

Session Types for Functional Multithreading*

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

November 9, 2004

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 system of our language, and prove subject reduction and runtime type safety theorems. We also present a type checking algorithm and show that it is correct with respect to the type system.

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

1 Introduction

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

System programming is traditionally performed in the C programming language. It thus comes as no surprise that many attempts to statically check protocols are based on this language: safe control of stateful resources is achieved

*A revised and extended version of the paper in Concur 2004, volume 3170 of LNCS, pages 497–511, Springer-Verlag, 2004.

[†]Departamento de Informática, Faculdade de Ciências, Universidade de Lisboa, 1749-016 Lisboa, Portugal.

[‡]Department of Computing Science, University of Glasgow, Glasgow G12 8QQ, UK.

[§]CLC and Departamento de Matemática, Instituto Superior Técnico, 1049-001 Lisboa, Portugal.

via type systems that either run on annotated C programs [7, 9], or on programs written in a type-safe variant of C [13, 14]. Another approach to proving properties of protocols comes from the general setting of the π -calculus [21, 27], and includes type and effect systems to check correspondence assertions [2, 12]; the approximation of the behaviour of π -processes by CCS terms [5, 25]; and session types to describe structured communication programming [11, 15, 16, 28].

Session types allow 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 for conventional languages, 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, but has not been studied theoretically in the context of a standard language paradigm, despite a few contributions which bridge session types and conventional languages. Session types have been used to add behavioural information to the interfaces of CORBA objects [29] using Gay and Hole's theory of subtyping [11] to formalise compatibility and substitutability of components. Session types have also been encoded in the Haskell programming language [22]. The former does not link the improved CORBA interfaces to actual programming languages; the latter does not address the correspondence between a session-based programming language and Haskell.

Our contribution to the problem of structured communication-based programming in general, and the verification of protocols in particular, is to transfer the concept of session types from the π -calculus to a multi-threaded functional language with side-effecting input/output operations. This shows that static checking of session types could be added to a language such as Concurrent ML [26] (at least without imperative features) or Concurrent Haskell [20] (cf. [22]). More generally, it is our view that by starting with typing concepts which are well-understood in the context of a theoretical calculus, and transferring them to a language which is closer to mainstream programming practice, we can achieve a powerful type system which is suited to practical programming while retaining the benefits of a sound foundation.

The key technical steps which we have undertaken, in order to address the differences between a conventional programming style and the programming notation of the π -calculus, are as follows:

- 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 naming and creation 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.

In previous work [10] we have presented a language supporting typed func-

tional programming with inter-process communication channels, but we only considered individual processes in isolation. In [30] we addressed collections of functional threads communicating via (session) channels created from shared names. Here we present the proofs for subject reduction and type safety, claimed in [30], and introduce a type checking algorithm for the language, together with a proof of its correctness.

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 to 5 define the syntax, operational semantics, and type system of our language. In Section 6 we present the runtime safety result. Section 7 presents a type checking algorithm. In Sections 8 and 9 we discuss related and future work. The appendix contains the proofs of all the results in the paper.

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:

```
server u = let x = receive u in
           let y = receive u in
           send x + y on u
```

or more concisely: `send (receive u) + (receive u) 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.

```
client u = send 2 on u
           send 3 on u
           let x = receive u in code
```

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:

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

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, *sqr*, 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.

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

This example shows that session types allow the description of protocols that cannot be easily 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:

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

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 [10].

Function Types. We have not mentioned the type of the *server* itself. Clearly, it accepts a channel (in state $\&\langle \text{add} : \dots, \text{neg} : \dots \rangle$), and returns nothing (described by the Unit type). The body of the function “consumes” the channel, leaving it in a state ready to be closed (described by type End). We write all this as follows, where c is the (runtime) channel denoted by the (program) variable u .

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

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

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

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

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

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

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

¹We often omit the empty channel environment on each side of the arrow.

Establishing a Connection. How do the client and the server reach a state in which they both know about channel c ? We follow Takeuchi, Kubo and Honda [28], 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. We may then write:

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

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.

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.

$$\begin{aligned} \text{system} &:: \text{Unit} \\ \text{system} &= \text{let } x = \text{new } \&\langle \text{add}: \dots, \text{neg}: \dots, \text{eval}: \dots \rangle \text{ in} \\ &\quad \text{fork } \text{negClient } x; \text{fork } \text{addClient } x; \text{fork } \text{server } x \end{aligned}$$

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.

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

Sending Channels on Channels. Imagine two clients that 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`, 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

²Alternatively, we may `send x` on an existing channel.

type $?(?Int.End).End$ (call this type S) and establish a connection for the sole purpose of transmitting the server channel.

$$\begin{array}{ll}
askNeg :: [\langle add: \dots \rangle] \rightarrow [S] \rightarrow Unit & getNeg :: [S] \rightarrow Unit \\
askNeg \ x \ y = \text{let } u = \text{request } x \text{ in} & getNeg \ y = \text{let } w = \text{accept } y \text{ in} \\
\quad \text{select } neg \text{ on } u; \text{send } 7 \text{ on } u & \quad \text{let } u = \text{receive } w \text{ in} \\
\quad \text{let } w = \text{request } y \text{ in} & \quad \text{let } i = \text{receive } u \text{ in} \\
\quad \text{send } u \text{ on } w; \text{close } w & \quad \text{close } u; \text{close } w; \text{code}
\end{array}$$

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.

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 $getNeg$.

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

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

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.

$$\begin{array}{l}
sendTwice :: c: !Int.!Int.End; Chan \ c \rightarrow Unit; c: End \\
sendTwice \ w = sendSend \ w \ w
\end{array}$$

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

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

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

$$\begin{array}{l}
sendFree :: c: !Int.!Int.End; Chan \ c \rightarrow Unit; c: End \\
sendSend :: c: !Int.!Int.End; Chan \ c \rightarrow Chan \ c \rightarrow Unit; c: End
\end{array}$$

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

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

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

in a typing environment associating the type $\text{Chan } c$ to the free variable u of sendFree . However, if aliasing u and v is not sought, then we must have

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

in a typing environment containing $u: \text{Chan } d$. Note how this type for sendFree captures channel changes, parameters to the function or not.

Polymorphism. We have seen that sendFree admits at least two different types. In order to allow for code reuse we work with a type-free syntax, and type our functions 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 the former 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 $!\text{Int}!\text{Int.End}$. Here sendTwice must be typed once with channel c , and another with channel d . For the latter we have:

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

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

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 Figures 1 and 2.⁴

⁴Cf. [19] on the use of threads to facilitate the presentation of the operational semantics.

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

Figure 2: Syntax of channel types, data types, types, and channel environments

$$\begin{aligned}
(C, |, \langle \mathbf{unit} \rangle) &\text{ is a commutative monoid} && \text{(S-MONOID)} \\
(\nu a : T)C_1 \mid C_2 &\equiv (\nu a : T)(C_1 \mid C_2) \text{ if } a \text{ not free in } C_2 && \text{(S-SCOPE)}
\end{aligned}$$

Figure 3: Structural congruence

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 definitions, we use $(\nu a : T)$ to stand for either $(\nu c : S)$ or $(\nu n : T)$.

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.

The syntax of types is described in figure 2. We define *session types* S , *data types* D , *term types* T , and *channel environments* Σ . The type $\mathbf{Chan} \ c$ 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_S , denotes a channel that is already in use by two threads, hence that cannot be used further. One of the two threads expects to see the channel at type S , the other at type dual of S (figure 5). Similarly to channel and name identifiers, \perp_S is not available at the top level syntax, arising only via the channel environment composition operator, $\Sigma_1 \bullet \Sigma_2$, defined in section 5. 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 .

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 **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**.

