ISRN SICS-R--97/01--SE

On the Verification of Open Distributed Systems

by

Mads Dam and Lars-åke Fredlund

On the Verification of Open Distributed Systems ¹

Mads Dam and Lars-åke Fredlund

Swedish Institute of Computer Science
Box 1263, S-164 28 Kista, Sweden
Email: mfd@sics.se and fred@sics.se

¹Work partially supported by the Computer Science Laboratory of Ericsson Telecom AB, Stockholm, EU Esprit BRA project 8130 LOMAPS, and a Swedish Foundation for Strategic Research Junior Individual Grant.

Abstract

A logic and proof system is introduced for specifying and proving properties of open distributed systems. Key problems that are addressed include the verification of process networks with a changing interconnection structure, and where new processes can be continuously spawned. To demonstrate the results in a realistic setting we consider a core fragment of the Erlang programming language. Roughly this amounts to a first-order actor language with data types, buffered asynchronous communication, and dynamic process spawning. Our aim is to verify quite general properties of programs in this fragment. The specification logic extends the first-order μ -calculus with Erlang-specific primitives. For verification we use an approach which combines local model checking with facilities for compositional verification. We give a specification and verification example based on a billing agent which controls and charges for user access to a given resource.

1 Introduction

A central feature of open distributed systems as opposed to concurrent systems in general is their reliance on modularity. Open distributed systems must accommodate addition of new components, modification of interconnection structure, and replacement of existing components without affecting overall system behaviour adversely. To this effect it is important that component interfaces are clearly defined, and that systems can be put together relying only on component behaviour along these interfaces. That is, behaviour specification, and hence verification, needs to be parametric on subcomponents. But almost all prevailing approaches to verification of concurrent and distributed systems rely on an assumption that process networks are static, or can safely be approximated as such, as this assumption opens up for the possibility of bounding the space of global system states. Clearly such assumptions square poorly with the dynamic and parametric nature of open distributed systems.

Our aim in this paper is to demonstrate an approach to system specification and verification that has the potentiality of addressing open distributed systems in general. We study the issue in terms of a core fragment of Ericsson's Erlang programming language [AVWW96] which we call Core Erlang. Core Erlang is essentially a first-order actor language (cf. [AMST97]). The language has primitives for local computation: data types, first-order abstraction and pattern matching, and sequential composition. In addition to this Core Erlang has a collection of primitives for component (process) coordination: sending and receiving values between named components by means of ordered message queues, and for dynamically creating new components. Similar to [AMST97] an operational semantics is given in terms of a two-level transition semantics: A transition relation on aggregate processes (or: composed system states) on top of one for local computation.

We use a temporal logic based on a first-order extension of the modal μ -calculus for the specification of component behaviour. In this logic it is possible to describe a wide range of important system properties, ranging from type-like assertions to complex interdependent safety and liveness properties. The development of this logic is actually fairly uncontroversial: To adequately describe component behaviour it is certainly needed to express potentialities of actions across interfaces and the (necessary and contingent) effects of these actions, to express properties of data types, to access component names, and to express properties of messages in transit.

The real challenge is to develop techniques that allow such temporal properties to be verified in a parametric fashion in face of the following basic difficulties:

1. Components can dynamically create other components.
2. Component names can be bound dynamically, thus dynamically changing component interconnection structure (similar to the case of the π -calculus [MPW92]).
3. Components are connected through unbounded message queues.
4. Through use of non-tail recursion components can give rise to local state spaces of unbounded size.
5. Basic data types such as natural numbers and lists are also unbounded.

We would expect some sort of uniformity in the answers to these difficulties. For instance techniques for handling dynamic process creation are likely to be adaptable to non-tail recursive constructions quite generally, and similarly message queues is just another unbounded data type.

In [Dam95] an answer to the question of dynamic process creation was suggested, cast in terms of CCS. Instead of closed correctness assertions of the shape $S : \phi$ (S is a system, ϕ its specification) which are the typical objects of state exploration based techniques the paper considered more general *open correctness assertions* of the shape $\Gamma \vdash S : \phi$ where Γ expresses assumptions $s : \psi$ on components s of S . Thus the behaviour of S is specified parametrically upon the behaviour of its component s . To address verification a sound and weakly complete proof system was presented, consisting of proof rules to reduce complex proof goals to (hopefully) simpler ones. Most importantly loop detection is applied to accommodate discharging of proof goals that are instances of proof goals already encountered during proof construction.

Our contribution in the present paper is to show how the approach of [Dam95] can be used to address the difficulties enumerated above for a fragment of a *real* programming language, and to show the utility of our approach on a concrete example exhibiting some of those difficulties. Thus many details of the proof system are kept very informal in the present paper, and we concentrate instead on an intuitive explanation of our approach, in particular with respect to the rule of discharge. The proof system captures both model checking-like reasoning (by state exploration) and parametric reasoning, and we motivate the discharge rule by first seeing how it reduces to well-known termination conditions in local model checking, and then generalising to open correctness assertions.

Example The formal machinery introduced in the paper will be illustrated by means of a running example, which is based on the following scenario: A user wants to access a resource paying for this using a given account. She therefore issues a request to a resource manager which responds by dynamically creating a billing agent to act as an intermediary between the user, the resource, and the users account. We view this scenario as quite typical of many security-critical mobile agent applications. The user is clearly taking a risk by exposing her account to the resource manager and the billing agent. One of these parties might violate the trust put in him eg. by charging for services not provided, or by passing information to third parties that should be kept confidential. Equally the resource manager need to trust the billing agent (and to some minor extent the user). We show how the system can be represented in Core Erlang, how a few critical properties can be expressed, and outline a proof of the desirable property that the number of transfers from the user account does not exceed the number of requests to use the resource.

2 Core Erlang

We introduce a core fragment of the Erlang programming language with dynamic networks of processes operating on data types such as natural numbers, lists, tuples, or process identifiers (pid's), using asynchronous, first-order call-by-value communication via unbounded ordered queues. Real Erlang has several additional features such as com-

munication guards, exception handling, modules, distribution extensions, and a host of built-in functions.

