Model Checking Erlang Programs – LTL-Propositions and Abstract Interpretation

Frank Huch

Christian-Albrechts-University of Kiel,Germany

`fhu@informatik.uni-kiel.de`

Abstract: We present an approach for the formal verification of Erlang programs using abstract interpretation and model checking. In previous work we defined a framework for the verification of Erlang programs using abstract interpretation and LTL model checking. The application of LTL model checking yields some problems in the verification of state propositions, because propositions are also abstracted in the framework. While state propositions must be satisfied, negated state propositions have to be refuted. We show how this can be decided by means of the abstract domain. The approach is implemented as a prototype and we are able to prove properties like mutual exclusion or the absence of deadlocks and lifelocks for some Erlang programs.

1 Introduction

The programming language Erlang [AWV93], developed by Ericsson, is successfully used for the implementation of distributed systems. In [Hu99] we have developed a framework for abstract interpretations [CC77, JN94, SS98] for the formal verification of Erlang programs by model checking. This framework guarantees safety of an abstract operational semantics (*AOS*) with respect to the standard operational semantics (*SOS*). Since the AOS can contain more paths than the SOS in this framework, it is only possible to prove properties that have to be satisfied on all paths, like in linear time logic (*LTL*). If the abstraction satisfies a property expressed in LTL, then also the program satisfies it, but not vice versa. For finite domain abstract interpretations and an additional flow-abstraction [Hu02b] we obtain a finite AOS, for which model checking is decideable [LP85, Va96].

However, the application of LTL model checking to the AOS is not that straightforward. LTL is usually defined by means of simple state propositions. In our approach of verification by model checking, the validity of a state proposition depends on the chosen abstraction. However, in the context of abstraction, the equivalence of not validating a state proposition and refuting it does not hold. We show how safe model checking with respect to the abstract interpretation can be implemented.

The paper is organized as follows: Section 2 introduces the programming language Erlang and Section 3 shortly presents our framework for abstract interpretation. In Section 4 we add propositions to Erlang programs which can be used for formal verification by LTL model checking, shortly introduced in Section 5. Then, Section 6 motivates how abstract state propositions can be decided, which is formalized in Section 7. Finally, we present a concrete verification in Section 9 and conclude in Section 10.

2 Erlang

Erlang is a strict functional programming language, extended with concurrent processes. In Erlang, variables start with an uppercase letter and can be bound to arbitrary values, with no type restrictions. Basic values are atoms (starting with a lowercase letter), e.g. `true,false`. More complex data structures can be created by constructors for tuples of any arity ($\{\ldots\}/n$) and constructors for lists (`[.|.]` and `[]`). It is also possible to use atoms and these constructors in pattern matching.

In Erlang, a new process can be created by $\text{spawn}(f, [a_1, \ldots, a_n])$. The process starts with the evaluation of $f(a_1, \ldots, a_n)$. The functional result of `spawn` is the process identifier (*pid*) of the newly created process. With $p!v$ arbitrary values (including pids) can be sent to other processes, addressed by the pid p. The messages sent to a process are stored in its mailbox and the process can access them conveniently with pattern matching in a `receive`-statement. Especially, it is possible to ignore some messages and fetch messages from further behind. For more details see [AWV93].

For the formal verification we restrict to a core fragment, called *Core Erlang*, which contains the main features of Erlang. Its syntax is defined as follows:

$$p \quad ::= f(X_1, \ldots, X_n) \text{ -> } e. \mid p\,p$$
$$e \quad ::= \phi(e_1, \ldots, e_n) \mid X \mid pat = e \mid e_1, e_2 \mid e_1 ! e_2 \mid \text{case } e \text{ of } m \text{ end} \mid$$
$$\quad\quad \text{receive } m \text{ end} \mid \text{spawn}(f, e)$$
$$m \quad ::= pat_1 \text{->} e_1; \ldots; pat_n \text{->} e_n$$
$$pat ::= c(pat_1, \ldots, pat_n) \mid X$$

where ϕ represents predefined functions, functions of the program and constructors (c).

Example 1 A database process for storing key-value pairs can be defined as follows:

```
main() -> DB = spawn(dataBase,[[]]), spawn(client,[DB]), client(DB).

dataBase(L) -> prop(top), receive
      {allocate,Key,P} ->
          prop({allocate,P}),
          case lookup(Key,L) of
             fail -> P!free, receive
                               {value,V,P} -> prop({value,P}),
                                              dataBase([{Key,V}|L])
                             end;
             {succ,V} -> P!prop(allocated), dataBase(L)
          end;
      {lookup,Key,P} -> prop(lookup), P!lookup(Key,L), dataBase(L)
   end.
```

At this point the reader should ignore the applications of the function `prop`. In Section 9 we will verify the database and we will need the state propositions introduced by `prop`.

