

Session Types for Functional Multithreading

Vasco Vasconcelos
António Ravara
Simon Gay

DI-FCUL

TR-05-9

May 20, 2005

Departamento de Informática
Faculdade de Ciências da Universidade de Lisboa
Campo Grande, 1749-016 Lisboa
Portugal

Technical reports are available at <http://www.di.fc.ul.pt/tech-reports>. The files are stored in PDF, with the report number as filename. Alternatively, reports are available by post from the above address.

Session Types for Functional Multithreading*

Vasco Vasconcelos[†] António Ravara[‡] Simon Gay[§]

May 20, 2005

Abstract

We define a language whose type system, incorporating session types, allows complex protocols to be specified by types and verified by static typechecking. A session type, associated with a communication channel, specifies the state transitions of a protocol and also the data types of messages associated with transitions; thus typechecking can verify both correctness of individual messages and correctness of sequences of transitions. Previously session types have mainly been studied in the context of the π -calculus; instead, our formulation is based on a multi-threaded functional language with side-effecting input/output operations. 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 channels. We formalize the syntax, semantics and typing rules of our language, and prove subject reduction and runtime type safety theorems.

Keywords: Session types, static typechecking, concurrent programming, specification of communication protocols.

1 Introduction

Communication in distributed systems is typically structured around protocols, which specify the sequence and form of messages passing over communication channels. Correctness of such systems implies that protocols are obeyed.

*A revised version of the paper in *Concur 2004*, volume 3170 of LNCS, pages 497–511, Springer-Verlag, 2004; includes proofs for the results. A forthcoming paper in *Theoretical Computer Science* describes a type inference algorithm for an explicitly typed version of the language.

[†]Departamento de Informática, Faculdade de Ciências, Universidade de Lisboa, 1749-016 Lisboa, Portugal. Email: <vv@di.fc.ul.pt>

[‡]CLC and Departamento de Matemática, Instituto Superior Técnico, 1049-001 Lisboa, Portugal. Email: <amar@math.ist.utl.pt>

[§]Department of Computing Science, University of Glasgow, Glasgow G12 8QQ, UK. Email: <simon@dcs.gla.ac.uk>

The theory of *session types* [9, 13, 14, 24] 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. Such a 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 does so in accordance with the protocol.

The theory of session types has been developed in the context of the π -calculus [18, 23], an idealized 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. [25] use session types to add behavioural information to the interfaces of CORBA objects, and use Gay and Hole's [9] theory of subtyping to formalize compatibility and substitutability of components, but they have not attempted to design a complete language.

The Vault [5] and Cyclone [12] languages extend C with facilities for safe control of stateful resources. In Cyclone, locks must be acquired and released; Vault goes further by allowing operations on a resource to be statically checked against an automaton which specifies valid transitions. In contrast, session types are specialized to communication channels as a particular kind of resource, but as a result they enable further typechecking in association with each state transition: typechecking verifies the types of individual messages, as well as verifying that a sequence of messages obeys a given protocol. (These languages are further discussed in Section 7.)

In previous work [8] we have presented a language supporting typed functional programming with inter-process communication channels, but we only considered individual processes in isolation. Here we address collections of functional threads communicating via channels. This formulation allows us to prove a runtime safety property: well-typed programs do not misuse channels.

By transferring the concept of session types from the π -calculus to a multi-threaded functional language with side-effecting input/output operations, we show that static checking of session types could be added to a language such as Concurrent ML [22], at least without imperative features. In particular we have addressed the key differences between a conventional programming style and the programming notation of the π -calculus:

- The operations on channels are 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 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.

The structure of the paper is as follows. In Section 2 we explain session types in connection with a progressively more sophisticated example. Sections 3, 4 and 5 define the syntax, operational semantics and type system of our language. In Section 6 we present the runtime safety result. In Sections 7 and 8 we discuss related and future work.

2 Session Types and the Maths Server

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 u ; we think of the client engaging in a *session* with the server, using the channel u for communication. In our language, the server looks like this:

$$\begin{aligned} \text{server } u = & \text{ let } x = \text{receive } u \text{ in} \\ & \text{ let } y = \text{receive } u \text{ in} \\ & \text{ send } x + y \text{ on } u \end{aligned}$$

or more concisely: $\text{send } ((\text{receive } u) + (\text{receive } u)) \text{ on } u$.

Interchanging $?$ and $!$ yields the type describing the client side of the protocol:

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

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

$$\begin{aligned} \text{client } u = & \text{ send } 2 \text{ on } u \\ & \text{ send } 3 \text{ on } u \\ & \text{ let } x = \text{receive } u \text{ in } \text{code} \end{aligned}$$

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:

$$\begin{aligned} \text{server } u = \text{case } u \text{ of } \{ \\ \quad \text{add} \Rightarrow \text{send } (\text{receive } u) + (\text{receive } u) \text{ on } u \\ \quad \text{neg} \Rightarrow \text{send } -(\text{receive } u) \text{ on } u \} \end{aligned}$$

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 client implementation makes a particular choice, for example:

$$\begin{aligned} \text{addClient } u = \text{select } \text{add} \text{ on } u \\ \quad \text{send } 2 \text{ on } u \\ \quad \text{send } 3 \text{ on } u \\ \quad \text{let } x = \text{receive } u \text{ in } \text{code} \\ \\ \text{negClient } u = \text{select } \text{neg} \text{ on } u \\ \quad \text{send } 7 \text{ on } u \\ \quad \text{let } x = \text{receive } u \text{ in } \text{code} \end{aligned}$$

Note that the type of the subsequent interaction depends on the label which is selected. 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.

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

This example shows that session types allow the description of protocols which cannot easily be accommodated with objects, that is, with sequences of the form: **select** a method; **send** the arguments; **receive** the result.

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:

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

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. To simplify the theory we decided not to include recursive types in this paper; the interested reader may refer to report [8].

Function Types. We have not mentioned the type of the *server* itself. Clearly, it accepts a channel and returns nothing. If *c* is the name of the channel, the input/output behaviour of the function is described by $\text{Chan } c \rightarrow \text{Unit}$. When control enters the function, channel *c* is in a state where it offers *add* and *neg* services. The function then “consumes” the channel, leaving it in a state ready to be closed. In order to correctly control channel usage, we annotate function types with the initial and the final type of all the channels used by the function. If *c* is the (runtime) channel denoted by the (program) variable *u*, we may assign the following type to *server*.

$$\text{server} :: (c: \&\langle \text{add}: \dots, \text{neg}: \dots \rangle; \text{Chan } c \rightarrow \text{Unit}; c: \text{End}) \\ \text{server } u = \text{case } u \text{ of } \{ \text{add} \Rightarrow \dots, \text{neg} \Rightarrow \dots \}$$

Note how the function type describes not only the type of the parameter and that of the result, but also its effect on channel *c*. It can also be useful

to send functions on channels. For example we could add the component¹

$$eval: ?(Int \rightarrow Bool).?Int.!Bool.End$$

to the branch type of the *server*, with corresponding server code, to be placed within the server's case above.

$$eval \Rightarrow \text{send } (receive\ u)(receive\ u) \text{ on } u$$

A client which requires a primality test service (perhaps the server has fast hardware) can be written as follows.

$$\begin{aligned} primeClient &:: (c: \oplus \langle add: \dots, neg: \dots, eval: \dots \rangle; Chan\ c \rightarrow Unit; c: End) \\ primeClient\ u &= \text{select } eval \text{ on } u \\ &\quad \text{send } isPrime \text{ on } u \\ &\quad \text{send } bigNumber \text{ on } u \\ &\quad \text{let } x = \text{receive } u \text{ in } code \end{aligned}$$

Establishing a Connection. How do the client and the server reach a state in which they both know about channel u ? We follow Takeuchi, Kubo and Honda [24], and propose a pair of constructs: *request* v for use by clients, and *accept* v for use by servers. In use, *request* and *accept* occur in separate threads, and interact with each other to create a new channel. The value v in both *request* and *accept*, denotes the common knowledge of the two threads: a *shared name* used solely for the creation of new channels. If S is the type of a channel, the type of a name used to create channels of type S is denoted by $[S]$. Functions *server* and *negClient* now receive a name of type $[\&\langle add: \dots, neg: \dots, eval: \dots \rangle]$, as shown in the following piece of code.

$$\begin{aligned} server &:: [\&\langle add: \dots, neg: \dots, eval: \dots \rangle] \rightarrow Unit \\ server\ x &= \text{let } u = \text{accept } x \text{ in } (\text{case } u \text{ of } \dots; \text{close } u) \\ negClient &:: [\&\langle add: \dots, neg: \dots, eval: \dots \rangle] \rightarrow Unit \\ negClient\ x &= \text{let } u = \text{request } x \text{ in } (\text{select } neg \text{ on } u \dots; \text{close } u) \end{aligned}$$

Note that the same type for the shared name x is used both for the server and for the client; it is the *accept/request* construct that distinguishes one from the other. This is also where we introduce the operation to close a channel: *accept/request* creates a channel; *close* destroys it.

¹We often omit the empty channel environment on each side of the arrow, so that $Int \rightarrow Bool$ is short for $\emptyset; Int \rightarrow Bool; \emptyset$.

Sharing Names. In order for a name to become known by a client and a server, it must be created somewhere and distributed to both. To create a new, potentially shared, name of type $[S]$, we write `new S`. To distribute it to a second thread, we `fork` a new thread, in whose code the name occurs.² Our complete system creates a name x and launches three threads (a server and two clients), all sharing the newly created name.

```

system :: Unit
system = let x = new &<add: ..., neg: ..., eval: ...> in
        fork negClient x; fork addClient x; fork server x

```

Given the above implementation of `server`, one of the clients will be forever requesting x . Fortunately, it is easy to extend the `server` to accept more than one connection in its life time.

```

server :: [&<add: ..., neg: ..., eval: ...>] → Unit
server x = let u = accept x in fork (case u of ...; close u)
server x

```

Sending Channels on Channels. Imagine two clients which need to cooperate in their interaction with the server: one client establishes a connection, selects the `neg` operation, and sends the argument; the second client receives the result. After selecting `neg` and sending the argument, the first client must provide the second with the channel to the server. In order to do so, both clients must share a name of type $?(?Int.End).End$ (called S below) and establish a connection for the sole purpose of transmitting the server channel.

```

askNeg :: [<add: ...>] → [S] → Unit      getNeg :: [S] → Unit
askNeg x y = let u = request x in        getNeg y = let w = accept y in
        select neg on u; send 7 on u      let u = receive w in
        let w = request y in             let i = receive u in
        send u on w; close w             close u; close w;
                                         code

```

It is instructive to follow the evolution of the state (the type) of channels c and d , connected to variables u and w , respectively. After the execution of the first line of `getNeg`, d has type $S = ?(?Int.End).End$; after the second line, d is reduced to `End`, but c shows up with type `?Int.End`; after the third line both channels are of type `End`, that is, ready to be closed. By the end of the fourth line, we gather no more information on channels c and d , for they are now closed. That is the sort of analysis our type system performs.

²Alternatively, we may send the name on an existing channel.

After sending a channel, no further interaction on the channel is possible. Note that *askNeg* cannot close *u*, for otherwise the channel's client side would be closed twice (in *askNeg* and in *getNeg*). On the other hand, channel *w* must be closed at both its ends, by *askNeg* and by *getNeg*.

The remainder of this section deals with further issues arising from the interaction between types and programming.

