Multiparty Asynchronous Session Types

KOHEI HONDA, Queen Mary University of London
NOBUKO YOSHIDA, Imperial College London
MARCO CARBONE, IT University of Copenhagen

Communication is becoming one of the central elements in software development. As a potential typed foundation for structured communication-centred programming, session types have been studied over the last decade for a wide range of process calculi and programming languages, focussing on binary (two-party) sessions. This work extends the foregoing theories of binary session types to multiparty, asynchronous sessions, which often arise in practical communication-centred applications. Presented as a typed calculus for mobile processes, the theory introduces a new notion of types in which interactions involving multiple peers are directly abstracted as a global scenario. Global types retain the friendly type syntax of binary session types while specifying dependencies and capturing complex causal chains of multiparty asynchronous interactions. A global type plays the role of a shared agreement among communication peers, and is used as a basis of efficient type checking through its projection onto individual peers. The fundamental properties of the session type discipline such as communication safety, progress and session fidelity are established for general n-party asynchronous interactions.

CCS Concepts: •**Theory of computation** → **Distributed computing models; Process calculi;** *Type theory; Type structures; Program analysis;* Operational semantics;•**Software and its engineering** → *Distributed programming languages; Concurrent programming structures;*

Additional Key Words and Phrases: Session Types, the Pi-Calculus, Projection, Global Types, Global Protocols, Progress

ACM Reference Format:
Multiparty Asynchronous Session Types 0, 0, Article 0 (0), 67 pages.
DOI: 0000001.0000001

1. INTRODUCTION

Background. Communication is becoming one of the central elements in software development, ranging from web services to parallel scientific computing to multi-core programming. One of the main application areas of communication-based systems is business protocols. A *business protocol* is a series of structured and automated interactions among two or more business entities. During the 1990s, there were many attempts at describing and modelling business protocols in order to achieve, e.g., automation, scalability and correctness of protocols. As a result, several institutions started investing heavily in distributed computing technologies, for the purpose of reducing the risk of centralised controls.

A preliminary version of this article appeared in Proceedings of 35th annual ACM SIGPLAN - SIGACT Symposium on Principles of Programming Languages (POPL 2008).
Yoshida was supported by EPSRC EPSRC EP/K011715/1, EP/K034413/1, and EP/L00058X/1, EU project FP7-612985 UpScale and COST Action IC1201 BETTY. Carbone was supported by the Chords (granted by the Danish Agency for Science, Technology and Innovation) and COST Action IC1201 BETTY.
Author's addresses: K. Honda, School of Electronic Engineering and Computer Science, Queen Mary, University of London, Mile End Road, London E1 4NS, United Kingdom; N. Yoshida, Department of Computing, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom; M. Carbone, IT University of Copenhagen, Rued Langgaards Vej 7, 2300 Copenhagen, Denmark.

DOI: 0000001.0000001

Against this background, the Web Services Choreography Description Language Working Group (WS-CDL WG) [WS-CDL 2003] was formed by W3C with the goal of defining a language standard for specifying web service business protocols by means of distributed interactions among peers (business entities). Recognising the need for a foundational theory on which the design and infrastructure of the language were to be built, the working group took a strong interest in the π-calculus, leading to the involvement of Robin Milner and the authors as official invited experts in the standardisation process. While WS-CDL's design is informed by the π-calculus in both communication primitives and structuring constructs, WS-CDL differs from the π-calculus in that it describes message flows among multiple participants *globally*.

Engineers have always found it essential to use various notations for describing interaction patterns globally, such as the notations for cryptographic protocols, Message Sequence Charts [International Telecommunication Union 1996], and UML sequence diagrams. This is because a global description presents information on the behaviour of systems which is not immediately available from the corresponding endpoint-based descriptions: how conversations among multiple participants evolve and interleave, what are the synchronisation/communication points among participants, and how they together induce a desired global scenario. More formally, under a certain well-formedness condition, a global protocol automatically ensures that interactions satisfy the safety and deadlock-freedom properties.

WS-CDL follows these preceding global notations: an underlying intuition of its term *choreography* may be summarised as:

> *"Dancers dance following a global scenario (choreography) without a single point of control."*

Once specified, this scenario is to be executed by individual distributed processes without orchestrating nodes. A global description is meant to be executed by distributed interactions among end-point processes. Thus each global description should be *projected* onto processes at each end-point, whose mutual communications precisely realise the original global scenario. This translation from a global description to end-point processes is called *end-point projection* (EPP), in the terminology of the WS-CDL WG. The theory of multiparty session types introduced in this article was born from the attempts to formalise the EPP in WS-CDL, applying the idea to *types*, in order to overcome a significant technical limitation of binary session types. This is one step beyond from WS-CDL and gives closer links to recent tools for web services such as BPMN 2.0 Choreography [BPMNC 2012], as explained in the next paragraph.

Session Types. Over the last decade, *session types*, introduced in the 1990s [Honda 1993; Takeuchi et al. 1994; Honda et al. 1998], have provided a potential typed foundation for the design of communication-based systems. The main intuition behind session types is that a communication-centred application often exhibits a highly structured sequence of interactions involving, e.g., sequencing and branching, which as a whole form a natural unit of conversation called *session*. The structure of a conversation is abstracted as a type through an intuitive syntax, which is then used as a basis for validating programs through associated language primitives and a typing discipline.

