TLP **15** (4–5): 711–725, 2015. © Cambridge University Press 2015
doi:10.1017/S1471068415000332

*Concolic testing in logic programming**

FRED MESNARD, ÉTIENNE PAYET
LIM - Université de la Réunion, France
(*e-mail:* {fred,epayet}@univ-reunion.fr)

GERMÁN VIDAL
MiST, DSIC, Universitat Politècnica de València
(*e-mail:* gvidal@dsic.upv.es)

submitted 29 April 2015; revised 3 July 2015; accepted 15 July 2015

Abstract

Software testing is one of the most popular validation techniques in the software industry. Surprisingly, we can only find a few approaches to testing in the context of logic programming. In this paper, we introduce a systematic approach for dynamic testing that combines both concrete and symbolic execution. Our approach is fully automatic and guarantees full path coverage when it terminates. We prove some basic properties of our technique and illustrate its practical usefulness through a prototype implementation.

KEYWORDS: Symbolic execution, logic programming, testing.

1 Introduction

Essentially, software validation aims at ensuring that the developed software complies with the original requirements. One of the most popular validation approaches is *software testing*, a process that involves producing a test suite and then executing the system with these test cases. The main drawback of this approach is that designing a test suite with a high code coverage —i.e., covering as many execution paths as possible— is a complex and time-consuming task. As an alternative, one can use a tool for the random generation of test cases, but then we are often faced with a poor code coverage. Some hybrid approaches exist where random generation is driven by the user, as in QuickCheck (Claessen and Hughes 2000), but then again the process may become complex and time-consuming.

Another popular, fully automatic approach to test case generation is based on *symbolic execution* (King 1976; Clarke 1976). Basically, symbolic execution considers unknown (symbolic) values for the input parameters and, then, explores all feasible execution paths in a non-deterministic way. Symbolic states include now a *path*

* This work has been partially supported by the EU (FEDER) and the Spanish *Ministerio de Economía y Competitividad* under grant TIN2013-44742-C4-1-R and by the *Generalitat Valenciana* under grant PROMETEOII/2015/013. Part of this research was done while the third author was visiting the University of Reunion; G. Vidal gratefully acknowledges their hospitality.

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

condition that stores the current constraints on symbolic values, i.e., the conditions that must hold to reach a particular execution point. For each final state, a test case is produced by solving the constraints in the associated path condition.

A drawback of the previous approach, though, is that the constraints in the path condition may become very complex. When these constraints are not solvable, the only sound way to proceed is to stop the execution path, often giving rise to a poor coverage. Recently, a new variant called *concolic execution* (Godefroid *et al.* 2005; Sen *et al.* 2005) that combines both *conc*rete and symb*olic* execution has been proposed as a basis for both model checking and test case generation. The main advantage is that, now, when the constraints in the symbolic execution become too complex, one can still take some values from the concrete execution to simplify them. This is sound and often allows one to explore a larger execution space. Some successful tools that are based on concolic execution are, e.g., SAGE (Godefroid *et al.* 2012) and Java Pathfinder (Pasareanu and Rungta 2010).

In the context of the logic programming paradigm, one can find a flurry of *static*, complete techniques for software analysis and verification. However, only a few dynamic techniques for program validation have been proposed. Dynamic, typically incomplete, techniques have proven very useful for software validation in other paradigms. In general, these techniques are sound so that they avoid *false positives*. This contrasts with typical static verification methods which may produce some false positives due to the abstraction techniques introduced to ensure completeness. Therefore, we expect concolic execution to complement existing analysis and verification techniques for logic programs.

In this paper, we introduce a new, fully automatic scheme for *concolic testing* in logic programming. As in other paradigms, concolic testing may help the programmer to systematically find program bugs and generate test cases with a good code coverage. As it is common, our approach is always sound but usually incomplete. In the context of logic programming, we consider that "full path coverage" involves calling each predicate in all possible ways. Consider, e.g., the logic program $P = \{p(a)., \ p(b).\}$. Here, one could assume that the execution of the goals in $\{p(a), p(b)\}$ is enough for achieving a full path coverage. However, in this paper we consider that full path coverage requires, e.g., the set $\{p(X), p(a), p(b), p(c)\}$ so that we have a goal that matches both clauses, one that only matches the first clause, one that only matches the second clause, and one that matches no clause. We call this notion *choice* coverage, and it is specific of logic programming. To the best of our knowledge, such a notion of coverage has not been considered before. Typically, only a form of *statement* coverage has been considered, where only the clauses used in the considered executions are taken into account. For guaranteeing choice coverage, a new type of unification problems must be solved: we have to produce goals in which the selected atom A matches the heads of some clauses, say $H_1, \ldots, H_n$, but does not match the heads of some other clauses, say $H'_1, \ldots, H'_m$. We provide a constructive algorithm for solving such unifiability problems.

A prototype implementation of the concolic testing scheme for pure Prolog, called contest, is publicly available from `http://kaz.dsic.upv.es/contest.html`. The results from an experimental evaluation point out the usefulness of the approach.

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

Besides logic programming and Prolog, our technique might also be useful for other programming languages since there exist several transformational approaches that "compile in" programs to Prolog, like, e.g., (Gómez-Zamalloa *et al.* 2010).

Omitted proofs as well as some extensions can be found in the online appendix.

2 Concrete semantics

The semantics of a logic program is usually given in terms of the SLD relation on goals (Lloyd 1987). In this section, we present instead a *local* semantics which is similar to that of Ströder *et al.* (2011). Basically, this semantics deals with states that contain all the necessary information to perform the next step (in contrast to the usual semantics, where the SLD tree built so far is also needed, e.g., for dealing with the cut). In contrast to (Ströder *et al.* 2011), for simplicity, in this paper we only consider definite logic programs. However, the main difference w.r.t. (Ströder *et al.* 2011) comes from the fact that our concrete semantics only considers the computation of the first solution for the initial goal. This is the way most Prolog applications are used and, thus, our semantics should consider this behaviour in order to measure the coverage in a realistic way.

Before presenting the transition rules of the concrete semantics, let us introduce some auxiliary notions and notations. We refer the reader to (Apt 1997) for the standard definitions and notations for logic programs. The semantics is defined by means of a transition system on *states* of the form $\langle \mathscr{B}^1_{\delta_1} | \ldots | \mathscr{B}^n_{\delta_n} \rangle$, where $\mathscr{B}^1_{\delta_1} | \ldots | \mathscr{B}^n_{\delta_n}$ is a sequence of goals labeled with substitutions (the answer computed so far, when restricted to the variables of the initial goal). We denote sequences with $S, S', \ldots$, where ϵ denotes the empty sequence. In some cases, we label a goal $\mathscr{B}$ both with a substitution and a program clause, e.g., $\mathscr{B}^{H \leftarrow \mathscr{B}}_{\delta}$, which is used to determine the next clause to be used for an SLD resolution step (see rules choice and unfold in Fig. 1). Note that the clauses of the program are not included in the state but considered a global parameter since they are static. In the following, given an atom A and a logic program P, clauses(A, P) returns the sequence of renamed apart program clauses $c_1, \ldots, c_n$ from P whose head unifies with A. A syntactic object s_1 is *more general* than a syntactic object s_2, denoted $s_1 \leqslant s_2$, if there exists a substitution θ such that $s_1\theta = s_2$. $\mathscr{V}ar(o)$ denotes the set of variables of the syntactic object o. For a substitution θ, $\mathscr{V}ar(\theta)$ is defined as $Dom(\theta) \cup Ran(\theta)$.