Channel Aliasing. As soon as we separate creation and naming of channels, aliasing becomes an issue. Consider the function below.

$$\mathit{sendSend} \ u \ v = \text{send } 1 \text{ on } u; \text{send } 2 \text{ on } v$$

Function *sendSend* can be used in a number of different ways, including the one where *u* and *v* become aliases for a single underlying channel.

$$\begin{aligned} \mathit{sendTwice} &:: c: !\text{Int}.\!\text{Int}.\text{End}; \text{Chan } c \rightarrow \text{Unit}; c: \text{End} \\ \mathit{sendTwice} \ w &= \mathit{sendSend} \ w \ w \end{aligned}$$

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 the body of function *sendSend*, the types of *u* and *v* are both *Chan c*. The state of *c*, initially *!Int.!Int.End*, is recorded separately.

Free Variables in Functions. If we write

$$\mathit{sendFree} \ v = \text{send } 1 \text{ on } u; \text{send } 2 \text{ on } v$$

then function *sendSend* becomes $\lambda u. \mathit{sendFree}$. In order to type *sendTwice*, thus effectively aliasing *u* and *v* in *sendSend*, we must have³

$$\begin{aligned} \mathit{sendFree} &:: c: !\text{Int}.\!\text{Int}.\text{End}; \text{Chan } c \rightarrow \text{Unit}; c: \text{End} \\ \mathit{sendSend} &:: c: !\text{Int}.\!\text{Int}.\text{End}; \text{Chan } c \rightarrow \text{Chan } c \rightarrow \text{Unit}; c: \text{End} \end{aligned}$$

in a typing environment associating the type *Chan c* to the free variable *u* of *sendFree*. However, if we do not want to alias *u* and *v*, then we must have

$$\begin{aligned} \mathit{sendFree} &:: c: !\text{Int}.\text{End}, d: !\text{Int}.\text{End}; \text{Chan } c \rightarrow \text{Unit}; c: \text{End}, d: \text{End} \\ \mathit{sendSend} &:: c: !\text{Int}.\text{End}, d: !\text{Int}.\text{End}; \text{Chan } c \rightarrow \text{Chan } d \rightarrow \text{Unit}; c: \text{End}, d: \text{End} \end{aligned}$$

in a typing environment containing *u: Chan d*. Note how the above type for *sendFree* captures changes to channels that are parameters (*c*) and to channels that occur free (*d*).

³We abbreviate $\Sigma; T \rightarrow (\Sigma; U \rightarrow V; \Sigma'); \Sigma' \rightarrow \Sigma; T \rightarrow U \rightarrow V; \Sigma'$.

$$\begin{aligned}
v &::= c \mid n \mid x \mid \lambda x.e \mid \text{rec } x.v \mid \text{true} \mid \text{false} \mid \text{unit} \\
e &::= t \mid vv \mid \text{if } v \text{ then } e \text{ else } e \mid \text{new} \mid \text{accept } v \mid \text{request } v \mid \\
&\quad \text{send } v \text{ on } v \mid \text{receive } v \mid \text{case } v \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} \mid \text{select } l \text{ on } v \mid \text{close } v \\
t &::= v \mid \text{let } x = e \text{ in } t \mid \text{fork } t; t \\
C &::= \langle t \rangle \mid (C \mid C) \mid (\nu n)C \mid (\nu c)C
\end{aligned}$$

Figure 1: Syntax of values, expressions, threads and configurations

Polymorphism. We have seen that *sendFree* admits at least two different types. In order to allow for code reuse we type our let-bound values as many times as needed, potentially with different types. The paragraph above showed a *share/not-share* kind of polymorphism. Other forms include *channel polymorphism* and *session polymorphism*. For an example of channel polymorphism, consider

$$\begin{aligned}
\text{sendTwiceSendTwice} &:: c: S, d: S; \text{Chan } c \rightarrow \text{Chan } d \rightarrow \text{Unit}; c: \text{End}, d: \text{End} \\
\text{sendTwiceSendTwice } x \ y &= \text{sendTwice } x; \text{sendTwice } y
\end{aligned}$$

where S is $!\text{Int}!\text{Int}.\text{End}$. Here *sendTwice* must be typed once with channel c , and another with channel d . For an example of session polymorphism, we have:

$$\begin{aligned}
\text{sendQuad} &:: c: !\text{Int}!\text{Int}!\text{Int}!\text{Int}.\text{End}; \text{Chan } c \rightarrow \text{Unit}; c: \text{End} \\
\text{sendQuad } x &= \text{sendTwice } x; \text{sendTwice } x
\end{aligned}$$

where *sendTwice* must be typed once with $c: !\text{Int}!\text{Int}!\text{Int}!\text{Int}.\text{End}$, and a second time with $c: !\text{Int}!\text{Int}.\text{End}$.

3 Syntax

Most of the syntax of our language has been illustrated in the previous section; here we define it formally by the grammar in Figure 1 (cf. [17]).

We use *channel identifiers* c, \dots , *name identifiers* n, \dots , *term variables* x, \dots , and *labels* l, \dots , and define *values* v , *expressions* e , *threads* t , and *configurations* C . To simplify some definitions, we use (νa) to stand for either (νc) or (νn) .

Channel identifiers and name identifiers are not available in the top-level syntax of threads; they arise only during reduction, in a *request/accept* synchronization, and in a *new* operation, respectively, as described in Section 4.

In Section 2 we used several derived constructors. An expression $e; t$ (sometimes implied in our examples by the indentation) is an abbreviation for *let* $y = e$ in t , provided y does not occur free in t . Idioms like

$(C, , \langle \text{unit} \rangle)$ is a commutative monoid	(S-MONOID)
$(\nu n)C_1 C_2 \equiv (\nu n)(C_1 C_2)$ if n not free in C_2	(S-SCOPEN)
$(\nu c)C_1 C_2 \equiv (\nu c)(C_1 C_2)$ if c not free in C_2	(S-SCOPEC)

Figure 2: Structural congruence

send (receive c)(receive c) on c need appropriate de-sugaring into consecutive lets, making the evaluation order explicit. We sometimes “terminate” threads with an expression rather than a value: a thread e is short for let $x = e$ in x . Recursive function definitions must be made explicit with `rec`.

4 Operational Semantics

The binding occurrences are x in $\lambda x.e$, `rec` $x.e$, let $x = e$ in t , n in $(\nu n)C$ and c in $(\nu c)C$. Free and bound identifiers are defined as usual and we work up to α -equivalence. Substitution, of values for variables, is defined as expected. We define a reduction semantics on configurations (Figure 3), making use of a simple structural congruence relation [18] (Figure 2), allowing for the rearrangement of threads in a configuration, so that reduction may happen.⁴

We now explain the reduction rules. R-INIT synchronizes two threads on a *shared name* n , creating a new channel c known to both threads. Rules R-COM, R-BRANCH, and R-CLOSE synchronize two threads on a *channel* c : R-COM transmits a value v from one thread to the other; R-BRANCH, rather than transmitting a value, chooses one of the branches in the `case` thread; and R-CLOSE closes a channel in *both* threads simultaneously. R-NEW creates a new name n , and records the fact that the name is potentially shared, by means of a (νn) in the resulting configuration. The last four rules allow reduction to happen underneath restriction, parallel composition, and structural congruence.

Unlike other thread models, the value a thread reduces to is not communicated back to its parent thread (the one that forked the terminating thread). Such behaviour would have to be programmed by arranging for both threads to share a channel and explicitly **sending** the result back to the parent.

⁴We could easily arrange for structural congruence to garbage collect all threads of the form $\langle v \rangle$, for v closed.

$$\begin{array}{l}
\langle \text{let } x = \text{request } n \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{accept } n \text{ in } t_2 \rangle \rightarrow \\
\quad (\nu c)(\langle \text{let } x = c \text{ in } t_1 \rangle \mid \langle \text{let } y = c \text{ in } t_2 \rangle) \quad (\text{R-INIT}) \\
\langle \text{let } x = \text{receive } c \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{send } v \text{ on } c \text{ in } t_2 \rangle \rightarrow \\
\quad \langle \text{let } x = v \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle \quad (\text{R-COM}) \\
\langle \text{let } x = \text{case } c \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{select } l_j \text{ on } c \text{ in } t_2 \rangle \rightarrow \\
\quad \langle \text{let } x = e_j \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle \quad (\text{R-BRANCH}) \\
\langle \text{let } x = \text{close } c \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{close } c \text{ in } t_2 \rangle \rightarrow \\
\quad \langle \text{let } x = \text{unit in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle \quad (\text{R-CLOSE}) \\
\langle \text{let } x = \text{new in } t \rangle \rightarrow (\nu n)\langle \text{let } x = n \text{ in } t \rangle \quad (\text{R-NEW}) \\
\langle \text{fork } t_1; t_2 \rangle \rightarrow \langle t_1 \rangle \mid \langle t_2 \rangle \quad (\text{R-FORK}) \\
\langle \text{let } x = \text{if true then } e \text{ else } e' \text{ in } t \rangle \rightarrow \langle \text{let } x = e \text{ in } t \rangle \quad (\text{R-IFT}) \\
\langle \text{let } x = \text{if false then } e \text{ else } e' \text{ in } t \rangle \rightarrow \langle \text{let } x = e' \text{ in } t \rangle \quad (\text{R-IFF}) \\
\langle \text{let } x = (\lambda y. e)v \text{ in } t \rangle \rightarrow \langle \text{let } x = e\{v/y\} \text{ in } t \rangle \quad (\text{R-APP}) \\
\langle \text{let } x = (\text{rec } y.v)u \text{ in } t \rangle \rightarrow \langle \text{let } x = (v\{\text{rec } y.v/y\})u \text{ in } t \rangle \quad (\text{R-REC}) \\
\langle \text{let } x = (\text{let } y = e \text{ in } t') \text{ in } t \rangle \rightarrow \langle \text{let } y = e \text{ in } (\text{let } x = t' \text{ in } t) \rangle \quad (\text{R-LET}) \\
\langle \text{let } x = v \text{ in } t \rangle \rightarrow \langle t\{v/x\} \rangle \quad (\text{R-BETA}) \\
\frac{C \rightarrow C'}{(\nu a)C \rightarrow (\nu a)C'} \quad \frac{C \rightarrow C'}{C \mid C'' \rightarrow C' \mid C''} \quad \frac{C \equiv \rightarrow \equiv C'}{C \rightarrow C'} \quad (\text{R-CONF})
\end{array}$$

In R-INIT, c is not free in t_1, t_2 ; in R-NEW, n is not free in t .

Figure 3: Reduction rules

5 Typing

The syntax of types is described in Figure 4. We define *session types* S , *channel environments* Σ , *data types* D , and *term types* T . The type $\text{Chan } c$ represents the type of the channel with identity c ; the session type associated with c is recorded separately in a channel environment Σ . Channel type *bottom*, \perp , denotes a channel that is already in use by two threads, hence that cannot be used further. Similarly to channel and name identifiers, \perp is not available at the top level syntax, arising only via the channel environment composition operator, $\Sigma_1 \bullet \Sigma_2$, defined below. Among datatypes we have channel-state annotated functional types $\Sigma; T \rightarrow T; \Sigma$, and types for names $[S]$ capable of establishing sessions of type S .