The program creates a database process holding a list `L` of key-value pairs. The communication interface of the database is given by the messages `{allocate,Key,P}` for allocating a new key and `{lookup,Key,P}` for retrieving the value of a stored key. In both patterns `P` is bound to the requesting client pid to which the answer is sent.

The allocation of a new key is done in two steps. First the key is received and checked. If there is no conflict, then the corresponding value can be received and stored in the database. This exchange of messages in more than one step has to guarantee mutual exclusion on the

439

database. Otherwise, it could be possible that two client processes send keys and values to the database and they are stored in the wrong combination. A client can be defined accordingly [Hu99]. In Section 9 we will verify this property for a database with two clients.

In [Hu99] we presented a formal semantics for Core Erlang. In the following we will refer to it as standard operational semantics (*SOS*). It is an interleaving semantics over a set of processes $\Pi \in States$. A process ($proc \in Proc$) consists of a pid ($\pi \in Pid := \{@n \mid n \in \mathbb{N}\}$), a Core Erlang evaluation term e and a word over constructor terms for the mailbox.

The semantics is a non-confluent transition system with interleaved evaluations of the processes. Only communication and process creation have side effects to the whole system. For modeling these actions two processes are involved. To give an impression of the semantics, we present the rule for sending a value to another process:

$$\frac{v_1 = \pi' \in Pid}{\Pi, (\pi, E[v_1 \, ! \, v_2], \mu)(\pi', e, \mu') \overset{! \, v_2}{\Longrightarrow} \Pi, (\pi, E[v_2], \mu)(\pi', e, \mu' : v_2)}$$

E is the context of the leftmost-innermost evaluation position, with v_1 a pid and v_2 an arbitrary constructor term. The value v_2 is added to the mailbox μ' of the process π' and the functional result of the send action is the sent value. For more details see [Hu02a].

3 Abstract Interpretation of Core Erlang Programs

In [Hu99] we developed a framework for abstract interpretations of Core Erlang programs. The abstract operational semantics (*AOS*) yields a transition system which includes all paths of the SOS. In an abstract interpretation $\widehat{\mathcal{A}} = (\widehat{A}, \widehat{\iota}, \sqsubseteq, \alpha)$ for Core Erlang programs $\widehat{A}$ is the abstract domain which should be finite for our application in model checking. The abstract interpretation function $\widehat{\iota}$ defines the semantics of predefined function symbols and constructors. Its codomain is $\widehat{A}$. Therefore, for example, it is not possible to interpret constructors freely in a finite domain abstraction. $\widehat{\iota}$ also defines the abstract behaviour of pattern matching in equations, case, and receive. Here the abstraction can yield additional non-determinism because branches can get undecidable in the abstraction. Hence, $\widehat{\iota}$ yields a set of results which defines possible successors. Furthermore, an abstract interpretation contains a partial order $\sqsubseteq$, describing which elements of $\widehat{A}$ are more precise than others. We do not need a complete partial order or a lattice because we do not compute any fixed point. Instead, we just evaluate the operational semantics with respect to the abstract interpretation[1]. This yields a finite transition system, which we use for (on the fly) model checking. An example for an abstraction of numbers with an ordering of the abstract representations is: $\mathbb{N} \sqsubseteq \{v \mid v \leq 10\} \sqsubseteq \{v \mid v \leq 5\}$. It is more precise to know, that a value is ≤ 5, than ≤ 10 than any number. The last component of $\widehat{\mathcal{A}}$ is the abstraction function α which maps every concrete value to its most precise abstract representation. Finally, the abstract interpretation has to fulfill five properties, which relate an abstract interpretation to the standard interpretation. They also guarantee that all paths of the SOS are represented

---

[1]Since we do not need a cpo, we use a partial order for the abstract domain. Therefore, the orientation of our abstract domain is upside down compared to standard frameworks using lattices. The least precise abstract value is the least element of the abstract domain.

in the AOS. An example for these properties is the following:

$$\text{For all } \phi/n \in \Sigma \cup C, v_1, \dots, v_n \in T_{\mathcal{C}}(Pid) \text{ and} \qquad \text{(P1)}$$

$$\widetilde{v}_i \sqsubseteq \alpha(v_i) \text{ it holds that } \phi_{\widehat{\mathcal{A}}}(\widetilde{v}_1, \dots, \widetilde{v}_n) \sqsubseteq \alpha(\phi_{\mathcal{A}}(v_1, \dots, v_n)).$$

It postulates, that evaluating a predefined function or a constructor on abstract values which are representations of some concrete constructor terms ($T_{\mathcal{C}}(Pid)$) yields abstractions of the evaluation of the same function on the concrete values. The other properties postulate correlating properties (P2-P5) for pattern matching in equations, `case`, and `receive`, and the pids represented by an abstract value. More details and some example abstractions can be found in [Hu99, Hu02a].

