

Concolic Execution and Test Case Generation in Prolog^{*}

Germán Vidal

MiST, DSIC, Universitat Politècnica de València
Camino de Vera, s/n, 46022 Valencia, Spain
gvidal@dsic.upv.es

Abstract. Symbolic execution extends concrete execution by allowing symbolic input data and then exploring all feasible execution paths. It has been defined and used in the context of many different programming languages and paradigms. A symbolic execution engine is at the heart of many program analysis and transformation techniques, like partial evaluation, test case generation or model checking, to name a few. Despite its relevance, traditional symbolic execution also suffers from several drawbacks. For instance, the search space is usually huge (often infinite) even for the simplest programs. Also, symbolic execution generally computes an overapproximation of the concrete execution space, so that false positives may occur. In this paper, we propose the use of a variant of symbolic execution, called concolic execution, for test case generation in Prolog. Our technique aims at full statement coverage. We argue that this technique computes an underapproximation of the concrete execution space (thus avoiding false positives) and scales up better to medium and large Prolog applications.

1 Introduction

There is a renewed interest in *symbolic execution* [9, 3], a well-known technique for program verification, testing, debugging, etc. In contrast to concrete execution, symbolic execution considers that the values of some input data are unknown, i.e., some input parameters x, y, \dots take *symbolic values* X, Y, \dots . Because of this, symbolic execution is often non-deterministic: at some control statements, we need to follow more than one execution path because the available information does not suffice to determine the validity of a control expression, e.g., symbolic execution may follow both branches of the conditional “`if ($x > 0$) then $exp1$ else $exp2$ ” when the symbolic value X of variable x is not constrained enough to imply neither $x > 0$ nor $\neg(x > 0)$. Symbolic states include a path condition that stores the current constraints on symbolic values, i.e., the`

^{*} This work has been partially supported by the EU (FEDER) and the Spanish *Ministerio de Economía y Competitividad (Secretaría de Estado de Investigación, Desarrollo e Innovación)* under grant TIN2013-44742-C4-1-R and by the *Generalitat Valenciana* under grant PROMETEO/2011/052.

conditions that must hold to reach a particular execution state. E.g., after symbolically executing the above conditional, the derived states for *exp1* and *exp2* would add the conditions $X > 0$ and $X \leq 0$, respectively, to their path conditions.

Traditionally, formal techniques based on symbolic execution have enforced *soundness*: if a symbolic state is reached and its path condition is satisfiable, there must be a concrete execution path that reaches the corresponding concrete state. In contrast, we say that symbolic execution is *complete* when every reachable state in a concrete execution is “covered” by some symbolic state. For the general case of infinite state systems, completeness usually requires some kind of *abstraction* (as in infinite state model checking).

In the context of logic programming, we can find many techniques that use some form of *complete* symbolic execution, like partial evaluation [12, 13, 4]. However, these overapproximations of the concrete semantics may have a number of drawbacks in the context of testing and debugging. On the one hand, one should define complex subsumption and abstraction operators since the symbolic search space is usually infinite. These abstraction operators, may introduce *false positives*, which is often not acceptable when debugging large applications. On the other hand, because of the complexity of these operators, the associated methods usually do not scale to medium and large applications.

In imperative programming, an alternative approach, called *concolic execution* [6, 16], has become popular in the last years. Basically, concolic execution proceeds as follows: first, a concrete execution using random input data is performed. In parallel to the concrete execution, a symbolic execution is also performed, but restricted to the same conditional choices of the concrete execution. Then, by negating one of the constraints in the symbolic execution, new input data are obtained, and the process starts again. Here, only concrete executions are considered and, thus, no false positives are produced. This approach has given rise to a number of powerful and scalable tools in the context of imperative and concurrent programming, like Java Pathfinder [14] and SAGE [7].

In this paper, we present a novel scheme for testing pure Prolog (without negation) based on a notion of concolic execution. To the best of our knowledge, this is the first approach to concolic execution in the context of a declarative programming paradigm.

2 Preliminaries

We assume some familiarity with the standard definitions and notations for logic programs [11]. Nevertheless, in order to make the paper as self-contained as possible, we present in this section the main concepts which are needed to understand our development.

In this work, we consider a first-order language with a fixed vocabulary of predicate symbols, function symbols, and variables denoted by Π , Σ and \mathcal{V} , respectively. In the following, we let $\overline{o_n}$ denote the sequence of syntactic objects o_1, \dots, o_n , and we let $\overline{o_{n,m}}$ denote the sequence o_n, o_{n+1}, \dots, o_m . Also, we often use \overline{o} when the number of elements in the sequence is irrelevant. We let $\mathcal{T}(\Sigma, \mathcal{V})$

denote the set of *terms* constructed using symbols from Σ and variables from \mathcal{V} . An *atom* has the form $p(\overline{t_n})$ with $p/n \in \Pi$ and $t_i \in \mathcal{T}(\Sigma, \mathcal{V})$ for $i = 1, \dots, n$. A *goal* is a finite sequence of atoms A_1, \dots, A_n , where the *empty goal* is denoted by *true*. A *clause* has the form $H \rightarrow \mathcal{B}$ where H is an atom and \mathcal{B} is a goal (note that we only consider *definite* programs). A logic *program* is a finite sequence of clauses. $\text{Var}(s)$ denotes the set of variables in the syntactic object s (i.e., s can be a term, an atom, a query, or a clause). A syntactic object s is *ground* if $\text{Var}(s) = \emptyset$. In this work, we only consider *finite* ground terms.