For simplicity, w.l.o.g., we only consider *atomic* initial goals. Therefore, given an atom A, an initial state has the form $\langle A_{id} \rangle$, where id denotes the identity substitution. The transition rules, shown in Figure 1, proceed as follows:

- In rules success and failure, we use fresh constants to denote a final state: $\langle \text{SUCCESS}_\delta \rangle$ denotes that a sucessful derivation ended with computed answer substitution δ, while $\langle \text{FAIL}_\delta \rangle$ denotes a finitely failing derivation; recording δ for failing computations might be useful for debugging purposes.
- Rule backtrack applies when the first goal in the sequence finitely fails, but there is at least one alternative choice.

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

$$\text{(success)}\ \frac{}{\langle \mathsf{true}_\delta \,|\, S\rangle \rightarrow \langle \textsc{success}_\delta\rangle} \qquad \text{(failure)}\ \frac{}{\langle(\mathsf{fail},\mathscr{B})_\delta\rangle \rightarrow \langle \textsc{fail}_\delta\rangle}$$

$$\text{(backtrack)}\ \frac{S \neq \epsilon}{\langle(\mathsf{fail},\mathscr{B})_\delta \,|\, S\rangle \rightarrow \langle S\rangle}$$

$$\text{(choice)}\ \frac{\mathsf{clauses}(A,\mathscr{P}) = (c_1,\ldots,c_n) \wedge n > 0}{\langle(A,\mathscr{B})_\delta \,|\, S\rangle \rightarrow \langle(A,\mathscr{B})_\delta^{c_1} \,|\ldots|\, (A,\mathscr{B})_\delta^{c_n} \,|\, S\rangle} \qquad \text{(choice_fail)}\ \frac{\mathsf{clauses}(A,\mathscr{P}) = \{\}}{\langle(A,\mathscr{B})_\delta \,|\, S\rangle \rightarrow \langle(\mathsf{fail},\mathscr{B})_\delta \,|\, S\rangle}$$

$$\text{(unfold)}\ \frac{\mathsf{mgu}(A,H_1) = \sigma}{\langle(A,\mathscr{B})_\delta^{H_1 \leftarrow \mathscr{B}_1} \,|\, S\rangle \rightarrow \langle(\mathscr{B}_1\sigma, \mathscr{B}\sigma)_{\delta\sigma} \,|\, S\rangle}$$

Fig. 1. Concrete semantics

- Rule choice represents the first stage of an SLD resolution step. If there is at least one clause whose head unifies with the leftmost atom, this rule introduces as many copies of a goal as clauses returned by function clauses. If there is at least one matching clause, unfolding is then performed by rule unfold. Otherwise, if there is no matching clause, rule choice_fail returns fail so that either rule failure or backtrack applies next.

Example 1
Consider the following logic program:

$$\begin{array}{lll} p(s(a)). & q(a). & r(a). \\ p(s(X)) \leftarrow q(X). & q(b). & r(c). \\ p(f(X)) \leftarrow r(X). & & \end{array}$$

Given the initial goal $p(f(X))$, we have the following successful computation (for clarity, we label each step with the applied rule):

$$\begin{array}{lllll} \langle p(f(X))_{id}\rangle & \rightarrow^{\mathsf{choice}} & \langle p(f(X))_{id}^{p(f(Y))\leftarrow r(Y)}\rangle & \rightarrow^{\mathsf{unfold}} & \langle r(X)_{id}\rangle \\ & \rightarrow^{\mathsf{choice}} & \langle r(X)_{id}^{r(a)} \,|\, r(X)_{id}^{r(c)}\rangle & \rightarrow^{\mathsf{unfold}} & \langle \mathsf{true}_{\{X/a\}} \,|\, r(X)_{id}^{r(c)}\rangle \\ & \rightarrow^{\mathsf{success}} & \langle \textsc{success}_{\{X/a\}}\rangle & & \end{array}$$

Therefore, we have a successful computation for $p(f(X))$ with computed answer $\{X/a\}$. Observe that only the first answer is considered.

We do not formally prove the correctness of the concrete semantics, but it is an easy consequence of the correctness of the semantics in (Ströder *et al.* 2011). Note that our rules can be seen as an instance for pure Prolog without negation, where only the computation of the first answer for the initial goal is considered.

3 Concolic execution semantics

In this section, we introduce a concolic execution semantics for logic programs that is a conservative extension of the concrete semantics of the previous section. In this semantics, *concolic states* have the form $\langle S \,][\, S'\rangle$, where S and S' are sequences of (possibly labeled) concrete and symbolic goals, respectively. In logic programming, the notion of *symbolic* execution is very natural: the structure of both S and S' is

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

$$\text{(success)}\ \frac{}{\langle \mathsf{true}_\delta \mid S \rangle\!\rangle\!\langle \mathsf{true}_\theta \mid S' \rangle \rightsquigarrow_\diamond \langle \textsc{success}_\delta \rangle\!\rangle\!\langle \textsc{success}_\theta \rangle}$$

$$\text{(failure)}\ \frac{}{\langle (\mathsf{fail}, \mathscr{B})_\delta \rangle\!\rangle\!\langle (\mathsf{fail}, \mathscr{B}')_\theta \rangle \rightsquigarrow_\diamond \langle \textsc{fail}_\delta \rangle\!\rangle\!\langle \textsc{fail}_\theta \rangle}$$

$$\text{(backtrack)}\ \frac{S \neq \epsilon}{\langle (\mathsf{fail}, \mathscr{B})_\delta \mid S \rangle\!\rangle\!\langle (\mathsf{fail}, \mathscr{B}')_\theta \mid S' \rangle \rightsquigarrow_\diamond \langle S \rangle\!\rangle\!\langle S' \rangle}$$

$$\text{(choice)}\ \frac{\mathsf{clauses}(A, \mathscr{P}) = \overline{c_n} \wedge n > 0 \wedge \mathsf{clauses}(A', \mathscr{P}) = \overline{d_k}}{\langle (A, \mathscr{B})_\delta \mid S \rangle\!\rangle\!\langle (A', \mathscr{B}')_\theta \mid S' \rangle \rightsquigarrow_{c(\ell(\overline{c_n}), \ell(\overline{d_k}))} \langle (A, \mathscr{B})_\delta^{c_1} \mid \ldots \mid (A, \mathscr{B})_\delta^{c_n} \mid S \rangle\!\rangle\!\langle (A', \mathscr{B}')_\theta^{c_1} \mid \ldots \mid (A', \mathscr{B}')_\theta^{c_n} \mid S' \rangle}$$