4 Adding Propositions

For a convenient specification of state propositions, we add propositions to our programs which are interpreted as state propositions in the AOS. As propositions we use arbitrary Core Erlang constructor terms which is very natural for Erlang programmers. For the definition of propositions we assume a predefined Core Erlang function `prop/1` with the identity as operational semantics. Hence, adding applications of `prop` does not effect the SOS nor the AOS. Nevertheless, as a kind of side-effect, the state in which `prop` is evaluated has the argument of `prop` as a valid state proposition, marked by the label prop in the AOS:

$$\Pi, (\pi, E[\texttt{prop}(v)], \mu) \overset{\textsf{prop}}{\longrightarrow}_{\widehat{\mathcal{A}}} \Pi, (\pi, E[v], \mu)$$

The valid state propositions[2] of a process and an abstract system state can be evaluated by the function $\textsf{prop}_{\widehat{\mathcal{A}}} : \widehat{Proc}_{\widehat{\mathcal{A}}} \longrightarrow \mathcal{P}(\widehat{A})$ where

$$\textsf{prop}_{\widehat{\mathcal{A}}}((\pi, E[e], \mu)) := \begin{cases} \{\widehat{v}\} \text{ , if } e = \texttt{prop}(\widehat{v}) \text{ and } \widehat{v} \in \widehat{A} \\ \emptyset \quad \text{ , otherwise} \end{cases}$$

The propositions of a state $\textsf{prop}_{\widehat{\mathcal{A}}} : \widehat{State}_{\widehat{\mathcal{A}}} \longrightarrow \mathcal{P}(\widehat{A})$ are defined as the union of all propositions of its processes. In Example 1 we have added four propositions to the database:

`top`	the main state of the database process
`{allocate,P}`	the process with pid `P` tries to allocate a key
`{value,P}`	the process with pid `P` enters a value into the database
`lookup`	a reading access to the database

In most cases, propositions will be added in a sequence as for example the proposition `top`. However, defining propositions by an Erlang function, it is also possible to mark existing (sub-)expressions as propositions. For example, the atom `allocated`.

5 Linear Time Logic

The abstract operational semantics is defined by a transition system. We want to prove properties (described in temporal logic) of this transition system using model checking. We use *linear time logic* (*LTL*) [Ga80] in which properties have to be satisfied on every path of a given transition system.

Defintion 2 (Linear Time Logic (LTL))
Let $Props$ be a set of state propositions. The set of *LTL-formulas* is defined as:

$$\varphi ::= P \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid X\varphi \mid \varphi_1 \, U \, \varphi_2$$

---

[2]For both functions we use the name $\textsf{prop}_{\widehat{\mathcal{A}}}$. The concrete instance of this overloading will be clear from the application of $\textsf{prop}_{\widehat{\mathcal{A}}}$. We will also omit the abstract interpretation in the index, if it is clear from the context.

An infinite word over sets of propositions $\pi = p_0 p_1 p_2 \ldots \in \mathcal{P}(Props)^\omega$ is called a path. A path π satisfies an LTL-formula φ ($\pi \models \varphi$) in the following cases:

$$p_0 \pi \models P \quad \text{iff } P \in p_0 \qquad\qquad \pi \models \neg\varphi \;\; \text{iff } \pi \not\models \varphi$$
$$\pi \models \varphi \wedge \psi \;\; \text{iff } \pi \models \varphi \text{ and } \pi \models \psi \qquad p_0 \pi \models X\varphi \;\; \text{iff } \pi \models \varphi$$
$$p_0 p_1 \ldots \models \varphi \, U \, \psi \;\; \text{iff } \exists i \in \mathbb{N} : p_i p_{i+1} \ldots \models \psi \text{ and } \forall j < i : p_j p_{j+1} \ldots \models \varphi$$

The propositional formulas $P \in Props$ are satisfied, if the first state of a path satisfies them. The next modality $X\varphi$ holds if φ holds in the continuation of the path. Finally, if φ holds until ψ holds and ψ finally holds, then $\varphi \, U \, \psi$ holds.