Substitutions and their operations are defined as usual. In particular, the set $\text{Dom}(\sigma) = \{x \in \mathcal{V} \mid \sigma(x) \neq x\}$ is called the *domain* of a substitution σ . We let *id* denote the empty substitution. The application of a substitution θ to a syntactic object s is usually denoted by juxtaposition, i.e., we write $s\theta$ rather than $\theta(s)$. A syntactic object s_1 is *more general* than a syntactic object s_2 , denoted $s_1 \leq s_2$, if there exists a substitution θ such that $s_2 = s_1\theta$. A *variable renaming* is a substitution that is a bijection on \mathcal{V} . Two syntactic objects t_1 and t_2 are *variants* (or equal up to variable renaming), denoted $t_1 \approx t_2$, if $t_1 = t_2\rho$ for some variable renaming ρ . A substitution θ is a unifier of two syntactic objects t_1 and t_2 iff $t_1\theta = t_2\theta$; furthermore, θ is the *most general unifier* of t_1 and t_2 , denoted by $\text{mgu}(t_1, t_2)$ if, for every other unifier σ of t_1 and t_2 , we have that $\theta \leq \sigma$.

The notion of *computation rule* \mathcal{R} is used to select an atom within a goal for its evaluation. Given a program P , a goal $\mathcal{G} \equiv A_1, \dots, A_n$, and a computation rule \mathcal{R} , we say that $\mathcal{G} \rightsquigarrow_{P, \mathcal{R}, \sigma} \mathcal{G}'$ is an *SLD resolution step* for \mathcal{G} with P and \mathcal{R} if

- $\mathcal{R}(\mathcal{G}) = A_i$, $1 \leq i \leq n$, is the selected atom,
- $H \rightarrow \mathcal{B}$ is a renamed apart clause of P (in symbols $H \rightarrow \mathcal{B} \ll P$),
- $\sigma = \text{mgu}(A, H)$, and
- $\mathcal{G}' \equiv (A_1, \dots, A_{i-1}, \mathcal{B}, A_{i+1}, \dots, A_n)\sigma$.

We often omit P , \mathcal{R} and/or σ in the notation of an SLD resolution step when they are clear from the context. An *SLD derivation* is a (finite or infinite) sequence of SLD resolution steps. We often use $\mathcal{G}_0 \rightsquigarrow_{\theta}^* \mathcal{G}_n$ as a shorthand for $\mathcal{G}_0 \rightsquigarrow_{\theta_1} \mathcal{G}_1 \rightsquigarrow_{\theta_2} \dots \rightsquigarrow_{\theta_n} \mathcal{G}_n$ with $\theta = \theta_n \circ \dots \circ \theta_1$ (where $\theta = \{\}$ if $n = 0$). An SLD derivation $\mathcal{G} \rightsquigarrow_{\theta}^* \mathcal{G}'$ is *successful* when $\mathcal{G}' = \text{true}$; in this case, we say that θ is the *computed answer substitution*. SLD derivations are represented by a (possibly infinite) finitely branching tree.

3 A Deterministic Semantics

In this section, we introduce a deterministic small-step semantics for pure Prolog (without negation). Basically, as we will see in the next section, we need to keep some information through the complete Prolog computation, and the usual semantics based on non-determinism and backtracking is not adequate for this purpose. Therefore, we propose the use of a stack to store alternative execution paths that are tried when a failure is reached. The resulting small-step semantics

$$\begin{array}{c}
\text{(unfolding)} \frac{\text{let } H_n \stackrel{\ell_n}{\leftarrow} \mathcal{B}_n \ll P \text{ be all the clauses such that} \\ \text{mgu}(A_1, H_i) = \sigma_i \neq \text{fail}, i = 1, \dots, n}{\langle \overline{A_m}; \sigma; \mathcal{S} \rangle \xrightarrow{u(\overline{\ell_n})} \langle (\mathcal{B}_1, \overline{A_{2,m}})\sigma_1; \sigma\sigma_1; [(\ell_2; \sigma\sigma_2; (\mathcal{B}_2, \overline{A_{2,m}})\sigma_2), \\ \dots, \\ (\ell_n; \sigma\sigma_n; (\mathcal{B}_n, \overline{A_{2,m}})\sigma_n)] ++ \mathcal{S} \rangle} \\
\text{(backtracking)} \frac{\text{there is no clause } H \stackrel{\ell'}{\leftarrow} \mathcal{B} \ll P \text{ such that } \text{mgu}(A_1, H) \neq \text{fail}}{\langle \overline{A_m}; \sigma'; [(\ell; \mathcal{G}; \sigma) | \mathcal{S}] \rangle \xrightarrow{b(\ell)} \langle \mathcal{G}; \sigma; \mathcal{S} \rangle} \quad (m \geq 1) \\
\text{(failure)} \frac{\text{there is no clause } H \stackrel{\ell}{\leftarrow} \mathcal{B} \ll P \text{ such that } \text{mgu}(A_1, H) \neq \text{fail}}{\langle \overline{A_m}; \sigma; [] \rangle \xrightarrow{f} \langle \text{fail}; \sigma; [] \rangle} \quad (m \geq 1)
\end{array}$$

Fig. 1. Deterministic small-step semantics

is clearly equivalent to the original one when a depth-first search is considered. Actually, our deterministic semantics is essentially equivalent to (a subset of) the *linear* semantics presented in [17].

In the following, we assume that the program clauses are labeled. In particular, given a program P , we use the notation $H \stackrel{\ell}{\leftarrow} \mathcal{B} \ll P$ to refer to a (renamed apart) labeled clause $H \leftarrow \mathcal{B}$ in P . Labels must be unique. Moreover, we only consider Prolog's left-to-right computation rule, and assume that only the computation of the first answer for the initial goal is relevant (as it is common in practical Prolog applications).

Our semantics deals with *states*, which are defined as follows:

Definition 1 (state). *A state is a tuple $\langle \mathcal{G}; \sigma; \mathcal{S} \rangle$, where \mathcal{G} is a goal, σ is a substitution—the (partial) answer computed so far—and \mathcal{S} , the stack, is a (possibly empty) list of tuples $(\ell; \mathcal{G}'; \sigma')$ with ℓ a clause label, \mathcal{G}' a goal and σ' a substitution.*

The small-step deterministic semantics is defined as the smallest relation that obeys the labeled transition rules shown in Figure 1, where $[H|R]$ denotes a list with head H and tail R , and “++” denotes list concatenation.

