 170 *of a trace τ_S is denoted $\tau_S \Downarrow_\sigma$ and inductively defined by*

$$\begin{aligned} 171 \quad [] \Downarrow_\sigma &= \sigma \\ 172 \quad \tau_S :: b \Downarrow_\sigma &= \tau_S \Downarrow_\sigma \\ 173 \quad \tau_S :: (x := e) \Downarrow_\sigma &= \sigma'[x := (e\sigma')] \text{ where } \sigma' = \tau_S \Downarrow_\sigma \end{aligned}$$

175 When $\sigma = id$ we omit it and write $\tau_S \Downarrow$

176 ► **Definition 2.7** (Final Valuation 🍷). *Given an initial valuation V , the final valuation of a*
 177 *trace τ_C is denoted $\tau_C \Downarrow_V$ and inductively defined by*

$$\begin{aligned} 178 \quad [] \Downarrow_V &= V \\ 179 \quad \tau_C :: (x := e) \Downarrow_V &= V'[x := V'(e)] \text{ where } V' = \tau_C \Downarrow_V \end{aligned}$$

181 Semantics can then be given by a simple reduction relation on atomic statements (Figure 3),
 182 which extends to the full language by S/C-IN-CONTEXT. The symbolic (resp. concrete)
 183 relation works on pairs of statements and symbolic (resp. concrete) traces to extend them
 184 with appropriate events.

185 ► **Definition 2.8** (Symbolic and Concrete Semantics). *The symbolic semantics \rightarrow between two*
 186 *symbolic configurations is given on the left of Fig. 3. The concrete semantics \Rightarrow between two*
 187 *concrete configurations is given on the right of Fig. 3.*

188 Both semantics are straightforward, we point out three details. First, the main difference is
 189 that the rules with branching (*-IF-T, *-IF-F, *-WHILE-T, *-WHILE-F) are non-deterministic
 190 and add an event in the symbolic case, but are deterministic in the concrete case.

191 Second, in order to concisely deal with both sequential and parallel composition, we use
 192 *contexts* [12]. A context C represents a statement with a “hole” (\square) in it and is generated by
 193 the grammar:

$$194 \quad C ::= \square \mid (C ; s) \mid (C \parallel s) \mid (s \parallel C)$$

195 Intuitively, the statement we are interested in may occur on its own, sequentially before some
 196 other statement, or on either side of a parallel operator. By $C[s]$ we denote the statement s
 197 in the hole in context C .

198 Finally, we point out that we model termination by reduction to **skip**.

199 ► **Example 2.9.** 🍷 Consider the program $s = y := 1 \parallel x := 3 \parallel \text{if } X \leq 1 \{ Y := 2 \} \{ Y := 3 \}$.
 200 We will show that $(s, []) \rightarrow^* (\text{skip}, [x := 3, y := 1, x > 1, y := 3])$. In other words that
 201 $[x := 3, y := 1, x > 1, y := 3]$ is one possible trace of the program.

23:6 Compositional Symbolic POR

$\text{S-ASGN} \frac{}{(x := e, \tau) \rightsquigarrow (\text{skip}, \tau :: (x := e))}$	$\frac{}{(x := e, \tau) \rightsquigarrow_V (\text{skip}, \tau :: (x := e))} \text{C-ASGN}$
$\text{S-IF-T} \frac{}{(\text{if } b \{s_1\}\{s_2\}, \tau) \rightsquigarrow (s_1, \tau :: b)}$	$\frac{\tau \Downarrow_V (b) = \text{true}}{(\text{if } b \{s_1\}\{s_2\}, \tau) \rightsquigarrow_V (s_1, \tau)} \text{C-IF-T}$
$\text{S-IF-F} \frac{}{(\text{if } b \{s_1\}\{s_2\}, \tau) \rightsquigarrow (s_2, \tau :: \neg b)}$	$\frac{\tau \Downarrow_V (b) = \text{false}}{(\text{if } b \{s_1\}\{s_2\}, \tau) \rightsquigarrow_V (s_2, \tau)} \text{C-IF-F}$
$\text{S-WHILE-T} \frac{}{(\text{while } b \{s\}, \tau) \rightsquigarrow (s ; \text{while } b \{s\}, \tau :: b)}$	$\frac{\tau \Downarrow_V (b) = \text{true}}{(\text{while } b \{s\}, \tau) \rightsquigarrow_V (s ; \text{while } b \{s\}, \tau)} \text{C-WHILE-T}$
$\text{S-WHILE-F} \frac{}{(\text{while } b \{s\}, \tau) \rightsquigarrow (\text{skip}, \tau :: \neg b)}$	$\frac{\tau \Downarrow_V (b) = \text{false}}{(\text{while } b \{s\}, \tau) \rightsquigarrow_V (\text{skip}, \tau)} \text{C-WHILE-F}$
$\text{S-SEQ} \frac{}{(\text{skip} ; s, \tau) \rightsquigarrow (s, \tau)}$	$\frac{}{(\text{skip} ; s, \tau) \rightsquigarrow_V (s, \tau)} \text{C-SEQ}$
$\text{S-PAR} \frac{}{(\text{skip} \parallel \text{skip}, \tau) \rightsquigarrow (\text{skip}, \tau)}$	$\frac{}{(\text{skip} \parallel \text{skip}, \tau) \rightsquigarrow_V (\text{skip}, \tau)} \text{C-PAR}$
$\text{S-IN-CONTEXT} \frac{(s, \tau) \rightsquigarrow (s', \tau')}{(C[s], \tau) \rightsquigarrow (C[s'], \tau')}$	$\frac{(s, \tau) \rightsquigarrow_V (s', \tau')}{(C[s], \tau) \Rightarrow_V (C[s'], \tau')} \text{C-IN-CONTEXT}$

■ **Figure 3** Reduction rules for symbolic and concrete semantics

202 First apply S-IN-CONTEXT with $C = y := 1 \parallel \square \parallel \text{if } X \leq 1 \{Y := 2\} \{Y := 3\}$ and S-ASGN
 203 to obtain

$$204 \quad (s, []) \rightarrow (y := 1 \parallel \text{skip} \parallel \text{if } X \leq 1 \{Y := 2\} \{Y := 3\}, [x := 3])$$

205 The second assignment is similar, followed by S-IF-F in the context $\text{skip} \parallel \text{skip} \parallel \square$ to obtain

$$206 \quad (\text{skip} \parallel \text{skip} \parallel \text{if } X \leq 1 \{Y := 2\} \{Y := 3\}, [x := 3, y := 1]) \rightarrow (\text{skip} \parallel \text{skip} \parallel Y := 3, [x := 3, y := 1, x > 1])$$

207 After the last assignment, the superfluous `skips` are dispensed with by S-PAR and putting
 208 the steps in sequence gives the desired

$$209 \quad (s, []) \rightarrow^* (\text{skip}, [x := 2, y := x, z := x])$$

210 Note that we could choose to apply the contexts in a different order, resulting in five
 211 other potential traces.

212 2.3 Correctness and Completeness

213 The value of symbolic execution comes from its ability to simultaneously capture many
 214 possible concrete execution paths. However, not all of these paths will be feasible for all initial
 215 valuations. The feasibility of any particular symbolic trace depends on its *path condition* —
 216 a conjunction of guards that allow execution to follow down this particular path — which is
 217 computed in a similar fashion to final substitutions.