$$\text{(choice_fail)}\ \frac{\mathsf{clauses}(A, \mathscr{P}) = \{\} \wedge \mathsf{clauses}(A', \mathscr{P}) = \overline{c_k}}{\langle (A, \mathscr{B})_\delta \mid S \rangle\!\rangle\!\langle (A', \mathscr{B}')_\theta \mid S' \rangle \rightsquigarrow_{c(\{\}, \ell(\overline{c_k}))} \langle (\mathsf{fail}, \mathscr{B})_\delta \mid S \rangle\!\rangle\!\langle (\mathsf{fail}, \mathscr{B}')_\theta \mid S' \rangle}$$

$$\text{(unfold)}\ \frac{\mathsf{mgu}(A, H_1) = \sigma \wedge \mathsf{mgu}(A', H_1) = \sigma'}{\langle (A, \mathscr{B})_\delta^{H_1 \leftarrow \mathscr{B}_1} \mid S \rangle\!\rangle\!\langle (A', \mathscr{B}')_\theta^{H_1 \leftarrow \mathscr{B}_1} \mid S' \rangle \rightsquigarrow_\diamond \langle (\mathscr{B}_1\sigma, \mathscr{B}\sigma)_{\delta\sigma} \mid S \rangle\!\rangle\!\langle (\mathscr{B}_1\sigma', \mathscr{B}'\sigma')_{\theta\sigma'} \mid S' \rangle}$$

Fig. 2. Concolic execution semantics

the same, and the only difference is that some atoms might be less instantiated in S' than in S.

In the following, we let $\overline{o_n}$ denote the sequence of syntactic objects $o_1, \ldots, o_n$. Given an atom A, we let $root(A) = p/n$ if $A = p(\overline{t_n})$. Now, given an atom A with $root(A) = p/n$, an *initial concolic state* has the form $\langle A_{id} \rangle\!\rangle\!\langle p(\overline{X_n})_{id} \rangle$, where $\overline{X_n}$ are different fresh variables. In the following, we assume that every clause c has a corresponding unique label, which we denote by $\ell(c)$. By abuse of notation, we also denote by $\ell(\overline{c_n})$ the set of labels $\{\ell(c_1), \ldots, \ell(c_n)\}$.

The semantics is given by the rules of the labeled transition relation $\rightsquigarrow$ shown in Figure 2. Here, we consider two kinds of labels for the transition relation:

- The empty label, $\diamond$, which is often implicit.
- A label of the form $c(\ell(\overline{c_n}), \ell(\overline{d_k}))$, which represents a choice step. Here, $\ell(\overline{c_n})$ are the labels of the clauses matching the selected atom in the concrete goal, while $\ell(\overline{d_k})$ are the labels of the clauses matching the selected atom in the corresponding symbolic goal. Note that $\ell(\overline{c_n}) \subseteq \ell(\overline{d_k})$ since the concrete goal is always an instance of the symbolic goal (see Theorem 1 below).

For each transition step $\mathscr{C}_1 \rightsquigarrow_{c(\mathscr{L}_1, \mathscr{L}_2)} \mathscr{C}_2$, the first set of labels, $\mathscr{L}_1$, is used to determine the execution *trace* of a concrete goal (see below). Traces are needed to keep track of the execution paths already explored. The second set of labels, $\mathscr{L}_2$, is used to compute new goals that follow alternative paths not yet explored, if any.

In the concolic execution semantics, we perform both concrete and symbolic execution steps in parallel. However, the symbolic execution does not explore all possible execution paths but only mimics the steps of the concrete execution; observe, e.g., rule choice in Figure 2, where the clauses labeling the copies of the symbolic goal are the same clauses $\overline{c_n}$ matching the concrete goal, rather than the set of clauses $\overline{d_k}$ (a superset of $\overline{c_n}$).

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

Example 2
Consider again the logic program of Example 1, now with clause labels:

$$\begin{array}{lll}(\ell_1)\ p(s(a)). & (\ell_4)\ q(a). & (\ell_6)\ r(a).\\ (\ell_2)\ p(s(X)) \leftarrow q(X). & (\ell_5)\ q(b). & (\ell_7)\ r(c).\\ (\ell_3)\ p(f(X)) \leftarrow r(X). & & \end{array}$$

Given the initial goal $p(f(X))$, we have the following concolic execution:

$$\begin{array}{ll}\langle p(f(X))_{id} \,][\, p(N)_{id}\rangle & \leadsto^{\mathsf{choice}}_{c(\mathscr{L}_1,\mathscr{L}'_1)} \langle p(f(X))^{p(f(Y))\leftarrow r(Y)}_{id} \,][\, p(N)^{p(f(Y))\leftarrow r(Y)}_{id}\rangle\\ & \leadsto^{\mathsf{unfold}}_{\diamond} \langle r(X)_{id} \,][\, r(Y)_{\{N/f(Y)\}}\rangle\\ & \leadsto^{\mathsf{choice}}_{c(\mathscr{L}_2,\mathscr{L}'_2)} \langle r(X)^{r(a)}_{id} \,|\, r(X)^{r(c)}_{id} \,][\, r(Y)^{r(a)}_{\{N/f(Y)\}} \,|\, r(Y)^{r(c)}_{\{N/f(Y)\}}\rangle\\ & \leadsto^{\mathsf{unfold}}_{\diamond} \langle \mathsf{true}_{\{X/a\}} \,|\, r(X)^{r(c)}_{id} \,][\, \mathsf{true}_{\{N/f(a)\}} \,|\, r(Y)^{r(c)}_{\{N/f(Y)\}}\rangle\\ & \leadsto^{\mathsf{success}}_{\diamond} \langle \textsc{success}_{\{X/a\}} \,][\, \textsc{success}_{\{N/f(a)\}}\rangle\end{array}$$

where $\mathscr{L}_1 = \{\ell_3\}$, $\mathscr{L}'_1 = \{\ell_1, \ell_2, \ell_3\}$, and $\mathscr{L}_2 = \mathscr{L}'_2 = \{\ell_6, \ell_7\}$.

In this paper, we only consider finite concolic executions for initial goals. This is a reasonable assumption since one can expect concrete goals to compute the first answer finitely (unless the program is erroneous). We associate a *trace* to each concolic execution as follows:

Definition 1 (*trace*)
Let P be a program and $\mathscr{C}_0$ an initial concolic state. Let $E = (\mathscr{C}_0 \leadsto_{l_1} \ldots \leadsto_{l_m} \mathscr{C}_m)$, $m > 0$, be a concolic execution for $\mathscr{C}_0$ in P. Let $c(\mathscr{L}_1, \mathscr{L}'_1), \ldots, c(\mathscr{L}_k, \mathscr{L}'_k)$, $k \leqslant m$, be the sequence of labels in $l_1, \ldots, l_m$ which are different from $\diamond$. Then, the trace associated to the concolic execution E is $\mathit{trace}(E) = \mathscr{L}_1, \ldots, \mathscr{L}_k$.

