

Multiparty session types, beyond duality

Alceste Scalas*, Nobuko Yoshida*

Imperial College London, UK



ARTICLE INFO

Article history:

Received 16 August 2017

Accepted 22 January 2018

Available online 10 March 2018

Keywords:

Concurrency

Process calculi

Multiparty session types

Duality

ABSTRACT

Multiparty Session Types (MPST) are a well-established typing discipline for message-passing processes interacting on *sessions* involving two or more participants. Session typing can ensure desirable properties: absence of communication errors and deadlocks, and protocol conformance.

We propose a novel MPST theory based on a *rely/guarantee typing system*, that checks (1) the *guaranteed* behaviour of the process being typed, and (2) the *relied upon* behaviour of other processes. Crucially, our theory achieves type safety by enforcing a typing context *liveness* invariant throughout typing derivations.

Unlike “classic” MPST works, our typing system does not depend on global session types, and does not use syntactic *duality* checks. As a result, our new theory can prove type safety for processes that implement protocols with complex inter-role dependencies, thus sidestepping an intrinsic limitation of “classic” MPST.

© 2018 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Session types are a typing discipline for *processes* that exchange messages via *channels*, and implement *sessions* with structured protocols featuring inputs/outputs, choices, and recursion. The goal of session typing is to statically ensure that a given process correctly implements a given protocol, without causing communication errors or deadlocks at run-time. Years of research have shown that the session-based framework is applicable to a wide range of calculi, programming languages, and computing environments [23,1,16].

Binary session types In their original formulation, session types describe *binary* sessions involving exactly two participants, or *roles* [20,19,35]. In binary session theories, types are similar to $S = \oplus m_1 . \& m_2$, meaning “send message m_1 , and then receive message m_2 .” Typing judgements have the following form:

$$\vdash P \triangleright \Delta \quad (1)$$

meaning that process P uses its communication channels according to the typing context Δ – which in turn maps channels to session types. Hence, when Δ contains the mapping $c:S$, the judgement (1) means that P uses the channel c first to send m_1 , and then to receive m_2 – as prescribed by S .

Notably, the typing judgement (1) does not only determine the behaviour of P : it also implicitly describes how *other* processes in P 's environment can send/receive messages to/from P . This is because in a well-typed system, each input/output

* Corresponding author.

E-mail addresses: alceste.scalas@imperial.ac.uk (A. Scalas), n.yoshida@imperial.ac.uk (N. Yoshida).

of P must match a corresponding output/input of some other process playing the unique “opposite” role in the corresponding binary session. This complementarity between session endpoints is formalised as *duality*: a fundamental element of the binary session types theory, inspired by linear logic [18] as originally stated in [19], and subsequently studied in various works (e.g., [9,36,37,10]). By duality, if P interacts with another process Q through some channels, then we can infer the following typing judgement:

$$\vdash Q \triangleright \overline{\Delta} \quad (2)$$

where $\overline{\Delta}$ is the *dual* of Δ : it contains the same channels, but mapped to dual types where each input is turned into a corresponding output, and each output into an input. For example, if Δ contains the mapping $c:S$ above, then we can infer that $\overline{\Delta}$ contains $c:\overline{S}$, where $\overline{S} = \&m_1.\oplus m_2$ – meaning “use channel c to receive message m_1 , and then send m_2 .” Now, since (2) guarantees that Q behaves dually w.r.t. P , we know that the parallel composition of P and Q will execute correctly: e.g., P and Q will use channel c to exchange messages m_1 and then m_2 , according to types S and \overline{S} .

Multiparty session types Binary session types do not allow to directly formalise and verify protocols with more than two interacting roles. This limitation has been addressed by *Multiparty Session Types (MPST)*, which were introduced in [21] and have been actively studied both in theory and applications (see related work in [22]). In a nutshell, the MPST framework foster a top-down approach where:

- (1) a *choreography* involving two or more interacting *roles* is formalised as a *global type* G ;
- (2) G is *projected* onto a set of (*local*) *session types* S_1, \dots, S_n (one per role of G), each one describing the expected inputs/outputs of a specific role from/to other roles in G ;
- (3) the session types S_1, \dots, S_n are assigned to *channels*, used by *MPST π -calculus processes* that are type-checked against those types (i.e., S_1, \dots, S_n).

In the multiparty setting, typing judgements still look like (1); however, types contain role annotations, and look like $S' = \text{alice}\&m_1.\text{bob}\&m_2.\text{carol}\oplus m_3$, meaning “receive m_1 from *alice*, then receive m_2 from *bob*, then send m_3 to *carol*.” As a consequence, the duality-based reasoning outlined above (for binary session types) does not work, and duality cannot be used directly. This is because in a multiparty session, P might interact with *more than one* other processes playing different roles in a session (e.g., *alice*, *bob* and *carol*) – and correspondingly, the multiparty typing context Δ would contain types requiring inputs/outputs from/to multiple roles (like S'). Hence, instead of a single “other process” Q as in (2), we might have n other typed processes Q_1, \dots, Q_n playing distinct roles:

$$\vdash Q_1 \triangleright \Gamma_1 \quad \vdash Q_2 \triangleright \Gamma_2 \quad \dots \quad \vdash Q_n \triangleright \Gamma_n \quad (3)$$

Furthermore, Q_1, \dots, Q_n might exchange messages between each other – and these inputs/outputs, albeit occurring in a multiparty session involving P , would not be “seen” by P . This means that such message exchanges would be described by the types in $\Gamma_1, \dots, \Gamma_n$, but could not be inferred by “dualising” Δ , unlike the binary case. This is problematic because such invisible communications might determine which future inputs/outputs Q_1, \dots, Q_n will be willing to receive/send from/to P .

In the MPST theory, this lack of contextual information causes difficulties when proving *subject reduction* – i.e., when proving that typed processes reduce type-safely. These difficulties are addressed by imposing a “generalised duality” restriction, called *consistency*, that must hold for *all* typing contexts $\Delta, \Gamma_1, \dots, \Gamma_n$ – taken both in isolation, or in combination. Consistency ensures that well-typed processes will not engage in incompatible input/outputs at run-time; but unfortunately, consistency also severely restricts the set of MPST processes that can be typed and/or proven type-safe: in particular, complex protocols with recursion and message dependencies are often non-consistent, and thus, unsupported by the “classic” MPST theory.¹

Contributions We present a new multiparty session typing system that sidesteps the difficulties and limitations of consistency in “classic” MPST. Its typing judgement has the following form:

$$\vdash P \triangleright \Delta \triangleleft \Gamma$$

and intuitively reads: P uses its channels as specified by the *guarantee context* Δ , assuming that P is composed with other processes that use their channels according to the *rely context* Γ . For example, to type Q_1 in (3), the guarantee context is Γ_1 , while the rely context is $\Gamma_2, \dots, \Gamma_n$, abstracting the behaviour of other processes composed with Q_1 (in this case, Q_2, \dots, Q_n). With this rely/guarantee approach, our typing judgement conveys a global picture of the possible interactions between a process P and its environment, that is used to verify the absence of communication errors, and determine which processes can safely interact with P .

¹ The reason for the consistency constraint, and its limitations, are rather technical; we discuss them in detail in Section 3.

Channels	$c, d ::= x \mid s[\mathfrak{p}]$	(variable, channel with role \mathfrak{p})
Declarations	$D ::= X(\tilde{x}) = P$	(declaration of process variable X as P)
Processes	$P, Q ::= c[\mathfrak{q}] \oplus_m \langle d \rangle . P$ $c[\mathfrak{q}] \sum_{i \in I} \{m_i(x_i) . P_i\}$ $P \mid Q$ $(\nu s) P$ def D in P $X(\tilde{c})$ $\mathbf{0}$	(selection towards role \mathfrak{q}) (branching from role \mathfrak{q} , with $I \neq \emptyset$) (parallel composition) (session restriction/creation) (process definition) (process call) (terminated process)

Fig. 1. Syntax of the standard MPST π -calculus. Here, \tilde{x} represents a sequence x_1, \dots, x_n , and \tilde{c} represents a sequence c_1, \dots, c_n (for some $n \geq 0$).

Crucially, our typing system replaces consistency with a *liveness* constraint for typing contexts. Since liveness is a *behavioural* (rather than syntactic) property, our theory does *not* depend on the existence of global types: hence, it does not depend on the associated technicalities (necessary, e.g., to compute *global/local type projections and merges*), and supports multiparty protocols that cannot be projected from any global type.

Outline of the paper In Section 2, we present the standard definition of multiparty session processes, types, and typing contexts. In Section 3 we reprise the discussion above about consistency in “classic” MPST, formally explaining its limitations. In Section 4 we present our novel rely/guarantee multiparty session typing system (Definition 4.5), and its subject reduction property (Theorem 4.9), that does *not* depend on consistency. In Section 5, we highlight the differences between our new theory and “classic” MPST: we show that we support all protocols projected from any global type – and for this purpose, in Section 5.1 we relate liveness with a property of Communicating Finite-State Machines [8] called *Multiparty Compatibility*; in addition, we show that we also support protocols that *cannot* be projected from any global type. In Section 6, we conclude and discuss related works. Detailed proofs are available in the appendices.

2. Multiparty sessions: standard calculus, types and typing contexts

In this section, we summarise some key elements of the “classic” Multiparty Session Types (MPST) theory. We describe the multiparty session π -calculus and types, adopting a notation roughly based on Coppo et al. [11] and Scalas et al. [33] (i.e., the most common formulation in MPST literature). We focus on a streamlined presentation, that will be sufficient for highlighting the limitations of “classic” MPST (Section 3) and for developing our new theory (Section 4).

We present the MPST π -calculus in Section 2.1, and session types and typing contexts in Section 2.2.

2.1. Multiparty session π -calculus

The multiparty session π -calculus is a variant of the π -calculus [30,29,32] where channels allow two *or more* roles to exchange messages.

Definition 2.1 (MPST π -calculus). MPST processes have the syntax shown in Fig. 1. Session restriction, branching, and declarations act as binders, as expected. We write $\text{fc}(P)$ for the *free channels with roles in* P , $\text{dpv}(D)$ for the *process variables declared in* D , $\text{fpv}(D)$ for the *free process variables in* D , and $\text{fpv}(P)$ for the *free process variables in* P . We write $s \notin \text{fc}(P)$ iff there is no \mathfrak{p} such that $s[\mathfrak{p}] \in \text{fc}(P)$.

A **channel** c or d can be either a **variable** x , or a **channel with role** $s[\mathfrak{p}]$: the latter represents a *multiparty channel endpoint* used to play the *role* \mathfrak{p} in the *session* s , and send/receive messages to/from other roles in s ; variables, instead, are placeholders for channels with role, and are instantiated at run-time (details below). A **message** consists of a *label* m and a *payload*: different labels allow to distinguish different kinds of messages, while the payload is the actual data transmitted between processes. Note that payloads are channels themselves (and with minor extensions, they can be values like strings or integers – see Remark 2.2 later on). A **process** P or Q can be either:

- a **selection** that uses c to send a message with label m and payload d to \mathfrak{q} , and then continues as P ;
- a **branching** that uses c to receive from \mathfrak{q} a message with label m_i (for any $i \in I$); then, depending on which m_i has been received, the process instantiates the variable x_i with the message payload, and continues as P_i ;
- a **parallel composition** where P and Q run concurrently, and possibly communicate through their channels;
- a **session restriction** (or **session creation**) delimiting the scope of s to P ;
- a **process definition** declaring X (with *parameters* \tilde{x}) as P , with scope Q ;
- a **process call** where X is invoked with *actual parameters* \tilde{c} ;
- a **terminated process** $\mathbf{0}$.

Remark 2.2 (Abbreviations and simplifications). For brevity, when writing MPST π -calculus processes we will often omit trailing $\mathbf{0}$ s and immaterial message payloads, i.e.:

(a) $c[\mathfrak{p}] \oplus_m . P$ stands for $(\nu s)(c[\mathfrak{p}] \oplus_m (s[\mathfrak{q}]) . P)$ with some \mathfrak{q} and $s \notin \text{fc}(P)$;

$$\begin{aligned}
& \text{[R-COMM]} \quad s[\mathbf{p}][\mathbf{q}] \sum_{i \in I} \{m_i(x_i).P_i \mid s[\mathbf{q}][\mathbf{p}] \oplus_{m_k} (s'[\mathbf{r}]).Q \rightarrow P_k\{s'[\mathbf{r}]/x_k\} \mid Q \quad (k \in I) \\
\text{[R-X]} \quad \mathbf{def} X(x_1, \dots, x_n) = P \mathbf{in} (X\{s_1[\mathbf{p}_1], \dots, s_n[\mathbf{p}_n]\} \mid Q) & \rightarrow \mathbf{def} X(x_1, \dots, x_n) = P \mathbf{in} (P\{s_1[\mathbf{p}_1]/x_1\} \dots \{s_n[\mathbf{p}_n]/x_n\} \mid Q) \\
\text{[R-] } P \rightarrow P' \text{ implies } P \mid Q & \rightarrow P' \mid Q \quad \text{[R-}\nu\text{]} P \rightarrow P' \text{ implies } (\nu s) P \rightarrow (\nu s) P' \\
\text{[R-def]} P \rightarrow P' \text{ implies } \mathbf{def} X(\tilde{x}) = Q \mathbf{in} P & \rightarrow \mathbf{def} X(\tilde{x}) = Q \mathbf{in} P' \\
\text{[R=]} P' \equiv P \rightarrow Q \equiv Q' \text{ implies } P' & \rightarrow Q' \\
\\
P \mid Q \equiv Q \mid P \quad (P \mid Q) \mid R \equiv P \mid (Q \mid R) \quad P \mid \mathbf{0} \equiv P \quad (\nu s) \mathbf{0} \equiv \mathbf{0} \quad (\nu s) (\nu s') P \equiv (\nu s') (\nu s) P \\
(\nu s) (P \mid Q) \equiv P \mid (\nu s) Q \quad \text{if } s \notin \text{fc}(P) \quad \mathbf{def} D \mathbf{in} \mathbf{0} \equiv \mathbf{0} \quad \mathbf{def} D \mathbf{in} (\nu s) P \equiv (\nu s) (\mathbf{def} D \mathbf{in} P) \quad \text{if } s \notin \text{fc}(D) \\
\mathbf{def} D \mathbf{in} (P \mid Q) \equiv (\mathbf{def} D \mathbf{in} P) \mid Q \quad \text{if } \text{dpv}(D) \cap \text{fpv}(Q) = \emptyset \\
\mathbf{def} D \mathbf{in} (\mathbf{def} D' \mathbf{in} P) \equiv \mathbf{def} D' \mathbf{in} (\mathbf{def} D \mathbf{in} P) \quad \text{if } (\text{dpv}(D) \cup \text{fpv}(D)) \cap \text{dpv}(D') = (\text{dpv}(D') \cup \text{fpv}(D')) \cap \text{dpv}(D) = \emptyset
\end{aligned}$$

Fig. 2. Standard MPST π -calculus: semantics (top) and structural congruence \equiv (bottom).

(b) $c[\mathbf{p}] \sum \{m_1.P_1, m_2(x).P_2\}$ stands for $c[\mathbf{p}] \sum \{m_1(z).P_1, m_2(x).P_2\}$ with $z \notin \text{fv}(P_1)$.

Furthermore, note that the calculus in Definition 2.1 can be easily extended with, e.g., basic values (strings, integers, booleans...) or conditionals. Such extensions are routine and orthogonal to our treatment; we will sometimes use them in examples, to make them more readable.

The abbreviations described in Remark 2.2 omit either (a) a channel with role $s[\mathbf{p}]$ that is sent just after s is created and just before s is discarded, or (b) a message payload variable z that is never referenced. We will see that the two abbreviations are often combined: in case (a) the message payload $s[\mathbf{p}]$ does not allow to interact with any other role (hence, the only “meaningful” element of the message is the label m , that can trigger a corresponding branch of the recipient process); correspondingly, the recipient of m can safely discard its payload, as in case (b).

Example 2.3. Below is an example of MPST π -calculus process:

$$\begin{aligned}
P_a &= s[\mathbf{alice}][\mathbf{bob}] \oplus \text{stop} \\
(\nu s) (P_a \mid P_b \mid P_c) \quad \text{where:} \quad P_b &= s[\mathbf{bob}][\mathbf{alice}] \sum \left\{ \begin{array}{l} m_1(x).s[\mathbf{bob}][\mathbf{carol}] \oplus m_2(\text{"Hi"}), \\ \text{stop}.s[\mathbf{bob}][\mathbf{carol}] \oplus \text{quit} \end{array} \right\} \\
P_c &= s[\mathbf{carol}][\mathbf{bob}] \sum \left\{ \begin{array}{l} m_2(y).s[\mathbf{carol}][\mathbf{alice}] \oplus m_3(\text{true}), \\ \text{quit} \end{array} \right\}
\end{aligned} \tag{4}$$

Here, $(\nu s) (P_a \mid P_b \mid P_c)$ is the parallel composition of processes P_a , P_b , and P_c in the scope of the *multiparty session* s . Inside such processes, $s[\mathbf{alice}]$ (resp. $s[\mathbf{bob}]$, $s[\mathbf{carol}]$) is a *channel with role*: it is used to send/receive messages on s , while playing the role of *alice* (resp. *bob*, *carol*). In the process P_a , the selection “ $s[\mathbf{alice}][\mathbf{bob}] \oplus \text{stop}$ ” means: the channel with role $s[\mathbf{alice}]$ is used to play role *alice* in session s , and send stop to *bob*. In process P_b , the channel with role $s[\mathbf{bob}]$ is used to play role *bob* in s and receive (\sum) either m_1 or stop from *alice*; then, in the first case, $s[\mathbf{bob}]$ is used to send m_2 to *carol*; in the second case, $s[\mathbf{bob}]$ is used to send quit to *carol*.

Note that in (4), the trailing $\mathbf{0}$ s and some message payloads are omitted, as explained in Remark 2.2.

We adopt the standard *reduction relation* \rightarrow for MPST π -calculus processes, formalised in Definition 2.4.

Definition 2.4 (MPST π -calculus congruence and semantics). The *structural congruence* \equiv between MPST π -calculus processes is inductively defined by the rules in Fig. 2 (bottom). The *reduction relation* \rightarrow between MPST π -calculus processes is inductively defined as shown in Fig. 2 (top). We write \rightarrow^* to denote the reflexive and transitive closure of \rightarrow .

Fig. 2 contains the following rules:

- **[R-COMM]** is the **communication rule**: it says that the parallel composition of a branching process and a selection process operating on the same session s , and playing respectively the roles \mathbf{p} and \mathbf{q} , can reduce provided that (1) the two processes are targeting each other (i.e., the branching is from \mathbf{q} and the selection is towards \mathbf{p}), (2) the branching process supports the selected label m_k , and (3) the selection payload is a channel with role $s'[\mathbf{r}]$ (not a variable). If these conditions hold, the two processes reduce to the respective continuations P_k and Q – with $s'[\mathbf{r}]$ replacing x_k in the former;
- **[R-X]** is the **process call rule**: inside the scope of the process declaration $X(\tilde{x}) = P$, calling X amounts to expanding P and replacing the formal parameters \tilde{x} with the actual parameters of the call (that must be channels with roles, and not variables);

- rules [R-|], [R-ν], and [R-def] are the **parallel, restriction, and process definition rules**: they say that reductions can occur respectively inside parallel compositions, session restrictions and process definitions;
- [R-≡] is the **up-to congruence rule**: it says that if P reduces to Q , then any $P' \equiv P$ reduces to any $Q' \equiv Q$. The **process congruence**, in turn, is standard: it says that $|$ is a commutative monoid with $\mathbf{0}$ as neutral element, and garbage-collects unnecessary session restrictions and process declarations, or reorders them without name capturing.

Example 2.5. Take the process in Example 2.3. By Definition 2.4, it reduces as follows:

$$(\nu s)(P_a | P_b | P_c) \rightarrow (\nu s)(\mathbf{0} | s[\text{bob}][\text{carol}] \oplus \text{quit} | P_c) \rightarrow (\nu s)(\mathbf{0} | \mathbf{0} | \mathbf{0}) \equiv (\nu s)\mathbf{0} \equiv \mathbf{0}$$

2.2. Multiparty session types and typing contexts

Similarly to “classic” MPST works, our new typing system (introduced later on, in Section 4) will assign *session types* (Definition 2.6 below) to channels, using *typing contexts* (Definition 2.9). For both, we adopt standard definitions from MPST literature.