As an example, consider a simple business protocol between a buyer (Buyer) and a seller (Seller) from Buyer's viewpoint: Buyer sends the title of a book (a string), Seller sends a quote (an integer). If Buyer is satisfied by the quote, he then sends his address (a string) and Seller sends back the delivery date (a date); otherwise he quits the conversation. This can be precisely described by the following session type:

$$!\texttt{string};\ ?\texttt{int};\ \oplus\{\ \texttt{ok} : !\texttt{string};\ ?\texttt{date};\ \texttt{end},\quad \texttt{quit} :\ \texttt{end}\ \} \tag{1}$$

The session type above denotes patterns of communication operations (where ; denotes sequencing and $\oplus$ choice) describing Buyer's communication behaviour in the business protocol. In particular, the term !`string` denotes an output of a value of type `string`, whereas ?`int` denotes an input of a value of type `int`. The choice $\oplus$ features the two options ok and quit. The term end represents the termination of the session. From Seller's viewpoint the same session is described by the dual type

$$?\texttt{string};\ !\texttt{int};\ \&\{\ \mathsf{ok} : ?\texttt{string};\ !\texttt{date};\ \mathsf{end},\quad \mathsf{quit} :\ \mathsf{end}\ \} \tag{2}$$

in which $\&$ means that a choice is offered.

Such an explicit representation of conversation structures allows us to deal with one of the most common bugs in communication-based programming, namely, the synchronisation bug. A programmer expects that communicating programs should together realise a consistent conversation, but, unfortunately, they can easily fail to handle a specific incoming message or to send a message at the correct timing, with no way to detect such errors before runtime. An explicit specification as in (1) guides to principled programming of communication behaviour and enables automatic protocol validation [WS-CDL 2003]. In addition, a clean separation between abstraction and implementation given by type-based abstraction and associated primitives leads to intelligible programs and flexible implementations [Hu et al. 2008; Hu et al. 2010]. Underlying these merits are the following central properties guaranteed by session types.

(1) Interactions within a session never incur a communication error (*communication safety*).
(2) Channels for a session are used linearly (linearity) and are deadlock-free in a single session (*progress*).
(3) The communication sequence in a session follows the scenario declared in the session type (*session fidelity, predictability*).

As a consequence of these properties, at each step in a session, a single input and a single output or a single selection and a single branching can take place via a session channel, moving to the next step.

Our previous research shows that the session-based programming framework is applicable to a wide range of calculi, programming languages and computing environments, including calculi of mobile processes [Takeuchi et al. 1994; Gay and Hole 2005; Honda et al. 1998; Bonelli and Compagnoni 2007; Mezzina 2008; Yoshida and Vasconcelos 2007; Gay 2008; Dezani-Ciancaglini et al. 2007; Carbone et al. 2008], higher-order processes [Mostrous and Yoshida 2007; 2009], Ambients [Garralda et al. 2006], multi-threaded ML [Vasconcelos et al. 2006; Gay and Vasconcelos 2009], multicore programming [Yoshida et al. 2008], Haskell [Neubauer and Thiemann 2004a; Pucella and Tov 2008], F# [Bhargavan et al. 2009; Swamy et al. 2011], operating systems [Fähndrich et al. 2006], Java [Dezani-Ciancaglini et al. 2006; Coppo et al. 2007; Dezani-Ciancaglini et al. 2009; Hu et al. 2008; Gay et al. 2010; Hu et al. 2010; Ng et al. 2011] and Web Services [Carbone et al. 2006; Carbone et al. 2007; WS-CDL 2003; Carbone et al. 2012; Sparkes 2006; Honda et al. 2007].

Multiparty Asynchronous Sessions. The foregoing studies on session types have focussed on binary (two-party) sessions. While many conversation patterns can be captured through a composition of binary sessions, there are cases where binary session types are not powerful enough for describing and validating interactions which involve more than two parties.

As an example, let us consider a simple refinement of the above Buyer-Seller protocol: consider two buyers, Buyer1 and Buyer2, who wish to buy an expensive book from Seller by combining their money. Buyer1 sends the title of the book to Seller, Seller

sends to both Buyer1 and Buyer2 its quote, Buyer1 tells Buyer2 how much she can pay, and Buyer2 either accepts the quote or rejects the quote by notifying Seller. It is extremely awkward (if logically possible) to decompose this scenario into three binary sessions, between Buyer1 and Seller, between Buyer2 and Seller, and between Buyer1 and Buyer2. Abstracting this protocol as three separate session types also means that our type abstraction loses essential sequencing information in this interaction scenario. For validating this conversation scenario as a whole, therefore, the conversation structure should be represented as a *single session*.

Many existing business protocols including financial protocols are written as a collaboration of several peers. Typical message-passing parallel algorithms also frequently demand distribution of a request to, and collection of the results from, many peers. All these usecases are most naturally abstracted as single multiparty sessions.

Furthermore, many of these applications are implemented with an *asynchronous transport* where the senders send the messages without being blocked (but often preserving their order), to avoid the heavy overhead of synchronisation. The widely used network transport, such as TCP, provides this mechanism through familiar APIs to alleviate the latency problem. Asynchronous message passing is also a standard assumption in financial messaging [AMQP 2015], parallel algorithms oand distributed objects and functions [Coppo et al. 2007; Hu et al. 2010; Hu et al. 2008; Ng et al. 2011; Neubauer and Thiemann 2004b; Fähndrich et al. 2006]. Thus, we ask:

> *Can we generalise the foregoing binary session types to multiparty asynchronous sessions preserving clarity and their key formal properties?*

Challenges of Multiparty Asynchronous Sessions. To answer this open question, we face two major technical difficulties. First, the simplicity and tractability of the theory of binary sessions come from a notion of *duality* in interactions found in Linear Logic [Girard 1987]. Consider the binary session type specified in (1) for Buyer. Not only Buyer's behaviour can be checked against the session type, but also the whole conversation structure is already represented in this single type, since the interaction pattern of Seller is fully given as this type's dual (exchanging input and output and branching and selection in the original type). When composing two parties, we only have to check they have mutually dual types. This framework based on duality is no longer effective in multiparty communication where the whole conversation cannot be constructed from only single behaviour. We need an effective means to abstract as a type a global scenario which a programmer wishes to realise through interacting programs (hence against which she would wish to check their correctness), and establish an effective method to ensure composability.

