

Bounded Symbolic Execution for Runtime Error Detection of Erlang Programs *

Emanuele De Angelis, Fabio Fioravanti

DEC, University “G. d’Annunzio” of Chieti-Pescara
Viale Pindaro 42, 65127 Pescara, Italy

{emanuele.deangelis, fabio.fioravanti}@unich.it

Adrián Palacios[†]

MiST, DSIC, Universitat Politècnica de València
Camino de Vera, s/n, 46022 València, Spain

apalacios@dsic.upv.es

Alberto Pettorossi

University of Roma Tor Vergata
Via del Politecnico 1, 00133 Roma, Italy

pettorossi@info.uniroma2.it

Maurizio Proietti

CNR-IASI
Via dei Taurini 19, 00185 Roma, Italy

maurizio.proietti@iasi.cnr.it

Dynamically typed languages, like Erlang, allow developers to quickly write programs without explicitly providing any type information on expressions or function definitions. However, this feature makes those languages less reliable than statically typed languages, where many runtime errors can be detected at compile time. In this paper, we present a preliminary work on a tool that, by using the well-known techniques of metaprogramming and symbolic execution, can be used to perform bounded verification of Erlang programs. In particular, by using Constraint Logic Programming, we develop an interpreter that, given an Erlang program and a symbolic input for that program, returns answer constraints that represent sets of concrete data for which the Erlang program generates a runtime error.

1 Introduction

Erlang [16] is a functional, message passing, concurrent language with dynamic typing. Due to this type discipline, Erlang programmers are quite familiar with typing and pattern matching errors at runtime, which normally appear during the first executions of freshly written programs. Less often, these errors will be undetected for a long time, until the user inputs a particular value that causes the program to crash or, in the case of concurrent programs, determines a particular interleaving that causes an error to occur.

In order to mitigate these problems, many static analysis tools have been proposed. Here let us recall:

- Dialyzer [13], which is a popular tool included in Erlang/OTP for performing type inference based on success typings, and
- SOTER [7], which is a tool that performs verification of Erlang programs by using model checking and abstract interpretation.

However, those tools are not all fully automatic, and they can only be used to cover either the sequential or the concurrent part of an Erlang program, but not both.

*This work has been partially supported by the EU (FEDER) and the *Spanish Ministerio de Ciencia, Innovación y Universidades/AEI* under grant TIN2016-76843-C4-1-R, the *Generalitat Valenciana* under grant PROMETEO-II/2015/013 (Smart-Logic), and INdAM-GNCS (Italy). Emanuele De Angelis, Fabio Fioravanti, and Alberto Pettorossi are research associates at CNR-IASI, Rome, Italy.

[†]Partially supported by the EU (FEDER) and the Spanish *Ayudas para contratos predoctorales para la formación de doctores* and *Ayudas a la movilidad predoctoral para la realización de estancias breves en centros de I+D*, MINECO (SEIDI), under FPI grants BES-2014-069749 and EEBB-I-17-12101.

In this paper we present a preliminary work on a technique, based on Constraint Logic Programming (CLP) [10], for detecting runtime errors in Erlang programs. In our approach, sequential Erlang programs are first translated into Prolog facts and then run by using an interpreter written in CLP. Our CLP interpreter is able to run programs on symbolic input data, and it can perform verification of Erlang programs up to a fixed bound on the number of execution steps.

The Erlang language. In this work we consider only sequential programs written in a first-order subset of the Core Erlang language¹. The syntax of this subset can be found in Figure 1.