To support typechecking, we annotate λ -abstractions with a channel environment as well as a typed argument.

$$\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 } S \text{ 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(\Sigma; y: T).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: T).v)u \text{ in } t \rangle \rightarrow \langle \text{let } x = (v\{\text{rec } (y: T).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: T)C \rightarrow (\nu a: T)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 4: Reduction rules

4 Operational Semantics

The binding occurrences are the variable x in $\lambda(\Sigma; x: T).e$, in $\text{rec } (x: T).e$, and in $\text{let } x = e \text{ in } t$; the name n in $(\nu n: D)C$; and the channel c in $(\nu c: S)C$. 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 4), making use of a simple structural congruence relation [21] (figure 3), allowing for the rearrangement of threads in a configuration, so that reduction may happen.⁵

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

⁵We could easily arrange for structural congruence to garbage collect all threads of the form $\langle v \rangle$, for v closed.

$$\begin{aligned} \overline{\text{End}} &= \text{End} & \overline{?D.\overline{S}} &= !D.\overline{S} & \overline{?S'.\overline{S}} &= !S'.\overline{S} & \overline{!D.\overline{S}} &= ?D.\overline{S} & \overline{!S'.\overline{S}} &= ?S'.\overline{S} \\ \overline{\&l_i: S_i}_{i \in I} &= \oplus \langle l_i: \overline{S_i} \rangle_{i \in I} & \oplus \langle l_i: S_i \rangle_{i \in I} &= \&l_i: \overline{S_i} \rangle_{i \in I} \end{aligned}$$

Figure 5: Duality on session types

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 explicitly programmed by arranging for both threads to share a channel and explicitly sending the result back to the parent.

Example 1. *We follow the execution of thread $\langle \text{system} \rangle$ in page 6. Types play no rôle in reduction; we omit them.*

$$\begin{aligned} & \langle \text{system} \rangle \rightarrow^* \\ & (\nu n) \langle \text{fork } \text{negClient } n; \text{fork } \text{addClient } n; \text{fork } \text{server } n \rangle \rightarrow^* \\ (\nu n) (\langle \text{let } u = \text{request } n \text{ in } \dots \rangle \mid \langle \text{let } u = \text{request } n \text{ in } \dots \rangle \mid \langle \text{let } u = \text{accept } n \text{ in } \dots \rangle) & \rightarrow^* \\ (\nu n) (\langle \text{let } u = \text{request } n \text{ in } \dots \rangle \mid (\nu c) (\langle \text{select } \text{neg} \text{ on } c; \dots \rangle \mid \langle \text{case } c \text{ of } \dots \rangle)) & \rightarrow^* \\ (\nu n) (\langle \text{let } u = \text{request } n \text{ in } \dots \rangle \mid (\nu c) (\langle \text{close } c \rangle \mid \langle \text{close } c \rangle)) & \rightarrow^* \\ & (\nu n) \langle \text{let } u = \text{request } n \text{ in } \dots \rangle \end{aligned}$$

In the third line we have two threads competing for requesting a session on name n . We have chosen the negClient to go ahead; a new channel c , known only to the negClient and to the server, is created. In line four we have a typical interaction between two threads on a common channel. Similar interactions continue until both channels are ready to be closed, in line five. Once closed, structural congruence allows to get rid of the terminated threads. The run ends in line six with the addClient still waiting for a server that will never come.

5 Typing

The type system is presented in figures 6 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 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 2). The difference between Σ and Σ' reflects the effect of an expression on the types of channels. For example

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

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

Figure 6: Typing rules for values

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

Figure 7: Typing rules for expressions I: Channel operations

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 6). 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 7, 8). 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

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

Figure 8: Typing rules for expressions II: Other rules

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 the continuation to hold both ends of the channel. The rules T-SEND 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* [11, 15, 16, 28, 29]. The dual of a session type S , denoted \overline{S} , is defined for all session types except \perp_S , 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 5.

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 S` has type $[S]$, denoting a name that, when shared by two threads, is able to produce (via `accept/request`) new channels of type S .

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} : $\perp_{\vec{S}}$

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

Figure 9: Typing rules for configurations

and $\vec{d}: \perp_{\vec{S}}$ in the antecedent, that each thread involved either consumes their channels (that is, sends or closes the channels), or uses them in dual mode.

The *composition of two channel environments*, $\Sigma_1 \bullet \Sigma_2$, is defined only when $\Sigma_1(c) = \Sigma_2(c)$, 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_S when $c \in \text{dom } \Sigma_1 \cap \text{dom } \Sigma_2$ and $\Sigma_1(c) = S$, 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 [24, 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 threads, since the type \perp_S 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 24.

Theorem 2 (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.

$$\begin{array}{c}
\Gamma; \text{true} \mapsto \text{Bool} \quad \Gamma; \text{false} \mapsto \text{Bool} \quad \Gamma; \text{unit} \mapsto \text{Unit} \quad (\text{C-CONST}) \\
\Gamma; c \mapsto \text{Chan } c \quad \Gamma, n: [S]; n \mapsto [S] \quad \Gamma, x: T; x \mapsto T \\
\hspace{15em} (\text{C-CHAN, C-NAME, C-VAR}) \\
\frac{\Gamma, x: T; \Sigma; e \mapsto \Sigma_1; U; \Sigma_2 \quad \forall c. \Sigma(c) \neq \perp}{\Gamma; \lambda(\Sigma; x: T).e \mapsto (\Sigma; T \rightarrow U; \Sigma_1, \Sigma_2)} \quad (\text{C-ABS}) \\
\frac{\Gamma, x: T; v \mapsto T \quad T = (\Sigma; U \rightarrow U'; \Sigma')}{\Gamma; \text{rec } (x: T).v \mapsto T} \quad (\text{C-REC})
\end{array}$$

Figure 10: Type checking values

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 $C \equiv (\nu\Gamma)(C_1 \mid C_2)$ 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 34.

Theorem 3 (Type Safety). *Typable configurations are not faulty.*

7 Type Checking

This section presents a type checking algorithm in the form of typing rules. Typing judgements for values are of the form $\Gamma; v \mapsto T$, where Γ and v describe the input, and T describes the output for the type checking function. Judgements for expressions are of the form $\Gamma; \Sigma; e \mapsto \Sigma'; T; \Sigma''$, where the symbols at the left of the arrow represent the input, and those on the right the output. Channel environment Σ' describes the *unused* part of Σ , whereas Σ'' represents the final types of the channels which are used by e . Channels which are created

$$\begin{array}{c}
\frac{\Gamma; v \mapsto D \quad \Gamma; v' \mapsto \text{Chan } c}{\Gamma; \Sigma, c: !D.S; \text{send } v \text{ on } v' \mapsto \Sigma; \text{Unit}; c: S} \\
\frac{\Gamma; v \mapsto D \quad \Gamma; v' \mapsto \text{Chan } c}{\Gamma; \Sigma, c: \perp_{!D.S}; \text{send } v \text{ on } v' \mapsto \Sigma, c: ?D.\bar{S}; \text{Unit}; c: S} \\
\frac{\Gamma; v \mapsto D \quad \Gamma; v' \mapsto \text{Chan } c}{\Gamma; \Sigma, c: \perp_{?D.S}; \text{send } v \text{ on } v' \mapsto \Sigma, c: ?D.S; \text{Unit}; c: \bar{S}} \quad (\text{C-SEND}) \\
\frac{\Gamma; v \mapsto \text{Chan } c \quad \Gamma; \Sigma, c: S_j; e_j \mapsto \Sigma_1; T; \Sigma_2 \quad \forall j \in I}{\Gamma; \Sigma, c: \&\langle l_i: S_i \rangle_{i \in I}; \text{case } v \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} \mapsto \Sigma_1; T; \Sigma_2} \\
\frac{\Gamma; v \mapsto \text{Chan } c \quad \Gamma; \Sigma, c: S_j; e_j \mapsto \Sigma_1; T; \Sigma_2 \quad \forall j \in I}{\Gamma; \Sigma, c: \perp_{\&\langle l_i: S_i \rangle_{i \in I}}; \text{case } v \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} \mapsto \Sigma_1, c: \oplus \langle l_i: \bar{S}_i \rangle_{i \in I}; T; \Sigma_2} \\
\frac{\Gamma; v \mapsto \text{Chan } c \quad \Gamma; \Sigma, c: S_j; e_j \mapsto \Sigma_1; T; \Sigma_2 \quad \forall j \in I}{\Gamma; \Sigma, c: \perp_{\oplus \langle l_i: S_i \rangle_{i \in I}}; \text{case } v \text{ of } \{l_i \Rightarrow e_i\}_{i \in I} \mapsto \Sigma_1, c: \&\langle l_i: \bar{S}_i \rangle_{i \in I}; T; \Sigma_2} \quad (\text{C-CASE})
\end{array}$$

Figure 11: Type checking expressions I: Channel operations

by e appear in Σ'' . If $c: S \in \Sigma$ and $S \neq \perp$, then we will have either $c: S \in \Sigma'$ if e does not use c , or $c: S' \in \Sigma''$ if e uses c and leaves it with type S' , or neither if $S = \text{End}$ and e closes c . If $c: \perp_S \in \Sigma$ or $c: \perp_{\bar{S}} \in \Sigma$, and e uses c , then it will be checked that e uses just one end of c , and we will have either $c: \bar{S} \in \Sigma'$ and $c: S' \in \Sigma''$, or $c: S \in \Sigma'$ and $c: S'' \in \Sigma''$. For example

$$\begin{array}{l}
x: \text{Chan } c; c: ?\text{Int.End}; \text{receive } x \mapsto \emptyset; \text{Int}; c: \text{End}. \\
x: \text{Chan } c; c: \perp_{?\text{Int.End}}; \text{receive } x \mapsto c: !\text{Int.End}; \text{Int}; c: \text{End}. \\
x: \text{Chan } c; c: \perp_{!\text{Int.End}}; \text{receive } x \mapsto c: !\text{Int.End}; \text{Int}; c: \text{End}.
\end{array}$$

Finally, typing judgements for configurations are of the form $\Delta; \Sigma; C \mapsto \Sigma'$ where Σ' describes the unused part of Σ .

