

# Session Types for Inter-Process Communication

Simon Gay<sup>1</sup>, Vasco Vasconcelos<sup>2</sup>, and António Ravara<sup>3</sup>

<sup>1</sup> Department of Computing Science, University of Glasgow, Glasgow G12 8QQ, UK.  
Email: <simon@dcs.gla.ac.uk>

<sup>2</sup> Departamento de Informática, Faculdade de Ciências da Universidade de Lisboa,  
1749-016 Lisboa, Portugal. Email: <vv@di.fc.ul.pt>

<sup>3</sup> Departamento de Matemática, Instituto Superior Técnico, 1049-001 Lisboa,  
Portugal. Email: <amar@math.ist.utl.pt>

**Abstract.** We define a language whose type system, incorporating session types, allows complex protocols to be specified by types and verified by static typechecking. Although session types are well understood in the context of the  $\pi$ -calculus, our formulation is based on  $\lambda$ -calculus with side-effecting input/output operations and is different in significant ways. Our typing judgements statically describe dynamic changes in the types of channels, our channel types statically track aliasing, and our function types not only specify argument and result types but also describe changes in channel types. After formalising the syntax, semantics and typing rules of our language, and proving a subject reduction theorem, we outline some possibilities for adding references, objects and concurrency.

**Keywords:** Session types, static typechecking, semantics, distributed programming, specification of communication protocols.

## 1 Introduction

Communication in distributed systems is typically structured by protocols, which specify the sequence and form of messages passing over communication channels between agents. In order for correct communication to occur, it is essential that protocols are obeyed.

The theory of *session types* [6, 12] allows the specification of a protocol to be expressed as a type; when a communication channel is created, a session type is associated with it. A session type specifies not only the data types of individual messages, but also the state transitions of the protocol and hence the allowable sequences of messages. By extending the standard methodology of static typechecking, it becomes possible to verify, at compile-time, that an agent using the channel do so in accordance with the protocol.

The theory of session types has been developed in the context of the  $\pi$ -calculus [8, 11], an idealised concurrent programming language which focuses on inter-process communication. Session types have not yet been incorporated into a mainstream programming language, or even studied theoretically in the context of a standard language paradigm: functional, imperative or object-oriented. Vallecillo *et al.* [13] use session types to add behavioural information to

the interfaces of CORBA objects, and use Gay and Hole's [3] theory of subtyping to formalise compatibility and substitutability of components, but they have not attempted to design a complete language.

In the absence of session types, current languages do little to assist the programmer in checking that a protocol has been implemented correctly. Although recent developments in programming languages have increasingly emphasised the benefits of static typechecking, programming with communication channels or streams remains largely untyped. In some respects the situation has deteriorated: Pascal provided static typechecking of file input/output, but the modern use of the stream abstraction for both files and network sockets has resulted in the loss of even this level of support.

The aim of this paper is to study session types in a language based on  $\lambda$ -calculus with side-effecting input/output operations, and to establish that compile-time typechecking of session types could feasibly be added to a mainstream imperative language. Although our language is somewhat idealised, we have attempted to organise it around realistic principles, and in particular to address the key differences between a conventional programming style and the programming notation of the  $\pi$ -calculus.

The structure of the paper is as follows. In Section 2 we explain session types in connection with a progressively more sophisticated server for mathematical operations. Sections 3, 4 and 5 define the syntax, operational semantics and type system of our language. In Section 6 we outline the proof of soundness of our type system. Section 7 concludes, and indicates some possibilities for extending the language with references, objects, and concurrency.

## 2 Session Types and the Maths Server

### 2.1 Input, Output and Sequencing Types

First consider a server which provides a single operation: addition of integers. A suitable protocol can be defined as follows.

The client sends two integers. The server sends an integer which is their sum, then closes the connection.

The corresponding session type, from the server's point of view, is

$$S = ?\text{Int}.\text{?Int}!\text{Int}.\text{End}$$

in which  $?$  means *receive*,  $!$  means *send*, dot ( $.$ ) is sequencing, and **End** indicates the end of the session. The type does not correspond precisely to the specification, because it does not state that the server calculates the sum. However, the type captures the parts of the specification which we can reasonably expect to verify statically.

The server communicates with a client on a channel called  $c$ ; we think of the client engaging in a *session* with the server, using the channel  $c$  for communication. We can view  $c$  as a TCP/IP socket connection, which can be used for

bidirectional communication; in practice it would typically be necessary to extract separate input and output streams from  $c$ . For the moment we will ignore the mechanism of establishing the connection.

In our language, which uses ML-style let-bindings (as syntactic sugar for applications of  $\lambda$ -abstractions), the server looks like this:

```
let  $x = \text{receive } c$ 
     $y = \text{receive } c$ 
in  $\text{send } x + y \text{ on } c$ 
```

or more concisely:

```
 $\text{send } ((\text{receive } c) + (\text{receive } c)) \text{ on } c$ 
```

In a language with assignable variables (which we have not yet included) it might look like this:

```
var  $x : \text{Int}$ 
     $y : \text{Int}$ 
begin
   $x := \text{receive } c$ 
   $y := \text{receive } c$ 
   $\text{send } x + y \text{ on } c$ 
end
```

Interchanging  $?$  and  $!$  yields the type describing the client side of the protocol<sup>1</sup>:

$$\bar{S} = !\text{Int}.\text{Int}.\text{?Int}.\text{End}$$

and a client implementation uses the server to add two particular integers; the *code* may use  $x$  but cannot use the channel  $c$  except to close it.

```
 $\text{send } 2 \text{ on } c$ 
 $\text{send } 3 \text{ on } c$ 
let  $x = \text{receive } c$ 
in code
```

An alternative possibility is to interpret the specification “the client sends two integers” as a type in which the client sends a single message consisting of a pair of integers (Honda *et al.* [6] call this “piggybacking”). In this case, the type of the server side would be

$$\text{?}(\text{Int} \times \text{Int}).!\text{Int}.\text{End}$$

and the type of the client side would again be obtained by exchanging  $?$  and  $!$ .