$$\begin{aligned}
\text{module} & ::= \text{module } atom = fun_1, \dots, fun_n \\
\text{fun} & ::= \text{fname} = \text{fundef} \\
\text{fname} & ::= atom / nat \\
\text{fundef} & ::= \text{fun } (X_1, \dots, X_n) \rightarrow expr \text{ end} \\
\text{lit} & ::= atom \mid int \mid [] \\
\text{expr} & ::= X \mid lit \mid \text{fname} \mid [expr_1 \mid expr_2] \mid \{expr_1, \dots, expr_n\} \\
& \quad \mid \text{let } X = expr_1 \text{ in } expr_2 \\
& \quad \mid \text{case } expr \text{ of } clause_1; \dots; clause_m \text{ end} \\
& \quad \mid \text{apply } \text{fname} (expr_1, \dots, expr_n) \\
& \quad \mid \text{call } atom : \text{fname} (expr_1, \dots, expr_n) \\
& \quad \mid \text{primop } atom (expr_1, \dots, expr_n) \\
& \quad \mid \text{try } expr_1 \text{ of } X_1 \rightarrow expr_2 \text{ catch } X_2 \rightarrow expr_3 \\
\text{clause} & ::= pat \text{ when } expr_1 \rightarrow expr_2 \\
\text{pat} & ::= X \mid lit \mid [pat_1 \mid pat_2] \mid \{pat_1, \dots, pat_n\}
\end{aligned}$$

Figure 1: Language syntax rules

Here, a module is a sequence of function declarations, where each function name has an associated definition of the form “fun $(X_1, \dots, X_n) \rightarrow expr$ end” (for simplicity, we assume that programs are made out of a single module). The body of a function is an expression $expr$, which can include variables, literals (atoms, integers, floats, or the empty list), list constructors, tuples, let expressions, case expressions, function applications, calls to built-in functions, and try/catch blocks.

In a case expression “case $expr$ of $clause_1; \dots; clause_m$ end”, the expression $expr$ is first reduced to a value v , and then v is matched against the clauses “ pat when $expr_1 \rightarrow expr_2$ ” of the case expression. The first clause to match this value (i.e., the first clause for which there exists a substitution σ such that $v = pat \sigma$ and $expr_1 \sigma$ reduces to true) is selected, and the computation continues with the evaluation of the clause body (after updating the environment with σ).

Let us remark that primop expressions of the form “primop $atom (expr_1, \dots, expr_n)$ ” are primitive operation calls. In general, their evaluation is implementation-dependent, and they may have side effects or raise exceptions. However, in our setting, these are mainly used for raising exceptions in pattern matching errors.

The Erlang program in Figure 2 will successfully compile with no warnings in Erlang/OTP and will correctly compute the sum of the elements in L provided that L is a list of numbers.

¹Core Erlang is the intermediate language used by the standard Erlang compiler, which removes most of the syntactic sugar present in Erlang programs.

```

-module(sum_list).
-export([sum/1]).

sum(L) ->
  case L of
    [] -> 0;
    [H|T] -> H + sum(T)
  end.

```

Figure 2: A program in Erlang that computes the sum of all numbers in the input list L.

analysis on this program. In the following, we will see how our tool lists all the potential runtime errors, together with the input types that can cause them.

Otherwise, the program generates a runtime error. For instance, if the input to `sum` is an atom, then the program crashes and outputs a pattern matching error (`match_fail`), because there are no patterns that match an atom. Similarly, if the input to `sum` is a list of values, where at least one element is an atom, the execution halts with a type error (`badarith`), when applying the function ‘+’ to a non-numerical argument.

The tool Dialyzer does not generate any warnings when analyzing this program. That is coherent with the Dialyze approach, which only complains about type errors that would guarantee the program to crash for all input values. However, it might be the case that we want to perform a more detailed

2 Symbolic Interpreter for Runtime Error Detection of Erlang Programs

The main component of the verifier is a CLP interpreter that defines the operational semantics of our language. This executable specification of the semantics enables the execution of Erlang programs represented as Prolog facts.

Therefore, we have defined a translation from Erlang programs to Prolog facts. More precisely, our translator generates one fact `fundef` for each function definition occurring in the Core Erlang code obtained from the original Erlang module. For instance, the translation of the function `sum/1` defined in the module `sum_list` corresponds to the fact which can be seen in Figure 3.

```

fundef(lit(atom,'sum_list'),var('sum',1),
  fun([var('@c0')],
    case(var('@c0'),[
      clause([lit(list,nil)],lit(atom,'true'),
        lit(int,0)),
      clause([cons(var('H'),var('T'))],lit(atom,'true'),
        let([var('@c1')],apply(var('sum',1),[var('T')]),
          call(lit(atom,'erlang'),lit(atom,'+'),[var('H'),var('@c1')])))],
      clause([var('@c2')],lit(atom,'true'),
        primop(lit(atom,'match_fail'),
          [tuple([lit(atom,'case_clause'),var('@c2')]))])))).