Definition 2.6 (Multiparty session types). *Session types* (ranged over by S, T) have the syntax:

$$S, T ::= \mathfrak{p} \&_{i \in I} \mathfrak{m}_i(S_i) \cdot S'_i \mid \mathfrak{p} \oplus_{i \in I} \mathfrak{m}_i(S_i) \cdot S'_i \mid \mathbf{end} \mid \mu \mathfrak{t}. S \mid \mathfrak{t} \quad \text{with } I \neq \emptyset$$

where \mathfrak{m}_i range over pairwise distinct *message labels*. We require recursive types to be *contractive*: i.e., in $\mu \mathfrak{t}. S$ we have $S \neq \mathfrak{t}' (\forall \mathfrak{t}')$. We define \equiv as the *coinductive equivalence between type trees*, such that $\mu \mathfrak{t}. S \equiv S\{\mu \mathfrak{t}. S/\mathfrak{t}\}$ (i.e., a recursive type is equivalent to its unfolding, cf. Pierce [31, §21.7 and §21.8]).

In Definition 2.6, the **branching type** (or *external choice*) $\mathfrak{p} \&_{i \in I} \mathfrak{m}_i(S_i) \cdot S'_i$ says that a channel must be used to receive from \mathfrak{p} one input of the form $\mathfrak{m}_i(S_i)$, for any $i \in I$ chosen by \mathfrak{p} ; then, the channel must be used following the *continuation type* S'_i . The **selection type** (or *internal choice*) $\mathfrak{p} \oplus_{i \in I} \mathfrak{m}_i(S_i) \cdot S'_i$, instead, says that a channel must be used to perform one output $\mathfrak{m}_i(S_i)$ towards \mathfrak{p} , for some $i \in I$, and then according to S'_i . The type **end** describes a **terminated** channel that cannot be further used for inputs/outputs. Finally, $\mu \mathfrak{t}. S$ is a **recursive** session type: μ binds the *recursion variable* \mathfrak{t} in S . We distinguish a recursive type from its unfolding; when necessary, we explicitly relate them with the equivalence \equiv .

Remark 2.7. For brevity, we will often omit the trailing **end** in types, and **end**-typed message payloads: e.g., $\mathfrak{p} \oplus \mathfrak{m}$ stands for $\mathfrak{p} \oplus \mathfrak{m}(\mathbf{end}) \cdot \mathbf{end}$.

For readability, we will sometimes write examples using ground types like *Int*, *Bool*, etc.: their addition to Definition 2.6 is standard, and we omit it for simplicity.

Example 2.8. Consider the following session types:

$$\begin{aligned} S_a &= \text{bob} \oplus \left\{ \begin{array}{l} \text{m1(Int) . carol} \& \text{m3(Bool) . end ,} \\ \text{stop . end} \end{array} \right\} & S_b &= \text{alice} \& \left\{ \begin{array}{l} \text{m1(Int) . carol} \oplus \text{m2(Str) . end ,} \\ \text{stop . carol} \oplus \text{quit . end} \end{array} \right\} \\ S_c &= \text{bob} \& \left\{ \begin{array}{l} \text{m2(Str) . alice} \oplus \text{m3(Bool) . end ,} \\ \text{quit . end} \end{array} \right\} \end{aligned} \quad (5)$$

Intuitively, S_a says that a process using an S_a -typed channel must send (\oplus) to *bob* a message that can be either m1(Int) , or stop ; in the first case, the process must then use the same channel to receive ($\&$) message m3(Bool) from *carol*, after which the session **ends**; otherwise, in the second case, the session **ends** without further interactions. The other types follow a similar intuition.

Note that in (5), **end**-typed message payloads are omitted, and we use ground types as explained in Remark 2.7.

We also adopt the standard definitions of *MPST typing context* (Definition 2.9) and typing context reduction (Definition 2.10).

Definition 2.9 (Typing context). A *multiparty session typing context*, denoted by Δ or Γ , is a partial mapping defined as:

$$\Delta, \Gamma ::= \emptyset \mid \Delta, x:S \mid \Delta, s[\mathfrak{p}]:S$$

- The *typing context composition* Δ_1, Δ_2 is defined iff $\text{dom}(\Delta_1) \cap \text{dom}(\Delta_2) = \emptyset$.
- We write $\Delta \equiv \Delta'$ iff $\text{dom}(\Delta) = \text{dom}(\Delta')$ and $\forall c \in \text{dom}(\Delta) : \Delta(c) \equiv \Delta'(c)$.
- We define $\Delta \setminus c$ so that $(\Delta \setminus c)(d) = \Delta(d)$ iff $d \neq c$ (and is undefined otherwise).

As anticipated in Section 1, a multiparty session typing context maps channels (either variables or channels with role) to session types, thus specifying how a channel is expected to be used. The typing context composition Δ_1, Δ_2 is defined iff the domains of Δ_1 and Δ_2 do not overlap.² The notation $\Delta \setminus c$ stands for the typing context defined as Δ , minus its entry for c (if present).

Multiparty session typing contexts are equipped with reduction semantics: this feature can be found in many typing systems for process calculi, when types abstract the behaviour of processes (for a survey, see [23]).

Definition 2.10 (Typing context reduction). The typing contexts reduction relation \rightarrow is inductively defined by the following rules:

$$\frac{k \in I \quad S_k \equiv T_k}{s[p]:q\oplus_{i \in I} m_i(S_i).S'_i, s[q]:p\&_{i \in I \cup J} m_i(T_i).T'_i \rightarrow s[p]:S'_k, s[q]:T'_k} \text{ [T-COMM]}$$

$$\frac{k \in I}{x:p\oplus_{i \in I} m_i(S_i).S'_i \rightarrow x:S'_k} \text{ [T-x}\oplus] \quad \frac{k \in I}{x:p\&_{i \in I} m_i(S_i).S'_i \rightarrow x:S'_k} \text{ [T-x}\&]$$

$$\frac{\Delta, c:S\{\mu t.S/t\} \rightarrow \Delta'}{\Delta, c:\mu t.S \rightarrow \Delta'} \text{ [T-}\mu] \quad \frac{\Delta \rightarrow \Delta'}{\Delta, c:S \rightarrow \Delta', c:S} \text{ [T-COMP]}$$

We write \rightarrow^* to denote the reflexive and transitive closure of \rightarrow .

The rules in Definition 2.10 have the following meaning:

- [T-COMM] says that two entries $s[p]:S_p, s[q]:S_q$ can “interact”, and reduce to their continuations, provided that (1) S_p is a selection towards q , (2) S_q is a branching from p , and (3) the message labels of S_q are a superset of those in S_p , with equivalent carried types. The last clause implies that if S_p has some output message that is not supported by S_q , then the pair $s[p]:S_p, s[q]:S_q$ is stuck;
- [T-x \oplus] and [T-x $\&$] says that an entry $x:S$ (i.e., mapping a variable to a type) can reduce autonomously. These rules are a minor extension w.r.t. “classic” MPST papers; they will be useful later on, to define our typing system (Section 4) and formalise some results (Section 5.3);
- [T- μ] says that a recursive type can take part in a reduction, provided that its unfolding can induce the same reduction;
- [T-COMP] says that a reduction can occur within a larger typing context, with each additional entry $c:S$ left unchanged.

Example 2.11. Consider the typing context:

$$\Delta_\mu = s[p]:\mu t.q\oplus m.t, s[q]:\mu t.p\& m.t$$

By Definition 2.10, it reduces as follows (for clarity, we restore some **end** payloads omitted above):

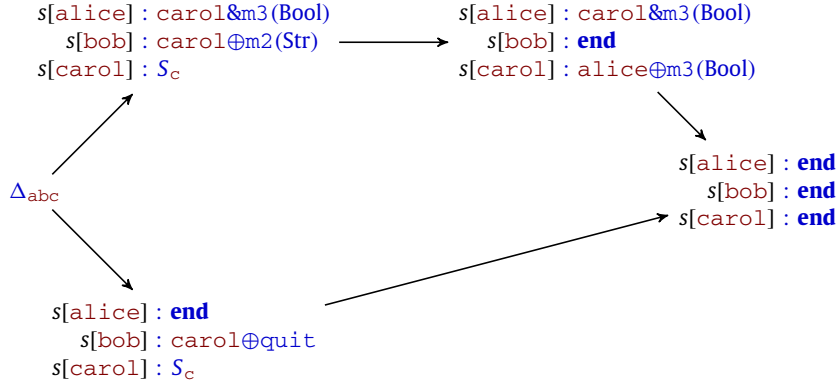
$$\frac{\frac{s[p]:q\oplus m(\mathbf{end}).\mu t.q\oplus m.t, s[q]:p\& m(\mathbf{end}).\mu t.p\& m.t \rightarrow \Delta_\mu \text{ [T-COMM]}}{s[p]:q\oplus m.\mu t.q\oplus m.t, s[q]:\mu t.p\& m.t \rightarrow \Delta_\mu \text{ [T-}\mu]}}{\Delta_\mu \rightarrow \Delta_\mu} \text{ [T-}\mu]$$

Example 2.12. Consider the types S_a, S_b , and S_c in Example 2.8, and the following typing context:

$$\Delta_{abc} = s[\mathbf{alice}]:S_a, s[\mathbf{bob}]:S_b, s[\mathbf{carol}]:S_c$$

By Definition 2.10, Δ_{abc} can reduce as follows, by rules [T-COMM] and [T-COMP]:

² As in many MPST works, we could slightly relax the condition by allowing Δ_1 and Δ_2 to overlap on **end**-typed entries (i.e., treat such entries as *unrestricted* rather than *linear*). For simplicity, we treat all typing context entries in the same way. Note that this does not impact the expressiveness of the typed fragment of our MPST calculus: in fact, **end**-typed entries can be created “for free” when necessary, using session restriction.



In the upper part of the transition diagram, $s[\text{alice}]$ sends m_1 to $s[\text{bob}]$, who sends m_2 to $s[\text{carol}]$, who sends m_3 to $s[\text{alice}]$. In the lower part of the diagram, $s[\text{alice}]$ sends stop to $s[\text{bob}]$, who sends quit to $s[\text{carol}]$. In both cases, the reductions reach a final typing context configuration where all entries are **end**-typed.

3. “Classic” multiparty session typing system and subject reduction

In this section, we outline the “classic” multiparty session typing system, and its limitations. We reprise our claims in Section 1 more formally: we discuss “classic” multiparty session typing (Section 3.1), subject reduction (Section 3.2), and the limitations introduced by consistency (Section 3.3). Moreover, in Section 3.4, we discuss further cases of protocols that are “safe” but non-consistent: the fact that “classic” MPST do *not* cover such simple examples suggests that a more flexible approach is desirable.

We will provide pointers to some definitions in Appendix A, that are technically necessary in “classic” MPST theory. We relegate such technicalities in the appendix to make this section more readable – and also because they will *not* be needed for our new typing system in Section 4.

3.1. “Classic” multiparty session typing, global types, and projections

The “classic” multiparty session typing judgements have the form shown in Section 1, eq. (1). We repeat it here for clarity:

$$\vdash P \triangleright \Delta \quad \text{with } \Delta \text{ from Definition 2.9} \quad (6)$$

which reads: “ P uses its channels as described in Δ ”.

Under the typical “top-down” approach outlined in Section 1, the typing context Δ usually contains multiparty session types that are *projected from some given global type*, whose syntax is formalised in Definition 3.1 below.

Definition 3.1 (Global types). The syntax of a global type G is:

$$G ::= p \rightarrow q : \{m_i(T_i) . G_i\}_{i \in I} \mid \mu t . G \mid \mathbf{t} \mid \mathbf{end} \quad \text{where } p \neq q, I \neq \emptyset, \text{ and } \forall i \in I : \text{fv}(T_i) = \emptyset$$

We require recursive global types to be *contractive*, i.e., in $\mu t . G$ we have $G \neq t'$ ($\forall t'$). We will sometimes omit **end**-typed payloads, and trailing **ends**: i.e., $p \rightarrow q : m$ stands for $p \rightarrow q : m(\mathbf{end}) . \mathbf{end}$.

We explain Definition 3.1 with an example. Consider the global type G below:

$$G = \text{alice} \rightarrow \text{bob} : \left\{ \begin{array}{l} m_1(\text{Int}) . \text{bob} \rightarrow \text{carol} : m_2(\text{Str}) . \text{carol} \rightarrow \text{alice} : m_3(\text{Bool}) . \mathbf{end} , \\ \text{stop} . \text{bob} \rightarrow \text{carol} : \text{quit} . \mathbf{end} \end{array} \right\} \quad (7)$$

The global type G involves three roles: alice , bob , and carol . It says that alice sends to bob *either* a message m_1 (carrying an **Integer**) *or* stop ; in the first case, bob sends m_2 to carol (carrying a **String**), then carol sends m_3 to alice (carrying a **Boolean**), and the session **ends**; otherwise, in the second case, bob sends quit to carol , and the session **ends**.

The *projections of G* (Definition A.1) describe the I/O actions expected from processes that play the roles in G . When G is projected onto alice , bob and carol , it yields respectively the types S_a , S_b and S_c shown in Example 2.8. Such types can be used in the typing judgement (6), in combination with the processes in Example 2.3, to yield a typing derivation similar to the following:

$$\frac{\frac{\vdots}{\vdash P_a \triangleright s[\text{alice}]:S_a} \quad \frac{\vdots}{\vdash P_b \triangleright s[\text{bob}]:S_b}}{\vdash P_a | P_b \triangleright s[\text{alice}]:S_a, s[\text{bob}]:S_b} \text{[STD-]} \quad \frac{\vdots}{\vdash P_c \triangleright s[\text{carol}]:S_c} \text{[STD-]}}{\vdash P_a | P_b | P_c \triangleright s[\text{alice}]:S_a, s[\text{bob}]:S_b, s[\text{carol}]:S_c} \text{[STD-]} \quad (8)$$

This derivation says that the parallel composition of processes $P_a | P_b | P_c$ uses the channels with role $s[\text{alice}]$, $s[\text{bob}]$, and $s[\text{carol}]$ respectively according to types S_a , S_b , and S_c . Notice that $P_a | P_b | P_c$ is typed by rule [STD-], that splits the typing context in the conclusion so that a channel with role is not used by two parallel sub-processes in its premises. This kind of context splitting implies that channels are used *linearly*, i.e., by one process a time, and cannot be shared.

In the omitted part of the derivation (8), the sub-processes P_a , P_b , and P_c are typed: this means, in particular, that they use respectively the channel with role $s[\text{alice}]$, $s[\text{bob}]$, and $s[\text{carol}]$ abiding by types S_a , S_b , and S_c . Hence, we know that the processes in Example 2.3 respect the local session types in Example 2.8, and thus, play the roles of the global type G in (7).

Inter-role dependencies Observing the processes in Example 2.3, we can notice that P_a sends `stop` to `bob`, and then terminates without receiving messages from `carol`. Indeed, this is what we expect from the type S_a in (5): P_a must input a message from `carol` if and only if it decides to send `m1` (instead of `stop`) to `bob`. However, if we observe the process P_c , we see that it *could* potentially try to send `m3` to `alice` – and this should be regarded as an error, because P_a would not be expecting such a message. Fortunately, this does *not* occur: by observing the whole ensemble $P_a | P_b | P_c$, we can see that when P_b receives `stop` from P_a , it “forwards” `quit` to `carol` (as per S_b); and when P_c receives `quit` from `bob`, it does not try to communicate again with `alice` (as per S_c). Intuitively, we can say that the process P_a “assumes” that by following the type S_a , and thus sending a certain message to the process playing role `bob`, it will also influence the behaviour of the process playing `carol`. This is a case of implicit *inter-role communication dependency*, that can be more explicitly observed in the global type G .

3.2. Subject reduction and consistency

Multiparty session typing ensures that processes enjoy properties such as protocol conformance (i.e., all sent/received messages comply with the session types, and ultimately with their originating global type G), and type safety (i.e., typed processes “never go wrong”).

As in most typing systems, type safety for MPST π -calculus processes is based on the *subject reduction* property: it guarantees that typed processes can only reduce to typed processes, and therefore, no untypable configurations can be reached. Among such untypable configurations there are, e.g., the undesired cases where a process is willing to send a message that is not supported by the intended recipient – as in the hypothetical error involving P_a and P_b discussed above. Hence, one might expect that the subject reduction theorem for MPST has a statement similar to:

$$\vdash P \triangleright \Delta \text{ and } P \rightarrow^* P' \text{ implies } \exists \Delta' \text{ such that } \Delta \rightarrow^* \Delta' \text{ and } \vdash P' \triangleright \Delta' \quad (9)$$

Then, one might also expect that the subject reduction statement in (9) can ensure type safety for our typed example in (8), by letting $P = P_a | P_b | P_c$ and $\Delta = s[\text{alice}]:S_a, s[\text{bob}]:S_b, s[\text{carol}]:S_c$. But quite surprisingly, this is *not* the case! The reason is that (9) is not quite accurate: it does *not* hold for all processes typed by any arbitrary Δ , as illustrated in Example 3.2.

Example 3.2. In “classic” MPST, an arbitrary typing context can type arguably incorrect processes. For example, we have:

$$\vdash s[p][q] \oplus_{m_\Delta} (s[r]) \mid s[q][p] \sum_{m_\Theta} (x) \triangleright s[p]:q \oplus_{m_1} (\mathbf{end}), s[q]:p \&_{m_2} (\mathbf{end}), s[r]:\mathbf{end} \quad (m_\Delta \neq m_\Theta) \quad (10)$$

Note that the typed processes in (10) are willing to interact, but are stuck (by Definition 2.4). Eq. (10) can be considered an error configuration, where a process expecting m_1 receives an unsupported message m_2 .

Another problematic example can be shown using ground types and expressions, e.g., as in [11, p. 163]:

$$\vdash s[p][q] \oplus_{m} (\text{"Hi"}) . s[p][q] \sum_{m} (y) \mid s[q][p] \sum_{m} (x) . s[q][p] \oplus_{m} (x + 42) \triangleright \frac{s[p]:q \oplus_{m} (\text{Str}) . q \&_{m} (\text{Int}),}{s[q]:p \&_{m} (\text{Int}) . p \oplus_{m} (\text{Int})} \quad (11)$$

By Definition 2.4, the process in (11) reduces to $s[p][q] \sum_{m} (y) \mid s[q][p] \oplus_{m} (\text{"Hi"} + 42)$, which is clearly untypable. This disproves the tentative subject reduction statement in (9).

To avoid the scenarios described in Example 3.2, in “classic” MPST works (e.g., by Coppo et al. [11]), the subject reduction statement is more restrictive, and reads:

$$\vdash P \triangleright \Delta \text{ with } \Delta \text{ consistent and } ; P \rightarrow^* P' \text{ implies } \exists \Delta' \text{ consistent such that } \Delta \rightarrow^* \Delta' \text{ and } \vdash P' \triangleright \Delta' \quad (12)$$

i.e., subject reduction is only proved for *consistent* (sometimes called *coherent*) typing contexts, as per Definition 3.3 below.

Definition 3.3 (Consistency). We say that Δ is consistent, written $\text{consistent}(\Delta)$, iff $\forall \Delta', s, p, q, S, T : \Delta = \Delta', s[p]:S, s[q]:T$ implies $S|q \leq T|p$.

Checking whether Definition 3.3 holds for some typing context Δ is quite cumbersome, as it requires to:

- (1) take each pair of channels with role in Δ belonging to the same session s – say, having roles p and q , and mapping respectively to session types S and T ;
- (2) further project S and T onto the “opposite” role of the pair (Definition A.4) – i.e., project S onto q and T onto p . The outcome of such projections are *partial session types* (Definition A.2), similar to *binary session types*, and only describing the interactions between p and q ;
- (3) check that such projected partial types are *dual* (Definition A.3) – modulo *partial session subtyping* \leq (Definition A.5).

These complications have a payoff: consistency ensures that all possible interactions between pairs of roles in a multiparty session are “correct,” and avoids situations like Example 3.2 above; in fact, the typing contexts in (10) and (11) are *not* consistent, because their partial type projections yield non-dual input/output message labels or payloads, violating Definition 3.3.

3.3. Limitations of consistency

Unfortunately, the consistency-based approach outlined above has a drawback: it is very restrictive, and does not hold for many “intuitively correct” processes – especially when they implement protocols with inter-role dependencies. For instance, the processes in Example 2.3 are “intuitively correct”, since they interact smoothly until they terminate, as shown in Example 2.5. However, as shown in (8), they are only typable under a typing context $\Delta = s[\text{alice}]:S_a, s[\text{bob}]:S_b, s[\text{carol}]:S_c$ that is *not* consistent, as explained in Example 3.4 below.

Example 3.4. Take S_a, S_b, S_c from Example 2.8, and $\Delta = s[\text{alice}]:S_a, s[\text{bob}]:S_b, s[\text{carol}]:S_c$. By Definition 3.3, in order to check whether Δ is consistent we need to compute the partial projections of all types (using Definition A.4), and check their duality. For example, we have the following partial projections:

$$S_a|_{\text{bob}} = \oplus\{\text{m1}(\text{Int}), \text{stop}\} \quad S_b|_{\text{alice}} = \&\{\text{m1}(\text{Int}), \text{stop}\}$$

and we also get $S_a|_{\text{bob}} = \overline{S_b|_{\text{alice}}}$ – i.e., the projections are dual (Definition A.3), because the inputs/outputs of one match the outputs/inputs of the other. This means that, by Definition 3.3, the pair $s[\text{alice}]:S_a, s[\text{bob}]:S_b$ is consistent.