Formulas are not only interpreted with respect to a single path. Their semantics is extended to Kripke Structures $\mathcal{K} = (S, Props, \longrightarrow, \tau, s_0)$ with S a set of states, $Props$ a set of propositions, $\longrightarrow \subseteq S \times S$ the transition relation, $\tau : S \longrightarrow \mathcal{P}(Props)$ a labeling function for the states, and $s_0 \in S$ the initial state. A Kripke structure satisfies an LTL formula φ ($\mathcal{K} \models \varphi$) iff for all paths π of $\mathcal{K}$: $\pi \models \varphi$. The technique of *model checking* automatically decides, whether a given Kripke structure satisfies a given formula. For finite Kripke structures and the logic LTL model checking is decidable [LP85]. Some useful abbreviations can be defined as follows:

$$ff := \neg P \wedge P \qquad tt := \neg ff \qquad \varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi) \qquad \varphi \rightarrow \psi := \neg\varphi \vee \psi$$
$$F\varphi := tt \, U \, \varphi \qquad G\varphi := \neg F\neg\varphi \qquad F^\infty \varphi := G \, F \, \varphi \qquad\qquad G^\infty \varphi := F \, G \, \varphi$$

The propositional abbreviations are standard. $F\varphi$ is satisfied if there exists a position in the path, where φ holds, whereas in $G\varphi$ the formula φ has to hold in every position. The property $F^\infty \varphi$ postulates, that φ holds infinitely often on a path, whereas $G^\infty \varphi$ is satisfied, if φ is satisfied with only finitely many exceptions. In other words there is a position, from where on φ always holds.

For the verification of Core Erlang programs we use the AOS respectively the SOS of a Core Erlang program as a Kripke structure. We use the transition system which is spawned from the initial state (@0,main(),()). As labeling function for the states we use the function prop from the previous section.

6 Abstraction of Propositions

We want to verify Core Erlang programs with model checking. The framework for abstract interpretations of Core Erlang programs guarantees, that every path of the SOS is also represented in the AOS. If the resulting AOS is finite, then we can use simple model checking algorithms to check, if it satisfies a property φ expressed in LTL. If φ is satisfied in the AOS, then φ also holds in the SOS. In the other case model checking yields a counter example which is a path in the AOS on which φ is not satisfied. Due to the fact that the AOS contains more paths than the SOS, the counter example must not be a counter example for the SOS. Therefore, in this case it only yields a hint, that the chosen abstraction is too coarse and must be refined. The application of model checking seems to be easy but proving state propositions yields problems as the following example shows:

```
main() -> prop(42).
```
A possible property of the program could be $F\,\mathbf{42}$ (finally 42). To prove this property we use the AOS with an abstract interpretation, for instance the even-odd interpretation, which only contains the values **even** (representing all even numbers), **odd** (representing all odd numbers), and **?** (representing all values). With this abstraction the AOS yields the following transition system:

$$(\texttt{@0}, \texttt{main()}, ()) \longrightarrow (\texttt{@0}, \texttt{prop(42)}, ()) \longrightarrow (\texttt{@0}, \texttt{prop}(\mathbf{even}), ()) \overset{\mathsf{prop}}{\longrightarrow} (\texttt{@0}, \mathbf{even}, ())$$

Only the state $(\texttt{@0}, \texttt{prop}(\textbf{even}), ())$ has a property, namely **even**. **42** is an even number, but it is not the only even number. For safeness, this property cannot be proven. For instance, we could otherwise also prove the property F **40**.

It is only possible to prove properties for which the corresponding abstract value exclusively represents this value. However, it does not make much sense to abstract from special values and express properties for these values afterwards. Therefore, we only use propositions of the abstract domain, like F **even** (finally even). In the AOS the state $(\texttt{@0},\texttt{prop(even)},())$ has the property **even**. Therefore, the program satisfies this property. Now we consider a more complicated example:

```
main() -> prop(84 div 2).
```

This system satisfies the property too because $(84 \div 2) = 42$. On the other hand, in the even-odd abstraction we only get:

$(\texttt{@0}, \texttt{main()}, ()) \longrightarrow (\texttt{@0}, \texttt{prop(84 div 2)}, ()) \longrightarrow (\texttt{@0}, \texttt{prop(even div 2)}, ())$
$\longrightarrow (\texttt{@0}, \texttt{prop(even div even)}, ()) \longrightarrow (\texttt{@0}, \texttt{prop(?)}, ()) \longrightarrow (\texttt{@0}, \textbf{?}, ())$

with $\mathsf{prop}((\texttt{@0},\texttt{prop(?)},()))= \{\textbf{?}\}$ and $\emptyset$ as propositions of the other states. The result of the division of two even values must not be even. In a safe abstraction we cannot be sure, that the property F **even** is satisfied. Hence, model checking must yield, that it does not hold. For instance, the program

```
main() -> prop(42 div 2).
```

has a similar AOS but the property is not satisfied ($42 \div 2 = 21$). Therefore, a property is satisfied in a state, if the property of the state is at least as precise, as the expected property: $p_0 p_1 \ldots \models \widetilde{v}$ iff $\exists \widetilde{v}' \in p_0$ with $\widetilde{v} \sqsubseteq \widetilde{v}'$. However, this is not correct in all cases, as the following example shows. We want to prove that the program satisfies the property $G\neg$**even** (always not even) Therefore, one point is to check that the state $(\texttt{@0},\texttt{prop(?)},())$ models $\neg$**even**. With the definition from above we can conclude

