

# Session Fidelity for ElixirST: A Session-Based Type System for Elixir Modules

Gerard Tabone

Adrian Francalanza

Computer Science Department  
University of Malta  
Msida, Malta

gerard.tabone.17@um.edu.mt

adrian.francalanza@um.edu.mt

This paper builds on prior work investigating the adaptation of session types to provide behavioural information about Elixir modules. A type system called ElixirST has been constructed to statically determine whether functions in an Elixir module observe their endpoint specifications, expressed as session types; a corresponding tool automating this typechecking has also been constructed. In this paper we formally validate this type system. An LTS-based operational semantics for the language fragment supported by the type system is developed, modelling its runtime behaviour when invoked by the module client. This operational semantics is then used to prove session fidelity for ElixirST.

## 1 Introduction

In order to better utilise recent advances in microprocessor design and architecture distribution, modern programming languages offer a variety of abstractions for the construction of concurrent programs. In the case of message-passing programs, concurrency manifests itself as spawned computation that exhibits *communication as a side-effect*, potentially influencing the execution of other (concurrent) computation. Such side-effects inevitably increase the complexity of the programs produced and lead to new sources of errors. As a consequence, program correctness becomes harder to verify and language support for detecting errors at the development stage can substantially decrease the number of concurrency errors.

Elixir [34], based on the actor model [1, 14], is one such example of a modern programming language for concurrency. As depicted in Figure 1, Elixir programs are structured as a collection of *modules* that contain *functions*, the basic unit of code decomposition in the language. A module only exposes a subset of these functions to external invocations by defining them as *public*; these functions act as the only entry points to the functionality encapsulated by a module. Internally, the bodies of these public functions may then invoke other functions, which can either be the *public* ones already exposed or the *private* functions that can only be invoked from within the same module. For instance, Figure 1 depicts a module  $m$  which contains several public functions (*i.e.*,  $f_1, \dots, f_n$ ) and private functions (*i.e.*,  $g_1, \dots, g_j$ ). For example, the public function  $f_1$  delegates part of its computation by calling the private functions  $g_1$  and  $g_j$ , whereas the body of the public function  $f_n$  invokes the other public function  $f_1$  when executed. Internally, the body of the private function  $g_1$  calls the other private function  $g_2$  whereas the private function  $g_j$  can recursively call itself.

A prevalent Elixir design pattern is that of a server listening for client requests. For each request, the server spawns a (public) function to execute independently and act as a dedicated client handler: after the respective process IDs of the client and the spawned handler are made known to each other, a session of interaction commences between the two concurrent entities (via message-passing). For instance, in Figure 1, a handler process running public function  $f_1$  is assigned to the session with client  $client_1$  whereas




Figure 1: An Elixir module consisting of public and private functions, interacting with client processes

the request from  $\text{client}_k$  is assigned a dedicated handler running function  $f_n$ . Although traditional interface elements such as function parameters (used to instantiate the executing function body with values such as the client process ID) and the function return value (reporting the eventual outcome of handled request) are important, the messages exchanged between the two concurrent parties within a session are equally important for software correctness. More specifically, communication incompatibilities between the interacting parties could lead to various runtime errors. For example, if in a session a message is sent with an unexpected payload, it could cause the receiver’s subsequent computation depending on it to crash (*e.g.* multiplying by a string when a number should have been received instead). Also, if messages are exchanged in an incorrect order, they may cause *deadlocks* (*e.g.* two processes waiting forever for one another to send messages of a particular kind when a message of a different kind has been sent instead).

In many cases, the expected protocol of interactions within a session can be statically determined from the respective endpoint implementations, namely the function bodies; for simplicity, our discussion assumes that endpoint interaction protocols are dual, *e.g.*  $S_1$  and  $\bar{S}_1$  in Figure 1. Although Elixir provides mechanisms for specifying (and checking) the parameters and return values of a function within a module, it does *not* provide any support for describing (and verifying) the interaction protocol of a function in terms of its communication side-effects. To this end, in earlier work [32] we devised the tool<sup>1</sup> ElixirST, assisting module construction in two ways: (a) it allows module designers to formalise the session endpoint protocol as a session type, and ascribe it to a public function; (b) it implements a type-checker that verifies whether the body of a function respects the ascribed session type protocol specification.

**Contribution.** This paper validates the underlying type system on which the ElixirST type-checker is built. More concretely, in Section 3 we formalise the runtime semantics of the Elixir language fragment supported by ElixirST as a labelled transition system (LTS), modelling the execution of a spawned handler interacting with a client within a session. This operational semantics then allows us to prove *session fidelity* for the ElixirST type system in Section 4. In Section 2 we provide the necessary background on the existing session type system from [32] to make the paper self-contained.

<sup>1</sup>ElixirST is available on GitHub: <https://github.com/gertab/ElixirST>

## 2 Preliminaries

We introduce a core Elixir subset and review the main typing rules for the ElixirST type system [32].

### 2.1 The Actor Model

Elixir uses the actor concurrency model [1, 14]. It describes computation as a group of concurrent processes, called *actors*, which do *not* share any memory and interact exclusively via asynchronous messages. Every actor is identified via a unique process identifier (*pid*) which is used as the address when sending messages to a specific actor. Messages are communicated asynchronously, and stored in the mailbox of the addressee actor. An actor is the only entity that can fetch messages from its mailbox, using mechanisms such as pattern matching. Apart from sending and reading messages, an actor can also spawn other actors and obtain their fresh *pid* as a result; this *pid* can be communicated as a value to other actors via messaging.

### 2.2 Session Types

The ElixirST type system [32] assumes the standard expression types, including basic types, such as boolean, number, atom and pid, and inductively defined types, such as tuples ( $\{T_1, \dots, T_n\}$ ) and lists ( $[T]$ ); these already exist in the Elixir language and they are dynamically checked. It extends these with (binary) session types, which are used to statically check the message-passing interactions.

Expression types	$T ::= \text{boolean} \mid \text{number} \mid \text{atom} \mid \text{pid} \mid \{T_1, \dots, T_n\} \mid [T]$		
Session types	$S ::= \&\{?l_i(\tilde{T}_i).S_i\}_{i \in I}$	Branch	$\mid \text{rec } X.S$ Recursion
	$\mid \oplus\{!l_i(\tilde{T}_i).S_i\}_{i \in I}$	Choice	$\mid X$ Variable
	$\mid \text{end}$	Termination	

The *branching* construct,  $\&\{?l_i(\tilde{T}_i).S_i\}_{i \in I}$ , requires the code to be able to receive a message that is labelled by any one of the labels  $l_i$ , with the respective list of values of type  $\tilde{T}_i$  (where  $\tilde{T}$  stands for  $T^1, \dots, T^k$  for some  $k \geq 0$ ), and then adhere to the continuation session type  $S_i$ . The *choice* construct is its dual and describes the range and format of outputs the code is allowed to perform. In both cases, the labels need to be pairwise distinct. Recursive types are treated equi-recursively [27], and used interchangeably with their unfolded counterparts. For brevity, the symbols  $\&$  and  $\oplus$  are occasionally omitted for singleton options, e.g.,  $\oplus\{!l(\text{number}).S_1\}$  is written as  $!l(\text{number}).S_1$ ; similarly end may be omitted as well, e.g.,  $?l()$  stands for  $?l().\text{end}$ . The *dual* of a session type  $S$  is denoted as  $\bar{S}$ .

### 2.3 Elixir Syntax

Elixir programs are organised as modules, *i.e.*, `defmodule  $m$  do  $\tilde{P} \tilde{D}$  end`. Modules are defined by their name,  $m$ , and contain two sets of public  $\tilde{D}$  and private  $\tilde{P}$  functions, declared as sequences. Public functions, `def  $f(y, \tilde{x})$  do  $t$  end`, are defined by the **def** keyword, and can be called from any module. In contrast, private functions, `defp  $f(y, \tilde{x})$  do  $t$  end`, can only be called from within the defining module. Functions are defined by their name,  $f$ , and their body,  $t$ , and parametrised by a sequence of *distinct* variables,  $y, \tilde{x}$ , the length of which,  $|y, \tilde{x}|$ , is called the *arity*. As an extension to [32], the first parameter

<p>Module <math>M ::= \text{defmodule } m \text{ do } \tilde{P} \tilde{D} \text{ end}</math></p> <p>Public fun. <math>D ::= K \ B \ \text{def } f(y, \tilde{x}) \text{ do } t \text{ end}</math></p> <p>Private fun. <math>P ::= B \ \text{defp } f(y, \tilde{x}) \text{ do } t \text{ end}</math></p> <p>Type ann. <math>B ::= @\text{spec } f(\tilde{T}) :: T</math></p> <p>Session ann. <math>K ::= @\text{session } "X = S"</math>  <math>\quad \quad \quad   @\text{dual } "X"</math></p> <p>Expressions <math>e ::= w</math>  <math>\quad \quad \quad   \text{not } e \mid e_1 \diamond e_2</math>  <math>\quad \quad \quad   [e_1 \mid e_2] \mid \{e_1, \dots, e_n\}</math></p> <p>Operators <math>\diamond ::= &lt; \mid &gt; \mid &lt;= \mid &gt;= \mid ==</math>  <math>\quad \quad \quad   != \mid + \mid - \mid * \mid / \mid \text{and} \mid \text{or}</math></p>	<p>Basic val. <math>b ::= \text{boolean} \mid \text{number} \mid \text{atom} \mid \text{pid}</math>  <math>\quad \quad \quad   []</math></p> <p>Values <math>v ::= b \mid [v_1 \mid v_2] \mid \{v_1, \dots, v_n\}</math></p> <p>Identifiers <math>w ::= b \mid x</math></p> <p>Patterns <math>p ::= w \mid [w_1 \mid w_2] \mid \{w_1, \dots, w_n\}</math></p> <p>Terms <math>t ::= e</math>  <math>\quad \quad \quad   x = t_1; t_2</math>  <math>\quad \quad \quad   \text{send}(w, \{:\!l, e_1, \dots, e_n\})</math>  <math>\quad \quad \quad   \text{receive do}</math>  <math>\quad \quad \quad \quad (\{:\!l_i, p_i^1, \dots, p_i^n\} \rightarrow t_i)_{i \in I} \text{end}</math>  <math>\quad \quad \quad   f(w, e_1, \dots, e_n)</math>  <math>\quad \quad \quad   \text{case } e \text{ do } (p_i \rightarrow t_i)_{i \in I} \text{end}</math></p>
--	---

Figure 2: Elixir syntax

(y), is reserved for the *pid* of the dual process. Although a module may contain functions with the same name, their arity must be different.