However, the following projections (also required by Definition 3.3) are *not* defined:

$$\begin{aligned} S_a|_{\text{carol}} &= \text{carol}\&\text{m3}(\text{Bool})|_{\text{carol}} \sqcap \text{end}|_{\text{carol}} = \&\text{m3}(\text{Bool}) \sqcap \text{end} \text{ (undefined)} \\ S_c|_{\text{alice}} &= \text{alice}\oplus\text{m3}(\text{Bool})|_{\text{alice}} \sqcap \text{end}|_{\text{alice}} = \oplus\text{m3}(\text{Bool}) \sqcap \text{end} \text{ (undefined)} \end{aligned} \quad (13)$$

because (by Definition A.4) they use the *partial types merging operator* \sqcap in a case where it is not defined.

Therefore, the pair $s[\text{alice}]:S_a, s[\text{carol}]:S_c$ is *not* consistent; hence, by Definition 3.3, we conclude that the typing context Δ is not consistent.

In Example 3.4, the partial projection $S_a|_{\text{carol}}$ (resp. $S_c|_{\text{alice}}$) is undefined because the session type S_a (resp. S_c) starts with an interaction with **bob**; as a consequence, in order to project the type onto **carol** (resp. **alice**), it is necessary to:

- (1) “skip” the interaction with **bob**,
- (2) take all possible continuations and project each one onto **carol** (resp. **alice**), and
- (3) *merge* such projected continuations, with the operator \sqcap (Definition A.4) that tries to combine them into a single partial type.

The merge operator between partial types is quite restrictive: it does *not* allow to combine a branching (resp. selection) type with **end**, as shown in (13). This means that the interaction between **alice** (resp. **carol**) and **bob** can only have a limited influence on the interactions between **alice** (resp. **carol**) and other roles. This curtails the variety of inter-role dependencies supported by the “classic” MPST theory.

The consequence of this limitation is that the subject reduction statement in (12) is vacuously true for the typed process $P = P_a | P_b | P_c$ from Example 2.3. Therefore, we cannot use (12) to formally prove that $P_a | P_b | P_c$ reduce in a type-safe way. Further limitations arise for recursive protocols: we illustrate them in Section 3.4 below.

3.4. More “correct” but non-consistent examples

We now describe more cases of “correct” choreographies that, when projected, yield typing contexts that are *not consistent* according to Definition 3.3. As a consequence, the “classic” MPST theory does *not* provide subject reduction guarantees for processes that implement such choreographies.

Note that all the following examples will be supported by our theory, presented in Section 4: in fact, they all yield *live* typing contexts (Definition 4.1) that can type processes enjoying subject reduction (Theorem 4.9).

3.4.1. Consistency vs. history-dependent branching

In Example 3.4, consistency does *not* hold because the partial projections $S_a \upharpoonright \text{carol}$ and $S_c \upharpoonright \text{alice}$ are undefined, due to the non-mergeability of internal/external choice and **end** (as per Definition A.4). We now try to “adjust” G in (7) and make it consistent, by adding another communication between **carol** and **alice**:

$$G' = \text{alice} \rightarrow \text{bob}: \left\{ \begin{array}{l} m1(\text{Int}).\text{bob} \rightarrow \text{carol}: m2(\text{Str}).\text{carol} \rightarrow \text{alice}: m3(\text{Bool}).\text{end}, \\ \text{stop}.\text{bob} \rightarrow \text{carol}: \text{quit}.\text{carol} \rightarrow \text{alice}: m4(\text{Int}).\text{end} \end{array} \right\}$$

The projections of G' onto **alice**, **bob** and **carol** are respectively:

$$S'_a = \text{bob} \oplus \left\{ \begin{array}{l} m1(\text{Int}).\text{carol} \& m3(\text{Bool}). \\ \text{stop}.\text{carol} \& m4(\text{Int}) \end{array} \right\} \quad S'_b = \text{alice} \& \left\{ \begin{array}{l} m1(\text{Int}).\text{carol} \oplus m2(\text{Str}). \\ \text{stop}.\text{carol} \oplus \text{quit} \end{array} \right\}$$

$$S'_c = \text{bob} \& \left\{ \begin{array}{l} m2(\text{Str}).\text{alice} \oplus m3(\text{Bool}). \\ \text{quit}.\text{alice} \oplus m4(\text{Int}) \end{array} \right\}$$

Let us now examine the partial projections $S'_a \upharpoonright \text{carol}$ and $S'_c \upharpoonright \text{alice}$:

$$S'_a \upharpoonright \text{carol} = \text{carol} \& m3(\text{Bool}) \upharpoonright \text{carol} \sqcap \text{carol} \& m4(\text{Int}) \upharpoonright \text{carol} = \& m3(\text{Bool}) \sqcap \& m4(\text{Int}) \quad (\text{undefined})$$

$$S'_c \upharpoonright \text{alice} = \text{alice} \oplus m3(\text{Bool}) \upharpoonright \text{alice} \sqcap \text{alice} \oplus m4(\text{Int}) \upharpoonright \text{alice} = \oplus m3(\text{Bool}) \sqcap \oplus m4(\text{Int}) = \oplus \{m3(\text{Bool}), m4(\text{Int})\}$$

Unfortunately, partial type merging does not allow to combine different external choices, thus making $S'_a \upharpoonright \text{carol}$ undefined. Hence, external choices cannot depend on the outcome of previous interactions with other roles: i.e., **alice** cannot wait for different messages from **carol**, depending on the interaction between **alice** and **bob**. Still, different *internal* choices can be merged: e.g., **carol** can send a different message to **alice**, depending on the previous interaction between **carol** and **bob**: for this reason, $S'_c \upharpoonright \text{alice}$ is defined.

3.4.2. Consistency vs. recursion

Consider the following global type:

$$G'' = \mu \mathbf{t}.\text{alice} \rightarrow \text{bob}: \left\{ \begin{array}{l} m1(\text{Int}).\text{bob} \rightarrow \text{carol}: m2(\text{Str}).\text{carol} \rightarrow \text{alice}: m3(\text{Bool}).\mathbf{t}, \\ \text{stop}.\text{bob} \rightarrow \text{carol}: \text{quit}.\text{carol} \rightarrow \text{alice}: m3(\text{Bool}).\text{end} \end{array} \right\}$$

The global type G'' above is again similar to G in (7) – but now, if **alice** sends $m1(\text{Int})$ to **bob**, the choreography loops: in fact, **bob** then sends $m2(\text{Str})$ to **carol**, which sends $m3(\text{Bool})$ to **alice**, that continues recursively. Instead, if **alice** chooses to send **stop** to **bob**, the latter notifies **quit** to **carol**, which sends $m3(\text{Bool})$ to **alice**, and the choreography **ends**. The projections of G'' onto **alice**, **bob** and **carol** are respectively:

$$S''_a = \mu \mathbf{t}.\text{bob} \oplus \left\{ \begin{array}{l} m1(\text{Int}).\text{carol} \& m3(\text{Bool}).\mathbf{t}, \\ \text{stop}.\text{carol} \& m3(\text{Bool}).\text{end} \end{array} \right\} \quad S''_b = \mu \mathbf{t}.\text{alice} \& \left\{ \begin{array}{l} m1(\text{Int}).\text{carol} \oplus m2(\text{Str}).\mathbf{t}, \\ \text{stop}.\text{carol} \oplus \text{quit}.\text{end} \end{array} \right\}$$

$$S''_c = \mu \mathbf{t}.\text{bob} \& \left\{ \begin{array}{l} m2(\text{Str}).\text{alice} \oplus m3(\text{Bool}).\mathbf{t}, \\ \text{quit}.\text{alice} \oplus m3(\text{Bool}).\text{end} \end{array} \right\}$$

Note that in all cases, **carol** sends to **alice** the message $m3(\text{Bool})$: this circumvents the merging problems described in Example 3.4 and Section 3.4.1. However, we stumble into another roadblock:

$$S''_a \upharpoonright \text{carol} = \mu \mathbf{t}.\left((\text{carol} \& m3(\text{Bool}).\mathbf{t}) \upharpoonright \text{carol} \sqcap (\text{carol} \& m3(\text{Bool}).\text{end}) \upharpoonright \text{carol} \right) \\ = \mu \mathbf{t}.\left(\& m3(\text{Bool}).(\mathbf{t} \sqcap \text{end}) \right) \quad (\text{undefined})$$

$$S''_c \upharpoonright \text{alice} = \mu \mathbf{t}.\left((\text{alice} \oplus m3(\text{Bool}).\mathbf{t}) \upharpoonright \text{alice} \sqcap (\text{alice} \oplus m3(\text{Bool}).\text{end}) \upharpoonright \text{alice} \right) \\ = \mu \mathbf{t}.\left(\oplus m3(\text{Bool}).(\mathbf{t} \sqcap \text{end}) \right) \quad (\text{undefined})$$

This is another case of (unsupported) inter-role dependency: in S''_a (resp. S''_c), the interaction with **bob** determines whether, after receiving (resp. sending) $m3(\text{Bool})$ from **carol** (resp. to **alice**), the type will (a) loop on \mathbf{t} , and thus, keep interacting with **carol** (resp. **alice**), or (b) **end** the session. Hence, the partial projection $S''_a \upharpoonright \text{carol}$ (resp. $S''_c \upharpoonright \text{alice}$) tries to merge \mathbf{t} (from case (a)) with **end** (from case (b)) – but the operation is undefined.

4. Multiparty session typing, beyond duality and consistency

We now introduce our new multiparty session typing system. Following the discussion in Section 3, our goal is to replace consistency with a more flexible typing context property that avoids the issues in Example 3.2, but supports protocols with complex inter-role dependencies. Our key idea is to avoid the “loss of contextual information” discussed in Section 1, by introducing a new *rely/guarantee* typing judgement.

In Section 4.1, we define a behavioural *liveness* invariant used by our novel typing system, that is presented in Section 4.2; then, in Section 4.3, we formalise and discuss its type safety guarantees.

4.1. Liveness

In order to replace consistency, we need an alternative typing context property that still allows to prove subject reduction. Technically, this means that the new property must satisfy the following requirements:

- (R1) it must ensure that the typing context is “safe” (e.g., if the type of $s'[x]$ sends a message to x' , then the type of $s'[x']$ can input such a message from x), thus avoiding cases like Example 3.2;
- (R2) it must be preserved when typed processes communicate and reduce; and
- (R3) it must be preserved when typing contexts are (de)composed by parallel composition (see $[\text{STD-}]$ in (8)).

To address requirements (R1) and (R2), we introduce a coinductive notion of *typing context liveness* (Definition 4.1). It is based on the reductions of typing contexts (Definition 2.10), and ensures that each selection is matched by a compatible branching, and each branching can be triggered by a compatible selection. We will address requirement (R3) later on, with our liveness-based typing system (Section 4.2).

Definition 4.1 (*Liveness*). The predicate $\text{live}(\Delta)$ (read “ Δ is live”) is the largest predicate such that:

- $[\text{L-}\&]$ $\text{live}(\Delta, s[p]:S)$ with $S = \mathbf{q}\&_{i \in I} \mathbf{m}_i(S_i) \cdot S'_i$ implies $\exists i \in I : \exists \Delta' : \Delta, s[p]:S \xrightarrow{*} \Delta', s[p]:S'_i$;
- $[\text{L-}\oplus]$ $\text{live}(\Delta, s[p]:S)$ with $S = \mathbf{q}\oplus_{i \in I} \mathbf{m}_i(S_i) \cdot S'_i$ implies $\forall i \in I : \exists \Delta' : \Delta, s[p]:S \xrightarrow{*} \Delta', s[p]:S'_i$;
- $[\text{L-}\mu]$ $\text{live}(\Delta, s[p]:\mu\mathbf{t}.S)$ implies $\text{live}(\Delta, s[p]:S\{\mu\mathbf{t}.S/\mathbf{t}\})$;
- $[\text{L-}\rightarrow]$ $\text{live}(\Delta)$ and $\Delta \rightarrow \Delta'$ implies $\text{live}(\Delta')$.

The crucial difference between liveness and consistency (Definition 3.3) is that the former is a *behavioural* property, and does not depend on syntactic projections and duality checks. If a typing context $\Delta, s[p]:S$ is live, then:

- $[\text{L-}\&]$ says that if S is an external choice, then Δ must be able to reduce until *some* branch of S is triggered;
- $[\text{L-}\oplus]$ says that if S is an internal choice, then Δ must be able to reduce allowing to send *each* message of S ;
- $[\text{L-}\mu]$ says that if S is recursive, then its unfolding must be live, too. This means that liveness inspects all step-by-step unfoldings of S , until the first non-recursive subterm is exposed,³ and thus checked by clause $[\text{L-}\oplus]/[\text{L-}\&]$;
- $[\text{L-}\rightarrow]$ says that if the typing context reduces, the reduct must be live, too.

Example 4.2. Consider the typing context Δ_μ from Example 2.11. We have $\text{live}(\Delta_\mu)$. In fact, consider the following predicate, containing Δ_μ and its unfoldings:

$$\mathbb{P} = \left\{ \begin{array}{l} \Delta_\mu, \\ (s[p]:\mathbf{q}\oplus\mathbf{m}.\mu\mathbf{t}.\mathbf{q}\oplus\mathbf{m}.\mathbf{t}, s[q]:\mu\mathbf{t}.\mathbf{p}\&\mathbf{m}.\mathbf{t}), \\ (s[p]:\mu\mathbf{t}.\mathbf{q}\oplus\mathbf{m}.\mathbf{t}, s[q]:\mathbf{p}\&\mathbf{m}.\mu\mathbf{t}.\mathbf{p}\&\mathbf{m}.\mathbf{t}), \\ (s[p]:\mathbf{q}\oplus\mu\mathbf{t}.\mathbf{q}\oplus\mathbf{m}.\mathbf{t}, s[q]:\mathbf{p}\&\mathbf{m}.\mu\mathbf{t}.\mathbf{p}\&\mathbf{m}.\mathbf{t}) \end{array} \right\}$$

By inspecting each $\Delta \in \mathbb{P}$, we can verify that it satisfies the coinductive clauses of Definition 4.1:

- each $\Delta \in \mathbb{P}$ reduces to Δ_μ , that belongs to \mathbb{P} – thus satisfying clause $[\text{L-}\rightarrow]$;
- if $\Delta = \Delta', s[x]:\mu\mathbf{t}.S$ (for some x), then $\Delta', s[x]:S\{\mu\mathbf{t}.S/\mathbf{t}\}$ belongs to \mathbb{P} – thus satisfying clause $[\text{L-}\mu]$;
- if some type in Δ exposes an internal/external choice, then it can be eventually triggered by the rest of the context (see Example 2.11) – thus satisfying clause $[\text{L-}\oplus]/[\text{L-}\&]$.

Now, by Definition 4.1, live is the largest predicate satisfying such clauses, and therefore we have that $\Delta \in \mathbb{P}$ implies $\text{live}(\Delta)$. Then, since $\Delta_\mu \in \mathbb{P}$, we conclude $\text{live}(\Delta_\mu)$.

³ Recall that session types are contractive, by Definition 2.6: hence, the unfolding of a sequence of nested recursions $\mu\mathbf{t}_1.\mu\mathbf{t}_2\dots$ must eventually yield a type with a top-level internal/external choice, or **end**.

$$\begin{array}{c}
\frac{\forall c \in \text{dom}(\Delta) : \Delta(c) \equiv \mathbf{end}}{\mathbf{end}(\Delta)} \quad [\text{T-END}] \quad \frac{\Delta \equiv c:S}{\Delta \vdash c:S} \quad [\text{T-CHAN}] \quad \frac{\Theta(X) \equiv S_1, \dots, S_n}{\Theta \vdash X : S_1, \dots, S_n} \quad [\text{T-PROC}] \quad \frac{\mathbf{end}(\Delta)}{\Theta \vdash \mathbf{0} \triangleright \Delta \triangleleft \Gamma} \quad [\text{T-0}] \\
\frac{\Theta, X:S_1, \dots, S_n \vdash P \triangleright x_1:S_1, \dots, x_n:S_n \triangleleft \emptyset \quad \Theta, X:S_1, \dots, S_n \vdash Q \triangleright \Delta \triangleleft \Gamma}{\Theta \vdash \mathbf{def} X(x_1, \dots, x_n) = P \mathbf{in} Q \triangleright \Delta \triangleleft \Gamma} \quad [\text{T-def}] \quad \frac{\Theta \vdash X : S_1, \dots, S_n \quad \mathbf{end}(\Delta_0) \quad \forall i \in 1..n \quad \Delta_i \vdash c_i:S_i}{\Theta \vdash X(c_1, \dots, c_n) \triangleright \Delta_0, \Delta_1, \dots, \Delta_n \triangleleft \Gamma} \quad [\text{T-X}] \\
\frac{\Delta_s = \{s[p]:S_p\}_{p \in I} \quad \Theta \vdash P \triangleright \Delta, \Delta_s \triangleleft \Gamma}{\Theta \vdash (\nu S:\Delta_s) P \triangleright \Delta \triangleleft \Gamma} \quad [\text{T-}\nu] \quad \frac{\Theta \vdash P_1 \triangleright \Delta_1 \triangleleft \Gamma, \Delta_2 \quad \Theta \vdash P_2 \triangleright \Delta_2 \triangleleft \Gamma, \Delta_1}{\Theta \vdash P_1 | P_2 \triangleright \Delta_1, \Delta_2 \triangleleft \Gamma} \quad [\text{T-}|] \\
\frac{S \equiv \mathbf{q}\&_{i \in I} m_i(S_i).S'_i \quad \forall i \in I \quad \forall \Gamma', \Gamma'' : \Gamma \xrightarrow{*} \Gamma' \text{ and } \Gamma', c:S \rightarrow \Gamma'', c:S'_i \quad \Theta \vdash P_i \triangleright \Delta, y_i:S_i, c:S'_i \triangleleft \Gamma''}{\Theta \vdash c[\mathbf{q}] \sum_{i \in I \cup J} \{m_i(y_i).P_i\} \triangleright \Delta, c:S \triangleleft \Gamma} \quad [\text{T-}\&] \\
\frac{S \equiv \mathbf{q}\oplus_{i \in I} m_i(S_i).S'_i \quad \exists k \in I \quad \Delta_d \vdash d:S_k \quad \forall \Gamma', \Gamma'' : \Gamma \xrightarrow{*} \Gamma' \text{ and } \Gamma', c:S \rightarrow \Gamma'', c:S'_k \quad \Theta \vdash P \triangleright \Delta, c:S'_k \triangleleft \Gamma', \Delta_d}{\Theta \vdash c[\mathbf{q}]\oplus_{m_k}(d).P \triangleright \Delta_d, \Delta, c:S \triangleleft \Gamma} \quad [\text{T-}\oplus]
\end{array}$$

Fig. 3. Rely/guarantee multiparty session typing system.

Example 4.3. Consider Δ_{abc} from Example 2.12, and the following predicate, containing Δ_{abc} and all its reducts:

$$\mathbb{P} = \{ \Delta' \mid \Delta_{abc} \xrightarrow{*} \Delta' \}$$

We can verify that each element of \mathbb{P} satisfies clauses [L- \rightarrow], [L- $\&$], [L- \oplus] and [L- μ] of Definition 4.1 (possibly vacuously). Hence, similarly to Example 4.2, we conclude $\text{live}(\Delta_{abc})$.

Notably, a live typing context never deadlocks: i.e., if it cannot further reduce, then all its entries must be **end**-typed. The *vice versa* is not true: a deadlock-free typing context might not be live. Moreover, liveness and consistency are incomparable, i.e., neither implies the other. This is shown in Example 4.4 below.

Example 4.4. The typing context $s[p]:S_a, s[q]:S_b, s[r]:S_c$ in (8) is live, but not consistent.

The context $s[p]:\mathbf{q}\&_{m_1}.r\oplus_{m_3}, s[q]:r\&_{m_2}.p\oplus_{m_1}, s[r]:p\&_{m_3}.q\oplus_{m_2}$ is consistent but not live (nor deadlock-free): its inputs/outputs occur in the wrong order, thus preventing reductions.

Finally, the typing context $s[p]:\mu\mathbf{t}.q\oplus_{m_1}.\mathbf{t}, s[q]:\mu\mathbf{t}.p\&_{m_1}.\mathbf{t}, s[r]:p\&_{m_2}$ is deadlock-free and consistent, but not live: $s[p]$ and $s[q]$ generate infinite reductions, but $s[r]$ cannot possibly perform its input (thus violating clause [L- $\&$] of Definition 4.1).

4.2. A rely/guarantee typing system for multiparty sessions

We can now introduce our new typing system.

