

Document downloaded from:

<http://hdl.handle.net/10251/66304>

This paper must be cited as:

Vidal Oriola, GF. (2015). Towards Symbolic Execution in Erlang. En Perspectives of System Informatics. Springer. 351-360. doi:10.1007/978-3-662-46823-4_28.



The final publication is available at

http://link.springer.com/chapter/10.1007%2F978-3-662-46823-4_28

Copyright Springer

Additional Information

The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-662-46823-4_28

Towards Symbolic Execution in Erlang^{*}

Germán Vidal

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

1 Introduction

The concurrent functional language Erlang [1] has a number of distinguishing features, like dynamic typing, concurrency via asynchronous message passing or hot code loading, that make it especially appropriate for distributed, fault-tolerant, soft real-time applications. The success of Erlang is witnessed by the increasing number of its industrial applications. For instance, Erlang has been used to implement Facebook’s chat back-end, the mobile application Whatsapp or Twitterfall—a service to view trends and patterns from Twitter—, to name a few. The success of the language, however, also requires the development of powerful testing and verification techniques.

Symbolic execution is at the core of many program analysis and transformation techniques, like partial evaluation, test-case generation or model checking. In this paper, we introduce a symbolic execution technique for Erlang. We discuss how both an overapproximation and an underapproximation of the concrete semantics can be obtained. We illustrate our approach through some examples. To the best of our knowledge, this is the first attempt to formalize symbolic execution in the context of this language, where previous approaches have only considered exploring different schedulings but have not dealt with symbolic data. More details can be found in the companion technical report [17].

2 Erlang Syntax

In this section, we present the basic syntax of a significant subset of Erlang. In particular, we consider a slightly simplified version of the language where some features are excluded (basically, we do not consider modules, exceptions, records, binaries, monitors, ports or process links, most of which are not difficult to deal with but would encumber the notations and definitions of this paper). Nevertheless, this is still a large subset of Erlang and covers its main distinguishing features, like pattern matching, higher-order functions, process creation, message sending and receiving, etc.

^{*} This work has been partially supported by 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.

$$\begin{aligned}
\text{Program } \ni \text{pgm} &::= f(\overline{X}_n) \rightarrow e. \mid \text{pgm pgm} \\
\text{Exp } \ni e &::= bv \mid [e_1|e_2] \mid \{\overline{e}_n\} \mid X \mid e(\overline{e}_n) & (n \geq 0) \\
&\mid \text{case } e \text{ of } cl \text{ end} \mid e_1 ! e_2 \mid \text{receive } cl \text{ end} \\
&\mid p = e \mid e_1, e_2 \\
\text{BasicValue } \ni bv &::= a \mid n \mid p \mid [] \mid \{ \} \\
\text{Value } \ni v &::= bv \mid [v_1|v_2] \mid \{\overline{v}_n\} & (n > 0) \\
\text{Pattern } \ni p &::= bv \mid X \mid [p_1|p_2] \mid \{\overline{p}_n\} & (n > 0) \\
\text{Clauses } \ni cl &::= p_1 \text{ when } g_1 \rightarrow e_1; \dots; p_n \text{ when } g_n \rightarrow e_n & (n > 0)
\end{aligned}$$

where $a \in \text{Atom}$, $n \in \text{Number}$, $p \in \text{Pid}$, $X \in \text{Var}$, $g \in \text{Guard}$

Fig. 1. Erlang syntax rules

The syntax of the language can be found in Figure 1. We denote by \overline{o}_n the sequence of syntactic objects o_1, \dots, o_n . Programs are sequences of function definitions, where each function f/n is defined by a rule $f(X_1, \dots, X_n) \rightarrow e$. with X_1, \dots, X_n distinct variables and the body of the function, e , an expression that might include basic values, lists, tuples, variables, function applications, case expressions, message sending and receiving, pattern matching and sequences.