In our formalisation, Elixir function parameters and return values are assigned a type using the `@spec` annotation,  $f(\tilde{T}) :: T$ , describing the parameter types,  $\tilde{T}$ , and the return type,  $T$ . This annotation is already used by Dialyzer for success typing [21]. In addition to this, we decorate public functions with session types, defined in Section 2.2, to describe their side-effect protocol. Public functions can be annotated directly using `@session "X = S"`, or indirectly using the dual session type, `@dual "X"`, where  $X = S$  is shorthand for `rec X.S`.

The body of a function consists of a term,  $t$ , which can take the form of an expression, a `let` statement, a `send` or `receive` construct, a `case` statement or a function call; see Figure 2. In the case of the `let` construct,  $x = t_1; t_2$ , the variable  $x$  is a *binder* for the variables in  $t_2$  acting as a placeholder for the value that the subterm  $t_1$  evaluates to. We write  $t_1; t_2$ , as *syntactic sugar* for  $x = t_1; t_2$  whenever  $x$  is not used in  $t_2$ . The *send* statement, `send(x, {:\!l, e1, ..., en})`, allows a process to send a message to the *pid* stored in the variable  $x$ , containing a message  $\{:\!l, e_1, \dots, e_n\}$ , where  $:\!l$  is the label. The *receive* construct, `receive do ( {:\!li, pi1, ..., pin } → ti )i ∈ I end`, allows a process to receive a message tagged with a label that matches one of the labels  $:\!l_i$  and a list of payloads that match the patterns  $p_i^1, \dots, p_i^n$ , branching to continue executing as  $t_i$ . Patterns,  $p$ , defined in Figure 2, can take the form of a variable, a basic value, a tuple or a list (e.g.  $[x \mid y]$ , where  $x$  is the head and  $y$  is the tail of the list). The remaining constructs are fairly standard. Variables in patterns  $p_i^1, \dots, p_i^n$  employed by the `receive` and `case` statements are binders for the respective continuation branches  $t_i$ . We assume standard notions of open (i.e.,  $\text{fv}(t) \neq \emptyset$ ) and closed (i.e.,  $\text{fv}(t) = \emptyset$ ) terms and work up to alpha-conversion of bound variables.

## 2.4 Type System

The session type system from [32] statically verifies that public functions within a module observe the communication protocols ascribed to them. It uses three environments:

$$\begin{array}{ll}
 \text{Variable binding env.} & \Gamma ::= \emptyset \mid \Gamma, x : T \\
 \text{Session typing env.} & \Delta ::= \emptyset \mid \Delta, f/n : S \\
 \text{Function inf. env.} & \Sigma ::= \emptyset \mid \Sigma, f/n : \left\{ \begin{array}{l} \text{params} = \tilde{x}, \text{param\_types} = \tilde{T}, \\ \text{body} = t, \text{return\_type} = T, \text{dual} = y \end{array} \right\}
 \end{array}$$


The *variable binding* environment,  $\Gamma$ , maps (data) variables to basic types ( $x : T$ ). We write  $\Gamma, x : T$  to extend  $\Gamma$  with the new mapping  $x : T$ , where  $x \notin \mathbf{dom}(\Gamma)$ . The *session typing* environment,  $\Delta$ , maps function names and arity pairs to their session type ( $f/n : S$ ). If a function  $f/n$  has a *known* session type, then it can be found in  $\Delta$ , *i.e.*,  $\Delta(f/n) = S$ . Each module has a static *function information* environment,  $\Sigma$ , that holds information related to the function definitions. For a function  $f$ , with arity  $n$ ,  $\Sigma(f/n)$  returns the tail list of parameters (`params`) and their types (`param_types`), the function's body (`body`), and its return type (`return_type`). In contrast to the original type system from [32],  $\Sigma(f/n)$  also returns the variable name that represents the interacting process' *pid*, *i.e.*, the option `dual`. We assume that *function information* environments,  $\Sigma$ , are *well-formed*, meaning that all functions mapped ( $f/n \in \mathbf{dom}(\Sigma)$ ) observe the following condition requiring that the body of function  $f/n$  is *closed*:

$$\mathbf{fv}(\Sigma(f/n).\text{body}) \setminus (\Sigma(f/n).\text{params} \cup \Sigma(f/n).\text{dual}) = \emptyset$$

Session typechecking is initiated by analysing an Elixir module, rule [TMODULE]. A module is type-checked by inspecting each of its public functions, ascertaining that they correspond and fully consume the session types ascribed to them. The rule uses three helper functions. The function **functions**( $\tilde{D}$ ) returns a list of all function names (and arity) of the public functions ( $\tilde{D}$ ) to be checked individually. The function **sessions**( $\tilde{D}$ ) obtains a mapping of all the public functions to their expected session types stored in  $\Delta$ . This ensures that when a function  $f$  with arity  $n$  executes, it adheres to the session type associated with it using either the `@session` or `@dual` annotations. The helper function **details** populates the *function information* environment ( $\Sigma$ ) with details about all the *public* ( $\tilde{D}$ ) and *private* functions ( $\tilde{P}$ ) within the module.

$$\begin{array}{c}
 \Delta = \mathbf{sessions}(\tilde{D}) \quad \Sigma = \mathbf{details}(\tilde{P}\tilde{D}) \\
 \forall f/n \in \mathbf{functions}(\tilde{D}) \cdot \left\{ \begin{array}{l} \Delta(f/n) = S \quad \Sigma(f/n) = \Omega \\ \Omega.\text{params} = \tilde{x} \quad \Omega.\text{param\_types} = \tilde{T} \\ \Omega.\text{body} = t \quad \Omega.\text{return\_type} = T \quad \Omega.\text{dual} = y \\ \Delta \cdot (y : \text{pid}, \tilde{x} : \tilde{T}) \vdash_{\Sigma}^y S \triangleright t : T \triangleleft \text{end} \end{array} \right. \\
 \text{[TMODULE]} \frac{}{\vdash \text{defmodule } m \text{ do } \tilde{P} \tilde{D} \text{ end}}
 \end{array}$$

For every public function  $f/n$  in **functions**( $\tilde{D}$ ), [TMODULE] checks that its body adheres to its session type using the **highlighted term typing judgement** detailed below:



This judgement states that “the term  $t$  can produce a value of type  $T$  after following an interaction protocol starting from the initial session type  $S$  up to the residual session type  $S'$ , while interacting with a dual process with  $pid$  identifier  $w$ . This typing is valid under some *session typing* environment  $\Delta$ , *variable binding* environment  $\Gamma$  and *function information* environment  $\Sigma$ .” Since the *function information* environment  $\Sigma$  is static for the whole module (and by extension, for all sub-terms), it is left implicit in the term typing rules. We consider four main rules, and relegate the rest to [Appendix B.1](#).

$$[\text{TBRANCH}] \frac{\forall i \in I \quad \forall j \in 1..n \quad \vdash_{\text{pat}}^w p_i^j : T_i^j \triangleright \Gamma_i^j \quad \Delta \cdot (\Gamma, \Gamma_i^1, \dots, \Gamma_i^n) \vdash^w S_i \triangleright t_i : T \triangleleft S'_i}{\Delta \cdot \Gamma \vdash^w \&\{?l_i(\tilde{T}_i).S_i\}_{i \in I} \triangleright \text{receive do } (\{ :l_i, \tilde{p}_i \} \rightarrow t_i)_{i \in I} \text{end} : T \triangleleft S'}$$

The `receive` construct is typechecked using the [TBRANCH] rule. It expects an (external) branching session type  $\&\{\dots\}$ , where each branch in the session type must match with a corresponding branch in the `receive` construct, where *both* the labels ( $l_i$ ) and payload types ( $\tilde{T}_i$ ) correspond. The types within each `receive` branch are computed using the pattern typing judgement,  $\vdash_{\text{pat}}^w p : T \triangleright \Gamma$ , which assigns types to variables present in patterns (see [Appendix B.3](#)). Each `receive` branch is then checked w.r.t. the common type  $T$  and a common residual session type  $S'$ .