218 ► **Definition 2.10** (Path Condition). *The path condition of a symbolic trace τ_S is denoted*
 219 *$pc(\tau_S)$ and defined by*

$$220 \quad pc([]) = \text{true}$$

$$221 \quad pc(\tau_S :: b) = pc(\tau_S) \wedge b(\tau_S \Downarrow)$$

$$222 \quad pc(\tau_S :: (x := e)) = pc(\tau_S)$$

224 Because it is a conjunction of terms, once a path condition becomes false, it cannot
 225 become true again. The following lemma captures the contrapositive: a model of a trace's
 226 path condition is also a model of any prefix's path condition.

227 ► **Lemma 2.11** (Path Condition Monotonicity 🍷). *If $V \models pc(\tau :: ev)$, then $V \models pc(\tau)$*

228 To relate the symbolic and concrete traces we define a notion of abstraction based on the
 229 correctness and completeness relations of de Boer and Bonsangue.

230 ► **Definition 2.12** (Trace abstraction [5] 🍷). *Given an initial valuation V , a symbolic trace
 231 τ_S and a concrete trace τ_C we say τ_S V -abstracts τ_C if $V \models pc(\tau_S)$ and $\tau_C \Downarrow_V = V \circ \tau_S \Downarrow$*

232 The steps of the symbolic and concrete systems correspond very closely. Every concrete
 233 step corresponds to a symbolic step whose path condition is satisfiable, and every symbolic
 234 step with a satisfiable path condition corresponds to a concrete step. In both cases the
 235 resulting final states are related by simple composition. This relationship is formalized in
 236 the following bisimulation result.

237 ► **Theorem 2.13** (Bisimulation 🍷). *For any initial valuation V and initial traces τ_0, τ'_0 such
 238 that τ_0 V -abstracts τ'_0 :*

- 239 ■ *if there is a concrete step $(s, \tau_0) \Rightarrow_V (s', \tau)$, then there exists a symbolic step $(s, \tau'_0) \rightarrow (s', \tau')$
 240 such that τ' V -abstracts τ , and*
- 241 ■ *if there is a symbolic step $(s, \tau'_0) \rightarrow (s', \tau')$ and $V \models pc(\tau')$, then there exists a concrete
 242 step $(s, \tau_0) \Rightarrow_V (s', \tau)$ such that $\tau \Downarrow_V = V \circ \tau' \Downarrow$*

243 By induction over the transitive closure and Lemma 2.11 we obtain correctness and
 244 completeness results. Intuitively, correctness means that each symbolic execution whose path
 245 condition is satisfied by some initial valuation V corresponds to a concrete execution with
 246 the same initial valuation. Additionally its trace abstracts the concrete trace in the sense
 247 that the final concrete state is the concretization of V by the final symbolic state. In other
 248 words the subset of states described by its path condition contains V , and there is a concrete
 249 execution corresponding to the transformation described by its final symbolic state.

250 ► **Corollary 2.14** (Trace Correctness 🍷). *If $(s, \tau_S) \rightarrow^* (s', \tau'_S)$, τ_S V -abstracts τ_C , and
 251 $V \models pc(\tau'_S)$, then there exists a concrete trace τ'_C s.t. $(s, \tau_C) \Rightarrow_V^* (s', \tau'_C)$ and $\tau'_C \Downarrow_V = \tau_C \Downarrow_V \circ (\tau'_S \Downarrow)$*

252 Completeness captures the opposite relationship: every concrete execution has a symbolic
 253 counterpart. Furthermore the symbolic trace recovers the concrete state, and its path
 254 condition is satisfied by the initial valuation.

255 ► **Corollary 2.15** (Trace Completeness 🍷). *If $(s, \tau_C) \Rightarrow_V^* (s', \tau'_C)$ and τ_S V -abstracts τ_C ,
 256 there exist τ'_S s.t. $(s, \tau_S) \rightarrow^* (s', \tau'_S)$ and τ'_S V -abstracts τ'_C .*

257 **3 Trace Equivalence**

258 In this section we introduce a notion of trace equivalence which will be used to formulate
 259 partial order reduction in Section 4. Intuitively two traces should be equivalent if execution
 260 could continue from either one, i.e., if partial order reduction would prune away one of them.

261 This is surely the case when their final states are the same. In the symbolic case their
 262 path conditions must also be equivalent. Additionally, we do not want to equate traces
 263 describing observably different behavior, so equivalent traces must contain the same events.
 264 These considerations motivate the following definition.

265 ► **Definition 3.1** (Symbolic Trace Equivalence 🍷). *Symbolic traces τ and τ' are equivalent*
 266 *(denoted $\tau \sim \tau'$) if*

- 267 ■ τ' is a permutation of τ ,
- 268 ■ $\tau \Downarrow_\sigma = \tau' \Downarrow_\sigma$ for all initial substitutions σ , and
- 269 ■ $V \models pc(\tau) \iff V \models pc(\tau')$ for all valuations V

270 ► **Definition 3.2** (Concrete Trace Equivalence 🍷). *Concrete traces τ and τ' are equivalent*
 271 *(denoted $\tau \simeq \tau'$) if*

- 272 ■ τ' is a permutation of τ ,
- 273 ■ $\tau \Downarrow_V = \tau' \Downarrow_V$ for all initial valuations V

274 ► **Example 3.3.** Let $\tau_1 = [y := x, z := x]$ and $\tau_2 = [z := x, y := x]$. It is both the case
 275 that $\tau_1 \sim \tau_2$ and $\tau_1 \simeq \tau_2$.¹ They evidently contain the same events and have the same
 276 (trivially true) path condition. Any initial substitution σ results in a final substitution

$$277 \sigma'(v) = \begin{cases} x, & v \in \{y, z\} \\ \sigma(v), & \text{otherwise} \end{cases} \quad \text{and any initial valuation } V \text{ results in } V'(v) = \begin{cases} V(x), & v \in \{y, z\} \\ V(v), & \text{otherwise} \end{cases}$$

278 Clearly, trace equivalence defines an equivalence relation. Furthermore it allows continued
 279 execution in the following sense: given a statement s and a trace τ , we can replace τ with
 280 an equivalent trace τ' , such that the next execution step will result in two different, but
 281 equivalent traces.