Processes A Core Erlang system consists of a number of processes computing in parallel, and communicating by means of ordered, unbounded message queues. Each process is named by a unique pid of which we assume an infinite supply. Associated with a process is an Erlang expression currently being evaluated and a queue of incoming messages, the process mailbox. Messages are sent by addressing a data value to a receiving process mailbox, identified using its pid. We use the notation $\{E, p, Q\}$ for a process with pid p , Erlang expression E , and queue Q . *System states*, or *aggregate processes*, S , are sets of processes, written using the grammar (where “ $\parallel$ ” expresses parallel execution):

$$S ::= \{E, p, Q\} \mid S \parallel S$$

Abstract Syntax We operate with the syntactical categories of *expressions* E , *matches* M , *patterns* P , and *values* V . The abstract syntax of Core Erlang is summarised as follows:

$$\begin{aligned} E &::= \Omega(E_1, \dots, E_n) \mid X \mid \text{self} \mid \text{case } E \text{ of } M \mid EE \mid \\ &\quad \text{spawn}(E, E) \mid E!E \mid \text{receive } M \text{ end} \mid E, E \\ M &::= P_1 \rightarrow E_1; \dots; P_n \rightarrow E_n \\ P &::= \Omega(P_1, \dots, P_n) \mid X \\ V &::= \Omega(V_1, \dots, V_n) \end{aligned}$$

X is an Erlang variable, and Ω ranges over a set of primitive constants and operations including zero 0, successor $E+1$, tupling $\{E_1, E_2\}$, the empty list $[]$, list prefix $[E_1|E_2]$, pid constants ranged over by p , and atom constants ranged over by a , f , and g . In addition we need constants and operations for message queues: ϵ is the empty queue and $Q_1 \cdot Q_2$ is queue concatenation.

Atoms are used to name functions. Expressions are interpreted relative to an environment of function definitions $f(X_1, \dots, X_n) \rightarrow E$. Such definitions should be considered sugared versions of definitions $f = \{X_1, \{\dots, X_n\}\} \rightarrow E$ assigning matches to function atoms. Each function atom f is assumed to be defined at most once in this environment.

Intuitive Semantics The intuitive behaviour of most operators, in the context of a pid p and queue Q , should not be too surprising:

- Ω is a data type constructor: To evaluate $\Omega(E_1, \dots, E_n)$, E_1 to E_n are evaluated in left-to-right order.
- **self** evaluates to p (the pid of the current process).
- **case** E of M is evaluated by first evaluating E to a value V , then matching V using M . If several patterns in M match the first one is chosen. Matching a pattern P_i of M against V can cause unbound variables to become bound in E_i ¹. In a function definition $f = M$ all free variables are considered as unbound.

¹This is not quite the binding the convention of Erlang proper: There the first occurrence of X in (case E_1 of $X \rightarrow E_2$), X can bind the second.

- $E_1 E_2$ is application: First E_1 is evaluated to a function atom f (there are no lambda abstractions in the language), then E_2 is evaluated to a value V , then the function definition of f is looked up and matched to V .
- $\text{spawn}(E_1, E_2)$ is evaluated by first evaluating E_1 to a function atom f , then E_2 to a value V , a new pid p' is generated, and a process with that pid is spawned evaluating $f V$. The value of the spawn expression is p' .
- $E_1!E_2$ evaluates E_1 to a pid p' , then E_2 to a value V , then V is sent to p' , resulting in V as the value of the send expression.
- $\text{receive } M \text{ end}$ inspects the process mailbox Q and retrieves the first element in Q that matches any pattern of M . Once such an element V has been found, evaluation proceeds analogously to $\text{case } V \text{ of } M$.
- E_1, E_2 is sequential composition: First E_1 is evaluated to a value, and then evaluation proceeds with E_2 .

Example: Billing Agents The introduction presented a scenario for securely accessing private resources. In this section we make the scenario more concrete by giving example Core Erlang functions that implements the proposed resource access scheme.

In the code examples below Erlang variables are upper-case while atoms are lower-case. Atoms are used to name functions (*rm*, *lookup* and *billagent*), but also for identifying message types (*contract*, *res*, *use* and *trans*) and message parameters (*ok* and *nok*).

```
rm(ResList, Bank, RAcc) →
  receive
    {contract, {Pu, UAcc}, From} →
      case lookup(Pu, ResList) of
        {ok, Pr} → From!{contract, ok, spawn(billagent, (Pr, Bank, RAcc, UAcc))};
        nok → From!{contract, nok}
      end
  end
end, rm(ResList, Bank, RAcc).
```

The *rm* function above manages accesses to private resources (a resource manager). As arguments it accepts a resource list, the pid of a trusted bank agent, and a private account as arguments. The resource list contains pairs of public and private “names”. The function *lookup*(*Pu*, *ResList*) (not shown here) searches for the public name *Pu* in *ResList*, and if found, returns the matching private name. The resource manager, after receiving a contract offer (identifying the paying account *UAcc*) and identifying the private name of the requested resource, spawns off a billing agent to mediate between the resource, the user and the bank.

```
billagent(Res, Bank, RAcc, UAcc) →
  receive
    {use, From} →
      Res!{use, self},
      receive
```

```

    {res, ok, Value} →
      Bank!({trans, UAcc, RAcc}, self),
    receive
      {{trans, UAcc, RAcc}, ok} → From!{use, ok, Value};
      {{trans, UAcc, RAcc}, nok} → From!{use, nok}
    end;
    {res, nok} → From!{use, nok}
  end
end, billagent(Res, Bank, RAcc, UAcc).

```

The billing agent coordinates access to the resource with withdrawals from the user's account. Upon receiving a request for the resource $\{use, From\}$ it attempts to acquire the resource, and if this succeeds (signalled by the reception of the tuple $\{use, ok, Value\}$), it attempts to transfer money from the user account to the resource manager account $Bank!(\{trans, UAcc, RAcc\}, self)$, and finally returns the resource to the user.