$$[\text{TCHOICE}] \frac{\exists i \in I \quad l = l_i \quad \forall j \in 1..n \quad \Gamma \vdash_{\text{exp}} e_j : T_i^j}{\Delta \cdot \Gamma \vdash^w \oplus\{!l_i(\tilde{T}_i).S_i\}_{i \in I} \triangleright \text{send}(w, \{ :l, e_1, \dots, e_n \}) : \{\text{atom}, T_i^1, \dots, T_i^n\} \triangleleft S_i}$$

The rule [TCHOICE] typechecks the sending of messages. This rule requires an internal choice session type  $\oplus\{\dots\}$ , where the label tagging the message to be sent must match with one of the labels ( $l_i$ ) offered by the session choice. The message payloads must also match with the corresponding types associated with the label ( $\tilde{T}_i$  of  $l_i$ ) stated via the expression typing judgement  $\Gamma \vdash_{\text{exp}} e : T$  (see [Appendix B.2](#)). The typing rule also checks the *pid* of the addressee of the `send` statement which must match with the dual *pid* ( $w$ ) states in the judgment itself to ensure that messages are only sent to the correct addressee.

$$[\text{TRECKNOWNCALL}] \frac{\Delta(f/n) = S \quad \forall i \in 2..n \cdot \{\Gamma \vdash_{\text{exp}} e_i : T_i\} \quad \Sigma(f/n) = \Omega \quad \Omega.\text{return\_type} = T \quad \Omega.\text{param\_types} = \tilde{T}}{\Delta \cdot \Gamma \vdash^w S \triangleright f(w, e_2, \dots, e_n) : T \triangleleft \text{end}}$$

Since public functions are decorated with a session type explicitly using the `@session` (or `@dual`) annotation, they are listed in  $\text{dom}(\Delta)$ . Calls to public functions are typechecked using the [TRECKNOWNCALL] rule, which verifies that the expected initial session type is equivalent to the function’s *known* session type ( $S$ ) obtained from the *session typing* environment, *i.e.*,  $\Delta(f/n) = S$ . Without typechecking the function’s body, which is done in rule [TMODULE], this rule ensures that the parameters have the correct types (using the expression typing rules). From the check performed in rule [TMODULE], it can also safely assume that this session type  $S$  is fully consumed, thus the residual type becomes `end`. Rule [TRECKNOWNCALL] also ensures that the *pid* ( $w$ ) is preserved during a function call, by requiring it to be passed as a parameter and comparing it to the expected dual *pid* (*i.e.*,  $\Delta \cdot \Gamma \vdash^w S \triangleright f(w, \dots) : T \triangleleft \text{end}$ ).

$$\begin{array}{c}
\Sigma(f/n) = \Omega \quad f/n \notin \mathbf{dom}(\Delta) \quad \Omega.\mathbf{dual} = y \\
\Omega.\mathbf{params} = \tilde{x} \quad \Omega.\mathbf{param\_type} = \tilde{T} \quad \Omega.\mathbf{body} = t \quad \Omega.\mathbf{return\_type} = T \\
[\mathbf{TRECUNKNOWNCALL}] \frac{(\Delta, f/n : S) \cdot (\Gamma, y : \mathbf{pid}, \tilde{x} : \tilde{T}) \vdash^y S \triangleright t : T \triangleleft S' \quad \forall i \in 2..n \cdot \{\Gamma \vdash_{\text{exp}} e_i : T_i\}}{\Delta \cdot \Gamma \vdash^w S \triangleright f(w, e_2, \dots, e_n) : T \triangleleft S'}
\end{array}$$

Contrastingly, a call to a (private) function,  $f/n$ , with an *unknown* session type associated to it is type-checked using the  $[\mathbf{TRECUNKNOWNCALL}]$  rule. As in the other rule, it ensures that parameters have the correct types ( $\Gamma \vdash_{\text{exp}} e_i : T_i$ ). However, it also analyses the function's body  $t$  (obtained from  $\Sigma$ ) with respect to the session type  $S$  inherited from the initial session type of the call. Furthermore, this session type is appended to the *session typing* environment  $\Delta$  for future reference, *i.e.*,  $\Delta' = (\Delta, f/n : S)$  which allows it to handle recursive calls to itself; should the function be called again, rule  $[\mathbf{TRECUNKNOWNCALL}]$  is used thus bypassing the need to re-analyse its body.

## 2.5 Elixir System

The ElixirST provides a bespoke spawning function called `session/4` which allows the initiation of two concurrent processes executing in tandem as part of a session. This `session/4` function takes two pairs of arguments: two references of function names (that will be spawned), along with their list of arguments. Its participant creation flow is shown in [Figure 3](#). Initially the actor (pre-server) is spawned, passing its *pid* ( $\iota_{\text{server}}$ ) to the second spawned actor (pre-client). Then, pre-client relays back its *pid* ( $\iota_{\text{client}}$ ) to pre-server. In this way, both actors participating in a session become aware of each other's *pids*. From this point onwards, the two actors execute their respective function to behave as the participants in the binary session; the first argument of each running function is initiated to the respective *pid* of the other participant. [Figure 3](#) shows that the server process executes the body  $t$ , where it has access to the mailbox  $\mathcal{M}$ . As it executes, messages may be sent or received (shown by the action  $\alpha$ ) and stored in the mailbox  $\mathcal{M}'$ . The specific working of these transitions is explained in the following section.




Figure 3: Spawning two processes (green boxes represent *spawned* concurrent processes)

## 3 Operational Semantics

We describe the operational semantics of the Elixir language subset of [Figure 2](#) as a *labelled transition system* (LTS) [18] describing how a handler process within a session executes while interacting with the session client, as outlined in [Figure 1](#). The transitions  $t \xrightarrow{\alpha} t'$  describes the fact that a handler process in state  $t$  performs an execution step to transition to the new state  $t'$  while producing action  $\alpha$  as a side-effect. External actions are visible by, and bear an effect on the client, whereas internal actions do not. In our case, an action  $\alpha$  can take the following forms:

$$\begin{array}{l}
\alpha \in \text{ACT} ::= \iota! \{ :l, \tilde{v} \} \quad \text{Output message to } \iota \text{ tagged as } :l \text{ with payload } \tilde{v} \\
\quad | ? \{ :l, \tilde{v} \} \quad \text{Input message tagged as } :l \text{ with payload } \tilde{v} \\
\quad | f/n \quad \text{Call function } f \text{ with arity } n \\
\quad | \tau \quad \text{Internal reduction step}
\end{array}
\left. \vphantom{\begin{array}{l} \alpha \in \text{ACT} \\ \tau \end{array}} \right\} \begin{array}{l} \text{external action} \\ \text{internal action} \end{array}$$

Both output and input actions constitute external actions that affect either party in a session; the type system from [Section 2.4](#) disciplines these external actions. Internal actions, include *silent* transition ( $\tau$ ) and function calls ( $f/n$ ); although the latter may be formalised as a silent action, the decoration facilitates our technical development. We note that, function calls can only transition subject to a well-formed *function information* environment ( $\Sigma$ ), which contains details about all the functions available in the module. Since  $\Sigma$  remains static during transitions, we leave it implicit in the transitions rules.

The transitions are defined by the *term* transition rules listed in [Figure 4](#). Rules [RLET<sub>1</sub>] and [RLET<sub>2</sub>] deal with the evaluation of a *let* statement,  $x = t_1; t_2$  modelling a *call-by-value* semantic, where the first term  $t_1$  has to transition fully to a value before being substituted for  $x$  in  $t_2$  denoted as  $[v/x]$  (or  $[v_1, v_2/x_1, x_2]$  for multiple substitutions). The *send* statement,  $\text{send}(\iota, \{ :l, e_1, \dots, e_n \})$ , evaluates by first reducing each

$$\boxed{t \xrightarrow[\Sigma]{\alpha} t'} \quad \begin{array}{l} \text{[RLET}_1\text{]} \frac{t_1 \xrightarrow{\alpha} t'_1}{x = t_1; t_2 \xrightarrow{\alpha} x = t'_1; t_2} \quad \text{[RLET}_2\text{]} \frac{}{x = v; t \xrightarrow{\tau} t[v/x]} \\ \\ \text{[RCHOICE}_1\text{]} \frac{e_k \rightarrow e'_k}{\text{send}(\iota, \{ :l, v_1, \dots, v_{k-1}, e_k, \dots, e_n \}) \xrightarrow{\tau} \text{send}(\iota, \{ :l, v_1, \dots, v_{k-1}, e'_k, \dots, e_n \})} \\ \\ \text{[RCHOICE}_2\text{]} \frac{}{\text{send}(\iota, \{ :l, v_1, \dots, v_n \}) \xrightarrow{\iota! \{ :l, v_1, \dots, v_n \}} \{ :l, v_1, \dots, v_n \}} \\ \\ \text{[RBRANCH]} \frac{\exists j \in I \quad 1_j = 1 \quad \text{match}(\tilde{p}_j, v_1, \dots, v_n) = \sigma}{\text{receive do } (\{ :l_i, \tilde{p}_i \} \rightarrow t_i)_{i \in I} \text{end} \xrightarrow{? \{ :l, v_1, \dots, v_n \}} t_j \sigma} \\ \\ \text{[RCALL}_1\text{]} \frac{e_k \rightarrow e'_k}{f(v_1, \dots, v_{k-1}, e_k, \dots, e_n) \xrightarrow{\tau} f(v_1, \dots, v_{k-1}, e'_k, \dots, e_n)} \\ \\ \text{[RCALL}_2\text{]} \frac{\Sigma(f/n) = \Omega \quad \Omega.\text{body} = t \quad \Omega.\text{params} = x_2, \dots, x_n \quad \Omega.\text{dual} = y}{f(\iota, v_2, \dots, v_n) \xrightarrow{f/n} t[\iota/y][v_2, \dots, v_n/x_2, \dots, x_n]} \\ \\ \text{[RCASE}_1\text{]} \frac{e \rightarrow e'}{\text{case } e \text{ do } (p_i \rightarrow t_i)_{i \in I} \text{end} \xrightarrow{\tau} \text{case } e' \text{ do } (p_i \rightarrow t_i)_{i \in I} \text{end}} \\ \\ \text{[RCASE}_2\text{]} \frac{\exists j \in I \quad \text{match}(p_j, v) = \sigma}{\text{case } v \text{ do } (p_i \rightarrow t_i)_{i \in I} \text{end} \xrightarrow{\tau} t_j \sigma} \quad \text{[REXPRESSION]} \frac{e \rightarrow e'}{e \xrightarrow{\tau} e'} \end{array}$$