The type inference system is presented in figures 10 to 13. Most of the rules are obtained from those in figures 6 to 9, by a suitable reading of the sequents.

Type Checking Values (figure 10). In rule C-ABS we require, just as in rule T-ABS, that Σ does not contain any \perp types. This means that a function can only use one end of a channel. The following lemma (Lemma 19, Appendix C) guarantees that Σ_1 and Σ_2 have disjoint domains, so that Σ_1, Σ_2 is defined and the overall transformation of channel types can be described by the function type $\Sigma; T \rightarrow U; \Sigma_1, \Sigma_2$.

Lemma If $\Gamma; \Sigma; e \mapsto \Sigma_1; T; \Sigma_2$ and Σ contains no \perp types, then $\text{dom}(\Sigma_1) \cap \text{dom}(\Sigma_2) = \emptyset$.

$$\begin{array}{c}
\frac{\Gamma; v \mapsto [S] \quad c \text{ fresh}}{\Gamma; \Sigma; \text{request } v \mapsto \Sigma; \text{Chan } c; c: \overline{S}} \quad \frac{\Gamma; v \mapsto [S] \quad c \text{ fresh}}{\Gamma; \Sigma; \text{accept } v \mapsto \Sigma; \text{Chan } c; c: S} \\
\text{(C-REQUEST, C-ACCEPT)} \\
\frac{\Gamma; \Sigma; t_1 \mapsto \Sigma_1; T_1; \emptyset \quad \Gamma; \Sigma_1; t_2 \mapsto \Sigma_2; T_2; \emptyset}{\Gamma; \Sigma; \text{fork } t_1; t_2 \mapsto \Sigma_2; T_2; \emptyset} \quad \text{(C-FORK)} \\
\Gamma; \Sigma; \text{new } S \mapsto \Sigma; [S]; \emptyset \quad \frac{\Gamma; v \mapsto (\Sigma; T \rightarrow U; \Sigma') \quad \Gamma; v' \mapsto T}{\Gamma; \Sigma; vv' \mapsto \emptyset; U; \Sigma'} \\
\text{(C-NEW, C-APP)} \\
\frac{\Gamma; v \mapsto T}{\Gamma; \Sigma; v \mapsto \Sigma; T; \emptyset} \quad \frac{\Gamma; v \mapsto \text{Bool} \quad \Gamma; \Sigma; e \mapsto \Sigma_1; T; \Sigma_2 \quad \Gamma; \Sigma; e' \mapsto \Sigma_1; T; \Sigma_2}{\Gamma; \Sigma; \text{if } v \text{ then } e \text{ else } e' \mapsto \Sigma_1; T; \Sigma_2} \\
\text{(C-VAL, C-IF)} \\
\frac{\Gamma; \Sigma; e \mapsto \Sigma_1; T_1; \Sigma'_1 \quad \Gamma, x: T_1; (\Sigma \cap \Sigma_1), \Sigma'_1; t \mapsto \Sigma_2; T_2; \Sigma'_2}{\Gamma; \Sigma; \text{let } x = e \text{ in } t \mapsto \Sigma_3; T_2; (\Sigma'_1 \cap \Sigma_2), \Sigma'_2} \quad \text{(C-LET)}
\end{array}$$

In rule C-LET, $\Sigma_3 = (\Sigma_1 \cap \Sigma_2) \cup \{c: S \mid (c: \perp_S \in \Sigma \text{ or } c: \perp_{\overline{S}} \in \Sigma) \text{ and } (c: S \in \Sigma_1 \text{ or } c: S \in \Sigma_2)\}$.

Figure 12: Type checking expressions II: other rules

Type Checking Expressions (figures 11, 12). Each rule in figure 7 gives rise to three rules in the type checking system. In figure 11 we present the rules corresponding to T-SENDD and to T-CASE; the remaining cases are easy to infer.

The three rules for `send-datatype` distinguish the three possible types for channel c in the input environment. If the type is output, $!D.S$, then the type-checking function consumes the type, yielding $c: S$ in the consumed channel-environment. When the type is bottom-output, $\perp_{!D.S}$, then, as above, the type-checking function returns $c: S$ in the consumed part. In this case, however, the function announces, via $c: ?D.\overline{S}$ in the unused part, that the dual type is available. When the type is bottom-input, the process is dual.

The three rules for `case` follow a similar pattern. For example, for the second rule, when the input channel environment Σ contains an entry of the form $c: \perp_{\& \langle l_i: S_i \rangle_{i \in I}}$, we call the type checking function for each branch e_j in the body of the `case` with $c: S_j$, and place $c: \oplus \langle l_i: \overline{S_i} \rangle_{i \in I}$ in the unused part.

In figure 12, contrast rule C-NEW with rule C-ABS. Since values do not change channels, the former rule consumes nothing, placing the input Σ in the unused position in the output. The latter consumes channels as described by Σ in the type for the function, leaving empty the unused component in the output.

Rule C-FORK requires that t_1 and t_2 fully consume all of the channels that they use, either by using them completely and closing them, or by sending them to other threads. This corresponds to the hypotheses of rule T-FORK. Of the

$$\begin{array}{c}
\frac{\Delta; \Sigma; t \mapsto \Sigma'; T; \emptyset}{\Delta; \Sigma; \langle t \rangle \mapsto \Sigma'} \quad \frac{\Delta; \Sigma; C_1 \mapsto \Sigma_1 \quad \Delta; \Sigma_1; C_2 \mapsto \Sigma_2}{\Delta; \Sigma; C_1 \mid C_2 \mapsto \Sigma_2} \quad \frac{\Delta, n: [S]; \Sigma; C \mapsto \Sigma'}{\Delta; \Sigma; (\nu n: S)C \mapsto \Sigma'} \\
\text{(C-THREAD, C-PAR, C-NEWN)} \\
\frac{\Delta; \Sigma, c: \perp_S; C \mapsto \Sigma' \quad c \text{ not in } \Delta, \Sigma}{\Delta; \Sigma; (\nu c: \perp_S)C \mapsto \Sigma'} \quad \frac{\Delta; \Sigma; C \mapsto \Sigma' \quad c \text{ not in } \Delta, \Sigma}{\Delta; \Sigma; (\nu c: \perp_S)C \mapsto \Sigma'} \\
\text{(C-NEWB, C-NEWC)}
\end{array}$$

Figure 13: Type checking configurations

initial channels Σ , some are consumed by t_1 and the remainder, Σ_1 , are used to typecheck t_2 . Any channels remaining after typechecking t_2 , i.e. Σ_2 , are returned as the channels which are not used by the `fork` expression.

The most complex typechecking rule is `C-LET`. The channel environment $(\Sigma \cap \Sigma_1), \Sigma'_1$ which is used when typechecking t consists of the channels $(\Sigma \cap \Sigma_1)$ which are not used by e and the updated types (Σ_1) of the channels which are used by e . Note that if $c: \perp_S \in \Sigma$ but e uses (one end of) c , so that $c: \bar{S} \in \Sigma$ and $c: S' \in \Sigma'_1$, then t is typechecked with respect to $c: S'$, which ensures that t can only use the same end of c which e uses. The final modified channel environment, $(\Sigma'_1 \cap \Sigma_2), \Sigma'_2$, has a similar form to the initial channel environment when typechecking t . The definition of the unused channels, Σ_3 , includes channels which are used by neither e nor t , together with the unused portions of channels $c: \perp \in \Sigma$ which are used by either e or t .