Roughly speaking, a trace is just a sequence with the sets of labels of the matching clauses in each choice step. For instance, the trace associated to the concolic execution of Example 2 is $(\{\ell_3\}, \{\ell_6, \ell_7\})$, i.e., we have two unfolding steps with matching clauses $\{\ell_3\}$ and $\{\ell_6, \ell_7\}$, respectively. Note that traces ending with $\{\,\}$ represent failing derivations.

The following result states an essential invariant for concolic execution:

Theorem 1
Let P be a program and $\mathscr{C}_0 = \langle p(\overline{t_n})_{id} \,][\, p(\overline{X_n})_{id}\rangle$ be an initial concolic state. Let $\mathscr{C}_0 \leadsto \ldots \leadsto \mathscr{C}_m$, $m \geqslant 0$, be a finite (possibly incomplete) concolic execution for $\mathscr{C}_0$ in P. Then, for all concolic states $C_i = \langle \mathscr{B}^c_\delta \,|\, S \,][\, \mathscr{D}^{c'}_\theta \,|\, S'\rangle$, $i = 0, \ldots, m$, the following invariant holds: $|S| = |S'|$, $\mathscr{D} \leqslant \mathscr{B}$, $c = c'$ (if any), and $p(\overline{X_n})\theta \leqslant p(\overline{t_n})\delta$.

4 Concolic testing

In this section, we introduce a concolic testing procedure for logic programs based on the concolic execution semantics of the previous section.

4.1 The procedure

As we have seen in Section 3, the concolic execution steps labeled with $c(\mathscr{L}_1, \mathscr{L}_2)$ give us a hint of (potential) alternative execution paths. Consider, for instance, the

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

concolic execution of Example 2. The first step is labeled with $c(\{\ell_3\},\{\ell_1,\ell_2,\ell_3\})$. This means that the selected atom in the concrete goal only matched clause ℓ_3, while the selected atom in the symbolic goal matched clauses ℓ_1, ℓ_2 and ℓ_3. In principle, there are as many alternative execution paths as elements in $\mathcal{P}(\{\ell_1,\ell_2,\ell_3\}) \setminus \{\ell_3\}$; e.g., $\{\}$ denotes an execution path where the selected atom matches no clause, $\{\ell_1\}$ another path in which the selected atom *only* matches clause ℓ_1, $\{\ell_1,\ell_2,\ell_3\}$ another path where the selected atom matches all three clauses ℓ_1, ℓ_2 and ℓ_3, and so forth.

When aiming at full choice coverage, we need to solve both unification and disunification problems. Consider, e.g., that A is the selected atom in a goal, and that we want it to unify with the head of clause ℓ_1 but not with the heads of clauses ℓ_2 and ℓ_3. For this purpose, we introduce the following auxiliary function alt, which also includes some groundness requirements (see below). In the following, we let $\approx$ denote the unifiability relation, i.e., given atoms A, B, $A \approx B$ holds if $\mathsf{mgu}(A,B) \neq \mathsf{fail}$; correspondingly, $\neg(A \approx B)$ holds if $\mathsf{mgu}(A,B) = \mathsf{fail}$.

Definition 2 (alt)
Let A be an atom and $\mathscr{L}, \mathscr{L}'$ be sets of clause labels. Let $\mathscr{V}$ be a set of variables. The function $\mathsf{alt}(A, \mathscr{L}, \mathscr{L}', \mathscr{V})$ returns a substitution θ such that the following holds:

$$A\theta \approx H_1 \wedge \ldots \wedge A\theta \approx H_n \wedge \neg(A\theta \approx H_{n+1}) \wedge \ldots \wedge \neg(A\theta \approx H_m) \wedge \mathscr{V}\theta \text{ are ground}$$

where $H_1,\ldots,H_n$ are the heads of the (renamed apart) clauses labeled by $\mathscr{L}$ and $H_{n+1},\ldots,H_m$ are the heads of the (renamed apart) clauses labeled by $\mathscr{L}' \setminus \mathscr{L}$, respectively. If such a substitution does not exist, then function alt returns fail.

When the considered signature is finite,[1] the following semi-algorithm is trivially sound and complete for solving the above unifiability problem: first, bind A with terms of depth 0.[2] If the condition above does not hold, then we try with terms of depth 1, and check it again. We keep increasing the considered term depth until a solution is found. If a solution exists, this naive semi-algorithm will find it (otherwise, it may run forever). In practice, however, it may be very inefficient.

Observe that, in general, there might be several most general solutions to the above problem. Consider, e.g., $A = p(X,Y)$, $\mathscr{H}_{pos} = \{p(Z,Z), p(a,b)\}$ and $\mathscr{H}_{neg} = \{p(c,c)\}$. Then, both $p(a,U)$ and $p(U,b)$ are most general solutions. In principle, any of them is equally good in our context. We postpone to the next section the introduction of a constructive algorithm for function alt. Here, we present an algorithm to systematically produce concrete initial goals so that all *feasible* choices in the execution paths are covered (unless the process runs forever). First, we introduce the following auxiliary definitions:

Definition 3 (conc, symb)
Let $\mathscr{C} = \langle \mathscr{B}^1_{\delta_1} | \ldots | \mathscr{B}^n_{\delta_n} \rangle\!\rangle\, \mathscr{D}^1_{\theta_1} | \ldots | \mathscr{D}^n_{\theta_n} \rangle$ be a concolic state. Then, we let $\mathsf{conc}(\mathscr{C}) = \mathscr{B}^1_{\delta_1}$ denote the first concrete goal and $\mathsf{symb}(\mathscr{C}) = \mathscr{D}^1_{\theta_1}$ the first symbolic goal.

[1] Full Prolog and infinite signatures like integers or real numbers are left as future work.
[2] The depth $depth(t)$ of a term t is defined as usual: $depth(t) = 0$ if t is a variable or a constant symbol, and $depth(f(t_1,\ldots,t_n)) = 1 + \max(depth(t_1),\ldots,depth(t_n))$, otherwise.