$(\texttt{@0},\texttt{prop(?)},()) \not\models \textbf{even}$ and hence $(\texttt{@0},\texttt{prop(?)},()) \models \neg\textbf{even}.$

which is wrong. In the program `main() -> prop(84 div 2)` the property is not satisfied (see above). The SOS has the property 42, which is an even value. The reason is the non-monotonicity of $\neg$. Considering abstraction, the equivalence $\pi \models \neg\varphi$ iff $\pi \not\models \varphi$ does not hold! $\pi \not\models \varphi$ only means that $\pi \models \varphi$ is not safe. In other words, there can be a concretization which satisfies φ but we cannot be sure that it holds for all concretizations. Therefore, negation has to be handled carefully.

Which value of our abstract domain would fulfill the negated proposition $\neg$**even**? Only the proposition **odd** does. The values **even** and **odd** are incomparable and no value exists, which is more precise than these two abstract values. This connection can be generalized as follows: $p_0 p_1 \ldots \models \neg\widetilde{v}$ if $\forall \widetilde{v}' \in p_0$ holds $\widetilde{v} \sqcup \widetilde{v}'$ does not exist.
Note, that this is no equivalence anymore. The non-existence of $\widetilde{v} \sqcup \widetilde{v}'$ does only imply that $p_0 p_1 \ldots \models \neg\widetilde{v}$. It does not give any information for the negation $p_0 p_1 \ldots \models \widetilde{v}$. This (double) negation holds, if $\exists \widetilde{v}' \in p_0$ with $\widetilde{v} \sqsubseteq \widetilde{v}'$.

On a first sight refuting a state proposition seems not to be correct for arbitrary abstract interpretations. Consider the abstract domain where the abstract value $\leq$**0** represents the concrete values $\{\textbf{0}, -\textbf{1}, -\textbf{2}, \ldots\}$, $\geq$**0** represents $\{\textbf{0}, \textbf{1}, \textbf{2}, \ldots\}$, and **num** represents $\mathbb{Z}$. The represented concrete values of $\leq$**0** and $\geq$**0** overlap in the value **0**. Therefore, it would be incorrect that a state with the proposition $\leq$**0** satisfies the formula $\neg\geq$**0**.

$$\leq\textbf{0} \qquad\qquad \geq\textbf{0}$$
$$\textbf{num}$$

However, this abstraction is not possible. Any abstraction function $\alpha : A \longrightarrow \widehat{A}$ yields one single abstract representation for a concrete value. Without loss of generality, let $\alpha(\mathbf{0}) = \geq \mathbf{0}$. Abstract values which represent the concrete value $\mathbf{0}$ can only be the result of the use of the abstract interpretation function $\widehat{\iota}$. However, all these results $\widetilde{v}$ must be less precise: $\widetilde{v} \sqsubseteq \alpha(\mathbf{0}) = \geq \mathbf{0}$ because of the properties claimed by our framework. Hence, this abstract domain can be defined but the value $\leq \mathbf{0}$ does only represent the values $\{-\mathbf{1}, -\mathbf{2}, \ldots\}$. The name of the abstract value is not relevant. However, for understandability it should be renamed to $< \mathbf{0}$.

Alternatively, the abstract domain can be refined. The two overlapping abstract values can be distinguished by a more precise abstract value:

$< \mathbf{0}$	$\mathbf{0}$	$> \mathbf{0}$			$\mathbf{0}$
	$\leq \mathbf{0}$		$\geq \mathbf{0}$	or $\leq \mathbf{0}$	$\geq \mathbf{0}$
	num				**num**
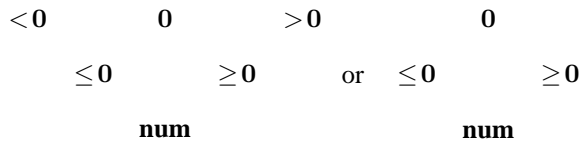
In both cases we must define $\alpha(\mathbf{0}) = \mathbf{0}$ because otherwise we have the same situation as before and the concrete value $\mathbf{0}$ is not represented by both abstract values $\leq \mathbf{0}$ and $\geq \mathbf{0}$.

