

# Specialization of Distributed Actors by Partial Evaluation

Germán Vidal

*MiST, DSIC, Universitat Politècnica de València*

Valencia, Spain

gvidal@dsic.upv.es

**Abstract**—Partial evaluation is a well-established technique for program specialization that might achieve dramatic runtime speedups. While it has been widely studied for sequential languages, partial evaluation of concurrent and distributed programs has received little attention. In this paper, we consider an asynchronous message passing language that can be seen as a simple but significant version of the concurrent and distributed language Erlang. We introduce a hybrid partial evaluation scheme for this language, and illustrate it with an example. A prototype tool has been implemented and is publicly available. To the best of our knowledge, this is the first approach to the partial evaluation of an asynchronous message passing language like Erlang.

**Index Terms**—partial evaluation, specialization, actor model, concurrency, distributed systems

## I. INTRODUCTION

The actor model [1] is a formalism for concurrent computation which is based on the notion of *actor*. In this model, actors can only communicate through messages (i.e., there is no shared memory). Every actor can send and receive messages (which are stored in a local mailbox until they are consumed), as well as spawn new actors. This formalism has been used for modeling distributed systems like, e.g., web services. Erlang [2] is a message passing concurrent and distributed functional programming language based on the actor model. It has been successfully used in the development of many applications where massive concurrency and fault-tolerance are essential. In this work, we consider a simplified but significant version of Erlang in order to focus on some interesting aspects of the language.

Partial evaluation is a well-known technique for program specialization [3]. Essentially, it takes a program and part of its input data, the so called *static* data, and produces a *residual* program which is specialized for these data. Following [3], let us consider that the *semantics* of a program  $P$  is denoted by  $\llbracket P \rrbracket$ , i.e.,  $\llbracket P \rrbracket$  is a function from  $P$ 's inputs to  $P$ 's outputs, and let  $\text{mix}$  be a partial evaluator. Given a program  $P$  with two inputs,  $in_1$  and  $in_2$ , the partial evaluation of  $P$  w.r.t.  $in_1$  (static data) is denoted by  $\llbracket \text{mix} \rrbracket(P, in_1)$ . Correctness is then stated as follows:

$$\llbracket \llbracket \text{mix} \rrbracket(P, in_1) \rrbracket(in_2) = \llbracket P \rrbracket(in_1, in_2)$$

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 and by the *Generalitat Valenciana* under grant PROMETEO-II/2015/013 (SmartLogic).

i.e., running the partially evaluated program with the remaining data  $in_2$ , the so called *dynamic* data, must return the same result as the original program with all input data  $in_1, in_2$ .

Performance improvements are basically achieved by propagating known values and unfolding operations when enough information is available. In some cases, dramatic speedups can be achieved (while performance penalties are unusual).

There are basically two types of partial evaluators, depending on when termination issues are considered: offline and online. *Offline* partial evaluators consider two separate stages: binding-time analysis (BTA) and proper partial evaluation. The BTA typically produces a program instrumented with termination annotations that ensure the termination of partial evaluation. In contrast, *online* partial evaluators are monolithic procedures that perform partial evaluation and deal with termination issues at the same time.

We can find many approaches to partial evaluation of functional languages in the literature (e.g., [4]–[6]). Most of these approaches consider an offline scheme (a notable exception being [7], which is online). Another close techniques are (positive) supercompilation [8] and narrowing-driven partial evaluation [9], [10], which are typically online and achieve more powerful optimizations than standard partial evaluation.

While partial evaluation has been widely studied for sequential languages, partial evaluation of concurrent and distributed programs has received little attention. Dealing with concurrent programs is difficult since the propagation of static data should also take into account shared data structures and/or the values sent through messages. One of the first approaches was introduced by Hosoya et al [11]. They present an online approach to partial evaluation of a concurrent process calculus, which is still far from a realistic programming language like Erlang. Most of the previous approaches, though, consider an offline partial evaluation scheme which includes a BTA that infers information about the possible runtime communications (e.g., [12]–[14]). While [12] considers a CSP-like language, [13] and [14] deal with concurrent ML. In all these cases, in contrast to our considered language, communication is synchronous. Moreover, they aim at removing communications at partial evaluation time, which may change the possible runtime behaviour of the programs.

In this paper, we present a novel approach to partial evaluation in Erlang. We consider a hybrid scheme where a static analysis is used to infer a sound underapproximation of the

<i>program</i>	::=	$fun_1 \cdots fun_n$
<i>fun</i>	::=	$f(X_1, \dots, X_n) \doteq e$
<i>lit</i>	::=	$Atom \mid Integer \mid Float \mid []$
$e \ni expr$	::=	$Var \mid lit \mid [e_1 e_2] \mid \{e_1, \dots, e_n\}$
		$\alpha(p_1, \dots, p_n) \mid f(p_1, \dots, p_n)$
		$let X = e_1 in e_2$
		$case e of cl_1; \dots; cl_m end$
		$spawn(f(p_1, \dots, p_n))$
		$send(p_1, p_2) \mid self()$
		$receive cl_1; \dots; cl_n end$
$cl \ni clause$	::=	$p \text{ when } e_1 \rightarrow e_2$
$p \ni pat$	::=	$Var \mid lit \mid [p_1 p_2] \mid \{p_1, \dots, p_n\}$

Fig. 1. Language syntax rules

messages sent at runtime (i.e., using the static data available at partial evaluation time). An online partial evaluator then takes this information into account in order to also specialize the statements for receiving messages. In this way, we can obtain a degree of specialization that previous approaches based solely on the specialization of sequential Erlang code (like [15]) cannot achieve. A prototype implementation of the partial evaluator, called `erlmix`, is publicly available from <https://github.com/mistupv/erlmix>.