---

<sup>1</sup> The duality operator  $\bar{\cdot} : S \mapsto \bar{S}$  [12, 6, 3, 13] is an important part of the theory of session types, but we do not need to discuss it in this paper because we only consider clients or servers in isolation.

## 2.2 Branching Types

Now let us modify the protocol and add a negation operation to the server.

The client selects one of two commands: `add` or `neg`. In the case of `add` the client then sends two integers and the server replies with an integer which is their sum. In the case of `neg` the client then sends an integer and the server replies with an integer which is its negation. In either case, the server then closes the connection.

The corresponding session type, for the server side, uses the constructor  $\&$  (*branch*) to indicate that a choice is offered.

$$S = \&\langle \text{add} : ?\text{Int}.\text{Int}!\text{Int}.\text{End}, \text{ neg} : ?\text{Int}!\text{Int}.\text{End} \rangle$$

Both services must be implemented. We introduce a `case` construct:

```
case c of {
  add ⇒ send ((receive c) + (receive c)) on c
  neg ⇒ send (-receive c) on c }
```

The type of the client side uses the dual constructor  $\oplus$  (*choice*) to indicate that a choice is made.

$$\bar{S} = \oplus\langle \text{add} : !\text{Int}!\text{Int}.\text{Int}.\text{End}, \text{ neg} : !\text{Int}.\text{Int}.\text{End} \rangle$$

A particular client implementation makes a particular choice, for example:

<pre>addclient select add on c send 2 on c send 3 on c let x = receive c in code</pre>	<pre>negclient select neg on c send 4 on c let x = receive c in code</pre>
--	--

Note that the type of the subsequent interaction depends on the label which is chosen. In order for typechecking to be decidable, it is essential that the label `add` or `neg` appears as a literal name in the program; labels cannot result from computations.

If we add a square root operation, `sqrt`, then as well as specifying that the argument and result have type `Real`, we must allow for the possibility of an error (resulting in the end of the session) if the client asks for the square root of a negative number. This is done by using the  $\oplus$  constructor on the server side, with options `ok` and `error`. The complete English description of the protocol is starting to become lengthy, so we will omit it and simply show the type of the server side.

$$S = \&\langle \text{add} : ?\text{Int}.\text{Int}!\text{Int}.\text{End}, \\ \text{neg} : ?\text{Int}!\text{Int}.\text{End}, \\ \text{sqrt} : ?\text{Real} . \oplus\langle \text{ok} : !\text{Real}.\text{End}, \text{ error} : \text{End} \rangle \rangle$$

In the type of the client side, the occurrence of  $\&$  indicates that the client must be prepared for both `ok` and `error` responses from the server.

$$\bar{S} = \oplus \langle \text{add} : !\text{Int} !\text{Int} ?\text{Int} . \text{End}, \\ \text{neg} : !\text{Int} ?\text{Int} . \text{End}, \\ \text{sqrt} : !\text{Real} . \& \langle \text{ok} : ?\text{Real} . \text{End}, \text{error} : \text{End} \rangle \rangle$$

More realistically we might like the operations of the server to allow both integer and real arguments and return results of the appropriate type. This is supported by the theory of subtyping for session types [3, 5] but we have not yet incorporated it into our present language.

### 2.3 Establishing a Connection

We have not yet considered the question of how the client and the server reach a state in which they both know about the channel  $c$ . In the pi calculus, it is natural to define a complete system consisting of a client and a server in parallel. Previous studies of session types in the  $\pi$ -calculus have suggested two mechanisms for creating a connection. Takeuchi, Kubo and Honda [12] propose a pair of constructs: `request`  $c$  in  $P$  for use by clients, and `accept`  $c$  in  $Q$  for use by servers. In use, `request` and `accept` occur in separate parallel processes, and interact with each other to create a new channel; this channel is bound to the name  $c$  in both  $P$  and  $Q$ . Gay and Hole [3] use the standard  $\pi$ -calculus `new` construct; the client creates a new channel and sends one end of it to the server along a public channel.

In both cases, the creation and naming of a connection are combined into a single operation. If we want to use session types in an imperative or object-oriented language, it is more realistic for the connection to be a value which is returned by an operator and which can then be assigned to a variable. Furthermore, we have not yet added concurrency to our language, and at this stage we are just considering clients or servers as isolated programs. We therefore use `chan`  $S$  to create a channel with type  $S$ , and view this operation as an abstraction of both requesting and accepting network connections. A complete client or server will have the form `let`  $x = \text{chan } S$  in  $\dots$ ; `close`  $x$

### 2.4 Function Types

Our maths server is bound to a particular channel,  $c$ . If we need the server to answer on different channels, we may abstract the channel  $c$ , transforming the maths server into a *function*.