Type Checking Configurations (figure 13). In rule `C-PAR`, we run the algorithm on C_1 and feed the output Σ_1 , the unused part of Σ , into another call, this time for C_2 . The output Σ_2 of the second call is the unused channel environment of the parallel composition $C_1 \mid C_2$ (cf. rule `C-FORK` in 12).

Correctness. The type checking algorithm is sound and complete with respect to the type system presented in section 5. To state soundness we define the partial operation $\Sigma - \Sigma'$ as follows.

$$\begin{aligned}
(\Sigma, c: S) - c: S &= \Sigma \\
(\Sigma, c: \perp_S) - c: S &= \Sigma, c: \bar{S} \\
(\Sigma, c: \perp_{\bar{S}}) - c: S &= \Sigma, c: \bar{S} \\
\Sigma - c: S &= \text{undefined, otherwise}
\end{aligned}$$

The definition is then extended inductively in the second argument. The formulation of correctness is standard; the proofs are in appendices C (page 36), and D (page 40).

Theorem 4 (Soundness). *If $\Delta; \Sigma; C \mapsto \Sigma'$ then $\Sigma - \Sigma'$ is defined and $\Delta \vdash \Sigma - \Sigma' \triangleright C$.*

Theorem 5 (Completeness). *If $\Delta \vdash \Sigma \triangleright C$ without using rule T-POLYLET, then for all Σ' such that $\Sigma \bullet \Sigma'$ is defined, $\Delta; \Sigma \bullet \Sigma'; C \mapsto \Sigma'$.*

8 Related Work

Cyclone [14] is a C-like type-safe polymorphic imperative language. It features region-based memory management, and more recently threads and locks [13], 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 [7] 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” [9], 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 [12] use one such system to prove progress properties of communication protocols written in π -calculus. Bonelli, Compagnoni, and Gunter [2] combine the language of Honda, Vasconcelos and Kubo [16] with the correspondence assertions of Gordon and Jeffrey, thus obtaining a setting where further properties can be proved about programs.

Rajamani *et al.*’s *Behave* [5, 25] 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) verification is 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 [18]. 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 in general. Walker, Crary, and Morrisett [32] present a language to describe region-based memory management together with a provably safe type system. Igarashi and Kobayashi [17] 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 [22], but the results are limited to type soundness.

9 Future Work

A prototype implementation of the language, targeted at the Multithreaded Intermediate Language [23], is under way. 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 [10]. We feel its incorporation in the present setting would not present major difficulties, although care must be taken in the definition of duality [29].

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.

Web services. Our work opens up the possibility of an application of session types to verification of web service implementations [6, 31]. Web services require a model for business interactions, which typically assume the form of sequences of peer-to-peer message exchanges, both synchronous and asynchronous, within stateful, long-running interactions involving two or more parties. Although some rigorous semantics have been developed (eg., [3]), there is still little assistance on the verification of the correctness of the protocol descriptions and their composition. (eg., [4, 8]). Session types may provide a useful static analysis tool.

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

Acknowledgements. This work was partially supported by the EU IST proactive initiative FET-Global Computing (projects Mikado, IST-2001-32222, and Profundis, IST-2001-33100), Fundação para a Ciência e a Tecnologia (via CLC,

CITI, and the project MIMO, POSI/CHS/39789/2001), and a Treaty of Windsor grant from the British Council in Portugal and the Portuguese Council of University Rectors.

References

- [1] H. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North-Holland, Amsterdam, 1984.
- [2] E. Bonelli, A. Compagnoni, and E. Gunter. Typechecking safe process synchronization. In *FGUC 2004. Third EATCS Workshop in the Foundations of Global Computing*, ENTCS. Elsevier Science, 2004. To appear.
- [3] M. J. Butler and C. Ferreira. A process compensation language. In *Proceedings of IFM*, pages 61–76, 2000.
- [4] M. J. Butler and C. Ferreira. Using SPIN and STeP to verify business processes specifications. In *Proceedings of Eeshov Memorial Conference*, pages 207–213, 2003.
- [5] S. Chaki, S. K. Rajamani, and J. Rehof. Types as models: model checking message-passing programs. In *POPL*, pages 45–57. ACM Press, 2002.
- [6] F. Curbera, Y. Golland, J. Klein, F. Leymann, S. T. D. Roller, and S. Weerawarana. Business process execution language for web services, version 1.1. Technical report, IBM, 2003. <http://www-106.ibm.com/developerworks/library/ws-bpel/>.
- [7] R. DeLine and M. Fähndrich. Enforcing high-level protocols in low-level software. In *PLDI*, pages 59–69. ACM Press, 2001.
- [8] C. Ferreira and M. J. Butler. Using b refinement to analyse compensating business processes. In *Proceedings of ZB*, pages 477–496, 2003.
- [9] J. S. Foster, T. Terauchi, and A. Aiken. Flow-Sensitive Type Qualifiers. In *PLDI*, pages 1–12, Berlin, Germany, June 2002.
- [10] S. Gay, V. T. Vasconcelos, and A. Ravara. Session types for inter-process communication. TR 2003–133, Department of Computing, University of Glasgow, March 2003.
- [11] S. J. Gay and M. J. Hole. Types and subtypes for client-server interactions. volume 1576 of *LNCS*, pages 74–90. Springer-Verlag, 1999.
- [12] A. Gordon and A. Jeffrey. Typing correspondence assertions for communication protocols. *Theoretical Computer Science*, 300:379–409, 2003.
- [13] D. Grossman. Type-safe multithreading in cyclone. In *ACM Workshop on Types in Language Design and Implementation*. ACM Press, 2003.

- [14] D. Grossman, G. Morrisett, T. Jim, M. Hicks, Y. Wang, and J. Cheney. Region-based memory management in cyclone. In *PLDI*, pages 282–293. ACM Press, 2002.
- [15] K. Honda. Types for dyadic interaction. In *CONCUR'93*, volume 715 of *LNCS*, pages 509–523. Springer-Verlag, 1993.
- [16] K. Honda, V. Vasconcelos, and M. Kubo. Language primitives and type discipline for structured communication-based programming. volume 1381 of *LNCS*, pages 122–138. Springer-Verlag, 1998.
- [17] A. Igarashi and N. Kobayashi. Resource usage analysis. In *POPL*, pages 331–342. ACM Press, 2002.
- [18] A. Igarashi and N. Kobayashi. A generic type system for the pi-calculus. *Theoretical Computer Science*, 311(1–3):121–163, 2003.
- [19] A. Jeffrey and J. 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.
- [20] S. P. Jones, A. Gordon, and S. Finne. Concurrent Haskell. In *Proceedings of the 23rd ACM Symposium on Principles of Programming Languages (POPL'96), St. Petersburg Beach, Florida, USA*, pages 295–308. ACM Press, 1996.
- [21] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, I and II. *Information and Computation*, 100(1):1–77, September 1992.
- [22] M. Neubauer and P. Thiemann. An implementation of session types. In *Practical Aspects of Declarative Languages (PADL'04)*, volume 3057 of *LNCS*, pages 56–70. Springer-Verlag, 2004.
- [23] H. Paulino, P. Marques, L. Lopes, V. T. Vasconcelos, and F. Silva. A multi-threaded asynchronous language. In *7th International Conference on Parallel Computing Technologies (PaCT'03)*, volume 2763 of *LNCS*, pages 316–323. Springer-Verlag, 2003.
- [24] B. Pierce. *Types and Programming Languages*. The MIT Press, 2002.
- [25] S. K. Rajamani and J. Rehof. A behavioral module system for the pi-calculus. volume 2126 of *LNCS*, pages 375–394. Springer-Verlag, 2001.
- [26] J. Reppy. CML: a higher order concurrent language. In *PLDI*, pages 293–305. ACM Press, 1991.
- [27] D. Sangiorgi and D. Walker. *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.

- [28] K. Takeuchi, K. Honda, and M. Kubo. An interaction-based language and its typing system. volume 817 of *LNCS*. Springer-Verlag, 1994.
- [29] A. Vallecillo, V. T. Vasconcelos, and A. 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.
- [30] V. T. Vasconcelos, A. Ravara, and S. Gay. Session types for functional multithreading. In *CONCUR'04*, volume 3170 of *LNCS*, pages 497–511. Springer-Verlag, 2004.
- [31] W3C. Web services choreography requirements, W3C working draft, 2004. <http://www.w3.org/TR/2004/WD-ws-chor-reqs-20040311/>.
- [32] D. Walker, K. Crary, and G. Morrisett. Typed memory management via static capabilities. *TOPLAS*, 22(4):701–771, 2000.