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

Definition 4 (alt_trace)

Let P be a program, $\mathscr{C}_0$ an initial concolic state, and $E = (\mathscr{C}_0 \rightsquigarrow_{l_1} \ldots \rightsquigarrow_{l_n} \mathscr{C}_n \rightsquigarrow_{c(\mathscr{L},\mathscr{L}')} \mathscr{C}_{n+1})$ be a (possibly incomplete) concolic execution for $\mathscr{C}_0$ in P. Then, the function alt_trace denotes the following set of (potentially) alternative traces:

$$\mathsf{alt_trace}(E) = \{\mathscr{L}_1, \ldots, \mathscr{L}_k, \mathscr{L}'' \mid \begin{array}{l} trace(\mathscr{C}_0 \rightsquigarrow_{l_1} \ldots \rightsquigarrow_{l_n} \mathscr{C}_n) = \mathscr{L}_1, \ldots, \mathscr{L}_k \\ \text{and } \mathscr{L}'' \in (\mathscr{P}(\mathscr{L}') \setminus \mathscr{L}) \end{array} \}$$

For instance, given the following (partial) concolic execution E from Example 2:

$$\begin{array}{lll} \langle p(f(X))_{id} \,[\![\, p(N)_{id} \rangle & \rightsquigarrow^{\mathsf{choice}}_{c(\mathscr{L}_1,\mathscr{L}'_1)} & \langle p(f(X))^{p(f(Y))\leftarrow r(Y)}_{id} \,[\![\, p(N)^{p(f(Y))\leftarrow r(Y)}_{id} \rangle \\ & \rightsquigarrow^{\mathsf{unfold}}_{\diamond} & \langle r(X)_{id} \,[\![\, r(Y)_{\{N/f(Y)\}} \rangle \\ & \rightsquigarrow^{\mathsf{choice}}_{c(\mathscr{L}_2,\mathscr{L}'_2)} & \langle r(X)^{r(a)}_{id} \,|\, r(X)^{r(c)}_{id} \,[\![\, r(Y)^{r(a)}_{\{N/f(Y)\}} \,|\, r(Y)^{r(c)}_{\{N/f(Y)\}} \rangle \end{array}$$

where $\mathscr{L}_1 = \{\ell_3\}$, $\mathscr{L}'_1 = \{\ell_1, \ell_2, \ell_3\}$, $\mathscr{L}_2 = \mathscr{L}'_2 = \{\ell_6, \ell_7\}$, we have $trace(E) = \mathscr{L}_1, \mathscr{L}_2$, $\mathscr{P}(\mathscr{L}'_2) \setminus \mathscr{L}_2 = \{\{\,\}, \{\ell_6\}, \{\ell_7\}\}$, and $\mathsf{alt_trace}(E) = \{(\mathscr{L}_1, \{\,\}), (\mathscr{L}_1, \{\ell_6\}), (\mathscr{L}_1, \{\ell_7\})\}$.

Now, we introduce our concolic testing procedure. It takes as input a program and a random —e.g., provided by the user— initial atomic goal rooted by the distinguished predicate $main/n$. In the following, we assume that each concrete initial goal $main(\overline{t_n})$ is existentially terminating w.r.t. Prolog's leftmost computation rule, i.e., either computes the *first* answer in a finite number of steps or finitely fails (Vasak and Potter 1986). For this purpose, we assume that $main/n$ has some associated *input* arguments, determined by a function input, so that an initial goal $main(\overline{t_n})$ existentially terminates if the terms $\mathsf{input}(main(\overline{t_n}))$ are ground. One could also consider that there are several combinations of input arguments that guarantee existential termination —this is similar to the modes of a predicate— but we only consider one set of input arguments for simplicity (extending the concolic testing algorithm would be straightforward). As mentioned before, assuming that concrete initial goals are existentially terminating is a reasonable assumption in practice.

Definition 5 (*concolic testing*)

Input: a logic program P and an atom $main(\overline{t_n})$ with $\mathsf{input}(main(\overline{t_n}))$ ground.

Output: a set TC of test cases.

1. Let $Pending := \{main(\overline{t_n})\}$, $TC := \{\}$, $Traces := \{\}$.
2. While $|Pending| \neq 0$ do

(a) Take $A \in Pending$, $Pending := Pending \setminus \{A\}$, $TC := TC \cup \{A\}$.

(b) Let $\mathscr{C}_0 = \langle A_{id} \,[\![\, main(\overline{X_n})_{id} \rangle$ and compute a successful or finitely failing derivation $E = (\mathscr{C}_0 \rightsquigarrow_{l_1} \ldots \rightsquigarrow_{l_m} \mathscr{C}_m)$.

(c) Let $Traces := Traces \cup \{trace(E)\}$.

(d) We update $Pending$ as follows:

- for each prefix $\mathscr{C}_0 \rightsquigarrow_{l_1} \ldots \rightsquigarrow_{l_j} \mathscr{C}_j \rightsquigarrow_{c(\mathscr{L},\mathscr{L}')} \mathscr{C}_{j+1}$ of E and
- for each (possibly partial) trace $\overline{\mathscr{L}_k}, \mathscr{L}_{k+1} \in \mathsf{alt_trace}(\mathscr{C}_0 \rightsquigarrow_{l_1} \ldots \rightsquigarrow_{l_j} \mathscr{C}_j \rightsquigarrow_{c(\mathscr{L},\mathscr{L}')} \mathscr{C}_{j+1})$ which is not the prefix of any trace in $Traces$,

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

- add $main(\overline{X_n})\theta\theta'$ to *Pending* if $\mathsf{alt}(A_1, \mathscr{L}_{k+1}, \mathscr{L}', G) = \theta' \neq \mathsf{fail}$, where $G = \mathscr{V}ar(\mathsf{input}(main(\overline{X_n})\theta))$ and $\mathsf{symb}(\mathscr{C}_j) = (A_1, \mathscr{B})_\theta$.[3]

3. Return the set TC of test cases

The soundness of concolic testing is immediate, since each atom from TC is indeed a test case of the form $main(\overline{s_n})$ with $\mathsf{input}(main(\overline{s_n}))$ ground. Completeness and termination are more subtle properties though.

In principle, one could argue that the concolic testing algorithm is a complete semi-algorithm in the sense that, if it terminates, the generated test cases cover all feasible paths. Our assumptions trivially guarantee that every considered concrete execution is finite (i.e., step (2b) in the loop of the concolic testing algorithm). Unfortunately, the algorithm will often run forever by producing infinitely many test cases. Consider, e.g., the following simple program:

$$(\ell_1)\ nat(0). \qquad (\ell_2)\ nat(s(X)) \leftarrow nat(X).$$

Even if every goal $nat(t)$ with t ground is terminating, our algorithm will still produce infinitely many test cases, e.g., $nat(0)$, $nat(s(0))$, $nat(s(s(0)))$, ..., since each goal will explore a different path (i.e., will produce a different execution trace: $(\{\ell_1\})$, $(\{\ell_2\}, \{\ell_1\})$, $(\{\ell_2\}, \{\ell_2\}, \{\ell_1\})$, etc). In practice, though, the quality of the generated test cases should be experimentally evaluated using a coverage tool.

Therefore, in general, we will sacrifice completeness in order to guarantee the termination of concolic testing. For this purpose, one can use a time limit, a bound for the length of concolic executions, or a maximum term depth for the arguments of the generated test cases. In this paper, we consider the last approach. Then, one can replace the use of a particular function alt in step (2d) of Definition 5 by a function alt_k with $\mathsf{alt}_k(A, \mathscr{L}, \mathscr{L}', G) = \mathsf{alt}(A, \mathscr{L}, \mathscr{L}', G) = \theta$ if $depth(t) \leqslant k$ for all $X/t \in \theta$, and $\mathsf{alt}_k(A, \mathscr{L}, \mathscr{L}', G) = \mathsf{fail}$ otherwise. This is the solution we implemented in the concolic testing tool described in Section 4.3.