$$\text{fun } \textit{serve } c = \text{case } c \text{ of } \{ \dots \}$$

Then we may have a server answering on  $\textit{myChannel}$ , by applying the function  $\textit{serve}$  to the channel  $\textit{myChannel}$ , as in  $\textit{serve } \textit{myChannel}$ . The type of our server

now reflects, not only the fact that it accepts a channel and returns nothing (that is, the type `Unit`), but also information on how the function uses the channel:

```
c: &<add: ..., neg: ..., sqrt...: > ▷ chan c → Unit ◁ c: End
```

It can also be useful to send functions on channels. For example we could add

```
eval: ?(▷ Int → Bool ◁).?Int.!Bool.End
```

to the type  $S$ , with corresponding server code

```
eval ⇒ send ((receive c) (receive c)) on c
```

and a client which requires a primality test service (perhaps the server has fast hardware):

```
select eval on c
send isPrime on c
send bignumber on c
let x = receive c in code
```

If it is possible to send channels on channels (we have not yet transferred this feature from the  $\pi$ -calculus to our language) then the server can release  $c$  earlier by accepting a channel on which to conduct the function application:

```
eval: ?(Chan ?(▷ Int → Bool ◁).?Int.!Bool.End).End
eval ⇒ let d = receive c in send ((receive d) (receive d)) on d
```

and the client would create a channel ( $d$ , here). Eventually the client uses `close d` to close the channel, changing its type from `End` to  $\perp$ .

```
select eval on c
let d = chan ?(▷ Int → Bool ◁).?Int.!Bool.End
in send d on c
  send isPrime on d
  send bignumber on d
  let x = receive d
  in close d
  code
```

## 2.5 Recursive Types

A more realistic server would allow a session to consist of a sequence of commands and responses. The corresponding type must be defined recursively, and it is useful to include a quit command. Here is the type of the server side:

```
S = &<add: ?Int.?Int.!Int.S,
  neg: ?Int.!Int.S,
  sqrt: ?Real . ⊕(ok: !Real.S, error: S),
  quit: End>
```

The server is now implemented by a recursive function, in which the positions of the recursive calls correspond to the recursive occurrences of  $S$  in the type definition. Our type system supports the use of recursive functions with any recursive type and any combination of terminating and recursive branches, and is flexible enough to allow some branches to work through the body of the recursive type more than once if desired.

```

fun serve c =
  case c of {
    add ⇒ send ((receive c) + (receive c)) on c
          serve c
    neg ⇒ send (-(receive c)) on c
          serve c
    sqrt ⇒ let x = receive c
            in if x < 0 then select error on c
              else select ok on c
                send  $\sqrt{x}$  on c
              serve c
    quit ⇒ close c }

```

## 2.6 Aliasing of Channels

As soon as we separate creation and naming of channels, aliasing becomes an issue. In this program:

```

let x = chan !Int.!Int.End
    y = x
in send 1 on x
   send 2 on y

```

$x$  and  $y$  are aliases for a single underlying channel, and the two `send` operations reduce the type to `End`. Clearly our type system must track aliases in order to be able to correctly typecheck programs such as this. Our approach is to introduce indirection into type environments. In this example, `chan` creates a channel with some identity,  $c$  say, and the types of both  $x$  and  $y$  are `Chan c`. The session type of  $c$ , initially `!Int.!Int.End`, is recorded separately.

## 3 Syntax

Most of the syntax of our language has been illustrated in the previous sections; here we define it formally by the grammar in Figure 1.

We define data types  $D$ , session types  $S$ , channel environments  $\Sigma$ , term types  $T$ , values  $v$  and terms  $e$ . We use channel identifiers  $c, \dots$ , term identifiers  $x, \dots$ , label identifiers  $l, \dots$ , and type variables  $X, \dots$ . Evaluation contexts  $E$  are used in the definition of the operational semantics; an evaluation context contains a single hole  $[ ]$  and we write  $E[e]$  for substitution of term  $e$  into the hole. Free

$$\begin{aligned}
D &::= \text{Int} \mid \text{Bool} \mid \text{Unit} \\
S &::= \perp \mid \text{End} \mid ?D.S \mid !D.S \mid \&\langle l_i : S_i \rangle_{i \in I} \mid \oplus \langle l_i : S_i \rangle_{i \in I} \mid X \mid \mu X \cdot S \\
&\quad \mid \Sigma \triangleright T \rightarrow T \triangleleft \Sigma \\
\Sigma &::= c_1 : S_1, \dots, c_n : S_n \quad (\text{the } c_i \text{ are pairwise distinct}) \\
T &::= D \mid \text{Chan } c \\
v &::= x \mid \lambda x.e \mid c \mid \text{unit} \mid \text{fix} \mid \text{unfold} \mid \text{receive} \mid \text{send} \mid \text{close} \\
&\quad \mid \text{true} \mid \text{false} \mid 0 \mid 1 \mid -1 \mid \dots \\
e &::= v \mid ee \mid \text{if } e \text{ then } e \text{ else } e \mid \text{case } e \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} \mid \text{select } l \text{ on } e \mid \text{chan } S \\
E &::= [] \mid Ee \mid vE \mid \text{if } E \text{ then } e \text{ else } e \mid \text{case } E \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} \mid \text{select } l \text{ on } E
\end{aligned}$$

**Fig. 1.** Syntax of types, terms, and contexts

and bound variables are defined as usual and we work up to  $\alpha$ -equivalence; the binding occurrences are  $x$  in  $\lambda x.e$ , and  $X$  in  $\mu X \cdot S$ . Substitution is defined as expected.