Figure 4: Term transition semantic rules



$$\boxed{e \rightarrow e'} \quad
\begin{array}{c}
\text{[REOPERATION}_1\text{]} \frac{e_1 \rightarrow e'_1}{e_1 \diamond e_2 \rightarrow e'_1 \diamond e_2} \quad
\text{[REOPERATION}_2\text{]} \frac{e_2 \rightarrow e'_2}{v_1 \diamond e_2 \rightarrow v_1 \diamond e'_2} \\
\text{[REOPERATION}_3\text{]} \frac{v = v_1 \diamond v_2}{v_1 \diamond v_2 \rightarrow v} \quad
\text{[RENOT}_1\text{]} \frac{e \rightarrow e'}{\text{not } e \rightarrow e'} \quad
\text{[RENOT}_2\text{]} \frac{v' = \neg v}{\text{not } v \rightarrow v'} \\
\text{[RELIST}_1\text{]} \frac{e_1 \rightarrow e'_1}{[e_1 \mid e_2] \rightarrow [e'_1 \mid e_2]} \quad
\text{[RELIST}_2\text{]} \frac{e_2 \rightarrow e'_2}{[v_1 \mid e_2] \rightarrow [v_1 \mid e'_2]} \\
\text{[RETUPLE]} \frac{e_k \rightarrow e'_k}{\{v_1, \dots, v_{k-1}, e_k, \dots, e_n\} \rightarrow \{v_1, \dots, v_{k-1}, e'_k, \dots, e_n\}}
\end{array}$$

Figure 5: Expression reduction rules

part of the message to a value from left to right. This is carried out via rule  $[\text{RCHOICE}_1]$  which produces no observable side-effects. When the whole message is reduced to a tuple of values  $\{:\mathbb{1}, v_1, \dots, v_n\}$ , rule  $[\text{RCHOICE}_2]$  performs the actual message sending operation. This transition produces an action  $u! \{:\mathbb{1}, v_1, \dots, v_n\}$ , where the message  $\{:\mathbb{1}, v_1, \dots, v_n\}$  is sent to the interacting process, which has a *pid* value of  $u$ . The operational semantics of the *receive* construct,  $\text{receive do } (\{:\mathbb{1}_i, \tilde{p}_i\} \rightarrow t_i)_{i \in I} \text{end}$ , is defined by rule  $[\text{RBRANCH}]$ . When a message is received (*i.e.*,  $\alpha = ?\{:\mathbb{1}, \tilde{v}\}$ ), it is matched with a valid branch from the *receive* construct, using the label  $:\mathbb{1}$ . Should one of the labels match ( $\exists j \in I$  such that  $:\mathbb{1}_j = :\mathbb{1}$ ), the payload of the message ( $\tilde{v}$ ) is compared to the corresponding patterns in the selected branch ( $\tilde{p}_j$ ) using  $\text{match}(\tilde{p}_j, \tilde{v})$ . If the values match with the pattern, the  $\text{match}$  function ([Definition 3.1](#)) produces the substitutions  $\sigma$ , mapping the matched variables in the pattern  $\tilde{p}_j$  to values from  $\tilde{v}$ . This substitution  $\sigma$  is then used to instantiate the free variables in continuation branch  $t_j$ .

**Definition 3.1** (*Pattern Matching*). The  $\text{match}$  function pairs patterns with a corresponding value, resulting in a sequence of substitutions (called  $\sigma$ ), *e.g.*,  $\text{match}(p, v) = [v_1/x_1][v_2/x_2] = [v_1, v_2/x_1, x_2]$ . Note that, a sequence of  $\text{match}$  outputs are combined together, where the empty substitutions (*i.e.*,  $[\ ]$ ) are ignored. The  $\text{match}$  function builds a meta-list of substitutions, which is a different form of lists defined by the Elixir syntax in [Figure 2](#).

$$\begin{array}{l}
\text{match}(\tilde{p}, \tilde{v}) \stackrel{\text{def}}{=} \text{match}(p_1, v_1), \dots, \text{match}(p_n, v_n) \\
\text{where } \tilde{p} = p_1, \dots, p_n \text{ and } \tilde{v} = v_1, \dots, v_n \\
\text{match}(p, v) \stackrel{\text{def}}{=} \begin{cases} [\ ] & p = b, v = b \text{ and } p = v \\ [v/x] & p = x \\ \text{match}(w_1, v_1), \text{match}(w_2, v_2) & p = [w_1 \mid w_2], v = [v_1 \mid v_2] \\ \text{match}(w_1, v_1), \dots, \text{match}(w_n, v_n) & p = \{w_1, \dots, w_n\} \text{ and} \\ & v = \{v_1, \dots, v_n\} \end{cases} \quad \blacksquare
\end{array}$$

**Example 3.1.** For the pattern  $p_1 = \{x, 2, y\}$  and the value tuple  $v_1 = \{8, 2, \text{true}\}$ ,  $\text{match}(p_1, v_1) = \sigma$  where  $\sigma = [8/x][\text{true}/y]$  (written also as  $\sigma = [8, \text{true}/x, y]$ ). However for pattern  $p_2 = \{x, 2, \text{false}\}$ , the operation

**match**( $p_2, v_1$ ) fails, since  $p_2$  expects a *false* value as the third element, but finds a *true* value instead. ■

Using rule  $[\text{RCALL}_1]$  from [Figure 4](#), a function call is evaluated by first reducing all of its parameters to a value, using the expression reduction rules in [Figure 5](#); again this models a call-by-value semantics. Once all arguments have been fully reduced,  $[\text{RCALL}_2]$ , the implicit environment  $\Sigma$  is queried for function  $f$  with arity  $n$  to fetch the function’s parameter names and body. This results in a transition to the function body with its parameters instantiated accordingly,  $t [t/y] [v_2, \dots, v_n/x_2, \dots, x_n]$ , decorated by the function name, *i.e.*,  $\alpha = f/n$ . Along the same lines a case construct first reduces the expression which is being matched using rule  $[\text{RCASE}_1]$ . Then, rule  $[\text{RCASE}_2]$  matches the value with the correct branch, using the **match** function, akin to  $[\text{RBRANCH}]$ . Whenever a term consists solely of an expression, it silently reduces using  $[\text{REXPRESSION}]$  using the expression reduction rules  $e \rightarrow e'$  of [Figure 5](#). These are fairly standard.

## 4 Session Fidelity

We validate the static properties imposed by the ElixirST type system [\[32\]](#), overviewed in [Section 2.4](#), by establishing a relation with the runtime behaviour of a typechecked Elixir program, using the transition semantics defined in [Section 3](#). Broadly, we establish a form of *type preservation*, which states that if a well-typed term transitions, the resulting term then remains well-typed [\[27\]](#). However, our notion of type preservation, needs to be stronger to also take into account (i) the side-effects produced by the execution; and (ii) the progression of the execution with respect to protocol expressed as a session type. Following the long-standing tradition in the session type community, these two aspects are captured by the refined preservation property called *session fidelity*. This property ensures that: (i) the communication action produced as a result of the execution of the typed process is one of the actions allowed by the current stage of the protocol; and that (ii) the resultant process following the transition is still well-typed w.r.t. the remaining part of the protocol that is still outstanding.

Before embarking on the proof for session fidelity, we prove an auxiliary proposition that acts as a sanity check for our operational semantics. We note that the operational semantics of [Section 3](#) assumes that only *closed* programs are executed; an *open* program (*i.e.*, a program containing free variables) is seen as an incomplete program that cannot execute correctly due to missing information. To this end, [Proposition 1](#) ensures that a closed term *remains closed* even after transitioning.

**Proposition 1** (Closed Term). *If  $\text{fv}(t) = \emptyset$  and  $t \xrightarrow{\alpha} t'$ , then  $\text{fv}(t') = \emptyset$*

*Proof.* By induction on the structure of  $t$ . □

The statement of the session fidelity property relies on the definition of a partial function called **after** ([Definition 4.1](#)), which takes a session type and an action as arguments and returns another session type as a result. This function serves two purposes: (a) the function **after**( $S, \alpha$ ) is only defined for actions  $\alpha$  that are (immediately) permitted by the protocol  $S$ , which allows us to verify whether a term transition step violated a protocol or not; and (b) since  $S$  describes the current stage of the protocol to be followed, we need a way to evolve this protocol to the next stage should  $\alpha$  be a permitted action, and this is precisely  $S'$ , the continuation session type returned where **after**( $S, \alpha$ ) =  $S'$ .