282 ► **Lemma 3.4** (🍷). *For equivalent traces $\tau \sim \tau'$, if $(s, \tau) \rightarrow (s', \tau_1)$ then there exists τ_2 such*
 283 *that $(s, \tau') \rightarrow (s', \tau_2)$ and $\tau_1 \sim \tau_2$.*

284 This lemma also holds for concrete traces with concrete equivalence and reduction system
 285 and underlies partial order reduction in both cases.

286 Crucially, the properties of trace equivalence ensure that it preserves abstraction. The
 287 following theorem shows that the notion of V-abstraction carries through trace equivalence,
 288 which will allow us to connect it with partial order reduction in the sequel.

289 ► **Theorem 3.5** (Abstraction Congruence 🍷). *For equivalent symbolic traces $\tau_S \sim \tau'_S$ and*
 290 *concrete traces $\tau_C \simeq \tau'_C$, if τ_S V-abstracts τ_C then τ'_S V-abstracts τ'_C*

291 ► **Example 3.6.** Continuing Example 3.3, the symbolic trace τ_1 V-abstracts the concrete
 292 trace τ_1 for every V , and so τ_1 also V-abstracts the equivalent concrete trace τ_2 .

293 In fact, every symbolic trace V-abstracts itself viewed as a concrete trace for any V .

294 3.1 Example: Interference Freedom

295 The reordering of independent events is the core of many POR approaches. In practice true
 296 independence is prohibitively expensive to compute, so some over-approximation is used.
 297 Interference freedom is a syntactic over-approximation of independence of events. We show
 298 that reordering interference free events is an instance of our notion of trace equivalence.

299 Interference freedom between ev_1 and ev_2 means that ev_1 does not read or write a variable
 300 written by ev_2 and vice versa. Formally:

¹ Recall that symbolic traces are also concrete traces if they contain no branching events (guards).

301 ► **Definition 3.7** (Interference Freedom). *Let ev be either a Boolean expression b or an*
 302 *assignment $(x := e)$. $R(ev)$ denotes the set of variables read by ev , ie. all the variables in b*
 303 *or e . $W(ev)$ denotes the set of variables written by ev , ie. x . Then ev_1, ev_2 are interference*
 304 *free iff*

$$305 \quad W(ev_1) \cap W(ev_2) = R(ev_1) \cap W(ev_2) = R(ev_2) \cap W(ev_1) = \emptyset$$

306 *Denote the interference freedom of ev_1 and ev_2 by $ev_1 \diamond ev_2$*

307 Interference freedom is an independence relation in the sense that if $ev_1 \diamond ev_2$, then the
 308 final state of $[ev_1, ev_2]$ is equal to that of $[ev_2, ev_1]$. The reason is that interference freedom
 309 allows for “simultaneous” updates without worrying about the order of operations in the
 310 assignment case, and the variables involved in a Boolean expression can not be changed in
 311 the guard case.

312 On the other hand, interference freedom is an over-approximation which is perhaps most
 313 easily seen by events like $(x := x)$ and $(x \leq 3)$. Clearly they are semantically independent
 314 since the value of x does not change, but they are not interference free.

315 Equipped with a concrete independence relation we can construct new traces by reordering
 316 adjacent independent events. Such a reordering is captured by the equivalence define above
 317 in the sense that it results in an equivalent trace.

318 ► **Theorem 3.8** (Interference free reordering is a trace equivalence 🍷). *Let \sim_{IF} be the smallest*
 319 *equivalence relation on symbolic traces such that $\tau \cdot [ev_1, ev_2] \cdot \tau' \sim_{IF} \tau \cdot [ev_2, ev_1] \cdot \tau'$ for all*
 320 *τ, τ' and $ev_1 \diamond ev_2$.*

321 *The equivalence relation \sim_{IF} is contained in \sim .*

322 The analogous result holds for concrete traces and \simeq 🍷.

323 This example shows that a POR scheme based on reordering of independent events is
 324 captured by trace equivalence.

325 **4 Correctness and Completeness for Symbolic Partial Order** 326 **Reduction**

327 We formulate POR in the present setting through the use of trace equivalence (defined above)
 328 and use it to define new PO-reduced reduction systems. These new systems bisimulate the
 329 non-reduced systems of Section 2, leading directly to correctness and completeness results.

330 At its core, partial order reduction works by observing that some events commute in
 331 the execution of a parallel program. These events can be reordered without affecting the
 332 final result, and so it is not necessary to explore *every* interleaving. The reduction is often
 333 formulated in terms of an (in)dependence relation that determines which events may be
 334 reordered. Such a relation must make sure that independent steps leave the system in
 335 equivalent states, regardless of the order they are performed in.

336 An independence relation lifts to an equivalence relation on traces by permuting adjacent
 337 independent events. POR approaches then employ some algorithm to compute the equivalence
 338 classes of such a relation and avoid exploring traces in the same class. In practice it is difficult
 339 to compute the independence of events, so a sound over-approximation is used instead.

340 We instead take a more high-level approach. Considering trace equivalence to be a
 341 fundamental semantic building block, we develop our POR semantics parametric in this
 342 notion. This gives us an abstract notion, independent of the specific algorithm for POR.

343 To take advantage of partial order reduction, we define new transition systems.

23:10 Compositional Symbolic POR

344 ► **Definition 4.1** (POR Semantics). *The transition rules for symbolic POR are:*

$$345 \frac{\tau_0 \sim \tau'_0 \quad (s, \tau_0) \rightsquigarrow (s', \tau)}{(s, \tau'_0) \rightsquigarrow_{POR} (s', \tau)} \quad \frac{(s, \tau) \rightsquigarrow_{POR} (s', \tau')}{(C[s], \tau) \rightarrow_{POR} (C[s'], \tau')}$$

346 *And the transition rules for concrete POR are:*

$$347 \frac{\tau_0 \simeq \tau'_0 \quad (s, \tau_0) \rightsquigarrow_V (s', \tau)}{(s, \tau'_0) \rightsquigarrow_{POR, V} (s', \tau)} \quad \frac{(s, \tau) \rightsquigarrow_{POR, V} (s', \tau')}{(C[s], \tau) \Rightarrow_{POR, V} (C[s'], \tau')}$$

348 This new reduction relation includes the steps of the symbolic case but requires only that
 349 the initial trace is *equivalent* in the sense defined in Section 3. Crucially, given a class of
 350 equivalent traces we may choose only one of them to continue execution. This is the source
 351 of *reduction*. Note that it is possible for (s, τ'_0) to be unreachable in the original semantics,
 352 however the following completeness and correctness results ensure that this does not affect
 353 the final result. This approach most closely resembles *sleep sets* [15, 17] which keeps track of
 354 equivalent traces that do not need to be explored.