For instance, by requiring a maximum term depth of 1, the generated test cases for the program *nat* above would be $nat(0)$, $nat(1)$, $nat(s(0))$ and $nat(s(1))$, where 1 is a fresh constant symbol, with associated traces $(\{\ell_1\})$, $(\{\ \})$, $(\{\ell_2\}, \{\ell_1\})$, and $(\{\ell_2\}, \{\ \})$, respectively.

Termination of the algorithm in Definition 5 is then guaranteed since only a finite number of new atoms can be added in step (2d) —up to variable renaming— and, moreover, only those (possibly partial) traces which are not a prefix of any trace already in the set *Traces* are considered. Observe that these facts suffice to ensure termination of the algorithm since one cannot have infinitely many traces with a finite number of atoms.

4.2 Solving unifiability problems

In this section, we present a constructive algorithm for function alt. Let us first reformulate our unification problem in slightly more general terms than in

[3] I.e., A_1 is the first atom of the symbolic goal $\mathsf{symb}(\mathscr{C}_j)$ of the concolic state $\mathscr{C}_j$, see Definition 3.

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

Definition 2. Let A be an atom and $\mathscr{H}_{pos}$, $\mathscr{H}_{neg}$ be two sets of atoms the elements of which are variable disjoint with A and unify with A, and a set of variables G. The problem consists in finding a substitution σ such that

$$\forall H^+ \in \mathscr{H}_{pos}.\ A\sigma \approx H^+ \wedge \forall H^- \in \mathscr{H}_{neg}.\ \neg(A\sigma \approx H^-),\ \text{and}\ G\sigma\ \text{is ground} \quad (*)$$

We introduce a stepwise method that, roughly speaking, proceeds as follows:

- First, we produce some "maximal" substitutions θ (called *maximal unifying substitution* below) for A such that $A\theta$ still unifies with the atoms in $\mathscr{H}_{pos}$. Here, we use a special set $\mathscr{U}$ of fresh variables with $\mathscr{V}ar(\{A\} \cup \mathscr{H}_{pos} \cup \mathscr{H}_{neg}) \cap \mathscr{U} = \{\}$. The elements of $\mathscr{U}$ are denoted by U, U', $U_1 \ldots$ Then, in θ, the variables from $\mathscr{U}$ (if any) denote positions where further binding will prevent $A\theta$ from unifying with some atom in $\mathscr{H}_{pos}$. In contrast, $A\theta\sigma'$ still unifies with all the atoms in $\mathscr{H}_{pos}$ as long as σ' does not bind any variable from $\mathscr{U}$. Roughly speaking, we apply some (minimal) generalizations to the atoms in $\mathscr{H}_{pos}$ so that they unify, and then return their most general unifier.
  For this stage, we use well known techniques like *variable elimination* (Martelli and Montanari 1982) and *generalization* (from the algorithm for most specific generalization (Plotkin 1970)); see Definition 6 below.
- In a second stage, we look for another substitution η such that $\theta\eta$ is a solution for $(*)$. Here, we basically follow a generate and test algorithm (as in the naive algorithm above), but it is now much more restricted thanks to θ.

4.2.1 The positive atoms

Here, we will use the variables from the special set $\mathscr{U}$ to replace *disagreement pairs* (see (Apt 1997) p. 27). Roughly speaking, given terms s and t, a subterm s' of s and a subterm t' of t form a disagreement pair if the root symbols of s' and t' are different, but the symbols from s' up to the root of s and from t' up to the root of t are the same. For instance, $X, g(a)$ and $b, h(Y)$ are disagreement pairs of the terms $f(X, g(b))$ and $f(g(a), g(h(Y)))$. A disagreement pair t, t' is called *simple* if one of the terms is a variable that does not occur in the other term and no variable of $\mathscr{U}$ occurs in t, t'. We say that the substitution $\{X/s\}$ is determined by t, t' if $\{X, s\} = \{t, t'\}$.

Basically, given an atom A and a set of atoms $\mathscr{H}_{pos}$, the following algorithm nondeterministically computes a substitution θ such that $A\theta\sigma'$ still unifies with all the atoms in $\mathscr{H}_{pos}$ as long as σ' does not bind any variable from $\mathscr{U}$.

Definition 6 (*maximal unifying substitution*)
Input: an atom A and a set of atoms $\mathscr{H}_{pos}$ such that $\mathscr{V}ar(\{A\} \cup \mathscr{H}_{pos}) \cap \mathscr{U} = \{\}$ and $A \approx B$ for all $B \in \mathscr{H}_{pos}$.
Output: a substitution θ.

1. Let $\mathscr{B} := \{A\} \cup \mathscr{H}_{pos}$.
2. While simple disagreement pairs occur in $\mathscr{B}$ do

(a) nondeterministically choose a simple disagreement pair X, t (resp. t, X) in $\mathscr{B}$ such that there is no other simple disagreement pair of the form X, t' (or t', X) with $t < t'$ (i.e., a strict instance);

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

(b) set $\mathscr{B}$ to $\mathscr{B}\eta$ where $\eta = \{X/t\}$.

3. While $|\mathscr{B}| \neq 1$ do

(a) nondeterministically choose a disagreement pair t, t' in $\mathscr{B}$;
(b) replace all disagrement pairs t, t' in $\mathscr{B}$ by a fresh variable of $\mathscr{U}$.

4. Return θ, where $\mathscr{B} = \{B\}$, $A\theta = B$, and $Dom(\theta) \subseteq \mathscr{V}ar(A)$.

We note that the algorithm assumes that the input atom A is always more general than the final atom B so that the last step is well defined. An invariant proving that this is indeed the case can be found in the online appendix (Appendix B).