The type system is presented in Figures 5 to 9. Typing judgements for constants are of the form $\Gamma \vdash v : T$, where Γ is a map from names and variables to types. Value judgements do not mention channel environments, for values, having no behaviour, do not change channels. Judgements for

$$\begin{aligned}
S &::= ?D.S \mid !D.S \mid ?S.S \mid !S.S \mid \&\langle l_i: S_i \rangle_{i \in I} \mid \oplus \langle l_i: S_i \rangle_{i \in I} \mid \text{End} \mid \perp \\
D &::= \text{Bool} \mid \text{Unit} \mid \Sigma; T \rightarrow T; \Sigma \mid [S] \\
\Sigma &::= \emptyset \mid \Sigma, c: S \quad (c: S \text{ not in } \Sigma) \\
T &::= D \mid \text{Chan } c
\end{aligned}$$

Figure 4: Syntax of types

$$\begin{array}{c}
\Gamma \vdash \text{true}: \text{Bool} \quad \Gamma \vdash \text{false}: \text{Bool} \quad \Gamma \vdash \text{unit}: \text{Unit} \quad (\text{T-CONST}) \\
\Gamma \vdash c: \text{Chan } c \quad \Gamma, n: [S] \vdash n: [S] \quad \Gamma, x: T \vdash x: T \\
\hspace{15em} (\text{T-CHAN, T-NAME, T-VAR}) \\
\frac{\Gamma, x: T \vdash \Sigma \triangleright e: U \triangleleft \Sigma'}{\Gamma \vdash \lambda x.e: (\Sigma; T \rightarrow U; \Sigma')} \quad \frac{\Gamma, x: T \vdash v: T \quad T = (\Sigma; U \rightarrow U'; \Sigma')}{\Gamma \vdash \text{rec } x.v: T} \\
\hspace{15em} (\text{T-ABS, T-REC})
\end{array}$$

Figure 5: Typing rules for values

expressions are of the form $\Gamma \vdash \Sigma \triangleright e: T \triangleleft \Sigma'$, where Σ is a channel environment (a map from channels into sorts, as in Figure 4). The difference between Σ and Σ' reflects the effect of an expression on the types of channels, for example

$$x: \text{Chan } c \vdash c: ?\text{Int.End} \triangleright \text{receive } x: \text{Int} \triangleleft c: \text{End}.$$

Finally, typing judgements for configurations are of the form $\Delta \vdash \Sigma \triangleright C$ where Δ is a map from names to datatypes of the form $[S]$.

Typing Values (Figure 5). T-CHAN says that a channel named c has type $\text{Chan } c$. The actual type (or state) of channel c is to be found in a channel environment Σ , in the rules for expressions. In T-ABS, the initial and final channel environments of the function body are recorded in the function type.

Typing Expressions (Figures 6, 7). There are two rules for **receive** and two rules for **send**, for these constructors are overloaded: they allow transmission of data as well as channels. In T-RECEIVED, the prefix $?D.$, of the type for channel c , is consumed, provided that we are receiving on a value aliased to channel c (of type $\text{Chan } c$). In T-RECEIVES, we receive a channel, that we decided to call d ; the type of the expression is $\text{Chan } d$, and we add a new entry to the final channel environment, where we record the type for d . The particular form of the final channel environment allows

$$\begin{array}{c}
\frac{\Gamma \vdash v : \text{Chan } c}{\Gamma \vdash \Sigma, c : ?D.S \triangleright \text{receive } v : D \triangleleft \Sigma, c : S} \quad (\text{T-RECEIVED}) \\
\frac{\Gamma \vdash v : \text{Chan } c \quad d \text{ fresh}}{\Gamma \vdash \Sigma, c : ?S'.S \triangleright \text{receive } v : \text{Chan } d \triangleleft \Sigma, d : S', c : S} \quad (\text{T-RECEIVES}) \\
\frac{\Gamma \vdash v : D \quad \Gamma \vdash v' : \text{Chan } c}{\Gamma \vdash \Sigma, c : !D.S \triangleright \text{send } v \text{ on } v' : \text{Unit} \triangleleft \Sigma, c : S} \quad (\text{T-SEND D}) \\
\frac{\Gamma \vdash v : \text{Chan } d \quad \Gamma \vdash v' : \text{Chan } c}{\Gamma \vdash \Sigma, c : !S'.S, d : S' \triangleright \text{send } v \text{ on } v' : \text{Unit} \triangleleft \Sigma, c : S} \quad (\text{T-SENDS}) \\
\frac{\Gamma \vdash v : \text{Chan } c \quad j \in I}{\Gamma \vdash \Sigma, c : \oplus \langle l_i : S_i \rangle_{i \in I} \triangleright \text{select } l_j \text{ on } v : \text{Unit} \triangleleft \Sigma, c : S_j} \quad (\text{T-SELECT}) \\
\frac{\Gamma \vdash v : \text{Chan } c \quad \Gamma \vdash \Sigma, c : S_j \triangleright e_j : T \triangleleft \Sigma' \quad \forall j \in I}{\Gamma \vdash \Sigma, c : \& \langle l_i : S_i \rangle_{i \in I} \triangleright \text{case } v \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} : T \triangleleft \Sigma'} \quad (\text{T-CASE}) \\
\frac{\Gamma \vdash v : \text{Chan } c}{\Gamma \vdash \Sigma, c : \text{End} \triangleright \text{close } v : \text{Unit} \triangleleft \Sigma} \quad (\text{T-CLOSE}) \\
\frac{\Gamma \vdash v : [S] \quad c \text{ fresh}}{\Gamma \vdash \Sigma \triangleright \text{request } v : \text{Chan } c \triangleleft \Sigma, c : \bar{S}} \quad \frac{\Gamma \vdash v : [S] \quad c \text{ fresh}}{\Gamma \vdash \Sigma \triangleright \text{accept } v : \text{Chan } c \triangleleft \Sigma, c : S} \\
\quad (\text{T-REQUEST, T-ACCEPT})
\end{array}$$

Figure 6: Typing rules for expressions I: Channel operations

the continuation to hold both ends of the channel. The rules T-SEND D and T-SENDS, for sending values and channels, are similar. In T-SELECT, the type for c in the final channel environment is that of branch l_i in the type for c in the source channel environment. 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. Rule T-CLOSE requires that the channel must be ready to be closed (of type `End`). We remove the closed channel from the environment.

Rules T-REQUEST and T-ACCEPT both introduce a new channel c in the channel environment, of *dual polarities* [9, 13, 14, 24, 25]. The dual of a session type S , denoted \bar{S} , is defined for all session types except \perp , and is obtained by interchanging output $!$ and input $?$, and by interchanging branching $\&$ and selection \oplus , and leaving S otherwise unchanged. The inductive definition of duality is in Figure 8.

In T-APP, the initial and final channel environments in the type of the function are released into the typing for the application. T-VAL says that constants do not affect the state of channels. Expression `new` has any type of the form $[S]$, denoting a name that, when shared by two threads, is able to produce (via `accept/request`) new channels of type S .

$$\begin{array}{c}
\frac{\Gamma \vdash \Sigma_1 \triangleright t_1 : T_1 \triangleleft \vec{c} : \perp \quad \Gamma \vdash \Sigma_2 \triangleright t_2 : T_2 \triangleleft \vec{d} : \perp}{\Gamma \vdash \Sigma_1 \bullet \Sigma_2 \triangleright \text{fork } t_1; t_2 : T_2 \triangleleft \vec{c} : \perp, \vec{d} : \perp} \quad (\text{T-FORK}) \\
\Gamma \vdash \Sigma \triangleright \text{new} : [S] \triangleleft \Sigma \quad \frac{\Gamma \vdash v : (\Sigma; T \rightarrow U; \Sigma') \quad \Gamma \vdash v' : T}{\Gamma \vdash \Sigma \triangleright vv' : U \triangleleft \Sigma'} \quad (\text{T-NEW, T-APP}) \\
\frac{\Gamma \vdash v : T}{\Gamma \vdash \Sigma \triangleright v : T \triangleleft \Sigma} \quad \frac{\Gamma \vdash v : \text{Bool} \quad \Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma' \quad \Gamma \vdash \Sigma \triangleright e' : T \triangleleft \Sigma'}{\Gamma \vdash \Sigma \triangleright \text{if } v \text{ then } e \text{ else } e' : T \triangleleft \Sigma'} \quad (\text{T-VAL, T-IF}) \\
\frac{\Gamma \vdash \Sigma \triangleright e : T \triangleleft \Sigma'' \quad \Gamma, x : T \vdash \Sigma'' \triangleright t : U \triangleleft \Sigma'}{\Gamma \vdash \Sigma \triangleright \text{let } x = e \text{ in } t : U \triangleleft \Sigma'} \quad \frac{\Gamma \vdash \Sigma \triangleright t\{v/x\} : T \triangleleft \Sigma'}{\Gamma \vdash \Sigma \triangleright \text{let } x = v \text{ in } t : T \triangleleft \Sigma'} \quad (\text{T-LET, T-POLYLET})
\end{array}$$

Figure 7: Typing rules for expressions II: Other rules

$$\begin{array}{c}
\overline{\text{End}} = \text{End} \quad \overline{?D.\bar{S}} = !D.\bar{S} \quad \overline{?S'.\bar{S}} = !S'.\bar{S} \quad \overline{!D.\bar{S}} = ?D.\bar{S} \quad \overline{!S'.\bar{S}} = ?S'.\bar{S} \\
\overline{\&\langle l_i : S_i \rangle_{i \in I}} = \oplus \langle l_i : \bar{S}_i \rangle_{i \in I} \quad \overline{\oplus \langle l_i : S_i \rangle_{i \in I}} = \&\langle l_i : \bar{S}_i \rangle_{i \in I}
\end{array}$$

Figure 8: Duality on session types

Rule T-FORK composes the initial channel environments of two configurations, by checking that the types of the channels occurring in both environments are dual. As for the final environment, the rule requires, via the $\vec{c} : \perp$ and $\vec{d} : \perp$ in the antecedent, that each thread involved either consumes their channels (that is sends or closes), or uses them in dual mode.

The *composition of two channel environments*, $\Sigma_1 \bullet \Sigma_2$, is defined only when $\Sigma_1(c) = \overline{\Sigma_2(c)}$, for all $c \in \text{dom } \Sigma_1 \cap \text{dom } \Sigma_2$. In this case $\text{dom}(\Sigma_1 \bullet \Sigma_2) = \text{dom } \Sigma_1 \cup \text{dom } \Sigma_2$, and $(\Sigma_1 \bullet \Sigma_2)(c)$ is \perp when $c \in \text{dom } \Sigma_1 \cap \text{dom } \Sigma_2$, and is $\Sigma_i(c)$ when $c \in \text{dom } \Sigma_i \setminus \text{dom } \Sigma_{3-i}$, for $i = 1, 2$.

Rule T-POLYLET types the various forms of polymorphism identified in Section 2, by separately typing different copies of the polymorphic value [20, Chapter 22].