**Definition 4.1** (*After Function*). The **after** function is partial function defined for the following cases:

$$\begin{aligned} \mathbf{after}(S, \tau) &\stackrel{\text{def}}{=} S \\ \mathbf{after}(S, f/n) &\stackrel{\text{def}}{=} S \\ \mathbf{after}(\oplus \{!1_i(\tilde{T}_i).S_i\}_{i \in I}, \iota! \{1_j, \tilde{v}\}) &\stackrel{\text{def}}{=} S_j \quad \text{where } j \in I \\ \mathbf{after}(\& \{?1_i(\tilde{T}_i).S_i\}_{i \in I}, ? \{1_j, \tilde{v}\}) &\stackrel{\text{def}}{=} S_j \quad \text{where } j \in I \end{aligned}$$

This function is undefined for all other cases. The **after** function is overloaded to range over *session typing* environments ( $\Delta$ ) in order to compute a new *session typing* environment given some action  $\alpha$  and session type  $S$ :

$$\begin{aligned} \mathbf{after}(\Delta, f/n, S) &\stackrel{\text{def}}{=} \Delta, f/n : S \\ \mathbf{after}(\Delta, \alpha, S) &\stackrel{\text{def}}{=} \Delta \quad \text{if } \alpha \neq f/n \end{aligned}$$

Intuitively, when the action produced by the transition is  $f/n$ , the *session typing* environment is extended by the new mapping  $f/n : S$ . For all other actions, the *session typing* environment remains unchanged. ■

Recall that module typechecking using rule [TMODULE] entails typechecking the bodies of all the public functions w.r.t. their ascribed session type,  $\Delta \cdot (y : \text{pid}, \tilde{x} : \tilde{T}) \vdash_{\Sigma}^y S \triangleright t : T \triangleleft S'$  (where  $S' = \text{end}$  for this specific case). At runtime, a spawned client handler process in a session starts running the function body term  $t$  where the parameter variables  $y, \tilde{x}$  are instantiated with the PID of the client, say  $\iota$ , and the function parameter values, say  $\tilde{v}$ , respectively,  $t[\iota/y][\tilde{v}/\tilde{x}]$ , as modelled in rule [RCALL<sub>2</sub>] from Figure 4. The instantiated function body is thus closed and can be typed w.r.t. an empty variable binding environment,  $\Gamma = \emptyset$ . Session fidelity thus states that if a closed term  $t$  is well-typed, *i.e.*,

$$\Delta \cdot \emptyset \vdash^w S \triangleright t : T \triangleleft S' \quad (1)$$

(where  $S$  and  $S'$  are initial and residual session types, respectively, and  $T$  is the basic expression type) and this term  $t$  transitions to a new term  $t'$  with action  $\alpha$ , *i.e.*,

$$t \xrightarrow{\alpha} t' \quad (2)$$


the new term  $t'$  remains well-typed, *i.e.*,

$$\Delta' \cdot \emptyset \vdash^w S'' \triangleright t' : T \triangleleft S' \quad (3)$$

where the evolved  $S''$  and  $\Delta'$  are computed as  $\mathbf{after}(S, \alpha) = S''$  and  $\mathbf{after}(\Delta, \alpha, S) = \Delta'$ . This ensures that the base type of the term is preserved (note the constant type  $T$  in eqs. (1) and (3)). Furthermore, it ascertains that the term  $t$  follows an interaction protocol starting from the initial session type  $S$  up to the residual session type  $S'$  (eq. (1)), since the updated session type  $S''$  is defined for  $\mathbf{after}(S, \alpha)$ .

**Theorem 2** (Session Fidelity). *If  $\Delta \cdot \emptyset \vdash_{\Sigma}^w S \triangleright t : T \triangleleft S'$  and  $t \xrightarrow{\alpha}_{\Sigma} t'$ , then there exists some  $S''$  and  $\Delta'$ , such that  $\Delta' \cdot \emptyset \vdash_{\Sigma}^w S'' \triangleright t' : T \triangleleft S'$  for  $\mathbf{after}(S, \alpha) = S''$  and  $\mathbf{after}(\Delta, \alpha, S) = \Delta'$*

*Proof.* By induction on the typing derivation  $\Delta \cdot \emptyset \vdash_{\Sigma}^w S \triangleright t : T \triangleleft S'$ . □

Figure 6: Repeated applications of *session fidelity*

As shown in [Figure 6](#), by repeatedly applying [Theorem 2](#), we can therefore conclude that all the (external) actions generated as a result of a computation (*i.e.*, sequence of transition steps) must all be actions that follow the protocol described by the initial session type. Since public functions are always typed with a residual session type `end`, certain executions could reach the case where the outstanding session is updated to `end` as well, *i.e.*,  $\mathbf{after}(S_n, \alpha_n) = \mathbf{end}$ . In such a case, we are guaranteed that the term will not produce further side-effects, as in the case of [Figure 6](#) where the term is reduced all the way down to some value,  $v$ .

**Example 4.1.** We consider a concrete example to show the importance of session fidelity. The function called `pinger/1` is able to send `ping` and receive `pong` repeatedly.

```


1 @session "X = !ping().?pong().X"
2 def pinger(pid) do
3   x = send(pid, {:ping})
4
5   receive do
6     {:pong} -> IO.puts("Received pong.")
7   end
8   pinger(pid)
9 end

```

This function adheres to the following protocol:

$$X = !\mathit{ping}().?\mathit{pong}().X$$

A process evaluating the function `pinger` executes by first sending a message containing a `ping` label to the interacting processes' `pid` ( $t_{pong}$ ), as shown below.



As the process evaluates, the initial term  $t$  transitions to  $t'$ , where it sends a message as a side-effect. This side-effect is denoted as an action  $\alpha$ , where  $\alpha = t_{pong} ! \{\mathit{ping}\}$ . By the [After Function Definition](#),  $X$  evolves to a new session type  $X'$ :

$$X' = \mathbf{after}(!\mathit{ping}().?\mathit{pong}().X, \alpha) = ?\mathit{pong}().X$$

For  $t'$  to remain well-typed, it must now match with the evolved session type  $X'$ , where it has to be able to receive a message labelled `pong`, before recursing. Although the process keeps executing

indefinitely, by the session fidelity property, we know that each step of execution will be in line with the original protocol. ■

## 5 Related Work

In this section, we compare ElixirST with other type systems and implementations.

**Type Systems for Elixir** Cassola *et al.* [4, 5] presented a gradual type system for Elixir. It statically typechecks the functional part of Elixir modules, using a gradual approach, where some terms may be left with an unknown expression type. In contrast to ElixirST, Cassola *et al.* analyse directly the unexpanded Elixir code which results in more explicit typechecking rules. Also, they focus on the static type system without formulating the operational semantics.

Another static type-checker for Elixir is *Gradient* [8]. It is a wrapper for its Erlang counterpart tool and takes a similar approach to [5], where gradual types are used. Another project, *TypeCheck* [35], adds dynamic type validations to Elixir programs. *TypeCheck* performs runtime typechecking by wrapping checks around existing functions. *Gradient* and *TypeCheck* are provided as an implementation only, without any formal analysis. In contrast to ElixirST, the discussed type-checkers [5, 8, 35] analyse the sequential part of the Elixir language omitting any checks related to message-passing between processes.

Some implementations aim to check issues related to message-passing. Harrison [11] statically checks Core Erlang for such issues. For instance, it detects orphan messages (*i.e.*, messages that will never be received) and unreachable receive branches. Harrison [12] extends [11] to add analyse Erlang/OTP behaviours (*e.g.*, `gen_server`, which structures processes in a hierarchical manner) by injecting runtime checks in the code. Compared to our work, [11, 12] perform automatic analysis of the implementation, however they do not verify communication with respect to a general protocol (*e.g.*, session types).

Another type system for Erlang was presented Svensson *et al.* [31]. Their body of work covers a larger subset of Erlang to what would be its equivalent in Elixir covered by our work. Moreover, its multi-tiered semantics captures an LTS defined over systems of concurrent actors. Although we opted for a smaller subset, we go beyond the pattern matching described by Svensson *et al.* since we perform a degree of typechecking for base types (*e.g.* in the premise of `[TBRANCH]`).

**Session Type Systems.** Closest to our work is [23], where Mostrous and Vasconcelos introduced session types to a fragment of Core Erlang, a dynamically typed language linked to Elixir. Their type system tags each message exchanged with a unique reference. This allows multiple sessions to coexist, since different messages could be matched to the corresponding session, using correlation sets. Mostrous and Vasconcelos take a more theoretic approach, so there is no implementation for [23]. Their type system guarantees *session fidelity* by inspecting the processes' mailboxes where, at termination, no messages should be left unprocessed in their mailboxes. Our work takes a more limited but pragmatic approach, where we introduce session types for functions within a module. Furthermore, we offer additional features, including variable binding (*e.g.*, in `let` statements), expressions (*e.g.*, addition operation), inductive types (*e.g.*, tuples and lists), infinite computation via recursion and explicit protocol definition.