## II. THE LANGUAGE

In this section, we introduce the syntax and (informal) semantics of the considered language. Its syntax can be found in Fig. 1. Here, a function is defined by a single rule, whose left-hand side contains different variables as arguments and whose right-hand side is an expression possibly containing atoms (i.e., user-defined constants), variables, numbers, tuples, lists,<sup>1</sup> calls to built-in operators (e.g., arithmetic or relational operators), function calls, let bindings, and case expressions. Atoms are typically denoted with roman letters, while variables start with an uppercase letter; built-in operators are denoted with the greek symbol  $\alpha$ , and user-defined functions with the letters  $f, g, h, \dots$ . Without loss of generality, we require the arguments of function and built-in calls to be patterns. This simplifies our developments, and functions with arbitrary arguments can still be transformed into the required form by using let bindings.

In addition to the previous constructs, we also have some basic non-functional operations with side-effects to spawn new processes, send messages, receive messages, and get the *process identifier* (pid) of a given process.

We consider the usual domains in this language: expressions, patterns (built from variables, literals, lists and tuples) and *values* (i.e., ground—without variables—patterns). Values are denoted by  $v, v', v_1, v_2, \dots$ . A *substitution*  $\theta$  is a mapping from variables to expressions, and  $Dom(\theta) = \{X \in Var \mid X \neq \theta(X)\}$  is its domain. Substitutions are usually denoted by sets of bindings, e.g.,  $\{X_1 \mapsto e_1, \dots, X_n \mapsto e_n\}$ . Substitutions are extended to morphisms from expressions

to expressions in the natural way. We consider a postfix notation for substitution application, i.e., given an expression  $e$  and a substitution  $\sigma$ , the application  $\sigma(e)$  is denoted by  $e\sigma$ . The identity substitution is denoted by  $id$ . Composition of substitutions is denoted by juxtaposition, i.e.,  $\theta\theta'$  denotes a substitution  $\theta''$  such that  $\theta''(X) = \theta'(\theta(X))$  for all  $X \in Var$ .

The intuitive semantics of the functional constructs is the usual one. Built-in's are evaluated using an external function `eval`. For user-defined function calls, we consider a call-by-value semantics where function parameters are evaluated before unfolding a call. In a let binding of the form `let X = e1 in e2`, we first evaluate  $e_1$  to a value, say  $v$ , and then proceed with the evaluation of  $e_2\{X \mapsto v\}$ . In a case expression “`case e of p1 when e1 → e'1; ...; pn when en → e'n end`” we first evaluate  $e$  to a value, say  $v$ ; then, we find (if it exists) the first clause  $p_i \text{ when } e_i \rightarrow e'_i$  such that  $v$  matches  $p_i$ , i.e., there exists a substitution  $\sigma$  for the variables of  $p_i$  such that  $v = p_i\sigma$ , and  $e_i\sigma$ —the *guard*—evaluates to *true*; we use the external function `eval` to evaluate guards, so the above condition is denoted by  $eval(e_i\sigma) = true$ . If these conditions hold, the case expression above reduces to  $e'_i\sigma$ . Guards can only contain calls to built-ins (typically, arithmetic and relational operators).

The execution of a program gives rise to a number of processes (actors) that interact only through message passing (i.e., there is no shared memory). Messages are stored in a FIFO queue, which is local to each process, until they are consumed by a receive expression. Message sending is asynchronous, while receive instructions block the execution of a process until a matching message reaches its local queue.

Let us now consider the functions with side-effects (most of which are built-in's in the Erlang programming language): `self`, `spawn`, `send`, and `receive`. The expression `self()` returns the pid of a process, while `send( $\pi, v$ )` sends a message  $v$  to the process with pid  $\pi$ , which will be eventually stored in  $\pi$ 's local queue. New processes are spawned with a call of the form `spawn( $f(v_1, \dots, v_n)$ )`, so that the new process begins with the evaluation of  $f(v_1, \dots, v_n)$ . Finally, an expression

$$\text{receive } p_1 \text{ when } e_1 \rightarrow e'_1; \dots; p_n \text{ when } e_n \rightarrow e'_n \text{ end}$$

traverses the messages in the process' queue until one of them matches a branch in the receive statement; i.e., it should find the *first* message  $v$  in the process' queue (if any) such that `case v of p1 when e1 → e'1; ...; pn when en → e'n end` can be reduced; then, the receive expression evaluates to the same expression to which the above case expression would be evaluated, with the additional side effect of deleting the message  $v$  from the process' queue. If there is no matching message in the queue, the process *suspends* its execution until a matching message arrives.