7 Concretization of Propositions

With the advisement of the previous section we can now formalize whether a state proposition is satisfied or refuted, respectively. Similar results have been found by Clark, Grumberg, and Long [CGL94] and Knesten and Pnueli [KP98]. In their results Knesten and Pnueli handle a simple toy language in which only propositions on integer values can be made. In our framework state propositions are arbitrary Erlang values and their validity has to be decided with respect to an arbitrary abstract domain. The same holds for the paper of Clark et. al which also does not consider a real programming language. In the following, we present how these ideas can be transfered to formal verification of the real programming language Erlang.

First we define the concretization of an abstract value. This is the set of all concrete values which have been abstracted to the value, or a more precise value.

Defintion 3 (Concretization of Abstract Values)
Let $\widehat{\mathcal{A}} = (\widehat{A}, \widehat{\iota}, \sqsubseteq, \alpha)$ be an abstract interpretation. The concretization function $\gamma : \widehat{A} \longrightarrow \mathcal{P}(T_{\mathcal{C}}(Pid))$ is defined as $\gamma(\widetilde{v}) = \{v \mid \widetilde{v} \sqsubseteq \alpha(v)\}$.

For the last example we get the concretizations: $\gamma(\mathbf{0}) = \{\mathbf{0}\}$, $\gamma(\geq \mathbf{0}) = \{\mathbf{0}, \mathbf{1}, \mathbf{2}, \ldots\}$, $\gamma(\leq \mathbf{0}) = \{\mathbf{0}, -\mathbf{1}, -\mathbf{2}, \ldots\}$, and $\gamma(\mathbf{num}) = \mathbb{Z}$.

Lemma 4 (Connections between γ and α)
Let $\widehat{\mathcal{A}} = (\widehat{A}, \widehat{\iota}, \sqsubseteq, \alpha)$ be an abstract interpretation and γ the corresponding concretization function. Then the following properties hold:

1. $\forall v \in \gamma(\widetilde{v}) : \widetilde{v} \sqsubseteq \alpha(v)$　　　　　*2. $\bigsqcap\{\alpha(v) \mid v \in \gamma(\widetilde{v})\} = \widetilde{v}$*

For all proofs of this paper see [Hu02a]. With the concretization function we can define whether a state proposition of a state satisfies a proposition in the formula or refutes it.

Defintion 5 (Semantics of a State Proposition)
Let $\widehat{\mathcal{A}} = (\widehat{A}, \widehat{\iota}, \sqsubseteq, \alpha)$ be an abstract interpretation. A set of abstract state propositions

satisfies or refutes a proposition of a formula in the following cases:

$$p \models \widetilde{v} \text{ if } \exists \widetilde{v}' \in p \text{ with } \gamma(\widetilde{v}') \subseteq \gamma(\widetilde{v}) \quad \text{and} \quad p \not\models \widetilde{v} \text{ if } \forall \widetilde{v}' \in p \text{ holds } \gamma(\widetilde{v}) \cap \gamma(\widetilde{v}') = \emptyset$$

Similarly to these definitions for the concretization, we can decide whether a state proposition if satisfied or refuted for abstract values. For finite domain abstractions, this can be decided automatically.

Lemma 6 (Deciding Propositions in the abstract domain)

Let $\widehat{\mathcal{A}} = (\widehat{A}, \widehat{\iota}, \sqsubseteq, \alpha)$ be an abstract interpretation. A set of abstract state propositions satisfies or refutes a proposition of a formula in the following cases:

$$p \models \widetilde{v} \text{ if } \exists \widetilde{v}' \in p \text{ with } \widetilde{v} \sqsubseteq \widetilde{v}' \quad \text{and} \quad p \not\models \widetilde{v} \text{ if } \forall \widetilde{v}' \in p \text{ holds } \widetilde{v} \sqcup \widetilde{v}' \text{ does not exist}$$

Note, that we only show an implication. We can define unnatural abstract domains in which a property is satisfied or refuted with respect to Definition 5 but using only the abstract domain, we cannot show this. We consider the following abstract domain:

$$\mathbf{num} \sqsubseteq \mathbf{zero} \sqsubseteq \mathbf{0} \qquad \text{with } \alpha(v) = \mathbf{0}, \text{ if } v = \mathbf{0} \text{ and } \alpha(v) = \mathbf{num} \text{ otherwise}$$

The abstract value **zero** is superfluous because it represents exactly the same values, as the abstract value **0**. However, this abstract domain is valid. Using the definition of the semantics of a state proposition from Definition 5, we can show that $\{\mathbf{zero}\} \models \mathbf{0}$ because $\gamma(\mathbf{zero}) = \gamma(\mathbf{0}) = \{\mathbf{0}\}$. However, $\mathbf{zero} \sqsubseteq \mathbf{0}$ and we cannot show that $\{\mathbf{zero}\} \models \mathbf{0}$ just using the abstract domain.