Second, linearity analysis of channels, which is the key for ensuring safety and progress, becomes highly involved under a combination of asynchrony and multiparty communication since a conflict of actions can arise more easily. A linearity property holds if a communication via the same channel of a global type does not break the order of messages as it is specified in the global description. This demands a precise causal analysis for correct sequencing of interactions distributed among multi-peers.

This Work. This paper presents a generalisation of binary session types to multiparty sessions for the π-calculus. We propose *three* major technical contributions in order to overcome the aforementioned challenges:

(1) A new notion of types which can directly abstract intended conversation structure among n-parties as *global scenarios*, retaining an intuitive type syntax.
(2) Consistency criteria for a conversation structure with respect to a protocol specification given as a causality analysis of actions in global types, modularly articulating different kinds of dependency.

(3) A type discipline for individual processes (programs) which uses global types through their *projections* onto individual end-point participants: the resulting end-point types are directly associated with individual processes for type checking.

The idea of type abstraction based on a global view (Point 1) comes from an *abstract version of "choreography"* developed in a W3C web services working group [Carbone et al. 2006; WS-CDL 2003]. Causality structures in asynchronous interactions are precisely and modularly captured in the abstract setting of global types, offering a foundation for the type discipline (Point 2). Through the use of global types, we propose a new effective method for designing, type-checking and developing programs based on multiparty sessions (Point 3).

Let us illustrate Point 3 in detail. First, we design and agree upon a global type G as an intended conversation scenario. A team of programmers then develops code, one for each participant, incrementally validating its conformance to (the projections of) G. When programs are executed, their interactions automatically follow the stipulated scenario. The projection can also be used as a hint for modelling, designing and debugging local behaviours of participants. After the development, a global type will serve as a basis of monitoring, maintenance and upgrade. For materialising this design framework, the proposed framework presents a type discipline which can validate whether a program is typable or not, given G (as a shared agreement) and an individual program (as its end-point realiser). The resulting type discipline guarantees all the original key properties of binary session types, such as communication error freedom, progress and session fidelity in a general n-party session, underpinning its practical use. For further discussions on this development framework and its applications developed in industry and academia, see §4.1, §6.1 and §7.

This paper is a full version of [Honda et al. 2008a], with detailed definitions and full proofs. It is also expanded with more examples and comparisons with recent related work. In the remainder, Section 2 gives the syntax and semantics of the calculus, and motivates the key ideas through business and streaming protocol examples and a usecase from [OOI 2015]. Section 3 explains the global types. Section 4 describes the typing system. Section 5 establishes the main results. Section 6 discusses extensions and related works. Section 7 concludes with future issues and a summary of applications, software and languages developed with the industry collaborators based on the multiparty session type theory. The appendix contains the proofs of the propositions, lemmas and theorems stated in the main sections.

2. MULTIPARTY ASYNCHRONOUS SESSIONS

2.1. Syntax for Multiparty Sessions

Several versions of π-calculi with session types have been proposed in the literature. A detailed survey can be found in [Dezani-Ciancaglini and de' Liguoro 2010]. In this work, we use a simple extension of the original language for binary sessions [Honda et al. 1998; Takeuchi et al. 1994] to multiparty sessions.

Informally, a *session* is a series of interactions which serve as a unit of conversation. A session is established among multiple parties via a *shared name*, which represents a public interaction point. Then, fresh *session channels* are generated and shared among all participants who can use them for communicating with each other.

In the remainder, we use the following base sets:

— *shared names* or *names*, ranged over by $a, b, x, y, z, \ldots$;
— *session channels* or *channels*, ranged over by $s, t, ...$;
— *labels*, ranged over by $l, l', \ldots$; and
— *process variables*, ranged over by $X, Y, \ldots$.

$$
\begin{array}{rll}
P & ::= \overline{a}_{[2..\mathtt{n}]}(\tilde{s}).P & \text{multicast session request} \\
& \mid a_{[\mathtt{p}]}(\tilde{s}).P & \text{session acceptance} \\
& \mid s!\langle \tilde{e} \rangle; P & \text{value sending} \\
& \mid s?(\tilde{x}); P & \text{value reception} \\
& \mid s!\langle\!\langle \tilde{s} \rangle\!\rangle; P & \text{session delegation} \\
& \mid s?(\!(\tilde{s})\!); P & \text{session reception} \\
& \mid s \lhd l; P & \text{label selection} \\
& \mid s \rhd \{l_i \colon P_i\}_{i \in I} & \text{label branching} \\
& \mid \mathsf{if}\ e\ \mathsf{then}\ P\ \mathsf{else}\ Q & \text{conditional branch} \\
& \mid P \mid Q & \text{parallel composition} \\
& \mid \mathbf{0} & \text{inaction} \\
& \mid (\boldsymbol{\nu}\, n)P & \text{hiding} \\
& \mid \mathsf{def}\ D\ \mathsf{in}\ P & \text{recursion} \\
& \mid X\langle \tilde{e}\tilde{s} \rangle & \text{process call} \\
& \mid s :: \tilde{h} & \text{message queue} \\
e & ::= v \mid e\ \mathsf{and}\ e' \mid \mathsf{not}\ e \quad \ldots & \text{expressions} \\
v & ::= a \mid \mathsf{true} \mid \mathsf{false} & \text{values} \\
h & ::= l \mid \tilde{v} \mid \tilde{s} & \text{messages-in-transit} \\
D & ::= \{X_i(\tilde{x}_i\tilde{s}_i) = P_i\}_{i \in I} & \text{declaration for recursion}
\end{array}
$$