*Example 1:* Figure 2 shows a program implementing a simple client-server scheme with one server and a generic number  $N$  of clients. Execution starts with a call to `main( $N, L$ )`, where  $N$  is a natural number and  $L$  is a list. It calls to function `init( $N, L, S$ )` in order to create the  $N$  clients, where  $S$  is the pid of the current process (i.e., the server), and then calls function `server()` in order to start the server. Every client

<sup>1</sup>Following Erlang, `[]` denotes an empty list while `[H|T]` denotes a list with head element  $H$  and tail  $T$  (a list).

```

main(N, L) ≐ let S = self()
  in let W = init(N, L, S) in server()
server() ≐ receive
  {C, L, N} → let V = {ack, rep(L, N)}
    in let W = send(C, V)
      in server();
  X → error
end
init(N, L, S) ≐ case N of
  0 → done;
  M when M > 0
    → let C = spawn(client(M, L, S))
      in init(M - 1, L, S)
end
client(N, L, S) ≐ let C = self() in let W = send(S, {C, L, N})
  in receive
    {ack, Res} → ok
  end
rep(L, N) ≐ case L of
  [] → [];
  [H|R] → [aux(H, N)|rep(R, N)]
end
aux(X, N) ≐ case N of
  0 → [];
  M when M > 0 → [X|aux(X, M - 1)]
end

```

Fig. 2. A simple message-passing program

performs a simple request of the form  $\{C, L, N\}$  where  $C$  is the pid of the client. For each request, the server performs a simple operation: it replaces every element  $X$  of  $L$  by a list with  $N$  repetitions of element  $X$ , returns this new list to the client, and calls to function *server* again in an endless loop. If the message does not have the right structure, the catch-all clause “ $X \rightarrow \text{error}$ ” returns error and stops.

Despite its simplicity, the interesting point in this example is that, even if we know some of the input parameters, say the value of  $N$  in the initial call to *main*, existing specialization techniques for (sequential) functional programs (and also for Erlang [15]) are not able to achieve any specialization since the value  $N$  is passed to the server through messages.  $\square$

By lack of space, we will not introduce the formal semantics of the language in this paper but refer the interested reader to [16] where a similar language is introduced.

### III. THE SPECIALIZATION SCHEME

In this section, we present a specialization scheme for our asynchronous message passing language.

In the following, we consider a fixed program  $Pgm$  and assume that partial evaluation begins with a call of the form  $main(p_1, \dots, p_n)$ , where  $p_i$  is either a value  $v_i$  (a static parameter) or a variable  $X_i$  (a dynamic parameter).

#### A. Message Analysis

In general, a message analysis computes an approximation of the messages sent at runtime. In [17], for instance, the authors present a message analysis for Erlang that *overapproximates* the values of runtime messages. In contrast to [17], we need an analysis that takes into account the static data; moreover, an underapproximation is more appropriate in our

```

mix(calls, seen, Δ) =
  if calls = [] then ∅
  else let [{π, f(p1, ..., pn)}|calls'] = calls in
    let e = [f(p1, ..., pn)]Δ(π) in
    let callse = succ(π, e) in
    let nseen = seen ++ [{π, f(p1, ..., pn)}] in
    let ncalls = (calls' ++ callse) \ nseen in
    [f(p1, ..., pn) ≐ e] ++ mix(ncalls, nseen, Δ)

```

Fig. 3. Partial evaluation procedure

setting (here, the degree of precision would only affect to the quality of the specialization but not to its correctness).

Essentially, we consider a symbolic execution version of the language semantics in order to simulate all possible executions with the static data. Termination can be ensured, e.g., using a depth bound or a timeout. In this way, we compute a sound *underapproximation* of the (partially known) messages sent at runtime. The output of this stage is a set,  $\Delta$ , with pairs of the form  $\{\pi, p\}$ , where  $\pi$  is a pid and  $p$  is pattern, i.e., a partial value denoting a message sent to process  $\pi$ .

*Example 2:* Let us consider, for instance, the program of Example 1 and the initial call  $main(2, L)$ . The execution of (any instance of) this call will create three processes, with pids, e.g.,  $s$  (for server),  $c1$  (for client #1) and  $c2$  (for client #2). For simplicity, we consider that the assignment of new pids is deterministic, so they will be the same in all the executions of the program. In practice, one can use the Erlang built-in “register” to assign a fixed name to every process (as we do in *erlmix*). Here, the message analysis returns the following sets of messages:

$$\Delta = \left\{ \begin{array}{l} (s, \{c1, L, 1\}), (s, \{c2, L, 2\}), \\ (c1, \{\text{ack}, Res\}), (c2, \{\text{ack}, Res'\}) \end{array} \right\}$$

#### B. Online Partial Evaluation

The next stage is the proper partial evaluation. It takes the program  $Pgm$ , the initial call  $main(p_1, \dots, p_n)$ , and the set of (possibly partial) messages,  $\Delta$ , and produces a residual, specialized program. In contrast to other approaches, we do not change the runtime communications of the program (as in, e.g., [12]–[14]). In other words, in our setting, the same messages will be sent in the original and specialized programs at runtime, although these messages will (hopefully) be processed in a more efficient way thanks to the specialization.