Besides the functions defined in the program, we consider some of the usual *built-in* functions (logical and relational operators, arithmetic operators, etc.), together with the functions `self`, that returns the pid of the current process, and `spawn`, that is used to create new processes. E.g., `spawn(foo, [a, 42])` creates a new process that starts calling the function `foo(a, 42)` and returns the new (fresh) pid assigned to this process. Only the concurrent actions have side effects. We assume that *guards* can only contain calls to built-in functions without side effects.

Example 1. Consider the program in Fig. 2 which follows a very simple client-server scheme. Here, the first process is called with `start(N)`, where N is the maximum number of requests accepted by the server. Then, it creates a client (a new concurrent process) and starts the server. A client request just includes its own pid and the request number. If the request number is smaller than N , the server answers “ok”; otherwise, it answers “last” and terminates. The client keeps asking the server with increasing numbers until it gets the reply “last”.

We do not consider I/O in this paper. Therefore, input parameters must be provided through the initial function.

3 Concrete Semantics

The semantics of Erlang is informally described, e.g., in [1]. The past years have witnessed an increasing number of works aimed at defining a formal semantics for the language. Some of the first attempts were done by Huch [9] and, more extensively, by Fredlund [6]. More recent approaches focus on the definition of the distributed aspects of the Erlang semantics, like [4]; this semantics was later refined in [15] and [14], where some assumptions on the future of the language

```

start(N) → S = self(), C = spawn(client, [1, S], server(N)).
server(N) → receive
    {Pid, M} when M < N → Pid!ok, server(N);
    {Pid, M} when M >= N → Pid!last
end.
client(N, Pid) → Pid! {self(), N},
    receive Atom → case Atom of
        ok → client(N + 1, Pid);
        last → ok
    end
end.

```

Fig. 2. Simple client-server example in Erlang

design are proposed. Other approaches have formalized the semantics of Erlang by defining its semantics in the framework of rewriting logic [10, 11].

Unfortunately, there is no commonly accepted semantics and, moreover, most of the above papers only cover part of the language semantics (e.g., [4, 15, 14] are mainly oriented towards the concurrent features of the language). Therefore, we have recently introduced a semantics for a subset of Erlang in [16]. In the following, we present a more elegant and general version of this semantics that follows some of the ideas in [14].

Erlang follows a leftmost innermost operational semantics. Following, e.g., [9, 6], every expression $C[e]$ can be decomposed into a context $C[\]$ with a (single) hole and a subexpression e where the next reduction can take place:

$$\begin{aligned}
C ::= & \square \mid C, e \mid \text{case } C \text{ of } cl \text{ end} \mid C!e \mid v!C \mid p = C \mid C(e_1, \dots, e_n) \\
& \mid f(v_1, \dots, v_i, C, e_{i+2}, \dots, e_n) \mid op(v_1, \dots, v_i, C, e_{i+2}, \dots, e_n) \\
& \mid [v_1, \dots, v_i, C|e] \mid \{v_1, \dots, v_i, C, e_{i+2}, \dots, e_n\}
\end{aligned}$$

An Erlang *process* is denoted by a tuple $\langle p; e; q \rangle$, where p is a the process identifier, e is the expression to be evaluated, and q is the process mailbox. An Erlang *system* is a pair (Π, \mathcal{Q}) , where Π is a pool of processes and \mathcal{Q} is the system mailbox (analogous to the *ether* in the semantics of [14]). We assume no order in Π since it is not relevant to our purposes (i.e., we will be interested in exploring *all* possible schedulings within symbolic execution). For implementing actual scheduling policies, an ordering would be required. The system mailbox \mathcal{Q} is a set of triples (p, p', q) , where q is a list of messages (values) sent from the process with pid p to the process with pid p' . The system mailbox is needed to correctly model a multi-node distributed system (see the discussion in [14]). Basically, Erlang only requires that the messages sent *directly* between two processes must arrive in the same order. However, if the messages follow different paths, say one message is sent directly from p to p'' , while another message is sent from p to p'' via p' , then there is no guarantee regarding which message arrives first to p'' .