355 ► **Example 4.2.** Consider again the program from Example 2.9 and note that $(y := 1)$
 356 and $(x := 3)$ are independent assignments. In the middle of some computation we are left
 357 with $\text{skip} \parallel \text{skip} \parallel \text{if } x \leq 1 \{Y := 2\}\{Y := 3\}$ and the trace $[x := 3, y := 1]$. However, we
 358 have previously explored a computation from the state

$$359 (\text{skip} \parallel \text{skip} \parallel \text{if } x \leq 1 \{Y := 2\}\{Y := 3\}, [y := 1, x := 3])$$

360 Now the POR semantics let us replace the equivalent traces and use this computation instead.

361 In order to utilize POR, we need to know that the reduced traces still model our programs'
 362 behavior. It should not throw away any important traces, nor should it invent new ones by
 363 taking unsound equivalence classes. Formally, we want the POR semantics to bisimulate
 364 their non-reduced counterpart up to trace equivalence.

365 ► **Theorem 4.3** (POR bisimulation). *For equivalent initial traces $\tau_0 \sim \tau'_0$:*

- 366 ■ *If $(s, \tau_0) \rightarrow_{POR} (s', \tau)$ then there exists $(s, \tau'_0) \rightarrow (s', \tau')$ such that $\tau \sim \tau'$, and*
- 367 ■ *If $(s, \tau_0) \rightarrow (s', \tau)$ then there exists $(s, \tau'_0) \rightarrow_{POR} (s', \tau')$ such that $\tau \sim \tau'$*

368 *For equivalent initial traces $\tau_0 \simeq \tau'_0$ and initial valuation V :*

- 369 ■ *If $(s, \tau_0) \Rightarrow_{POR, V} (s', \tau)$ then there exists $(s, \tau'_0) \Rightarrow_V (s', \tau')$ such that $\tau \simeq \tau'$, and*
- 370 ■ *If $(s, \tau_0) \Rightarrow_V (s', \tau)$ then there exists $(s, \tau'_0) \Rightarrow_{POR, V} (s', \tau')$ such that $\tau \simeq \tau'$*

371 From these bisimulation results, correctness and completeness follow by induction. Cor-
 372 rectness captures the intuition that every PO-reduced execution corresponds to a non-reduced
 373 execution with equivalent final traces. This means that partial order reduction is precise in
 374 the sense that it does not introduce new traces with different final states.


375 Completeness is the opposite relationship: every direct execution has a corresponding
 376 reduced execution with equivalent traces. Since equivalent traces result in the same final
 377 state, completeness means that we do not lose any possible states when performing partial
 378 order reduction.

379 ► **Corollary 4.4** (Correctness and Completeness). *For two equivalent symbolic traces $\tau_0 \sim \tau'_0$:*

- 380 **Completeness** *If $(s, \tau_0) \rightarrow^*_{POR} (s', \tau)$ then there exists $(s, \tau'_0) \rightarrow^* (s', \tau')$ with $\tau \sim \tau'$*
- 381 **Correctness** *If $(s, \tau_0) \rightarrow^* (s', \tau)$ then there exists $(s, \tau'_0) \rightarrow^*_{POR} (s', \tau')$ with $\tau \sim \tau'$*

382 *For two equivalent concrete traces $\tau_0 \simeq \tau'_0$ and initial valuation V :*

- 383 **Completeness** *If $(s, \tau_0) \Rightarrow^*_{POR, V} (s', \tau)$ then there exists $(s, \tau'_0) \Rightarrow^*_V (s', \tau')$ with $\tau \simeq \tau'$*
- 384 **Correctness** *If $(s, \tau_0) \Rightarrow^*_V (s', \tau)$ then there exists $(s, \tau'_0) \Rightarrow^*_{POR, V} (s', \tau')$ with $\tau \simeq \tau'$*



■ **Figure 4** Overview of the correctness and completeness results

5 Composition of SE and POR

In this section we show that the bisimulation results of Section 2 and 4 compose naturally. We use this composition to fill in the remaining edges of Fig. 1, resulting in Fig. 4. This leads to the main result: a bisimulation relation between direct concrete semantics and symbolic POR semantics. Importantly, this allows reasoning about program analysis using *both* SE and POR with the symbolic trace abstracting the concrete trace.

The results are parametric in abstraction and trace equivalence in the following sense. Any equivalence relation on traces which is contained in ours — that is, whose equivalent traces have equivalent final states and path conditions — can be used to perform partial order reduction. Additionally, any symbolic abstraction satisfying Theorem 3.5 can be used for the symbolic execution. The result is a complete and correct *symbolic partial order reduction* where completeness and correctness follows from the respective completeness and correctness results of SE and POR semantics.

First we relate symbolic and concrete POR by combining Theorem 2.13 and Theorem 4.3.

► **Theorem 5.1** (POR-POR Bisimulation 🍷). *For initial traces τ_S, τ_C such that τ_S V -abstracts τ_C :*

- *If $(s, \tau_C) \Rightarrow_{\text{POR},V} (s', \tau'_C)$, then there exists $(s, \tau_S) \rightarrow_{\text{POR}} (s', \tau'_S)$ such that τ'_S V -abstracts τ'_C*
- *If $(s, \tau_S) \rightarrow_{\text{POR}} (s', \tau'_S)$ and $V \models pc(\tau'_S)$, then there exists $(s, \tau_C) \Rightarrow_{\text{POR},V} (s', \tau'_C)$ and $\tau'_C \Downarrow_V = V \circ (\tau'_S \Downarrow)$*

From this bisimulation, correctness and completeness relations are obtained by induction. These results are analogous to the direct relationships in Section 2, which shows that the correctness and completeness of symbolic execution is maintained through partial order reduction. In particular we may work with representatives of an *equivalence class* of traces rather than one single trace — which may greatly reduce the state space — and then perform symbolic execution in this new setting.