Operational Semantics The intuitive meaning of the operators is easily formalised as a Plotkin-style SOS semantics. Transitions are labelled by triples “ $pre, \alpha, post$ ” where pre is a necessary precondition for performing the transition (formulated in the logic introduced in Section 3), α is the action causing the transition and $post$ is the result (in terms of variable bindings etc.) of taking the transition. To facilitate the definition of the operational semantics (with regards to process spawning) we introduce the notion of a *process configuration* $\langle E, p, Q \rangle$, which has the same semantics as a process $\{E, p, Q\}$, except at a few key points. For instance we have rules like

$$input \quad \frac{X \text{ fresh}}{\langle E, p, Q \rangle \xrightarrow{true, p?X, true} \langle E, p, Q \cdot X \rangle}$$

reflecting that input transitions are always enabled,

$$spawn1 \quad \frac{\langle E_1, p, Q \rangle \xrightarrow{pre, \alpha, post} \langle E'_1, p, Q' \rangle}{\langle spawn(E_1, E_2), p, Q \rangle \xrightarrow{pre, \alpha, post} \langle spawn(E'_1, E_2), p, Q' \rangle}$$

to allow for expression evaluation in contexts, and

$$spawn3 \quad \frac{p' \text{ fresh}}{\langle spawn(f, V), p, Q \rangle \xrightarrow{true, spawn(f, V, p'), true} \langle p', p, Q \rangle}$$

Here α is a metavariable over transition types, either τ (corresponding to an internal, spontaneous computation step), output $p!V$, input $p?X$, or spawn, $spawn(f, V, p')$. Rules for the remaining operators are left out for concerns of space. They are, however, rather trivial given the intuition stated above.

It remains to state the rules governing computation steps for (aggregate) processes. These we state in full:

$$proc1 \quad \frac{\langle E, p, Q \rangle \xrightarrow{pre, \tau, post} \langle E', p, Q' \rangle}{\{E, p, Q\} \xrightarrow{pre, \tau, post} \{E', p, Q'\}}$$

$$\begin{array}{c}
\text{proc2} \quad \frac{\langle E, p, Q \rangle \xrightarrow{pre, p!V, post} \langle E', p, Q' \rangle}{\{E, p, Q\} \xrightarrow{pre, p!V, post} \{E', p, Q'\}} \\
\text{proc3} \quad \frac{\langle E, p, Q \rangle \xrightarrow{pre, \text{spawn}(f, V, p'), post} \langle E', p, Q' \rangle}{\{E, p, Q\} \xrightarrow{pre, \tau, post} \{E', p, Q'\} \parallel \{f \ V, p', \epsilon\}} \\
\text{com} \quad \frac{S_1 \xrightarrow{pre_1, p!V, post_1} S'_1 \quad S_2 \xrightarrow{pre_2, p?X, post_2} S'_2}{S_1 \parallel S_2 \xrightarrow{p = p' \wedge pre_1 \wedge pre_2, \tau, post_1 \wedge post_2 \wedge X = V} S'_1 \parallel S'_2} \\
\text{interleave1} \quad \frac{S_1 \xrightarrow{pre, \tau, post} S'_1}{S_1 \parallel S_2 \xrightarrow{pre, \tau, post} S'_1 \parallel S_2} \\
\text{interleave2} \quad \frac{S_1 \xrightarrow{pre, p!V, post} S'_1}{S_1 \parallel S_2 \xrightarrow{\text{foreign}(p)(S_2) \wedge pre, p!V, post} S'_1 \parallel S_2}
\end{array}$$

Here we use $p!V$ as a wildcard among $\{p!V, p?V\}$, and $\text{foreign}(p)$ for the predicate on system states S stating that p is not a pid of a process in S .

3 The Property Specification Logic

In this section we introduce a specification logic for Core Erlang. The logic is based on a first-order μ -calculus, extended with Erlang-specific features. Thus the logic is based on the first-order language of equality, extended with recursive (minimal and maximal) definitions, modalities reflecting the transition capabilities of processes and process configurations, along with a few additional primitives.

Syntax The abstract syntax of formulas is given as follows:

$$\begin{aligned}
\phi ::= & P = P \mid P \neq P \mid \text{term}(P) = P \mid \text{queue}(P) = P \mid \\
& \text{local}(P) \mid \text{foreign}(P) \mid \text{atom}(P) \mid \text{unevaluated}(P) \mid \\
& f(P_1, \dots, P_n) \mid \phi \wedge \phi \mid \phi \vee \phi \mid \forall X. \phi \mid \exists X. \phi \mid \\
& \langle \rangle \phi \mid \langle P!P \rangle \phi \mid \langle \phi \rangle \phi \mid [] \phi \mid [P!P] \phi \mid [\phi] \phi
\end{aligned}$$

Formula atoms are defined by parametrised recursive definitions. As is by now well known monotone recursive definitions in a complete boolean lattice have minimal and maximal solutions. For readability we use the notation $f(X_1, \dots, X_n) \Rightarrow \phi$ for maximal definitions (the greatest fixed point of a possibly recursive formula definition with formal parameters $X_1, \dots, X_n$), and $f(X_1, \dots, X_n) \Leftarrow \phi$ for minimal ones. Maximal solutions are used, typically, for safety, or invariant properties, while minimal solutions are used for liveness, or eventuality properties. We also use standard abbreviations like *true*, *false*, $\forall X_1, \dots, X_n. \phi$, etc.

Intuitive Semantics To limit space requirements we refrain from giving a formal semantics of formulas and make do instead with an informal one:

- Equality, inequality, boolean operators and the quantifiers take their usual meaning.
- The purposes of $\text{term}(P_1) = P_2$ and $\text{queue}(P_1) = P_2$ are to pick up the values of terms and queues associated with given pid's. $\text{term}(P_1) = P_2$ requires P_1 to be the pid of a process which is part of the system state being predicated (alt. is the pid of the predicated process configuration), and the Erlang expression associated with that pid to be identical to P_2 . Similarly $\text{queue}(P_1) = P_2$ holds if the queue associated with P_1 is P_2 .
- $\text{atom}(P)$ holds if P is an atom.
- $\text{local}(P)$ holds if P is the pid of a process in the system state being predicated, and analogously $\text{foreign}(P)$ holds if P is a pid and there is no process with pid P in the predicated system state.
- $\text{unevaluated}(P)$ holds if $\text{local}(P)$ does and the Erlang expression associated with P is not a ground value.
- $\langle \rangle \phi$ holds if an internal transition is enabled to a system state (process configuration) satisfying ϕ . $[] \phi$ is the dual of $\langle \rangle \phi$ (ie. all states following an internal transition satisfy ϕ). $\langle P_1!P_2 \rangle \phi$ holds if an output transition with appropriate parameters is enabled to a state (configuration) satisfying ϕ . $\langle P_1?P_2 \rangle \phi$ is used for input transitions. Finally $\langle \phi_1 \rangle \phi_2$ predicates process configurations only and holds if a process satisfying ϕ_1 can be spawned such that the ensuing configuration satisfies ϕ_2 . The boxed modalities, as before, are dual.