The operational semantics is defined by the labelled transition relation \rightarrow shown in Fig. 3. Here, we use the notation $\langle p; e; q \rangle \& \Pi$ to denote an arbitrary

(seq)	$\frac{}{\langle\langle p; C[v, e]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; C[e]; q \rangle \& \Pi, \mathcal{Q} \rangle}$
(self)	$\frac{}{\langle\langle p; C[\text{self}()]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; C[p]; q \rangle \& \Pi, \mathcal{Q} \rangle}$
(builtin)	$\frac{\text{eval}(\text{op}(\overline{v_n})) = v}{\langle\langle p; C[\text{op}(\overline{v_n})]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; C[v]; q \rangle \& \Pi, \mathcal{Q} \rangle}$
(fun)	$\frac{f(\overline{X_n}) \rightarrow e. \in \text{pgm}}{\langle\langle p; C[f(\overline{v_n})]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; C[\widehat{e}\{\overline{X_n} \mapsto \overline{v_n}\}]; q \rangle \& \Pi, \mathcal{Q} \rangle}$
(match)	$\frac{\exists \sigma. p\sigma = v}{\langle\langle p; C[p = v]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; (C[v])\sigma; q \rangle \& \Pi, \mathcal{Q} \rangle}$
(case)	$\frac{\text{match}(v, cl) = (e, \sigma)}{\langle\langle p; C[\text{case } v \text{ of } cl \text{ end}]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; (C[e])\sigma; q \rangle \& \Pi, \mathcal{Q} \rangle}$
(receive)	$\frac{\text{matchrec}(q, cl) = (e, \sigma, q')}{\langle\langle p; C[\text{receive } cl \text{ end}]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; (C[e])\sigma; q' \rangle \& \Pi, \mathcal{Q} \rangle}$
(spawn)	$\frac{p' \text{ is a fresh pid}}{\langle\langle p; C[\text{spawn}(f, \overline{v_n})]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; C[p']; q \rangle \& \langle p', f(\overline{v_n}), [] \rangle \& \Pi, \mathcal{Q} \rangle}$
(send)	$\frac{v_1 = p' \in \text{Pid} \wedge \text{add_msg}(p, p', v_2, \mathcal{Q}) = \mathcal{Q}'}{\langle\langle p; C[v_1 ! v_2]; q \rangle \& \Pi, \mathcal{Q} \rangle \xrightarrow{\tau} \langle\langle p; C[v_2]; q \rangle \& \Pi, \mathcal{Q}' \rangle}$
(sched)	$\frac{(p, p') \in \text{sched}(\Pi, \mathcal{Q}) \wedge \text{delivery}(p, p', \Pi, \mathcal{Q}) = (\Pi', \mathcal{Q}')}{(\Pi, \mathcal{Q}) \xrightarrow{\tau} (\Pi', \mathcal{Q})}$

Fig. 3. Concrete Semantics

pool of processes that contains the process $\langle p; e; q \rangle$. The initial system has the form $(\langle p_0; e; [] \rangle, [])$. Most rules are self-explanatory. Let us just explain the more involved ones:

In rule **builtin**, we assume a function *eval* that evaluates all built-in's without side effects (i.e., arithmetic or relational expressions, etc).

In rule **fun**, we assume that the program *pgm* is a global parameter of the transition system. Moreover, we let \widehat{e} denote a copy of *e* with local variables renamed with fresh names. The notation $\{\overline{X_n} \mapsto \overline{v_n}\}$ denotes a substitution binding variables X_1, \dots, X_n to values v_1, \dots, v_n . The application of a substitution σ to an expression *e* is denoted by $e\sigma$.

In rule **case**, we assume an auxiliary function *match* that takes a value *v* and the clauses $p_1 \text{ when } g_1 \rightarrow e_1; \dots; p_n \text{ when } g_n \rightarrow e_n$ and returns a pair (e_i, σ) if *i* is the smaller number such that $p_i\sigma = v$ and $\text{eval}(g_i\sigma) = \text{true}$.