```

Figure 3: Prolog fact generated by the Erlang-to-Prolog translation for the `sum` function definition.

This translation is quite straightforward, since the standard compilation from Erlang to Core Erlang greatly simplifies the code. Note that, since we generate Prolog facts, we have used the `cons` predicate for Erlang list constructors to distinguish them from Prolog list constructors. Note also that an additional *catch-all* clause has been added for the case in which the argument does not match any of the clauses from the case expression (i.e., for pattern matching errors). This transformation and similar ones are automatically made by the standard compilation from Erlang to Core Erlang.

The CLP interpreter provides a flexible means to perform the *bounded verification* of Erlang programs. By using a symbolic representation of the input data, the interpreter allows the exhaustive exploration of the program computations without explicitly enumerating all the concrete input values. In particular, the interpreter can run on input terms containing logic variables, and it uses constraint solvers to manipulate expressions with variables ranging over integer or real numbers. By fixing a bound to limit the number of computation steps performed by the interpreter, we force the exploration process to be finite, and hence either we detect a runtime error or we prove that the program is error-free up to the given bound.

Let us consider an Erlang program *Prog*, which is translated into Prolog facts as shown in Figure 3. In order to perform the bounded verification of the Erlang program *Prog*, the interpreter provides the predicate `run(FName/Arity,Bound,In,Out)`, whose execution evaluates the application of the function *FName* of arity *Arity* to the input arguments *In* in at most *Bound* steps. The arguments *In* are represented as a Prolog list (written using square brackets) of length *Arity*. *Out* is the result of the function application. If the evaluation of the function application generates an error, then *Out* is bound to a term of the form `error(Err)`, where *Err* is an error name (e.g., `match_fail`, indicating a match failure, or `badarith`, indicating an attempt to evaluate an arithmetic function on a non-arithmetic input), meaning that the specific error *Err* had occurred. Hence, the bounded verification of a given Erlang program can be performed by executing a query of the form:

```
?- run(FName/Arity,Bound,In,error(Err)).
```

where *FName* is a constant, *Arity* and *Bound* are non-negative integers, and *In* and *Err* are, possibly non-ground, terms.

Any answer to the query is a successful detection of the error *Err* generated by evaluating the application of the function *FName* to the input *In*. If no answer is found, then it means that no error is generated by exploring the computation of *FName* up to the value of *Bound*, and we say that the program *Prog* is correct up to the given bound.

Now let us see the bounded verifier in action by considering the `sum_list` program of the previous section and the following query:

```
?- run(sum/1,20,In,error(Err)).
```

Among the answers to the query, we get the following constraints relative to the input *In* and the error *Err*:

```
In=[cons(lit(Type,_V),lit(list,nil))],
Err=badarith,
dif(Type,int), dif(Type,float)
```

meaning that if `sum/1` takes as input a list (represented as a Prolog term of the form `cons(Head,Tail)`) whose head is not a numeric literal (denoted by the constraints `dif(Type,int)` and `dif(Type,float)`), then a `badarith` error occurs, that is, a non-numerical argument is given as input to an arithmetic operator. Another answer we get is:

```
In=[L],
Err=match_fail,
dif(L,cons(_Head,_Tail)), dif(L,lit(list,nil))
```

meaning that if `sum/1` takes as input an argument *L* which is neither a `cons` nor a `nil` term, then a `match_fail` error occurs. Note that, due to the recursive definition of `sum`, the bound is essential to detect this error.

Now let us introduce the predicate `int_list(L,M)` that generates lists `L` of integers of length up to `M`. For instance, the query

```
?- int_list(L,100).
```

generates the answer

```
L=cons(lit(int,N1),cons(lit(int,N2),...))
```

where `L` is a list of length 100 and `N1,N2,...,N100` are variables. If we give `L` as input to `sum` as follows:

```
?- int_list(L,100), run(sum/1,100,L,error(Err)).
```

the bounded verifier terminates after 0.347 seconds² with answer `false`, meaning that if the input to `sum` is any list of 100 integers, then the program is correct up to the bound 100. Note that no concrete integer element of the list is needed for the verification of this property.