Fig. 1. Syntax

We use n for either a single shared name or a vector of session channels. The symbol $\tilde{s}$ denotes the vector of session channels $s_1, \ldots, s_k$ for some k. Similarly for other names, channels and variables. Then, *processes*, ranged over by $P, Q \ldots$, and *expressions*, ranged over by $e, e', \ldots$, are given by the grammar in Figure 1. Except for the first two primitives for session initiation and the final message queue, all constructs are from the binary session calculi [Honda et al. 1998]. Session initiation is introduced to establish a session between multiple processes, while message queues are added to model asynchronous session communication, as explained later.

Among the primitives for session initiation, the prefix process $\overline{a}_{[2..\mathtt{n}]}(\tilde{s}).P$ initiates a new session through a shared interaction point a, by distributing a vector of freshly generated session channels $\tilde{s}$ to the remaining $\mathtt{n} - 1$ participants, each of shape $a_{[\mathtt{p}]}(\tilde{s}).Q_{\mathtt{p}}$ for $2 \leq \mathtt{p} \leq \mathtt{n}$. All receive $\tilde{s}$, over which the actual session communications can now take place among the $\mathtt{n}$ parties. $\mathtt{p}$, $\mathtt{q}$,... range over natural numbers called *participants* of a session. As we shall formalise later through operational semantics, these primitives offer a distilled syntactic presentation of "sharing of a fresh context for a new session" among multiple parties.

Session communications are performed using the next three pairs of primitives: sending and receiving, session delegation and reception (the former delegates to the latter the capability to participate in a session by passing channels associated with the session), and selection and branching (the former chooses one of the branches offered by the latter). Branching and selection constructs correspond to external and internal choices.

The next three (the conditional, parallel and inaction) are standard. $(\boldsymbol{\nu} a)P$ makes a local to P while $(\boldsymbol{\nu}\tilde{s})P$ makes $\tilde{s}$ local to P. The recursion and process call primitives

$$P \mid \mathbf{0} \equiv P \qquad P \mid Q \equiv Q \mid P \qquad (P \mid Q) \mid R \equiv P \mid (Q \mid R)$$

$$(\boldsymbol{\nu}\, n)P \mid Q \equiv (\boldsymbol{\nu}\, n)(P \mid Q) \quad \text{if } n \notin \mathsf{fn}(Q)$$

$$(\boldsymbol{\nu}\, nn')P \equiv (\boldsymbol{\nu}\, n'n)P \qquad (\boldsymbol{\nu}\, n)\mathbf{0} \equiv \mathbf{0} \qquad (\boldsymbol{\nu}\, s_1 ... s_n)(s_1 :: \emptyset \mid \cdots s_n :: \emptyset) \equiv \mathbf{0} \qquad \mathsf{def}\ D\ \mathsf{in}\ \mathbf{0} \equiv \mathbf{0}$$

$$\mathsf{def}\ D\ \mathsf{in}\ (\boldsymbol{\nu}\, n)P \equiv (\boldsymbol{\nu}\, n)\mathsf{def}\ D\ \mathsf{in}\ P \quad \text{if } n \notin \mathsf{fn}(D)$$

$$(\mathsf{def}\ D\ \mathsf{in}\ P) \mid Q \equiv \mathsf{def}\ D\ \mathsf{in}\ (P \mid Q) \quad \text{if } \mathsf{dpv}(D) \cap \mathsf{fpv}(Q) = \emptyset$$

$$\mathsf{def}\ D\ \mathsf{in}\ (\mathsf{def}\ D'\ \mathsf{in}\ P) \equiv \ \mathsf{def}\ D \ \cup \ D'\ \mathsf{in}\ P \quad \text{if } \mathsf{dpv}(D) \cap \mathsf{dpv}(D') = \emptyset$$

Fig. 2. Structural congruence.

realise recursive behaviour. $s :: \tilde{h}$ is a *message queue* representing ordered messages in transit $\tilde{h}$ with destination s (which may be considered as a network pipe in a TCP-like transport). $(\boldsymbol{\nu}\tilde{s})P$ and $s :: \tilde{h}$ only appear at runtime. We often omit trailing $\mathbf{0}$ and write $s!$ and $s?.P$, omitting the arguments if unnecessary. Informally speaking, if we map our syntax to TCP, each queue corresponds to the TCP FIFO channel. Then a shared name correspond to a pair of an IP and a port name to initiate the session, while each session name is mapped to a pair of freshly generated IP and a port name which connects to the pair of the other side.