We do not allow channels to be transmitted on channels. It should be possible to remove this restriction, as the necessary typing techniques are well-understood in the  $\pi$ -calculus, but we leave it for future work. The type  $\perp$  describes a channel which has been fully used and explicitly closed. The type  $\text{Chan } c$  represents the type of the channel with identity  $c$ ; the actual session type is associated with  $c$  in the typing environment, as will become clear in Section 5. Channel identifiers  $c$  are not available in the top-level syntax of terms; they arise only during reduction of terms.

In Section 2 we use two *derived constructors*. A term  $\text{let } x = e \text{ in } e'$  stands for  $(\lambda x.e')e$ , and  $e; e'$  (implied in our examples by the indentation) is an abbreviation for  $(\lambda y.e')e$ , provided  $y$  does not occur free in  $e'$ .

## 4 Operational Semantics

We define a small step operational semantics, making use of evaluation contexts to simplify the definitions. In order to be able to prove a Subject Reduction theorem, we include channel environments  $\Sigma$  in reductions. This is simply to track the changing types of channels; there are no runtime type checks.

We first define the relation  $\Sigma, e \longrightarrow_{\nu} \Sigma', e'$  by the axiom schemas in Figure 2. We abbreviate an axiom schema  $\Sigma, e \longrightarrow_{\nu} \Sigma', e'$  by  $e \longrightarrow_{\nu} e'$ . The  $+$  operation is disjoint union on maps. We then define the relation  $\Sigma, e \longrightarrow \Sigma', e'$  by the rule R-CONTEXT. Some comments on the reduction rules:

- R-CHAN creates a new channel with an arbitrary identity. This models both requesting and accepting a connection.
- R-CASE allows any permitted label to be chosen. This use of  $\Sigma$  in order to define reductions is only necessary because we are working with programs



$\text{if true then } e \text{ else } e' \longrightarrow_v e$	(R-IFT)
$\text{if false then } e \text{ else } e' \longrightarrow_v e'$	(R-IFF)
$(\lambda x.e)v \longrightarrow_v e\{v/x\}$	(R-BETA)
$\text{fix } v \longrightarrow_v v(\text{fix } v)$	(R-FIX)
$\Sigma + c : \mu X \cdot S, \text{ unfold } c \longrightarrow_v \Sigma + c : S\{(\mu X \cdot S)/X\}, \text{ unit}$	(R-UNFOLD)
$\Sigma, \text{ chan } S \longrightarrow_v \Sigma + c : S, c$	(R-CHAN)
$\Sigma + c : !D.S, \text{ send } v \text{ on } c \longrightarrow_v \Sigma + c : S, \text{ unit}$	(R-SEND)
$\Sigma + c : ?D.S, \text{ receive } c \longrightarrow_v \Sigma + c : S, v \quad \text{if } \text{typeof}(v) = D$	(R-RECEIVE)
$\Sigma + c : \text{End}, \text{ close } c \longrightarrow_v \Sigma + c : \perp, \text{ unit}$	(R-CLOSE)
$\Sigma + c : \&\langle l_i : S_i \rangle_{i \in I}, \text{ case } c \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} \longrightarrow_v \Sigma + c : S_j, e_j \quad \text{if } j \in I$	(R-CASE)
$\Sigma + c : \oplus \langle l_i : S_i \rangle_{i \in I}, \text{ select } l_j \text{ on } c \longrightarrow_v \Sigma + c : S_j, \text{ unit} \quad \text{if } j \in I$	(R-SELECT)
$\frac{\Sigma, e \longrightarrow_v \Sigma', e'}{\Sigma, E[e] \longrightarrow_v \Sigma', E[e']}$	(R-CONTEXT)

**Fig. 2.** Reduction rules

which only use one end of a channel. In effect, our Subject Reduction theorem (Section 6) is proved relative to the assumption that the program at the other end of the channel does not introduce any type violations.

- Similarly R-RECEIVE allows `receive c` to evaluate to any value permitted by the type of `c`.

## 5 The Type System

Typing judgements are of the form

$$\Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma'$$

where  $\Gamma$  is a map from variables to types and  $\Sigma, \Sigma'$  are channel environments as in Section 3. The difference between  $\Sigma$  and  $\Sigma'$  reflects the effect of a term on the types of channels, for example

$$x : \text{Chan } c \vdash c : !\text{Int.End} \triangleright \text{send } 2 \text{ on } x : \text{Unit} \triangleleft c : \text{End}$$

The assignment of types to values is shown in Figure 3. The values `send`, `receive`, `close`, `unfold` and `fix` have multiple types: for example, the type of `receive` is a type schema representing the set of all types of the form  $c : ?D.S \triangleright \text{Chan } c \rightarrow D \triangleleft c : S$  where  $D$  is an arbitrary data type,  $S$  is an arbitrary session type and  $c$  is an arbitrary channel identifier. We write  $\text{typeof}(v)$  for the set of types assigned to value  $v$ . To simplify the theory we explicitly `unfold` recursive types; an implementation would insert `unfold` automatically where necessary.