Typing Configurations (Figure 9). Rule T-THREAD requires that threads either consume their channels or use them in dual mode, similarly to T-FORK. The Δ in the antecedent of rule T-THREAD ensures that threads are closed for variables, for the domain of Δ does not include variables. Rule T-PAR is similar to T-FORK. T-NEWN discards information on the bound name. There are two rules for channel creation. Rule T-NEWB says that a newly created channel must be used with dual modes by exactly two

$$\begin{array}{c}
\frac{\Delta \vdash \Sigma \triangleright t: _ \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle t \rangle} \quad \frac{\Delta \vdash \Sigma_1 \triangleright C_1 \quad \Delta \vdash \Sigma_2 \triangleright C_2}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright C_1 \mid C_2} \quad \frac{\Delta, n: [_] \vdash \Sigma \triangleright C}{\Delta \vdash \Sigma \triangleright (\nu n)C} \\
\text{(T-THREAD, T-PAR, T-NEWN)} \\
\frac{\Delta \vdash \Sigma, c: \perp \triangleright C \quad c \text{ not in } \Delta, \Sigma}{\Delta \vdash \Sigma \triangleright (\nu c)C} \quad \frac{\Delta \vdash \Sigma \triangleright C \quad c \text{ not in } \Delta, \Sigma}{\Delta \vdash \Sigma \triangleright (\nu c)C} \\
\text{(T-NEWB, T-NEWC)}
\end{array}$$

Figure 9: Typing rules for configurations

threads, since the type \perp usually arises from the \bullet operator in rules T-PAR or T-FORK. Rule T-NEWC allows to garbage collect unused channels.

The formulation of Subject Reduction is standard; the proof is in Appendix A, page 21.

Theorem 1 (Subject Reduction). *If $\Delta \vdash \Sigma \triangleright C$ and $C \rightarrow C'$, then $\Delta \vdash \Sigma \triangleright C'$.*

6 Type Safety

In our language of functional communicating threads different sorts of problems may occur at runtime, ranging from the traditional error of testing, in a conditional expression, a value that is not true or false; through applying close to a value that is not a channel; to the most relevant to our work: having one thread trying to send on a given channel, and another trying to select on the same channel, or having three or more threads trying to synchronize on the same channel.

In order to define what we mean by a faulty configuration, we start by calling a *c-thread* any thread ready to perform an operation on channel c , that is a thread of the form $\langle \text{let } x = \text{receive } c \text{ in } t \rangle$, and similarly for send, case, select, and close. A *c-redex* is the parallel composition of two threads ready to communicate on channel c , that is $\langle \text{let } x = \text{send } v \text{ on } c \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{receive } c \text{ in } t_2 \rangle$, and similarly for case/select, close/close. A *configuration C is faulty* when $C \equiv (\nu \vec{a})(C_1 \mid C_2)$ and C_1 is

1. the thread $\langle \text{let } x = e \text{ in } t \rangle$, where e is i) if v then $_$ else $_$ with $v \neq \text{true, false}$, or is ii) $v _$ with $v \neq \lambda y. e'$ and $v \neq \text{rec } y. e'$; or is
2. the thread $\langle \text{let } x = \text{accept/request } v \text{ in } t \rangle$, where v is not a name; or is
3. the thread $\langle \text{let } x = e \text{ in } t \rangle$, where e is i) receive/close v , or ii) send $_$ on v , or iii) case v of $_$, or iv) select $_$ on v , with v not a channel; or is
4. the parallel composition of two c -threads that do not form a c -redex; or is

5. the parallel composition of three or more c -threads.

The main property of this section says that typable configurations are not faulty; the proof is in Appendix B, page 31.

Theorem 2 (Type Safety). *Typable configurations are not faulty.*

7 Related Work

Cyclone [12] is a C-like type-safe polymorphic imperative language. It features region-based memory management, and more recently threads and locks [11], via a sophisticated type system. The multithreaded version requires “a lock name for every pointer and lock type, and an effect for every function”. Our locks are channels; but more than mutual exclusion, channels also allow a precise description of the protocol “between” acquiring and releasing the lock. In Cyclone a thread acquires a lock for a resource, uses the resource in whichever way it needs, and then releases the lock. Using our language a thread acquires the lock via a `request` operation, and then follows the protocol for the resource, before closing the channel obtained with `request`.

In the *Vault* system [5] 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 type system describes protocols in less detail than we can achieve with session types. “Adoption and Focus” [7], by the same authors, is a type system able to track changes in the state of objects; the system handles aliasing, and includes a form of polymorphism in functions. In contrast, our system checks the types of individual messages, as well as tracking the state of the channel. Our system is more specialized, but the specialization allows more type checking in the situation that we handle.

Type and effect systems can be used to prove properties of protocols. Gordon and Jeffrey [10] use one such system to prove progress properties of communication protocols written in π -calculus. Bonelli, Compagnoni, and Gunter [2, 3] combine the language of Honda, Vasconcelos and Kubo [14] with the correspondence assertions of Gordon and Jeffrey, thus obtaining a setting where further properties can be proved about programs. Adding correspondence assertions to session types increases the expressiveness of the system in two ways. Although session types only specify the structure of interactions between pairs of participants of a possibly multiparty protocol, the new setting makes it possible to specify and check that the interactions between participants in different pairs respect the overall protocol. Furthermore, the integrity and correct propagation of data is also verifiable.

However, this is a different kind of extension of session types than our work; their language does not include function types.

Rajamani *et al.*'s *Behave* [4, 21] uses CCS to describe properties of π -calculus programs, verified via a combination of type and model checking. Since our system is purely type checking (not model checking) we expect verification to be more efficient and easier to implement. Igarashi and Kobayashi have developed a generic framework in which a range of π -calculus type systems can be defined [16]. Although able to express sequencing of input and output types similarly to session types, it cannot express branching types.

A somewhat related line of research addresses resource access. Walker, Crary, and Morrisett [26] present a language to describe region-based memory management together with a provably safe type system. Igarashi and Kobayashi [15] present a general framework comprising a language with primitives for creating and accessing resources, and a type inference algorithm that checks whether programs access resources in a disciplined manner. Although types for resources in this latter work are similar in spirit to session types, we work in a much simpler setting.

Neubauer and Thiemann encoded a version of session types in the Haskell programming language, and proved that the embedding preserves typings [19], but the results are limited to type soundness.

Very recently, Dezani-Ciancaglini *et al.* [6] have proposed a minimal distributed object-oriented language with session types. Apart from the use of objects, the main difference between their version of session types and ours seems to be that they do not allow channels to be sent along channels (although objects containing *names* of potential channels may be sent). A more detailed comparison is a subject for future work.

8 Future Work

We outline some of the issues involved in extending our language to include a wider range of standard features.

Recursive Session Types. We have introduced recursive session types in a previous work [8]. We feel its incorporation in the present setting would not present major difficulties, although care must be taken in the definition of duality [25].

Principal Typings. For practical type inference, for separate compilation and modularity, one needs a notion of principal typings for the language. Particularly challenging is the share/not-share kind of polymorphism identified in section 2.

Type Inference. We are working on a constraint-based type inference algorithm for (the monomorphic fragment of) the language.

ML-style references and assignment. This would introduce further issues of aliasing. We do not yet know whether our present infrastructure for typechecking in the presence of aliasing would be sufficient for this extension.

Acknowledgements. This work was partially supported by the EU via FEDER and the IST proactive initiative FET-Global Computing (projects Mikado, IST-2001-32222, and Profundis, IST-2001-33100), by the Fundação para a Ciência e a Tecnologia (via CLC, CITI, and the project MIMO, POSI/CHS/39789/2001), and by a Treaty of Windsor grant from the British Council in Portugal and the Portuguese Council of University Rectors.

References

- [1] H.P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North-Holland, Amsterdam, 1984.
- [2] Eduardo Bonelli, Adriana Compagnoni, and Elsa Gunter. Correspondence assertions for process synchronization in concurrent communication. *Journal of Functional Programming*, 15(2):219–247, 2005.
- [3] Eduardo Bonelli, Adriana Compagnoni, and Elsa Gunter. Typechecking safe process synchronization. In *Workshop on the Foundations of Global Ubiquitous Computing*, Electronic Notes on Theoretical Computer Science. Elsevier Science, 2005. To appear.
- [4] Sagar Chaki, Sriram K. Rajamani, and Jakob Rehof. Types as models: model checking message-passing programs. In *POPL*, pages 45–57. ACM Press, 2002.
- [5] Robert DeLine and Manuel Fähndrich. Enforcing high-level protocols in low-level software. In *PLDI*, pages 59–69. ACM Press, 2001.
- [6] Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, Alexander Ahern, and Sophia Drossopolou. A distributed object-oriented language with session types. In *Symposium on Trustworthy Global Computing*. Springer-Verlag, 2005. To appear.
- [7] Jeffrey S. Foster, Tachio Terauchi, and Alex Aiken. Flow-Sensitive Type Qualifiers. In *PLDI*, pages 1–12, Berlin, Germany, June 2002.
- [8] Simon Gay, Vasco T. Vasconcelos, and António Ravara. Session types for inter-process communication. TR 2003-133, Department of Computing, University of Glasgow, March 2003.

- [9] Simon J. Gay and Malcolm J. Hole. Types and subtypes for client-server interactions. In S. Doaitse Swierstra, editor, *ESOP'99*, volume 1576 of *LNCS*, pages 74–90. Springer-Verlag, 1999.
- [10] A.D. Gordon and A.S.A. Jeffrey. Typing correspondence assertions for communication protocols. *Theoretical Computer Science*, 300:379–409, 2003.
- [11] Dan Grossman. Type-safe multithreading in cyclone. In *ACM Workshop on Types in Language Design and Implementation*. ACM Press, 2003.
- [12] Dan Grossman, Greg Morrisett, Trevor Jim, Michael Hicks, Yanling Wang, and James Cheney. Region-based memory management in cyclone. In *PLDI*, pages 282–293. ACM Press, 2002.
- [13] K. Honda. Types for dyadic interaction. In *CONCUR'93*, volume 715 of *LNCS*, pages 509–523. Springer-Verlag, 1993.
- [14] Kohei Honda, Vasco Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In Chris Hankin, editor, *ESOP'98*, volume 1381 of *LNCS*, pages 122–138. Springer-Verlag, 1998.
- [15] Atsushi Igarashi and Naoki Kobayashi. Resource usage analysis. In *POPL*, pages 331–342. ACM Press, 2002.
- [16] Atsushi Igarashi and Naoki Kobayashi. A generic type system for the pi-calculus. *Theoretical Computer Science*, 311(1–3):121–163, 2003.
- [17] Alan Jeffrey and Julian Rathke. A fully abstract may testing semantics for concurrent objects. In *Proc. Lics2002, 17th Annual Symposium on Logic in Computer Science, Copenhagen*, pages 101–112. IEEE Computer Society Press, 2002.
- [18] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I and II. *Information and Computation*, 100(1):1–77, September 1992.
- [19] Matthias Neubauer and Peter Thiemann. An implementation of session types. In *Practical Aspects of Declarative Languages*, volume 3057 of *Lecture Notes in Computer Science*, pages 56–70. Springer-Verlag, 2004.
- [20] B. Pierce. *Types and Programming Languages*. The MIT Press, 2002.
- [21] Sriram K. Rajamani and Jakob Rehof. A behavioral module system for the pi-calculus. In Patrick Cousot, editor, *Static Analysis: 8th International Symposium, SAS 2001*, volume 2126 of *LNCS*, pages 375–394. Springer-Verlag, 2001.

- [22] J.H. Reppy. CML: a higher order concurrent language. In *PLDI*, pages 293–305. ACM Press, 1991.
- [23] Davide Sangiorgi and David Walker. *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
- [24] Kaku Takeuchi, Kohei Honda, and Makoto Kubo. An interaction-based language and its typing system. In Constantine Halatsis, Dimitris G. Maritsas, George Philokyprou, and Sergios Theodoridis, editors, *PARLE '94*, volume 817 of *LNCS*. Springer-Verlag, 1994.
- [25] Antonio Vallecillo, Vasco T. Vasconcelos, and António Ravara. Typing the behavior of objects and components using session types. In *FOCLASA 2002*, volume 68 of *Electronic Notes in Theoretical Computer Science*. Elsevier, August 2002.
- [26] David Walker, Karl Cray, and Greg Morrisett. Typed memory management via static capabilities. *TOPLAS*, 22(4):701–771, 2000.