Binders are $\tilde{s}$ in $\overline{a}_{[2..\mathtt{n}]}(\tilde{s}).P$, $a_{[\mathtt{p}]}(\tilde{s}).P$ and $s?((\tilde{s}\,)); P$, $\tilde{x}$ in $s?(\tilde{x}); P$, $\tilde{x}\tilde{s}$ in $X(\tilde{x}\tilde{s}) = P$, n in $(\boldsymbol{\nu}\, n)P$ and process variables in $\mathsf{def}\ D\ \mathsf{in}\ P$. The notions of bound and free identifiers, channels, alpha equivalence $\equiv_\alpha$ and substitution are standard. The functions $\mathsf{fpv}(P)$ and $\mathsf{fn}(P)$ denote the sets of *free process variables* and *free identifiers*, respectively. Function $\mathsf{dpv}(\{X_i(\tilde{x}_i\tilde{s}_i) = P_i\}_{i\in I})$ denotes the set of *process variables* $\{X_i\}_{i\in I}$ introduced in $\{X_i(\tilde{x}_i\tilde{s}_i) = P_i\}_{i\in I}$. The notation $\Pi_i P_i$ denotes the parallel composition of zero or more processes P_i.

Structural congruence $\equiv$ over processes is the smallest congruence relation on processes that includes the equations given in Figure 2. These are standard except that we allow a vector of session channels in hiding, which is convenient for some proofs in the typing system (no substantial difference arises regarding the nature of the calculus by hiding channels one by one).

DEFINITION 2.1 (PROGRAM PHRASE AND PROGRAM). A process P is a *program phrase* if P has no queues and no ν-bound session channels (up to $\equiv$). P is a *program* (up to $\equiv$) if P is a program phrase in which no free session channels and process variables occur.

In the examples in §2.3, processes such as Buyer1, Buyer2, Seller, Kernel, DataProducer and Consumer are programs, hence they are also program phrases.

2.2. Operational Semantics

The operational semantics is given by the *reduction relation*, denoted $P \to Q$, which is the smallest relation on processes generated by the rules in Figure 3. In the figure, $e \downarrow v$ says that expression e evaluates to values v, but we leave its formal definition unspecified. We now explain each rule.

Rule [LINK] describes a session initiation among $\mathtt{n}$-parties through n-party synchronisation, generating m fresh session channels and the associated m empty queues ($\emptyset$ denotes the empty string). Each fresh channel is given a new empty queue. As a result $\mathtt{n}$ participants now share the newly generated m channels, hence their queues. Note the number of participants ($\mathtt{n}$) can be different from that of session channels (m), giving

$$\overline{a}_{[2..n]}(\tilde{s}).P_1 \mid a_{[2]}(\tilde{s}).P_2 \mid \cdots \mid a_{[n]}(\tilde{s}).P_n \;\to\; (\boldsymbol{\nu}\,\tilde{s})(P_1 \mid P_2 \mid ... \mid P_n \mid s_1 :: \emptyset \mid \cdots \mid s_m :: \emptyset) \qquad \text{[LINK]}$$

$$s!\langle \tilde{e} \rangle; P \mid s :: \tilde{h} \;\to\; P \mid s :: \tilde{h} \cdot \tilde{v} \quad (\tilde{e} \downarrow \tilde{v}) \qquad \text{[SEND]}$$

$$s!\langle\!\langle \tilde{t} \rangle\!\rangle; P \mid s :: \tilde{h} \;\to\; P \mid s :: \tilde{h} \cdot \tilde{t} \qquad \text{[DELEG]}$$

$$s \lhd l; P \mid s :: \tilde{h} \;\to\; P \mid s :: \tilde{h} \cdot l \qquad \text{[LABEL]}$$

$$s?(\tilde{x}); P \mid s :: \tilde{v} \cdot \tilde{h} \;\to\; P[\tilde{v}/\tilde{x}] \mid s :: \tilde{h} \qquad \text{[RECV]}$$

$$s?(\!(\tilde{t})\!); P \mid s :: \tilde{t} \cdot \tilde{h} \;\to\; P \mid s :: \tilde{h} \qquad \text{[SREC]}$$

$$s \rhd \{l_i \colon P_i\}_{i \in I} \mid s :: l_j \cdot \tilde{h} \;\to\; P_j \mid s :: \tilde{h} \qquad (j \in I) \qquad \text{[BRANCH]}$$

$$\mathsf{if}\ e\ \mathsf{then}\ P\ \mathsf{else}\ Q \;\to\; P \qquad (e \downarrow \mathsf{true}) \qquad \text{[IFT]}$$

$$\mathsf{if}\ e\ \mathsf{then}\ P\ \mathsf{else}\ Q \;\to\; Q \qquad (e \downarrow \mathsf{false}) \qquad \text{[IFF]}$$

$$\mathsf{def}\ D\ \mathsf{in}\ (X\langle \tilde{e}\tilde{s} \rangle \mid Q) \;\to\; \mathsf{def}\ D\ \mathsf{in}\ (P[\tilde{v}/\tilde{x}] \mid Q) \qquad (\tilde{e} \downarrow \tilde{v}, X(\tilde{x}\tilde{s}) = P \in D) \qquad \text{[DEF]}$$

$$P \;\to\; P' \;\Rightarrow\; (\boldsymbol{\nu}\, n)P \;\to\; (\boldsymbol{\nu}\, n)P' \qquad \text{[SCOP]}$$

$$P \;\to\; P' \;\Rightarrow\; P \mid Q \;\to\; P' \mid Q \qquad \text{[PAR]}$$

$$P \;\to\; P' \;\Rightarrow\; \mathsf{def}\ D\ \mathsf{in}\ P \;\to\; \mathsf{def}\ D\ \mathsf{in}\ P' \qquad \text{[DEFIN]}$$

$$P \equiv P' \text{ and } P' \;\to\; Q' \text{ and } Q' \equiv Q \;\Rightarrow\; P \;\to\; Q \qquad \text{[STR]}$$

Fig. 3. Reduction

flexibility and maintaining linearity in channel usage. The use of the n-party synchronisation in this rule captures, albeit abstractly, an n-party handshake which would be necessary for establishing an n-party link in real-world protocols. For example, we can create an arbitrary number of queues which can be dequeued and enqueued by all parties in that session.