Definition 4.5 (Rely/guarantee typing system). Let Θ be a partial mapping defined as:

$$\Theta ::= \emptyset \mid \Theta, X:\tilde{S}$$

A rely/guarantee multiparty session typing judgement has the form:

$$\Theta \vdash P \triangleright \Delta \triangleleft \Gamma \quad \text{with } \text{live}(\Delta, \Gamma)$$

and is inductively defined by the rules in Fig. 3. Above, Δ and Γ are called respectively *guarantee* and *rely* contexts.

In Definition 4.5, similarly to most MPST works, Θ assigns an n -uple of types to each process variable X (one type per argument). Intuitively, the judgement $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$ says: given the process definitions typed in Θ , P “guarantees” to use its channels *linearly* according to Δ , by “relying” on the assumption that other processes use their channels according to Γ . The additional Γ is the most visible departure from standard MPST typing rules (although we subsume them under some conditions, as we will show in Section 5.3). Note that the requirement $\text{live}(\Delta, \Gamma)$ implies that Δ, Γ must be defined: therefore, by Definition 2.9, we must have $\text{dom}(\Delta) \cap \text{dom}(\Gamma) = \emptyset$.

In Fig. 3, the first three rules define auxiliary judgements:

- [T-END] says that $\mathbf{end}(\Delta)$ holds if all entries of Δ are **end**-typed;
- [T-CHAN] says that the judgement $\Delta \vdash c:S$ (read “ c has type S in Δ ”) holds if Δ contains exactly one entry $c:S'$, with $S' \equiv S$. Note that this judgement allows to consider S' up-to unfolding: this makes the typing system *equi-recursive* [31, Chapter 21];
- [T-PROC] says that the judgement $\Theta \vdash X : S_1, \dots, S_n$ (read “process X has parameter types S_1, \dots, S_n in Θ ”) holds if Θ maps X to an n -uple of types equivalent to S_1, \dots, S_n (again, up-to unfolding).

Using Lemma 4.8, we prove the *Subject Reduction Theorem 4.9* below. Note that its statement is similar to the (tentative) statement (9) in Section 3, but without the counterexamples shown in Example 3.2: this is because in Definition 4.5, we “embed” the liveness requirement in the typing judgement, so that Theorem 4.9 holds for *all* typed processes (as one might usually expect).

Theorem 4.9 (*Subject reduction*). *If $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$ and $P \rightarrow^* P'$, then $\exists \Delta'$ such that $\Delta \rightarrow^* \Delta'$ and $\Theta \vdash P' \triangleright \Delta' \triangleleft \Gamma$.*

Proof. See Appendix F. The proof crucially exploits the fact that, in our rely/guarantee typing system, liveness satisfies the requirements (R1)–(R3) outlined in Section 4.1. \square

5. Comparison: rely/guarantee vs. “classic” multiparty session typing

In this section, we compare our new rely/guarantee multiparty session typing system with the “classic” MPST theory. In Example 4.4, we have shown that liveness and consistency are incomparable; therefore:

1. on the one hand, using our theory we can prove type safety in cases that are not covered in “classic” MPST, as shown in Example 4.7;
2. but on the other hand, the “classic” MPST subject reduction statement in (12) allows to prove type safety for processes that would be rejected by our theory, because their typing context is consistent but not live (see Example 4.4).

Regarding item 2, it is important to notice that if a consistent typing context Δ is not live, then (by Definition 4.1) it contains some internal/external choice that can be reached but never fired; therefore, Δ types processes that might wait forever on some output/input that cannot be possibly triggered by another process (see Example 4.4). For this reason, we argue that non-live typing contexts should be considered ill-formed, and can be reasonably rejected. Indeed, in “classic” MPST, well-formed global types and projections are used to rule out this kind of ill-formed protocols.

We now further investigate the differences between our new theory and the “classic” one. In Section 5.1, we show that our typing system supports all multiparty protocols projected from some global type, *plus* other “live” protocols that cannot be projected from any global type; in Section 5.2, we show that our typing system tolerates some (safe) mismatches between processes and types; finally, in Section 5.3, we show how our typing rules can be simplified (in some cases) to match the “classic” ones.

5.1. Global types, and beyond

Our theory covers all typing contexts obtained by projecting some global type G . This is implied by Lemma 5.1 below – where $\text{roles}(G)$ is the set of roles occurring in G .

Lemma 5.1. *Take any global type G , and let $\Delta_G = \{s[x]:S_x\}_{x \in \text{roles}(G)}$ such that $\forall x \in \text{roles}(G) : S_x = G \upharpoonright x$. Then, we have $\text{live}(\Delta_G)$.*

Proof. (*Sketch*) This is a consequence of the correspondence between multiparty session types and *Communicating Finite-State Machines (CFSMs)* [8], originally established by Deniérou and Yoshida [14], and later reprised by Bocchi et al. [6].

Deniérou and Yoshida [14] define the projection of a global type G onto a *system of CFSMs* (one per role of G), that we denote here as \mathcal{S}_G . The projection corresponds to the usual global-to-local type projection (Definition A.1), and generates restricted CFSMs whose states are *non-mixed* and *directed* – i.e., each state q is either an input state (with receive transitions, only) or an output state (with send transitions, only), and each transition from q is an input/output from/to the same participant. With these restrictions, each CFSM in \mathcal{S}_G corresponds to a local multiparty session type (Definition 2.6), and the system \mathcal{S}_G is homomorphic to the projected typing context Δ_G (since it contains the same projected session types).

Then, by Theorem 4.3 in Deniérou and Yoshida [14], we know that \mathcal{S}_G is *Multiparty Compatible (MC)* ([14, Def. 4.2]; [6, Def. 4]). In our setting, MC (with the formulation of Bocchi et al. [6]) can be translated as follows:

(MC1) $\Delta_G \rightarrow^* \Delta', s[p]:S$ with $S \equiv \mathfrak{q} \&_{i \in I} m_i(S_i).S'_i$ implies $\exists k \in I : \exists \Delta'', \Delta''': \Delta' \rightarrow^* \Delta''$ and $\Delta'', s[p]:S \rightarrow \Delta''', s[p]:S'_k$;
(MC2) $\Delta_G \rightarrow^* \Delta', s[p]:S$ with $S \equiv \mathfrak{q} \oplus_{i \in I} m_i(S_i).S'_i$ implies $\forall k \in I : \exists \Delta'', \Delta''': \Delta' \rightarrow^* \Delta''$ and $\Delta'', s[p]:S \rightarrow \Delta''', s[p]:S'_k$.

Clause (MC1) says that whenever Δ_G reduces to a configuration with a type/CFSM for p that inputs from \mathfrak{q} , then it can eventually receive a corresponding message from \mathfrak{q} . Clause (MC2) says that whenever Δ_G reduces to a configuration with a type/CFSM for p that outputs some message towards \mathfrak{q} , then the message can be eventually received by \mathfrak{q} . Besides the use of session types instead of non-mixed and directed CFSMs, the above translation of MC has two differences w.r.t. Deniérou and Yoshida [14] and Bocchi et al. [6]:

1. our typing context reductions (Definition 2.10) are synchronous. Instead, Deniérou and Yoshida [14] and Bocchi et al. [6] have *asynchronous* communication via message queues. Still, MC is based on a notion of synchronous reduction: it only considers restricted pairs of “alternating” transitions where each message is first queued, and then received, without other transitions in between;

2. our typing context reductions (Definition 2.10) compare *both* message labels *and* message payloads – and the latter are compared with the up-to unfolding equivalence \equiv (Definition 2.6). Instead, Deniérou and Yoshida [14] and Bocchi et al. [6] do not have payloads, and use CFSM transitions that only compare message labels. Still, extending CFSMs to support and compare message payloads with the same treatment of Definition 2.10 is not difficult, and has no consequence on MC and its properties.

Now, summing up: if we have Δ_G defined as in the hypotheses, by Deniérou and Yoshida [14] and Bocchi et al. [6] we know that Δ_G satisfies clauses (MC1) and (MC2) above. Then, we can show that Δ_G is live: it suffices to define a predicate \mathbb{P} containing Δ_G and its reducts and unfoldings (similarly to Examples 4.2 and 4.3), and show that each element of \mathbb{P} satisfies the clauses of Definition 4.1; in particular, clause $[L-\&]$ holds by (MC1), and $[L-\oplus]$ holds by (MC2). Hence the result. \square

Lemma 5.1 means that our typing judgement (Definition 4.5) supports all multiparty protocols typically considered in “classic” MPST theory. Moreover, for all typed processes, it provides a subject reduction result that is not restricted by consistency (Theorem 4.9): this is *not* the case in “classic” MPST, as discussed in Section 3.3.

Additionally, our rely/guarantee typing system goes beyond global types: in Examples 5.2 and 5.3 below, we show two typing contexts that are live (and thus supported by our theory), but cannot be obtained by projecting any global type.

Example 5.2. Consider the following typing context (that is not consistent):

$$\Delta = s[p]:\mu t.q\oplus\{m_1.t, m_2\}, s[q]:\mu t.p\&\{m_1.t, m_2.r\oplus m_3\}, s[r]:q\&m_3$$

It reduces as follows:

$$\begin{array}{c} \curvearrowright \\ \Delta \end{array} \longrightarrow \begin{array}{l} s[p]:\text{end} \\ s[q]:r\oplus m_3 \\ s[r]:q\&m_3 \end{array} \longrightarrow \begin{array}{l} s[p]:\text{end} \\ s[q]:\text{end} \\ s[r]:\text{end} \end{array}$$

In this diagram, Δ loops when $s[p]$ sends m_1 to $s[q]$; otherwise, if $s[p]$ sends m_2 to $s[q]$, then $s[q]$ sends m_3 to $s[r]$, and the final configuration is reached. By inspecting these reductions, we can verify that Δ is live.

However, Δ is not projectable from any global type. In fact, observe that the type of $s[q]$ recursively waits to receive either m_1 or m_2 from p , and only in the second case, $s[q]$ sends m_3 to r . This means that an hypothetical G projecting into Δ should look like $G = \mu t.p \rightarrow q:\{m_1.t, m_2.q \rightarrow r:m_3\}$, i.e., the global type syntactically involves role r in one branch, but *not* in the other – and this makes the projection $G \upharpoonright r$ undefined (see Definition A.1).

Example 5.3. Consider the following typing context:

$$\Delta'_\mu = s[p]:S, s[q]:\mu t'.p\&m(S).t' \quad \text{where } S = \mu t.q\oplus m(t).t$$

Here, S is the type of a channel that must be used to recursively send a message m to a process playing role q ; moreover, m carries an S -typed payload channel that the recipient must use accordingly.

Note that Δ'_μ is not projectable from any global type, because projection (Definition A.1) cannot yield session types with recursion variables occurring as payloads. Moreover, Δ'_μ is *not* consistent: the projection $S \upharpoonright q = \mu t'.\oplus m(t).t'$ yields an invalid partial type, with an open session type as message payload (Definition A.2).

By Definition 2.10, Δ'_μ reduces as follows:

$$\frac{\frac{S \equiv S}{s[p]:q\oplus m(S).S, s[q]:p\&m(S).\mu t'.p\&m(S).t' \rightarrow \Delta'_\mu} [T-\text{COMM}]}{s[p]:q\oplus m(S).S, s[q]:\mu t'.p\&m(S).t' \rightarrow \Delta'_\mu} [T-\mu]}{\Delta'_\mu \rightarrow \Delta'_\mu} [T-\mu]$$

and we can prove $\text{live}(\Delta'_\mu)$ similarly to Example 4.2, by (1) defining a predicate \mathbb{P} containing Δ'_μ and its unfoldings, and (2) showing that the elements of \mathbb{P} satisfy the clauses of Definition 4.1.

Note that in our rely/guarantee session typing theory, the following judgement holds:

$$\begin{array}{l} \emptyset \vdash Q_p \mid Q_q \triangleright \Delta'_\mu \triangleleft \emptyset \quad \text{where :} \\ Q_p = \mathbf{def} X(x) = \left(\mathbf{def} Z(z) = z[p] \sum \{m(x').(Z(z) \mid X(x'))\} \right. \\ \quad \left. \mathbf{in} ((\nu s') (x[q] \oplus m(s'[p])).X(x) \mid Z(s'[q]))) \right) \mathbf{in} X(s[p]) \\ Q_q = \mathbf{def} Z(z) = \left(\mathbf{def} X(x) = ((\nu s') (x[q] \oplus m(s'[p])).X(x) \mid Z(s'[q])) \right) \mathbf{in} z[p] \sum \{m(x).(Z(z) \mid X(x))\} \end{array}$$

and therefore, by Theorem 4.9, we know that $Q_p \mid Q_q$ interact type-safely.

Remark 5.4. Example 5.3 shows that by using liveness, our rely/guarantee typing system does not incur in the limitations of duality applied to recursive types, that have been variously addressed in several works. We will reprise this topic in Section 6.

5.2. Leveraging typing context reductions

Unlike “classic” MPST typing systems, our rules $[T-\&]$ and $[T-\oplus]$ (Fig. 3) inspect typing context reductions. This adds a certain flexibility: the contextual information can be exploited to type processes in ways that would be forbidden in “classic” MPST, as shown in Example 5.5 below.

Example 5.5. Consider the following variation of the session types in Example 2.8:

$$S'_a = \text{bob} \oplus \text{stop} . \text{end} \quad S'_b = \text{alice} \& \left\{ \begin{array}{l} m_1(\text{Int}) . \text{alice} \oplus m_2(\text{Str}) . \text{end} , \\ \text{stop} . \text{carol} \oplus \text{quit} . \text{end} \end{array} \right\} \quad S'_c = \text{bob} \& \left\{ \begin{array}{l} m_2(\text{Str}) . \text{end} , \\ \text{quit} . \text{end} \end{array} \right\}$$

The differences with Example 2.8 are:

1. S'_a must send `stop` to `bob`, and has no other choice;
2. if S'_b receives m_1 from `alice`, then it sends m_2 to `alice` (instead of `carol`);
3. if S'_c receives m_2 from `bob`, then it terminates without sending m_3 to `alice`.

Note that the typing context $s[\text{alice}]:S'_a, s[\text{bob}]:S'_b, s[\text{carol}]:S'_c$ is live, because the mismatching outputs in the altered branches of S'_b and S'_c are never reached, and therefore, do not violate clauses $[L-\&]$ and $[L-\oplus]$ of Definition 4.1.

Now, consider the processes in Example 2.3: with the variations 1–3 above, the process P_a still matches S'_a , but P_b and P_c do *not* match S'_b and S'_c . Nonetheless, the following typing derivation holds (here, we abbreviate `alice`, `bob`, `carol` resp. as a , b , c):

$$\frac{\frac{\frac{\vdots}{\emptyset \vdash P_a \triangleright s[a]:S'_a \triangleleft s[b]:S'_b, s[c]:S'_c} \quad \frac{\vdots}{\emptyset \vdash P_b \triangleright s[b]:S'_b \triangleleft s[a]:S'_a, s[c]:S'_c}}{\emptyset \vdash P_a \mid P_b \triangleright s[a]:S'_a, s[b]:S'_b \triangleleft s[c]:S'_c} \quad \frac{\vdots}{\emptyset \vdash P_c \triangleright s[c]:S'_c \triangleleft s[a]:S'_a, s[b]:S'_b}}{\emptyset \vdash P_a \mid P_b \mid P_c \triangleright s[a]:S'_a, s[b]:S'_b, s[c]:S'_c \triangleleft \emptyset} \quad [T-\mid]}{\Delta' = s[a]:S'_a, s[b]:S'_b, s[c]:S'_c \quad \emptyset \vdash (v_S:\Delta')(P_a \mid P_b \mid P_c) \triangleright \emptyset \triangleleft \emptyset} \quad [T-v]$$

The reason is that, when typing P_b and P_c , the rely context does *not* trigger the altered branches of S'_b and S'_c , and correspondingly, the typing rule $[T-\&]$ in Fig. 3 does *not* check its inductive clauses for the corresponding process continuations. As a consequence, the typing system (safely) accepts this process/type discrepancy.

5.3. From “new” to simpler “old” rules

The typing rules $[T-\&]/[T-\oplus]$ (Fig. 3) can be demanding: the quantification “ $\forall \Gamma' : \Gamma \rightarrow^* \Gamma' \dots$ ” can yield many premises (but their number is always finite, cf. Remark 4.6). Fortunately, Lemma 5.6 below allows to reduce the number of premises that must be checked explicitly.

Lemma 5.6 (Rely context strengthening). *If $\emptyset \vdash P \triangleright \Delta \triangleleft \Gamma$ and $\Gamma \rightarrow^* \Gamma'$, then $\emptyset \vdash P \triangleright \Delta \triangleleft \Gamma'$.*

Intuitively, Lemma 5.6 holds because if a typing context reduces, it becomes “more deterministic”; hence, if P is typed relying on Γ , it can also rely on its reductions.

Additionally, type checks can be simplified using Proposition 5.7 below (where x is a variable, and *not* a channel with role, by Definition 2.1).

Proposition 5.7. *Assume $x \notin \text{dom}(\Delta) \cup \text{dom}(\Gamma)$. Then, $\emptyset \vdash P \triangleright \Delta \triangleleft \Gamma, x:S$ iff $\emptyset \vdash P \triangleright \Delta \triangleleft \Gamma$.*

By Proposition 5.7, variables in the rely context Γ are immaterial – intuitively, because they do *not* influence liveness, nor the possible reductions of Δ and Γ in the statement.

Lemma 5.6 and Proposition 5.7 become particularly useful when typing a process that branches/selects from/to a variable – and (in the latter case) sends a variable. In this case, rules $[T-\&]$ and $[T-\oplus]$ become the standard MPST typing rules for branch/input and select/output (see [11]), plus an *unused* rely context:

$$\frac{S \equiv \mathfrak{q}\&_{i \in I} m_i(S_i).S'_i \quad \forall i \in I \quad \Theta \vdash P_i \triangleright \Delta, y_i : S_i, x : S'_i \triangleleft \Gamma}{\Theta \vdash x[\mathfrak{q}] \sum_{i \in I \cup J} \{m_i(y_i).P_i\} \triangleright \Delta, x : S \triangleleft \Gamma} \text{[T\&x]}$$

$$\frac{S \equiv \mathfrak{q}\oplus_{i \in I} m_i(S_i).S'_i \quad k \in I \quad \Delta_y \vdash y : S_k \quad \Theta \vdash P \triangleright \Delta, x : S'_k \triangleleft \Gamma}{\Theta \vdash x[\mathfrak{q}] \oplus_{m_k}(y).P \triangleright \Delta_y, \Delta, x : S \triangleleft \Gamma} \text{[T\oplusx]}$$

This is because in $[\text{T}\oplus]$, by Proposition 5.7 we can omit Δ_y from the rely context of the rightmost premise. Further, since $x:S$ reduces autonomously (cf. Definition 2.10), in both $[\text{T}\&]$ and $[\text{T}\oplus]$ the quantification “ $\forall \Gamma', \Gamma'' : \Gamma \rightarrow^* \Gamma'$ and $\Gamma', x:S \rightarrow \Gamma'', x:S'_i$ ” amounts to “ $\forall \Gamma' : \Gamma \rightarrow^* \Gamma'$ and $\Gamma', x:S \rightarrow \Gamma', x:S'_i$ ”; hence, by Lemma 5.6, checking the rightmost premise *once* with Γ is sufficient to know that it holds $\forall \Gamma'$ such that $\Gamma \rightarrow^* \Gamma'$.

This shows that in most cases, our rely/guarantee type checking can be simplified, becoming similar to the “classic” MPST rules.

6. Conclusions and related work

We presented a new multiparty session typing system (Definition 4.5), based on a *rely/guarantee* typing strategy, and a *liveness* invariant on typing contexts (Definition 4.1). Our typing system provides a subject reduction result (Theorem 4.9) that overcomes the consistency restrictions of “classic” MPST works. By directly exploiting the typing context semantics, our typing rules provide flexible type checks supporting all protocols projectable from global types, *plus* various kinds of live protocols that cannot be projected from any global type (Section 5). As a bonus, the independence from global types means that our theory is simpler than the “classic” one, in the sense that it does *not* depend on its syntactic type manipulations (i.e., the type projection/merging/duality machinery of “classic” MPST, summarised in Appendix A).

Inter-role dependencies Our rely/guarantee typing system allows to prove type safety of processes implementing protocols with complex inter-role dependencies: this is generally not supported by “classic” MPST, as discussed in Sections 1, 3.4 and 5, and Examples 3.4 and 4.7. In this respect, to the best of our knowledge, the only work with comparable features is the one by Dezani-Ciancaglini et al. [15]: it provides a “non-classic” MPST theory where processes can only interact on *one* session – and this limitation is crucially exploited to type parallel compositions without splitting their typing context (cf. Table 8, rule $[\text{T}\text{-SESS}]$). Their approach directly tracks a global type along process reductions, without introducing consistency restrictions. However, unlike the present work, Dezani-Ciancaglini et al. [15] do not support multiple sessions and delegation – and extending their work in that direction appears challenging. Moreover, our approach does not depend on the existence of global types: therefore, it is arguably simpler (since it only uses local types), and supports protocols for which no corresponding global type exists, as shown in Section 5.1.