The kernel of our specialization scheme is the partial evaluation procedure *mix* shown in Fig. 3. It takes a list of calls to be partially evaluated with their associated pids (initially  $\{s, main(p_1, \dots, p_n)\}$ ), a list of calls already partially evaluated (initially  $[\ ]$ ), and a set of messages  $\Delta$ . It then proceeds as follows:

- The first function call in *calls* is specialized using the partial evaluation semantics shown in Fig. 4 (see below), thus producing a residual expression  $e$ . The partial evaluation semantics has two additional parameters: the list of unfolded calls (initially empty), which is used to

ensure the so called *local* termination, and the list of messages for the current process, which is needed to produce specialized clauses in receive statements.

- Then, we update the list *seen* by adding the pair  $\{\pi, f(p_1, \dots, p_n)\}$  just partially evaluated, resulting in the new list *nseen*.
- Function *succ* is used to extract the function calls in *e*. The set *ncalls* removes from  $calls' ++ succ(\pi, e)$  those calls already partially evaluated (i.e., *nseen*). Here, we consider that  $++$  and  $\setminus$  denote list concatenation and list subtraction modulo *variants* (i.e., considering all variables as a single value).
- Finally, we produce a residual, specialized function,  $f(p_1, \dots, p_n) \doteq e$ , and call recursively to function *mix* with the updated parameters until there are no more calls to be specialized.

The partial evaluation semantics is shown in Fig. 4.

In the following, we let  $\llbracket e \rrbracket_q^A \downarrow$  denote the normal (or irreducible) form of *e* according to the partial evaluation semantics, i.e., an expression *e'* such that  $\llbracket e \rrbracket_q^A \Rightarrow \dots \Rightarrow e' \not\Rightarrow$ . By definition, *e'* cannot contain occurrences of  $\llbracket \cdot \rrbracket$ . Moreover, we consider a well-founded order [18], denoted by  $\triangleright$ , which guarantees that we cannot have an infinite sequence of the form  $e_1 \triangleright e_2 \triangleright \dots$ . Given a list of calls  $\mathcal{A}$  and another call *e*, we let  $\mathcal{A} \triangleright e$  denote that *e* is smaller than the last call with the same function symbol in  $\mathcal{A}$  (this is safe since the number of different function symbols in a program is finite). Let us now briefly explain the rules of the partial evaluation semantics:

- The first rules are self-explanatory: patterns are not partially evaluated (they are considered *values* at partial evaluation time), while tuples and lists are partially evaluated by recursively partially evaluating their elements.
- In rule **Built-in**, we first partially evaluate the parameters of the operation. If all of them can be reduced to a value, then we execute the built-in using the auxiliary call *eval*; otherwise, we residualize it.
- The partial evaluation of a function call starts by partially evaluating its arguments. Then, we distinguish two cases: if the new call is “smaller” than the last call with the same function symbol (if any) in  $\mathcal{A}$ , we unfold the call and continue partially evaluating it; we also add this call to the current set of unfolded calls. Otherwise, we residualize the call and partially evaluate its arguments.
- In a let expression of the form  $let\ X = e_1\ in\ e_2$ , we first partially evaluate  $e_1$  up to its normal form. If we get a pattern,  $p_1$ , then we continue with the partial evaluation of  $e_2\{X \mapsto p_1\}$ . Otherwise, we partially evaluate both  $e_1$  and  $e_2$  but residualize the let statement.
- Case expressions are partially evaluated in a similar way. First, the case argument is partially evaluated to its normal form. If we get a pattern, we execute the case statement and continue with the partial evaluation of the selected branch. Otherwise, we residualize the case statement and partially evaluate its clauses.
- The concurrent actions *spawn*, *send* and *self* are always

residualized (and their arguments partially evaluated).

- Finally, given a receive statement, we keep adding new, specialized clauses as long as there are messages in *q* that match (in a non-trivial way) some pattern and the corresponding guard, if any, evaluates to *true*. Once no more specialized clauses are added, the second rule residualizes the receive statement<sup>2</sup> and continue partially evaluating its clauses.

Termination of the partial evaluation semantics is guaranteed by the fact that no concurrent action is executed and the use of a well-founded order to avoid unfolding infinite calls.

### C. Post-Processing Renaming

Once the partial evaluation procedure of Fig. 3 terminates computing a residual program  $Pgm'$  for the set of calls  $\mathcal{C}$  (i.e., the left-hand sides of the program rules or, equivalently, the value of *seen* in the last iteration of the partial evaluation procedure), we apply a standard post-processing of renaming. Renaming is necessary for  $Pgm'$  to respect the syntax of Fig. 1. It also helps to further optimize the code, removing unnecessary data and avoiding overlapping functions.