$\text{true, false: Bool}$   
 $0, 1, -1, \dots: \text{Int}$   
 $\text{unit: Unit}$   
 $\text{close: } c: \text{End} \triangleright \text{Chan } c \rightarrow \text{Unit} \triangleleft c: \perp$   
 $\text{receive: } c: ?D.S \triangleright \text{Chan } c \rightarrow D \triangleleft c: S$   
 $\text{send: } D \rightarrow (c: !D.S \triangleright \text{Chan } c \rightarrow \text{Unit} \triangleleft c: S)$   
 $\text{fix: } \emptyset \triangleright (T \rightarrow T) \rightarrow T \triangleleft \Sigma$   
 $\text{unfold: } c: \mu X \cdot S \triangleright \text{Chan } c \rightarrow \text{Unit} \triangleleft c: S\{(\mu X \cdot S)/X\}$   
 $+: \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}$

**Fig. 3.** Types for constants

The typing rules are shown in Figure 4. The rules with no typing judgements in their hypotheses have side conditions of the form  $\Sigma \vdash \Gamma$  (reading:  $\Gamma$  is well-formed with respect to  $\Sigma$ ), defined by

$$\Sigma \vdash \Gamma \iff x : \text{Chan } c \in \Gamma \implies c \in \text{dom}(\Sigma)$$

which guarantees that  $\Gamma$  contains no references to non-existent channels. Some comments on the typing rules:

- T-CHAN links the identity of a channel to its type in  $\Sigma$ .
- In T-NEW, a new channel is created with an arbitrary identity.
- In T-ABS, the initial and final channel environments of the function body go directly into the function type. The function itself, being a value, cannot affect channels, hence the empty environments both on the left and on the right.
- In T-APP, the final channel environment  $\Sigma''$  of function  $e$  must match the initial environment of the argument  $e'$ . The final environment of the argument is split into two:  $\Sigma_1$ , that satisfies the channel requirements in the function type; and  $\Sigma'$  that must go (together with the final environment in the function type) into the final environment of the application.
- In T-CASE, all branches must produce the same final channel environment. This enables us to know the environment for any code following the `case`, independently of which branch is chosen at runtime. The same applies to the two branches of the conditional in T-IF.

A complete program  $e$  is well-typed if there exist  $T$  and  $\Sigma$  such that

$$\emptyset \vdash \emptyset \triangleright e : T \triangleleft \Sigma$$

Unclosed channels can be detected by examining  $\Sigma$  for types which are not  $\perp$ .

$$\begin{array}{c}
\Gamma \vdash \Sigma \triangleright v : T \triangleleft \Sigma \quad \text{if } \Sigma \vdash \Gamma \text{ and } T \in \text{typeof}(v) \quad (\text{T-CONST}) \\
\Gamma + x : T \vdash \Sigma \triangleright x : T \triangleleft \Sigma \quad \text{if } \Sigma \vdash \Gamma + x : T \quad (\text{T-VAR}) \\
\Gamma \vdash \Sigma + c : S \triangleright c : \text{Chan } c \triangleleft \Sigma + c : S \quad \text{if } \Sigma + c : S \vdash \Gamma \quad (\text{T-CHAN}) \\
\Gamma \vdash \Sigma \triangleright \text{chan } S : \text{Chan } c \triangleleft \Sigma + c : S \quad \text{if } \Sigma \vdash \Gamma \quad (\text{T-NEW}) \\
\frac{\Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma'}{\Gamma \vdash \Sigma + \Sigma'' \triangleright e : T \triangleleft \Sigma' + \Sigma''} \quad (\text{T-WEAK}) \\
\frac{\Gamma \vdash \Sigma \triangleright e : \text{Bool} \triangleleft \Sigma' \quad \Gamma \vdash \Sigma' \triangleright e' : T \triangleleft \Sigma'' \quad \Gamma \vdash \Sigma' \triangleright e'' : T \triangleleft \Sigma''}{\Gamma \vdash \Sigma \triangleright \text{if } e \text{ then } e' \text{ else } e'' : T \triangleleft \Sigma''} \quad (\text{T-IF}) \\
\frac{\Gamma + x : T \vdash \Sigma \triangleright e : U \triangleleft \Sigma' \quad \emptyset \vdash \Gamma}{\Gamma \vdash \emptyset \triangleright \lambda x. e : (\Sigma \triangleright T \rightarrow U \triangleleft \Sigma') \triangleleft \emptyset} \quad (\text{T-ABS}) \\
\frac{\Gamma \vdash \Sigma \triangleright e : (\Sigma_1 \triangleright T \rightarrow U \triangleleft \Sigma_2) \triangleleft \Sigma'' \quad \Gamma \vdash \Sigma'' \triangleright e' : T \triangleleft \Sigma_1 + \Sigma'}{\Gamma \vdash \Sigma \triangleright e e' : U \triangleleft \Sigma_2 + \Sigma'} \quad (\text{T-APP}) \\
\frac{\Gamma \vdash \Sigma \triangleright e : \text{Chan } c \triangleleft \Sigma' + c : \&\langle l_i : S_i \rangle_{i \in I} \quad \forall i. (\Gamma \vdash \Sigma' + c : S_i \triangleright e_i : T \triangleleft \Sigma'')}{\Gamma \vdash \Sigma \triangleright \text{case } e \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} : T \triangleleft \Sigma''} \quad (\text{T-CASE}) \\
\frac{\Gamma \vdash \Sigma \triangleright e : \text{Chan } c \triangleleft \Sigma' + c : \oplus\langle l_i : S_i \rangle_{i \in I}}{\Gamma \vdash \Sigma \triangleright \text{select } l_i \text{ on } e : T \triangleleft \Sigma' + c : S_i} \quad (\text{T-SELECT})
\end{array}$$