Global types and liveness Our type system does *not* assume that a session is typed by projecting some global type: it only (implicitly) requires that a session is typed by a live composition of *local* session types, by rule $[\text{T}\text{-v}]$ (Fig. 3). If needed, a set of local types can be related to a choreography either via “top-down” projection (Section 3.1), or via “bottom-up” synthesis (studied, e.g., by Lange and Tuosto [26] and Lange et al. [27]): these concerns are orthogonal to our typing system. Our Definition 4.1 (liveness) is inspired by (1) the notion of *safety* for Communicating Finite-State Machines (CFSMs) Brand and Zafropulo [8] (no deadlocks, orphan messages, unspecified receptions) and (2) the idea of “multiparty session types as CFSMs” outlined by Deniérou and Yoshida [14]. In particular, liveness (Definition 4.1) is closely related to the notion of *Multiparty Compatibility* for CFSMs, as illustrated in Section 5.1 (proof of Lemma 5.1). Notably, Deniérou and Yoshida [14] focus on projecting CFSMs and studying their interactions, and do *not* technically develop the intuition of “CFSMs as types” into an actual typing system. For this reason, they do not realise that the “classic” MPST theory cannot fully support their intuition, due to the consistency requirement and the resulting restrictions. Our work provides an actual multiparty session typing system that sidesteps these limitations, and comes closer to using “CFSMs as types.”

Deadlock freedom Similarly to most previous MPST works, our approach ensures that a typed ensemble of processes interacting on a single session (i.e., a typed $(\nu s)(\prod_{p \in I} P_p)$, with each P_p only interacting on $s[p]$) is deadlock-free; however, it does not guarantee deadlock freedom in presence multiple interleaved sessions. This topic is studied by Bettini et al. [5] and Coppo et al. [12]: as future work, we plan to investigate how to reuse their results in our theory.

Recursion and duality Various works on binary session types have investigated the subtle interplay between recursion and duality. Bono and Padovani [7] and Bernardi and Hennessy [3] independently noticed that the “standard” duality proposed by Honda et al. [20] (corresponding to Definition A.3) does *not* commute with the unfolding of recursion for non-tail-recursive types (i.e., with type variables as payloads). An example of such types is $\mu t. \oplus_{m(t)}. t$, similar to our Example 5.3. This implies that some “safe” process compositions cannot be typed in binary session theories – and this limitation can also be found in “classic” MPST works, due to their consistency requirement (that is rooted in binary duality, cf. Definition 3.3). To solve this issue, Bono and Padovani [7] and Bernardi and Hennessy [3] defined a new notion of duality, called *complement* by the latter, and further studied by Bernardi et al. [2]. However, [4, Example 5.21] later noticed that even complement does *not* commute, e.g., when unfolding $\mu t. \mu t'. \&_{m(t')}. t$. This issue was later solved by Lindley and Morris [28, Section 5.2]:

they add *dualised type variables* to session types, with specialised duality and unfolding definitions that commute as desired. Our approach sidesteps these issues by *not* requiring syntactic duality, but using a liveness property that unfolds types as necessary, and (coinductively) compares the trees of payload types (see Example 5.3 and its discussion).

Encodability and implementations Scalas et al. [33] have proven that the consistency requirement of “classic” MPST creates a strong link between MPST and linear π -calculus with *binary* channels, allowing to encode the former into the latter. They also exploit this result to provide an implementation of “classic” MPST in Scala [34]. Our new rely/guarantee typing system does not depend on consistency: therefore, the encoding strategy of Scalas et al. [33] can only be applied to the consistent fragment of our theory, and implementing the rest is now an open problem.

Future work We plan to investigate typed behavioural congruences and bisimulations for our new theory, as done by Kouzapas and Yoshida [24,25] for “classic” MPST.

We also plan to extend our approach to *asynchronous* multiparty session types [21,22] where process interaction is mediated by message queues (similarly to the TCP/IP protocol). The main differences and challenges we foresee stem from the fact that, in the present work, the *synchronous* typing context reduction semantics (Definition 2.10) induce *finite* state transition systems – whereas the addition of types for message queues can yield infinite-state systems, impacting decidability of liveness and type checking.

Acknowledgements

This work was partially supported by EPSRC grants EP/K011715/1, EP/K034413/1, EP/L00058X/1, EP/N028201/1, and EP/N027833/1. We thank the anonymous reviewers for their valuable comments and suggestions, and Raymond Hu, Julien Lange, and Bernardo Toninho for the fruitful discussion.

Appendix A. “Classic” multiparty session types: definitions

This appendix contains standard definitions from MPST literature. They are mainly referenced from Section 3, but are *not* necessary for our rely/guarantee typing system (Section 4): it uses a *behavioural* liveness property (Definition 4.1), and does not depend on the existence of global types and their projections.

A.1. Partial session types and projections

The *global and local type projections* in Definition A.1 and Definition A.4 below, with their respective *merging operators*, are based on [38] and [13]. They are more flexible than the simpler projection/merging used, e.g., by Coppo et al. [11]: hence, they allow for more global types to be projectable.

Intuitively, the projection and merging in Definition A.1 below allow to reconstruct the expected behaviour of a role in a global type, by collecting and combining its internal/external choices along different execution paths.

Definition A.1 (*Projection of a global type*). The *projection of G onto role p* , written $G \upharpoonright p$, is:

$$(\mathfrak{q} \rightarrow \mathfrak{r} : \{m_i(T_i) \cdot G_i\}_{i \in I}) \upharpoonright p = \begin{cases} \mathfrak{r} \oplus_{i \in I} m_i(T_i) \cdot (G_i \upharpoonright p) & \text{if } p = \mathfrak{q} \\ \mathfrak{q} \&_{i \in I} m_i(T_i) \cdot (G_i \upharpoonright p) & \text{if } p = \mathfrak{r} \\ \sqcap_{i \in I} G_i \upharpoonright p & \text{if } \mathfrak{q} \neq p \neq \mathfrak{r} \end{cases} \quad (\mu \mathfrak{t}.G) \upharpoonright p = \begin{cases} \mu \mathfrak{t} \cdot (G \upharpoonright p) & \text{if } G \upharpoonright p \neq \mathfrak{t}' \ (\forall \mathfrak{t}') \\ \mathbf{end} & \text{otherwise} \end{cases}$$

$$\mathfrak{t} \upharpoonright p = \mathfrak{t} \quad \mathbf{end} \upharpoonright p = \mathbf{end}$$

where \sqcap is the *merge operator for (local) session types*, defined as:

$$p \&_{i \in I} m_i(S_i) \cdot S'_i \sqcap p \&_{j \in J} m_j(S_j) \cdot T'_j = p \&_{k \in I \cap J} m_k(S_k) \cdot (S'_k \sqcap T'_k) \ \& \ p \&_{i \in I \setminus J} m_i(S_i) \cdot S'_i \ \& \ p \&_{j \in J \setminus I} m_j(S_j) \cdot T'_j$$

$$p \oplus_{i \in I} m_i(S_i) \cdot S'_i \sqcap p \oplus_{i \in I} m_i(S_i) \cdot S'_i = p \oplus_{i \in I} m_i(S_i) \cdot S'_i$$

$$\mu \mathfrak{t}.S \sqcap \mu \mathfrak{t}.T = \mu \mathfrak{t} \cdot (S \sqcap T) \quad \mathfrak{t} \sqcap \mathfrak{t} = \mathfrak{t} \quad \mathbf{end} \sqcap \mathbf{end} = \mathbf{end}$$

For example, if we have a global type $G = p \rightarrow \mathfrak{q} : \{m_i(T_i) \cdot \mathbf{end}\}_{i \in I}$, then the projection $G \upharpoonright p$ describes p 's behaviour as an internal choice towards \mathfrak{q} , with labels m_i and payload T_i ($i \in I$); dually, the projection $G \upharpoonright \mathfrak{q}$ is an external choice from p , with labels m_i and payloads T_i . Instead, the projection of G onto a role \mathfrak{r} that is neither p nor \mathfrak{q} requires to skip the first interaction (since it does not involve \mathfrak{r}), compute the projections onto \mathfrak{r} along each branch m_i (for each $i \in I$), and then combine the resulting projections using the *merge operator* \sqcap : in this case, $G \upharpoonright \mathfrak{r} = \sqcap_{i \in I} (\mathbf{end} \upharpoonright \mathfrak{r}) = \sqcap_{i \in I} \mathbf{end} = \mathbf{end}$, i.e., \mathfrak{r} does not perform any input/output in G .

Note that projection and merging are sometimes undefined: for example, it is possible to merge different external choices (by merging the continuation types that follow common message label), but *not* different internal choices. This rules out meaningless global types like $G' = p \rightarrow \mathfrak{q} : \{m_1 \cdot \mathfrak{r} \rightarrow p : m_1, m_2 \cdot \mathfrak{r} \rightarrow p : m_2\}$: here, \mathfrak{r} is supposed to send either m_1 or m_2 to p , depending on what message p sent to \mathfrak{q} earlier – but \mathfrak{r} is not privy to the interactions between other roles, and correspondingly, the projection $G' \upharpoonright \mathfrak{r} = p \oplus m_1 \sqcap p \oplus m_2$ is undefined.

A.2. Partial session types and duality

A session type S can be further projected onto a role p , to “isolate” the interactions of S involving p from those involving other roles (Definition A.4 below). The projection mechanics remind of those in Definition A.1 – but in this case, the result is a *partial session type* (Definition A.2 below); moreover, the merging operator is different (details below). Partial session types feature branches, selections and recursion – but no role annotations: therefore, they are similar to binary session types (except that their payload types are multiparty), and endowed with a notion of duality (Definition A.3 below). As shown in Section 3.2, these technicalities are necessary to define consistency (Definition 3.3).

Definition A.2 (*Partial session types*). *Partial session types*, ranged over by H , have the syntax:

$$H ::= \&_{i \in I} m_i(S_i).H_i \mid \oplus_{i \in I} m_i(S_i).H_i \mid \mathbf{end} \mid \mu t.H \mid \mathbf{t} \quad \text{with } I \neq \emptyset \text{ and } \forall i \in I : \text{fv}(S_i) = \emptyset$$

where m_i range over pairwise distinct *message labels*.

Definition A.3 (*Duality of partial session types*). The *dual* of a partial type H , written \bar{H} , is:

$$\overline{\&_{i \in I} m_i(S_i).H_i} = \oplus_{i \in I} m_i(S_i).\bar{H}_i \quad \overline{\oplus_{i \in I} m_i(S_i).H_i} = \&_{i \in I} m_i(S_i).\bar{H}_i \quad \overline{\mu t.H} = \mu \bar{t}.\bar{H} \quad \overline{\mathbf{t}} = \mathbf{t} \quad \overline{\mathbf{end}} = \mathbf{end}$$

We say that H and H' are dual iff $\bar{H} = H'$.

Definition A.4 (*Partial projection*). The *projection* of S onto p , written $S|_p$, is a partial type defined as:

$$\begin{aligned} (\&_{i \in I} m_i(S_i).S'_i)|_p &= \begin{cases} \&_{i \in I} m_i(S_i).(S'_i|_p) & \text{if } p = q \\ \sqcap_{i \in I} S'_i|_p & \text{if } p \neq q \end{cases} & (q \oplus_{i \in I} m_i(S_i).S'_i)|_p &= \begin{cases} \oplus_{i \in I} m_i(S_i).(S'_i|_p) & \text{if } p = q \\ \sqcap_{i \in I} S'_i|_p & \text{if } p \neq q \end{cases} \\ (\mu t.S)|_p &= \begin{cases} \mu t.(S|_p) & \text{if } S|_p \neq \mathbf{t}' \ (\forall \mathbf{t}') \\ \mathbf{end} & \text{otherwise} \end{cases} & \mathbf{t}|_p &= \mathbf{t} \quad \mathbf{end}|_p = \mathbf{end} \end{aligned}$$

where \sqcap is the *merge operator for partial session types*, defined as:

$$\begin{aligned} \&_{i \in I} m_i(S_i).H_i \sqcap \&_{i \in I} m_i(S_i).H'_i &= \&_{i \in I} m_i(S_i).(H_i \sqcap H'_i) \\ \oplus_{i \in I} m_i(S_i).H_i \sqcap \oplus_{j \in J} m_j(S_j).H'_j &= \oplus_{k \in I \cap J} m_k(S_k).(H_k \sqcap H'_k) \oplus \oplus_{i \in I \setminus J} m_i(S_i).H_i \oplus \oplus_{j \in J \setminus I} m_j(S_j).H'_j \\ \mu t.H \sqcap \mu t'.H' &= \mu t.(H \sqcap H') \quad \mathbf{t} \sqcap \mathbf{t}' = \mathbf{t} \quad \mathbf{end} \sqcap \mathbf{end}' = \mathbf{end} \end{aligned}$$

For example, if we have a session type $S = p \oplus_{i \in I} m_i(S_i).\mathbf{end}$, its projection $S|_p$ is the partial session type $\oplus_{i \in I} m_i(S_i).\mathbf{end}$, reflecting the fact that S requires to perform an internal choice towards p . Instead, if S is projected onto a role $q \neq p$, we need to skip the first output towards p , compute the projections onto q along each branch m_i ($i \in I$), and merge them, using the partial session type merging operator \sqcap . Therefore, we have $S|_q = \sqcap_{i \in I} (\mathbf{end}|_q) = \sqcap_{i \in I} \mathbf{end} = \mathbf{end}$ – i.e., q is not involved in any interaction in S .

Note that, unlike the session type merging operator in Definition A.1, the partial type merging in Definition A.4 allows to combine different *internal* choices, but *not* different *external* choices. Therefore, we have the following examples:

$$\begin{aligned} S' &= p \& \left\{ \begin{array}{l} m_1(S_1).q \oplus m_1(S_1) \\ m_2(S_2).q \oplus m_2(S_2) \end{array} \right\} & S'|_q &= (q \oplus m_1(S_1))|_q \sqcap (q \oplus m_2(S_2))|_q = \oplus_{m_1(S_1)} \sqcap \oplus_{m_2(S_2)} = \oplus \left\{ \begin{array}{l} m_1(S_1) \\ m_2(S_2) \end{array} \right\} \\ S'' &= p \& \left\{ \begin{array}{l} m_1(S_1).q \& m_1(S_1) \\ m_2(S_2).q \& m_2(S_2) \end{array} \right\} & S''|_q &= (q \& m_1(S_1))|_q \sqcap (q \& m_2(S_2))|_q = \&_{m_1(S_1)} \sqcap \&_{m_2(S_2)} \quad (\text{undefined}) \end{aligned}$$

The partial projection $S'|_q$ is defined, and says that S' might send either m_1 or m_2 to q ; by inspecting S' , we can further notice that m_1 (resp. m_2) should be sent to q only if m_1 (resp. m_2) is first received from p – but this dependency from p 's messages is safely approximated as an internal choice.

Instead, the partial projection $S''|_q$ is undefined: this is because S'' might be willing to receive either m_1 or m_2 from q , depending on which message is first received from p – but this might lead to incorrect interactions, as shown below:

$$s[x]:S'', s[p]:x \oplus m_1(S_1), s[q]:x \oplus m_2(S_2) \rightarrow s[x]:q \& m_1(S_1), s[p]:\mathbf{end}, s[q]:x \oplus m_2(S_2) \not\rightarrow \quad (\text{A.1})$$

Since the projection $S''|_q$ is undefined, the typing context in (A.1) is not consistent (Definition 3.3). By imposing such restrictions, the “classic” MPST theory excludes typing contexts where, e.g., an output from q to x might not be matched by a corresponding input.⁴ For more examples and discussion, see Sections 3.3 and 3.4.

⁴ Notably, the typing context in (A.1) is also *not* live: we can easily verify that it does not satisfy Definition 4.1.

The final ingredient for Definition 3.3 is partial session subtyping, formalised in Definition A.5 below.

Definition A.5 (Partial session subtyping). The subtyping relation \leq for partial types is coinductively defined by the rules:

$$\frac{\forall i \in I \quad S_i \leq T_i \quad H'_i \leq H''_i}{\&_{i \in I} \mathfrak{M}_i(S_i).H'_i \leq \&_{i \in I \cup J} \mathfrak{M}_i(T_i).H''_i} \quad \frac{\forall i \in I \quad T_i \leq S_i \quad H'_i \leq H''_i}{\oplus_{i \in I \cup J} \mathfrak{M}_i(S_i).H'_i \leq \oplus_{i \in I} \mathfrak{M}_i(T_i).H''_i}$$

$$\frac{}{\mathbf{end} \leq \mathbf{end}} \quad \frac{H\{\mu\mathbf{t}.H/t\} \leq H'}{\mu\mathbf{t}.H \leq H'} \quad \frac{H \leq H'\{\mu\mathbf{t}.H'/t\}}{H \leq \mu\mathbf{t}.H'}$$

Definition A.5 is similar to binary session subtyping, studied by Gay and Hole [17]: it says that a smaller type is “less demanding”, as it allows to choose between more internal choices, and requires to support less external choices.

Appendix B. Typing context reductions

Proposition B.1. For all $\Delta, \Delta', \Delta_0$ such that $\text{dom}(\Delta) \cap \text{dom}(\Delta_0) = \emptyset$: $\Delta \rightarrow \Delta'$ if and only if $\Delta, \Delta_0 \rightarrow \Delta', \Delta_0$, with Δ_0 not reducing.

Proof. (\implies) By induction on the size of Δ_0 , and by rule [T-COMP] of Definition 2.10.

(\impliedby) By induction on the size of Δ_0 , and by inversion of rule [T-COMP] of Definition 2.10. \square

Proposition B.2. For all $\Delta, \Delta', \Delta_0$ such that $\text{dom}(\Delta) \cap \text{dom}(\Delta_0) = \emptyset$: $\Delta \xrightarrow{n \text{ steps}} \Delta'$ if and only if $\Delta, \Delta_0 \xrightarrow{n \text{ steps}} \Delta', \Delta_0$, with Δ_0 not reducing.

Proof. By induction on n , using (B.1) in the inductive case. \square

Appendix C. Liveness

Proposition C.1. If $\text{live}(\Delta)$ and $\Delta \xrightarrow{*} \Delta'$, then $\text{live}(\Delta')$.

Proof. By induction on the number of reductions in $\Delta \xrightarrow{*} \Delta'$. The base case (0 reductions) is trivial, as it implies $\Delta' = \Delta$. The inductive case ($n+1$ reductions) is proved by the induction hypothesis, and clause [L- \rightarrow] of Definition 4.1 (liveness). \square

Proposition C.2. Assume $\text{live}(\Delta_0)$. Then, $\Delta, \Delta_0 \rightarrow \Delta'$ implies $\Delta' = \Delta'', \Delta''_0$ such that either:

- (a) $\Delta'' = \Delta$ and $\Delta_0 \rightarrow \Delta''_0$, or
- (b) $\Delta \rightarrow \Delta''$ and $\Delta''_0 = \Delta_0$.

Proof. Notice that we must have $\text{dom}(\Delta) \cap \text{dom}(\Delta_0) = \emptyset$ (otherwise, Δ, Δ_0 would be undefined, by Definition 2.9). By contradiction, assume that neither (a) nor (b) holds. Then, one entry of Δ and one of Δ_0 interact, via rule [T-COMM] of Definition 2.10. Hence, Δ_0 contains some (possibly recursive) internal/external choice that cannot be triggered by the other elements of Δ_0 : this violates clause [L- \oplus]/[L- $\&$] of Definition 4.1, and therefore, Δ_0 is not live – contradiction. We conclude that either (a) or (b) holds. \square

Proposition C.3. Assume $\text{live}(\Delta_0)$. Then, $\Delta, \Delta_0 \xrightarrow{n \text{ steps}} \Delta'$ implies $\Delta' = \Delta'', \Delta''_0$ such that, for some $m \leq n$, $\Delta \xrightarrow{m \text{ steps}} \Delta''$ and $\Delta_0 \xrightarrow{n-m \text{ steps}} \Delta''_0$; moreover, $\Delta, \Delta_0 \xrightarrow{m \text{ steps}} \Delta'', \Delta_0$.

Proof. By induction on n using Proposition C.2. The “moreover...” part of the statement is an immediate consequence of the first part, and Proposition B.2. \square

Lemma C.4. Assume $\text{live}(\Delta_0)$ and $\text{dom}(\Delta) \cap \text{dom}(\Delta_0) = \emptyset$. Then, $\text{live}(\Delta, \Delta_0)$ if and only if $\text{live}(\Delta)$.