► **Corollary 5.2** (Trace POR Correctness 🍷). *If $(s, \tau_S) \rightarrow_{\text{POR}}^* (s', \tau'_S)$, τ_S V -abstracts τ_C , and $V \models pc(\tau'_S)$, then there exists a concrete trace τ'_C s.t. $(s, \tau_C) \Rightarrow_{\text{POR},V}^* (s', \tau'_C)$ and $\tau'_C \Downarrow_V = \tau_C \Downarrow_V \circ (\tau'_S \Downarrow)$*

► **Corollary 5.3** (Trace POR Completeness 🍷). *If $(s, \tau_C) \Rightarrow_{\text{POR},V}^* (s', \tau'_C)$ and τ_S V -abstracts τ_C , there exist τ'_S s.t. $(s, \tau_S) \rightarrow_{\text{POR}}^* (s', \tau'_S)$ and τ'_S V -abstracts τ'_C .*

We are now ready to state our main result, filling in the diagonal and connecting concrete semantics directly to PO-reduced symbolic semantics. Formally, Theorem 2.13 and

418 Theorem 4.3 can be combined to obtain bisimulation of the basic concrete semantics and
 419 PO-reduced symbolic semantics.

420 ► **Theorem 5.4** (Total Bisimulation 🍷). *For initial traces τ_S, τ_C such that τ_S V -abstracts τ_C :*
 421 ■ *If $(s, \tau_C) \Rightarrow_V (s', \tau'_C)$, then there exists $(s, \tau_S) \rightarrow_{POR} (s', \tau'_S)$ such that τ'_S V -abstracts*
 422 *τ'_C*
 423 ■ *If $(s, \tau_S) \rightarrow_{POR} (s', \tau'_S)$ and $V \models pc(\tau'_S)$, then there exists $(s, \tau_C) \Rightarrow_V (s', \tau'_C)$ and*
 424 *$\tau'_C \Downarrow_{V=V} (\tau'_S \Downarrow)$*

425 ► **Corollary 5.5** (Total Correctness 🍷). *If $(s, \tau_0) \rightarrow_{POR}^* (s', \tau)$, τ_0 V -abstracts τ'_0 and $V \models pc(\tau)$,*
 426 *then there exists τ' such that $(s, \tau'_0) \Rightarrow_V^* (s', \tau')$ and τ V -abstracts τ' .*

427 ► **Corollary 5.6** (Total Completeness 🍷). *If $(s, \tau_0) \Rightarrow_V^* (s', \tau)$ and τ'_0 V -abstracts τ_0 , there*
 428 *exist τ' s.t $(s, \tau'_0) \rightarrow_{POR}^* (s', \tau')$ and τ' V -abstracts τ .*

429 Figure 4 shows all four reduction systems — symbolic and concrete, with and without
 430 POR. Each double arrow denotes a notion of bisimulation, and we obtain the properties
 431 shown: both symbolic and concrete traces are equivalent across POR, and V -abstraction is
 432 maintained across the symbolic/concrete divide as well as their composition. Additionally we
 433 show the relationships between the four traces — the symbolic traces abstract their concrete
 434 counterparts, and the POR traces are equivalent — although by Theorem 3.5 it suffices to
 435 know the equivalences and one of the abstractions.

436 5.1 Discussion

437 The bisimulations compose naturally. As an example, consider Theorem 5.4 which is obtained
 438 by composing the symbolic/concrete bisimulation of Theorem 2.13 and the direct/reduced
 439 bisimulation of Theorem 4.3. Starting with a concrete execution with trace τ_C we first obtain
 440 a symbolic execution with trace τ_S such that τ_S V -abstracts τ_C . Then the POR-bisimulation
 441 of Theorem 4.3 gives a symbolic POR-computation with an equivalent trace τ_S . Since trace
 442 equivalence is a congruence for abstraction (Theorem 3.5) and τ_C is equivalent to itself, this
 443 final trace also abstracts τ_C .

444 The ease of this composition is not unexpected, since both abstraction and trace equivalence
 445 were explicitly formulated to preserve the relevant parts of the program state. The result
 446 is that any partial order reduction which picks equivalent traces in this sense preserves the
 447 correctness and completeness properties of the symbolic execution. Explicitly, if the notion
 448 of trace equivalence is contained in ours and the symbolic abstraction can be transported
 449 along this equivalence in the sense of Theorem 3.5 then the techniques can be composed.

450 5.2 Mechanization

451 In this section we cover some of the details of the mechanization in Coq.

452 The basic building blocks of program state are simple. Both substitutions and valuations
 453 are implemented as total maps from strings, parameterized by a result type. Updates,
 454 notation and several useful lemmas about maps can be proven generically and the notation
 455 mirrors that of Pierce et al. [25]. Similarly traces are an inductive type, parametric in the
 456 type of events. In essence they are lists, but extended to the right for convenience, with the
 457 expected operations and properties.

458 Trace *equivalence* is defined as a relation. Then we show necessary properties of this
 459 relation, in particular Lemma 3.4 and Theorem 3.5 which are used in proofs. Additionally, we
 460 implement an equivalence by permuting independent events and show that it satisfies the same

461 properties if the independence relation does. This part is parametric in the independence
 462 relation and serves as an example of a POR relation. The example at the end of Section 3 is
 463 an instance with interference freedom as the independence relation 🗨️.

464 Expressions (both arithmetic and Boolean) and statements are inductive types. As an
 465 example, the type of statements is given by:

```
466 Inductive Stmt : Type :=
467   | SAsgn (x:Var) (e:Aexpr)
468   | SPar (s1 s2:Stmt)
469   | SIf (b:Bexpr) (s1 s2:Stmt)
470   | ...
471
```

473 To give semantics to this language, we define a *head reduction* relation and a type of
 474 contexts. The head reduction describes the single step reductions for each atomic and how it
 475 transforms the current trace. For example an assignment reduces to `skip` and appends the
 476 assignment to the current trace. Here `<{ _ }>` encloses language statements and `Asgn__S x e`
 477 represents the symbolic event $(x := e)$.

```
478 Variant head_red__S: (trace__S * Stmt) → (trace__S * Stmt) → Prop :=
479   | head_red_asgn__S: ∀t x e,
480     head_red__S (t, <{ x := e }>) (t :: Asgn__S x e, SSkip)
481   | ...
482
483
```

484 Note that `Variant` is a version of `Inductive` that does not include recursive constructors.

485 Contexts are implemented as functions `Stmt → Stmt` along with an inductive relation
 486 `is_context: (Stmt → Stmt) → Prop` — an approach inspired by Xavier Leroy [21]. This
 487 approach allows us to define transition relation semantics parametric in both the type of
 488 contexts and the head reduction relation. The following generalizes the `*-IN-CONTEXT` rules
 489 for any type of state `X`. In our case, `X` will be a type of traces, but note that `X` appears on the
 490 left — this makes the rule amenable to states represented by product types due to the way
 491 parentheses associate.

```
492 Variant context_red
493   (is_cont: (Stmt → Stmt) → Prop) (head_red: relation (X * Stmt))
494   : relation (X * Stmt) :=
495   | ctx_red_intro: ∀C x x' s s',
496     head_red (x, s) (x', s') → is_context C →
497     context_red is_cont head_red (x, C s) (x', C s').
498
499
```

500 Having used `context_red` with the appropriate `is_context` and `head_red` we obtain the full
 501 transition relation by stepwise reflexive-transitive closure to the right (`clos_refl_trans_n1`)
 502 from the `Relations` library.