Some Simple Examples The combination of recursive definitions with data types makes the logic very expressive. For instance the type of natural numbers is defined as a minimal recursion, following the inductive definition of the set of natural numbers being the *least* set containing 0 and closed under the successor operation:

$$\text{nat}(P) \Leftarrow P = 0 \vee \exists X.(P = X + 1 \wedge \text{nat}(X))$$

Using this idea all involved data types can be defined quite easily. This in turn can be used to define a De-Morganised negation *not*. We can define “weak” modalities that are insensitive to the specific number of internal transitions in the following style:

$$\begin{aligned} [] \phi &\Rightarrow \phi \wedge [] [] \phi & [[X!Y]] \phi &= [] [X!Y] [] \phi \\ \langle \langle \rangle \rangle \phi &\Leftarrow \phi \vee \langle \rangle \langle \langle \rangle \rangle \phi & \langle \langle X!Y \rangle \rangle \phi &= \langle \langle \rangle \rangle \langle X!Y \rangle \langle \langle \rangle \rangle \phi \end{aligned}$$

Observe the use of formula parameters, and the use of “=” for non-recursive definitions. A temporal property like *always* is also easily defined:

$$\begin{aligned} \text{always}(\phi) &\Rightarrow \phi \\ &\wedge [] \text{always}(\phi) \\ &\wedge \forall X, Y. [X?Y] \text{always}(\phi) \\ &\wedge \forall X, Y. [X!Y] \text{always}(\phi) \end{aligned} \tag{1}$$

Eventuality operators are more delicate as progress is in general only made when internal or output transitions are taken. This can easily be handled, though, by nesting minimal and maximal definitions.

Billing Agents: Specification We enumerate some desired properties of the billing agent system introduced on Page 4.

Disallowing spontaneous account withdrawals. The first correctness property forbids spontaneous withdrawals from the user account by the billing agent. That is, the number of transfers from the user account should be less than or equal to the number of requests to use the resource.

$$\begin{aligned}
& \text{safe}(Ag, Bank, UAcc, N) \Rightarrow \\
& \quad [] \text{safe}(Ag, Bank, UAcc, N) \\
& \quad \wedge \quad \forall P, V. [P?V] \\
& \quad \quad \left(\begin{array}{l} P = Ag \wedge \exists Pid. V = \{use, Pid\} \wedge \text{safe}(Ag, Bank, UAcc, N + 1) \\ \vee \quad \text{safe}(Ag, Bank, UAcc, N) \vee \text{contains}(V, UAcc) \end{array} \right) \\
& \quad \wedge \quad \forall P, V. [P!V] \\
& \quad \quad \left(\begin{array}{l} \left(\begin{array}{l} P = Bank \wedge \exists Pid, Acc. V = \{\{trans, UAcc, Acc\}, Pid\} \\ \wedge \quad N > 0 \wedge \text{safe}(Ag, Bank, UAcc, N - 1) \end{array} \right) \\ \vee \left(\begin{array}{l} (P \neq Bank \vee \text{not}(\exists Pid, Acc. V = \{\{trans, UAcc, Acc\}, Pid\})) \\ \wedge \quad \text{safe}(Ag, Bank, UAcc, N) \end{array} \right) \end{array} \right)
\end{aligned}$$

So, a system consisting of a resource manager and a billing agent with pid Ag is considered *safe* if the difference N between the number of requests to use the resource ($\{use, Pid\}$) and the number of attempts to transfer money from the user account ($\{\{trans, UAcc, RAcc\}, Pid\}$) is always non-negative. The predicate $\text{contains}(Y, A)$ is defined via structural induction over an Erlang value expression Y and holds if a subexpression of Y is equal to A (definition omitted).

Expected Service is Received. Other interesting properties concern facts like: Denial of service responses correspond to failed money transfers, and returning the resource to the proper user. These sorts of properties are not hard to formalise in a style similar to the first example.

Preventing Abuse by a Third Party. The payment scheme presented here depends crucially on the non-communication of private names. For instance, even if we can prove that a resource manager or billing agent does not make illegal withdrawals nothing stops the resource manager from communicating the user account key to a third party, that can then access the account in non-approved ways.

Thus we need to prove at least that the system communicates neither the user account key nor the agent process identifier. Perhaps the service user also requests that her identity not be known outside of the system. In such a case the return process identifiers may not be communicated either. As an example, the property that the system does not communicate the user account key (unless it is also received by the process) is captured by $\text{notrans}(UAcc)$ given the definition below.

$$\begin{aligned}
notrans(A) \Rightarrow & \quad []notrans(A) \\
& \wedge \quad \forall X, Y. [X?Y] (contains(Y, A) \vee notrans(A)) \\
& \wedge \quad \forall X, Y. [X!Y] (not(contains(Y, A)) \wedge notrans(A))
\end{aligned}$$

4 Toward Parametric Verification

Consider a correctness assertion of the shape

$$\Gamma \vdash S : \phi. \quad (2)$$

where S is an arbitrary system state, ϕ an arbitrary specification, and Γ consists of first-order assumptions (like: $V_1 = V_2$) on value variables V in S and ϕ (a restriction that will be lifted in later sections). We wish to devise an adequate and general set of proof rules that allow us to prove such sequents. A natural starting point for this is the classical sequent calculus with equality. A general sequent will have the shape $\Gamma \vdash S : \Delta$ where Δ is a finite set, interpreted disjunctively, of formulas. Classical sequent calculus provides us with a standard collection of rules, easily adaptable to the sequent format of (2), for manipulating sequents, and for introducing connectives to the left and the right of the turnstile. For instance the rule for introducing conjunction to the right becomes

$$\text{AndR} \quad \frac{\Gamma \vdash S : \phi, \Delta \quad \Gamma \vdash S : \psi, \Delta}{\Gamma \vdash S : \phi \wedge \psi, \Delta}$$