Basically, we map every call  $f(p_1, \dots, p_n)$  in  $\mathcal{C}$  to an expression of the form  $f'(X_1, \dots, X_m)$ , where  $X_1, \dots, X_m$  are the variables of  $f(p_1, \dots, p_n)$ . We then use this mapping to rename the calls to these functions in the bodies of the program functions. By construction, we know that these functions must be instances of some call in  $\mathcal{C}$ . Therefore, given a call  $f(p'_1, \dots, p'_n)$  such that  $f(p'_1, \dots, p'_n) = f(p_1, \dots, p_n)\sigma$ , we replace it by  $f'(X_1, \dots, X_n)\sigma$ . See, e.g., [19], for more details on the post-processing of renaming.

### D. The Partial Evaluation Scheme in Practice

Let us now illustrate the specialization process with an example. Consider the program from Example 1 (Fig. 2). Given the initial call  $main(2, L)$ , the message analysis computes the set  $\Delta$  shown in Example 2. Thus, the partial evaluation procedure starts with the call  $mix(\llbracket \{s, main(2, L)\} \rrbracket, [], \Delta)$ .

Therefore, our first call to the partial evaluation semantics is  $\llbracket main(2, L) \rrbracket_q^{\Delta}$ , with  $q = \Delta(s) = [\{c1, L, 1\}, \{c2, L', 2\}]$ .

In the following, for clarity, we do not unfold all function calls even if they are smaller than previous calls; in this way, the partially evaluated program is kept as close to the original one as possible.<sup>3</sup> The associated derivation is as follows:

$$\begin{aligned}
& \llbracket main(2, L) \rrbracket_q^{\Delta} \\
& \Rightarrow \llbracket let\ S = self()\ in\ \dots \rrbracket_q^{A_1} \\
& \Rightarrow let\ S = self()\ in\ \llbracket let\ W = init(2, L, S)\ in\ \dots \rrbracket_q^{A_1} \\
& \Rightarrow let\ S = self()\ in\ let\ W = init(2, L, S)\ in\ \llbracket server() \rrbracket_q^{A_2} \\
& \Rightarrow let\ S = self()\ in\ let\ W = init(2, L, S)\ in\ server()
\end{aligned}$$

<sup>2</sup>Note that receive statements are never removed, in contrast to case expressions, for instance.

<sup>3</sup>The implemented tool allows the user to add annotations on some functions to achieve this effect.

Pattern	$\llbracket p \rrbracket_q^A \Rightarrow p$
List	$\llbracket [e_1   e_2] \rrbracket_q^A \Rightarrow \llbracket [e_1] \rrbracket_q^A \llbracket [e_2] \rrbracket_q^A$
Tuple	$\llbracket \{e_1, \dots, e_n\} \rrbracket_q^A \Rightarrow \{\llbracket e_1 \rrbracket_q^A, \dots, \llbracket e_n \rrbracket_q^A\}$
Built-in	$\llbracket \alpha(p_1, \dots, p_n) \rrbracket_q^A \Rightarrow v$ if $\llbracket p_i \rrbracket_q^A \Downarrow = v_i, i = 1, \dots, n$ and $\text{eval}(\alpha(v_1, \dots, v_n)) = v$ $\llbracket \alpha(p_1, \dots, p_n) \rrbracket_q^A \Rightarrow \alpha(\llbracket p_1 \rrbracket_q^A, \dots, \llbracket p_n \rrbracket_q^A)$ otherwise
Fun call	$\llbracket f(p_1, \dots, p_n) \rrbracket_q^A \Rightarrow \llbracket e\sigma \rrbracket_q^{A'}$ if $\llbracket p_i \rrbracket_q^A \Downarrow = p'_i, i = 1, \dots, n, \mathcal{A} \triangleright f(p'_1, \dots, p'_n), \mathcal{A}' = [f(p'_1, \dots, p'_n)   \mathcal{A}],$ $f(X_1, \dots, X_n) \rightarrow e \lll Pgm, \text{ and } \sigma = \{X_1 \mapsto p'_1, \dots, X_n \mapsto p'_n\}$ $\llbracket f(p_1, \dots, p_n) \rrbracket_q^A \Rightarrow f(\llbracket p_1 \rrbracket_q^A, \dots, \llbracket p_n \rrbracket_q^A)$
Let	$\llbracket \text{let } X = e_1 \text{ in } e_2 \rrbracket_q^A \Rightarrow \llbracket e_2\sigma \rrbracket_q^A$ if $\llbracket e_1 \rrbracket_q^A \Downarrow = p_1$ and $\sigma = \{X \mapsto p_1\}$ $\llbracket \text{let } X = e_1 \text{ in } e_2 \rrbracket_q^A \Rightarrow \text{let } X = \llbracket e_1 \rrbracket_q^A \text{ in } \llbracket e_2 \rrbracket_q^A$ otherwise
Case	$\llbracket \text{case } e \text{ of } cl_1; \dots; cl_n \text{ end} \rrbracket_q^A \Rightarrow \llbracket e'_i\sigma \rrbracket_q^A$ if $\llbracket e \rrbracket_q^A \Downarrow = p, cl_i = (p_i \text{ when } e_i \rightarrow e'_i)$ is the first clause such that $p_i\sigma = p$ and $\text{eval}(e_i\sigma) = \text{true}$ $\llbracket \text{case } e \text{ of } cl_1; \dots; cl_n \text{ end} \rrbracket_q^A \Rightarrow \text{case } \llbracket e \rrbracket_q^A \text{ of } \llbracket cl_1 \rrbracket_q^A; \dots; \llbracket cl_n \rrbracket_q^A \text{ end}$ otherwise
Spawn	$\llbracket \text{spawn}(f(p_1, \dots, p_n)) \rrbracket_q^A \Rightarrow \text{spawn}(f(\llbracket p_1 \rrbracket_q^A, \dots, \llbracket p_n \rrbracket_q^A))$
Send	$\llbracket \text{send}(p_1, p_2) \rrbracket_q^A \Rightarrow \text{send}(\llbracket p_1 \rrbracket_q^A, \llbracket p_2 \rrbracket_q^A)$
Receive	$\llbracket \text{receive } cl_1; \dots; cl_n \text{ end} \rrbracket_q^A \Rightarrow \llbracket \text{receive } cl_1; \dots; cl_{i-1}; p \text{ when } e_i\sigma \rightarrow e'_i\sigma; cl_i; \dots; cl_n \text{ end} \rrbracket_q^A$ if $p \in q, cl_i = (p_i \text{ when } e_i \rightarrow e'_i)$ is the first clause such that $p_i\sigma = p,$ $\text{eval}(e_i\sigma) = \text{true}$ and $\sigma$ is not a renaming $\llbracket \text{receive } cl_1; \dots; cl_n \text{ end} \rrbracket_q^A \Rightarrow \text{receive } \llbracket cl_1 \rrbracket_q^A; \dots; \llbracket cl_n \rrbracket_q^A \text{ end}$ otherwise
Self	$\llbracket \text{self}() \rrbracket_q^A \Rightarrow \text{self}()$
Clause	$\llbracket p \text{ when } e \rightarrow e' \rrbracket_q^A \Rightarrow p \text{ when } \llbracket e \rrbracket_q^A \rightarrow \llbracket e' \rrbracket_q^A$