The case of rule **receive** uses a similar auxiliary function *matchrec* that takes a mailbox queue *q* and the clauses *cl*, determines the first message *v* such that $\text{match}(v, cl) = (e, \sigma)$, and returns (e, σ, q') , where q' is obtained from *q* by deleting message *v*.

In rule **send**, the message is stored in the system mailbox, together with the source and target pids, using the auxiliary function *add_msg*, whose definition is straightforward. Note that the message is not actually delivered to the process with pid p' until the **sched** rule is applied (see below).

Finally, rule **sched** uses the auxiliary function *sched* to model a particular scheduling policy. Basically, it selects two pids (p, p') from *II* (source and target

$$\begin{aligned}
& \langle \langle p_0; \text{start}(1); [] \rangle, [] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; S = \underline{\text{self}()} \rangle, C = \text{spawn}(\text{client}, [1, S]), \text{server}(1); [] \rangle, [] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; S = p_0 \rangle, C = \text{spawn}(\text{client}, [1, S]), \text{server}(1); [] \rangle, [] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; p_0, C = \text{spawn}(\text{client}, [1, p_0]) \rangle, \text{server}(1); [] \rangle, [] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; C = \text{spawn}(\text{client}, [1, p_0]) \rangle, \text{server}(1); [] \rangle, [] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; C = p_1, \text{server}(1); [] \rangle \& \langle p_1; \text{client}(1, p_0); [] \rangle, [] \rangle \\
& \dots \\
& \xrightarrow{\tau} \langle \langle p_0; \text{receive} \dots \text{end}; [] \rangle \& \langle p_1; \text{receive} \dots \text{end}; [] \rangle, [(p_1, p_0, \{\{p_1, 1\}\})] \rangle \\
& \xrightarrow{\alpha} \langle \langle p_0; \underline{\text{receive} \dots \text{end}}; \{\{p_1, 1\}\} \rangle \& \langle p_1; \text{receive} \dots \text{end}; [] \rangle, [(p_1, p_0, [])] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; p_1 ! \text{last}; [] \rangle \& \langle p_1; \text{receive} \dots \text{end}; [] \rangle, [(p_1, p_0, [])] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; \text{last}; [] \rangle \& \langle p_1; \text{receive} \dots \text{end}; [] \rangle, [(p_1, p_0, []), (p_0, p_1, [\text{last}])] \rangle \\
& \xrightarrow{\alpha} \langle \langle p_0; \text{last}; [] \rangle \& \langle p_1; \underline{\text{receive} \dots \text{end}}; [\text{last}] \rangle, [(p_1, p_0, []), (p_0, p_1, [])] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; \text{last}; [] \rangle \& \langle p_1; \text{case last of} \dots \text{end}; [] \rangle, [(p_1, p_0, []), (p_0, p_1, [])] \rangle \\
& \xrightarrow{\tau} \langle \langle p_0; \text{last}; [] \rangle \& \langle p_1; \text{ok}; [] \rangle, [(p_1, p_0, []), (p_0, p_1, [])] \rangle
\end{aligned}$$

Fig. 4. Computation for the program of Example 1

processes, which might be the same) such that $(p, p', q) \in \mathcal{Q}$ and q is not empty. Then, function *delivery* moves the first message of q to the local mailbox of the process with pid p' , thus returning a new pair $(\mathcal{I}', \mathcal{Q}')$.

Observe that all rules are labeled with τ except for the last one. This is explained by the fact that we are interested in a particular type of computations that we call *normalized* computations. In the following, given a state s , we denote by $s \downarrow^\tau$ the state that results from s by only applying transitions labeled with τ until no more transitions labeled with τ are possible, i.e., if $s \equiv s_0 \xrightarrow{\tau} s_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} s_n \not\xrightarrow{\tau}$, then $s \downarrow^\tau = s_n$.