Rules [SEND], [DELEG] and [LABEL] respectively enqueue values, channels and a label at the tail of the queue for s. In rule [SEND], $\tilde{e} \downarrow \tilde{v}$ evaluates each expression e_i to a value v_i and its definition is left unspecified. Symmetrically, rules [RECV], [SREC] and [BRANCH] dequeue, from the head of the queue, values, session channels and a label, respectively. Rules [RECV] further instantiates the received value in the continuation P, while rule [BRANCH] selects, from its continuation, the branch corresponding to the received label. The reduction rules [DELEG] and [SREC] are often called (session) *delegation* or *higher-order session passing*.

In these communication rules, sending and receiving are mediated by a queue: only when a message sent by (say) Alice is received by (say) Bob through a queue, we can say that an interaction between Alice and Bob has taken place. Since [LINK] generates a queue for each channel, these rules entail that:

(1) A sending action is never blocked (communication asynchrony); and that
(2) two messages from the same sender to the same channel arrive in the sending order (message order preservation).

As we discussed in §1, these are among the main features of the well-known transport mechanisms TCP, and the message queue is introduced for modelling these transports.

All other rules are standard: for reference, we briefly describe them. The two rules for conditional, [IFT] and [IFF], reduce to one of the branches, depending on the evaluation of the guard. Rule [DEF] performs unfolding of recursion. Rules [SCOP], [PAR] and [DEFIN] close the reduction relation under hiding, parallel composition and definition, respectively. Finally, rule [STR] says that the reduction relation is defined over processes up to $\equiv$.

REMARK 2.2. The rule for delegation [SREC], originally introduced in [Honda et al. 1998] for the π-calculus with sessions, uses the same session name $\tilde{t}$ without substitution for a simpler presentation. However, having [SREC] with substitution as in [RECV] breaks the subject reduction theorem and requires either two endpoint channels or bidirectional buffers [Gay and Hole 2005; Gay and Vasconcelos 2009; Yoshida and Vasconcelos 2007]. The reader can find a more detained explanation in [Yoshida and Vasconcelos 2007]. Roughly speaking, a substitution creates a self-delegation where the receiver gets his own session by which the shape of the session type is changed and the subject reduction is broken. Hence we require additional queues and restrictions on the form of the communication. The technical development of this work does not depend on this choice, see also § 6.2.

2.3. Examples

We now report two examples that have been used for discussion within the W3C WS-CDL working group [WS-CDL 2003]. Further large examples and applications of multiparty session types are listed in [Honda et al. 2008b].

EXAMPLE 2.3 (TWO BUYER PROTOCOL). We describe the two-buyers-protocol from the Introduction first by a sequence diagram, then by processes.

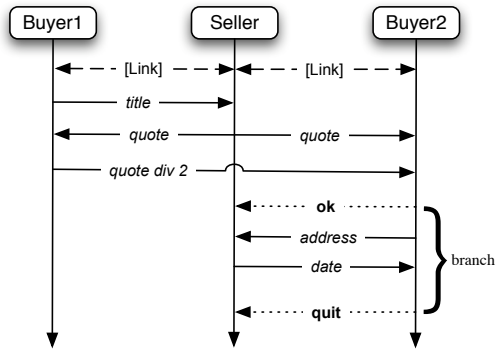

First Buyer1 sends a book title to Seller; then, Seller sends back a quote to Buyer1 and Buyer2; Buyer1 tells Buyer2 how much she is willing to contribute; and, finally, Buyer2 notifies Seller whether it accepts the quote or not. We can describe the behaviour of Buyer1 as with the following process:

$$\begin{aligned}\text{Buyer1} \stackrel{\text{def}}{=} \overline{a}_{[2,\,3]}\,(b_1, b_2, b_2', s).\;\; & s!\langle\text{“War and Peace”}\rangle;\\ & b_1?(quote);\;\; b_2'!\langle quote \text{ div } 2\rangle; P_1\end{aligned}$$

Channel b_1 is for Buyer1 to receive messages: b_2 and b_2' for Buyer2 and s for Seller (we discuss soon why Buyer2 needs two receiving channels). Buyer1 above is willing to contribute to half of the quote. In P_1, Buyer1 may perform the remaining transactions

with Seller and Buyer2. The remaining participants follow.

$$\begin{aligned}\mathbf{Buyer2} \stackrel{\text{def}}{=}\ & a_{[2]}(b_1, b_2, b_2', s).\ b_2?(quote);\ b_2'?(contrib);\\ & \textsf{if}\ (quote - contrib \leq 99)\\ & \quad \textsf{then}\quad s \lhd \mathsf{ok}; s!\langle address\rangle; b_2?(x); P_2\\ & \quad \textsf{else}\quad s \lhd \mathsf{quit}; \mathbf{0}\\ \mathbf{Seller} \stackrel{\text{def}}{=}\ & a_{[3]}(b_1, b_2, b_2', s).\ s?(title);\quad b_1, b_2!\langle quote\rangle;\\ & s \rhd \{\mathsf{ok}\colon s?(x); b_2!\langle date\rangle; Q\,,\quad \mathsf{quit}\colon \mathbf{0}\}\end{aligned}$$