Fig. 4. Partial evaluation semantics for actors

with  $\mathcal{A}_1 = [main(2, L)], \mathcal{A}_2 = [init(2, L, S), main(2, L)].$  Therefore, the associated residual rule is

$$main(2, L) \doteq \text{let } S = \text{self}() \text{ in} \\ \text{let } W = \text{init}(2, L, S) \text{ in } server()$$

Then, function succ returns  $[\{s, init(2, L, S)\}, \{s, server()\}],$  which are added to the set of calls to be partially evaluated. Let us first consider  $init(2, L, S)$ :

$$\llbracket init(2, L, S) \rrbracket_q^{A_1} \\ \Rightarrow \llbracket \text{case } 2 \text{ of } \dots \rrbracket_q^{A_1} \\ \Rightarrow \llbracket \text{let } C = \text{spawn}(client(2, L, S)) \text{ in } \dots \rrbracket_q^{A_1} \\ \Rightarrow \text{let } C = \text{spawn}(client(2, L, S)) \text{ in } \llbracket init(2-1, L, S) \rrbracket_q^{A_1} \\ \Rightarrow \text{let } C = \text{spawn}(client(2, L, S)) \text{ in } \llbracket \text{case } 1 \text{ of } \dots \rrbracket_q^{A_2} \\ \Rightarrow \dots \\ \Rightarrow \text{let } C = \text{spawn}(client(2, L, S)) \text{ in} \\ \text{let } C' = \text{spawn}(client(1, L, S)) \text{ in done}$$

with  $\mathcal{A}_1 = [init(2, L, S), \mathcal{A}_2 = [init(1, L, S), init(2, L, S)].$  Here,  $\mathcal{A}_1 \triangleright init(1, L, S)$  and, thus, the function call is unfolded. The associated residual rule is as follows:

$$init(2, L, S) \doteq \text{let } C = \text{spawn}(client(2, L, S)) \text{ in} \\ \text{let } C' = \text{spawn}(client(1, L, S)) \text{ in done}$$