Given a goal \mathcal{G}_0 , the initial state has the form $\langle \mathcal{G}_0; id; [] \rangle$. The transition relation is labeled with $u(\overline{\ell_n})$, denoting an unfolding step with the clauses labeled with $\overline{\ell_n}$, $b(\ell)$, denoting a backtracking step that tries a clause labeled with ℓ , or f , denoting a failing derivation.

Let us briefly explain the rules of the small-step semantics:

- The **unfolding** rule proceeds as in standard SLD resolution, but considers all matching clauses, so that all SLD resolution steps are performed in one go. The first unfolding step is used to replace the goal component of the state,

while the remaining ones (if any) are added on top of the stack (thus we mimic the usual depth-first search of Prolog). Here, the labels of the clauses and the partial computed answers are also stored in the stack in order to recover this information when a backtracking step is performed.

- The **backtracking** rule applies when no further unfolding is possible and the goal component is not *true* (the empty goal). In this case, we discard the current goal and consider the first goal in the stack, extracting the clause label and the partial answer that are needed for the transition step.
- Finally, the **failure** rule is used to terminate a computation that reaches a goal in which the selected atom does not match any rule and, moreover, there are no alternatives in the stack.

A *successful* computation has the form $\langle \mathcal{G}_0; id; [] \rangle \xrightarrow{s_1} \langle \mathcal{G}_1; \sigma_1; \mathcal{S}_1 \rangle \xrightarrow{s_2} \dots \xrightarrow{s_n} \langle true; \sigma_n; \mathcal{S}_n \rangle$, where σ_n (restricted to the variables of \mathcal{G}_0) is the computed answer substitution. A *failing* computation has the form $\langle \mathcal{G}_0; id; [] \rangle \xrightarrow{s_1} \langle \mathcal{G}_1; \sigma_1; \mathcal{S}_1 \rangle \xrightarrow{s_2} \dots \xrightarrow{s_n} \langle \mathcal{G}_n; \sigma_n; \mathcal{S}_n \rangle \xrightarrow{f_{id}} \langle fail; \sigma_n; [] \rangle$; we keep σ_n in the last state since it might be useful for analyzing finite failure derivations.

Now, we introduce the following notion of execution *trace*, that will be used in the next section to steer the symbolic execution.

Definition 2 (trace). Let $\langle \mathcal{G}_0; id; \mathcal{S}_0 \rangle \xrightarrow{s_1} \langle \mathcal{G}_1; \sigma_1; \mathcal{S}_1 \rangle \xrightarrow{s_2} \dots \xrightarrow{s_n} \langle \mathcal{G}_n; \sigma_n; \mathcal{S}_n \rangle$ be a computation. The associated trace is the list $[s_1, s_2, \dots, s_n]$, where each s_i is either of the form $u(\overline{\ell}_m)$, $b(\ell)$ or f .

Example 1. Consider the `rev_acc_type` program to reverse a list using an accumulator and also checking the type of the input parameter (from the DPPD library [10]), extended with predicates `main`, `length`, and `foo`:

```

(1) main(L,N,R) :-
    length(L,N),
    rev(L, [], R),
    foo(a).
(2) main(_L,_N,error).

(3) rev([],A,A).
(4) rev([H|T],Acc,Res) :-
    is_list(Acc),
    rev(T, [H|Acc], Res).

(5) is_list([]).
(6) is_list([_H|_T]) :-
    is_list(T).

(7) length([],0).
(8) length([_H|_R],s(N)) :-
    length(R,N).

(9) foo(b).

```

Here, we use natural numbers as clause labels. Predicate `main` considers two cases: if the input list `L` has length `N` (the length is represented using natural numbers built from 0 and `s(-)` to avoid the use of built-ins), the reverse of `L` is computed; otherwise, we assume that an error occurs. The computation for the initial goal `main([a,b],s(s(0)),R)` is shown in Figure 2, where only the relevant computed substitutions are shown. The trace associated to the computation is

$[u(1,2), u(8), u(8), u(7), u(4), u(5), u(4), u(6), u(5), u(3), b(2)]$

```

⟨main([a, b], s(s(0)), R); id; []⟩
   $\xrightarrow{u(1,2)}$  ⟨length([a, b], s(s(0))), rev([a, b], [], R), foo(a); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(8)}$  ⟨length([b], s(0)), rev([a, b], [], R), foo(a); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(8)}$  ⟨length([], 0), rev([a, b], [], R), foo(a); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(7)}$  ⟨rev([a, b], [], R), foo(a); id; [2; ({R/error}; true)]⟩
   $\xrightarrow{u(4)}$  ⟨is_list([], rev([b], [a], R), foo(a); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(5)}$  ⟨rev([b], [a], R), foo(a); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(4)}$  ⟨is_list([a], rev([], [b, a], R), foo(a); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(6)}$  ⟨is_list([], rev([], [b, a], R), foo(a); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(5)}$  ⟨rev([], [b, a], R), foo(a); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(3)}$  ⟨foo(a); {R/[b, a]}; [(2; {R/error}; true)]⟩
   $\xrightarrow{b(2)}$  ⟨true; {R/error}; []⟩

```

Fig. 2. Successful computation for $\text{main}([a, b], s(s(0)), R)$.

4 Concolic Execution

In this section, we introduce the semantics of concolic execution. Essentially, it deals with symbolic input data (free variables in our context), as in standard symbolic execution, but is driven by a concrete execution. Often, a single algorithm mixing both concrete and symbolic execution is introduced. In contrast, for clarity, we prefer to keep both calculi independent: the concrete semantics produces a trace, which is then used to steer the symbolic execution.¹

The symbolic states for concolic execution are defined as follows:

Definition 3 (symbolic state). *A symbolic state is a tuple $\langle \tau; \mathcal{L}; \mathcal{G}; \sigma; \mathcal{S}; T \rangle$, where*

- τ is a computation trace,
- \mathcal{L} is a list of clause labels (namely, a stack that keeps track of the current clause environment),²
- \mathcal{G} is a goal,
- σ is the partial answer computed so far (a substitution),
- \mathcal{S} is a (possibly empty) list of tuples, $(\ell; \mathcal{L}'; \sigma; \mathcal{G}')$, where \mathcal{L}' is also a list of clause labels, and
- T is a set of clause labels (the labels of those clauses not yet completely evaluated).