We thus need only consider additions needed for the Erlang-specific primitives, for formula definitions, and for the modal operators. For the former we give a single example for the term construction:

$$\text{Term} \quad \frac{\Gamma \vdash \{V, P_1, Q\} \parallel S : P_1 = P_2, \Delta \quad \Gamma \vdash \{V, P_1, Q\} \parallel S : V = P_3, \Delta}{\Gamma \vdash \{V, P_1, Q\} \parallel S : \text{term}(P_2) = P_3, \Delta}$$

For recursively defined formulas we need rules for definition unfolding, for instance:

$$\text{UnfoldR} \quad \frac{\Gamma \vdash S : \phi[A_1/X_1, \dots, A_n/X_n], \Delta \quad f(X_1, \dots, X_n) \Rightarrow \phi}{\Gamma \vdash S : f(A_1, \dots, A_n), \Delta}$$

This leaves the modal operators. For these one option is to explore the transition relation, by rules of the following shape²:

$$\text{Diamond} \quad \frac{\Gamma, pre, post \vdash S' : \phi \quad \Gamma \vdash pre \quad S \xrightarrow{pre, \alpha, post} S'}{\Gamma \vdash S : <\alpha>\phi}$$

$$\text{Box} \quad \frac{\{\Gamma, pre, post \vdash S' : \phi_1, \dots, \phi_n \mid S \xrightarrow{pre, \alpha, post} S'\}}{\Gamma \vdash S : [\alpha]\phi_1, \dots, [\alpha]\phi_n}$$

Proofs built up using the proof rules introduced so far will rarely terminate, as both processes and their specifications are given recursively. We thus need a way of safely discharging an assumption once it is seen to be an instance of a proof goal which has already been encountered during proof construction. We discuss the ideas on the basis of an example.

²Another option is exemplified by the DiaSeq1 proof rule on Page 13.

Example Consider the following Core Erlang function:

$$\text{stream}(N, \text{Out}) \rightarrow \text{Out!}N, \text{stream}(N + 1, \text{Out}).$$

which outputs the increasing stream $N, N + 1, N + 2, \dots$ along Out . The specification of the program could be that it can always eventually output a value

$$\text{stream_spec}(\text{Out}) = \text{always}(\exists X. \ll \text{Out!}X \gg \text{true})$$

The goal sequent takes the shape

$$\text{Out} \neq P \vdash \{\text{stream}(N, \text{Out}), P, Q\} : \text{stream_spec}(\text{Out}). \quad (3)$$

That is, assuming $\text{Out} \neq P$ (since otherwise output will go to the stream itself), and started with pid P and any input queue Q , the property $\text{stream_spec}(\text{Out})$ will hold. The first step is to unfold the formula definition. Then, using the proof rules (by unfolding *always* and performing some trivial proof steps) goal (3) is easily reduced to the following subgoals:

$$\text{Out} \neq P \vdash \{\text{stream}(N, \text{Out}), P, Q\} : \ll \text{Out!}N \gg \text{true} \quad (4)$$

$$\text{Out} \neq P \vdash \{\text{stream}((N + 1), \text{Out}), P, Q\} : \text{stream_spec}(\text{Out}) \quad (5)$$

$$\text{Out} \neq P \vdash \{\text{stream}(N, \text{Out}), P, Q :: Y\} : \text{stream_spec}(\text{Out}). \quad (6)$$

Discharge Proving (4) is straightforward and does not require recourse to discharge. However, in subgoals (5) and (6) we have arrived at sequents which are both instances of a sequent already “reduced”, namely (3). Discharge at these points is indeed sound, for reasons explained below.

Consider a proof constructed using the proof rules introduced so far. Suppose an internal node of the proof is labelled $\Gamma \vdash S : \phi$, and suppose the proof has a leaf which is labelled $\Gamma' \vdash S' : \phi'$ such that $S' : \phi'$ is a substitution instance of $S : \phi$ under a substitution, say, ρ . That is, we have a scenario like the one in the stream example above. Suppose also that $\Gamma' \vdash \rho(\phi)$ whenever $\phi \in \Gamma$. That is, with the leaf node $\Gamma' \vdash S' : \phi'$ we have arrived at a situation which is a special case of a situation we have already considered. We wish to devise safe conditions for discharging the leaf node $\Gamma' \vdash S' : \phi'$. Informally it suffices, during proof construction, to tag each formula identifier with a unique label whenever it is looked up by applications of *UnfoldR* while not yet having been assigned a tag, and to assign another unique tag, say, a *colour*, to each member of the right-hand formula set Δ in a manner which is preserved by all applications of proof rules. Then, if substitutions are required to respect colours, we can introduce a notion of *formula regeneration path* by tracking, from any given $\phi \in \Delta$, using the colouring information, a member of any subsequent right-hand formula set with the same colour, until the leaf has been reached. In that situation we know for the set Δ' that we have tracked a formula ψ which, up to tagging of formula identifiers, is identical to $\rho(\phi)$. We can now inspect the formula regeneration path to pick the outermost tagged formula identifier which is unfolded along the regeneration path and which occurs, with that same tag, in ψ . If, for some such ϕ we can find such a formula identifier which is defined using a maximal definition then discharge is permitted, otherwise it is not.

What we have described is a proof system for local model checking in the style of the tableau systems of eg. [SW91] or [Win91], extended to cater for symbolic reasoning on value parameters. A more formal treatment of a similar proof system can be found in [Dam95].

Parametricity The proof rules introduced so far are unfortunately incapable of dealing satisfactorily with the complications discussed in the introduction. Features like dynamic process spawning or non-tail recursion cause state spaces to grow in an essentially unbounded manner which can not in general be captured by the discharge conditions given in the previous section. This applies, for instance, to the resource manager on Page 4. Moreover, many proof goals require data type induction, mechanisms for which have not been considered yet. In this section we introduce a generalisation of the model checking-like approach of the previous section that permits us to address these issues in a very general way.