A Proof of Theorem 2, Subject Reduction

We start with a few auxiliary results; the proof of Subject Reduction is on page 29. 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 15.⁶

Lemma 6 (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, are straightforward. □

Lemma 7 (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, by mutual induction, are straightforward. □

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

Lemma 8 (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. □

Lemma 9 (Name Narrowing). *Suppose that n is not in C, e, v .*

1. *If $\Delta, n: [S] \vdash \Sigma \triangleright C$, then $\Delta \vdash \Sigma \triangleright C$.*

⁶In 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.

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

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

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 10 (Channel Weakening). *Suppose that c does not occur in any of $\Delta, \Gamma, \Sigma, \Sigma', C, e, v$.*

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 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 11 (Channel Narrowing). *Suppose that c does not occur in any of $\Delta, \Gamma, \Sigma, \Sigma', C, e, v$.*

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:

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

We analyse the possibilities for splitting the environment $(\Sigma_1 \bullet \Sigma_2), c: \perp$. 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 $\Sigma_2^* = \Sigma_2$, $c: \overline{S_1}$

Each case follows easily by induction.

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 $\vec{d}: \perp = \Sigma', c: S$ (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 7 (except T-CASE), as well as T-NEW and T-VAL in figure 8 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:

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

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.

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 12 (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 13 (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 12. For the scope extrusion rules S-SCOPE_N and S-SCOPE_C we must consider each rule in both directions; for S-SCOPE_C we must consider two cases, depending on whether the typing derivation uses T-NEWB or T-NEWC.

S-Scope_N, left-to-right, T-NewB. When reading the rule left-to-right we use name weakening (lemma 8). In the other direction we use name narrowing (lemma 9). In both cases, we use the hypothesis (in the congruence rule) that n is not in C_2 .

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

S-Scope_C, 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 we build the following derivation.

$$\frac{\frac{\Delta \vdash \Sigma_1 \triangleright C_1 \quad c \notin \Delta, \Sigma_1}{\Delta \vdash \Sigma_1 \triangleright (\nu c)C_1} \text{T-NEWC} \quad \frac{\Delta \vdash \Sigma_2, c: \perp \triangleright C_2 \quad c \notin \Delta, \Sigma_2, C_2}{\Delta \vdash \Sigma_2 \triangleright C_2} \text{Lemma 11}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright (\nu c)C_1 \mid C_2} \text{T-PAR}$$

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 11 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, throughout a derivation tree. We use it in Subject Reduction, rule R-INIT, to unify the two fresh channels in the hypothesis.

Lemma 14 (Channel replacement). *Suppose that neither d nor c occurs in any of $\Gamma, \Sigma, \Sigma', T, e, v$.*

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 15 (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 7. 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 2, page 14. The proof proceeds by induction on the derivation of $C \rightarrow C'$. We analyse each reduction rule in figure 4, page 10, 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

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

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 tree

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

where (1^{*}) is the tree

$$\frac{\frac{\frac{\Delta, x : \text{Chan } d_1 \vdash \Sigma_1, d_1 : S \triangleright t_1 : T \triangleleft \vec{c} : \perp}{\Delta \vdash c : \text{Chan } c} \text{T-CHAN} \quad \frac{\Delta, x : \text{Chan } c \vdash \Sigma_1, c : S \triangleright t_1 : T \triangleleft \vec{c} : \perp}{\Delta \vdash \Sigma_1, c : \bar{S} \triangleright \text{let } x = c \text{ in } t_1 : T \triangleleft \vec{c} : \perp} \text{Lemma 14}}{\Delta \vdash \Sigma_1, c : \bar{S} \triangleright \langle \text{let } x = c \text{ in } t_1 \rangle} \text{T-LET}}{\Delta \vdash \Sigma_1, c : \bar{S} \triangleright \langle \text{let } x = c \text{ in } t_1 \rangle} \text{T-THREAD}$$

and (2^{*}) is the tree

$$\frac{\frac{\frac{\Delta, y : \text{Chan } d_2 \vdash \Sigma_2, d_2 : S \triangleright t_2 : U \triangleleft \vec{d} : \perp}{\Delta \vdash c : \text{Chan } c} \text{T-CHAN} \quad \frac{\Delta, y : \text{Chan } c \vdash \Sigma_2, c : S \triangleright t_2 : U \triangleleft \vec{d} : \perp}{\Delta \vdash \Sigma_2, c : S \triangleright \text{let } y = c \text{ in } t_2 : U \triangleleft \vec{d} : \perp} \text{Lemma 14}}{\Delta \vdash \Sigma_2, c : S \triangleright \langle \text{let } y = c \text{ in } t_2 \rangle} \text{T-LET}}{\Delta \vdash \Sigma_2, c : S \triangleright \langle \text{let } y = c \text{ in } t_2 \rangle} \text{T-THREAD}$$

Notice that we are in the conditions of lemma 14, 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_2 and for U .

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).

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

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

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

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

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

where Σ'_2 is $\Sigma_2, c: \bar{S}$.

From the assumptions in the above tree, one may build

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

where (1*) is the tree

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

and (2*) is the tree

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

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

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

where (1) is the tree

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

and (2) is the tree below.

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

From the assumptions in the above tree, one may build

$$\frac{\frac{\frac{}{(1^*)} \quad \frac{}{(2^*)}}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright \langle \text{let } x = \text{unit in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle} \text{T-PAR}}{\Delta \vdash (\Sigma_1 \bullet \Sigma_2), c: \perp \triangleright \langle \text{let } x = \text{unit in } t_1 \rangle \mid \langle \text{let } y = \text{unit in } t_2 \rangle} \text{Lemma 10}}$$

where (1*) is the tree

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

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

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

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

$$\frac{\frac{\frac{\overline{\Delta, n : [S] \vdash n : [S]}}{\Delta, n : [S], x : [S] \vdash \Sigma \triangleright t : T \triangleleft \vec{c} : \perp} \text{Lemma 8}}{\Delta, n : [S] \vdash \Sigma \triangleright \text{let } x = n \text{ in } t : T \triangleleft \vec{c} : \perp} \text{T-LET}}{\frac{\Delta, n : [S] \vdash \Sigma \triangleright \langle \text{let } x = n \text{ in } t \rangle}{\Delta \vdash \Sigma \triangleright (\nu n) \langle \text{let } x = n \text{ in } t \rangle} \text{T-THREAD}} \text{T-NEWN}$$

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.

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

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

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

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

$$\frac{\frac{\Delta \vdash v: T \quad \Delta, y: T \vdash \Sigma \triangleright e: U \triangleleft \Sigma'}{\Delta \vdash \Sigma \triangleright e\{v/y\}: U \triangleleft \Sigma'} \text{Lemma 15} \quad \Delta, x: U \vdash \Sigma' \triangleright t: T' \triangleleft \vec{c}: \perp}{\frac{\Delta \vdash \Sigma \triangleright \text{let } x = e\{v/y\} \text{ in } t: T' \triangleleft \vec{c}: \perp}{\Delta \vdash \Sigma \triangleright \langle \text{let } x = e\{v/y\} \text{ in } t \rangle} \text{T-THREAD}} \text{T-LET}}$$

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

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

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.

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

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 15.

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

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

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

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

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

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

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

where (1) is the tree

$$\frac{\Delta, y: T \vdash \Sigma_1 \triangleright t': T_1 \triangleleft \Sigma'_1 \quad \frac{\Delta, x: T_1 \vdash \Sigma'_1 \triangleright t: U \triangleleft \vec{c}: \perp}{\Delta, y: T, x: T_1 \vdash \Sigma'_1 \triangleright t: U \triangleleft \vec{c}: \perp} \text{Lemma 6}}{\Delta, y: T \vdash \Sigma_1 \triangleright \text{let } x = t' \text{ in } t: U \triangleleft \vec{c}: \perp} \text{T-LET}}$$

R-IfT/R-IfF. 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 13. \square

B Proof of Theorem 3, Type Safety

We start with a couple of easy results.