A Proof of Theorem 1, Subject Reduction

We start with a few auxiliary results; the proof of Subject Reduction is on page 25. To simplify the proofs, we make use of the *variable convention* [1], allowing, for example, to assume that, in sequent $\Delta \vdash \Sigma \triangleright (\nu c)C$, channel c does not occur in either Δ or Σ . Relatedly, when we say that c does not occur in C , we mean that it does not occur free in C and, by the variable convention, that it does not occur bound either.

The following easy results allow to grow and shrink the variable environment of an expression. Weakening is used in Subject Reduction (rule R-LET) and narrowing in the Substitution Lemma 12.⁵

Lemma 3 (Variable Weakening). *Suppose that x does not occur in e, v .*

1. *If $\Gamma \vdash \Sigma \triangleright e: U \triangleleft \Sigma'$, then $\Gamma, x: T \vdash \Sigma \triangleright e: U \triangleleft \Sigma'$.*
2. *If $\Gamma \vdash v: U$, then $\Gamma, x: T \vdash v: U$.*

Proof. The proofs, by mutual induction on the derivation of the judgements, are straightforward. \square

Lemma 4 (Variable Narrowing). *Suppose that x does not occur in e, v .*

1. *If $\Gamma, x: T \vdash \Sigma \triangleright e: U \triangleleft \Sigma'$, then $\Gamma \vdash \Sigma \triangleright e: U \triangleleft \Sigma'$.*
2. *If $\Gamma, x: T \vdash v: U$, then $\Gamma \vdash v: U$.*

Proof. The proofs follow the pattern of the ones above. \square

The following two unchallenging results allow to grow and shrink, this time, the name environment of a configuration. They are both used in the proofs of Subject Congruence (rule S-SCOPEN), and Weakening is also used in the proofs of Subject Reduction (rule R-NEW).

Lemma 5 (Name Weakening). *Suppose that n does not occur in C, e, v .*

1. *If $\Delta \vdash \Sigma \triangleright C$, then $\Delta, n: [S] \vdash \Sigma \triangleright C$.*
2. *If $\Gamma \vdash \Sigma \triangleright e: T \triangleleft \Sigma'$, then $\Gamma, n: [S] \vdash \Sigma \triangleright e: T \triangleleft \Sigma'$.*
3. *If $\Gamma \vdash v: T$, then $\Gamma, n: [S] \vdash v: T$.*

Proof. The proof for configurations is a straightforward induction on the derivation of the judgement, using the result for expressions when the last rule in the derivation tree is T-THREAD. The proofs for expressions and for values are by mutual induction. \square

Lemma 6 (Name Narrowing). *Suppose that n is not in C, e, v .*

1. *If $\Delta, n: [S] \vdash \Sigma \triangleright C$, then $\Delta \vdash \Sigma \triangleright C$.*
2. *If $\Gamma, n: [S] \vdash \Sigma \triangleright e: T \triangleleft \Sigma'$, then $\Gamma \vdash \Sigma \triangleright e: T \triangleleft \Sigma'$.*
3. *If $\Gamma, n: [S] \vdash v: T$, then $\Gamma \vdash v: T$.*

⁵In the formulation of the lemma, we have omitted the hypothesis that x is not in the domain of Γ (for otherwise $\Gamma, x: T$ would not be defined in the conclusion). We henceforth follow this convention for all sorts of environments.

Proof. The proofs follow the pattern of the ones for the previous lemma. \square

The following two results allow to grow and shrink the channel environment of a configuration. Weakening is needed in Subject Reduction (rule R-CLOSE); Narrowing in Subject Congruence (channel extrusion using rule T-NEWB).

Lemma 7 (Channel Weakening). *Suppose that c does not occur in any of Σ, Σ', C, e, v .*

1. *If $\Delta \vdash \Sigma \triangleright C$, then $\Delta \vdash \Sigma, c: \perp \triangleright C$.*
2. *If $\Gamma \vdash \Sigma \triangleright e: T \triangleleft \Sigma'$, then $\Gamma \vdash \Sigma, c: \perp \triangleright e: T \triangleleft \Sigma', c: \perp$.*
3. *If $\Gamma \vdash v: (\Sigma; T \rightarrow U; \Sigma')$, then $\Gamma \vdash v: (\Sigma, c: \perp; T \rightarrow U; \Sigma', c: \perp)$.*

Proof. The proofs follow the pattern of the previous two lemmas. \square

Lemma 8 (Channel Narrowing). *Suppose that c does not occur in any of Γ, C, e, v .*

1. *If $\Delta \vdash \Sigma, c: S \triangleright C$, then $S = \perp$ and $\Delta \vdash \Sigma \triangleright C$.*
2. *If $\Gamma \vdash \Sigma, c: S \triangleright e: T \triangleleft \Sigma'$, then $\Sigma' = \Sigma'', c: S$ and $\Gamma \vdash \Sigma \triangleright e: T \triangleleft \Sigma''$.*
3. *If $\Gamma \vdash v: (\Sigma, c: S; T \rightarrow U; \Sigma', c: S)$, then $\Gamma \vdash v: (\Sigma; T \rightarrow U; \Sigma')$.*

Proof. The proofs are by induction on the possible derivation trees for the judgement.

1. Configurations. The only cases worth mentioning are when the last rule is T-PAR or T-THREAD; the four remaining cases follow directly by induction.

T-Par. By hypothesis we have a tree of the form:

$$\frac{\Delta \vdash \Sigma_1^* \triangleright C_1 \quad \Delta \vdash \Sigma_2^* \triangleright C_2}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: S \triangleright C_1 \mid C_2} \text{T-PAR}$$

We analyse the possibilities for splitting the environment $(\Sigma_1 \bullet \Sigma_2), c: S$. There are three cases.

1. $\Sigma_1^* = \Sigma_1, c: S$ and $\Sigma_2^* = \Sigma_2$
2. $\Sigma_1^* = \Sigma_1$ and $\Sigma_2^* = \Sigma_2, c: S$
3. $S = \perp, \Sigma_1^* = \Sigma_1, c: S_1$ and $\Sigma_2^* = \Sigma_2, c: \overline{S_1}$

The first two cases follow directly by induction. In the third case, the induction hypothesis implies that $S_1 = \perp$ and $\overline{S_1} = \perp$, meaning that $\Sigma_1^* \bullet \Sigma_2^*$ is not defined. So the third case cannot arise.

T-Thread. By hypothesis we have a tree of the form:

$$\frac{\Delta \vdash \Sigma, c: S \triangleright t: T \triangleleft \vec{d}: \perp}{\Delta \vdash \Sigma, c: S \triangleright \langle t \rangle}$$

and use the clause for expressions in this lemma, to conclude that $\vec{d}: \perp = \Sigma', c: S$ (and therefore $S = \perp$) and that $\Delta \vdash \Sigma \triangleright t: T \triangleleft \Sigma'$. The result follows by induction.

2. Expressions. The cases for all the rules in Figure 6 (except T-CASE), as well as T-NEW and T-VAL in Figure 7 are direct, for the antecedents (if any) do not mention channel environments. The case for rule T-APP uses the clause for

functional values in this lemma. All other cases, except T-FORK follow directly by induction.

T-Fork. By hypothesis we have a proof tree of the form:

$$\frac{\Delta \vdash \Sigma_1^* \triangleright t_1 : T_1 \triangleleft \vec{c} : \perp \quad \Delta \vdash \Sigma_2^* \triangleright t_2 : T_2 \triangleleft \vec{d} : \perp}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c : S \triangleright \text{fork } t_1; t_2 : T_2 \triangleleft \vec{c} : \perp, \vec{d} : \perp} \text{T-FORK}$$

A similar analysis to that done above in the case of rule T-PAR of configurations concludes the proof.

3. Values. Rules T-CONST, T-CHAN, and T-NAME do not apply, for the types in the axioms are not of the required form. Rule T-VAR does not apply, for c is not in Γ , by hypothesis. Rule T-REC follows by induction, and rule T-ABS uses the clause for expressions in this lemma. \square

The following result accounts for the monoidal structure of configurations; it is used in the proof of Subject Congruence.

Lemma 9 (Channel environment monoid). *Consider the commutative monoid axioms expressed in terms of $(\Sigma, \bullet, \emptyset)$, each in the form LHS = RHS. For each axiom, LHS is defined if and only if RHS is defined, and then they are equal.*

Proof. Directly from the definition of channel environment composition, on page 14. \square

Congruent configurations share the same typings. This result is used in the proof of Subject Reduction, rule R-CONF.

Lemma 10 (Subject Congruence). *If $\Delta \vdash \Sigma \triangleright C$ and $C \equiv C'$, then $\Delta \vdash \Sigma \triangleright C'$.*

Proof. The proof proceeds by induction on the derivation of $C \equiv C'$. The inductive cases (the congruence rules) are straightforward. We now consider the base cases.

When the last rule applied is the commutative monoid rule, we use Lemma 9. For the scope extrusion rules S-SCOPEOPEN and S-SCOPECLOSE we must consider each rule in both directions; for S-SCOPECLOSE we must consider two cases, depending on whether the typing derivation uses T-NEWB or T-NEWC.

S-ScopeN. When reading the rule left-to-right we use name weakening (Lemma 5). In the other direction we use name narrowing (Lemma 6). In both cases, we use the hypothesis (in the congruence rule) that n is not in C_2 .

S-ScopeC, left-to-right, T-NewB. By hypothesis, we have

$$\frac{\frac{\Delta \vdash \Sigma_1, c : \perp \triangleright C_1 \quad c \text{ not in } \Delta, \Sigma_1}{\Delta \vdash \Sigma_1 \triangleright (\nu c)C_1} \text{T-NEWB} \quad \Delta \vdash \Sigma_2 \triangleright C_2}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright (\nu c)C_1 \mid C_2} \text{T-PAR}$$

From the assumptions in the above tree, we build the following derivation, where we crucially use the variable convention to ensure that c is not in Σ_2 .

$$\frac{\frac{\Delta \vdash \Sigma_1, c : \perp \triangleright C_1 \quad \Delta \vdash \Sigma_2 \triangleright C_2}{\Delta \vdash \Sigma_1 \bullet \Sigma_2, c : \perp \triangleright C_1 \mid C_2} \text{T-PAR} \quad c \text{ not in } \Delta, \Sigma_1, \Sigma_2}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright (\nu c)(C_1 \mid C_2)} \text{T-NEWB}$$

S-ScopeC, left-to-right, T-NewC. Similar to the previous case, again using the variable convention.

S-ScopeC, right-to-left, T-NewB. By hypothesis, we have a proof tree of the form:

$$\frac{\frac{\Delta \vdash \Sigma_1^* \triangleright C_1 \quad \Delta \vdash \Sigma_2^* \triangleright C_2}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp \triangleright C_1 \mid C_2} \text{T-PAR} \quad c \text{ not in } \Delta, \Sigma_1, \Sigma_2}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2) \triangleright (\nu c)(C_1 \mid C_2)} \text{T-NEWB}$$

We analyse the possibilities for splitting environment $(\Sigma_1 \bullet \Sigma_2), c: \perp$. There are three cases.