Consider again a proof goal $\Gamma \vdash S : \phi$. Assume that S contains a recursively defined expression which when executed can spawn off some process E . Then, attempting to build a proof using the rules already introduced, we will eventually encounter a sequent of the shape

$$\Gamma \vdash S \parallel \{E, p, \epsilon\} : \psi. \quad (7)$$

Clearly we have no hope of terminating proof construction with only the proof rules introduced so far. The solution we propose is to replace components of the agent being predicated by variables of which sufficiently powerful properties can be assumed, for the entire system to be verifiable. This amounts to a cut, using a rule of the following shape:

$$\text{Cut} \quad \frac{\Gamma \vdash S_1 : \gamma' \quad \Gamma, s : \gamma' \vdash s \parallel S_2 : \gamma}{\Gamma \vdash S_1 \parallel S_2 : \gamma}$$

This rule allows us, intuitively, to abstract from computational details of a component process that are irrelevant for the desired property of the system as a whole. Now, to handle (7) we can then try to guess properties γ_1 and γ_2 for each of the component processes, and then cut out these components. First cutting out S results in the following two subgoals:

$$\Gamma \vdash S : \gamma_1 \quad (8)$$

$$\Gamma, s_1 : \gamma_1 \vdash s_1 \parallel \{E, p, \epsilon\} : \psi. \quad (9)$$

Observe that if the outermost specification ϕ is well chosen it will often be possible to choose γ_1 identical to ϕ itself. Moving on we would then use another cut to reduce (9) to two further subgoals

$$\Gamma, s_1 : \gamma_1 \vdash \{E, p, \epsilon\} : \gamma_2 \quad (10)$$

$$\Gamma, s_1 : \gamma_1, s_2 : \gamma_2 \vdash s_1 \parallel s_2 : \psi. \quad (11)$$

We have thus arrived in a situation where a proof goal involving a composite system (7) has been reduced to proof goals for the components, along with an open correctness assertion (9) which does not depend on the system components, only on their properties.

Proving properties of open correctness assertions, however, is a very different task from the problem of proving closed ones in that eg. the modal rules **Diamond** and **Box** do not apply. Moreover, the rule of discharge has to be revised in light of the more general sequent shape, and new sets of proof rules are needed.

Discharge Revisited Observe that in case S can recursively spawn new processes, as we assumed, the reduction of (7) to (8) may not have actually achieved anything, as S can just continue to spawn off new processes. The task is to terminate the loop. Assume for simplicity that we could choose $\gamma_1 = \phi$. In this case (8) and (2) are identical, and we would then hope to be able to discharge (8) for the same reasons as discharge for (5) was seen to be admissible (even though the path from (2) up to (8) arises for very different reasons (structural ones, the cut) than the corresponding path for (5) which uses the modal rules introduced above). In fact this hope turns out to be justified.

Since formulas are defined recursively also the subgoal (11) may give rise to a recursive proof structure. That is, in building a proof of (11) we may arrive at a situation where the goal sequent has the shape $\Gamma', s_1 : \gamma'_1, s_2 : \gamma'_2 \vdash s_1 \parallel s_2 : \psi'$ which is an instance of (11). In this case we will have the possibility of discharging when ψ was regenerated because of a maximal formula identifier, but in addition we will have the possibility of discharging when one of the γ'_i are regenerated because of a minimal formula identifier.

The difficulty is to devise side conditions that ensure that this intuition remains sound even when loops are nested. This issue was addressed in the paper [Dam95] to which we refer for details. For the examples given here, however, the basic intuition is sufficient and sound, whence we leave aside the wider issue for now.

To illustrate the power of the discharge principles notice how, using the representation of data types of section 3, it is possible to capture data type induction by unfolding minimally defined formula identifiers such as

$$\text{nat}(P) \Leftarrow P = 0 \vee \exists X.(P = X + 1 \wedge \text{nat}(X))$$

to the left of the turnstile.

Proof Rules for Open Correctness Assertions So far we have not discussed proof rules that handle the existence of left hand hypotheses of the shape $s : \phi$. What remains, besides the classical sequent calculus rules, are rules to handle modal operators to the left of the turnstile, and to unfold formula definitions (similar to **UnfoldR**). We need two monotonicity rules:

$$\begin{array}{c} \text{Mon1} \quad \frac{\Gamma, s : \phi, s : \phi_1, \dots, s : \phi_n \vdash s : \psi_1, \dots, \psi_m}{\Gamma, s : \langle \alpha \rangle \phi, s : [\alpha] \phi_1, \dots, s : [\alpha] \phi_n \vdash s : \langle \alpha \rangle \psi_1, \dots, \langle \alpha \rangle \psi_m} \\ \\ \text{Mon2} \quad \frac{\Gamma, s : \phi_1, \dots, s : \phi_n \vdash s : \psi}{\Gamma, s : [\alpha] \phi_1, \dots, s : [\alpha] \phi_n \vdash s : [\alpha] \psi} \end{array}$$

and then the following rules for combinations of $\parallel$ with the modalities:

$$\text{DiaPar1} \quad \frac{\Gamma, s_1 : \phi, s_2 : \psi \vdash P_1 = P_3 \quad \Gamma, s_1 : \phi, s_2 : \psi \vdash P_2 = P_4 \quad \Gamma, s_1 : \phi, s_2 : \psi \vdash s_1 \parallel s_2 : \gamma}{\Gamma, s_1 : \langle P_1!P_2 \rangle \phi, s_2 : \langle P_3?P_4 \rangle \psi \vdash s_1 \parallel s_2 : \langle \rangle \gamma}$$