Then, function succ extracts the following new calls to be partially evaluated:  $[\{c1, client(1, L, S)\}, \{c2, client(2, L, S)\}].$  The partial evaluation of these calls is very simple and we omit

it.<sup>4</sup> The next call to be partially evaluated—the most interesting one—is the following:

$$\llbracket server() \rrbracket_q^{A_1} \\ \Rightarrow \llbracket \text{receive } \{C, L, N\} \rightarrow \dots; X \rightarrow \dots \text{ end} \rrbracket_q^{A_1} \\ \Rightarrow \llbracket \text{receive } \{C, L, 1\} \rightarrow \dots; \{C, L, N\} \rightarrow \dots; \\ X \rightarrow \dots \text{ end} \rrbracket_q^{A_1} \\ \Rightarrow \llbracket \text{receive } \{C, L, 1\} \rightarrow \dots; \{C, L, 2\} \rightarrow \dots; \\ \{C, L, N\} \rightarrow \dots; X \rightarrow \dots \text{ end} \rrbracket_q^{A_1} \\ \Rightarrow \text{receive } \{C, L, 1\} \rightarrow \llbracket \dots \rrbracket_q^{A_1}; \{C, L, 2\} \rightarrow \llbracket \dots \rrbracket_q^{A_1}; \\ \{C, L, N\} \rightarrow \llbracket \dots \rrbracket_q^{A_1}; X \rightarrow \llbracket \dots \rrbracket_q^{A_1} \text{ end} \\ \Rightarrow \dots \\ \Rightarrow \text{receive} \\ \{C, L, 1\} \rightarrow \text{let } V = \{\text{ack}, rep(L, 1)\} \text{ in} \\ \text{let } W = \text{send}(C, V) \text{ in } server() \\ \{C, L, 2\} \rightarrow \text{let } V = \{\text{ack}, rep(L, 2)\} \text{ in} \\ \text{let } W = \text{send}(C, V) \text{ in } server() \\ \dots \\ \text{end}$$

with  $\mathcal{A}_1 = [server()].$  The associated residual rule is shown in Fig. 5. Here, function succ returns the following function calls:  $[\{s, rep(L, 1)\}, \{s, rep(L, 2)\}, \{s, rep(L, N)\}, \{s, server()\}].$  Observe that the last one is removed from the set in the

<sup>4</sup>Note that the receive statements are not modified since the messages gathered by the analysis are renamings of the receive's pattern.

```

main2(L) ≐ let S = self()
  in let W = init2(L, S) in server()

server() ≐ receive
  {C, L, 1} → let V = {ack, rep1(L)}
    in let W = send(C, V)
      in server();
  {C, L, 2} → let V = {ack, rep2(L)}
    in let W = send(C, V)
      in server();
  {C, L, N} → let V = {ack, rep(L, N)}
    in let W = send(C, V)
      in server();
  X → error
end

init2(L, S) ≐ let C = spawn(client2(L, S)) in
  let C' = spawn(client1(L, S)) in done

client1(L, S) ≐ let C = self() let W = send(S, {C, L, 1}) in
  receive
  {ack, Res} → ok
end

client2(L, S) ≐ let C = self() let W = send(S, {C, L, 2}) in
  receive
  {ack, Res} → ok
end

rep1(L) ≐ case L of [] → []; [H|R] → [[X]|rep1(R)] end
rep2(L) ≐ case L of [] → []; [H|R] → [[X, X]|rep2(R)] end

```

Fig. 5. Renamed partially evaluated program

algorithm of Fig. 3 since it has been already partially evaluated. Let us proceed with the first call:

$$\begin{aligned}
& \llbracket rep(L, 1) \rrbracket_q^1 \\
& \Rightarrow \llbracket case L of \dots \rrbracket_q^{A_1} \\
& \Rightarrow case L of [] \rightarrow \llbracket [] \rrbracket_q^{A_1}; [H|R] \rightarrow \llbracket [aux(H, 1)|rep(R, 1)] \rrbracket_q^{A_1} \\
& \Rightarrow case L of [] \rightarrow \llbracket [] \rrbracket_q^{A_1}; [H|R] \rightarrow \llbracket [aux(H, 1)|rep(R, 1)] \rrbracket_q^{A_1} \\
& \Rightarrow case L of [] \rightarrow \llbracket [] \rrbracket_q^{A_1}; [H|R] \rightarrow \llbracket [case 1 of \dots] \rrbracket_q^{A_2} \llbracket [rep(R, 1)] \rrbracket_q^{A_1} \\
& \Rightarrow \dots \\
& \Rightarrow case L of [] \rightarrow \llbracket [] \rrbracket_q^{A_1}; [H|R] \rightarrow \llbracket [X]|rep(R, 1) \rrbracket_q^{A_1}
\end{aligned}$$

with  $A_1 = [rep(L, 1)]$ ,  $A_2 = [aux(H, 1), rep(L, 1)]$ . Note that  $A_2 \triangleright aux(H, 0)$  but  $A_1 \not\triangleright rep(R, 1)$ . The partial evaluation of  $rep(L, 2)$  proceeds analogously, thus producing the residual rules

$$\begin{aligned}
rep(L, 1) & \doteq case L of [] \rightarrow \llbracket [] \rrbracket_q^{A_1}; [H|R] \rightarrow \llbracket [X]|rep(R, 1) \rrbracket_q^{A_1} \\
rep(L, 2) & \doteq case L of [] \rightarrow \llbracket [] \rrbracket_q^{A_1}; [H|R] \rightarrow \llbracket [X, X]|rep(R, 2) \rrbracket_q^{A_1}
\end{aligned}$$

and function succ does not return any new call from this code. The last call,  $rep(L, N)$ , returns the original functions  $rep$  and  $aux$ , so we omit it.