Above $s_1..s_m!\langle v\rangle; P$ stands for $s_1!\langle v\rangle; ..s_m!\langle v\rangle; P$, assuming $s_1..s_m$ are pairwise distinct.[1] We can now explain why Buyer2 needs to use two input channels, b_2 and b_2'. The first input (for *quote*) is from Seller, while the second one (for *contrib*) is from Buyer1. Hence there is no guarantee that they arrive in a fixed order, as can be easily seen by analysing reduction paths (this is Lamport's principle [Lamport 1978]). Thus if we were to use b_2 for both actions, the two messages can be confused, losing linear usage of a channel. The problem becomes visible after the fifth step of the following reduction. If b_2 and b_2' were the same then the contribution of the Buyer1 could be queued before the price of the book and therefore received before at Buyer2. In § 4, we use our type discipline to detect this kind of error.

We now show an example of reductions. Let us define:

$$\begin{aligned}P \triangleq\ & \textsf{if}\ (quote - contrib \leq 99)\\ & \quad \textsf{then}\ \ s \lhd \mathsf{ok}; s!\langle address\rangle; b_2?(x); P_2\\ & \quad \textsf{else}\ \ s \lhd \mathsf{quit}; \mathbf{0}\end{aligned}\qquad S \triangleq s \rhd \{\mathsf{ok}\colon s?(x); b_2!\langle date\rangle; Q\,,\ \ \mathsf{quit}\colon \mathbf{0}\}$$

Below, a tag denotes the name of the rule from Figure 3 we apply. For simplicity, we omit [PAR] and [SCOP] after the second reduction.

$$\begin{aligned}& \qquad\qquad \mathbf{Buyer1}\ \mid\ \mathbf{Buyer2}\ \mid\ \mathbf{Seller}\\ \rightarrow\ & \stackrel{[\text{LINK}]}{} (\boldsymbol{\nu}\, b_1, b_2, b_2', s)(\ s!\langle\text{“War and Peace”}\rangle; b_1?(quote); b_2'!\langle quote\ \mathsf{div}\ 2\rangle; P_1\\ & \qquad\qquad\qquad\qquad \mid\ b_2?(quote); b_2'?(contrib); P\\ & \qquad\qquad\qquad\qquad \mid\ s?(title); b_1, b_2!\langle quote\rangle; S\\ & \qquad\qquad\qquad\qquad \mid\ b_1 :: \emptyset \mid b_2 :: \emptyset \mid b_2' :: \emptyset \mid s :: \emptyset)\\ \rightarrow\ & \stackrel{[\text{SEND}],[\text{PAR}],[\text{SCOP}]}{} (\boldsymbol{\nu}\, b_1, b_2, b_2', s)(\ b_2'!\langle quote\ \mathsf{div}\ 2\rangle; P_1\\ & \qquad\qquad\qquad\qquad \mid\ b_2?(quote); b_2'?(contrib); P\\ & \qquad\qquad\qquad\qquad \mid\ s?(title); b_1, b_2!\langle quote\rangle; S\\ & \qquad\qquad\qquad\qquad \mid\ b_1 :: \emptyset \mid b_2 :: \emptyset \mid b_2' :: \emptyset \mid s :: \text{“War and Peace”})\end{aligned}$$

[1] Due to asynchrony there is in effect no order among the sending actions at $s_1..s_m$.

$$\begin{array}{lll}
\rightarrow^{[\text{RCV}]} & (\boldsymbol{\nu}\, b_1, b_2, b_2', s)(& b_1?(quote); b_2'!\langle quote \text{ div } 2\rangle; P_1 \\
& & \mid b_2?(quote); b_2'?(contrib); P \\
& & \mid b_1, b_2!\langle quote\rangle; S \\
& & \mid b_1 :: \emptyset \mid b_2 :: \emptyset \mid b_2' :: \emptyset \mid s :: \emptyset) \\
\rightarrow^{[\text{SEND}]} & (\boldsymbol{\nu}\, b_1, b_2, b_2', s)(& b_1?(quote); b_2'!\langle quote \text{ div } 2\rangle; P_1 \\
& & \mid b_2?(quote); b_2'?(contrib); P \\
& & \mid b_2!\langle quote\rangle; S \\
& & \mid b_1 :: quote \mid b_2 :: \emptyset \mid b_2' :: \emptyset \mid s :: \emptyset) \\
\rightarrow^{[\text{RCV}]} & (\boldsymbol{\nu}\, b_1, b_2, b_2', s)(& b_2'!\langle quote \text{ div } 2\rangle; P_1 \\
& & \mid b_2?(quote); b_2'?(contrib); P \\
& & \mid b_2!\langle quote\rangle; S \\
& & \mid b_1 :: \emptyset \mid b_2 :: \emptyset \mid b_2' :: \emptyset \mid s :: \emptyset)
\end{array}$$

...

EXAMPLE 2.4 (STREAMING PROTOCOL). We next consider a simple protocol for the standard stream cipher [Schneier 1993].

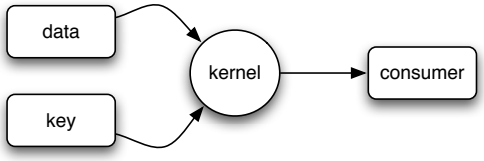

Data Producer and Key Producer continuously send a data stream and a key stream respectively to Kernel. Kernel calculates their XOR and sends the result to Consumer.