Now we sketch the implementation of the operational semantics of Erlang expressions. The semantics is given in terms of a transition relation of the form `tr(Bound,ICfg,FCfg)`, which defines how to get the final configuration `FCfg` from the initial configuration `ICfg` in `Bound` computation steps. Configurations are pairs of the form `cf(Env,Exp)`, where `Env` is the environment mapping program variables to values and `Exp` is a term representing an Erlang expression.

The environment is extended with a boolean flag that keeps track of the occurrence of any runtime error during program execution. The value of the error flag `Flag` in the environment `Env` can be retrieved by using the predicate `lookup_error_flag(Env,Flag)`. The value of the flag in a given environment `EnvIn` can be updated by using the predicate `update_error_flag(EnvIn,Flag,EnvOut)`, thereby deriving the environment `EnvOut` whose error flag is set to `Flag`. If the evaluation of `IExp` generates the error `Err`, then `FExp` is a term of the form `error(Err)` and the error flag is set to `true`.

In Figure 4 we present the clause for `tr/3` that implements the semantics of function applications represented using terms of the form `apply(FName/Arity,IExps)`, where `FName` is the name of a function of arity `Arity` applied to the expressions `IExps`. The transition performed by `tr/3` only applies if:

1. no error has occurred so far, that is, `lookup_error_flag(IEnv,false)`, and
2. the bound has not been exceeded, that is, `Bound > 0`.

Then, the following operations are performed:

3. the value of the bound `Bound` is decreased,
4. the parameters `FParams` and the body `FBody` of the function `FName` of arity `Arity` are retrieved (the predicate `lookup_fun/3` is responsible for extracting `FParams` and `FBody` from the `fundef` fact representing `FName/Arity`),
5. the list of the actual parameters `IExps` is evaluated in `IEnv`, thereby deriving the list of expressions `EExps` and the new environment `EEnv` (it may differ from `IEnv` in the error flag and new variables occurring in the expressions `IExps` may have been added),
6. the formal parameters `FParams` are bound to the expressions `IExps`, thereby deriving the new environment `AEnv`,
7. the error flag in `AEnv` is updated to the value `F1` in `EEnv`, thereby deriving the environment `BEnv`,

²The query has been executed using SWI-Prolog v7.6.4 (<http://www.swi-prolog.org/>) on an Intel Core i5-2467M 1.60GHz processor with 4GB of memory under GNU/Linux OS

```

tr(Bound, cf(IEnv, IExp), cf(FEnv, FExp)) :-
    IExp = apply(FName/Arity, IExps),
    lookup_error_flag(IEnv, false),           % 1
    Bound > 0,                               % 2
    Bound1 is Bound - 1,                     % 3
    lookup_fun(FName/Arity, FPars, FBody),   % 4
    tr_list(Bound1, IEnv, IExps, EEnv, EExps), % 5
    bind(FPars, EExps, AEnv),                % 6
    lookup_error_flag(EEnv, F1),
    update_error_flag(AEnv, F1, BEnv),       % 7
    tr(Bound1, cf(BEnv, FBody), cf(CEnv, FExp)), % 8
    lookup_error_flag(CEnv, F2),
    update_error_flag(EEnv, F2, FEnv).       % 9

```

Figure 4: Definition of the operational semantics for a function application `apply/2`.

8. the body `FBody` is evaluated in `BEnv` to get the final expression `FExp`, and
9. the final environment `FEnv` is obtained from `EEnv` by setting the error flag to the value `F2` obtained from the callee function.

Each rule of the operational semantics for Erlang programs is translated into a clause for the predicate `tr/3`. These clauses are omitted.