Definition 1 (normalized computation). *Let s_0 be the initial system. Then, we say that a computation is normalized if it has the form*

$$s_0 \xrightarrow{\tau^*} s_0 \downarrow^\tau \xrightarrow{\alpha} s_1 \xrightarrow{\tau^*} s_1 \downarrow^\tau \xrightarrow{\alpha} s_2 \xrightarrow{\tau^*} s_2 \downarrow^\tau \xrightarrow{\alpha} s_3 \dots$$

In the following, we only consider normalized computations in order to reduce the search space.

Example 2. Consider again the program of Ex. 1. A computation with this program is shown in Fig. 4, where the expression selected for reduction is underlined.

4 Symbolic Execution Semantics

In this section, we introduce a symbolic execution semantics for Erlang. Firstly, one could consider the semantics in Fig. 3 and just define a function *sched* that returns all feasible combinations of processes in the considered system. This is useful to explore all possible schedulings and detect errors (e.g., deadlocks) that only occur in a particular scheduling. This is the aim, e.g., of the model checker McErlang [7]. Basically, McErlang is today a mature tool that combines the use of random test cases (using, e.g., a tool like QuickCheck [2]) with a semantics that explores possible schedulings.

(seq)	$\frac{}{\langle p; C[p, e]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; C[e]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C}}$
(self)	$\frac{}{\langle p; C[\text{self}()]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; C[p]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C}}$
(builtin1)	$\frac{\text{eval}(\text{op}(\overline{v_n})) = v}{\langle p; C[\text{op}(\overline{v_n})]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; C[v]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C}}$
(builtin2)	$\frac{\exists i. p_i \text{ is not a value, } X \text{ is a fresh variable}}{\langle p; C[\text{op}(\overline{p_n})]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; C[X]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \wedge (X = \text{op}(\overline{p_n}))}$
(fun)	$\frac{f(\overline{X_n}) \rightarrow e. \in \text{pgm}}{\langle p; C[f(\overline{p_n})]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; C[e\{X_n \mapsto p_n\}]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C}}$
(match)	$\frac{\exists \sigma. p_1 \sigma = p_2 \sigma}{\langle p; C[p_1 = p_2]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; (C[p_2])\sigma; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C}}$
(case)	$\frac{(e, \sigma, \mathcal{C}') \in \text{unify}(C, p, cl), \mathcal{C}'' = \widehat{\sigma} \wedge \mathcal{C}'}{\langle p; C[\text{case } p \text{ of } cl \text{ end}]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; (C[e])\sigma; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \wedge \mathcal{C}''}$
(receive)	$\frac{(e, \sigma, q', \mathcal{C}') \in \text{unifyrec}(C, q, cl), \mathcal{C}'' = \widehat{\sigma} \wedge \mathcal{C}'}{\langle p; C[\text{receive } cl \text{ end}]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; (C[e])\sigma; q' \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \wedge \mathcal{C}''}$
(spawn)	$\frac{p' \text{ is a fresh pid}}{\langle p; C[\text{spawn}(f, \overline{p_n})]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; C[p']; q \rangle \& \langle p', f(\overline{p_n}), [] \rangle \& \Pi, \mathcal{Q}, \mathcal{C}}$
(send)	$\frac{v = p' \in \text{Pid} \wedge \text{add_msg}(p, p', p, \mathcal{Q}) = \mathcal{Q}'}{\langle p; C[v ! p]; q \rangle \& \Pi, \mathcal{Q}, \mathcal{C} \xrightarrow{\tau} \langle p; C[p]; q \rangle \& \Pi, \mathcal{Q}', \mathcal{C}}$
(sched)	$\frac{(p, p') \in \text{sched}(\Pi, \mathcal{Q}) \wedge \text{delivery}(p, p', \Pi, \mathcal{Q}) = (\Pi', \mathcal{Q}')}{(\Pi, \mathcal{Q}, \mathcal{C}) \xrightarrow{\alpha} (\Pi', \mathcal{Q}', \mathcal{C})}$