**Fig. 4.** Typing rules

### 5.1 Example: Typing the Maths Server

We simplify the maths server from Section 2.5, so that its channel type is

$$T = \mu X \cdot \&\langle \text{neg} : ?\text{Int}.\!\!\text{Int}.X, \text{quit} : \text{End} \rangle$$

and the program, removing some syntactic sugar, is

$$\begin{array}{l}
(\text{fix } (\lambda f. \lambda x. \\
\quad \text{unfold } x \\
\quad \text{case } x \text{ of } \{ \\
\quad \quad \text{neg} \Rightarrow \text{send } (-(\text{receive } x)) \text{ on } x \\
\quad \quad \quad f \ x \\
\quad \quad \text{quit} \Rightarrow \text{close } x \}) \text{ (chan } T)
\end{array}$$

Writing  $F = c : T \triangleright \text{Chan } c \rightarrow \text{Unit} \triangleleft c : \perp$  and setting  $\Gamma = f : F, x : \text{Chan } c$ , we apply T-APP to the following judgements

$$\Gamma \vdash c : \text{End} \triangleright \text{close} : (c : \text{End} \triangleright \text{Chan } c \rightarrow \text{Unit} \triangleleft c : \perp) \triangleleft c : \text{End}$$

$$\Gamma \vdash c : \text{End} \triangleright x : \text{Chan } c \triangleleft c : \text{End}$$

to obtain:

$$\Gamma \vdash c : \text{End} \triangleright \text{close } x : \text{Unit} \triangleleft c : \perp \quad (1)$$

Similarly the types of `receive` and `send`, and two applications of T-APP, give:

$$\Gamma \vdash c : ?\text{Int}!\text{Int}.T \triangleright \text{send } (\text{receive } x) \text{ on } x : \text{Unit} \triangleleft c : T \quad (2)$$

We use T-APP to type the recursive call; note that the type  $F$  was chosen to make (1) and (3) have the same final channel environment:

$$\Gamma \vdash c : T \triangleright f \ x : \text{Unit} \triangleleft c : \perp \quad (3)$$

Applying the expansion  $e;e' \mapsto (\lambda y.e')e$  we can use T-ABS, T-APP and T-WEAK with (2) and (3) to obtain: (call these steps T-SEQ)

$$\Gamma \vdash c : ?\text{Int}!\text{Int}.T \triangleright \text{send } (\text{receive } x) \text{ on } x ; f \ x : \text{Unit} \triangleleft c : \perp \quad (4)$$

and now (4) and (1) are the branches of the case, so we have

$$\Gamma \vdash c : \&\langle \text{neg} : ?\text{Int}!\text{Int}.T, \text{quit} : \text{End} \rangle \triangleright \text{case } x \text{ of } \{ \dots \} : \text{Unit} \triangleleft c : \perp \quad (5)$$

The typing derivation for `fix` is completed as follows:

$$\frac{\Gamma \vdash c : T \triangleright \text{unfold } x : \text{Unit} \triangleleft c : \&\langle \text{neg} : ?\text{Int}!\text{Int}.T, \text{quit} : \text{End} \rangle \quad (5)}{\frac{\frac{\Gamma \vdash c : T \triangleright \text{unfold } x ; \text{case } x \text{ of } \dots : \text{Unit} \triangleleft c : \perp}{\vdash \emptyset \triangleright \lambda f. \lambda x. \dots : (\emptyset \triangleright F \rightarrow F \triangleleft \emptyset) \triangleleft \emptyset} \text{T-ABS (twice)}}{\vdash \emptyset \triangleright \text{fix } \dots : (c : T \triangleright \text{Chan } c \rightarrow \text{Unit} \triangleleft c : \perp) \triangleleft \emptyset} \text{T-APP, fix}} \text{T-SEQ}$$

T-NEW gives  $\emptyset \vdash \emptyset \triangleright \text{chan } T : \text{Chan } c \triangleleft c : T$  and using T-APP again we have

$$\emptyset \vdash \emptyset \triangleright (\text{fix } \dots) (\text{chan } T) : \text{Unit} \triangleleft c : \perp$$

as the typing of the complete program, showing that a channel is created and eventually closed.

## 6 Type Safety

We state the key lemmas leading to the Subject Reduction theorem, and sketch their proofs.

**Lemma 1 (Typability of Subterms in  $E$ ).** *If  $\Gamma \vdash \Sigma \triangleright E[e] : T \triangleleft \Sigma'$  then there exist  $\Sigma''$  and  $U$  such that  $\Gamma \vdash \Sigma \triangleright e : U \triangleleft \Sigma''$ .*

*Proof.* By induction on the structure of  $E$ . The possible positions of the hole mean that in a derivation of  $\Gamma \vdash \Sigma \triangleright E[e] : T \triangleleft \Sigma'$ , there is a typing derivation for  $e$  at the left of the tree (with the same  $\Sigma$ ).

**Lemma 2 (Replacement in  $E$ ).** *If*

1.  $\mathcal{D}$  is a typing derivation concluding  $\Gamma \vdash \Sigma \triangleright E[e] : T \triangleleft \Sigma'$
2.  $\mathcal{D}'$  is a subderivation of  $\mathcal{D}$  concluding  $\Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma_2$

3.  $\mathcal{D}'$  occurs in  $\mathcal{D}$  in the position corresponding to the hole in  $E$
4.  $\Gamma \vdash \Sigma_1 \triangleright e' : T \triangleleft \Sigma_2$

then  $\Gamma \vdash \Sigma_1 \triangleright E[e'] : T \triangleleft \Sigma'$ .

*Proof.* Replace  $\mathcal{D}'$  in  $\mathcal{D}$  by a derivation of  $\Gamma \vdash \Sigma_1 \triangleright e' : T \triangleleft \Sigma_2$ . The structure of  $E$  is essential; it means that the typings of  $e'$  and  $E[e']$  have the same leftmost environment  $\Sigma'$ .