1. $\Sigma_1^* = \Sigma_1, c: \perp$ and $\Sigma_2^* = \Sigma_2$.
2. $\Sigma_1^* = \Sigma_1$ and $\Sigma_2^* = \Sigma_2, c: \perp$.
3. $\Sigma_1^* = \Sigma_1, c: S$ and $\Sigma_2^* = \Sigma_2, c: \bar{S}$.

In case 1 we build the following derivation.

$$\frac{\frac{\Delta \vdash \Sigma_1, c: \perp \triangleright C_1 \quad c \text{ not in } \Delta, \Sigma_1}{\Delta \vdash \Sigma_1 \triangleright (\nu c)C_1} \text{T-NEWB} \quad \Delta \vdash \Sigma_2 \triangleright C_2}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright (\nu c)C_1 \mid C_2} \text{T-PAR}$$

In case 2, as by hypothesis of rule S-SCOPEC c does not occur free in C_2 (and thus, by the variable convention, it does not occur bound neither), we build the following derivation.

$$\frac{\frac{\Delta \vdash \Sigma_1 \triangleright C_1 \quad c \notin \Delta, \Sigma_1}{\Delta \vdash \Sigma_1 \triangleright (\nu c)C_1} \text{T-NEWC} \quad \frac{\Delta \vdash \Sigma_2, c: \perp \triangleright C_2 \quad c \notin \Delta, \Sigma_2, C_2}{\Delta \vdash \Sigma_2 \triangleright C_2} \text{Lemma 8}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright (\nu c)C_1 \mid C_2} \text{T-PAR}$$

In case 3 the typing derivation gives us the additional hypothesis that $(\Sigma_1, c: S) \bullet (\Sigma_2, c: \bar{S})$ is defined. As we have $\Delta \vdash \Sigma_2, c: \bar{S} \triangleright C_2$ and we know that c is not in C_2 , Lemma 8 applies and this judgement is $\Delta \vdash \Sigma_2, c: \perp \triangleright C_2$, contradicting the assumption that $(\Sigma_1, c: S) \bullet (\Sigma_2, c: \bar{S})$ is defined. Therefore this case cannot arise.

S-ScopeC, right-to-left, T-NewC. Similar to case 1 of the previous argument. \square

The following result allow to replace a given channel for a another one, through-out a derivation tree. We use it in Subject Reduction, rule R-INIT, to unify the two fresh channels in the hypothesis.

Lemma 11 (Channel replacement). *Suppose that d does not occur in any of $\Gamma, \Sigma, \Sigma', T, e, v$, and c does not occur in e, v .*

1. *If $\Gamma \vdash \Sigma \triangleright e: T \triangleleft \Sigma'$, then $\Gamma\{d/c\} \vdash \Sigma\{d/c\} \triangleright e: T\{d/c\} \triangleleft \Sigma'\{d/c\}$.*
2. *If $\Gamma \vdash v: T$, then $\Gamma\{d/c\} \vdash v: T\{d/c\}$.*

Proof. The proof of the two results, by mutual induction, is straightforward. \square

The following lemma accounts for all cases in Subject Reduction where substitution is needed, namely, in rules R-APP, R-REC, and R-BETA.

Lemma 12 (Substitution). *Suppose that $\Gamma \vdash v : T$.*

1. *If $\Gamma, x : T \vdash \Sigma \triangleright e : U \triangleleft \Sigma'$ then $\Gamma \vdash \Sigma \triangleright e\{v/x\} : U \triangleleft \Sigma'$.*
2. *If $\Gamma, x : T \vdash u : U$ then $\Gamma \vdash u\{v/x\} : U$.*

Proof. The proof of the two results is by mutual induction on the derivation of the judgement.

1. Expressions. The result follows easily using the result for values and induction.

2. Values. The cases of rules T-CONST, T-CHAN, and T-NAME follow easily, observing that x does not occur in u , and applying Lemma 4. The case of rule T-VAR follows trivially, as $u = x$. The case of rule T-ABS uses the result for expressions, and that of rule T-REC follows by induction. \square

We are finally in a position to prove Subject Reduction.

Proof of Theorem 1, page 15. The proof proceeds by induction on the derivation of $C \rightarrow C'$. We analyse each reduction rule in Figure 3, page 11, in turn.

R-Init. By hypothesis, we have

$$\langle \text{let } x = \text{request } n \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{accept } n \text{ in } t_2 \rangle \rightarrow (\nu c)(\langle \text{let } x = c \text{ in } t_1 \rangle \mid \langle \text{let } y = c \text{ in } t_2 \rangle).$$

and $\Delta \vdash \Sigma \triangleright \langle \text{let } x = \text{request } n \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{accept } n \text{ in } t_2 \rangle$. The only proof tree for this sequent is of the form

$$\frac{\text{(1)} \quad \text{(2)}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright \langle \text{let } x = \text{request } n \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{accept } n \text{ in } t_2 \rangle} \text{T-PAR}$$

where (1) is the tree

$$\frac{\frac{\Delta \vdash n : [S] \quad d_1 \text{ fresh}}{\Delta \vdash \Sigma_1 \triangleright \text{request } n : \text{Chan } d_1 \triangleleft \Sigma_1, d_1 : \vec{S}} \quad \Delta, x : \text{Chan } d_1 \vdash \Sigma_1, d_1 : \vec{S} \triangleright t_1 : T \triangleleft \vec{c} : \perp}{\Delta \vdash \Sigma_1 \triangleright \text{let } x = \text{request } n \text{ in } t_1 : T \triangleleft \vec{c} : \perp} \text{T-THREAD}}{\Delta \vdash \Sigma_1 \triangleright \langle \text{let } x = \text{request } n \text{ in } t_1 \rangle} \text{T-THREAD}$$

and (2) is the tree

$$\frac{\frac{\Delta \vdash n : [S] \quad d_2 \text{ fresh}}{\Delta \vdash \Sigma_2 \triangleright \text{accept } n : \text{Chan } d_2 \triangleleft \Sigma_2, d_2 : S} \quad \Delta, y : \text{Chan } d_2 \vdash \Sigma_2, d_2 : S \triangleright t_2 : U \triangleleft \vec{d} : \perp}{\Delta \vdash \Sigma_2 \triangleright \text{let } y = \text{accept } n \text{ in } t_2 : U \triangleleft \vec{d} : \perp} \text{T-THREAD}}{\Delta \vdash \Sigma_2 \triangleright \langle \text{let } y = \text{accept } n \text{ in } t_2 \rangle} \text{T-THREAD}$$

From the assumptions in the above tree we may build the following derivation tree, since as c does not occur in t_1, t_2 , by the variable convention it does not occur neither in $\Delta, \Sigma_1, \Sigma_2$. Thus, we are in the conditions of Lemma 11, since d_1 and d_2 are fresh in the assumptions of tree (1) and (2). By the same reason $\Delta\{c/d_i\} = \Delta$, and similarly for Σ_1, Σ_2 and for T, U .

$$\frac{\frac{\text{(1*)}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2, c: \perp \triangleright \langle \text{let } x = c \text{ in } t_1 \rangle} \quad \frac{\text{(2*)}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2, c: \perp \triangleright \langle \text{let } y = c \text{ in } t_2 \rangle}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2, c: \perp \triangleright \langle \text{let } x = c \text{ in } t_1 \rangle \mid \langle \text{let } y = c \text{ in } t_2 \rangle} \text{T-PAR}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright (\nu c)(\langle \text{let } x = c \text{ in } t_1 \rangle \mid \langle \text{let } y = c \text{ in } t_2 \rangle)} \text{T-NEWB}$$

where (1*) is the tree

$$\frac{\frac{\frac{}{\Delta \vdash c: \text{Chan } c} \text{T-CHAN} \quad \frac{\Delta, x: \text{Chan } d_1 \vdash \Sigma_1, d_1: S \triangleright t_1: T \triangleleft \vec{c}: \perp}{\Delta, x: \text{Chan } c \vdash \Sigma_1, c: S \triangleright t_1: T \triangleleft \vec{c}\{c/d_1\}: \perp} \text{Lemma 11}}{\Delta \vdash \Sigma_1, c: \vec{S} \triangleright \text{let } x = c \text{ in } t_1: T \triangleleft \vec{c}\{c/d_1\}: \perp} \text{T-LET}}{\Delta \vdash \Sigma_1, c: \vec{S} \triangleright \langle \text{let } x = c \text{ in } t_1 \rangle} \text{T-THREAD}$$

and (2*) is the tree

$$\frac{\frac{\frac{}{\Delta \vdash c: \text{Chan } c} \text{T-CHAN} \quad \frac{\Delta, y: \text{Chan } d_2 \vdash \Sigma_2, d_2: S \triangleright t_2: U \triangleleft \vec{d}: \perp}{\Delta, y: \text{Chan } c \vdash \Sigma_2, c: S \triangleright t_2: U \triangleleft \vec{d}\{c/d_2\}: \perp} \text{Lemma 11}}{\Delta \vdash \Sigma_2, c: S \triangleright \text{let } y = c \text{ in } t_2: U \triangleleft \vec{d}\{c/d_2\}: \perp} \text{T-LET}}{\Delta \vdash \Sigma_2, c: S \triangleright \langle \text{let } y = c \text{ in } t_2 \rangle} \text{T-THREAD}$$

R-Com. By hypothesis, we have

$$\langle \text{let } x = \text{receive } c \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{send } v \text{ on } c \text{ in } t_2 \rangle \rightarrow \langle \text{let } x = v \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle$$

There are two possible derivations to consider, depending on the kind of value v carried by channel c . Let us consider the case where v is a channel (the other case is similar—and simpler).

$$\frac{\frac{\text{(1)}}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = \text{receive } c \text{ in } t_1 \rangle} \quad \frac{\text{(2)}}{\Delta \vdash \Sigma \triangleright \langle \text{let } y = \text{send } v \text{ on } c \text{ in } t_2 \rangle}}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = \text{receive } c \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{send } v \text{ on } c \text{ in } t_2 \rangle} \text{T-PAR}$$

where Σ is $(\Sigma_1 \bullet \Sigma_2), c: \perp, d: S'$, and (1) is the tree

$$\frac{\frac{\frac{}{\Delta \vdash c: \text{Chan } c} \text{T-CHAN} \quad d \text{ fresh}}{\Delta \vdash \Sigma_1, c: ?S'.S \triangleright \text{receive } c: \text{Chan } d \triangleleft \Sigma'_1} \quad \frac{\Delta, x: \text{Chan } d \vdash \Sigma'_1 \triangleright t_1: T \triangleleft \vec{c}: \perp, c: \perp}{\Delta \vdash \Sigma_1, c: ?S'.S \triangleright \text{let } x = \text{receive } c \text{ in } t_1: T \triangleleft \vec{c}: \perp, c: \perp}}{\Delta \vdash \Sigma_1, c: ?S'.S \triangleright \langle \text{let } x = \text{receive } c \text{ in } t_1 \rangle} \text{T-THREAD}$$

where Σ'_1 is $\Sigma_1, c: S, d: S'$. (2) is the tree

$$\frac{\frac{\Delta \vdash v: \text{Chan } d \quad \overline{\Delta \vdash c: \text{Chan } c}}{\Delta \vdash \Sigma_2, c: !S'.\bar{S}, d: S' \triangleright \text{send } v \text{ on } c: \text{Unit} \triangleleft \Sigma'_2 \quad \Delta, y: \text{Unit} \vdash \Sigma'_2 \triangleright t_2: U \triangleleft \vec{d}: \perp}}{\Delta \vdash \Sigma_2, c: !S'.\bar{S}, d: S' \triangleright \text{let } y = \text{send } v \text{ on } c \text{ in } t_2: U \triangleleft \vec{d}: \perp}}{\Delta \vdash \Sigma_2, c: !S'.\bar{S}, d: S' \triangleright \langle \text{let } y = \text{send } v \text{ on } c \text{ in } t_2 \rangle} \text{T-THREAD}$$

where Σ'_2 is $\Sigma_2, c: \bar{S}$.