The same holds for refuting state propositions. Consider the following abstract domain with $\alpha(v) = \geq \mathbf{0}$ if $v \geq \mathbf{0}$ and $\alpha(v) = \leq \mathbf{0}$ otherwise. In this domain the abstract value **0** is superfluous. Its concretization is empty. Hence, $\gamma(\leq \mathbf{0}) = \{-\mathbf{1}, -\mathbf{2}, \ldots\}$ and $\gamma(\geq \mathbf{0}) = \{\mathbf{0}, \mathbf{1}, \mathbf{2}, \ldots\}$. $\gamma(\leq \mathbf{0}) \cap \gamma(\geq \mathbf{0}) = \emptyset$ and $\leq \mathbf{0} \models \neg \geq \mathbf{0}$. However, this proposition cannot be refuted in this abstract domain since $\leq \mathbf{0} \sqcup \geq \mathbf{0} = \mathbf{0}$ exists.

These examples are unnatural because the domains contain superfluous abstract values. Nobody will define domains like these. Usually, the concretization of an abstract value is nonempty and differs from the concretizations of all other abstract values. In this case, deciding propositions in the abstract domain is complete with respect to the semantics of propositions. Although it is not complete in general, but it is safe.

8 Proving LTL Formulas

So far, we have discussed whether a state proposition is satisfied or refuted. However, in LTL negation is not only allowed in front of state propositions. Arbitrary sub-formulas can be negated. The basic idea is pushing negations into the formulas, like in [Va96]. We decide the state propositions in dependence of the negation in front. We distinguish positive (marked with $+$) and negative (marked with $-$) state propositions in dependence on an even respectively odd number of negations in front of them and decide whether they are satisfied respectively refuted by means of Lemma 6.

We implemented this model checker as a prototype, by which we are able to analyze small systems with up to 30000 states in the abstract model (AOS). The prototype also includes some standard abstractions (for which we proved correctness with respect to the claimed properties P1-P5 of [Hu99]) and a simple partial order reduction optimization. For the

verification of real systems, this prototype is not sufficient and a connection to an efficient model checker like Spin [Ho97] is necessary.

9   Verification of the Database

Now we want to verify the database process of Example 1, with two clients. We want to guarantee, that the process which allocates a key also sets the value:

> If a process π allocates a key, then no other process π' sets a value before π sets a value, or the key is already allocated.

This can for arbitrary processes be expressed in LTL as follows:

$$\bigwedge_{\pi \in Pid, \pi' \neq \pi} G\left(^-\{\mathbf{allocate}, \pi\} \longrightarrow (\neg\,^-\{\mathbf{value}, \pi'\})\, U\, (^+\{\mathbf{value}, \pi\} \vee\, ^+\mathbf{allocated})\right)$$

In our system only a finite number of pids occurs. Therefore, this formula can be translated into a pure LTL-formula as a conjunction of all possible permutations of possible pids which satisfy the condition. We have already marked the propositions in the formula with respect to the negations in front of them. Considering this marking the proposition $^-\{\mathbf{allocate}, \pi\}$ must be refuted. For instance, this is the case for the abstract values **top** and $\{\mathbf{lookup,?}\}$. However, **?** and $\{\mathbf{allocate},p\}$ with p the pid of the accessing client do not refute the proposition and the right side of the implication must be satisfied.

We can automatically verify this property using a finite domain abstraction, in which only the top-parts of depth 2 of the constructor terms are considered. The deeper parts of a constructor term are cut off and replaced by **?**. For more details see [Hu02a]. Our framework guarantees that the property also holds in the SOS and we have proven mutual exclusion.

10   Related Work and Conclusion

There exist other approaches for the formal verification of Erlang: EVT [NFG01] is a theorem prover especially tailored for Erlang. The main disadvantage of this approach is the complexity of proves. Especially, induction on infinite calculations is very complicated and automatization is not support yet. Therefore, a user must be an expert in theorem proving and EVT to prove system properties. For users it is easier to use our approach, since it is much closer to a push-button technique. This approach is also pursued by part of the EVT group [AE01]. They verified a distributed resource locker written in Erlang with a standard model checker. The disadvantage of this approach is that they can only verify finite state systems, while in practice, many systems have an infinite (or for model checkers too large) state space. As a solution, we think abstraction is needed to verify larger distributed systems. A similar approach for Java is made by the Bandera tool [HDL98]. They provide abstraction, but the user is responsible for consistent abstractions by explicitly defining abstracted methods in the source code. Hence, the user can define wrong abstractions and the correctness of their approach is not guaranteed. Furthermore, they have the same problems with negation as we discussed in this paper, but they do not explicitly discuss them and the user is responsible to solve them.