Observe that the step (2a) is nondeterministic since there may exist several disagreement pairs X, t (or t, X) for the same variable X. Consider the atom $A = p(X, Y)$ and the set $\mathscr{H}_{pos} = \{p(a, b), p(Z, Z)\}$. Then, both $\{X/a, Y/U\}$ and $\{X/U, Y/b\}$ are maximal unifying substitutions, as the following example illustrates:

Example 3
Let $A = p(X, Y)$ and $\mathscr{H}_{pos} = \{p(a, b), p(Z, Z)\}$, with $\mathscr{B} := \{p(X, Y), p(a, b), p(Z, Z)\}$. The algorithm then considers the simple disagreement pairs in $\mathscr{B}$. From X, a, we get $\eta_1 := \{X/a\}$ and the action (2b) sets $\mathscr{B}$ to $\mathscr{B}\eta_1 = \{p(a, Y), p(a, b), p(Z, Z)\}$. The substitution $\eta_2 := \{Y/b\}$ is determined by Y, b and the action (2b) sets $\mathscr{B}$ to $\mathscr{B}\eta_2 = \{p(a, b), p(Z, Z)\}$. Now, we have two non-deterministic possibilities:

- If we consider the disagreement pair a, Z, we have a substitution $\eta_3 := \{Z/a\}$ and Action (2b) then sets $\mathscr{B}$ to $\mathscr{B}\eta_3 = \{p(a, b), p(a, a)\}$. Now, no simple disagreement pair occurs in $\mathscr{B}$, hence the algorithm jumps to the loop at line 3. Action (3b) replaces the disagreement pair b, a with a fresh variable $U \in \mathscr{U}$, hence $\mathscr{B}$ is set to $\{p(a, U)\}$. As $|\mathscr{B}| = 1$ the loop at line 3 stops and the algorithm returns the substitution $\{X/a, Y/U\}$.
- If we consider the disagreement pair b, Z instead, we have a substitution $\eta_3' := \{Z/b\}$, and Action (2b) sets $\mathscr{B}$ to $\mathscr{B}\eta_3' = \{p(a, b), p(b, b)\}$. Now, by proceeding as in the previous case, the algorithm returns $\{X/U, Y/b\}$.

4.2.2 The negative atoms

Now we deal with the negative atoms by means of the following algorithm which is the basis of our implementation of function alt:

Definition 7 (*PosNeg*)
Input: an atom A and two sets of atoms $\mathscr{H}_{pos}$, $\mathscr{H}_{neg}$, the elements of which are variable disjoint with A and unify with A, and a set of variables G.
Output: fail or a substitution $\theta\eta$ (restricted to the variables of A).

1. Let θ be the substitution returned by the algorithm of Definition 6 with input A and $\mathscr{H}_{pos}$.
2. Let η be an idempotent substitution such that $G\theta\eta$ is ground.
3. Check that $Dom(\eta) \subseteq \mathscr{V}ar(A\theta)$ and $\mathscr{V}ar(\eta) \cap \mathscr{U} = \{\}$, otherwise return fail.

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

4. Check that for each $H^- \in \mathscr{H}_{neg}$, $\neg(A\theta\eta \approx H^-)$, otherwise return fail.
5. Return $\theta\eta$ (restricted to the variables of A).

The correctness of this algorithm is stated as follows:

Theorem 2
Let A be an atom and $\mathscr{H}_{pos}, \mathscr{H}_{neg}$ be two sets of atoms such that $\mathscr{V}ar(\{A\} \cup \mathscr{H}_{pos} \cup \mathscr{H}_{neg}) \cap \mathscr{U} = \{\}$ and $A \approx B$ for all $B \in \mathscr{H}_{pos} \cup \mathscr{H}_{neg}$, and a set of variables G. The algorithm in Definition 7 always terminates and, if it returns a substitution σ, then $\bigwedge_{H \in \mathscr{H}_{pos}} A\sigma \approx H \wedge \bigwedge_{H' \in \mathscr{H}_{neg}} \neg(A\sigma \approx H')$ holds and $G\sigma$ is ground.

Example 4
Let $A := p(X)$, $\mathscr{H}_{pos} := \{p(s(Y))\}$, $\mathscr{H}_{neg} := \{p(s(0))\}$, and $G := \{X\}$. The algorithm of Definition 6 returns $\theta = \{X/s(Y)\}$. We take $\eta = \{Y/s(0)\}$, it is idempotent and $G\theta\eta$ is ground. $Dom(\eta) \subseteq \mathscr{V}ar(A\theta)$ and $\mathscr{V}ar(\eta) = \{Y\}$ does not intersect with $\mathscr{U}$. Finally, $A\theta\eta = p(s(s(0)))$ does not unify with $p(s(0))$. The algorithm thus returns $\theta\eta = \{X/s(s(0)), Y/s(0)\}$ restricted to the variables of A, i.e., $\{X/s(s(0))\}$.

Example 5
Let $A := p(X)$, $\mathscr{H}_{pos} := \{p(a), p(b)\}$, $\mathscr{H}_{neg} := \{p(f(Z))\}$, and $G := \{\}$. The algorithm of Definition 6 applied to A and $\mathscr{H}_{pos}$ returns $\theta = \{X/U\}$. However, we cannot find η such that $A\theta\eta$ does not unify with $p(f(Z))$ without binding U. The algorithm thus returns fail.

Theorem 2 states the soundness of our procedure for computing function alt. As for completeness, we claim that binding an atom A with all possible maximal unifying substitutions for A and $\mathscr{H}_{pos}$ does not affect to the existence of a solution to the unification problem (*) above (see the online appendix (Appendix B) for more details).

4.3 A tool for concolic testing

In this section, we present a prototype implementation of the concolic testing scheme. The tool, called contest, is publicly available from the following URL

http://kaz.dsic.upv.es/contest.html

It consists of approx. 1000 lines of Prolog code and implements the concolic testing algorithm of Definition 5 with function alt as described in Section 4.2 and a maximum term depth that can be fixed by the user in order to guarantee the termination of the process. Moreover, we also introduced a bound for the number of alternatives when computing function alt_trace (see Definition 4). Roughly speaking, when the number of alternatives is too high, we give up aiming at full *choice* coverage and return sets with only one clause label (which suffice for clause coverage).

Table 1 shows a summary of the coverage achieved by the test cases automatically generated using contest. The complete benchmarks –including the source code, initial goal, input arguments and maximum term depth– can be found in the above URL. We used the coverage analysis tool of SICStus Prolog 4.3.1, which basically

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

Table 1. *Clause coverage analysis results (SICStus Prolog)*

paper	100%	*paper2*	100%	*nat*	100%	*advisor*	100%	*applast*	100%
depth	88%	*regexp*	86%	*relative*	100%	*rotateprune*	100%	*transpose*	100%
mult	100%	*hanoi*	100%	*automaton*	100%	*qsort*	95%	*inclist*	100%
doubleflip	100%	*recacctype*	100%	*ackermann*	100%	*fibonacci*	100%	*preorder*	100%

measures the number of times each clause is used. The results are very satisfactory, achieving a full coverage in most of the examples.

The current version is a proof-of-concept implementation and only deals with pure Prolog without negation. We plan to extend it to cope with full Prolog. The concrete semantics can be extended following (Ströder *et al.* 2011), and concolic execution is in general a natural extension of the semantics in Figure 2. For relational built-in's or equalities, we should label the execution step with an associated constraint, which can then be used to produce alternative execution paths by solving its negation. In this context, our tool will be useful not only for test case generation, but also to detect program errors during concolic testing (e.g., negated atoms which are not instantiated enough, incorrect calls to arithmetic built-in's, etc). See the online appendix (Appendix A) for more details on extending concolic execution to full Prolog.