¹ Nevertheless, an implementation of this technique may as well combine both calculi into a single algorithm to improve efficiency.

² The usefulness of keeping the clause stack will become clear in the next section. Basically, it is needed to know which other clauses can be completely evaluated when a given clause—the one that is on top of the stack—is completely evaluated.

$$\begin{array}{l}
\text{(unfolding)} \frac{\overline{\text{if } H_n \stackrel{\ell_n}{\leftarrow} \mathcal{B}_n \ll P \text{ with } \text{mgu}(A_1, H_i) = \sigma_i, i = 1, \dots, n \text{ and} \\
H'_k \stackrel{\ell'_k}{\leftarrow} \mathcal{B}'_k \ll P \text{ with } \text{mgu}(A_1, H'_j) = \sigma'_j, j = 1, \dots, k \\
\text{and } \{\ell_n\} \cap \{\ell'_k\} = \{\}}}{\langle [\mathbf{u}(\overline{\ell_n}) | \tau]; \mathcal{L}; \overline{A_m}; \sigma; \mathcal{S}; T \rangle \xrightarrow{\{c(\ell'_k, \sigma'_k)\}} \langle \tau; [\ell_1 | \mathcal{L}]; \\
(\mathcal{B}_1, \mathbf{e}(\ell_1), \overline{A_{2,m}}) \sigma_1; \sigma \sigma_1; \\
[(\ell_2; [\ell_2 | \mathcal{L}]; \sigma \sigma_2; (\mathcal{B}_2, \mathbf{e}(\ell_2), \overline{A_{2,m}}) \sigma_2), \\
\dots, \\
(\ell_n; [\ell_n | \mathcal{L}]; \sigma \sigma_n; (\mathcal{B}_n, \mathbf{e}(\ell_n), \overline{A_{2,m}}) \sigma_n)] ++ \mathcal{S}; T \rangle} \\
\\
\text{(exit)} \frac{A_1 = \mathbf{e}(\ell)}{\langle \tau; [\ell | \mathcal{L}]; \overline{A_m}; \sigma; \mathcal{S}; T \rangle \xrightarrow{\{\}} \langle \tau; \mathcal{L}; \overline{A_{2,m}}; \sigma; \mathcal{S}; T \setminus \{\ell\} \rangle} \\
\\
\text{(backtracking)} \frac{\overline{A_1 \neq \mathbf{e}(-), H_n \stackrel{\ell_n}{\leftarrow} \mathcal{B}_n \ll P \text{ s.t. } \text{mgu}(A_1, H_i) = \sigma \neq \text{fail}, i = 1, \dots, n} \quad (m \geq 1)}{\langle [\mathbf{b}(\ell) | \tau]; \mathcal{L}; \overline{A_m}; \sigma; [(\ell; \mathcal{L}'; \sigma'; \mathcal{G}) | \mathcal{S}]; T \rangle \xrightarrow{\{c(\ell_n, \sigma_n)\}} \langle \tau; \mathcal{L}'; \mathcal{G}; \sigma'; \mathcal{S}; T \rangle} \\
\\
\text{(failure)} \frac{\overline{\text{if } H_n \stackrel{\ell_n}{\leftarrow} \mathcal{B}_n \ll P \text{ such that } \text{mgu}(A_1, H_i) = \sigma \neq \text{fail}, i = 1, \dots, n} \quad (m \geq 1)}{\langle [\mathbf{f} | \tau]; \mathcal{L}; \overline{A_m}; \sigma; []; T \rangle \xrightarrow{\{c(\ell_n, \sigma_n)\}} \langle \tau; \mathcal{L}; \text{fail}; \sigma; []; T \rangle}
\end{array}$$

Fig. 3. Concolic execution semantics

The concolic execution semantics is defined as the smallest relation that obeys the labeled transition rules shown in Figure 3. Given a trace τ , the initial symbolic state has the form $\langle \tau; []; \mathcal{G}_0; id; []; T \rangle$, where \mathcal{G}_0 is a goal with the same predicates as in the concrete execution, but with fresh variables as arguments, and T is a set with the labels of all program clauses. The transition relation is labeled with a (possibly empty) list of terms of the form $c(\ell, \theta)$, which denote possible alternatives for unfolding that concrete execution did not consider. Missing alternatives will be used to generate new input data that explore different execution paths.

Let us briefly explain the rules of concolic execution semantics:

- The **unfolding** rule follows the trace of the concrete execution and applies the same unfolding step. Here, a call of the form $\mathbf{e}(\ell)$ is added to the end of the clause bodies to mark when the clauses are completely evaluated. This is required in our context since we only consider that a clause is *covered* when

all body atoms are successfully executed.³ This will be a useful information for test case generation, as we will see in the next section. Moreover, in this rule, we label the step with the information regarding the remaining clauses whose head unifies with the selected atom (and did not match with it in the concrete execution). Finally, we add ℓ_1 to the stack of clause labels (the current environment).

- The exit rule applies when the selected atom has the form $e(\ell)$. In this case, we remove ℓ from the top of the environment stack, and also delete ℓ from the set of clause labels T (i.e., clause ℓ has been completely evaluated).
- The **backtracking** and **failure** rules proceed analogously to the **unfolding** rule by labeling the step with the information regarding the additional clauses whose head unify with the selected atom (if any).

When the set labeling rules **unfolding**, **backtracking** and **failure** is not empty, we have identified situations in which the symbolic state can follow an execution path that is not possible with the concrete goal. Therefore, they allow us to construct new input data for the initial goal so that a different execution path is followed.