Our approach for the formal verification of Erlang programs uses abstract interpretation and LTL model checking. The main idea is the construction of a finite model in the AOS by means of a finite domain abstract interpretation. The abstraction is safe in the sense,

that all paths of the SOS are also represented in the AOS. For convenient verification we have added state propositions to the AOS. We showed how these state propositions can be decided by means of the abstract domain, with respect to the number of negations in front of them. Finally, we used this technique in the formal verification of the database process: we proved mutual exclusion for two accessing clients.

For future work, we want to transfer the presented and prototypically implemented results to a tool using existing model checkers like Spin [Ho97]. Here a problem is the integration of verifying or refuting properties in dependence of their positive or negative occurrences in formulas. α-Spin [Ga02] is a first approach tackling this problem. However, using their abstraction techniques will be difficult, because Spins specification language Promela and Erlang are quite different (imperative vs functional, different kinds of data structures).

Literatur

[AWV93] Joe Armstrong, Robert Virding, and Mike Williams. *Concurrent Programming in Erlang*. Prentice-Hall, Englewood Cliffs, NJ, 1993.

[AE01]   Thomas Arts and Clara Benac Earle. Development of a verified Erlang program for resource locking. In *Formal Methods in Industrial Critical Systems*, Paris, France, July 2001.

[CGL94] Edmund Clarke, Orna Grumberg, and David E. Long. Model checking and abstraction. *ACM Transactions on Programming Languages and Systems*, 16(5):1512–1542, September 1994.

[CC77]   Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction of approximation of fixed points. In *Proc. of the 4th ACM Symp. on Principles of Programming Languages*, pages 238–252, New York, NY, 1977. ACM.

[Ga02]   María del Mar Gallardo, Jesús Martínez, Pedro Merino, and Ernesto Pimentel. αSPIN: Extending SPIN with abstraction. In *Proc. of the 9th Int. SPIN Workshop on Model Checking of Software*, volume 2318 of *LNCS*, pages 254–258, 2002.

[Ga80]   Dov Gabbay, Amir Pnueli, Saharon Shelah, and Jonathan Stavi. On the temporal analysis of fairness. In *Conference Record of the Seventh Annual ACM Symposium on Principles of Programming Languages*, pages 163–173. ACM SIGACT and SIGPLAN, ACM Press, 1980.

[HDL98] John Hatcliff, Matthew B. Dwyer, and Shawn Laubach. Staging static analyses using abstraction-based program specialization. *LNCS*, 1490:134–148, 1998.

[Ho97]   Gerard J. Holzmann. The model checker SPIN. *IEEE Trans. on Software Engineering*, 23(5):279–295, May 1997. Special issue on Formal Methods in Software Practice.

[Hu99]   Frank Huch. Verification of Erlang programs using abstract interpretation and model checking. *ACM SIGPLAN Notices*, 34(9):261–272, September 1999. Proceedings of the ACM SIGPLAN Int. Conference on Functional Programming (ICFP '99).

[Hu02a] Frank Huch. Verification of Erlang programs using abstract interpretation and model checking. PhD thesis, RWTH Aachen, 2002.

[Hu02b] Frank Huch. Model checking Erlang programs - abstracting recursive function calls. In *Proceedings of WFLP'01, ENTCS*, volume 64. Elsevier, 2002.

[JN94]   Neil D. Jones and Flemming Nielson. Abstract interpretation: a semantics-based tool for program analysis. In *Handbook of Logic in Computer Science*. Oxford University Press, 1994. 527–629.

[KP98]   Yonit Kesten and Amir Pnueli. Modularization and abstraction: The keys to practical
         formal verification. In L. Brim, J. Gruska, and J. Zlatuska, editors, *The 23rd Int. Symp.
         on Mathematical Foundations of Computer Science*, volume 1450 of *LNCS*, pages 54–71.
         Springer, 1998.

[LP85]   Orna Lichtenstein and Amir Pnueli. Checking that finite state concurrent programs satisfy
         their linear specification. In *Proc. of the 12th ACM Symposium on Principles of Program-
         ming Languages*, pages 97–107, New Orleans, Louisiana, 1985. ACM SIGPLAN, ACM
         Press.

[NFG01]   Thomas Noll, Lars-åke Fredlund, and Dilian Gurov. The Erlang verification tool. In
         Proceedings of TACAS'01, volume 2031 of *LNCS*, pages 582–585. Springer, 2001.

[SS98]   David Schmidt and Bernhard Steffen. Program analysis as model checking of abstract
         interpretations. *LNCS*, 1503:351–380, 1998.

[Va96]   Moshe Y. Vardi. *An Automata-Theoretic Approach to Linear Temporal Logic*, volume
         1043 of *LNCS*, pages 238–266. Springer, New York, NY, USA, 1996.