Proof. Assuming the hypotheses, we have:

$$\text{live}(\Delta_0) \quad (\text{by hypothesis}) \quad (\text{C.1})$$

$$\text{dom}(\Delta) \cap \text{dom}(\Delta_0) = \emptyset \quad (\text{by hypothesis}) \quad (\text{C.2})$$

(\implies) We show that the following predicate:

$$\mathbb{P} = \{\Delta \mid \text{live}(\Delta, \Delta_0)\} \quad (\text{C.3})$$

satisfies the clauses of Definition 4.1 (liveness). We proceed by examining each $\Delta \in \mathbb{P}$, and we have the following (non-mutually exclusive) cases:

- $\Delta = \Delta', s[p]:S$ with $S = \mathfrak{q}\oplus_{i \in I} m_i(S_i).S'_i$. We need to show that Δ satisfies clause $[L-\oplus]$. We have:

$$\text{live}(\Delta, \Delta_0) \quad \text{and thus} \quad \text{live}(\Delta', s[p]:S, \Delta_0) \quad (\text{by (C.3)}) \quad (\text{C.4})$$

$$\forall i \in I : \exists \Delta'' : \Delta', s[p]:S, \Delta_0 \rightarrow^* \Delta'', s[p]:S'_i, \Delta_0 \quad (\text{by (C.4), [L-}\oplus\text{], (C.1) and Proposition C.3}) \quad (\text{C.5})$$

$$\forall i \in I : \exists \Delta'' : \Delta', s[p]:S \rightarrow^* \Delta'', s[p]:S'_i \quad (\text{by (C.5) and Proposition B.2}) \quad (\text{C.6})$$

and by the shape of Δ and (C.6), we conclude that Δ satisfies clause $[L-\oplus]$;

- $\Delta = \Delta', s[p]:S$ with $S = \mathfrak{q}\&_{i \in I} m_i(S_i).S'_i$. We need to show that Δ satisfies clause $[L-\&]$. We have:

$$\text{live}(\Delta, \Delta_0) \quad \text{and thus} \quad \text{live}(\Delta', s[p]:S, \Delta_0) \quad (\text{by (C.3)}) \quad (\text{C.7})$$

$$\exists i \in I : \exists \Delta'' : \Delta', s[p]:S, \Delta_0 \rightarrow^* \Delta'', s[p]:S'_i, \Delta_0 \quad (\text{by (C.7), [L-}\&\text{], (C.1) and Proposition C.3}) \quad (\text{C.8})$$

$$\exists i \in I : \exists \Delta'' : \Delta', s[p]:S \rightarrow^* \Delta'', s[p]:S'_i \quad (\text{by (C.8) and Proposition B.2}) \quad (\text{C.9})$$

and by the shape of Δ and (C.9), we conclude that Δ satisfies clause $[L-\&]$;

- $\Delta = \Delta', s[p]:\mu t.S$. We need to show that Δ satisfies clause $[L-\mu]$. We have:

$$\text{live}(\Delta, \Delta_0) \quad \text{and thus} \quad \text{live}(\Delta', s[p]:\mu t.S, \Delta_0) \quad (\text{by (C.3)}) \quad (\text{C.10})$$

$$\text{live}(\Delta', s[p]:S\{\mu t.S/t\}, \Delta_0) \quad (\text{by (C.10) and [L-}\mu\text{]}) \quad (\text{C.11})$$

$$\Delta', s[p]:S\{\mu t.S/t\} \in \mathbb{P} \quad (\text{by (C.11) and (C.3)}) \quad (\text{C.12})$$

and by the shape of Δ and (C.12), we conclude that Δ satisfies clause $[L-\mu]$;

- $\Delta \rightarrow \Delta'$. We need to show that Δ satisfies clause $[L-\rightarrow]$. We have:

$$\Delta, \Delta_0 \rightarrow \Delta', \Delta_0 \quad (\text{by (C.2) and Proposition B.1}) \quad (\text{C.13})$$

$$\text{live}(\Delta, \Delta_0) \quad (\text{by (C.3)}) \quad (\text{C.14})$$

$$\text{live}(\Delta', \Delta_0) \quad (\text{by (C.14), (C.13) and [L-}\rightarrow\text{]}) \quad (\text{C.15})$$

$$\Delta' \in \mathbb{P} \quad (\text{by (C.15) and (C.3)}) \quad (\text{C.16})$$

and by the hypothesis $\Delta \rightarrow \Delta'$ and (C.16), we conclude that Δ satisfies clause $[L-\rightarrow]$.

We have proved that \mathbb{P} satisfies all clauses of Definition 4.1 (liveness). Since live is the largest predicate satisfying such clauses, we have that $\forall \Delta : \Delta \in \mathbb{P}$ implies $\text{live}(\Delta)$; and since $\forall \Delta : \text{live}(\Delta, \Delta_0)$ implies $\Delta \in \mathbb{P}$, we conclude that $\forall \Delta : \text{live}(\Delta, \Delta_0)$ implies $\text{live}(\Delta)$.

(\longleftarrow) We show that the following predicate:

$$\mathbb{P} = \{\Delta, \Delta_0 \mid \text{live}(\Delta)\} \quad (\text{C.17})$$

satisfies the clauses of Definition 4.1 (liveness). We proceed similarly to the case above (\implies): by examining each $\Delta, \Delta_0 \in \mathbb{P}$. The difference is that we now use Proposition B.2 to add (rather than remove) Δ_0 when examining the corresponding element of \mathbb{P} yielded by the set comprehension (C.17). Then, we know that live is the largest predicate satisfying all clauses of Definition 4.1 (liveness); therefore, we have that $\forall \Gamma : \Gamma \in \mathbb{P}$ implies $\text{live}(\Gamma)$; and since $\forall \Delta : \text{live}(\Delta)$ implies $(\Delta, \Delta_0) \in \mathbb{P}$, we conclude that $\forall \Delta : \text{live}(\Delta)$ implies $\text{live}(\Delta, \Delta_0)$. \square

Proposition C.5. $\text{live}(\Delta, x:S)$ if and only if $\text{live}(\Delta)$.

Proof. Straightforward by Definition 4.1 (liveness), noticing that $x:S$ is not relevant for the definition (which only considers channels with roles). \square

Proposition C.6. If $\text{live}(\Delta, c:S)$ and $S \equiv S'$, then $\text{live}(\Delta, c:S')$.

Proof. By showing that the following predicate:

$$\mathbb{P} = \{ \Delta, c : S' \mid \text{live}(\Delta, c : S) \text{ and } S \equiv S' \}$$

satisfies the clauses of Definition 4.1 (liveness). Then, since live is the largest predicate satisfying such clauses, we have that $\forall \Gamma : \Gamma \in \mathbb{P}$ implies $\text{live}(\Gamma)$; and since $\forall \Delta, S' : \text{live}(\Delta, c : S)$ and $S \equiv S'$ implies $(\Delta, c : S') \in \mathbb{P}$, we conclude that $\forall \Delta, S' : \text{live}(\Delta, c : S)$ and $S \equiv S'$ implies $\text{live}(\Delta, c : S')$. \square

Proposition C.7. *If $\text{live}(\Delta, \Delta')$ and $\Delta' \equiv \Delta''$, then $\text{live}(\Delta, \Delta'')$.*

Proof. By induction on the size of Δ' : the base case $\Delta' = \emptyset$ is trivial, and the inductive case is proved using Proposition C.6. \square

Proposition C.8. *If $\text{end}(\Delta)$, then $\text{live}(\Delta)$.*

Proof. Take Δ, Δ' such that $\text{dom}(\Delta) = \text{dom}(\Delta')$ and $\forall c \in \text{dom}(\Delta) : \Delta'(c) = \mathbf{end}$. By Definition 4.5, the singleton predicate $\{ \Delta' \}$ (vacuously) satisfies clauses [L- $\&$], [L- \oplus] [L- μ], and [L- \rightarrow] of Definition 4.1 (liveness): hence, $\text{live}(\Delta')$. Now, assume $\text{end}(\Delta)$: by [T-END] in Fig. 3, we have $\Delta \equiv \Delta'$; therefore, by Proposition C.7, we conclude $\text{live}(\Delta)$. \square

Appendix D. Basic properties of typing rules

Lemma D.1 (Typing inversion). *Assume $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. Then, exactly one of the following holds:*

- [T-0] $P = \mathbf{0}$ and $\text{end}(\Delta)$
- [T-def] $P = \mathbf{def} X(x_1, \dots, x_n) = P' \mathbf{in} Q$ and $\exists S_1, \dots, S_n$ such that:
 - $\Theta, X : S_1, \dots, S_n \vdash P' \triangleright x_1 : S_1, \dots, x_n : S_n \triangleleft \emptyset$
 - $\Theta, X : S_1, \dots, S_n \vdash Q \triangleright \Delta \triangleleft \Gamma$
- [T-X] $P = X(c_1, \dots, c_n)$ and $\exists S_1, \dots, S_n$ such that:
 - $\Delta = \Delta_0, \Delta_1, \dots, \Delta_n$
 - $\Theta \vdash X : S_1, \dots, S_n$
 - $\text{end}(\Delta_0)$
 - $\forall i \in 1..n : \Delta_i \vdash c_i : S_i$
- [T-v] $P = (v s : \Delta_s) Q$ and:
 - $\Delta_s = \{ s[p] : S_p \}_{p \in I}$
 - $\Theta \vdash Q \triangleright \Delta, \Delta_s \triangleleft \Gamma$
- [T-|] $P = P_1 \mid P_2$ and $\exists \Delta_1, \Delta_2$ such that:
 - $\Delta = \Delta_1, \Delta_2$
 - $\Theta \vdash P_1 \triangleright \Delta_1 \triangleleft \Gamma, \Delta_2$
 - $\Theta \vdash P_2 \triangleright \Delta_2 \triangleleft \Gamma, \Delta_1$
- [T- $\&$] $P = c[\mathbf{q}] \sum_{i \in I \cup J} \{ m_i(y_i) . P_i \}$ and $\exists \Delta_0, S$ such that:
 - $\Delta = \Delta_0, c : S$
 - $S \equiv \mathbf{q} \&_{i \in I} m_i(S_i) . S'_i$
 - $\forall i \in I$ and $\forall \Gamma', \Gamma''$ such that $\Gamma \rightarrow^* \Gamma'$ and $\Gamma', c : S \rightarrow \Gamma'', c : S'_i$:
 - * $\Theta \vdash P_i \triangleright \Delta, y_i : S_i, c : S'_i \triangleleft \Gamma''$
- [T- \oplus] $P = c[\mathbf{q}] \oplus_{m_k} \langle d \rangle . P'$ and $\exists \Delta_d, \Delta_0, S$ such that:
 - $\Delta = \Delta_d, \Delta_0, c : S$
 - $S \equiv \mathbf{q} \oplus_{i \in I} m_i(S_i) . S'_i$
 - $\exists k \in I$ such that:
 - * $\Delta_d \vdash d : S_k$
 - * $\forall \Gamma', \Gamma''$ such that $\Gamma \rightarrow^* \Gamma'$ and $\Gamma', c : S \rightarrow \Gamma'', c : S'_k$:
 - $\Theta \vdash P' \triangleright \Delta, c : S'_k \triangleleft \Gamma'', \Delta_d$

Proof. Straightforward by observing that the typing rules in Fig. 3 can be deterministically inverted: i.e., each process production in Fig. 1 can be typed by at most one rule. \square

Lemma 5.6 (Rely context strengthening). *If $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$ and $\Gamma \rightarrow^* \Gamma'$, then $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma'$.*

Proof. By induction on the derivation of $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. The crucial cases are [T- $\&$] and [T- \oplus]: they are both proved by inverting the typing rule (using Lemma D.1) and observing that its premises require the continuations of P to be typed using all reducts of Γ – which include all reducts of Γ' . This allows to apply the induction hypothesis, and conclude by the same rule. \square

Proposition D.2. Assume $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. Take some Γ' such that $\text{live}(\Gamma')$ and $(\text{dom}(\Delta) \cup \text{dom}(\Gamma)) \cap \text{dom}(\Gamma') = \emptyset$. Then, $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma, \Gamma'$.

Proof. By induction on the derivation of $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. \square

Proposition D.3. Assume $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. Take some Δ' such that $\text{end}(\Delta')$ and $(\text{dom}(\Delta) \cup \text{dom}(\Gamma)) \cap \text{dom}(\Delta') = \emptyset$. Then, $\Theta \vdash P \triangleright \Delta, \Delta' \triangleleft \Gamma$.

Proof. By induction on the derivation of $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. \square

Proposition D.4. Assume $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. If $X \notin \text{dom}(\Theta)$, then $\Theta, X:\tilde{S} \vdash P \triangleright \Delta \triangleleft \Gamma$.

Proof. By induction on the derivation of $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. \square

Lemma D.5 (Subject congruence). If $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$ and $P \equiv P'$, then $\Theta \vdash P' \triangleright \Delta \triangleleft \Gamma$.

Proof. Assuming, $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$, the proof proceeds by cases on P , examining the possible shapes of P' allowed by \equiv in Definition 2.4, and using Lemma D.1 to reconstruct their typing. \square

Appendix E. Substitution lemma

Proposition 5.7. Assume $x \notin \text{dom}(\Delta) \cup \text{dom}(\Gamma)$. Then, $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma, x:S$ iff $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$.

Proof. (\implies) By induction on the derivation of $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma, x:S$.

(\impliedby) By induction on the derivation of $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. \square

Corollary E.1. $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$ iff $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma \setminus x$.

Proof. If $x \notin \text{dom}(\Gamma)$, then the statement holds vacuously. Otherwise, if $x \in \text{dom}(\Gamma)$, it is a direct consequence of Proposition 5.7. \square

Proposition E.2. Assume $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$ and $x \in \text{fv}(P)$. Then, $x \in \text{dom}(\Delta)$.

Proof. By induction on the derivation of $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$. \square

Lemma 4.8 (Substitution). If $\Theta \vdash P \triangleright \Delta, x:S \triangleleft \Gamma, \Gamma'$ and $\Gamma \vdash s[p]:S$, then $\Theta \vdash P\{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma'$.

Proof. Assume the hypotheses. We have:

$$\Theta \vdash P \triangleright \Delta, x:S \triangleleft \Gamma \quad (\text{by hypothesis}) \quad (\text{E.1})$$

$$\Gamma = s[p]:S' \text{ with } S \equiv S' \quad (\text{by hypothesis and [T-CHAN] in Fig. 3}) \quad (\text{E.2})$$

$$\text{live}(\Delta, x:S, \Gamma, \Gamma') \quad (\text{by Definition 4.5}) \quad (\text{E.3})$$

$$\text{live}(\Delta, \Gamma, \Gamma') \quad (\text{by (E.3) and Proposition 5.7}) \quad (\text{E.4})$$

Note that, in the judgement in the conclusion, the liveness requirement on the guarantee/rely contexts (Definition 4.5) is trivially satisfied:

$$\text{live}(\Delta, s[p]:S, \Gamma') \quad (\text{by (E.4), (E.2) and Proposition C.7}) \quad (\text{E.5})$$

We proceed by induction on the derivation of $\Theta \vdash P \triangleright \Delta, x:S \triangleleft \Gamma, \Gamma'$:

- Base case [T-0]. We have:

$$P = \mathbf{0} \text{ and } \text{end}(\Delta, x:S) \quad (\text{by conclusion of [T-0] and Lemma D.1, case [T-0]}) \quad (\text{E.6})$$

$$S \equiv \mathbf{end} \quad (\text{by (E.6) and [T-END] in Fig. 3}) \quad (\text{E.7})$$

$$\text{end}(\Delta) \quad (\text{by (E.6) and [T-END] in Fig. 3}) \quad (\text{E.8})$$

$$\text{end}(\Delta, s[p]:S) \quad (\text{by (E.8), (E.7) and [T-END] in Fig. 3}) \quad (\text{E.9})$$

$$P\{s[p]/x\} = \mathbf{0}\{s[p]/x\} = \mathbf{0} \quad (\text{from (E.6)}) \quad (\text{E.10})$$

$$\frac{\text{end}(\Delta, s[p]:S)}{\Theta \vdash P\{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma'} \quad [\text{T-0}] \quad (\text{by (E.9), (E.10) and (E.5)})$$

- Base case [T-X]. We have:

$$\frac{\Theta \vdash X : S_1, \dots, S_n \quad \text{end}(\Delta_0) \quad \forall i \in 1..n \quad \Delta_i \vdash c_i : S_i}{\Theta \vdash P \triangleright \Delta, x : S \triangleleft \Gamma} \quad [\text{T-X}] \quad (\text{by Lemma D.1, case [T-X]}) \quad (\text{E.11})$$

$$P = X(c_1, \dots, c_n) \quad (\text{from the conclusion of (E.11)}) \quad (\text{E.12})$$

$$\Delta, x : S = \Delta_0, \Delta_1, \dots, \Delta_n \quad (\text{by (E.11)}) \quad (\text{E.13})$$

Now, from (E.13), we have two sub-cases, depending on whether x is used as parameter for X (i.e., whether $c_i = x$, for some $i \in 1..n$):

– x is *not* used as parameter for calling X , i.e., $\forall i \in 1..n : c_i \neq x$. Then:

$$P\{s[p]/x\} = P \quad (\text{from the hypothesis and (E.12)}) \quad (\text{E.14})$$

$$\Delta_0 \vdash x : S \quad (\text{from the hypothesis, (E.11) and (E.13)}) \quad (\text{E.15})$$

$$\Delta_0 = x : S \text{ and } S \equiv \mathbf{end} \quad (\text{by (E.15), (E.11) and [T-END] in Fig. 3}) \quad (\text{E.16})$$

$$\text{end}(s[p]:S) \quad (\text{by (E.16) and [T-END] in Fig. 3}) \quad (\text{E.17})$$

$$\frac{\Theta \vdash X : S_1, \dots, S_n \quad \text{end}(s[p]:S) \quad \forall i \in 1..n \quad \Delta_i \vdash c_i : S_i}{\Theta \vdash P\{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma'} \quad [\text{T-X}] \quad (\text{by (E.16), (E.17), (E.14)})$$

– x is used as parameter for calling X , i.e., $c_k = x$ for some $k \in 1..n$. Then:

$$P\{s[p]/x\} = X(c_1, \dots, c_{k-1}, s[p], c_{k+1}, \dots, c_n) \quad (\text{from the hypothesis and (E.12)}) \quad (\text{E.18})$$

$$\Delta_k \equiv x : S \quad (\text{by (E.11) and [T-CHAN] in Fig. 3}) \quad (\text{E.19})$$

$$\frac{\Theta \vdash X : S_1, \dots, S_n \quad \text{end}(\Delta_0) \quad s[p]:S \vdash s[p]:S \quad \forall i \in \{1..n\} \setminus \{k\} \quad \Delta_i \vdash c_i : S_i}{\Theta \vdash P\{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma'} \quad [\text{T-X}] \quad (\text{by (E.11), (E.19) and (E.18)})$$

- Inductive case [T-&]. In this case, P is a branching on some channel. We have two sub-cases:

(a) P is a branching on some $c \neq x$. Then, the substitution in the statement does not involve c , and the statement holds by the induction hypothesis and rule [T-&]; or

(b) P is a branching on x , which in turn is substituted by a channel with role $s[p]$.