Lemma 16. *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-SEND, 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 3, page 15. 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 ($i = 1, 2$):

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

where $\Sigma, \vec{c}: \vec{\Gamma} = \Sigma_1 \bullet \Sigma_2$. 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.

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

Analysing the rules for values (figure 6, 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,

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

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 $\oplus(l: S)$ is not the dual of **End**), hence $(\nu\bar{a})(C_1 \mid C_2)$ 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:

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

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 16, it must be the case that $c: \perp$ is part of Σ' . Since $\langle t_3 \rangle$ is a c -thread, by Lemma 16, c is in the domain of Σ'' . But then $\Sigma' \bullet \Sigma''$ is not defined (for \perp is dual to no type), and $(\nu\bar{a})(C_1 \mid C_2)$ is not typable. \square

C Proof of Theorem 4, Soundness of Typechecking

The first two lemmas deal with the first conclusion of the soundness result: that $\Sigma - \Sigma'$ is defined.

Lemma 17. *If $\Gamma; \Sigma; e \mapsto \Sigma_1; T; \Sigma_2$ then $\Sigma - \Sigma_1$ is defined.*

Proof. By induction on the derivation of $\Gamma; \Sigma; e \mapsto \Sigma_1; T; \Sigma_2$. In most cases the conclusion follows immediately from the inference rule and the definition of $\Sigma - \Sigma'$. Two cases are worthy of comment.

C-Fork. We have a derivation ending with

$$\frac{\Gamma; \Sigma; t_1 \mapsto \Sigma_1; T_1; \emptyset \quad \Gamma; \Sigma_1; t_2 \mapsto \Sigma_2; T_2; \emptyset}{\Gamma; \Sigma; \text{fork } t_1; t_2 \mapsto \Sigma_2; T_2; \emptyset}$$

By induction, $\Sigma - \Sigma_1$ and $\Sigma_1 - \Sigma_2$ are defined. Therefore $\Sigma - \Sigma_2$ is defined, because of the following reasoning. If $c: S \in \Sigma_2$ and $\Sigma_1 - \Sigma_2$ is defined then either $c: S \in \Sigma_1$ or $c: \perp_S \in \Sigma_1$ or $c: \perp_{\bar{S}} \in \Sigma_1$. Because $\Sigma - \Sigma_1$ is defined, either $c: S \in \Sigma$ or $c: \perp_S \in \Sigma$ or $c: \perp_{\bar{S}} \in \Sigma$. In each of these cases, $\Sigma - c: S$ is defined. Applying the same argument to each component of Σ_2 shows that $\Sigma - \Sigma_2$ is defined.

C-Let. We have a derivation ending with

$$\frac{\Gamma; \Sigma; e \mapsto \Sigma_1; T_1; \Sigma'_1 \quad \Gamma, x: T_1; (\Sigma \cap \Sigma_1), \Sigma'_1; t \mapsto \Sigma_2; T_2; \Sigma'_2}{\Gamma; \Sigma; \text{let } x = e \text{ in } t \mapsto \Sigma_3; T_2; (\Sigma'_1 \cap \Sigma_2), \Sigma'_2}$$

where

$$\Sigma_3 = (\Sigma_1 \cap \Sigma_2) \cup \{c: S \mid (c: \perp_S \in \Sigma \text{ or } c: \perp_{\bar{S}} \in \Sigma) \text{ and } (c: S \in \Sigma_1 \text{ or } c: S \in \Sigma_2)\}.$$

By induction, $\Sigma - \Sigma_1$ is defined. We argue that $\Sigma - \Sigma_3$ is defined, as follows. $\Sigma_1 \cap \Sigma_2$ can be subtracted from Σ because it is a subset of Σ_1 and we know that $\Sigma - \Sigma_1$ is defined. The set

$$\{c: S \mid (c: \perp_S \in \Sigma \text{ or } c: \perp_{\bar{S}} \in \Sigma) \text{ and } (c: S \in \Sigma_1 \text{ or } c: S \in \Sigma_2)\}$$

(assuming that $c: S \notin \Sigma_1 \cap \Sigma_2$) can be subtracted from $\Sigma - \Sigma_1 \cap \Sigma_2$ because in this case either $c: \perp_S \in \Sigma$ or $c: \perp_{\bar{S}} \in \Sigma$, and therefore $c: S$ can be subtracted. \square

Lemma 18. *If $\Delta; \Sigma; C \mapsto \Sigma'$ then $\Sigma - \Sigma'$ is defined.*

Proof. By induction on the derivation of $\Delta; \Sigma; C \mapsto \Sigma'$. The case of rule C-THREAD uses Lemma 17. The case of rule C-PAR uses the same argument as the case of rule C-FORK in the proof of Lemma 17. In the other cases the induction hypothesis is used directly. \square

Lemma 19. *If $\Gamma; \Sigma; e \mapsto \Sigma_1; T; \Sigma_2$ and Σ contains no \perp types, then $\text{dom}(\Sigma_1) \cap \text{dom}(\Sigma_2) = \emptyset$.*

Proof. Straightforward induction on the derivation. \square

Lemmas 20 and 21 are proved by simultaneous induction on the size of the derivation of the typechecking judgement in question. For clarity we present the proofs separately, noting that the use of each lemma in the proof of the other is valid.

Lemma 20 (Soundness for Values). *If $\Gamma; v \mapsto T$ then $\Gamma \vdash v: T$.*

Proof. The cases for literal values, channels, names and variables are trivial.

C-Abs. We have a derivation finishing with

$$\frac{\Gamma, x: T; \Sigma; e \mapsto \Sigma_1; U; \Sigma_2 \quad \forall c. \Sigma(c) \neq \perp}{\Gamma; \lambda(\Sigma; x: T).e \mapsto (\Sigma; T \rightarrow U; \Sigma_1, \Sigma_2)}$$

Applying soundness for expressions (Lemma 21) to the hypothesis, we have

$$\Gamma, x: T \vdash \Sigma - \Sigma_1 \triangleright e: U \triangleleft \Sigma_2$$

Using Weakening, we have

$$\Gamma, x: T \vdash (\Sigma - \Sigma_1), \Sigma_1 \triangleright e: U \triangleleft \Sigma_1, \Sigma_2$$

Because Σ does not contain any \perp types, $(\Sigma - \Sigma_1), \Sigma_1$ is defined and equal to Σ . Now we use rule T-ABS to conclude

$$\Gamma \vdash \lambda(\Sigma; x: T).e: (\Sigma; T \rightarrow U; \Sigma_1, \Sigma_2)$$

as required.

C-Rec. We have a derivation finishing with

$$\frac{\Gamma, x: T; v \mapsto T \quad T = (\Sigma; U \rightarrow U'; \Sigma')}{\Gamma; \text{rec}(x: T).v \mapsto T}$$

By induction we have

$$\Gamma, x: T \vdash v: T$$

from which we use rule T-REC to conclude

$$\Gamma \vdash \text{rec}(x: T).v: T$$

as required. \square

Lemma 21 (Soundness for Expressions). *If $\Gamma; \Sigma; e \mapsto \Sigma_1; T; \Sigma_2$ then $\Sigma - \Sigma_1$ is defined and $\Gamma \vdash \Sigma - \Sigma_1 \triangleright e : T \triangleleft \Sigma_2$.*

Proof. By induction on the derivation of $\Gamma; \Sigma; e \mapsto \Sigma_1; T; \Sigma_2$. Lemma 17 guarantees that $\Sigma - \Sigma_1$ is defined. We consider cases for each of the inference rules in the typechecking algorithm.

C-Send. In the first case we have

$$\frac{\Gamma; v \mapsto D \quad \Gamma; v' \mapsto \text{Chan } c}{\Gamma; \Sigma, c : !D.S; \text{send } v \text{ on } v' \mapsto \Sigma; \text{Unit}; c : \bar{S}}$$

Lemma 20 implies that $\Gamma \vdash v : D$ and $\Gamma \vdash v' : \text{Chan } c$. Using rule T-SENDDD we build the following derivation, noting that $\Sigma, c : !D.S - \Sigma = c : !D.S$.

$$\frac{\Gamma \vdash v : D \quad \Gamma \vdash v' : \text{Chan } c}{\Gamma \vdash c : !D.S \triangleright \text{send } v \text{ on } v' : \text{Unit} \triangleleft c : \bar{S}}$$

The two other cases are similar. For the former, we note that $\Sigma, c : !D.S - \Sigma, c : ?D.\bar{S} = c : !D.S$, and for the latter: $\Sigma, c : ?D.S - \Sigma, c : ?D.\bar{S} = c : !D.S$.