A session-based runtime monitoring tool for python was initially presented by Neykova and Yoshida [24, 25]. They use the Scribble [15] language to write *multiparty* session type (MPST) [16] protocols, which are then used to monitor the processes' actions. Different processes are ascribed a

role (defined in the MPST protocol) using function decorators (akin to our function annotations). Similar to [24, 25], Fowler [9] presented an MPST implementation for Erlang. This implementation uses Erlang/OTP behaviours (*e.g.*, `gen_server`), which take into account Erlang’s *let it crash* philosophy, where processes may fail while executing. In contrast, although our work accepts a more limited language, ElixirST provides static guarantees where issues are flagged at pre-deployment stages, rather than flagging them at runtime.

Scalas and Yoshida [29] applied binary session types to the Scala language, where session types are abstracted as Scala classes. Session fidelity is ensured using Scala’s compiler, which complains if an implementation does not follow its ascribed protocol. Linearity checks are performed at runtime, which ensure that an implementation fully exhausts its protocol exactly once. Bartolo Burlò *et al.* [3] extended the aforementioned work [29], to monitor one side of an interaction statically and the other side dynamically using runtime monitors.

Harvey *et al.* [13] presented a new actor-based language, called EnsembleS, which offers session types as a native feature of the language. EnsembleS statically verifies implementations with respect to session types, while still allowing for adaptation of *new* actors at runtime, given that the actors obey a known protocol. Thus, actors can be terminated and discovered at runtime, while still maintaining static correctness.

There have been several binary [17, 19] and multiparty [6, 20] session type implementations for Rust. These implementations exploit Rust’s affine type system to guarantee that channels mirror the actions prescribed by a session type. Padovani [26] created a binary session type library for OCaml to provide static communication guarantees. This project was extended [22] to include dynamic contract monitoring which flags violations at runtime. The approaches used in the Rust and OCaml implementations rely heavily on type-level features of the language, which do not readily translate to the dynamically typed Elixir language.

## 6 Conclusion

In this work we established a correspondence between the ElixirST type system [32] and the runtime behaviour of a client handler running an Elixir module function that has been typechecked w.r.t. its session type protocol. In particular, we showed that this session-based type system observes the standard *session fidelity* property, meaning that processes executing a typed function *always* follow their ascribed protocols at runtime. This property provides the necessary underlying guarantees to attain various forms of communication safety, whereby should two processes following mutually compatible protocols (*e.g.*  $S$  and its dual  $\bar{S}$ ), they avoid certain communication errors (*e.g.*, a send statement without a corresponding receive construct). An extended version of this work can be found in the technical report [33].

**Future work.** There are a number of avenues we intend to pursue. One line of investigation is the augmentation of protocols that talk about multiple entry points to a module perhaps from the point of view of a client that is engaged in multiple sessions at one time, possibly involving multiple modules. The obvious starting points to look at here are the well-established notions of multiparty session types [16, 30] or the body of work on intuitionistic session types organising processes hierarchically [2, 28]. Another natural extension to our work would be to augment our session type protocol in such a way to account for process failure and supervisors, which is a core part of the Elixir programming model. For this, we will look at previous work on session type extensions that account for failure [13]. Finally, we also plan to augment our session typed protocols to account for resource usage and cost, along the lines of [7, 10].

**Acknowledgements.** We thank Simon Fowler and Matthew Alan Le Brun for their encouragement and suggestions for improvement on earlier versions of the work. This work has been supported by the MoVeMnt project (No: 217987) of the Icelandic Research Fund and the BehAPI project funded by the EU H2020 RISE of the Marie Skłodowska-Curie action (No: 778233).

## References

- [1] Gul A. Agha (1990): *ACTORS - a model of concurrent computation in distributed systems*. MIT Press series in artificial intelligence, MIT Press.
- [2] Stephanie Balzer & Frank Pfenning (2017): *Manifest sharing with session types*. *Proc. ACM Program. Lang.* 1(ICFP), pp. 37:1–37:29, doi:[10.1145/3110281](https://doi.org/10.1145/3110281).
- [3] Christian Bartolo Burlò, Adrian Francalanza & Alceste Scalas (2021): *On the Monitorability of Session Types, in Theory and Practice*. In Anders Møller & Manu Sridharan, editors: *35th European Conference on Object-Oriented Programming, ECOOP 2021, July 11-17, 2021, Aarhus, Denmark (Virtual Conference), LIPIcs 194*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 20:1–20:30, doi:[10.4230/LIPIcs.ECOOP.2021.20](https://doi.org/10.4230/LIPIcs.ECOOP.2021.20).
- [4] Mauricio Cassola, Agustín Talagorria, Alberto Pardo & Marcos Viera (2020): *A Gradual Type System for Elixir*. *Proceedings of the 24th Brazilian Symposium on Context-Oriented Programming and Advanced Modularity*, doi:[10.1145/3427081.3427084](https://doi.org/10.1145/3427081.3427084).
- [5] Mauricio Cassola, Agustín Talagorria, Alberto Pardo & Marcos Viera (2022): *A gradual type system for Elixir*. *Journal of Computer Languages* 68, p. 101077, doi:[10.1016/j.col.2021.101077](https://doi.org/10.1016/j.col.2021.101077).
- [6] Zak Cutner & Nobuko Yoshida (2021): *Safe Session-Based Asynchronous Coordination in Rust*. In Ferruccio Damiani & Ornella Dardha, editors: *Coordination Models and Languages - 23rd IFIP WG 6.1 International Conference, COORDINATION 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, June 14-18, 2021, Proceedings, Lecture Notes in Computer Science 12717*, Springer, pp. 80–89, doi:[10.1007/978-3-030-78142-2\\_5](https://doi.org/10.1007/978-3-030-78142-2_5).
- [7] Ankush Das, Jan Hoffmann & Frank Pfenning (2018): *Work Analysis with Resource-Aware Session Types*. In Anuj Dawar & Erich Grädel, editors: *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, ACM, pp. 305–314, doi:[10.1145/3209108.3209146](https://doi.org/10.1145/3209108.3209146).
- [8] Erlang Solutions: *Gradient*. Available at <https://github.com/esl/gradient>.
- [9] Simon Fowler (2016): *An Erlang Implementation of Multiparty Session Actors*. In Massimo Bartoletti, Ludovic Henrio, Sophia Knight & Hugo Torres Vieira, editors: *Proceedings 9th Interaction and Concurrency Experience, ICE 2016, Heraklion, Greece, 8-9 June 2016, EPTCS 223*, pp. 36–50, doi:[10.4204/EPTCS.223.3](https://doi.org/10.4204/EPTCS.223.3).
- [10] Adrian Francalanza, Edsko de Vries & Matthew Hennessy (2014): *Compositional Reasoning for Explicit Resource Management in Channel-Based Concurrency*. *Log. Methods Comput. Sci.* 10(2), doi:[10.2168/LMCS-10\(2:15\)2014](https://doi.org/10.2168/LMCS-10(2:15)2014).
- [11] Joseph Harrison (2018): *Automatic detection of Core Erlang message passing errors*. In Natalia Chechina & Adrian Francalanza, editors: *Proceedings of the 17th ACM SIGPLAN International Workshop on Erlang, ICFP 2018, St. Louis, MO, USA, September 23-29, 2018*, ACM, pp. 37–48, doi:[10.1145/3239332.3242765](https://doi.org/10.1145/3239332.3242765).