We detail the proof for the latter (and most interesting) case. We have:

$$P = x[q] \sum_{i \in I \cup J} \{m_i(y_i) \cdot P_i\} \quad (\text{conclusion of rule [T-&]}) \quad (\text{E.20})$$

$$\frac{\begin{array}{l} S \equiv q \&_{i \in I} m_i(S_i) \cdot S'_i \quad \forall i \in I \\ \forall \Gamma'', \Gamma''' : \Gamma, \Gamma' \xrightarrow{*} \Gamma'' \text{ and } \Gamma'', x : S \rightarrow \Gamma''', x : S'_i \\ \Theta \vdash P_i \triangleright \Delta, y_i : S_i, x : S'_i \triangleleft \Gamma'''' \end{array}}{\Theta \vdash x[q] \sum_{i \in I \cup J} \{m_i(y_i) \cdot P_i\} \triangleright \Delta, x : S \triangleleft \Gamma, \Gamma'} \quad [\text{T-&}] \quad (\text{by Lemma D.1, case [T-&]}) \quad (\text{E.21})$$

$$\Gamma = s[p]:S' \text{ with } S' \equiv S \equiv q \&_{i \in I} m_i(S_i) \cdot S'_i \quad (\text{by (E.2), and (E.21)}) \quad (\text{E.22})$$

We can now observe that in (E.21), by Definition 2.10, we have $\forall i \in I : \Gamma'', x : S \rightarrow \Gamma''', x : S'_i$ – i.e., the typing context entry for x reduces “in isolation”, and does not involve the rest of the typing context Γ'' . Therefore, from the reductions in (E.21) we have:

$$\forall i \in I : \forall \Gamma'', \Gamma''' : \Gamma, \Gamma' \multimap^* \Gamma'' \text{ and } \Gamma'', x:S \rightarrow \Gamma''', x:S'_i \quad (\text{E.23})$$

$$\text{implies } \Gamma''' = \Gamma'' \text{ and } \Theta \vdash P_i \triangleright \Delta, y_i:S_i, x:S'_i \triangleleft \Gamma'' \quad (\text{from (E.21), premises of [T-\&]}) \quad (\text{E.24})$$

In order to apply the induction hypothesis, we focus on a *subset* of the reductions in (E.23) that involves the entries for x and $s[p]$ as specified in (E.25) below (recall that $\Gamma = s[p]:S' \equiv s[p]:S$, by (E.22)):

$$\forall i \in I : \forall \Gamma'', \Gamma''' \text{ such that } \begin{cases} \Gamma' \multimap^* \Gamma'' \text{ and} \\ \Gamma'', s[p]:S \rightarrow \Gamma''', s[p]:S'_i \text{ and} \\ \Gamma''', s[p]:S'_i, x:S \rightarrow \Gamma''', s[p]:S'_i, x:S'_i : \end{cases} \quad (\text{E.25})$$

$$\Theta \vdash P_i \triangleright \Delta, y_i:S_i, x:S'_i \triangleleft \Gamma''', s[p]:S'_i \quad (\text{by (E.25) and (E.24)}) \quad (\text{E.26})$$

$$s[p]:S'_i \vdash s[p]:S'_i \quad (\text{by Definition 4.5}) \quad (\text{E.27})$$

$$\Theta \vdash P_i\{s[p]/x\} \triangleright \Delta, y_i:S_i, s[p]:S'_i \triangleleft \Gamma'''' \quad (\text{by (E.26), (E.27) and induction hyp.}) \quad (\text{E.28})$$

By exploiting (E.28), we can now conclude:

$$P\{s[p]/x\} = s[p][q] \sum_{i \in I \cup J} \{m_i(y_i). P_i\{s[p]/x\}\} \quad (\text{by (E.20)}) \quad (\text{E.29})$$

$$\frac{\begin{array}{l} S \equiv q \&_{i \in I} m_i(S_i). S'_i \quad \forall i \in I \\ \forall \Gamma'', \Gamma'''' : \Gamma' \multimap^* \Gamma'' \text{ and } \Gamma'', s[p]:S \rightarrow \Gamma''', s[p]:S'_i \\ \Theta \vdash P_i\{s[p]/x\} \triangleright \Delta, y_i:S_i, s[p]:S'_i \triangleleft \Gamma'''' \end{array}}{\Theta \vdash P\{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma'} \quad [\text{T-\&}] \quad (\text{by (E.29), (E.28)})$$

- Inductive case [T- \oplus]. In this case, P is a selection on some variable. We have two sub-cases:
 - (a) P is a selection on some $c \neq x$. Then, the substitution in the statement does not involve c , and the statement holds by the induction hypothesis and rule [T- \oplus]; or
 - (b) P is a selection on x , which in turn is substituted by a channel with role $s[p]$.
 We detail the proof for the latter (and most interesting) case. The strategy is similar to the one we adopted for proving the case [T-\&]. We have:

$$\exists k \in I : P = x[q] \oplus_{m_k} \langle d \rangle . P' \quad (\text{conclusion of rule [T-\oplus]}) \quad (\text{E.30})$$

$$\frac{\begin{array}{l} S \equiv q \oplus_{i \in I} m_i(S_i). S'_i \quad \exists k \in I \quad \Delta_d \vdash d:S_k \\ \forall \Gamma'', \Gamma'''' : \Gamma, \Gamma' \multimap^* \Gamma'' \text{ and } \Gamma'', x:S \rightarrow \Gamma''', x:S'_k \\ \Theta \vdash P' \triangleright \Delta_0, x:S'_k \triangleleft \Gamma''', \Delta_d \end{array}}{\Delta = \Delta_0, \Delta_d \text{ and } \Theta \vdash x[q] \oplus_{m_k} \langle d \rangle . P' \triangleright \Delta, x:S \triangleleft \Gamma, \Gamma'} \quad [\text{T-\oplus}] \quad (\text{by Lemma D.1, case [T-\oplus]}) \quad (\text{E.31})$$

$$\Gamma = s[p]:S' \text{ with } S' \equiv S \equiv s[p]:q \oplus_{i \in I} m_i(S_i). S'_i \quad (\text{by (E.2), and (E.31)}) \quad (\text{E.32})$$

$$\forall \Gamma'', \Gamma'''' : \Gamma \multimap^* \Gamma'' \text{ and } \Gamma'', x:S \rightarrow \Gamma''', x:S'_k \text{ implies } \Gamma'''' = \Gamma'' \text{ and } \Theta \vdash P \triangleright \Delta_0, x:S'_k \triangleleft \Gamma'', \Delta_d \quad (\text{from (E.31), premises of [T-\oplus]}) \quad (\text{E.33})$$

$$\forall \Gamma'', \Gamma'''' \text{ such that } \begin{cases} \Gamma' \multimap^* \Gamma'' \text{ and} \\ \Gamma'', s[p]:S \rightarrow \Gamma''', s[p]:S'_k \text{ and} \\ \Gamma''', s[p]:S'_k, x:S \rightarrow \Gamma''', s[p]:S'_k, x:S'_k : \end{cases} \quad (\text{E.34})$$

$$\Theta \vdash P \triangleright \Delta_0, x:S'_k \triangleleft \Gamma''', s[p]:S'_k, \Delta_d \quad (\text{by (E.34) and (E.33)}) \quad (\text{E.35})$$

$$s[p]:S'_k \vdash s[p]:S'_k \quad (\text{by Definition 4.5}) \quad (\text{E.36})$$

$$\Theta \vdash P'\{s[p]/x\} \triangleright \Delta_0, s[p]:S'_k \triangleleft \Gamma''', \Delta_d \quad (\text{by (E.35), (E.36) and induction hyp.}) \quad (\text{E.37})$$

$$P\{s[p]/x\} = s[p][q] \oplus_{m_k} \langle d \rangle . P'\{s[p]/x\} \quad (\text{by (E.30)}) \quad (\text{E.38})$$

$$\frac{\begin{array}{l} S \equiv q \oplus_{i \in I} m_i(S_i). S'_i \quad \exists k \in I \quad \Delta_d \vdash d:S_k \\ \forall \Gamma'', \Gamma'''' : \Gamma' \multimap^* \Gamma'' \text{ and } \Gamma'', s[p]:S \rightarrow \Gamma''', s[p]:S'_k \\ \Theta \vdash P'\{s[p]/x\} \triangleright \Delta_0, s[p]:S'_k \triangleleft \Gamma''', \Delta_d \end{array}}{\Theta \vdash P\{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma'} \quad [\text{T-\oplus}] \quad (\text{by (E.38), (E.37), (E.31)})$$

- Inductive case $[T-|]$. We have:

$$P = P_1 | P_2 \quad (\text{conclusion of rule } [T-|]) \quad (\text{E.39})$$

$$\Delta = \Delta_1, \Delta_2 \text{ and } \Theta \vdash P_1 | P_2 \triangleright \Delta_1, \Delta_2, x:S \triangleleft \Gamma, \Gamma' \quad (\text{by (E.39), Lemma D.1, case } [T-|]) \quad (\text{E.40})$$

Now, the typing judgement in (E.40) implies that, in the corresponding instance of $[T-|]$, $x:S$ could be used to type either P_1 or P_2 in the rule premises – i.e., we can have either:

$$\left\{ \begin{array}{l} \Theta \vdash P_1 \triangleright \Delta_1, x:S \triangleleft \Gamma, \Gamma', \Delta_2 \\ \Theta \vdash P_2 \triangleright \Delta_2 \triangleleft \Gamma, \Gamma', \Delta_1, x:S \\ x \notin \text{fv}(P_2) \end{array} \right. \quad (\text{by Proposition E.2}) \quad (\text{E.41})$$

or

$$\left\{ \begin{array}{l} \Theta \vdash P_1 \triangleright \Delta_1 \triangleleft \Gamma, \Gamma', \Delta_2, x:S \\ \Theta \vdash P_2 \triangleright \Delta_2, x:S \triangleleft \Gamma, \Gamma', \Delta_1 \\ x \notin \text{fv}(P_1) \end{array} \right. \quad (\text{by Proposition E.2}) \quad (\text{E.42})$$

We proceed assuming (E.41) (the proof for (E.42) is symmetric). We have:

$$\Theta \vdash P_1 \{s[p]/x\} \triangleright \Delta_1, s[p]:S \triangleleft \Gamma', \Delta_2 \quad (\text{by (E.41), (E.2) and induction hyp.}) \quad (\text{E.43})$$

$$\Theta \vdash P_2 \triangleright \Delta_2 \triangleleft \Gamma', \Delta_1, s[p]:S \quad (\text{by (E.41), (E.2) and Proposition 5.7}) \quad (\text{E.44})$$

$$P \{s[p]/x\} = (P_1 | P_2) \{s[p]/x\} = P_1 \{s[p]/x\} | P_2 \{s[p]/x\} \quad (\text{by (E.39)}) \quad (\text{E.45})$$

$$P_2 = P_2 \{s[p]/x\} \quad (\text{by (E.41), since } x \notin \text{fv}(P_2)) \quad (\text{E.46})$$

$$\frac{\begin{array}{l} \Theta \vdash P_1 \{s[p]/x\} \triangleright \Delta_1, s[p]:S \triangleleft \Gamma', \Delta_2 \\ \Theta \vdash P_2 \{s[p]/x\} \triangleright \Delta_2 \triangleleft \Gamma', \Delta_1, s[p]:S \end{array}}{\Theta \vdash P \{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma'} \quad [T-|] \quad (\text{by (E.45), (E.43), (E.44), (E.46), (E.40)}) \quad (\text{E.47})$$

- Inductive case $[T\text{-def}]$. We have:

$$P = \mathbf{def} X(x_1, \dots, x_n) = Q \mathbf{in} R \quad (\text{conclusion of rule } [T\text{-def}]) \quad (\text{E.48})$$

$$\frac{\begin{array}{l} \Theta, X:S_1, \dots, S_n \vdash Q \triangleright x_1:S_1, \dots, x_n:S_n \triangleleft \emptyset \\ \Theta, X:S_1, \dots, S_n \vdash R \triangleright \Delta, x:S \triangleleft \Gamma, \Gamma' \end{array}}{\Theta \vdash P \triangleright \Delta, x:S \triangleleft \Gamma, \Gamma'} \quad [T\text{-def}] \quad (\text{by (E.48) and Lemma D.1, case } [T\text{-def}]) \quad (\text{E.49})$$

$$\text{fv}(Q) \subseteq \{x_1, \dots, x_n\} \quad (\text{by (E.49) (typing of } Q) \text{ and Proposition E.2 } \times n) \quad (\text{E.50})$$

$$\text{fv}(X(x_1, \dots, x_n) = Q) = \emptyset \quad (\text{by (E.50)}) \quad (\text{E.51})$$

$$P \{s[p]/x\} = \mathbf{def} X(x_1, \dots, x_n) = Q \mathbf{in} (R \{s[p]/x\}) \quad (\text{by (E.48) and (E.51)}) \quad (\text{E.52})$$

$$\Theta, X:S_1, \dots, S_n \vdash R \{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma' \quad (\text{by (E.49) (typing of } R) \text{, induction hyp.}) \quad (\text{E.53})$$

$$\frac{\begin{array}{l} \Theta, X:S_1, \dots, S_n \vdash Q \triangleright x_1:S_1, \dots, x_n:S_n \triangleleft \emptyset \\ \Theta, X:S_1, \dots, S_n \vdash R \{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma' \end{array}}{\Theta \vdash P \{s[p]/x\} \triangleright \Delta, s[p]:S \triangleleft \Gamma'} \quad [T\text{-def}] \quad (\text{by (E.52), (E.49) and (E.53)}) \quad \square$$

Appendix F. Subject reduction

Theorem 4.9 (Subject reduction). *If $\Theta \vdash P \triangleright \Delta \triangleleft \Gamma$ and $P \rightarrow^* P'$, then $\exists \Delta'$ such that $\Delta \rightarrow^* \Delta'$ and $\Theta \vdash P' \triangleright \Delta' \triangleleft \Gamma$.*

Proof. Assume:

$$\Theta \vdash P \triangleright \Delta \triangleleft \Gamma \quad (\text{by hypothesis}) \quad (\text{F.1})$$

Note that by (F.1) and Definition 4.5, we have $\text{live}(\Delta, \Gamma)$, and to obtain a valid typing judgement for P' , we must ensure $\text{live}(\Delta', \Gamma)$: by Proposition C.1, this is guaranteed once we prove $\Delta \rightarrow^* \Delta'$, since it implies $\Delta, \Gamma \rightarrow^* \Delta', \Gamma$.

First, we prove the following property, for *one* reduction step of P :

$$P \rightarrow P' \text{ implies } \exists \Delta' \text{ such that } \Delta \rightarrow^* \Delta' \text{ and } \Theta \vdash P' \triangleright \Delta' \triangleleft \Gamma \quad (\text{F.2})$$

Assuming the transition $P \rightarrow P'$, we proceed by induction on its derivation.

• Base case [R-COMM]. We have:

$$P = s[\mathbf{p}][\mathbf{q}] \sum_{i \in I} \{m_i(y_i) \cdot P_i\} \mid s[\mathbf{q}][\mathbf{p}] \oplus_{m_k} (s'[\mathbf{r}]) \cdot Q \quad (k \in I) \quad (\text{by inversion of the rule}) \quad (\text{F.3})$$

$$P' = P_k \{s'[\mathbf{r}]/y_k\} \mid Q \quad (\text{by inversion of the rule}) \quad (\text{F.4})$$

$\exists \Delta_1, \Delta_2 : \Delta = \Delta_1, \Delta_2$ and

$$\Theta \vdash s[\mathbf{p}][\mathbf{q}] \sum_{i \in I} \{m_i(y_i) \cdot P_i\} \triangleright \Delta_1 \triangleleft \Gamma, \Delta_2$$

$$\frac{\Theta \vdash s[\mathbf{q}][\mathbf{p}] \oplus_{m_k} (s'[\mathbf{r}]) \cdot Q \triangleright \Delta_2 \triangleleft \Gamma, \Delta_1}{\Theta \vdash P \triangleright \Delta \triangleleft \Gamma} \quad [\text{T-}|]$$

(by (F.3), Lemma D.1, case [T-|]) (F.5)

$\exists \Gamma_1 = \Gamma, \Delta_2$ and $\exists \Gamma_2 = \Gamma, \Delta_1$

(by (F.5)) (F.6)

$\exists J : \emptyset \neq J \subseteq I \quad \exists \Delta'_1, S : \Delta_1 = \Delta'_1, s[\mathbf{p}]:S$ such that

$$S \equiv \mathbf{q} \&_{i \in J} m_i(S_i) \cdot S'_i \quad \forall i \in J$$

$\forall \Gamma', \Gamma'' : \Gamma_1 \rightarrow^* \Gamma'$ and $\Gamma', s[\mathbf{p}]:S \rightarrow \Gamma'', s[\mathbf{p}]:S'_i$ (by (F.5), Lemma D.1, case [T-&]) (F.7)

$$\Theta \vdash P_i \triangleright \Delta'_1, y_i : S_i, s[\mathbf{p}]:S'_i \triangleleft \Gamma''$$

$$\frac{\Theta \vdash P_i \triangleright \Delta'_1, y_i : S_i, s[\mathbf{p}]:S'_i \triangleleft \Gamma''}{\Theta \vdash s[\mathbf{p}][\mathbf{q}] \sum_{i \in I} \{m_i(y_i) \cdot P_i\} \triangleright \Delta_1 \triangleleft \Gamma_1} \quad [\text{T-&}]$$

$\exists I' \neq \emptyset \quad \exists \Delta_{2s'}, \Delta'_2, T : \Delta_2 = \Delta_{2s'}, \Delta'_2, s[\mathbf{q}]:T$ such that

$$T \equiv \mathbf{p} \oplus_{i \in I'} m_i(T_i) \cdot T'_i$$

$$\exists k \in I' \quad \Delta_{2s'} \vdash s'[\mathbf{r}]:T_k$$

$\forall \Gamma', \Gamma'' : \Gamma_2 \rightarrow^* \Gamma'$ and $\Gamma', s[\mathbf{q}]:T \rightarrow \Gamma'', s[\mathbf{q}]:T'_k$ (by (F.5), Lemma D.1, case [T- \oplus]) (F.8)

$$\Theta \vdash Q \triangleright \Delta'_2, s[\mathbf{q}]:T'_k \triangleleft \Gamma'', \Delta_{2s'}$$

$$\frac{\Theta \vdash Q \triangleright \Delta'_2, s[\mathbf{q}]:T'_k \triangleleft \Gamma'', \Delta_{2s'}}{\Theta \vdash s[\mathbf{q}][\mathbf{p}] \oplus_{m_k} (s'[\mathbf{r}]) \cdot Q \triangleright \Delta_2 \triangleleft \Gamma_2} \quad [\text{T-}\oplus]$$

$$s[\mathbf{q}]:\mathbf{p} \oplus_{i \in I'} m_i(T_i) \cdot T'_i \equiv s[\mathbf{q}]:T \in \Delta_2 \subsetneq \Gamma_1 \quad (\text{by (F.8) and (F.6)}) \quad (\text{F.9})$$

$$s[\mathbf{p}]:\mathbf{q} \&_{i \in J} m_i(S_i) \cdot S'_i \equiv s[\mathbf{p}]:S \in \Delta_1 \subsetneq \Gamma_2 \quad (\text{by (F.7) and (F.6)}) \quad (\text{F.10})$$

$$I' \subseteq J \text{ and therefore } k \in J \quad (\text{by (F.9), (F.6), (F.7), live}(\Delta_1, \Gamma_1) \text{ and Definition 4.1}) \quad (\text{F.11})$$

$$\forall i \in J : S_i \equiv T_i \quad (\text{by (F.11), live}(\Delta_1, \Gamma_1) \text{ and Definition 4.1}) \quad (\text{F.12})$$

$$\Delta_{2s'} \vdash s'[\mathbf{r}]:S_k \quad (\text{by (F.8) and (F.12)}) \quad (\text{F.13})$$

$$s'[\mathbf{r}]:S_k \equiv \Gamma_1 \quad (\text{by (F.13), (F.8) and (F.6)}) \quad (\text{F.14})$$

$$\forall i \in J : s[\mathbf{p}]:S, s[\mathbf{q}]:T \rightarrow \equiv s[\mathbf{p}]:S'_i, s[\mathbf{q}]:T_i \quad ((\text{F.11}), (\text{F.9}), \text{live}(\Delta_1, \Gamma_1), \text{Definition 4.1}) \quad (\text{F.15})$$

$$\exists \Gamma'_1 : \Gamma_1, s[\mathbf{p}]:S \rightarrow \equiv \Gamma'_1, s[\mathbf{p}]:S'_k, s[\mathbf{q}]:T'_k \quad (\text{by (F.9) and (F.15)}) \quad (\text{F.16})$$

$$\Theta \vdash P_k \triangleright \Delta'_1, y_k : S_k, s[\mathbf{p}]:S'_k \triangleleft \Gamma'_1 \quad (\text{by (F.16) and (F.7)}) \quad (\text{F.17})$$

$$s'[\mathbf{r}]:S_k \equiv \Gamma'_1 \quad (\text{by (F.14) and (F.16)}) \quad (\text{F.18})$$

$$\Theta \vdash P_k \{s'[\mathbf{r}]/y_k\} \triangleright \Delta'_1, s'[\mathbf{r}]:S_k, s[\mathbf{p}]:S'_k \triangleleft \Gamma'_1 \setminus s'[\mathbf{r}] \quad ((\text{F.17}), (\text{F.18}), \text{Lemma 4.8}) \quad (\text{F.19})$$

$$\Gamma'_1 \setminus s'[\mathbf{r}] \equiv \Gamma, \Delta'_2, s[\mathbf{q}]:T'_k \quad (\text{by (F.6), (F.8), (F.13), and (F.16)}) \quad (\text{F.20})$$

$$\exists \Gamma'_2 : \Gamma_2, s[\mathbf{q}]:T \rightarrow \equiv \Gamma'_2, s[\mathbf{p}]:S'_k, s[\mathbf{q}]:T'_k \quad (\text{by (F.10) and (F.15)}) \quad (\text{F.21})$$

$$\Theta \vdash Q \triangleright \Delta'_2, s[\mathbf{q}]:T'_k \triangleleft \Gamma'_2, \Delta_{2s'} \quad (\text{by (F.21) and (F.8)}) \quad (\text{F.22})$$

$$\Gamma'_2, \Delta_{2s'} \equiv \Gamma, \Delta'_1, \Delta_{2s'}, s[\mathbf{p}]:S'_k \quad (\text{by (F.6), (F.7), (F.21), (F.13) and Definition 4.5}) \quad (\text{F.23})$$