$$\begin{array}{c}
\text{DiaPar2} \quad \frac{\Gamma, s : \phi \vdash s \parallel S : \psi}{\Gamma, s : \langle \rangle \phi \vdash s \parallel S : \langle \rangle \psi} \\
\\
\text{DiaPar3} \quad \frac{\Gamma, s : \phi \vdash s \parallel S : \psi \quad \Gamma \vdash S : \text{foreign}(P_1)}{\Gamma, s : \langle P_1!P_2 \rangle \phi \vdash s \parallel S : \langle P_1!P_2 \rangle \psi} \\
\\
\text{BoxPar1} \quad \frac{\begin{array}{c} \Gamma, s_1 : \phi_1, s_2 : []\psi_1, s_2 : \forall X, Y. [X!Y]\psi_2, s_2 : \forall X, Y. [X?Y]\psi_3 \vdash s_1 \parallel s_2 : \gamma \\ \Gamma, s_1 : []\phi_1, s_1 : \forall X, Y. [X!Y]\phi_2, s_1 : \forall X, Y. [X?Y]\phi_3, s_2 : \psi_1 \vdash s_1 \parallel s_2 : \gamma \\ \Gamma, s_1 : \phi_2, s_2 : \psi_3 \vdash s_1 \parallel s_2 : \gamma \\ \Gamma, s_1 : \phi_3, s_2 : \psi_2 \vdash s_1 \parallel s_2 : \gamma \end{array}}{\begin{array}{c} \Gamma, s_1 : []\phi_1, s_1 : \forall X, Y. [X!Y]\phi_2, s_1 : \forall X, Y. [X?Y]\phi_3, \\ s_2 : []\psi_1, s_2 : \forall X, Y. [X!Y]\psi_2, s_2 : \forall X, Y. [X?Y]\psi_3 \vdash s_1 \parallel s_2 : []\gamma \end{array}} \\
\\
\text{BoxPar2} \quad \frac{\begin{array}{c} \Gamma, s_1 : \phi, s_2 : \text{foreign}(P_1), s_2 : [P_1!P_2]\psi \vdash s_1 \parallel s_2 : \gamma \\ \Gamma, s_1 : [P_1!P_2]\phi, s_1 : \text{foreign}(P_1), s_2 : \psi \vdash s_1 \parallel s_2 : \gamma \end{array}}{\Gamma, s_1 : [P_1!P_2]\phi, s_2 : [P_1!P_2]\psi \vdash s_1 \parallel s_2 : [P_1!P_2]\gamma}
\end{array}$$

For these five rules we assume that Γ mentions neither s_1 , s_2 , nor s . A similar assumption applies to the monotonicity rules. In actual proofs we tend to use slight generalisations of the last two proof rules where the left hand assumptions like $s_1 : []\phi_1$ are allowed to be vectors.

So far we have presented only rules to reason compositionally about system states. To prove general properties also of non-tail recursive processes another layer is needed in the proof system, for sequents of the form $\Gamma \vdash \langle E, p, Q \rangle : \Delta$. This extension is quite voluminous though it offers no significant new problems. As an example one of the rules for handling sequences is shown below.

$$\text{DiaSeq1} \quad \frac{\Gamma, \langle x', p, q' \rangle : \phi \vdash \langle (x', E), p, q' \rangle : \psi}{\Gamma, \langle x, p, q \rangle : \langle \alpha \rangle \phi \vdash \langle (x, E), p, q \rangle : \langle \alpha \rangle \psi}$$

5 Verifying the Resource Manager

In this section the proof system is demonstrated by outlining a proof that the resource manager on Page 4 satisfies the *safe* specification defined on Page 8. For simplicity it is assumed that the manager knows of only one resource, with public name P_u and private P_r . The list $\{P_u, P_r\}$ is denoted with R_L . Furthermore R_P denotes the process identifier of the resource manager process, R_Q its input queue.

Since the definition of *safe* accepts a billing agent pid and a user account as parameters, the formula must be preceded by an initialisation phase:

$$\begin{array}{c}
\forall \text{PubRes}, U\text{Acc}, \text{From}, \text{Agent}. \\
[R_P? \{ \text{contract}, \{ \text{PubRes}, U\text{Acc} \}, \text{From} \}] \\
[[\text{From}! \{ \text{contract}, \text{ok}, \text{Agent} \}]] \text{safe}(\text{Agent}, \text{Bank}, U\text{Acc}, 0)
\end{array}$$

So we set out to prove the following sequent:

$$\begin{array}{c}
\Gamma \vdash \{ \text{rm}(R_L, \text{Bank}, R\text{Acc}), R_P, \epsilon \} \\
: \quad \forall \text{PubRes}, U\text{Acc}, \text{From}, \text{Agent}. [R_P? \{ \text{contract}, \{ \text{PubRes}, U\text{Acc} \}, \text{From} \}] \dots \quad (1)
\end{array}$$

The necessary assumptions on the non-equivalence of process identifiers (e.g., $R_P \neq P_r$) are collected in Γ . By application of simple proof steps – four applications of *ForallR* and then repeated applications of the rules for unfolding, elimination of conjunctions, and the rules for the box modality – the following proof state is reached:

$$\begin{aligned} \Gamma' \vdash & \{rm(R_L, Bank, RAcc), R_P, \epsilon\} || \{billagent(P_u, Bank, RAcc, UAcc), B_P, \epsilon\} \\ & : safe(B_P, Bank, UAcc, 0) \end{aligned} \quad (2)$$

where Γ' is Γ extended with the fact that B_P is a fresh process identifier. This is a critical proof state, where we must come up with a property of the resource manager, and a property of the billing agent, that are sufficiently strong to prove that their parallel composition satisfies the *safe* property. In general such a proof step may be very difficult, but here the choice is relatively simple:

- ϕ_a : The billing agent satisfies the *safe* property, i.e., $safe(B_P, Bank, UAcc, 0)$.
- ϕ_b : The billing agent communicates the user account to no process except the bank.
- ϕ_c : The resource manager does not communicate the user account, unless some other process first sends it the account. This property can be formulated as *notrans*($UAcc$) given the definition of the *notrans* property at the end of Section 3.
- ϕ_d : The resource manager does not send a message containing an atom *use* (a usage request to a billing agent).
- ϕ_e : The resource manager cannot receive messages sent to the Bank process.