From the assumptions in the above tree, one may build

$$\frac{\frac{(1^*)}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp, d: S' \triangleright \langle \text{let } x = v \text{ in } t_1 \rangle} \quad \frac{(2^*)}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp, d: S' \triangleright \langle \text{let } y = \text{unit in } t_2 \rangle}}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp, d: S' \triangleright \langle \text{let } x = v \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle} \text{T-PAR}$$

where (1*) is the tree

$$\frac{\frac{\Delta \vdash v: \text{Chan } d \quad \Delta, x: \text{Chan } d \vdash \Sigma_1, c: S, d: S' \triangleright t_1: T \triangleleft \vec{c}: \perp, c: \perp}{\Delta \vdash \Sigma_1, c: S, d: S' \triangleright \text{let } x = v \text{ in } t_1: T \triangleleft \vec{c}: \perp, c: \perp}}{\Delta \vdash \Sigma_1, c: S, d: S' \triangleright \langle \text{let } x = v \text{ in } t_1 \rangle} \text{T-LET}$$

and (2*) is the tree

$$\frac{\frac{\overline{\Delta \vdash \text{unit}: \text{Unit}}} \text{T-UNIT} \quad \Delta, y: \text{Unit} \vdash \Sigma_2, c: \bar{S} \triangleright t_2: U \triangleleft \vec{d}: \perp}{\Delta \vdash \Sigma_2, c: \bar{S} \triangleright \text{let } y = \text{unit in } t_2: U \triangleleft \vec{d}: \perp}}{\Delta \vdash \Sigma_2, c: \bar{S} \triangleright \langle \text{let } y = \text{unit in } t_2 \rangle} \text{T-LET}$$

Notice that the type environment $(\Sigma_1, c: S, d: S') \bullet (\Sigma_2, c: \bar{S})$ in the conclusion of rule T-PAR above is defined, since $(\Sigma_1, c: ?S'.S) \bullet (\Sigma_2, c: !S'.\bar{S}, d: S')$ is defined (in the tree for the hypothesis) and d is fresh (in tree (1)).

R-Close. By hypothesis, we have

$$\langle \text{let } x = \text{close } c \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{close } c \text{ in } t_2 \rangle \rightarrow \langle \text{let } x = \text{unit in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle$$

and

$$\frac{\frac{(1)}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp \triangleright \langle \text{let } x = \text{close } c \text{ in } t_1 \rangle} \quad \frac{(2)}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp \triangleright \langle \text{let } y = \text{close } c \text{ in } t_2 \rangle}}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp \triangleright \langle \text{let } x = \text{close } c \text{ in } t_1 \rangle \mid \langle \text{let } y = \text{close } c \text{ in } t_2 \rangle} \text{T-PAR}$$

where (1) is the tree

$$\frac{\frac{\dots \quad \Delta, x: \text{Unit} \vdash \Sigma_1 \triangleright t_1: T \triangleleft \vec{c}: \perp_1}{\Delta \vdash \Sigma_1, c: \text{End} \triangleright \text{let } x = \text{close } c \text{ in } t_1: T \triangleleft \vec{c}: \perp_1}}{\Delta \vdash \Sigma_1, c: \text{End} \triangleright \langle \text{let } x = \text{close } c \text{ in } t_1 \rangle} \text{T-LET}$$

and (2) is the tree below.

$$\frac{\dots \quad \Delta, y: \text{Unit} \vdash \Sigma_2 \triangleright t_2: U \triangleleft \vec{c}: \perp_2}{\Delta \vdash \Sigma_2, c: \text{End} \triangleright \text{let } y = \text{close } c \text{ in } t_2: U \triangleleft \vec{c}: \perp_2} \text{T-LET}$$

$$\frac{\Delta \vdash \Sigma_2, c: \text{End} \triangleright \text{let } y = \text{close } c \text{ in } t_2: U \triangleleft \vec{c}: \perp_2}{\Delta \vdash \Sigma_2, c: \text{End} \triangleright \langle \text{let } y = \text{close } c \text{ in } t_2 \rangle} \text{T-THREAD}$$

From the assumptions in the above tree, one may build

$$\frac{\frac{\text{(1*)}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright \langle \text{let } x = \text{unit in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle} \text{T-PAR}}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp \triangleright \langle \text{let } x = \text{unit in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle} \text{Lemma 7}}$$

where (1*) is the tree

$$\frac{\frac{\frac{\Delta \vdash \text{unit}: \text{Unit}}{\Delta \vdash \Sigma_1 \triangleright \text{unit}: \text{Unit} \triangleleft \Sigma_1} \text{T-VAL}}{\Delta \vdash \Sigma_1 \triangleright \text{let } x = \text{unit in } t_1: T \triangleleft \vec{c}: \perp_1} \text{T-LET}}{\Delta \vdash \Sigma_1 \triangleright \langle \text{let } x = \text{unit in } t_1 \rangle} \text{T-THREAD}$$

and (2*) is a similar tree.

R-New. By hypothesis, we have

$$\langle \text{let } x = \text{new in } t \rangle \rightarrow (\nu n) \langle \text{let } x = n \text{ in } t \rangle$$

and

$$\frac{\frac{\frac{\Delta \vdash \text{new}: [S]}{\Delta \vdash \Sigma \triangleright \text{new}: [S] \triangleleft \Sigma} \text{T-VAL}}{\Delta \vdash \Sigma \triangleright \text{let } x = \text{new in } t: T \triangleleft \vec{c}: \perp} \text{T-LET}}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = \text{new in } t \rangle} \text{T-THREAD}$$

From the hypothesis in the above tree, we build a tree to complete the proof. Notice that, by the hypothesis of rule R-NEW, n is not free in t . Thus, Lemma 5 is applicable to the premise of rule T-LET above, and hence,

$$\frac{\frac{\frac{\Delta, x: [S] \vdash \Sigma \triangleright t: T \triangleleft \vec{c}: \perp}{\Delta, n: [S] \vdash n: [S]} \text{Lemma 5}}{\Delta, n: [S], x: [S] \vdash \Sigma \triangleright t: T \triangleleft \vec{c}: \perp} \text{T-LET}}{\Delta, n: [S] \vdash \Sigma \triangleright \text{let } x = n \text{ in } t: T \triangleleft \vec{c}: \perp} \text{T-THREAD}$$

$$\frac{\Delta, n: [S] \vdash \Sigma \triangleright \langle \text{let } x = n \text{ in } t \rangle}{\Delta \vdash \Sigma \triangleright (\nu n) \langle \text{let } x = n \text{ in } t \rangle} \text{T-NEWN}$$

R-Fork. By hypothesis, we have

$$\langle \text{fork } t_1; t_2 \rangle \rightarrow \langle t_1 \rangle \mid \langle t_2 \rangle$$

and

$$\frac{\frac{\Delta \vdash \Sigma_1 \triangleright t : _ \triangleleft \vec{c} : \perp \quad \Delta \vdash \Sigma_2 \triangleright t' : _ \triangleleft \vec{d} : \perp}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright \text{fork } t; t' : _ \triangleleft \vec{c} : \perp, \vec{d} : \perp} \text{T-FORK}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright \langle \text{fork } t; t' \rangle} \text{T-THREAD}$$

From the hypotheses in the above tree, we build a tree to complete the proof.

$$\frac{\frac{\Delta \vdash \Sigma_1 \triangleright t : _ \triangleleft \vec{c} : \perp}{\Delta \vdash \Sigma_1 \triangleright \langle t \rangle} \text{T-THREAD} \quad \frac{\Delta \vdash \Sigma_2 \triangleright t' : _ \triangleleft \vec{d} : \perp}{\Delta \vdash \Sigma_2 \triangleright \langle t' \rangle} \text{T-THREAD}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright \langle t \rangle \mid \langle t' \rangle} \text{T-PAR}$$

R-App. By hypothesis, we have

$$\langle \text{let } x = (\lambda y.e)v \text{ in } t \rangle \rightarrow \langle \text{let } x = e\{v/y\} \text{ in } t \rangle$$

and

$$\frac{\frac{\frac{\Delta, y : T \vdash \Sigma \triangleright e : U \triangleleft \Sigma'}{\Delta \vdash \lambda y.e : (\Sigma; T \rightarrow U; \Sigma')} \text{T-ABS} \quad \Delta \vdash v : T}{\Delta \vdash \Sigma \triangleright (\lambda y.e)v : U \triangleleft \Sigma'} \text{T-APP} \quad \Delta, x : U \vdash \Sigma' \triangleright t : T \triangleleft \vec{c} : \perp}{\frac{\Delta \vdash \Sigma \triangleright \text{let } x = (\lambda y.e)v \text{ in } t : T \triangleleft \vec{c} : \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = (\lambda y.e)v \text{ in } t \rangle} \text{T-THREAD}} \text{T-THREAD}$$

Then, one may build the following derivation to complete the proof.

$$\frac{\frac{\Delta \vdash v : T \quad \Delta, y : T \vdash \Sigma \triangleright e : U \triangleleft \Sigma'}{\Delta \vdash \Sigma \triangleright e\{v/y\} : U \triangleleft \Sigma'} \text{Lemma 12} \quad \Delta, x : U \vdash \Sigma' \triangleright t : T \triangleleft \vec{c} : \perp}{\frac{\Delta \vdash \Sigma \triangleright \text{let } x = e\{v/y\} \text{ in } t : T \triangleleft \vec{c} : \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = e\{v/y\} \text{ in } t \rangle} \text{T-THREAD}} \text{T-LET}$$

R-Rec. By hypothesis, we have

$$\langle \text{let } x = (\text{rec } y.v)u \text{ in } t \rangle \rightarrow \langle \text{let } x = (v\{\text{rec } y.v/y\})u \text{ in } t \rangle$$

and

$$\frac{\frac{\frac{\Delta, y : (\Sigma; T \rightarrow U; \Sigma') \vdash v : (\Sigma; T \rightarrow U; \Sigma')}{\Delta \vdash \text{rec } y.v : (\Sigma; T \rightarrow U; \Sigma')} \text{T-REC} \quad \Delta \vdash u : T}{\Delta \vdash \Sigma \triangleright (\text{rec } y.v)u : U \triangleleft \Sigma'} \text{T-APP} \quad (1)}{\frac{\Delta \vdash \Sigma \triangleright \text{let } x = (\text{rec } y.v)u \text{ in } t : T \triangleleft \vec{c} : \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = (\text{rec } y.v)u \text{ in } t \rangle} \text{T-LET}} \text{T-THREAD}$$

where (1) is $\Delta, x : U \vdash \Sigma' \triangleright t : T \triangleleft \vec{c} : \perp$.