Fig. 5. Symbolic Execution

Here, we plan to also cope with missing input data (analogously to the tool Java Pathfinder [12] for model checking of Java bytecode). Our *symbolic systems* are now triples of the form $(\Pi, \mathcal{Q}, \mathcal{C})$, where the new element \mathcal{C} is the so called *path constraint* (initialized to *true*). Loosely speaking, \mathcal{C} contains some constraints on the symbolic values that represent the missing input data, such that the system (Π, \mathcal{Q}) is reachable (using the concrete semantics) when the input data in the initial system satisfies the constraint \mathcal{C} .

An Overapproximation. First, we consider that symbolic execution must *overapproximate* the concrete semantics. This is useful, e.g., in the context of partial evaluation or when a property that holds for *all* states must be verified. The symbolic execution semantics is shown in Fig. 5. Let us briefly explain the main differences w.r.t. the concrete semantics:

Rule **builtin** considers now two cases: **builtin1**, which is equivalent to the previous rule in the concrete semantics, and **builtin2** that considers the case when some argument is not a value. In the latter case, the built in function cannot be evaluated and we reduce it to a fresh variable and add the corresponding constraint to the system. E.g., given the expression $3 + Y$, we reduce it to a fresh variable X and add the constraint $X = 3 + Y$ to the system constraint.

Rule **fun** remains unchanged. Applications of the form $X(p_1, \dots, p_n)$ are not considered since it would involve calling every function and built-in of the pro-

$$\begin{array}{l}
\langle (p_0; \text{start}(K); []), [], \text{true} \rangle \\
\xrightarrow{\tau} \langle (p_0; S = \underline{\text{self}}(), C = \text{spawn}(\text{client}, [1, S]), \text{server}(K); []), [], \text{true} \rangle \\
\cdots \\
\overset{\alpha}{\rightarrow} \langle (p_0; \text{receive} \dots \text{end}; [\{p_1, 1\}]) \& \langle p_1; \text{receive} \dots \text{end}; [] \rangle, [(p_1, p_0, [])], \text{true} \rangle \\
\overset{\tau}{\rightarrow} \langle (p_0; p_1 ! \text{ok}, \text{server}(K); []) \& \langle p_1; \text{receive} \dots \text{end}; [] \rangle, [(p_1, p_0, [])], 1 < K \rangle \\
\overset{\tau}{\rightarrow} \langle (p_0; \text{ok}, \text{server}(K); []) \& \langle p_1; \text{receive} \dots \text{end}; [] \rangle, [(p_1, p_0, []), (p_0, p_1, [\text{ok}])], 1 < K \rangle \\
\overset{\tau}{\rightarrow} \langle (p_0; \text{receive} \dots \text{end}; []) \& \langle p_1; \text{receive} \dots \text{end}; [] \rangle, [(p_1, p_0, []), (p_0, p_1, [\text{ok}])], 1 < K \rangle \\
\cdots
\end{array}$$

Fig. 6. Partial symbolic execution for the program of Example 1

gram to keep the symbolic execution complete, which is not acceptable. If such an expression is reached, we give up and stop symbolic execution with a failure.

Rule `match` is similar to the original rule in the concrete semantics but replaces matching with unification. Analogously, rules `case` and `receive` mainly replaces the auxiliary functions `match` and `matchrev` with `unify` and `unifyrev` where unification replaces matching as follows. Function `unify` takes a constraint \mathcal{C} , a pattern p and the clauses $p_1 \text{ when } g_1 \rightarrow e_1; \dots; p_n \text{ when } g_n \rightarrow e_n$ and returns a triple $(e_i, \sigma, \mathcal{C}')$ for each i such that $p_i\sigma = p\sigma$ (i.e., σ is a *unifier* of p_i and p) and $\mathcal{C} \Rightarrow \neg g_i\sigma$ cannot be proved (i.e., the unsatisfiability of $g_i\sigma$ cannot be proved); here, \mathcal{C}' is the constraint $\mathcal{C} \wedge g_i\sigma$ (when $g_i\sigma$ is different from *true*). Function `unifyrec` proceeds analogously. Note that we also add the computed unifier to the path constraint (where $\hat{\sigma}$ denotes the equational representation of a substitution σ). This will be required in the next section. The new functions return a *set* since the pattern might unify with more than one clause whose guard is also satisfiable. Note that this strategy is complete but typically not sound since (besides the limitations of the constraint solver) we might follow several paths while the original, concrete semantics only considers the first clause even if a value matches several clauses.