The remaining cases in figure 11 are similar.

C-Request. We have

$$\frac{\Gamma; v \mapsto [S] \quad c \text{ fresh}}{\Gamma; \Sigma; \text{request } v \mapsto \Sigma; \text{Chan } c; c : \bar{S}}$$

Lemma 20 implies that $\Gamma \vdash v : [S]$. Using rule T-REQUEST and noting that $\Sigma - \Sigma = \emptyset$, we obtain

$$\frac{\Gamma \vdash v : [S] \quad c \text{ fresh}}{\Gamma \vdash \emptyset \triangleright \text{request } v : \text{Chan } c \triangleleft c : \bar{S}}$$

The case of C-ACCEPT is similar.

C-Fork. We have

$$\frac{\Gamma; \Sigma; t_1 \mapsto \Sigma_1; T_1; \emptyset \quad \Gamma; \Sigma_1; t_2 \mapsto \Sigma_2; T_2; \emptyset}{\Gamma; \Sigma; \text{fork } t_1; t_2 \mapsto \Sigma_2; T_2; \emptyset}$$

The induction hypothesis yields the hypotheses in the following instance of rule T-FORK:

$$\frac{\Gamma \vdash \Sigma - \Sigma_1 \triangleright t_1 : T_1 \triangleleft \emptyset \quad \Gamma \vdash \Sigma_1 - \Sigma_2 \triangleright t_2 : T_2 \triangleleft \emptyset}{\Gamma \vdash (\Sigma - \Sigma_1) \bullet (\Sigma_1 - \Sigma_2) \triangleright \text{fork } t_1; t_2 : T_2 \triangleleft \emptyset}$$

and by considering the possible cases for occurrences of $c : S$ or $c : \perp_S$ in Σ and Σ_1 , we can show that $(\Sigma - \Sigma_1) \bullet (\Sigma_1 - \Sigma_2) = \Sigma - \Sigma_2$.

C-New. We have

$$\Gamma; \Sigma; \text{new } S \mapsto \Sigma; [S]; \emptyset$$

and, noting that $\Sigma - \Sigma = \emptyset$, we obtain

$$\Gamma \vdash \emptyset \triangleright \text{new } S : [S] \triangleleft \emptyset$$

directly from rule T-NEW.

C-App. We have

$$\frac{\Gamma; v \mapsto (\Sigma; T \rightarrow U; \Sigma') \quad \Gamma; v' \mapsto T}{\Gamma; \Sigma; vv' \mapsto \emptyset; U; \Sigma'}$$

We have $\Sigma - \emptyset = \Sigma$, and Lemma 20 gives the hypotheses of the following instance of rule T-APP:

$$\frac{\Gamma \vdash v: (\Sigma; T \rightarrow U; \Sigma') \quad \Gamma \vdash v': T}{\Gamma \vdash \Sigma \triangleright vv': U \triangleleft \Sigma'}$$

C-Val. We have

$$\frac{\Gamma; v \mapsto T}{\Gamma; \Sigma; v \mapsto \Sigma; T; \emptyset}$$

Lemma 20 gives $\Gamma \vdash v: T$, and $\Sigma - \Sigma = \emptyset$, so we use rule T-VAL to obtain

$$\frac{\Gamma \vdash v: T}{\Gamma \vdash \emptyset \triangleright v: T \triangleleft \emptyset}$$

C-Let. We have

$$\frac{\Gamma; \Sigma; e \mapsto \Sigma_1; T_1; \Sigma'_1 \quad \Gamma, x: T_1; (\Sigma \cap \Sigma_1), \Sigma'_1; t \mapsto \Sigma_2; T_2; \Sigma'_2}{\Gamma; \Sigma; \text{let } x = e \text{ in } t \mapsto \Sigma_3; T_2; (\Sigma'_1 \cap \Sigma_2), \Sigma'_2}$$

where

$$\Sigma_3 = (\Sigma_1 \cap \Sigma_2) \cup \{c: S \mid (c: \perp_S \in \Sigma \text{ or } c: \perp_{\bar{S}} \in \Sigma) \text{ and } (c: S \in \Sigma_1 \text{ or } c: S \in \Sigma_2)\}.$$

The induction hypothesis gives

$$\Gamma \vdash \Sigma - \Sigma_1 \triangleright e: T_1 \triangleleft \Sigma'_1$$

and

$$\Gamma, x: T_1 \vdash ((\Sigma \cap \Sigma_1), \Sigma'_1) - \Sigma_2 \triangleright t: T_2 \triangleleft \Sigma'_2$$

Using Weakening (lemma 6) we get

$$\Gamma \vdash (\Sigma - \Sigma_1), \Sigma_4 \triangleright e: T_1 \triangleleft \Sigma'_1, \Sigma_4$$

and

$$\Gamma, x: T_1 \vdash (((\Sigma \cap \Sigma_1), \Sigma'_1) - \Sigma_2), (\Sigma'_1 \cap \Sigma_2) \triangleright t: T_2 \triangleleft (\Sigma'_1 \cap \Sigma_2), \Sigma'_2$$

where

$$\begin{aligned} \Sigma_4 &= \{c: S \mid c: S \in \Sigma_1 \cap \Sigma \text{ and } c \notin \text{dom}(\Sigma_2)\} \\ &\cup \{c: S \mid c: \perp_S \in \Sigma_1 \cap \Sigma_2 \text{ and } c: \bar{S} \in \Sigma_2\} \end{aligned}$$

By considering all the cases for a typed channel, we can show that

$$\Sigma'_1, \Sigma_4 = (((\Sigma \cap \Sigma_1), \Sigma'_1) - \Sigma_2), (\Sigma'_1 \cap \Sigma_2)$$

and

$$(\Sigma - \Sigma_1), \Sigma_4 = \Sigma - \Sigma_3$$

Therefore, using rule T-LET we get

$$\Gamma \vdash \Sigma - \Sigma_3 \triangleright \text{let } x = e \text{ in } t: T_2 \triangleleft (\Sigma'_1 \cap \Sigma_2), \Sigma'_2.$$

□

We are finally in a position to prove soundness.

Proof of theorem 4, page 18. By induction on the derivation of $\Delta; \Sigma; C \mapsto \Sigma'$. Lemma 18 guarantees that $\Sigma - \Sigma'$ is defined. We consider cases for each of the inference rules in the typechecking algorithm.

C-Thread. We have

$$\frac{\Delta; \Sigma; t \mapsto \Sigma'; T; \vec{c}: \perp_{\vec{g}}}{\Delta; \Sigma; \langle t \rangle \mapsto \Sigma'}$$

Lemma 21 gives $\Delta \vdash \Sigma - \Sigma' \triangleright t: T \triangleleft \vec{c}: \perp_{\vec{g}}$ and we use rule T-THREAD to obtain

$$\frac{\Delta \vdash \Sigma - \Sigma' \triangleright t: T \triangleleft \vec{c}: \perp_{\vec{g}}}{\Delta \vdash \Sigma - \Sigma' \triangleright \langle t \rangle}$$

C-Par. We have

$$\frac{\Delta; \Sigma; C_1 \mapsto \Sigma_1 \quad \Delta; \Sigma_1; C_2 \mapsto \Sigma_2}{\Delta; \Sigma; C_1 \mid C_2 \mapsto \Sigma_2}$$

The induction hypothesis gives the hypotheses in the following instance of rule T-PAR:

$$\frac{\Delta \vdash \Sigma - \Sigma_1 \triangleright C_1 \quad \Delta \vdash \Sigma_1 - \Sigma_2 \triangleright C_2}{\Delta \vdash (\Sigma - \Sigma_1) \bullet (\Sigma_1 - \Sigma_2) \triangleright C_1 \mid C_2}$$

and by the same reasoning as for case C-FORK in the proof of Lemma 21 we have $(\Sigma - \Sigma_1) \bullet (\Sigma_1 - \Sigma_2) = \Sigma - \Sigma_2$.

C-NewN, C-NewB, C-NewC follow directly by using the induction hypothesis. □

D Proof of Theorem 5, Completeness of Type-checking

Lemma 22 (Completeness for Values). *If $\Gamma \vdash v: T$ then $\Gamma; v \mapsto T$.*

Proof. The cases for literal values, channels, names and variables are trivial.