Then, one may build the following derivation to complete the proof.

$$\begin{array}{c}
\frac{\Delta, y: (\Sigma; T \rightarrow U; \Sigma') \vdash v: (\Sigma; T \rightarrow U; \Sigma')}{\Delta \vdash \text{rec } y.v: (\Sigma; T \rightarrow U; \Sigma')} \text{T-REC} \\
\frac{\Delta \vdash \text{rec } y.v: (\Sigma; T \rightarrow U; \Sigma')}{\Delta \vdash v\{\text{rec } y.v/y\}: (\Sigma; T \rightarrow U; \Sigma')} \text{Lemma 12} \\
\frac{\Delta \vdash v\{\text{rec } y.v/y\}: (\Sigma; T \rightarrow U; \Sigma') \quad \Delta \vdash u: T}{\Delta \vdash \Sigma \triangleright (v\{\text{rec } y.v/y\})u: U \triangleleft \Sigma'} \text{T-APP} \quad (1) \\
\frac{\Delta \vdash \Sigma \triangleright (v\{\text{rec } y.v/y\})u: U \triangleleft \Sigma'}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = (v\{\text{rec } y.v/y\})u \text{ in } t \rangle: T \triangleleft \vec{c}: \perp} \text{T-LET} \\
\frac{\Delta \vdash \Sigma \triangleright \langle \text{let } x = (v\{\text{rec } y.v/y\})u \text{ in } t \rangle: T \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = (v\{\text{rec } y.v/y\})u \text{ in } t \rangle} \text{T-THREAD}
\end{array}$$

R-Beta. By hypothesis, we have

$$\langle \text{let } x = v \text{ in } t \rangle \rightarrow \langle t\{v/x\} \rangle.$$

There are two possible derivations for $\langle \text{let } x = v \text{ in } t \rangle$; we analyse each in turn. When the derivation uses rule T-LET, the result follows by Lemma 12.

$$\begin{array}{c}
\frac{\Delta \vdash v: U}{\Delta \vdash \Sigma \triangleright v: U \triangleleft \Sigma} \text{T-VAL} \\
\frac{\Delta \vdash \Sigma \triangleright v: U \triangleleft \Sigma \quad \Delta, x: U \vdash \Sigma \triangleright t: T \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = v \text{ in } t \rangle: T \triangleleft \vec{c}: \perp} \text{T-LET} \\
\frac{\Delta \vdash \Sigma \triangleright \langle \text{let } x = v \text{ in } t \rangle: T \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = v \text{ in } t \rangle} \text{T-THREAD}
\end{array}$$

When the derivation uses rule T-POLYLET, the result is immediate.

$$\begin{array}{c}
\frac{\Delta \vdash \Sigma \triangleright t\{v/x\}: T \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = v \text{ in } t \rangle: T \triangleleft \vec{c}: \perp} \text{T-POLYLET} \\
\frac{\Delta \vdash \Sigma \triangleright \langle \text{let } x = v \text{ in } t \rangle: T \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = v \text{ in } t \rangle} \text{T-THREAD}
\end{array}$$

R-Let. By hypothesis, we have

$$\langle \text{let } x = (\text{let } y = e \text{ in } t') \text{ in } t \rangle \rightarrow \langle \text{let } y = e \text{ in } (\text{let } x = t' \text{ in } t) \rangle$$

and

$$\begin{array}{c}
\frac{\Delta \vdash \Sigma \triangleright e: T \triangleleft \Sigma_1 \quad \Delta, y: T \vdash \Sigma_1 \triangleright t': T_1 \triangleleft \Sigma'_1}{\Delta \vdash \Sigma \triangleright \text{let } y = e \text{ in } t': T_1 \triangleleft \Sigma'_1} \\
\frac{\Delta \vdash \Sigma \triangleright \text{let } y = e \text{ in } t': T_1 \triangleleft \Sigma'_1 \quad \Delta, x: T_1 \vdash \Sigma'_1 \triangleright t: U \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \text{let } x = (\text{let } y = e \text{ in } t') \text{ in } t: U \triangleleft \vec{c}: \perp} \text{T-THREAD} \\
\frac{\Delta \vdash \Sigma \triangleright \text{let } x = (\text{let } y = e \text{ in } t') \text{ in } t: U \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = (\text{let } y = e \text{ in } t') \text{ in } t \rangle}
\end{array}$$

Then,, using Lemma 3 one may build the following derivation to complete the proof. Notice that, by the variable convention, y is not free in t , since it is bound in $\text{let } y = e \text{ in } t'$.

$$\begin{array}{c}
\frac{\Delta \vdash \Sigma \triangleright e: T \triangleleft \Sigma_1 \quad (1)}{\Delta \vdash \Sigma \triangleright \text{let } y = e \text{ in } (\text{let } x = t' \text{ in } t): U \triangleleft \vec{c}: \perp} \text{T-LET} \\
\frac{\Delta \vdash \Sigma \triangleright \text{let } y = e \text{ in } (\text{let } x = t' \text{ in } t): U \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } y = e \text{ in } (\text{let } x = t' \text{ in } t) \rangle} \text{T-THREAD}
\end{array}$$

where (1) is the tree

$$\frac{\frac{\Delta, x: T_1 \vdash \Sigma'_1 \triangleright t: U \triangleleft \vec{c}: \perp}{\Delta, y: T \vdash \Sigma_1 \triangleright t': T_1 \triangleleft \Sigma'_1} \text{Lemma 3} \quad \Delta, y: T, x: T_1 \vdash \Sigma'_1 \triangleright t: U \triangleleft \vec{c}: \perp}{\Delta, y: T \vdash \Sigma_1 \triangleright \text{let } x = t' \text{ in } t: U \triangleleft \vec{c}: \perp} \text{T-LET}$$

R-IfT/R-IffF. Follows the pattern in all the above cases.

R-Branch. Follows the pattern in all the above cases.

R-Conf. The three cases follow directly by induction. For the rule that uses structural congruence, we use Lemma 10. \square

B Proof of Theorem 2, Type Safety

We start with a couple of easy results.

Lemma 13. *Suppose that $\Delta \vdash \Sigma \triangleright C$.*

1. *If C is a c -thread, then c is in the domain of Σ .*
2. *If C is a c -redex, then Σ is of the form $\Sigma', c: \perp$.*

Proof. 1. A simple analysis of the conclusions of the last rule applied in the derivation of the sequent for c -threads, namely T-SENDD, T-SENDS, T-RECEIVED, T-RECEIVES, T-CASE, T-SELECT, and T-CLOSE.

2. A simple analysis of the possible derivation trees for the three possible c -redex cases. \square

Proof of Theorem 2, page 16. By contradiction, assuming faulty configurations typable and performing a case analysis on the possible forms of the faulty configurations.

Assume $\Delta \vdash \Sigma \triangleright (\nu \vec{a})(C_1 \mid C_2)$. Without loss of generality, assume that $\vec{a} = \vec{n} \vec{c} \vec{d}$, where \vec{d} are the channels that do not occur in Δ, Σ . Build the only possible proof tree for the above sequent, first using rule T-NEWN as many times as there are names in \vec{n} , then proceeding similarly with rules T-NEWB and T-NEWC, a finally with rule T-PAR, to obtain two subtrees ending with the sequents ($i = 1, 2$):

$$\Delta, \vec{n}: [\vec{S}] \vdash \Sigma_i \triangleright C_i \tag{1}$$

where $\Sigma, \vec{c}: \vec{\perp} = \Sigma_1 \bullet \Sigma_2$. We now analyse each of the five possible cases of faulty configurations defined in Section 6, where we let $\Delta' = \Delta, \vec{n}: [\vec{S}]$.

1. The three cases are similar. We analyse the conditional expression. The only derivation tree for sequent (1) above is of the form below.

$$\frac{\frac{\Delta' \vdash v: \text{Bool} \quad \dots \quad \dots}{\Delta' \vdash \Sigma_1 \triangleright \text{if } v \text{ then } e_1 \text{ else } e_2: _ \triangleleft \vec{f}: \perp} \text{T-IF} \quad \dots}{\Delta' \vdash \Sigma_1 \triangleright \text{let } x = \text{if } v \text{ then } e_1 \text{ else } e_2 \text{ in } t: _ \triangleleft \vec{f}: \perp} \text{T-LET} \quad \dots}{\Delta' \vdash \Sigma_1 \triangleright \langle \text{let } x = \text{if } v \text{ then } e_1 \text{ else } e_2 \text{ in } t \rangle} \text{T-THREAD}$$

Analysing the rules for values (Figure 5, page 12), one realises that v can only be true or false, for the T-VAR does not apply since variables are not in the domain of

Δ' , and the type in the conclusion of the remaining rules (T-ABS, T-REC, T-CHAN, T-NAME, Unit) is not **Bool**.

2. As above, analyse the lower part of the only proof tree for, say,

$$\Delta' \vdash \Sigma_1 \triangleright \langle \text{let } x = \text{accept } v \text{ in } t \rangle$$

to obtain a tree for

$$\Delta' \vdash v : [S].$$

Once again, among the rules for values, only T-NAME applies. Then, v is a name.

3. As above, analyse the lower part of the only proof tree for, say,

$$\Delta' \vdash \Sigma_1 \triangleright \langle \text{let } x = \text{receive } v \text{ in } t \rangle$$

to obtain a tree for

$$\Delta' \vdash v : \text{Chan } c.$$

Once again, among the rules for values, only T-CHAN applies. Clearly v can only be the channel c .

4. There are several cases to check in this point; they are all similar. Pick, for example, the pair **select/close**, and expand the lower part of the proof tree, until obtaining subtrees for the following two sequents,

$$\Delta' \vdash \Theta_1 \triangleright \text{select } l \text{ on } c : T_1 \triangleleft \Theta'_1 \qquad \Delta' \vdash \Theta_2 \triangleright \text{close } c : T'_2 \triangleleft \Theta'_2$$

where $\Sigma_1 = \Theta_1 \bullet \Theta_2$. Analysing the rule for **select**, one finds that $c : \oplus \langle l : S \rangle$ must be in Θ_1 . Similarly, analysing the rule for **close** one realises that $c : \text{End}$ must be in Θ_2 . Then, $\Theta_1 \bullet \Theta_2$ is not defined (for $\oplus \langle l : S \rangle$ is not the dual of **End**), hence $(\nu \vec{a})(C_1 \mid C_2)$ is not typable.

5. We check the case for three c -threads $\langle t_1 \rangle \mid \langle t_2 \rangle \mid \langle t_3 \rangle$, the others reduce to this. We have:

$$\frac{\Delta' \vdash \Sigma' \triangleright \langle t_1 \rangle \mid \langle t_2 \rangle \qquad \Delta' \vdash \Sigma'' \triangleright \langle t_3 \rangle}{\Delta' \vdash \Sigma_1 \triangleright \langle t_1 \rangle \mid \langle t_2 \rangle \mid \langle t_3 \rangle} \text{T-PAR}$$

with $\Sigma_1 = \Sigma' \bullet \Sigma''$. If $\langle t_1 \rangle \mid \langle t_2 \rangle$ is not a c -redex, then we use the previous case. Otherwise, by Lemma 13, it must be the case that $c : \perp$ is part of Σ' . Since $\langle t_3 \rangle$ is a c -thread, by Lemma 13, c is in the domain of Σ'' . But then $\Sigma' \bullet \Sigma''$ is not defined (for \perp is dual to no type), and $(\nu \vec{a})(C_1 \mid C_2)$ is not typable. \square