$\exists \Delta' = \Delta'_1, \Delta'_2$ such that:

$$\Delta'_1 = \Delta'_1, \Delta_{2s'}, s[\mathbf{p}]:S'_k \quad (\text{F.24})$$

$$\Delta'_2 = \Delta'_2, s[\mathbf{q}]:T'_k \quad (\text{F.25})$$

$$\Delta \xrightarrow{*} \Delta' \quad (\text{by (F.5), (F.7), (F.8), (F.24), (F.25)}) \quad (\text{F.26})$$

$$\frac{\frac{\Theta \vdash P_k \{s'[\mathbf{x}]/y_k\} \triangleright \Delta'_1 \triangleleft \Gamma, \Delta'_2}{\Theta \vdash Q \triangleright \Delta'_2 \triangleleft \Gamma, \Delta'_1}}{\Theta \vdash P_k \{s'[\mathbf{x}]/y_k\} \mid Q \triangleright \Delta' \triangleleft \Gamma} \text{[T-]} \quad \left(\begin{array}{l} \text{by (F.19), (F.24), (F.20), (F.25);} \\ \text{by (F.22), (F.25), (F.23), (F.24)} \end{array} \right) \quad (\text{F.27})$$

• Base case [R-X]. We have:

$$P = \mathbf{def} X(x_1, \dots, x_n) = Q \mathbf{in} (X(s_1[\mathbf{p}_1], \dots, s_1[\mathbf{p}_1]) \mid R) \quad (\text{by inversion of the rule}) \quad (\text{F.28})$$

$$P' = \mathbf{def} X(x_1, \dots, x_n) = Q \mathbf{in} (Q \{s_1[\mathbf{p}_1]/x_1\} \cdots \{s_n[\mathbf{p}_n]/x_n\} \mid R) \quad (\text{by inversion of the rule}) \quad (\text{F.29})$$

$$\Delta = \Delta_*, \Delta_R \text{ and } \Delta_* = \Delta_0, \Delta_1, \dots, \Delta_n \text{ and } \Theta' = \Theta, X:S_1, \dots, S_n \text{ such that:} \quad (\text{F.30})$$

$$\frac{\frac{\frac{\Theta' \vdash X:S_1, \dots, S_n}{\text{end}(\Delta_0) \quad \forall i \in 1..n \quad \Delta_i \vdash s_i[\mathbf{p}_i]:S_i} \text{[T-X]} \quad \frac{\Theta' \vdash X(s_1[\mathbf{p}_1], \dots, s_1[\mathbf{p}_1]) \triangleright \Delta_* \triangleleft \Gamma, \Delta_R}{\Theta' \vdash R \triangleright \Delta_R \triangleleft \Gamma, \Delta_*} \text{[T-]} \quad \frac{\Theta' \vdash Q \triangleright x_1:S_1, \dots, x_n:S_n \triangleleft \emptyset}{\Theta' \vdash X(s_1[\mathbf{p}_1], \dots, s_1[\mathbf{p}_1]) \mid R \triangleright \Delta \triangleleft \Gamma} \text{[T-def]} \quad \frac{\Theta' \vdash X(s_1[\mathbf{p}_1], \dots, s_1[\mathbf{p}_1]) \mid R \triangleright \Delta \triangleleft \Gamma}{\Theta \vdash P \triangleright \Delta \triangleleft \Gamma} \text{[T-def]} \quad (\text{by (F.28), (F.1), Lemma D.1, cases [T-def],[T-],[T-X]}) \quad (\text{F.31})$$

$$\text{live}(\Delta_0) \quad (\text{by (F.31) and Proposition C.8}) \quad (\text{F.32})$$

$$\text{live}(\Delta_1, \dots, \Delta_n, \Gamma, \Delta_R) \quad (\text{by (F.31), Definition 4.5, (F.32), Lemma C.4}) \quad (\text{F.33})$$

$$\text{live}((\Delta_1, \dots, \Delta_n, \Gamma, \Delta_R) \setminus x_1, \dots, x_n) \quad (\text{by (F.33) and Proposition C.5} \times n) \quad (\text{F.34})$$

$$\text{live}(\Delta_1, \dots, \Delta_n, ((\Gamma, \Delta_R) \setminus x_1, \dots, x_n)) \quad (\text{by (F.34), (F.31) and Definition 4.5}) \quad (\text{F.35})$$

$$\Theta' \vdash Q \triangleright x_1:S_1, \dots, x_n:S_n \triangleleft \Delta_1, \dots, \Delta_n, ((\Gamma, \Delta_R) \setminus x_1, \dots, x_n) \quad ((\text{F.35}), (\text{F.31}), \text{Proposition D.2}) \quad (\text{F.36})$$

$$\Theta' \vdash Q \{s_1[\mathbf{p}_1]/x_1\} \cdots \{s_n[\mathbf{p}_n]/x_n\} \triangleright \Delta_1, \dots, \Delta_n \triangleleft (\Gamma, \Delta_R) \setminus x_1, \dots, x_n \quad \left(\begin{array}{l} (\text{F.36}), (\text{F.31}), \\ \text{Lemma 4.8} \times n \end{array} \right) \quad (\text{F.37})$$

$$\Theta' \vdash Q \{s_1[\mathbf{p}_1]/x_1\} \cdots \{s_n[\mathbf{p}_n]/x_n\} \triangleright \Delta_* \triangleleft \Gamma, \Delta_R \quad \left(\begin{array}{l} (\text{F.37}), (\text{F.30}), (\text{F.31}), \\ \text{Proposition D.3, Corollary E.1} \times n \end{array} \right) \quad (\text{F.38})$$

$$\text{Let } \Delta' = \Delta = \Delta_*, \Delta_R = \Delta_0, \Delta_1, \dots, \Delta_n, \Delta_R \text{ (hence, } \Delta \xrightarrow{*} \Delta') \quad (\text{by (F.30)}) \quad (\text{F.39})$$

$$\frac{\frac{\Theta' \vdash Q \{s_1[\mathbf{p}_1]/x_1\} \cdots \{s_n[\mathbf{p}_n]/x_n\} \triangleright \Delta_* \triangleleft \Gamma, \Delta_R}{\Theta' \vdash R \triangleright \Delta_R \triangleleft \Gamma, \Delta_*} \text{[T-]} \quad \frac{\Theta' \vdash Q \triangleright x_1:S_1, \dots, x_n:S_n \triangleleft \emptyset}{\Theta' \vdash Q \{s_1[\mathbf{p}_1]/x_1\} \cdots \{s_n[\mathbf{p}_n]/x_n\} \mid R \triangleright \Delta' \triangleleft \Gamma} \text{[T-]} \quad \frac{\Theta' \vdash Q \{s_1[\mathbf{p}_1]/x_1\} \cdots \{s_n[\mathbf{p}_n]/x_n\} \mid R \triangleright \Delta' \triangleleft \Gamma}{\Theta \vdash P' \triangleright \Delta' \triangleleft \Gamma} \text{[T-def]} \quad (\text{by (F.38), (F.39), (F.29), (F.31)})$$

• Inductive case [R-]. We have:

$$P = P_1 \mid P_2 \quad (\text{by inversion of the rule}) \quad (\text{F.40})$$

$$\exists \Delta_1, \Delta_2: \Delta = \Delta_1, \Delta_2 \quad (\text{by (F.40), (F.1) and Lemma D.1, case [T-]}) \quad (\text{F.41})$$

$$\Theta \vdash P_1 \triangleright \Delta_1 \triangleleft \Gamma, \Delta_2 \quad (\text{by (F.40), (F.41) and Lemma D.1, case [T-]}) \quad (\text{F.42})$$

$$\Theta \vdash P_2 \triangleright \Delta_2 \triangleleft \Gamma, \Delta_1 \quad (\text{by (F.40), (F.41) and Lemma D.1, case [T-]}) \quad (\text{F.43})$$

$$P_1 \rightarrow P'_1 \text{ and } P' = P'_1 \mid P_2 \quad (\text{by inversion of [R-]}) \quad (\text{F.44})$$

$$\exists \Delta'_1: \Delta_1 \xrightarrow{*} \Delta'_1 \text{ and } \Theta \vdash P'_1 \triangleright \Delta'_1 \triangleleft \Gamma, \Delta_2 \quad (\text{by (F.42), (F.44) and i.h.}) \quad (\text{F.45})$$

$$\Gamma, \Delta_1 \xrightarrow{*} \Gamma, \Delta'_1 \quad (\text{by (F.45) and Definition 2.10}) \quad (\text{F.46})$$

$$\Theta \vdash P_2 \triangleright \Delta_2 \triangleleft \Gamma, \Delta'_1 \quad (\text{by (F.43), (F.46) and Lemma 5.6}) \quad (\text{F.47})$$

$$\exists \Delta' = \Delta'_1, \Delta_2 \text{ such that } \Delta \xrightarrow{*} \Delta' \quad (\text{by (F.41), (F.45) and Definition 2.10}) \quad (\text{F.48})$$

$$\frac{\Theta \vdash P'_1 \triangleright \Delta'_1 \triangleleft \Gamma, \Delta_2 \quad \Theta \vdash P_2 \triangleright \Delta_2 \triangleleft \Gamma, \Delta'_1}{\Theta \vdash P' \triangleright \Delta' \triangleleft \Gamma} \text{[T-]} \quad (\text{by (F.44), (F.45), (F.47), (F.48)})$$

• Inductive case [R- ν]. We have:

$$\exists Q : P = (\nu s) Q \quad (\text{by inversion of the rule}) \quad (\text{F.49})$$

$$P \rightarrow P' \text{ implies } \exists Q' : Q \rightarrow Q' \text{ and } P' = (\nu s) Q' \quad (\text{by (F.49), inv. of [R-}\nu\text{]}) \quad (\text{F.50})$$

$$\frac{\Delta_s = \{c[p]:S_p\}_{p \in I} \quad \Theta \vdash Q \triangleright \Delta, \Delta_s \triangleleft \Gamma}{\Theta \vdash (\nu s : \Delta_s) Q \triangleright \Delta \triangleleft \Gamma} \text{[T-}\nu\text{]} \quad (\text{by (F.49), Lemma D.1, case [T-}\nu\text{]}) \quad (\text{F.51})$$

$$\exists \Delta'' : \Delta, \Delta_s \xrightarrow{*} \Delta'' \text{ and } \Theta \vdash Q' \triangleright \Delta'' \triangleleft \Gamma \quad (\text{by (F.51), (F.50), i.h.}) \quad (\text{F.52})$$

$$\text{live}(\Delta_s) \quad (\text{by (F.51) (premise and conclusion of [T-}\nu\text{]) and Lemma C.4}) \quad (\text{F.53})$$

$$\exists \Delta', \Delta'_s : \Delta'' = \Delta', \Delta'_s \text{ and } \Delta \xrightarrow{*} \Delta' \text{ and } \Delta_s \xrightarrow{*} \Delta'_s \quad ((\text{F.51}), (\text{F.52}), (\text{F.53}), \text{Proposition C.3}) \quad (\text{F.54})$$

$$\Delta, \Gamma \xrightarrow{*} \Delta', \Gamma \quad (\text{by (F.54) and Definition 2.10}) \quad (\text{F.55})$$

$$\text{live}(\Delta', \Gamma) \quad (\text{by (F.51), (F.55) and Proposition C.1}) \quad (\text{F.56})$$

$$\frac{\Delta'_s = \{s[p]:S'_p\}_{p \in I} \quad \Theta \vdash Q' \triangleright \Delta', \Delta'_s \triangleleft \Gamma}{\Theta \vdash P' \triangleright \Delta' \triangleleft \Gamma} \text{[T-}\nu\text{]} \quad (\text{by (F.50), (F.52), (F.54), (F.56)})$$

• Inductive case [R-def]. We have:

$$P = \mathbf{def} X(\tilde{x}) = Q \mathbf{in} R \quad (\text{by inversion of the rule}) \quad (\text{F.57})$$

$$P' = \mathbf{def} X(\tilde{x}) = Q \mathbf{in} R' \text{ with } R \rightarrow R' \quad (\text{from rule premises}) \quad (\text{F.58})$$

$$\frac{\Theta, X:S_1, \dots, S_n \vdash Q \triangleright x_1:S_1, \dots, x_n:S_n \triangleleft \emptyset \quad \Theta, X:S_1, \dots, S_n \vdash R \triangleright \Delta \triangleleft \Gamma}{\Theta \vdash P \triangleright \Delta \triangleleft \Gamma} \text{[T-def]} \quad (\text{by (F.57), (F.1), Lemma D.1, case [T-def]}) \quad (\text{F.59})$$

$$\exists \Delta' : \Delta \xrightarrow{*} \Delta' \text{ and } \Theta, X:S_1, \dots, S_n \vdash R' \triangleright \Delta' \triangleleft \Gamma \quad (\text{by (F.59), (F.58), i.h.}) \quad (\text{F.60})$$

$$\frac{\Theta, X:S_1, \dots, S_n \vdash Q \triangleright x_1:S_1, \dots, x_n:S_n \triangleleft \emptyset \quad \Theta, X:S_1, \dots, S_n \vdash R' \triangleright \Delta' \triangleleft \Gamma}{\Theta \vdash P' \triangleright \Delta' \triangleleft \Gamma} \text{[T-def]} \quad (\text{by (F.60) and (F.58)})$$

• Inductive case [R= \equiv]. We have:

$$P \equiv Q \rightarrow Q' \equiv P' \text{ (for some } Q, Q') \quad (\text{by inversion of the rule}) \quad (\text{F.61})$$

$$\Theta \vdash Q \triangleright \Delta \triangleleft \Gamma \quad (\text{by (F.1), (F.61) and Lemma D.5}) \quad (\text{F.62})$$

$$\exists \Delta' : \Delta \xrightarrow{*} \Delta' \text{ and } \Theta \vdash Q' \triangleright \Delta' \triangleleft \Gamma \quad (\text{by (F.62), (F.61) and induction hyp.}) \quad (\text{F.63})$$

$$\Theta \vdash P' \triangleright \Delta' \triangleleft \Gamma \quad (\text{by (F.63), (F.61) and Lemma D.5})$$

We have thus proved (F.2). Then, we prove the main statement by induction on the number of reductions in $P \xrightarrow{*} P'$: the base case is trivial (0 reductions, and thus $P = P'$); in the inductive case ($n + 1$ reductions), we apply the induction hypothesis and (F.2). \square

References

- [1] D. Ancona, V. Bono, M. Bravetti, J. Campos, G. Castagna, P.-M. Deniérou, S.J. Gay, N. Gesbert, E. Giachino, R. Hu, E.B. Johnsen, F. Martins, V. Mascardi, F. Montesi, R. Neykova, N. Ng, L. Padovani, V.T. Vasconcelos, N. Yoshida, Behavioral types in programming languages, *Foundations and Trends in Programming Languages* 3 (2–3) (2017) 95–230.
- [2] G. Bernardi, O. Dardha, S.J. Gay, D. Kouzapas, On duality relations for session types, in: *Trustworthy Global Computing (TGC)*, 2014, pp. 51–66.
- [3] G. Bernardi, M. Hennessy, Using higher-order contracts to model session types (extended abstract), in: *CONCUR*, in: *Lect. Notes Comput. Sci.*, vol. 8704, Springer, 2014, pp. 387–401.
- [4] G. Bernardi, M. Hennessy, Using higher-order contracts to model session types, *Log. Methods Comput. Sci.* 12 (2) (2016).
- [5] L. Bettini, M. Coppo, L. D’Antoni, M. De Luca, M. Dezani-Ciancaglini, N. Yoshida, Global progress in dynamically interleaved multiparty sessions, in: *CONCUR*, 2008, pp. 418–433.
- [6] L. Bocchi, J. Lange, N. Yoshida, Meeting deadlines together, in: *CONCUR*, in: *LIPICs. Leibniz Int. Proc. Inform.*, vol. 42, 2015, pp. 283–296.
- [7] V. Bono, L. Padovani, Typing copyless message passing, *Log. Methods Comput. Sci.* 8 (1) (2012).
- [8] D. Brand, P. Zafirovulo, On communicating finite-state machines, *J. ACM* 30 (2) (Apr. 1983) 323–342.
- [9] L. Caires, F. Pfenning, Session types as intuitionistic linear propositions, in: *CONCUR*, 2010, pp. 222–236.
- [10] L. Caires, F. Pfenning, B. Toninho, Linear logic propositions as session types, *Math. Struct. Comput. Sci.* 26 (3) (2016) 367–423.
- [11] M. Coppo, M. Dezani-Ciancaglini, L. Padovani, N. Yoshida, A gentle introduction to multiparty asynchronous session types, in: *Formal Methods for Multicore Programming*, 2015, pp. 146–178.
- [12] M. Coppo, M. Dezani-Ciancaglini, N. Yoshida, L. Padovani, Global progress for dynamically interleaved multiparty sessions, *Math. Struct. Comput. Sci.* 26 (2) (2016) 238–302.
- [13] P. Deniérou, N. Yoshida, A. Bejleri, R. Hu, Parameterised multiparty session types, *Log. Methods Comput. Sci.* 8 (4) (2012).
- [14] P.-M. Deniérou, N. Yoshida, Multiparty compatibility in communicating automata: characterisation and synthesis of global session types, in: *ICALP*, 2013, pp. 174–186.
- [15] M. Dezani-Ciancaglini, S. Ghilezan, S. Jakšić, J. Pantović, N. Yoshida, Precise subtyping for synchronous multiparty sessions, in: *PLACES 2015*, vol. 203, 2016, pp. 29–43.
- [16] S. Gay, A. Ravara (Eds.), *Behavioural Types: From Theory to Tools*, River Publishers Series in Automation, Control and Robotics, 2017.
- [17] S.J. Gay, M. Hole, Subtyping for session types in the pi calculus, *Acta Inform.* 42 (2–3) (2005) 191–225.
- [18] J.-Y. Girard, Linear logic, *Theor. Comput. Sci.* 50 (1987) 1–102.
- [19] K. Honda, Types for dyadic interaction, in: *CONCUR*, 1993, pp. 509–523.
- [20] K. Honda, V.T. Vasconcelos, M. Kubo, Language primitives and type discipline for structured communication-based programming, in: *ESOP*, 1998, pp. 122–138.
- [21] K. Honda, N. Yoshida, M. Carbone, Multiparty asynchronous session types, in: *POPL*, 2008, pp. 273–284, full version: [22].
- [22] K. Honda, N. Yoshida, M. Carbone, Multiparty asynchronous session types, *J. ACM* 63 (1) (2016) 9:1–9:67.
- [23] H. Hüttel, I. Lanese, V.T. Vasconcelos, L. Caires, M. Carbone, P.-M. Deniérou, D. Mostrous, L. Padovani, A. Ravara, E. Tuosto, H.T. Vieira, G. Zavattaro, Foundations of session types and behavioural contracts, *ACM Comput. Surv.* 49 (1) (2016) 3:1–3:36.
- [24] D. Kouzapas, N. Yoshida, Globally governed session semantics, in: *CONCUR*, 2013, pp. 395–409.
- [25] D. Kouzapas, N. Yoshida, Globally governed session semantics, *Log. Methods Comput. Sci.* 10 (4) (2014).
- [26] J. Lange, E. Tuosto, Synthesising choreographies from local session types, in: *CONCUR*, 2012, pp. 225–239.
- [27] J. Lange, E. Tuosto, N. Yoshida, From communicating machines to graphical choreographies, in: *POPL*, 2015, pp. 221–232.
- [28] S. Lindley, J.G. Morris, Talking bananas: structural recursion for session types, in: *ICFP*, 2016, pp. 434–447.
- [29] R. Milner, *Communicating and Mobile Systems: The π -Calculus*, Cambridge University Press, 1999.
- [30] R. Milner, J. Parrow, D. Walker, A calculus of mobile processes, parts I and II, *Inf. Comput.* 100 (1) (1992) 1–77.
- [31] B.C. Pierce, *Types and Programming Languages*, MIT Press, MA, USA, 2002.
- [32] D. Sangiorgi, D. Walker, *The π -Calculus: A Theory of Mobile Processes*, Cambridge University Press, 2001.
- [33] A. Scalas, O. Dardha, R. Hu, N. Yoshida, A linear decomposition of multiparty sessions for safe distributed programming, in: *ECOOP*, 2017, pp. 24:1–24:31.
- [34] A. Scalas, O. Dardha, R. Hu, N. Yoshida, A linear decomposition of multiparty sessions for safe distributed programming (artifact), *Dagstuhl Artifacts Ser.* 3 (1) (2017) 3:1–3:2.
- [35] K. Takeuchi, K. Honda, M. Kubo, An interaction-based language and its typing system, in: *PARLE*, 1994, pp. 398–413.
- [36] P. Wadler, Propositions as sessions, in: *ICFP*, 2012, pp. 273–286.
- [37] P. Wadler, Propositions as sessions, *J. Funct. Program.* 24 (2–3) (2014) 384–418.
- [38] N. Yoshida, P. Deniérou, A. Bejleri, R. Hu, Parameterised multiparty session types, in: *FoSSaCS*, 2010, pp. 128–145.