5 Related work and concluding remarks

Mera *et al.* (2009) present a framework unifying unit testing and run-time verification for the Ciao system (Hermenegildo *et al.* 2012). The ECL^{i}PS^e constraint programming system (Schimpf and Shen 2012) and SICStus Prolog (Carlsson and Mildner 2012) both provide tools which run a given goal and compute how often program points in the code were executed. SWI-Prolog (Wielemaker *et al.* 2012) offers a unit testing tool associated to an optional interactive generation of test cases. It also includes an experimental coverage analysis which runs a given goal and computes the percentage of the used clauses and failing clauses. Belli and Jack (1993) and Degrave *et al.* (2008) consider automatic generation of test inputs for strongly typed and moded logic programming languages like the Mercury programming language (Somogyi *et al.* 1996), whereas we only require moding the top-level predicate of the program.

One of the closest approaches to our work is the test case generation technique by (Albert *et al.* 2014). The main difference, though, is that their technique is based solely on traditional symbolic execution. As mentioned before, concolic testing may scale better since one can deal with more complex constraints by using data from the concrete component of the concolic state. Another difference is that we aim at full path coverage (i.e., choice coverage), and not only a form of statement coverage.

Another close approach is (Vidal 2015), where a concolic execution semantics for logic programs is presented. However, this approach only considers a simpler statement coverage and, thus, it can be seen as a particular instance of the technique in the present paper. Another significant difference is that, in (Vidal 2015), concolic

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

execution proceeds in a stepwise manner: first, concrete execution produces an execution *trace*, which is then used to drive concolic execution. Although this scheme is conceptually simpler, it may give rise to poorer results in practice since one cannot use concrete values in symbolic executions, one of the main advantages of concolic execution over traditional symbolic execution. Moreover, Vidal (2015) presents no formal results nor an implementation of the concolic execution technique.

Summarizing the paper, we have introduced a novel scheme for concolic testing in logic programming. It offers a sound and fully automatic technique for test case generation with a good code coverage. We have stated the particular type of unification problems that should be solved to produce new test cases. We have proposed a correct algorithm for such unification problems. Furthermore, we have developed a publicly available proof-of-concept implementation of the concolic testing scheme, contest, that shows the usefulness of our approach. To the best of our knowledge, this is the first fully automatic testing tool for Prolog that aims at full path coverage (here called choice coverage).

As future work, we plan to extend the scheme to full Prolog (see the remarks in Section 4.3). Another interesting subject for further research is the definition of appropriate heuristics to drive concolic testing w.r.t. a given coverage criterion. This might have a significant impact on the quality of the test cases when the process is incomplete. Finally, from the experimental evaluation, we observed that the results could be improved by introducing type information, so that the generated values are restricted to the right type. Hence, improving concolic testing with type annotations is also a promising line of future work.

References

Albert, E., Arenas, P., Gómez-Zamalloa, M., and Rojas, J. 2014. Test Case Generation by Symbolic Execution: Basic Concepts, a CLP-Based Instance, and Actor-Based Concurrency. In *SFM 2014*. Springer LNCS 8483, 263–309.

Apt, K. 1997. *From Logic Programming to Prolog*. Prentice Hall.

Belli, F. and Jack, O. 1993. Implementation-based analysis and testing of Prolog programs. In *ISSTA*. 70–80.

Carlsson, M. and Mildner, P. 2012. SICStus Prolog - the first 25 years. *Theory and Practice of Logic Programming 12,* 1-2, 35–66.

Claessen, K. and Hughes, J. 2000. QuickCheck: a lightweight tool for random testing of Haskell programs. In *Proc. of (ICFP 2000)*. ACM, 268–279.

Clarke, L. 1976. A program testing system. In *Proceedings of the 1976 Annual Conference (ACM'76)*. 488–491.

Degrave, F., Schrijvers, T., and Vanhoof, W. 2008. Automatic generation of test inputs for Mercury. In *Logic-Based Program Synthesis and Transformation, 18th International Symposium, LOPSTR 2008*. 71–86.

Godefroid, P., Klarlund, N., and Sen, K. 2005. DART: directed automated random testing. In *Proc. of PLDI'05*. ACM, 213–223.

Godefroid, P., Levin, M., and Molnar, D. 2012. Sage: whitebox fuzzing for security testing. *CACM 55,* 3, 40–44.

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press

Gómez-Zamalloa, M., Albert, E., and Puebla, G. 2010. Test case generation for object-oriented imperative languages in CLP. *TPLP 10,* 4-6, 659–674.

Hermenegildo, M. V., Bueno, F., Carro, M., López-García, P., Mera, E., Morales, J. F., and Puebla, G. 2012. An overview of Ciao and its design philosophy. *TPLP 12,* 1-2, 219–252.

King, J. C. 1976. Symbolic execution and program testing. *CACM 19,* 7, 385–394.

Lloyd, J. 1987. *Foundations of Logic Programming*. Springer-Verlag, Berlin. 2nd Ed.

Martelli, A. and Montanari, U. 1982. An Efficient Unification Algorithm. *ACM Transactions on Programming Languages and Systems 4*, 258–282.

Mera, E., López-García, P., and Hermenegildo, M. V. 2009. Integrating software testing and run-time checking in an assertion verification framework. In *25th International Conference on Logic Programming, ICLP 2009, Pasadena*. 281–295.

Pasareanu, C. and Rungta, N. 2010. Symbolic PathFinder: symbolic execution of Java bytecode. In *ASE*, C. Pecheur, J. Andrews, and E. D. Nitto, Eds. ACM, 179–180.

Plotkin, G. 1970. A note on inductive generalization. *Machine intelligence 5*, 153–163.

Schimpf, J. and Shen, K. 2012. ECL^{i}PS^e - from LP to CLP. *Theory and Practice of Logic Programming 12,* 1-2, 127–156.

Sen, K., Marinov, D., and Agha, G. 2005. CUTE: a concolic unit testing engine for C. In *Proc. of ESEC/SIGSOFT FSE 2005*. ACM, 263–272.

Somogyi, Z., Henderson, F., and Conway, T. 1996. The execution algorithm of Mercury, an efficient purely declarative Logic Programming language. *The Journal of Logic Programming 29,* 1–3, 17–64.

Ströder, T., Emmes, F., Schneider-Kamp, P., Giesl, J., and Fuhs, C. 2011. A Linear Operational Semantics for Termination and Complexity Analysis of ISO Prolog. In *LOPSTR'11*. Springer LNCS 7225, 237–252.

Vasak, T. and Potter, J. 1986. Characterization of terminating logic programs. In *Proc. of the 1986 Intl. Symp. on Logic Programming*. IEEE, 140–147.

Vidal, G. 2015. Concolic Execution and Test Case Generation in Prolog. In *Proc. of the 24th International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR'14)*, M. Proietti and H. Seki, Eds. Springer LNCS 8981, 167–181.

Wielemaker, J., Schrijvers, T., Triska, M., and Lager, T. 2012. SWI-Prolog. *Theory and Practice of Logic Programming 12,* 1-2, 67–96.

https://doi.org/10.1017/S1471068415000332 Published online by Cambridge University Press