Now we can present the definition of `run/4`, which depends on `tr/3`:

```

run(FName/Arity, Bound, In, Out) :-
    lookup_fun_pars(FName/Arity, FPars),
    bind(FPars, In, IEnv),
    tr(Bound, cf(IEnv, apply(FName/Arity, FPars)), cf(FEnv, Out)).

```

The predicate `run` retrieves the formal parameters `FPars` of `FName/Arity` and creates an environment `IEnv` where those parameters are bound to the input values `In`. Then, it evaluates the application of `FName` to its parameters, thereby producing the final expression `Out`.

3 Conclusions and future work

We have presented a work in progress for the development of a CLP interpreter for detecting runtime errors of Erlang programs. An Erlang program is first translated into a set of Prolog facts, then the CLP interpreter is run using this translation together with symbolic input data. At present, our interpreter is able to deal with first-order sequential Erlang programs, but we think that the extension to higher-order functions can be achieved by following a similar approach. In the future, we also plan to consider concurrency with an appropriate technique for handling the state explosion problem. For instance, we may employ a partial order reduction technique [1] to obtain the minimal set of concurrent behaviours for a given program, and then generate the associated executions using our interpreter.

Let us briefly compare our work with the static analysis tools available for Erlang. Unlike Dyalizer [13], our tool computes answer constraints that describe type-related input patterns which lead to runtime errors. However, as already mentioned, due to the bounded symbolic execution, our interpreter may terminate with no answer, even if runtime errors are possible for concrete runs which exceed the given bound. One of the weaknesses of Dyalizer is that it is hard to know where typing errors come from. An extension of Dyalizer that provides an explanation for the cause of typing errors has been proposed to overcome this problem [15]. We believe that we are able to provide a similar information if we include debugging information in the clauses generated by our Erlang-to-CLP translation.

Unlike SOTER [7], which is based on abstract interpretation, our CLP interpreter provides full support to arithmetic operations through the use of constraint solvers. Moreover, the symbolic interpreter does not require any user intervention (except for the bound), while in SOTER the user is responsible for providing a suitable abstraction.

Besides being useful on its own for bounded verification, the CLP interpreter for Erlang may be the basis for more sophisticated analysis techniques. In particular, by following an approach developed in the case of imperative languages, we intend to apply CLP transformation techniques to specialize the interpreter with respect to a given Erlang program and its symbolic input [6]. The specialized CLP clauses may enable more efficient bounded verification, and they can also be used as input to other tools for analysis and verification (such as constraint-based analyzers [3, 11] and SMT solvers [8, 14]), which have already been shown to be effective in other contexts [2, 4, 5]. Moreover, the specialized clauses can be used to apply backward analysis techniques for CLP programs based on abstract interpretation (see, for instance, [9, 12]). Backward analysis aims at deriving from a property that is expected to hold at the end of the execution of a program, conditions on the query which guarantee that the desired property indeed holds. In our context, backward analysis can be applied to deduce those conditions that may cause a runtime error, and then use them to improve the forward symbolic execution of the program.

Acknowledgements

We would like to thank the referees of HCVS 2018 for their useful and constructive comments. We also thank John Gallagher for stimulating discussions on the subject.

References

- [1] Parosh Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas (2014): *Optimal dynamic partial order reduction*. *ACM SIGPLAN Notices* 49(1), pp. 373–384, doi:10.1145/2535838.2535845.
- [2] Emanuele De Angelis, Fabio Fioravanti, Alberto Pettorossi, and Maurizio Proietti (2014): *Program verification via iterated specialization*. *Science of Computer Programming* 95, Part 2, pp. 149–175, doi:10.1016/j.scico.2014.05.017.
- [3] Emanuele De Angelis, Fabio Fioravanti, Alberto Pettorossi, and Maurizio Proietti (2014): *VeriMAP: A Tool for Verifying Programs through Transformations*. In: *Proceedings of the 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2014*, Lecture Notes in Computer Science, vol 8413, Springer, pp. 568–574, doi:10.1007/978-3-642-54862-8_47. Available at: <http://www.map.uniroma2.it/VeriMAP>.
- [4] Emanuele De Angelis, Fabio Fioravanti, Alberto Pettorossi, and Maurizio Proietti (2015): *A Rule-based Verification Strategy for Array Manipulating Programs*. *Fundamenta Informaticae* 140(3–4), pp. 329–355, doi:10.3233/FI-2015-1257.