**Lemma 3. (Values do not use channels)** *If  $\Gamma \vdash \Sigma \triangleright v : T \triangleleft \Sigma'$ , then  $\Sigma' = \Sigma$ .*  
**(Well-formedness)** *If  $\Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma'$ , then  $\Sigma \vdash \Gamma$  and  $\Sigma' \vdash \Gamma$ .*  
**(Narrowing)** *If  $\Gamma \vdash \Sigma + \Sigma'' \triangleright e : T \triangleleft \Sigma' + \Sigma''$  and  $\Sigma \vdash \Gamma$ , then  $\Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma'$ .*  
**( $\Gamma$ -Weakening)** *If  $\Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma'$  and  $\Gamma' \vdash \Sigma$ , then  $\Gamma + \Gamma' \vdash \Sigma \triangleright e : T \triangleleft \Sigma'$ .*

*Proof.* The first item follows from the typing rules for values; the others require induction on typing derivations.

**Lemma 4 (Substitution).** *If  $\Gamma \vdash \Sigma \triangleright v : T \triangleleft \Sigma$  and  $\Gamma + x : T \vdash \Sigma \triangleright e : U \triangleleft \Sigma'$ , then  $\Gamma \vdash \Sigma \triangleright e\{v/x\} : U \triangleleft \Sigma'$ .*

*Proof.* By induction on the derivation of  $\Gamma + x : T \vdash \Sigma \triangleright e : U \triangleleft \Sigma'$ .

Our Subject Reduction theorem describes the evolution of the channel environment as a program is executed. The invariance of  $\Sigma'$  during reduction steps reflects the fact that  $\Sigma'$  is the final channel environment of a program.

**Theorem 1 (Subject Reduction).** *If  $\Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma'$  and  $\Sigma, e \longrightarrow \Sigma'', e'$ , then  $\Gamma \vdash \Sigma'' \triangleright e' : T \triangleleft \Sigma'$ .*

*Proof.* By induction on the derivation of  $\Sigma, e \longrightarrow \Sigma'', e'$  (which means induction on the structure of evaluation contexts  $E[\ ]$ ), first considering  $\Sigma, e \longrightarrow_v \Sigma'', e'$  by a simple case analysis.

## 7 Conclusions

We have transferred the concept of session types from the  $\pi$ -calculus to a language based on the  $\lambda$ -calculus. This is a first step towards the application of session types to the specification and verification of protocols implemented in mainstream programming languages. The main differences between the  $\pi$ -calculus with session types, and our language, are as follows. The operations on channels are now independent terms, rather than prefixes of processes, so we have introduced a new form of typing judgement which describes the effect of a term on the channel environment. We have separated creation and naming of channels, and because this introduces the possibility of aliasing, we represent the types of channels by indirection from the main type environment to the channel environment. We have introduced a new form of function type, which in addition to specifying parameter and result types, describes changes in the channel types of both parameters and free variables.

We have defined a static type system which guarantees that channels are not misused, even in the presence of aliasing, and we have proved this property with respect to a formal operational semantics of the language. We have therefore established a sound basis for the design of a more complete language with session types.

*Related Work.* In the *Vault* system [2] annotations are added to C programs, in order to describe protocols that a compiler can statically enforce. Similarly to our approach, individual runtime objects are tracked by associating keys (channels, in our terminology) with resources, and function types describe the effect of the function on the keys. Although incorporating a form of selection ( $\oplus$ ) the language of types seems quite simple in comparison with session types, and the annotations to C types can appear awkward.

A somewhat related line of research uses *flow-sensitive type qualifiers* to prove that programs access resources in a disciplined manner. Walker, Crary, and Morrisett [14] present a language to describe region-based memory management together with a provably safe type system. Igarashi and Kobayashi [7] present a general framework comprising a language with primitives for creating and accessing resources, and a type inference algorithm that checks whether programs are resource-safe. Although it might be possible to formulate operations on channels as resource use in these frameworks, our work focuses on a more specific problem: we transfer the concept of session types into a language which is closer to programming practice, and we do it purely with a straightforward type system.

*Type and effect* systems can be used to prove properties of protocols. Gordon and Jeffrey [4] use one such system to prove progress properties of communication protocols written in  $\pi$ -calculus. Rajamani *et al.*'s *Behave* [1, 9] uses CCS to describe properties of  $\pi$ -calculus programs, verified via a combination of type and model checking. In contrast, our system, while embodying more sophisticated protocols (using branch and select, for example), does not attempt to prove correctness of the contents of messages; only correctness of the types and sequence of messages.

*Future Work.* Finally, we outline some of the issues involved in extending our language to include a wider range of standard features.

Adding ML-style references and assignment seems straightforward, as we have already defined the infrastructure for typechecking in the presence of aliasing. It is important to dereference types sufficiently: we should have

$$x : \text{Ref Chan } c, y : \text{Chan } d \vdash c : T, d : T \triangleright x := y : \text{Unit} \triangleleft c : T, d : T$$

even though the types of  $x$  and  $y$  do not match directly.

Our function types have a limitation: we are not able to type functions which use free channel variables. We do not yet know how to generalize the rule T-ABS so that a channel identifier can enter the function type and also remain in the environment. This creates a problem with curried functions, and so we cannot

type this program:

```

fun sendTwice x y =
  send 1 on x
  send 2 on y

```

It might be possible to make this program typable by allowing multi-argument functions (via tuple types, for example). However, there are more serious implications for the use of session types in an object oriented language: it is typical and important for the methods of an object to use variables which, although bound in the object, are free in the method.