- [12] Joseph Harrison (2019): *Runtime Type Safety for Erlang/OTP Behaviours*. In Adrian Francalanza & Viktória Fördös, editors: *Proceedings of the 18th ACM SIGPLAN International Workshop on Erlang, Erlang@ICFP 2019, Berlin, Germany, August 18, 2019*, ACM, pp. 36–47, doi:[10.1145/3331542.3342571](https://doi.org/10.1145/3331542.3342571).
- [13] Paul Harvey, Simon Fowler, Ornela Dardha & Simon J. Gay (2021): *Multiparty Session Types for Safe Runtime Adaptation in an Actor Language*. In Anders Møller & Manu Sridharan, editors: *35th European Conference on Object-Oriented Programming, ECOOP 2021, July 11-17, 2021, Aarhus, Denmark (Virtual Conference), LIPIcs 194, Schloss Dagstuhl - Leibniz-Zentrum für Informatik*, pp. 10:1–10:30, doi:[10.4230/LIPIcs.ECOOP.2021.10](https://doi.org/10.4230/LIPIcs.ECOOP.2021.10).
- [14] Carl Hewitt, Peter Boehler Bishop & Richard Steiger (1973): *A Universal Modular ACTOR Formalism for Artificial Intelligence*. In Nils J. Nilsson, editor: *Proceedings of the 3rd International Joint Conference on Artificial Intelligence. Stanford, CA, USA, August 20-23, 1973*, William Kaufmann, pp. 235–245.
- [15] Kohei Honda, Aybek Mukhamedov, Gary Brown, Tzu-Chun Chen & Nobuko Yoshida (2011): *Scribbling Interactions with a Formal Foundation*. In Raja Natarajan & Adegboyega K. Ojo, editors: *Distributed Computing and Internet Technology - 7th International Conference, ICDCIT 2011, Bhubaneswar, India, February 9-12, 2011. Proceedings, Lecture Notes in Computer Science 6536*, Springer, pp. 55–75, doi:[10.1007/978-3-642-19056-8\\_4](https://doi.org/10.1007/978-3-642-19056-8_4).
- [16] Kohei Honda, Nobuko Yoshida & Marco Carbone (2016): *Multiparty Asynchronous Session Types*. *J. ACM* 63(1), pp. 9:1–9:67, doi:[10.1145/2827695](https://doi.org/10.1145/2827695).
- [17] Thomas Bracht Laumann Jespersen, Philip Munksgaard & Ken Friis Larsen (2015): *Session types for Rust*. In Patrick Bahr & Sebastian Erdweg, editors: *Proceedings of the 11th ACM SIGPLAN Workshop on Generic Programming, WGP@ICFP 2015, Vancouver, BC, Canada, August 30, 2015*, ACM, pp. 13–22, doi:[10.1145/2808098.2808100](https://doi.org/10.1145/2808098.2808100).
- [18] Robert M. Keller (1976): *Formal Verification of Parallel Programs*. *Commun. ACM* 19(7), pp. 371–384, doi:[10.1145/360248.360251](https://doi.org/10.1145/360248.360251).
- [19] Wen Kokke (2019): *Rusty Variation: Deadlock-free Sessions with Failure in Rust*. In Massimo Bartoletti, Ludovic Henrio, Anastasia Mavridou & Alceste Scalas, editors: *Proceedings 12th Interaction and Concurrency Experience, ICE 2019, Copenhagen, Denmark, 20-21 June 2019, EPTCS 304*, pp. 48–60, doi:[10.4204/EPTCS.304.4](https://doi.org/10.4204/EPTCS.304.4).
- [20] Nicolas Lagailardie, Romyana Neykova & Nobuko Yoshida (2020): *Implementing Multiparty Session Types in Rust*. In Simon Bliudze & Laura Bocchi, editors: *Coordination Models and Languages - 22nd IFIP WG 6.1 International Conference, COORDINATION 2020, Held as Part of the 15th International Federated Conference on Distributed Computing Techniques, DisCoTec 2020, Valletta, Malta, June 15-19, 2020, Proceedings, Lecture Notes in Computer Science 12134*, Springer, pp. 127–136, doi:[10.1007/978-3-030-50029-0\\_8](https://doi.org/10.1007/978-3-030-50029-0_8).
- [21] Tobias Lindahl & Konstantinos Sagonas (2006): *Practical type inference based on success typings*. In Annalisa Bossi & Michael J. Maher, editors: *Proceedings of the 8th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming, July 10-12, 2006, Venice, Italy*, ACM, pp. 167–178, doi:[10.1145/1140335.1140356](https://doi.org/10.1145/1140335.1140356).
- [22] Hernán C. Melgratti & Luca Padovani (2017): *Chaperone Contracts for Higher-Order Sessions*. *Proc. ACM Program. Lang.* 1(ICFP), pp. 35:1–35:29, doi:[10.1145/3110279](https://doi.org/10.1145/3110279).



- [23] Dimitris Mostrous & Vasco Thudichum Vasconcelos (2011): *Session Typing for a Featherweight Erlang*. In Wolfgang De Meuter & Gruija-Catalin Roman, editors: *Coordination Models and Languages - 13th International Conference, COORDINATION 2011, Reykjavik, Iceland, June 6-9, 2011. Proceedings, Lecture Notes in Computer Science 6721*, Springer, pp. 95–109, doi:[10.1007/978-3-642-21464-6\\_7](https://doi.org/10.1007/978-3-642-21464-6_7).
- [24] Romyana Neykova & Nobuko Yoshida (2014): *Multiparty Session Actors*. In Alastair F. Donaldson & Vasco T. Vasconcelos, editors: *Proceedings 7th Workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software, PLACES 2014, Grenoble, France, 12 April 2014, EPTCS 155*, pp. 32–37, doi:[10.4204/EPTCS.155.5](https://doi.org/10.4204/EPTCS.155.5).
- [25] Romyana Neykova & Nobuko Yoshida (2017): *Multiparty Session Actors*. *Log. Methods Comput. Sci.* 13(1), doi:[10.23638/LMCS-13\(1:17\)2017](https://doi.org/10.23638/LMCS-13(1:17)2017).
- [26] Luca Padovani (2017): *A Simple Library Implementation of Binary Sessions*. *J. Funct. Program.* 27, p. e4, doi:[10.1017/S0956796816000289](https://doi.org/10.1017/S0956796816000289).
- [27] Benjamin C. Pierce (2002): *Types and Programming Languages*. MIT Press.
- [28] Klaas Pruiksma & Frank Pfenning (2022): *Back to futures*. *Journal of Functional Programming* 32, p. e6, doi:[10.1017/S0956796822000016](https://doi.org/10.1017/S0956796822000016).
- [29] Alceste Scalas & Nobuko Yoshida (2016): *Lightweight Session Programming in Scala*. In Shriram Krishnamurthi & Benjamin S. Lerner, editors: *30th European Conference on Object-Oriented Programming, ECOOP 2016, July 18-22, 2016, Rome, Italy, LIPIcs 56*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 21:1–21:28, doi:[10.4230/LIPIcs.ECOOP.2016.21](https://doi.org/10.4230/LIPIcs.ECOOP.2016.21).
- [30] Alceste Scalas, Nobuko Yoshida & Elias Benussi (2019): *Verifying message-passing programs with dependent behavioural types*. In Kathryn S. McKinley & Kathleen Fisher, editors: *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, ACM, pp. 502–516, doi:[10.1145/3314221.3322484](https://doi.org/10.1145/3314221.3322484).
- [31] Hans Svensson, Lars-Åke Fredlund & Clara Benac Earle (2010): *A unified semantics for future Erlang*. In Scott Lystig Fritchie & Konstantinos Sagonas, editors: *Proceedings of the 9th ACM SIGPLAN workshop on Erlang, Baltimore, Maryland, USA, September 30, 2010*, ACM, pp. 23–32, doi:[10.1145/1863509.1863514](https://doi.org/10.1145/1863509.1863514).
- [32] Gerard Tabone & Adrian Francalanza (2021): *Session Types in Elixir*. In Elias Castegren, Joeri De Koster & Simon Fowler, editors: *Proceedings of the 11th ACM SIGPLAN International Workshop on Programming Based on Actors, Agents, and Decentralized Control, AGERE 2021, Virtual Event / Chicago, IL, USA, 17 October 2021*, ACM, pp. 12–23, doi:[10.1145/3486601.3486708](https://doi.org/10.1145/3486601.3486708).
- [33] Gerard Tabone & Adrian Francalanza (2022): *Static Checking of Concurrent Programs in Elixir Using Session Types*. Technical Report, University of Malta, Msida, Malta. Available at <https://gtabone.page.link/V9Hh>.
- [34] Dave Thomas (2018): *Programming Elixir: Functional, Concurrent, Pragmatic, Fun*. Pragmatic Bookshelf.
- [35] Wiebe-Marten Wijnja: *TypeCheck: Fast and flexible runtime type-checking for your Elixir projects*. Available at [https://github.com/Qqwy/elixir-type\\_check](https://github.com/Qqwy/elixir-type_check).

# Appendix

## A Additional Definitions

In this appendix, we formalise some auxiliary definitions that were used in [Sections 2–4](#).

**Definition A.1** (*Type*).

$$\begin{aligned} \mathbf{type}(boolean) &\stackrel{\text{def}}{=} \text{boolean} & \mathbf{type}(number) &\stackrel{\text{def}}{=} \text{number} \\ \mathbf{type}(atom) &\stackrel{\text{def}}{=} \text{atom} & \mathbf{type}(t) &\stackrel{\text{def}}{=} \text{pid, where } t \text{ is a pid instance} \quad \blacksquare \end{aligned}$$

**Definition A.2** (*Function Details*). We can extract function details (*i.e.*, `params`, `body`, `param_types`, `return_type`, `dual`) from a list of functions ( $\tilde{Q}$ ) and build a mapping, using set-comprehension, as follows. The list of functions ( $\tilde{Q}$ ) may consist of public ( $D$ ) and private ( $P$ ) functions.

$$\mathbf{details}(\tilde{Q}) \stackrel{\text{def}}{=} \left\{ f/n : \begin{bmatrix} \text{dual} = y, \text{params} = \tilde{x}, \\ \text{param\_types} = \tilde{T}, \\ \text{return\_type} = T, \text{body} = t \end{bmatrix} \mid \begin{bmatrix} [\text{@session "S"}] \\ \text{@spec } f(\text{pid}, \tilde{T}) :: T \\ \text{def [p]} f(y, \tilde{x}) \text{ do } t \text{ end} \end{bmatrix} \in \tilde{Q} \right\}$$

**Definition A.3** (*Functions Names and Arity*). This definition takes the set of all public function ( $\tilde{D}$ ) as input, and returns a set of all public function names and their arity.

$$\mathbf{functions}(\tilde{D}) \stackrel{\text{def}}{=} \left\{ f/n \mid \begin{bmatrix} \text{@session...; @spec...} \\ \text{def } f(y, x_2, \dots, x_n) \text{ do } t \text{ end} \end{bmatrix} \in \tilde{D} \right\}$$

**Definition A.4** (*All Session Types*). The function  $\mathbf{sessions}(\tilde{D})$ , returns the session type corresponding to each annotated public function.

$$\mathbf{sessions}(\tilde{D}) \stackrel{\text{def}}{=} \left\{ f/n : S \mid \begin{bmatrix} \text{@session "S"; @spec...} \\ \text{def } f(y, x_2, \dots, x_n) \text{ do } t \text{ end} \end{bmatrix} \in \tilde{D} \right\}$$

In case the `@dual` annotation is used instead of `@session`, the dual session type is computed automatically. ■

## B Type System Rules

In this appendix, we present the full typing rules of the type system, adapted from [32], which were omitted from the **Preliminaries** Section.