503 The proofs are performed in two steps. Induction on the transition relation leaves us
 504 with either a reflexive step or an induction hypothesis and some sequence followed by a
 505 step. Then unfolding and dependent destruction (from `Program.Equality`) can be used on the
 506 step to unpack `ctx_red_intro` and split on the head reduction rule while remembering the
 507 ultimate and penultimate traces.

508 6 Related Work

509 We focus on a simple formal model that permits reasoning about symbolic execution and
 510 partial order reduction. De Boer and Bonsangue [5] lay the foundations of our work — a
 511 symbolic execution model based on transition systems and symbolic substitutions which
 512 may be composed with concrete valuations. They do not consider parallelism, but do apply

513 their model to languages with other features including recursive function calls and dynamic
 514 object creation. They also explore a kind of trace semantics for the latter extension, but it
 515 differs from the semantics considered herein. Extending the current work with more language
 516 features, including procedure calls and synchronization tools would be interesting.

517 SymPaths [6] explores the use of POR for SE in a manner very similar to ours, but
 518 does not explicitly compose the correctness and completeness of SE and POR, nor treat
 519 the relationship to partial order reduction in the non-symbolic case. Additionally, their
 520 treatment of trace equivalence focuses on one specific independence relation while we take a
 521 more abstract view.

522 Other formal approaches to symbolic execution have also been considered in the literature.
 523 Steinhöfel [31] focuses on the semantics of the SE system and uses a concretization function to
 524 relate sets of symbolic and concrete states. The Gillian platform [14,23] and related work [28]
 525 uses separation logic to construct a SE system that is parametric in the target memory model.
 526 Rosu et al. [22,27,30] develop reachability logic to present symbolic execution parameterized
 527 by the semantics of the target language. These all present alternative approaches to the left
 528 edge of Figure 4.

529 There are also other approaches to partial order reduction. In particular, dynamic or
 530 stateless POR (DPOR) [1,13,16,26] avoids exploring equivalent future traces by identifying
 531 backtracking points. Additionally the *unfolding* approach explores partial orders more directly
 532 as a tree-like event structure [26]. Unfolding has been fruitfully combined with symbolic
 533 execution in practice [29].

534 **7 Conclusion**

535 POR and SE are fundamental abstraction techniques in program analysis. SE is particu-
 536 larly useful as a state abstraction technique for sequential programs, while POR addresses
 537 equivalent interleavings in the execution of concurrent programs. In this paper, we study
 538 the foundations of both techniques based on transition systems and trace semantics, in the
 539 context of a core imperative language with parallelism. The formalization provides a unified
 540 view of concrete and symbolic semantics with and without partial order reduction. We
 541 further formalize correctness and completeness relations for both POR and SE, and compose
 542 these relations to study how SE and POR can be combined while preserving correctness
 543 and completeness. Our work shows that the framework of correctness and completeness
 544 relations between symbolic and concrete transition systems, introduced by de Boer and
 545 Bonsangue, extends to parallelism and trace semantics, and provides a natural setting to
 546 study formalizations of abstraction techniques for SE, such as POR.

547 In addition, our formal development of correctness and completeness relations of SE and
 548 POR has been fully mechanized using Coq². We believe the mechanization of this framework
 549 in Coq can be useful to the community to study further formalizations of abstraction
 550 techniques for symbolic execution and their correctness. In particular, in future work, we
 551 plan to extend the framework developed in this paper to understand relations between
 552 concrete SE frameworks typically used for software testing [9], such as Klee [8], in which
 553 states are described using symbolic stores as in this paper, and abstract SE frameworks
 554 typically used for deductive verification, such as KeY [2], in which states are described using
 555 predicates.

² Provided as supplementary material at <https://github.com/Aqissiaq/symex-formally-formalized>
 and <https://zenodo.org/record/8070170>

556 — References —

- 557 1 Parosh Aziz Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas. Optimal
558 dynamic partial order reduction. In Suresh Jagannathan and Peter Sewell, editors, *Proc.*
559 *41st Annual Symposium on Principles of Programming Languages (POPL'14)*, pages 373–384.
560 ACM, 2014. doi:10.1145/2535838.2535845.
- 561 2 Wolfgang Ahrendt, Bernhard Beckert, Richard Bubel, Reiner Hähnle, Peter H. Schmitt,
562 and Mattias Ulbrich, editors. *Deductive Software Verification - The KeY Book - From*
563 *Theory to Practice*, volume 10001 of *Lecture Notes in Computer Science*. Springer, 2016.
564 doi:10.1007/978-3-319-49812-6.
- 565 3 Roberto Baldoni, Emilio Coppa, Daniele Cono D'elia, Camil Demetrescu, and Irene Finocchi.
566 A survey of symbolic execution techniques. *ACM Computing Surveys (CSUR)*, 51(3):1–39,
567 2018.
- 568 4 Yves Bertot and Pierre Castéran. *Interactive theorem proving and program development:*
569 *Coq'Art: the calculus of inductive constructions*. Springer, 2013.
- 570 5 Frank S de Boer and Marcello Bonsangue. Symbolic execution formally explained. *Formal*
571 *Aspects of Computing*, 33(4):617–636, 2021.
- 572 6 Frank S de Boer, Marcello Bonsangue, Einar Broch Johnsen, Violet Ka I Pun, S Lizeth
573 Tapia Tarifa, and Lars Tveito. SymPaths: Symbolic execution meets partial order reduction.
574 In *Deductive Software Verification: Future Perspectives*, pages 313–338. Springer, 2020.
- 575 7 Robert S. Boyer, Bernard Elspas, and Karl N. Levitt. SELECT - a formal system for testing
576 and debugging programs by symbolic execution. In Martin L. Shooman and Raymond T. Yeh,
577 editors, *Proc. International Conference on Reliable Software 1975*, pages 234–245. ACM, 1975.
578 doi:10.1145/800027.808445.
- 579 8 Cristian Cadar, Daniel Dunbar, and Dawson R. Engler. KLEE: unassisted and automatic
580 generation of high-coverage tests for complex systems programs. In Richard Draves and
581 Robbert van Renesse, editors, *Proc. 8th USENIX Symposium on Operating Systems Design*
582 *and Implementation (OSDI 2008)*, pages 209–224. USENIX Association, 2008. URL: http://www.usenix.org/events/osdi08/tech/full_papers/cadar/cadar.pdf.
- 583 9 Cristian Cadar and Koushik Sen. Symbolic execution for software testing: three decades later.
584 *Commun. ACM*, 56(2):82–90, 2013. doi:10.1145/2408776.2408795.
- 585 10 Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José
586 Meseguer, and Carolyn L. Talcott, editors. *All About Maude - A High-Performance Logical*
587 *Framework, How to Specify, Program and Verify Systems in Rewriting Logic*, volume 4350 of
588 *Lecture Notes in Computer Science*. Springer, 2007. doi:10.1007/978-3-540-71999-1.
- 589 11 Crystal Chang Din, Reiner Hähnle, Ludovic Henrio, Einar Broch Johnsen, Violet Ka I Pun,
590 and Silvia Lizeth Tapia Tarifa. LAGC semantics of concurrent programming languages. *arXiv*
591 *preprint arXiv:2202.12195*, 2022.
- 592 12 Matthias Felleisen and Robert Hieb. The revised report on the syntactic theories of se-
593 quential control and state. *Theoretical Computer Science*, 103(2):235–271, 1992. doi:
594 10.1016/0304-3975(92)90014-7.
- 595 13 Cormac Flanagan and Patrice Godefroid. Dynamic partial-order reduction for model checking
596 software. In *Proc. 32nd Symposium on Principles of Programming Languages (POPL'05)*,
597 pages 110–121. ACM, 2005. doi:10.1145/1040305.1040315.
- 598 14 José Fragoso Santos, Petar Maksimović, Sacha-Élie Ayoun, and Philippa Gardner. Gillian,
599 part i: a multi-language platform for symbolic execution. In *Proc. 41st ACM Conference on*
600 *Programming Language Design and Implementation (PLDI'20)*, pages 927–942, 2020.
- 601 15 Patrice Godefroid. Using partial orders to improve automatic verification methods. In
602 Edmund M. Clarke and Robert P. Kurshan, editors, *Computer-Aided Verification*, pages
603 176–185. Springer, 1991.
- 604 16 Patrice Godefroid. *Partial-order methods for the verification of concurrent systems: an*
605 *approach to the state-explosion problem*. Springer, 1996.
- 606