Rule `spawn`, analogously to the case of rule `fun`, does not consider an expression like `spawn(X, [\overline{p_n}])`, which will be considered a failure. A similar situation happens with rule `send`. Here, we consider the case where the message is a pattern and, thus, might be a variable. However, we do not consider that the pid of the target process is a variable, since it would involve broadcasting the message to all processes to keep the symbolic execution complete, which is not acceptable.

Finally, rule `sched` just considers a scheduling function `sched` that returns all possible combinations in order to explore all feasible schedulings.

We assume that the system constraint is checked for *unsatisfiability* at every step. When unsatisfiability cannot be proved we continue with the symbolic execution (which is complete, but a potential source of unsoundness).

As in the previous case, only *normalized* symbolic executions are considered.

Example 3. Consider again the program of Ex. 1. Now, Fig. 6 shows a normalized symbolic execution starting with an unknown number K of maximum requests.

An Underapproximation. So far, we have put the emphasis on completeness (i.e., producing an overapproximation of the original Erlang computations). For this purpose, we had to take a number of decisions that make the resulting search

space too huge to scale to real world Erlang applications with thousands or millions of processes. Moreover, there are a number of situations in which we have to give up (i.e., variable applications, process spawning with an unknown function or sending a message to an unknown pid) because dealing with them is simply intractable.

As an alternative, we propose in this section a *sound* symbolic execution that computes an *underapproximation* of the concrete semantics. This is useful for many applications (like test case generation or model checking), and it is often more scalable and avoids false positives. Here, we follow the approach of [8, 13] to so called *concolic execution* and consider the following scheme:

Processes are slightly extended as follows: $\langle p, e_c, e_s, q \rangle$, where p is a pid, e_c is a concrete expression, e_s is a symbolic expression, and q is the mailbox queue. The symbolic expression is only used to compute the corresponding path constraint.

Now, one starts the execution with a random test input data and execute the program using basically the symbolic execution semantics of Fig. 5 using an initial system like $\langle p_0, start(1), start(K), [], true \rangle$.

Then, when the computation terminates, we produce a sequence of the form $E_0, E_1, E_2, \dots, E_n$ where each E_i is either a constraint C_i (associated to the i -th computation step) or the symbol α denoting one application of the sched rule. We now traverse this sequence starting from the last element and either negate a constraint or consider alternative schedulings, depending on the type of the considered element. In the case of a negated constraint, we use a constraint solver to produce a new set of input data. Either way, a new concolic execution is considered and the process starts again. Usually, backtracking can be used to explore all possibilities.

If the algorithm terminates and the constraint solver is always able to generate a new set of input data, concolic execution is both sound and complete; otherwise, it is only sound (an underapproximation). Termination can be ensured using, e.g., a maximum depth for symbolic execution.

Example 4. Consider again the program of Example 1 and the initial call $start(1)$. The initial system is thus $(\langle p_0, start(1), start(K), [] \rangle, [], true)$. Here, we would basically perform the same computation shown in Example 2 but using the rules of Fig. 5 to also obtain the following sequence of constraints and scheduling steps: $(\alpha, 1 \geq K)$ (only the constraints relevant to the symbolic input data, K , have to be considered). Now, by negating the constraint $1 \geq K$, we produce a new value, e.g., $K = 5$, and consider a new symbolic execution starting from the system $(\langle p_0, start(5), start(K), [] \rangle, [], true)$. Finally, one should consider alternative schedulings (because we reach a symbol α) but no alternative exists. Therefore, we conclude that executing $start(1)$ and $start(5)$ is sufficient to cover all possible execution paths for the source program.