Let us now show a simple computation with the concolic execution semantics. We postpone to the next section the algorithm for test case generation.

Example 2. Consider the following simple program:

- (1) $p(X) :- q(X), r(X).$
- (2) $q(X) :- s(X).$
- (3) $s(a).$
- (4) $s(b).$
- (5) $r(b).$

where we again consider natural numbers as clause labels. The concrete execution for the initial goal $p(a)$ is as follows:⁴

$$\langle p(a); id; [] \rangle \xrightarrow{u(1)} \langle q(a), r(a); id; [] \rangle \xrightarrow{u(2)} \langle s(a), r(a); id; [] \rangle \xrightarrow{u(3)} \langle r(a); id; [] \rangle \xrightarrow{f} \langle fail; id; [] \rangle$$

Therefore, its associated trace is $\tau = [u(1), u(2), u(3), f]$. Now, for concolic execution, we consider the trace τ and the initial goal $p(X)$. The concolic execution is shown in Figure 4. As can be seen, the execution of clauses 1, 4 and 5 has not been completed. Moreover, we can observe that there was only one missing alternative when unfolding $s(X)$. In the next section, we show how this information can be used for test case generation.

³ Observe that other, more relaxed, notions of clause covering are possible; e.g., consider that a clause is covered as soon as the clause is used in an unfolding step. Also, see [2] for a more declarative notion of test coverage.

⁴ In the examples, we restrict the (partial) computed answers to the variables of the initial goal.

	$\langle [u(1), u(2), u(3), f]; []; \mathbf{p}(\mathbf{X}); id; []; [1, 2, 3, 4, 5] \rangle$		
\Downarrow	$\langle [u(2), u(3), f]; [1]; \mathbf{q}(\mathbf{X}), \mathbf{r}(\mathbf{X}), \mathbf{e}(1); id; []; [1, 2, 3, 4, 5] \rangle$		
\Downarrow	$\langle [u(3), f]; [2, 1]; \mathbf{s}(\mathbf{X}), \mathbf{e}(2), \mathbf{r}(\mathbf{X}), \mathbf{e}(1); id; []; [1, 2, 3, 4, 5] \rangle$		
$\{c(4, \{x/b\})\}$	$\langle [f]; [3, 2, 1]; \mathbf{e}(3), \mathbf{e}(2), \mathbf{r}(\mathbf{a}), \mathbf{e}(1); \{\mathbf{X}/\mathbf{a}\}; []; [1, 2, 3, 4, 5] \rangle$		
\Downarrow	$\langle [f]; [2, 1]; \mathbf{e}(2), \mathbf{r}(\mathbf{a}), \mathbf{e}(1); \{\mathbf{X}/\mathbf{a}\}; []; [1, 2, 4, 5] \rangle$		
\Downarrow	$\langle [f]; [1]; \mathbf{r}(\mathbf{a}), \mathbf{e}(1); \{\mathbf{X}/\mathbf{a}\}; []; [1, 4, 5] \rangle$		
\Downarrow	$\langle []; [1]; \mathbf{fail}; \{\mathbf{X}/\mathbf{a}\}; []; [1, 4, 5] \rangle$		

Fig. 4. Concolic execution for $[u(1), u(2), u(3), f]$ and $\mathbf{p}(\mathbf{X})$

5 Test Case Generation

In this section, we present an algorithm for test case generation using concolic execution. In contrast to previous approaches for Prolog testing, our technique considers an underapproximation, i.e., only actual executions are considered (since there is no abstraction involved). Therefore, no false positives may occur. If a test case shows an error, this is an actual error in the considered program.

5.1 The Algorithm

In this section, we assume that the program contains a single predicate that starts the execution, which we denote with *main*. This is not unusual for real applications. Moreover, we consider a particular *mode* for *main*,⁵ where $in(main/n) = \{i_1, \dots, i_m\}$ denotes the set of input parameters of *main*.

The algorithm for test case generation proceeds as follows:

1. First, a random goal of the form $main(\overline{t}_n)$ is produced, where at least the input arguments (according to $in(main/n)$) must be ground.
2. Now, we use the concrete semantics to execute the goal $main(\overline{t}_n)$, thus obtaining an associated trace τ . We assume that this execution terminates, which is reasonable since the input arguments are ground. In practice, one can use a timeout and report a warning when the execution takes more time.
3. Then, we use concolic execution to run an initial symbolic state of the form

$$\langle \tau; []; main(\overline{X}_n); id; []; T \rangle$$

where T is a set with the labels of all program clauses. Since the concrete execution was finite, so is the concolic execution (since it performs exactly the same steps). Let us consider that it has the following form:

$$\langle \tau_0; \mathcal{L}_0; \mathcal{G}_0; \sigma_0; \mathcal{S}_0; T_0 \rangle \xrightarrow{c_1} \dots \xrightarrow{c_m} \langle \tau_m; \mathcal{L}_m; \mathcal{G}_m; \sigma_m; \mathcal{S}_m; T_m \rangle$$

⁵ Extending our approach to multiple modes would not be difficult, but would introduce another source of nondeterminism when grounding an input goal.

where $\tau_0 = \tau$, $\mathcal{L}_0 = []$, $\mathcal{G}_0 = \text{main}(\overline{X_n})$, $\sigma_0 = \text{id}$, $\mathcal{S}_0 = []$, $T_0 = T$, and \mathcal{G}_m is either *true* or *fail*.