- [5] Emanuele De Angelis, Fabio Fioravanti, Alberto Pettorossi, and Maurizio Proietti (2018): *Predicate Pairing for program verification*. *Theory and Practice of Logic Programming* 18(2), pp. 126–166, doi:10.1017/S1471068417000497.
- [6] Emanuele De Angelis, Fabio Fioravanti, Alberto Pettorossi, and Maurizio Proietti (2017): *Semantics-based generation of verification conditions via program specialization*. *Science of Computer Programming* 147, pp. 78–108, doi:10.1016/j.scico.2016.11.002.
- [7] Emanuele D’Osualdo, Jonathan Kochems, and C-H Luke Ong (2013): *Automatic verification of Erlang-style concurrency*. In: *Proceedings of the 20th International Static Analysis Symposium, SAS 2013*, Lecture Notes in Computer Science, vol 7935, Springer, pp. 454–476, doi:10.1007/978-3-642-38856-9_24.
- [8] Hossein Hojjat, Filip Konecný, Florent Garnier, Radu Iosif, Viktor Kuncak, and Philipp Rümmer (2012): *A Verification Toolkit for Numerical Transition Systems*. In *Proceedings of the 18th International Symposium on Formal Methods, FM 2012*, Lecture Notes in Computer Science, vol 7436, Springer, pp. 247–251, doi:10.1007/978-3-642-32759-9_21.
- [9] Jacob M. Howe, Andy King, and Lunjin Lu (2004): *Analysing Logic Programs by Reasoning Backwards*, In *Program Development in Computational Logic*, Lecture Notes in Computer Science, vol 3049, Springer, pp. 152–188. doi:10.1007/978-3-540-25951-0_6.
- [10] Joxan Jaffar, and Michael J. Maher (1994): *Constraint Logic Programming: A Survey*. *Journal of Logic Programming* 19/20, pp. 503–581, doi:10.1016/0743-1066(94)90033-7.
- [11] Bishoksan Kafle, John P. Gallagher, and José F. Morales (2016): *RAHFT: A Tool for Verifying Horn Clauses Using Abstract Interpretation and Finite Tree Automata*. In: *Proceedings of the 28th International Conference on Computer Aided Verification, CAV 2016*, Lecture Notes in Computer Science, vol 9779, Springer, pp. 261–268, doi:10.1007/978-3-319-41528-4_14.
- [12] Bishoksan Kafle, John P. Gallagher, Graeme Gange, Peter Schachte, Harald Søndergaard, and Peter J. Stuckey (2018): *An iterative approach to precondition inference using constrained Horn clauses*. CoRR abs/1804.05989 (To appear in *Theory and Practice of Logic Programming*). Available at <http://arxiv.org/abs/1804.05989>.
- [13] Tobias Lindahl, and Konstantinos Sagonas (2006): *Practical type inference based on success typings*. In: *Proceedings of the 8th ACM SIGPLAN international conference on Principles and practice of declarative programming, PPDP 2006*, ACM, pp. 167–178, doi:10.1145/1140335.1140356.
- [14] Leonardo M. de Moura, and Nikolaj Bjørner (2008): *Z3: An Efficient SMT Solver*. In: *Proceedings of the 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2008*, Lecture Notes in Computer Science, vol 4963, Springer, pp. 337–340, doi:10.1007/978-3-540-78800-3_24.
- [15] Konstantinos Sagonas, Josep Silva, and Salvador Tamarit (2013): *Precise explanation of success typing errors*. In: *Proceedings of the ACM SIGPLAN 2013 workshop on Partial evaluation and program manipulation, PEPM 2013*, ACM, pp. 33–42, doi:10.1145/2426890.2426897.
- [16] Robert Virding, Claes Wikström, and Mike Williams (1996): *Concurrent Programming in ERLANG (2nd Ed.)*. Prentice Hall International (UK) Ltd., ISBN-10: 013508301X, ISBN-13: 978-0135083017.