Properties ϕ_b , ϕ_d and ϕ_e can easily be formulated in a manner similar to ϕ_c . Essentially these conditions guarantee that Bank transfers are the result of user requests, rather than incorrectly programmed billing agents or resource managers that exchange information with each other. The result of applying the Cut rule twice is the following proof obligations:

$$\Gamma' \vdash \{billagent(P_u, Bank, RAcc, UAcc), B_P, \epsilon\} : safe(B_P, Bank, UAcc, 0) \wedge \phi_b \quad (3)$$

$$\Gamma' \vdash \{rm(R_L, Bank, RAcc), R_P, \epsilon\} : notrans(UAcc) \wedge \phi_d \wedge \phi_e \quad (4)$$

$$\begin{aligned} \Gamma', s_1 : safe(B_P, Bank, UAcc, 0), s_2 : notrans(UAcc) \vdash \\ s_1 || s_2 : safe(B_P, Bank, UAcc, 0) \end{aligned} \quad (5)$$

The proof of (3) involves essentially well-known techniques for proving correctness of sequential programs. Instead we concentrate on on proving the first part of the conjunction in (4), or rather, the stronger statement that *rm* satisfies *notrans* even when its input queue is non-empty, as long as no element in the queue contains *UAcc*:

$$\Gamma', not(contains(R_Q, UAcc)) \vdash \{rm(R_L, Bank, RAcc), R_P, R_Q\} : notrans(UAcc) \quad (6)$$

To prove (6) we first unfold the definition of *notrans* and eliminate the conjunctions. In case of an input step ($[X?Y]$) either we are done immediately (if $contains(Y, P)$). Otherwise the resulting proof state is

$$\begin{aligned} \Gamma', not(contains(R_Q, UAcc)), not(contains(Y, UAcc)) \vdash \\ \{rm(R_L, Bank, RAcc), R_P, R_Q \cdot Y\} : notrans(UAcc) \end{aligned} \quad (7)$$

which can be rewritten into a proof state that can be discharged (see discussion on Page 10) against (6). The *rm* process can clearly not perform any output step so that part of the conjunction is trivially true. Thus only the internal step remains, and such a step must correspond to retrieving a value $\{contract, \{P_u', UAcc'\}, From\}$ from the queue. The resulting proof state is:

$$\begin{aligned} \Gamma', not(contains(R_Q, UAcc)) \vdash \\ \{case\ lookup(P_u', R_L) \dots, R_P, R_Q'\} : notrans(UAcc) \end{aligned} \quad (8)$$

where R_Q' is the queue resulting from the reception of the value. By repeating the above steps, i.e., handling input, output and internal steps eventually one reaches the goal:

$$\begin{aligned} \Gamma'' \vdash \{rm(R_L, Bank, RAcc), R_P, R_Q''\} \parallel \{billagent(P_u, Bank, RAcc, UAcc'), B_P', \epsilon\} \\ : notrans(UAcc) \end{aligned} \quad (9)$$

where Γ'' is $\Gamma', not(contains(R_Q, UAcc))$ together with inequalities involving the new process identifier B_P' . This goal is handled by applying Cut to the parallel composition using *notrans* as the cut formula both to the left and to the right. The resulting goals are:

$$\Gamma'' \vdash \{rm(R_L, Bank, RAcc), R_P, R_Q''\} : notrans(UAcc) \quad (10)$$

$$\Gamma'' \vdash \{billagent(P_u, Bank, RAcc, UAcc'), B_P', \epsilon\} : notrans(UAcc) \quad (11)$$

$$\begin{aligned} \Gamma'', Z : notrans(UAcc), V : notrans(UAcc) \vdash \\ Z \parallel V : notrans(UAcc) \end{aligned} \quad (12)$$

Goal (10) can be discharged against (6). Goal (11) is easy to prove, since no new processes are created (proof sketch omitted). Finally goals (5) and (12) are handled by applications of BoxPar1, BoxPar2, Unfold and discharging against previously seen goals.

The remaining goal is then to prove

$$\begin{aligned} \Gamma', not(contains(R_Q, UAcc)) \vdash \\ \{rm(R_L, Bank, RAcc), R_P, R_Q\} : \phi_d \wedge \phi_e \end{aligned} \quad (13)$$

which can be achieved in a manner similar to the proof of (6).

6 Concluding Remarks

We have introduced a specification logic and proof system for the verification of programs in a core fragment of Erlang, and illustrated its application on a small, but quite delicate, agent-based example. Our approach is quite general both regarding the kinds of languages and models that can be addressed, and the kinds of assertions that can be formulated³. In addition our approach permits the treatment of programming language constructs such as dynamic process creation, non-tail recursion and inductive data type definitions in a uniform way, via a powerful rule of discharge.

As the goal of this work is quite ambitious, i.e., to enable verification of open distributed systems as implemented using real programming languages, there remains a

³For instance we are not restricted, as in many other approaches to compositional verification, to linear-time logic, neither does the proof system rely on auxiliary features like history or prophecy variables.

number of shortcomings in the current work. First of all, we need to investigate the proof system in real applications. For this purpose a prototype proof checking tool has been produced based on the approach presented here, that can handle programs of a moderate size. Some support for automation of proof steps already exists (e.g., for the model checking fragment), but we also need to identify other classes of sequents that can be proved algorithmically. Other ongoing work focuses on integrating the operational semantics of Erlang more tightly with the proof systems (along the lines of [Sim95]), to improve the handling of process identifier scoping (but see [AD96] for an approach to this in the context of the π -calculus), and to handle fairness in a more elegant manner.

References

- [AD96] R. Amadio and M. Dam. A modal theory of types for the π -calculus. In *Proc. FTRTFT'96*, Lecture Notes in Computer Science, 1135:347–365, 1996.
- [AMST97] G. Agha, I. A. Mason, S. F. Smith, and C. L. Talcott. A foundation for actor computation. *J. Functional Programming*, 7:1–72, 1997.
- [AVWW96] J. Armstrong, R. Virding, C. Wikström, and M. Williams. *Concurrent Programming in Erlang (Second Edition)*. Prentice-Hall International (UK) Ltd., 1996.
- [Dam95] M. Dam. Compositional proof systems for model checking infinite state processes. In *Proc. CONCUR'95*, Lecture Notes in Computer Science, 962:12–26, 1995.
- [MPW92] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, I and II. *Information and Computation*, 100(1):1–40 and 41–77, 1992.
- [Sim95] A. Simpson. Compositionality via cut-elimination: Hennessy-Milner logic for an arbitrary GSOS. In *Proceedings, Tenth Annual IEEE Symposium on Logic in Computer Science*, pages 420–430, San Diego, California, 26–29 June 1995. IEEE Computer Society Press.
- [SW91] C. Stirling and D. Walker. Local model checking in the modal μ -calculus. *Theoretical Computer Science*, 89:161–177, 1991.
- [Win91] G. Winskel. A note on model checking the modal ν -calculus. *Theoretical Computer Science*, 83:157–187, 1991.