4. Now, we check the value of T_m . If $T_m = \{ \}$, the algorithm terminates since all clauses have been completely executed. Otherwise, we identify the *last* state $\langle \tau_i; \mathcal{L}_i; \mathcal{G}_i; \sigma_i; \mathcal{S}_i; T_i \rangle$ in the above concolic execution such that
 - the previous transition $\xrightarrow{c_i}$ is labeled with $c_i = \{c(\ell_k, \theta_k)\}$, $k > 0$, and
 - there exists $j \in \{1, \dots, k\}$ such that either $\ell_j \in T_m$ or \mathcal{L}_i contains (not necessarily in a top position) some labels from T_m ; the reason to also consider the labels from \mathcal{L}_i is that considering an alternative clause may help to complete the execution of *all* the clauses in the current clause stack. Either way, we choose a clause j in a non-deterministic way.
- Therefore, we have a prefix of the complete concolic execution of the form:

$$\langle \tau_0; \mathcal{L}_0; \mathcal{G}_0; \sigma_0; \mathcal{S}_0; T_0 \rangle \xrightarrow{c_1} \dots \xrightarrow{c_{i-1}} \langle \tau_{i-1}; \mathcal{L}_{i-1}; \mathcal{G}_{i-1}; \sigma_{i-1}; \mathcal{S}_{i-1}; T_{i-1} \rangle$$

$$\xrightarrow{c_i} \langle \tau_i; \mathcal{L}_i; \mathcal{G}_i; \sigma_i; \mathcal{S}_i; T_i \rangle$$

and we are interested in the (possibly) partial computed answer $\sigma_{i-1}\theta_j$, since it will allow us to explore a different execution path, possibly covering some more program clauses.

Hence, we have a second test case: $\mathcal{G}'_0 = \text{main}(\overline{X_n})\sigma_{i-1}\theta_j\gamma$, where γ is a substitution that is only aimed at grounding the input parameters *in(main/n)* of *main* using arbitrary values (of the right type, preferably minimal ones).

5. Finally, we consider the initial state $\langle \mathcal{G}'_0; \text{id}; [] \rangle$ and obtain a new trace using the concrete execution semantics τ' , so that a new initial symbolic state is defined as follows: $\langle \tau'; []; \text{main}(\overline{X_n}); \text{id}; []; T_m \rangle$ and the process starts again (i.e., we jump again to step 3). Observe that the initial state includes the set of clause labels T_m obtained in the last state of the previous concolic execution, in order to avoid producing new tests for clauses that are already covered by some previous test case.

Let us now illustrate the complete test case generation process with an example.

5.2 Test Case Generation in Practice

In this section, we illustrate the generation of test cases using a slight modification of the program in Example 1:

<pre>(1) main(L,N,R) :- length(L,N), rev(L, [], R). (2) main(_L,_N,error). (3) rev([],A,A). (4) rev([H T],Acc,Res) :- is_list(Acc), rev(T, [H Acc], Res).</pre>	<pre>(5) is_list([]). (6) is_list([_H T]) :- is_list(T). (7) length([],0). (8) length([_H R],s(N)) :- length(R,N).</pre>
--	---

$$\begin{array}{l}
\langle [u(1, 2), u(8), b(2)]; []; \text{main}(L, N, R); \text{id}; []; \{1, \dots, 8\} \rangle \\
\stackrel{\{\}}{\rightarrow} \langle [u(8), b(2)]; [1]; \left(\begin{array}{l} \text{length}(L, N), \\ \text{rev}(L, [], R), e(1) \end{array} \right); \text{id}; \mathcal{S}; \{1, \dots, 8\} \rangle \\
\langle [b(2)]; [8, 1]; \left(\begin{array}{l} \text{length}(L', N'), e(8), \\ \text{rev}([X|L'], [], R), e(1) \end{array} \right); \{L/[X|L'], N/s(N')\}; \mathcal{S}; \{1, \dots, 8\} \rangle \\
\langle [1]; [2]; e(2); \{\mathbf{R/error}\}; []; \{1, \dots, 8\} \rangle \\
\stackrel{\{\}}{\rightarrow} \langle [1]; [1]; \text{true}; \{\mathbf{R/error}\}; []; \{1, 3, \dots, 8\} \rangle \\
\text{with } \mathcal{S} = [(2; [2]; \{\mathbf{R/error}\}; e(2))];
\end{array}$$

Fig. 5. Concolic execution for $\langle [u(1, 2), u(8), b(2)]; []; \text{main}(L, N, R); \text{id}; []; \{1, 2, \dots, 8\} \rangle$

Observe that, in this example, using a random generation of test cases would be useless since the length of the generated list and the second argument would hardly coincide. Also, using standard symbolic execution might be difficult too since the search space is infinite and, moreover, due to the use of predicate `rev` that includes an accumulating parameter, there are goals that are not instances of any previous goal, thus requiring some powerful abstraction operators.

Using concolic execution, though, we can easily generate appropriate test cases.

First iteration. We start with a random initial goal, e.g., $\text{main}([a, b], s(0), R)$, where the input arguments $[1, 2]$ are assumed ground. The associated concrete execution is the following:

$$\begin{array}{l}
\langle \text{main}([a, b], s(0), R); \text{id}; [] \rangle \\
\stackrel{u(1,2)}{\rightarrow} \langle \text{length}([a, b], s(0)), \text{rev}([a, b], [], R); \text{id}; [(2; \{\mathbf{R/error}\}; \text{true})] \rangle \\
\stackrel{u(8)}{\rightarrow} \langle \text{length}([b], 0), \text{rev}([a, b], [], R); \text{id}; [(2; \{\mathbf{R/error}\}; \text{true})] \rangle \\
\stackrel{b(2)}{\rightarrow} \langle \text{true}; \{\mathbf{R/error}\}; [] \rangle
\end{array}$$

and its associated trace is thus $\tau = [u(1, 2), u(8), b(2)]$.

Now, we use concolic execution and produce the computation shown in Figure 5. Therefore, by executing $\text{main}([a, b], s(0), R)$ only clause (2) is completely evaluated. According to the previous algorithm for test case generation, we now consider the following prefix of the concolic execution:

$$\langle [u(1, 2), u(8), b(2)]; []; \text{main}(L, N, R); \text{id}; []; \{1, \dots, 8\} \rangle \stackrel{\{\}}{\rightarrow} \dots \langle [c(7, \{L'/[1], N'/0\})] \rangle \langle \dots \rangle$$

and the associated substitution $\{L/[X], N/s(0)\}$.