In the case of rule T-ABS we have a derivation finishing with

$$\frac{\Gamma, x: T \vdash \Sigma \triangleright e: U \triangleleft \Sigma'}{\Gamma \vdash \lambda(\Sigma; x: T).e: (\Sigma; T \rightarrow U; \Sigma')}$$

Applying completeness for expressions (Lemma 23) to the hypothesis, we have

$$\Gamma, x: T; \Sigma; e \mapsto \Sigma_2; U; \Sigma_3$$

where $\Sigma' = \Sigma_2, \Sigma_3$. Using C-ABS we obtain

$$\Gamma; \lambda(\Sigma; x: T).e \mapsto (\Sigma; T \rightarrow U; \Sigma_2, \Sigma_3)$$

as required.

The case of rule T-REC follows directly from the induction hypothesis. \square

Lemma 23 (Completeness for Expressions). *If $\Gamma \vdash \Sigma \triangleright e: T \triangleleft \Sigma_1$ and $\Sigma \bullet \Sigma'$ is defined then $\Gamma; \Sigma \bullet \Sigma'; e \mapsto \Sigma_2 \bullet \Sigma'; T; \Sigma_3$ where $\Sigma_1 = \Sigma_2, \Sigma_3$.*

Proof. By induction on the derivation of $\Gamma \vdash \Sigma \triangleright e: T \triangleleft \Sigma_1$ with a case for each inference rule.

For T-SENDD we have

$$\frac{\Gamma \vdash v: D \quad \Gamma \vdash v': \text{Chan } c}{\Gamma \vdash \Sigma, c: !D.S \triangleright \text{send } v \text{ on } v': \text{Unit} \triangleleft \Sigma, c: S}$$

There are two cases to consider, depending on the structure of Σ' .

1. If $c \notin \text{dom}(\Sigma')$ then $(\Sigma, c: !D.S) \bullet \Sigma' = (\Sigma \bullet \Sigma'), c: !D.S$ and by using the first form of C-SEND we obtain

$$\Gamma; (\Sigma \bullet \Sigma'), c: !D.S; \text{send } v \text{ on } v' \mapsto \Sigma \bullet \Sigma'; \text{Unit}; c: S.$$

2. If $c: ?D.\bar{S} \in \Sigma'$ then $\Sigma' = \Sigma'', c: ?D.\bar{S}$ and $(\Sigma, c: !D.S) \bullet \Sigma' = (\Sigma \bullet \Sigma''), c: \perp$ where \perp is either $\perp_{!D.S}$ or $\perp_{?D.\bar{S}}$. By using the second or third form of C-SEND we obtain

$$\Gamma; (\Sigma \bullet \Sigma''), c: \perp; \text{send } v \text{ on } v' \mapsto (\Sigma \bullet \Sigma''), c: ?D.\bar{S}; \text{Unit}; c: S$$

and we have $\Sigma \bullet \Sigma' = (\Sigma \bullet \Sigma''), c: ?D.\bar{S}$.

For T-FORK we have

$$\frac{\Gamma \vdash \Sigma_1 \triangleright t_1: T_1 \triangleleft \emptyset \quad \Gamma \vdash \Sigma_2 \triangleright t_2: T_2 \triangleleft \emptyset}{\Gamma \vdash \Sigma_1 \bullet \Sigma_2 \triangleright \text{fork } t_1; t_2: T_2 \triangleleft \emptyset}$$

By the induction hypothesis,

$$\Gamma; \Sigma_1 \bullet \Sigma_2 \bullet \Sigma'; t_1 \mapsto \Sigma_2 \bullet \Sigma'; T_1; \emptyset$$

and

$$\Gamma; \Sigma_2 \bullet \Sigma'; t_2 \mapsto \Sigma'; T_2; \emptyset$$

so by rule C-FORK we obtain

$$\Gamma; \Sigma_1 \bullet \Sigma_2 \bullet \Sigma'; \text{fork } t_1; t_2 \mapsto \Sigma_2; T_2; \emptyset$$

The case of T-LET is the most complex. We have

$$\frac{\Gamma \vdash \Sigma \triangleright e: T_1 \triangleleft \Sigma'' \quad \Gamma, x: T_1 \vdash \Sigma'' \triangleright t: T_2 \triangleleft \Sigma'}{\Gamma \vdash \Sigma \triangleright \text{let } x = e \text{ in } t: T_2 \triangleleft \Sigma'}$$

Changing notation, for clarity, we must show that

$$\Gamma; \Sigma \bullet \Phi'; \text{let } x = e \text{ in } t \mapsto \Psi_3 \bullet \Phi'; T_2; \Phi'_3$$

where $\Sigma' = \Psi_3, \Phi'_3$.

Rule C-LET shows the recursive calls to the typechecker, and their results:

$$\frac{\Gamma; \Sigma \bullet \Phi'; e \mapsto \Phi_1; T_1; \Phi'_1 \quad \Gamma, x: T_1; ((\Sigma \bullet \Phi') \cap \Phi_1), \Phi'_1; t \mapsto \Phi_2; T_2; \Phi'_2}{\Gamma; \Sigma \bullet \Phi'; \text{let } x = e \text{ in } t \mapsto \Phi_3; T_2; (\Phi'_1 \cap \Phi_2), \Phi'_2}$$

where $\Phi_3 = (\Phi_1 \cap \Phi_2) \cup \{c: S \mid (c: \perp_S \in \Sigma \bullet \Phi' \text{ or } c: \perp_{\bar{S}} \in \Sigma \bullet \Phi') \text{ and } (c: S \in \Phi_1 \text{ or } c: S \in \Phi_2)\}$.

Applying the induction hypothesis to the first hypothesis gives $\Phi_1 = \Psi_1 \bullet \Phi'$ with $\Sigma'' = \Psi_1, \Phi'_1$.

In order to apply the induction hypothesis to the second hypothesis, we must find Φ'' such that $((\Sigma \bullet \Phi') \cap \Phi_1), \Phi'_1 = \Sigma'' \bullet \Phi'$. We then get $\Phi_2 = \Psi_2 \bullet \Phi''$ with $\Sigma' = \Psi_2, \Phi'_2$.

We take $\Phi'' = \Phi' - \{c: \Phi'(c) \mid c \in \text{dom}(\Phi') \cap \text{dom}(\Phi'_1)\}$.

By considering all possible cases for a typed channel $c: S$, we can show the necessary condition on Φ'' , and then prove that $\Phi_3 = \Psi_3 \bullet \Phi'$ with $\Sigma' = \Psi_3, (\Phi'_1 \cap \Phi_2), \Phi'_2$ as required. \square

Theorem 5 (Completeness). *If $\Delta \vdash \Sigma \triangleright C$, then for all Σ' such that $\Sigma \bullet \Sigma'$ is defined, $\Delta; \Sigma \bullet \Sigma'; C \mapsto \Sigma'$.*

Proof. By induction on the derivation of $\Delta \vdash \Sigma \triangleright C$, considering a case for each typing rule.

For T-THREAD we have

$$\frac{\Delta \vdash \Sigma \triangleright t: T \triangleleft \emptyset}{\Delta \vdash \Sigma \triangleright \langle t \rangle}$$

By Lemma 23 we have

$$\Delta; \Sigma \bullet \Sigma'; t \mapsto \Sigma'; T; \emptyset$$

and rule C-THREAD gives the required

$$\Delta; \Sigma \bullet \Sigma'; \langle t \rangle \mapsto \Sigma'.$$

For T-PAR we have

$$\frac{\Delta \vdash \Sigma_1 \triangleright C_1 \quad \Delta \vdash \Sigma_2 \triangleright C_2}{\Delta \vdash \Sigma_1 \bullet \Sigma_2 \triangleright C_1 \mid C_2}$$

By the induction hypothesis we have

$$\Delta; \Sigma_1 \bullet \Sigma_2 \bullet \Sigma'; C_1 \mapsto \Sigma_2 \bullet \Sigma'$$

and

$$\Delta; \Sigma_2 \bullet \Sigma'; C_2 \mapsto \Sigma'$$

and by using rule C-PAR we obtain

$$\Delta; \Sigma_1 \bullet \Sigma_2 \bullet \Sigma'; C_1 \mid C_2 \mapsto \Sigma'$$

as required.

The cases of T-NEWN, T-NEWB and T-NEWC follow directly from the induction hypothesis. \square