Finally, using a straightforward renaming for the residual rules, we get the specialized program shown in Fig. 5, together with the original definitions of functions  $rep$  and  $aux$ . Observe that there are several performance improvements: function  $init$  is not recursive anymore and, more importantly, the specialized functions  $rep1$  and  $rep2$  have completely unrolled the calls to the original function  $aux$ .

## IV. DISCUSSION

We have presented a novel approach to the specialization of distributed actors using a hybrid partial evaluation scheme. The first stage is a message analysis that computes a sound underapproximation of the messages sent starting from some initial call with partially known arguments. The second stage is an (online) partial evaluation procedure that uses not only the static data but also the output of the message analysis to further specialize the code. Local termination of the partial evaluation semantics is ensured using a well-founded order.

As future work, we plan to focus on formalizing and proving the correctness the message analysis. In principle, one can use a symbolic execution extension of the standard semantics with some mechanism to ensure termination.

## ACKNOWLEDGMENT

We would like to thank Adrián Palacios for his support with the implementation of `erlmix`.

## REFERENCES

- [1] C. Hewitt, P. B. Bishop, and R. Steiger, "A universal modular ACTOR formalism for artificial intelligence," in *Proc. of IJCAI'73*, N. J. Nilsson, Ed. William Kaufmann, 1973, pp. 235–245.
- [2] F. Cesarini and S. Thompson, *Erlang Programming - A Concurrent Approach to Software Development*. O'Reilly, 2009.
- [3] N. Jones, C. Gomard, and P. Sestoft, *Partial Evaluation and Automatic Program Generation*. Prentice-Hall, Englewood Cliffs, NJ, 1993.
- [4] N. D. Jones, P. Sestoft, and H. Søndergaard, "Mix: A self-applicable partial evaluator for experiments in compiler generation," *Lisp and Symbolic Computation*, vol. 2, no. 1, pp. 9–50, 1989.
- [5] A. Bondorf and O. Danvy, "Automatic autoprojection of recursive equations with global variables and abstract data types," *Science of Computer Programming*, vol. 16, no. 2, pp. 151–195, 1991.
- [6] J. Launchbury, "A strongly-typed self-applicable partial evaluator," in *Proc. of FPCA'91*, vol. 523. Springer LNCS, 1991, pp. 145–164.
- [7] D. Weise, R. Conybeare, E. Ruf, and S. Seligman, "Automatic online partial evaluation," in *Proc. of FPCA'91*, vol. 523. Springer LNCS, 1991, pp. 165–191.
- [8] M. Sørensen, R. Glück, and N. Jones, "A Positive Supercompiler," *Journal of Functional Programming*, vol. 6, no. 6, pp. 811–838, 1996.
- [9] E. Albert and G. Vidal, "The narrowing-driven approach to functional logic program specialization," *New Generation Computing*, vol. 20, no. 1, pp. 3–26, 2002.
- [10] G. Vidal, "A partial evaluation tool for multi-paradigm declarative languages," in *Proc. of SMC 2002*, vol. 1, 2002, pp. 194–199.
- [11] H. Hosoya, N. Kobayashi, and A. Yonezawa, "Partial evaluation scheme for concurrent languages and its correctness," in *Proc. of Euro-Par'96*, vol. 1123. Springer LNCS, 1996, pp. 625–632.
- [12] M. Marinescu and B. Goldberg, "Partial-evaluation techniques for concurrent programs," in *Proc. of PEPM'97*, ACM, 1997, pp. 47–62.
- [13] M. Martel and M. Gengler, "Partial evaluation of concurrent programs," in *Proc. of Euro-Par'01*, vol. 2150. Springer LNCS, 2001, pp. 504–513.
- [14] J. H. Reppy and Y. Xiao, "Specialization of CML message-passing primitives," in *Proc. of POPL 2007*, ACM, 2007, pp. 315–326.
- [15] G. Weinholt, "Supercompiling Erlang," Master's thesis, Chalmers University of Technology, 2013.
- [16] I. Lanese, N. Nishida, A. Palacios, and G. Vidal, "A theory of reversibility for Erlang," *Journal of Logical and Algebraic Methods in Programming*, 2018. *To appear*.
- [17] R. Carlsson, K. Sagonas, and J. Wilhelmsson, "Message analysis for concurrent programs using message passing," *ACM Transactions on Programming Languages and Systems*, vol. 28, no. 4, pp. 715–746, 2006.
- [18] B. Martens, D. D. Schreye, and T. Horváth, "Sound and complete partial deduction with unfolding based on well-founded measures," *Theoretical Computer Science*, vol. 122, no. 1, pp. 97–117, 1994.
- [19] J. G. Ramos, J. Silva, and G. Vidal, "Fast Narrowing-Driven Partial Evaluation for Inductively Sequential Systems," in *Proc. of ICFP'05*, ACM Press, 2005, pp. 228–239.