Second iteration. Now, we consider the goal $\text{main}([X], s(0), R)$. Since the first two arguments must be ground, as mentioned before, we apply a minimal ground-

```

⟨main([a], s(0), R); id; []⟩
   $\xrightarrow{u(1,2)}$  ⟨length([a], s(0)), rev([a], [], R); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(8)}$  ⟨length([], 0), rev([a], [], R); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(7)}$  ⟨rev([a], [], R); id; [2; ({R/error}; true)]⟩
   $\xrightarrow{u(4)}$  ⟨is_list([], rev([], [a], R)); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(5)}$  ⟨rev([], [a], R); id; [(2; {R/error}; true)]⟩
   $\xrightarrow{u(3)}$  ⟨true; {R/[a]}; [(2; {R/error}; true)]⟩

```

Fig. 6. Concrete execution for $\text{main}([a], s(0), R)$.

ing substitution and get, e.g., $\text{main}([a], s(0), R)$. The concrete execution, which is shown in Figure 6, computes the following trace:

$$\tau'' = [u(1, 2), u(8), u(7), u(4), u(5), u(3)]$$

Then, we use concolic execution again as shown in Figure 7. Therefore, according to the algorithm, we consider the following prefix of the concolic execution:

$$\langle [u(1, 2), u(8), \dots]; []; \text{main}(L, N, R); []; \{4, 5, 6, 8\} \rangle \xrightarrow{\{\}} \langle \dots \rangle \xrightarrow{\{c(7, \theta_1)\}} \langle \dots \rangle \xrightarrow{\{c(8, \theta_2)\}} \langle \dots \rangle$$

with the associated substitution $\sigma_1\theta_2 = \{L/[X, Y|L''], N/s(s(N''))\}$.

Third (and last) iteration As in the previous case, the instantiated goal, $\text{main}([X, Y|L''], s(s(N'')), R)$, is not ground enough according to its input mode and, thus, we apply a minimal grounding substitution. In this case, we get the initial goal $\text{main}([a, b], s(s(0)), R)$. Here, the concrete execution is basically the same shown in Figure 2, except for the last (backtracking) step. Therefore, the associated trace is

$$\tau''' = [u(1, 2), u(8), u(8), u(7), u(4), u(5), u(4), u(6), u(5), u(3)]$$

Now, concolic execution from the initial state

$$\langle [u(1, 2), u(8), u(8), u(7), u(4), u(5), u(4), u(6), u(5), u(3)]; []; \text{main}(L, N, R); id; []; \{6\} \rangle$$

proceeds similarly to the derivation shown in Figure 7, but now clause (6) is also completely evaluated, which means that the algorithm terminates successfully.

To summarize, concolic testing generated four test cases:

$$\begin{array}{ll} \text{main}([a, b], s(0), R) & \text{main}([a], s(0), R) \\ \text{main}([], 0, R) & \text{main}([a, b], s(s(0)), R) \end{array}$$

which suffice to cover the complete evaluation of all program clauses.

	$\langle [u(1, 2), u(8), \dots]; []; \text{main}(L, N, R); \text{id}; []; \{4, 5, 6, 8\} \rangle$
\Downarrow	$\langle [u(8), u(7), \dots]; [1]; \left(\begin{array}{l} \text{length}(L, N), \\ \text{rev}(L, [], R), e(1) \end{array} \right); \text{id}; \mathcal{S}; \{4, 5, 6, 8\} \rangle$
$\xrightarrow{\{c(7, \theta_1)\}}$	$\langle [u(7), u(4), \dots]; [8, 1]; \left(\begin{array}{l} \text{length}(L', N'), e(8), \\ \text{rev}([X]L', [], R), e(1) \end{array} \right); \sigma_1; \mathcal{S}; \{4, 5, 6, 8\} \rangle$
$\xrightarrow{\{c(8, \theta_2)\}}$	$\langle [u(4), u(5), u(3)]; [7, 8, 1]; \left(\begin{array}{l} e(7), e(8), \\ \text{rev}([X]L', [], R), e(1) \end{array} \right); \sigma_1\sigma_2; \mathcal{S}; \{4, 5, 6, 8\} \rangle$
\Downarrow	$\langle [u(4), u(5), u(3)]; [8, 1]; e(8), \text{rev}([X]L', [], R), e(1); \sigma_1\sigma_2; \mathcal{S}; \{4, 5, 6, 8\} \rangle$
\Downarrow	$\langle [u(4), u(5), u(3)]; [1]; \text{rev}([X]L', [], R), e(1); \sigma_1\sigma_2; \mathcal{S}; \{4, 5, 6\} \rangle$
\Downarrow	$\langle [u(5), u(3)]; [4, 1]; \left(\begin{array}{l} \text{is_list}([], \text{rev}(L', [X], R)), \\ e(4), e(1) \end{array} \right); \sigma_1\sigma_2; \mathcal{S}; \{4, 5, 6\} \rangle$
\Downarrow	$\langle [u(3)]; [5, 4, 1]; e(5), \text{rev}(L', [X], R), e(4), e(1); \sigma_1\sigma_2; \mathcal{S}; \{4, 5, 6\} \rangle$
\Downarrow	$\langle [u(3)]; [4, 1]; \text{rev}(L', [X], R), e(4), e(1); \sigma_1\sigma_2; \mathcal{S}; \{4, 6\} \rangle$
\Downarrow	$\langle []; [3, 4, 1]; e(3), e(4), e(1); \sigma_1\sigma_2; \mathcal{S}; \{4, 6\} \rangle$
\Downarrow	$\langle []; [4, 1]; e(4), e(1); \sigma_1\sigma_2; \mathcal{S}; \{4, 6\} \rangle$
\Downarrow	$\langle []; [1]; e(1); \sigma_1\sigma_2; \mathcal{S}; \{6\} \rangle$
\Downarrow	$\langle []; []; \text{true}; \sigma_1\sigma_2; \mathcal{S}; \{6\} \rangle$

with $\mathcal{S} = [(2; [2]; \{\mathbf{R}/\text{error}\}; e(2))];$

Fig. 7. Concolic execution for $\langle [u(1, 2), u(8), \dots]; \text{id}; []; \text{main}(L, N, R); []; \{4, 5, 6, 8\} \rangle$