If we can extend our type system so that functions can use free channel variables, we would expect a function type to describe the effect of a function on the types of any free channel variables which it uses. This will induce constraints on the sequence of function calls within a program: calling a function might change the channel environment in such a way that the same function cannot be called again immediately. Now consider, in an object-oriented language with session types, an object containing a channel which its methods can use. The session type of the channel would constrain the use of the methods, perhaps imposing a particular sequence of method calls. This opens up the possibility of a connection with type systems for non-uniform objects [10]: maybe a non-uniform object type could be generated by the session type of a hidden channel. If we add objects to our language then it will also be interesting to investigate how the existing theory of subtyping for session types [3] relates to subclasses and inheritance.

Our language has no concurrency constructs, and our examples are individual components of a distributed system. By introducing a parallel composition operator we could describe complete systems, and strengthen our type safety results. This would involve transferring more of the machinery of session types from the  $\pi$ -calculus to our language. An alternative way of introducing concurrency within our present language would be to allow certain functions to execute as separate threads; we could then implement a multi-threaded server, for example. If  $f$   $c$  starts a new thread then from the caller's point of view,  $f$  returns immediately, but the caller cannot use  $c$ . So the type of  $f$  must be  $c : S \triangleright T \rightarrow \mathbf{Unit} \triangleleft c : \perp$ . More generally we could introduce a new typing rule for thread-starting (*runnable*, in Java's terminology) abstraction:

$$\frac{\Gamma + x : T \vdash \Sigma + \Sigma_1 \triangleright e : \mathbf{Unit} \triangleleft \Sigma_2 + \Sigma \quad \Sigma \vdash \Gamma \quad \Sigma_1 \vdash x : T}{\Gamma \vdash \Sigma \triangleright \lambda x. e : (\Sigma_1 \triangleright T \rightarrow \mathbf{Unit} \triangleleft \perp) \triangleleft \Sigma} \quad (\text{T-ABS'})$$

where  $\perp$  in the conclusion means  $\Sigma_1$  with all types replaced by  $\perp$ .

We would like to allow channels to be first-class values which can be sent and received. It should not be too difficult to transfer the necessary machinery from the theory of session types in the pi calculus. This would allow us to implement delegation of sessions, which is necessary in a multi-threaded server.

*Acknowledgements.* The research reported in this paper was supported by a Treaty of Windsor grant from the British Council in Portugal and the Portuguese

Council of University Rectors. Work partially supported by the EU FEDER, the EU projects of the proactive initiative FET-Global Computing (Mikado and Profundis), and the Portuguese Fundação para a Ciência e Tecnologia, namely via the project MIMO (POSI/CHS/39789/2001).

## References

1. S. Chaki, S. K. Rajamani, and J. Rehof. Types as models: model checking message-passing programs. In *Proceedings, 29th ACM Symposium on Principles of Programming Languages*, pages 45–57. ACM Press, 2002.
2. R. DeLine and M. Fähndrich. Enforcing high-level protocols in low-level software. In *Proceedings of the PLDI conference*, pages 59–69. ACM Press, 2001.
3. S. J. Gay and M. J. Hole. Types and subtypes for client-server interactions. In S. D. Swierstra, editor, *ESOP'99: Proceedings of the European Symposium on Programming Languages and Systems*, number 1576 in Lecture Notes in Computer Science, pages 74–90. Springer-Verlag, 1999.
4. A. Gordon and A. Jeffrey. Typing correspondence assertions for communication protocols. In Elsevier, editor, *17th Conference on the Mathematical Foundations of Programming Semantics*, number 45 in ENTCS, 2001.
5. M. J. Hole and S. J. Gay. Bounded polymorphism in session types, 2002. Submitted for publication.
6. K. Honda, V. Vasconcelos, and M. Kubo. Language primitives and type discipline for structured communication-based programming. In *Proceedings of the European Symposium on Programming*, Lecture Notes in Computer Science. Springer-Verlag, 1998.
7. A. Igarashi and N. Kobayashi. Resource usage analysis. In *Proceedings, 29th ACM Symposium on Principles of Programming Languages*, pages 331–342. ACM Press, 2002.
8. R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, I and II. *Information and Computation*, 100(1):1–77, September 1992.
9. S. K. Rajamani and J. Rehof. A behavioral module system for the pi-calculus. In *Proceedings of the Static Analysis Symposium*, volume 2126 of LNCS, pages 375–394. Springer-Verlag, 2001.
10. A. Ravara and V. T. Vasconcelos. Typing non-uniform concurrent objects. In *CONCUR'00*, volume 1877 of LNCS, pages 474–488. Springer-Verlag, 2000.
11. D. Sangiorgi and D. Walker. *The  $\pi$ -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
12. K. Takeuchi, K. Honda, and M. Kubo. An interaction-based language and its typing system. In *Proceedings of the 6th European Conference on Parallel Languages and Architectures*, number 817 in Lecture Notes in Computer Science. Springer-Verlag, 1994.
13. A. Vallecillo, V. T. Vasconcelos, and A. Ravara. Typing the behavior of objects and components using session types. In *1st International Workshop on Foundations of Coordination Languages and Software Architectures (Foclasa 2002)*, Electronic Notes in Theoretical Computer Science. Elsevier, August 2002.
14. D. Walker, K. Cray, and G. Morrisett. Typed memory management via static capabilities. *ACM Transactions on Programming Languages and Systems*, 22(4):701–771, 2000.