### B.1 Term Typing

In **Section 2.4**, we explained a few term typing rules, including [TBRANCH] and [TCHOICE]. In **Figure 7**, we present the full list of term typing rules.

$$\boxed{\Delta \cdot \Gamma \vdash_{\Sigma}^w S \triangleright t : T \triangleleft S'}$$

$$\text{[TEXPRESSION]} \frac{\Gamma \vdash_{\text{exp}} e : T}{\Delta \cdot \Gamma \vdash^w S \triangleright e : T \triangleleft S}$$

$$\text{[TLET]} \frac{\Delta \cdot \Gamma \vdash^w S \triangleright t_1 : T' \triangleleft S'' \quad \Delta \cdot (\Gamma, x : T') \vdash^w S'' \triangleright t_2 : T \triangleleft S' \quad x \neq w}{\Delta \cdot \Gamma \vdash^w S \triangleright x = t_1; t_2 : T \triangleleft S'}$$

$$\text{[TBRANCH]} \frac{\forall i \in I. \left\{ \begin{array}{l} \forall j \in 1..n. \left\{ \vdash_{\text{pat}}^w p_i^j : T_i^j \triangleright \Gamma_i^j \right\} \\ \Delta \cdot (\Gamma, \Gamma_i^1, \dots, \Gamma_i^n) \vdash^w S_i \triangleright t_i : T \triangleleft S' \end{array} \right.}{\Delta \cdot \Gamma \vdash^w \&\{?l_i(\tilde{T}_i).S_i\}_{i \in I} \triangleright \text{receive do } (\{!l_i, \tilde{p}_i\} \rightarrow t_i)_{i \in I} \text{end} : T \triangleleft S'}$$

$$\text{[TCHOICE]} \frac{\exists i \in I \quad l = l_i \quad \forall j \in 1..n. \left\{ \Gamma \vdash_{\text{exp}} e_j : T_i^j \right\}}{\Delta \cdot \Gamma \vdash^w \oplus\{!l_i(\tilde{T}_i).S_i\}_{i \in I} \triangleright \text{send}(w, \{!l, e_1, \dots, e_n\}) : \{\text{atom}, T_i^1, \dots, T_i^n\} \triangleleft S_i}$$

$$\text{[TRECKNOWNCALL]} \frac{\Delta(f/n) = S \quad \forall i \in 2..n. \left\{ \Gamma \vdash_{\text{exp}} e_i : T_i \right\} \quad \Sigma(f/n) = \Omega \quad \Omega.\text{return\_type} = T \quad \Omega.\text{param\_types} = \tilde{T}}{\Delta \cdot \Gamma \vdash^w S \triangleright f(w, e_2, \dots, e_n) : T \triangleleft \text{end}}$$

$$\text{[TRECUNKNOWNCALL]} \frac{\Sigma(f/n) = \Omega \quad f/n \notin \mathbf{dom}(\Delta) \quad \Omega.\text{dual} = y \quad \Omega.\text{params} = \tilde{x} \quad \Omega.\text{param\_type} = \tilde{T} \quad \Omega.\text{body} = t \quad \Omega.\text{return\_type} = T}{(\Delta, f/n : S) \cdot (\Gamma, y : \text{pid}, \tilde{x} : \tilde{T}) \vdash^y S \triangleright t : T \triangleleft S' \quad \forall i \in 2..n. \left\{ \Gamma \vdash_{\text{exp}} e_i : T_i \right\}}{\Delta \cdot \Gamma \vdash^w S \triangleright f(w, e_2, \dots, e_n) : T \triangleleft S'}$$

$$\text{[TCASE]} \frac{\Gamma \vdash_{\text{exp}} e : U \quad \forall i \in I \quad \vdash_{\text{pat}}^w p_i : U \triangleright \Gamma_i' \quad \Delta \cdot (\Gamma, \Gamma_i') \vdash^w S \triangleright t_i : T \triangleleft S'}{\Delta \cdot \Gamma \vdash^w S \triangleright \text{case } e \text{ do } (p_i \rightarrow t_i)_{i \in I} \text{end} : T \triangleleft S'}$$

Figure 7: Term typing rules

## B.2 Expression Typing

Expressions are typechecked using the  $\Gamma \vdash_{\text{exp}} e : T$  judgement, which states that “an expression  $e$  has type  $T$ , subject to the *variable binding* environment  $\Gamma$ .” The expression typing rules are listed in [Figure 8](#).

$$\boxed{\Gamma \vdash_{\text{exp}} e : T}$$

$$\begin{array}{c}
 [\text{TTUPLE}] \frac{\forall i \in 1..n \quad \Gamma \vdash_{\text{exp}} e_i : T_i}{\Gamma \vdash_{\text{exp}} \{e_1, \dots, e_n\} : \{T_1, \dots, T_n\}} \\
 \\
 [\text{TLITERAL}] \frac{\text{type}(b) = T \quad b \neq []}{\Gamma \vdash_{\text{exp}} b : T} \qquad [\text{TVARIABLE}] \frac{\Gamma(x) = T}{\Gamma \vdash_{\text{exp}} x : T} \\
 \\
 [\text{TLIST}] \frac{\Gamma \vdash_{\text{exp}} e_1 : T \quad \Gamma \vdash_{\text{exp}} e_2 : [T]}{\Gamma \vdash_{\text{exp}} [e_1 \mid e_2] : [T]} \qquad [\text{TELIST}] \frac{}{\Gamma \vdash_{\text{exp}} [] : [T]} \\
 \\
 [\text{TARITHMETIC}] \frac{\Gamma \vdash_{\text{exp}} e_1 : \text{number} \quad \Gamma \vdash_{\text{exp}} e_2 : \text{number} \quad \diamond \in \{+, -, *, /\}}{\Gamma \vdash_{\text{exp}} e_1 \diamond e_2 : \text{number}} \\
 \\
 [\text{TBOOLEAN}] \frac{\Gamma \vdash_{\text{exp}} e_1 : \text{boolean} \quad \Gamma \vdash_{\text{exp}} e_2 : \text{boolean} \quad \diamond \in \{\text{and}, \text{or}\}}{\Gamma \vdash_{\text{exp}} e_1 \diamond e_2 : \text{boolean}} \\
 \\
 [\text{TCOMPARISONS}] \frac{\diamond \in \{<, >, <=, >=, ==, !=\} \quad \Gamma \vdash_{\text{exp}} e_1 : T \quad \Gamma \vdash_{\text{exp}} e_2 : T}{\Gamma \vdash_{\text{exp}} e_1 \diamond e_2 : \text{boolean}} \qquad [\text{TNOT}] \frac{\Gamma \vdash_{\text{exp}} e : \text{boolean}}{\Gamma \vdash_{\text{exp}} \text{not } e : \text{boolean}}
 \end{array}$$

Figure 8: Expression typing rules

## B.3 Pattern Typing

New variables may be created using patterns in the  $[\text{TBRANCH}]$  and  $[\text{TCASE}]$  rules. These variables are matched to a type using the judgement,  $\vdash_{\text{pat}}^w p : T \triangleright \Gamma$ . This judgement states that “a pattern  $p$  is matched to type  $T$ , where it produces new variables and their types are collected  $\Gamma$ ; under the assumption that the variable containing the dual  $pid$ ,  $w$ , remains unchanged.” The pattern typing rules are found in [Figure 9](#).

$$\boxed{\vdash_{\text{pat}}^w p : T \triangleright \Gamma}$$

$$\begin{array}{c}
 [\text{TPLITERAL}] \frac{\emptyset \vdash_{\text{exp}} b : T \quad b \neq []}{\vdash_{\text{pat}}^w b : T \triangleright \emptyset} \qquad [\text{TPVARIABLE}] \frac{x \neq w}{\vdash_{\text{pat}}^w x : T \triangleright x : T} \\
 \\
 [\text{TPTUPLE}] \frac{\forall i \in 1..n \quad \vdash_{\text{pat}}^w w_i : T_i \triangleright \Gamma_i}{\vdash_{\text{pat}}^w \{w_1, \dots, w_n\} : \{T_1, \dots, T_n\} \triangleright \Gamma_1, \dots, \Gamma_n} \\
 \\
 [\text{TPLIST}] \frac{\vdash_{\text{pat}}^w w_1 : T \triangleright \Gamma_1 \quad \vdash_{\text{pat}}^w w_2 : [T] \triangleright \Gamma_2}{\vdash_{\text{pat}}^w [w_1 \mid w_2] : [T] \triangleright \Gamma_1, \Gamma_2} \qquad [\text{TPELIST}] \frac{}{\vdash_{\text{pat}}^w [] : [T] \triangleright \emptyset}
 \end{array}$$

Figure 9: Pattern typing rules