In general, when the test case generation algorithm terminates, concolic testing is *sound* (i.e., there are no false positives since only concrete executions are considered) and *complete* (in the sense that all clauses are completely evaluated when using the computed test cases, i.e., we get a 100% coverage). When the process is stopped (e.g., because it does not terminate or takes too much time), our test case generation is only sound. Note that this contrasts with other approaches to test case generation in Prolog (and CLP), e.g., [8, 15], where full coverage is not considered. In [1], however, the authors consider a refined notion of coverage criterion: a pair $\langle TC, SC \rangle$, where TC is a termination criterion (e.g., the maximum number of recursive calls to a predicate allowed in symbolic execution) and SC is the selection criterion, used to determine which test cases must be produced. In particular, using the SC `program-points(P)`, where P includes all program points, amounts to the statement coverage that we consider in this paper. On the other hand, the authors have introduced a refinement for driving symbolic execution by means of “trace terms” (terms representing the shape of

a particular subset of the SLD search space). The idea behind this technique is closer to that of concolic execution, although they do not have the possibility of using concrete data in symbolic executions (which is one of the main advantages of concolic execution).

6 Concluding Remarks and Future Work

We have introduced a novel approach to Prolog testing and debugging. The so called concolic execution that mixes concrete and symbolic execution has been shown quite successful in other programming paradigms, especially when dealing with large applications. Therefore, it might have a great potential for Prolog testing, too.

In this paper, we have only considered a simple form of coverage, *statement* coverage, and a limited scenario: pure Prolog without negation. Nevertheless, the main distinctive features of the Prolog programming language—i.e., unification, non-determinism and backtracking—are present here, so adapting the standard concolic execution approach [6, 16] to Prolog was not trivial. The challenge, now, is experimentally verifying the effectiveness and scalability of our approach with real Prolog programs. For this purpose, though, we first need to extend concolic execution to deal with negation, built-in's, extra-logical features, etc. For this purpose, we will consider the linear operational semantics of [17], and its symbolic version [5], as a promising starting point.

Acknowledgements

The author gratefully acknowledges the anonymous referees and the participants of LOPSTR 2014 for many useful comments and suggestions. I would also like to thank Fred Mesnard and Etienne Payet for their remarks to improve the paper.

References

1. E. Albert, P. Arenas, M. Gómez-Zamalloa, and J.M. Rojas. Test Case Generation by Symbolic Execution: Basic Concepts, a CLP-Based Instance, and Actor-Based Concurrency. In *SFM 2014*, Springer LNCS 8483, pages 263–309, 2014.
2. F. Belli and O. Jack. Implementation-Based Analysis and Testing of Prolog Programs. In *ISSTA*, pages 70–80. ACM, 1993.
3. L.A. Clarke. A program testing system. In *Proceedings of the 1976 Annual Conference (ACM'76)*, pages 488–491, 1976.
4. D. De Schreye, R. Glück, J. Jørgensen, M. Leuschel, B. Martens, and M.H. Sørensen. Conjunctive Partial Deduction: Foundations, Control, Algorithms, and Experiments. *Journal of Logic Programming*, 41(2&3):231–277, 1999.
5. J. Giesl, T. Ströder, P. Schneider-Kamp, F. Emmes, and C. Fuhs. Symbolic evaluation graphs and term rewriting: a general methodology for analyzing logic programs. In *PPDP'12*, pages 1–12. ACM, 2012.
6. P. Godefroid, N. Klarlund, and K. Sen. DART: directed automated random testing. In *Proc. of PLDI'05*, pages 213–223. ACM, 2005.

7. P. Godefroid, M.Y. Levin, and D.A. Molnar. Sage: whitebox fuzzing for security testing. *Commun. ACM*, 55(3):40–44, 2012.
8. M. Gómez-Zamalloa, E. Albert, and G. Puebla. Test case generation for object-oriented imperative languages in CLP. *TPLP*, 10(4-6):659–674, 2010.
9. James C. King. Symbolic execution and program testing. *Commun. ACM*, 19(7):385–394, 1976.
10. M. Leuschel. The DPPD (Dozens of Problems for Partial Deduction) Library of Benchmarks, 2007. <http://www.ecs.soton.ac.uk/~mal/systems/dppd.html>.
11. J.W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, Berlin, 1987. Second edition.
12. J.W. Lloyd and J.C. Shepherdson. Partial Evaluation in Logic Programming. *Journal of Logic Programming*, 11:217–242, 1991.
13. B. Martens and J. Gallagher. Ensuring Global Termination of Partial Deduction while Allowing Flexible Polyvariance. In *Proc. of ICLP'95*, pages 597–611. MIT Press, 1995.
14. C.S. Pasareanu and N. Rungta. Symbolic PathFinder: symbolic execution of Java bytecode. In Charles Pecheur, Jamie Andrews, and Elisabetta Di Nitto, editors, *ASE*, pages 179–180. ACM, 2010.
15. J.M. Rojas and M. Gómez-Zamalloa. A Framework for Guided Test Case Generation in Constraint Logic Programming. In *Proc. of LOPSTR'12*, pages 176–193. Springer LNCS 7844, 2012.
16. K. Sen, D. Marinov, and G. Agha. CUTE: a concolic unit testing engine for C. In *Proc. of ESEC/SIGSOFT FSE 2005*, pages 263–272. ACM, 2005.
17. T. Ströder, F. Emmes, P. Schneider-Kamp, J. Giesl, and C. Fuhs. A Linear Operational Semantics for Termination and Complexity Analysis of ISO Prolog. In *LOPSTR'11*, pages 237–252. Springer LNCS 7225, 2011.