- 607 17 Patrice Godefroid and Pierre Wolper. Using partial orders for the efficient verification of
608 deadlock freedom and safety properties. *Formal Methods in System Design*, 2:149–164, 1993.
- 609 18 Shmuel Katz and Zohar Manna. Towards automatic debugging of programs. *ACM SIGPLAN*
610 *Notices*, 10(6):143–155, 1975.
- 611 19 James C King. Symbolic execution and program testing. *Communications of the ACM*,
612 19(7):385–394, 1976.
- 613 20 Åsmund Aqissiaq Arild Kløvstad. Compositional correctness and completeness for symbolic
614 partial order reduction: Concur23, June 2023. doi:10.5281/zenodo.8070170.
- 615 21 Xavier Leroy. Mechanized semantics. Course materials, 2020. URL: [https://github.com/
616 xavierleroy/cdf-mech-sem](https://github.com/xavierleroy/cdf-mech-sem).
- 617 22 Dorel Lucanu, Vlad Rusu, and Andrei Arusoae. A generic framework for symbolic execution:
618 A coinductive approach. *Journal of Symbolic Computation*, 80:125–163, 2017. SI: Program
619 Verification. doi:10.1016/j.jsc.2016.07.012.
- 620 23 Petar Maksimović, Sacha-Élie Ayoun, José Fragoso Santos, and Philippa Gardner. Gillian,
621 part ii: Real-world verification for JavaScript and C. In Alexandra Silva and K. Rustan M.
622 Leino, editors, *Computer Aided Verification*, pages 827–850. Springer, 2021.
- 623 24 Antoni Mazurkiewicz. Concurrent program schemes and their interpretations. *DAIMI Report*
624 *Series*, 6(78), Jul. 1977. URL: <https://tidsskrift.dk/daimipb/article/view/7691>, doi:
625 10.7146/dpb.v6i78.7691.
- 626 25 Benjamin C Pierce, A Azevedo de Amorim, Chris Casinghino, Marco Gaboardi, Michael
627 Greenberg, Cătălin Hrițcu, Vilhelm Sjöberg, A Tolmach, and B Yorgey. Software foundations,
628 volume 2: Programming language foundations. 2017.
- 629 26 César Rodríguez, Marcelo Sousa, Subodh Sharma, and Daniel Kroening. Unfolding-based
630 partial order reduction. In Luca Aceto and David de Frutos-Escrig, editors, *26th International*
631 *Conference on Concurrency Theory, CONCUR 2015*, volume 42 of *LIPICs*, pages 456–469.
632 Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.CONCUR.2015.
633 456.
- 634 27 Grigore Rosu, Andrei Stefanescu, Stefan Ciobăcă, and Brandon M. Moore. One-path reach-
635 ability logic. In *Proc. 28th Annual ACM/IEEE Symposium on Logic in Computer Science*
636 *(LICS'13)*, pages 358–367, 2013. doi:10.1109/LICS.2013.42.
- 637 28 José Fragoso Santos, Petar Maksimović, Théotime Grohens, Julian Dolby, and Philippa
638 Gardner. Symbolic execution for JavaScript. In *Proc. 20th International Symposium on*
639 *Principles and Practice of Declarative Programming, PPDP '18*. ACM, 2018. doi:10.1145/
640 3236950.3236956.
- 641 29 Daniel Schemmel, Julian Büning, César Rodríguez, David Laprell, and Klaus Wehrle. Symbolic
642 partial-order execution for testing multi-threaded programs. In *International Conference on*
643 *Computer Aided Verification*, pages 376–400. Springer, 2020.
- 644 30 Andrei Ștefănescu, Ștefan Ciobăcă, Radu Mereuta, Brandon M. Moore, Traian Florin Șer-
645 bănută, and Grigore Roșu. All-path reachability logic. In Gilles Dowek, editor, *Rewriting and*
646 *Typed Lambda Calculi*, pages 425–440. Springer, 2014.
- 647 31 Dominic Steinhöfel. *Abstract execution: automatically proving infinitely many programs*. PhD
648 thesis, Technische Universität Darmstadt, 2020.
- 649 32 Dominic Steinhöfel and Reiner Hähnle. The trace modality. In *Dynamic Logic. New Trends*
650 *and Applications*, pages 124–140. Springer, 2020. doi:10.1007/978-3-030-38808-9_8.
- 651 33 The Coq Development Team. The Coq proof assistant, September 2022. doi:10.5281/zenodo.
652 7313584.