5 Discussion

In this paper, we have introduced a high-level concrete semantics for the functional and concurrent language Erlang, and have explored the definition of an as-

sociated symbolic execution technique. We proposed both an overapproximation and an underapproximation—based on a variant of symbolic execution called concolic execution—. In principle, it seems that the underapproximation will be more practical and scalable in order to design a tool for model checking and/or test case generation. We are only aware of the approach of [3] to symbolic execution in Erlang, though no formalization is introduced in this paper (it is only explained informally). Hence we think that our approach is a promising step towards defining a practical symbolic execution technique for Erlang, which can be used in different contexts like model checking or test case generation.

References

1. J. Armstrong, R. Virding, and M. Williams. *Concurrent programming in Erlang (2nd edition)*. Prentice Hall, 1996.
2. T. Arts, J. Hughes, J. Johansson, and U.T. Wiger. Testing telecoms software with quviq QuickCheck. In *Proc. of the Erlang Workshop*, pages 2–10. ACM, 2006.
3. C. Benac Earle. Symbolic program execution using the Erlang verification tool. In *Proc. of the 9th Int'l Workshop on Functional and Logic Programming (WFLP 2000)*, pages 42–55, 2000.
4. K. Claessen and H. Svensson. A semantics for distributed Erlang. In *Proc. of the 2005 ACM SIGPLAN Workshop on Erlang*, pages 78–87. ACM, 2005.
5. M. Felleisen, D.P. Friedman, E.E. Kohlbecker, and B.F. Duba. A syntactic theory of sequential control. *Theor. Comput. Sci.*, 52:205–237, 1987.
6. L.-A. Fredlund. *A framework for reasoning about Erlang code*. PhD thesis, The Royal Institute of Technology, Sweden, 2001.
7. L.-A. Fredlund and H. Svensson. McErlang: a model checker for a distributed functional programming language. In *Proc. of ICFP'07*, pp. 125–136. ACM, 2007.
8. P. Godefroid, N. Klarlund, and K. Sen. DART: directed automated random testing. In *Proc. of PLDI'05*, pages 213–223. ACM, 2005.
9. F. Huch. Verification of Erlang Programs using Abstract Interpretation and Model Checking. In *Proc. of ICFP '99*, pages 261–272. ACM, 1999.
10. T. Noll. A Rewriting Logic Implementation of Erlang. *Electr. Notes Theor. Comput. Sci.*, 44(2):206–224, 2001.
11. T. Noll. Equational Abstractions for Model Checking Erlang Programs. *Electr. Notes Theor. Comput. Sci.*, 118:145–162, 2005.
12. C.S. Pasareanu, W. Visser, D.H. Bushnell, J. Geldenhuys, P.C. Mehlitz, and Neha Rungta. Symbolic PathFinder: integrating symbolic execution with model checking for Java bytecode analysis. *Autom. Softw. Eng.*, 20(3):391–425, 2013.
13. 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.
14. H. Svensson, L.-A. Fredlund, and C. Benac Earle. A unified semantics for future Erlang. In *Proc. of the Erlang Workshop*, pages 23–32. ACM, 2010.
15. H. Svensson and L.-A. Fredlund. A more accurate semantics for distributed Erlang. In *Proc. of the Erlang Workshop*, pages 43–54. ACM, 2007.
16. G. Vidal. Towards Erlang verification by term rewriting. In *Proc. of LOPSTR'13*, pages 161–178. Technical Report TR-11-13, Universidad Complutense de Madrid, 2013. Available from <http://users.dsic.upv.es/~gvidal/>.
17. G. Vidal. Symbolic execution in Erlang. Technical report, DSIC, UPV, 2014. Available from <http://users.dsic.upv.es/~gvidal/>.