653 **A** Auxiliary Lemmas

654 ► **Lemma A.1** (Substitution [5]). $V \circ \sigma(e) = V(e\sigma)$

655 ► **Corollary A.2** (Soundness of Assignment [5]). $V \circ (\sigma[x := a\sigma]) = V'[x := V'(a)]$, where
656 $V' = V \circ \sigma$

B Selected Proofs

B.1 Theorem 2.13: Bisimulation of concrete and symbolic semantics

For any initial valuation V and initial traces τ_0, τ'_0 such that τ_0 V -abstracts τ'_0 :

- if there is a concrete step $(s, \tau_0) \Rightarrow_V (s', \tau)$, then there exists a symbolic step $(s, \tau'_0) \rightarrow (s', \tau')$ such that τ' V -abstracts τ , and
- if there is a symbolic step $(s, \tau'_0) \rightarrow (s', \tau')$ and $V \models pc(\tau')$, then there exists a concrete step $(s, \tau_0) \Rightarrow_V (s', \tau)$ such that $\tau \Downarrow_V = V \circ \tau' \Downarrow$

Proof. Both directions are proven by case analysis and applying the corresponding rule (ie. the matching one from Figure 3). The `skip`-cases follow immediately, but the others take a bit more work.

For the conditional steps (choice and iteration), the concrete-to-symbolic direction requires splitting on the evaluation of the guard and picking the right symbolic branch. Then Lemma A.1 ensures that the resulting trace is the correct one. The symbolic-to-concrete direction requires showing that if the true (resp. false) branch is chosen, then the guard evaluates to true (resp. false) — which, again, holds by Lemma A.1.

Finally, both directions of the assignment step follow directly from Corollary A.2. ◀

B.2 Lemma 3.4: Equivalent steps

For equivalent traces $\tau \sim \tau'$, if $(s, \tau) \rightarrow (s', \tau_1)$ then there exists τ_2 so that $(s, \tau') \rightarrow (s', \tau_2)$ and $\tau_1 \sim \tau_2$

Proof. First note that

$$\text{If } \tau \sim \tau', \text{ then } \tau :: e \sim \tau' :: e \text{ for any event } e \quad (1)$$

Then we proceed by case analysis of the reduction step. In the `skip` cases, we pick $\tau_2 = \tau'$ and the equivalence follows by assumption. In the other cases, τ' is extended by some event, but then the result is equivalent by 1. ◀

B.3 Theorem 3.5: Abstraction Congruence

For equivalent symbolic traces $\tau_S \sim \tau'_S$ and concrete traces $\tau_C \simeq \tau'_C$, if τ_S V -abstracts τ_C then τ'_S V -abstracts τ'_C .

Proof. Since τ_S V -abstracts τ_C , V satisfies the path condition of τ_S and $\tau_C \Downarrow_V$ is the composition of V and $\tau_S \Downarrow$. But since it is trace equivalent, τ'_S has equivalent path condition, so it must also be satisfied. Furthermore its final substitution is the same as τ_S and likewise the final valuation of τ'_C is equal to that of τ_C , so $\tau'_C \Downarrow_V = V \circ \tau'_S \Downarrow$ as required. ◀

B.4 Theorem 4.3: POR Bisimulation

For equivalent initial traces $\tau_0 \sim \tau'_0$:

- If $(s, \tau_0) \rightarrow_{POR} (s', \tau)$ then there exists $(s, \tau'_0) \rightarrow (s', \tau')$ such that $\tau \sim \tau'$, and
- If $(s, \tau_0) \rightarrow (s', \tau)$ then there exists $(s, \tau'_0) \rightarrow_{POR} (s', \tau')$ such that $\tau \sim \tau'$

For equivalent initial traces $\tau_0 \simeq \tau'_0$ and initial valuation V :

- If $(s, \tau_0) \Rightarrow_{POR, V} (s', \tau)$ then there exists $(s, \tau'_0) \Rightarrow_V (s', \tau')$ such that $\tau \simeq \tau'$, and
- If $(s, \tau_0) \Rightarrow_V (s', \tau)$ then there exists $(s, \tau'_0) \Rightarrow_{POR, V} (s', \tau')$ such that $\tau \simeq \tau'$

Proof. Both results follow directly from unpacking contexts and applying Lemma 3.4. ◀

696 **B.5 Theorem 5.1: POR-POR Bisimulation**697 For initial traces τ_S, τ_C such that τ_S V -abstracts τ_C :

- 698 ■ If $(s, \tau_C) \Rightarrow_{\text{POR}, V} (s', \tau'_C)$, then there exists $(s, \tau_S) \rightarrow_{\text{POR}} (s', \tau'_S)$ such that τ'_S V -
699 abstracts τ'_C
- 700 ■ If $(s, \tau_S) \rightarrow_{\text{POR}} (s', \tau'_S)$ and $V \models pc(\tau'_S)$, then there exists $(s, \tau_C) \Rightarrow_{\text{POR}, V} (s', \tau'_C)$ and
701 $\tau'_C \Downarrow_V = V \circ (\tau'_S \Downarrow)$

702 **Proof.** Let us consider the first (completeness) direction. By the concrete version of Theo-
703 rem 4.3 we obtain the concrete step $(s, \tau) \Rightarrow_V (s', \tau''_C)$ with $\tau''_C \simeq \tau'_C$. Then, by Theorem 2.13
704 we have a symbolic step $(s, \tau) \rightarrow (s', \tau''_S)$ such that τ''_S V -abstracts τ''_C . Finally, the symbolic
705 half of Theorem 4.3 provides the desired $(s, \tau_S) \rightarrow_{\text{POR}} (s', \tau'_S)$ with $\tau'_S \sim \tau''_S$ and abstraction
706 congruence (Theorem 3.5) shows τ'_S V -abstracts τ'_C . The other direction is analogous, using
707 the other half of each bisimulation. ◀

708 **B.6 Theorem 5.4: Total Bisimulation**709 For initial traces τ_S, τ_C such that τ_S V -abstracts τ_C :

- 710 ■ If $(s, \tau_C) \Rightarrow_V (s', \tau'_C)$, then there exists $(s, \tau_S) \rightarrow_{\text{POR}} (s', \tau'_S)$ such that τ'_S V -abstracts
711 τ'_C
- 712 ■ If $(s, \tau_S) \rightarrow_{\text{POR}} (s', \tau'_S)$ and $V \models pc(\tau'_S)$, then there exists $(s, \tau_C) \Rightarrow_V (s', \tau'_C)$ and
713 $\tau'_C \Downarrow_V = V \circ (\tau'_S \Downarrow)$

714 **Proof.** The proof follows a similar structure to the above, composing the concrete/symbolic
715 (Theorem 2.13) and symbolic POR (Theorem 4.3) bisimulations. ◀