Assuming streams are sent block by block (say as large arrays), we can realise this protocol as communicating processes. We focus only on communication behaviour. The kernel initiates a session:

$$\begin{array}{rcl}
\text{Kernel} & \stackrel{\text{def}}{=} & \mathsf{def}\ \mathbf{K}(d,k,c) = d?(x);\ k?(y);\ c!\langle x \text{ xor } y\rangle;\ \mathbf{K}\langle d,k,c\rangle \\
& & \mathsf{in}\ \ \overline{a}_{[2,3,4]}\,(d,k,c).\mathbf{K}\langle d,k,c\rangle
\end{array}$$

The channels d and k are used for Kernel to receive data and keys from Data Producer and Key Producer, respectively, while c is used for Consumer to receive the encrypted data from Kernel. Data Producer and Consumer can be given as:[2]

$$\begin{array}{rcl}
\text{DataProducer} & \stackrel{\text{def}}{=} & \mathsf{def}\ \mathbf{P}(d,k,c) = d!\langle data\rangle; \mathbf{P}\langle d,k,c\rangle\ \mathsf{in}\ a_{[2]}\,(d,k,c).\mathbf{P}\langle d,k,c\rangle \\
\text{Consumer} & \stackrel{\text{def}}{=} & \mathsf{def}\ \mathbf{C}(d,k,c) = c?(data); \mathbf{C}\langle d,k,c\rangle\ \mathsf{in}\ a_{[3]}\,(d,k,c).\mathbf{C}\langle d,k,c\rangle
\end{array}$$

Key Producer is identical to Data Producer except it outputs at k instead of d. When three processes are composed, we can verify that messages are always consumed in the order they are produced, an essential requirement for correctness of the protocol (although processes repeatedly send and receive data using the same channel). This is because each channel is used exactly by one sender. We show how this argument can be cleanly represented and validated through session types in the next two sections.

[2]For simplicity our description lets both Data Producer and Consumer repeatedly send the same data: practically this is not the case but this simplified form is enough for our current concern, i.e. validation of communication behaviour.

$$
\begin{array}{lrll}
\text{Global} & G & ::= & \mathtt{p} \to \mathtt{p}' : k\,\langle U \rangle.\, G' & \text{values} \\
 & & \mid & \mathtt{p} \to \mathtt{p}' : k\,\{l_j : G_j\}_{j \in J} & \text{branching} \\
 & & \mid & G \mid G' & \text{parallel} \\
 & & \mid & \mu \mathbf{t}.G & \text{recursive} \\
 & & \mid & \mathbf{t} & \text{variable} \\
 & & \mid & \mathsf{end} & \text{end} \\
\text{Value} & U & ::= & \tilde{S} & \text{sorts} \\
 & & \mid & T@\mathtt{p} & \text{located session} \\
\text{Sort} & S & ::= & \mathsf{bool} \mid \mathsf{nat} \mid \ldots \mid \langle G \rangle &
\end{array}
$$

Fig. 4. Syntax of Global Types

3. GLOBAL TYPES AND CAUSAL ANALYSIS

Developing programs for multiparty sessions demands a clear formal design, since we need to program global interactions where multiple participants communicate and synchronise with each other. Programming individual participants without such a design and hope they somehow realise a meaningful and error-free conversation is hardly practical, especially when the implementation is done by a team of several programmers. In binary session types, the type for an endpoint also served as the description of the whole conversation between two parties, but this is no longer possible for multiparty sessions. This is why we require the type abstraction to describe global conversation scenarios of multiparty sessions: the global types introduced in this section extend binary session types to be able to directly express dependencies between communications among multiple peers.

3.1. Syntax of Global Types

A global type abstracts global multiparty conversations as a type signature. It takes a similar form to cryptography protocols where a message exchange from participant $\mathtt{p}$ to participant $\mathtt{p}'$ is specified as $\mathtt{p} \to \mathtt{p}'$. For example, the protocol "`Alice` sends a message with type nat to `Bob` via channel k, then terminates the interaction" is simply described as $\mathtt{Alice} \to \mathtt{Bob}\colon k\,\langle \mathsf{nat} \rangle.\ \mathsf{end}$. *Unlike the standard types of process calculi, the syntax no longer describes the input and output types separately: the information exchange between two parties is directly abstracted as one interaction.*

The full syntax of *global session types*, or *global types*, denoted by $G, G', \ldots$, is given in Figure 4. In a global type, we refer to session channels with a number, denoted by $k, k', \ldots$, which corresponds to the index of a vector of session channels: if we want to refer to the k-th session channel s_k of $s_1..s_n$ (such a vector is created by a session initiation), we write k in the global type. By writing number k (like de Bruijn notation), instead of channel s_k, we avoid including binding in the syntax of global types. We call k a *session channel index*.

$U, U', \ldots$ range over *value types*, denoting types for message values. Each value type is either a vector of *sorts* or a *located type*. Sorts, written $S, S', \ldots$, are types for shared names, where $\langle G \rangle$ means communicating a shared name typed by $\langle G \rangle$. A located type $T@\mathtt{p}$ denotes the communication (delegation) of a session channel of type T (called *endpoint type*) with role $\mathtt{p}$. Both of these types (T and S) are discussed in detail in §4.2. For understanding this section, it suffices to assume U as a single base type, i.e., only nat or bool. We often write $\mathtt{p} \to \mathtt{p}' : k.\ G'$ for $\mathtt{p} \to \mathtt{p}' : k\,\langle \rangle.\ G'$, i.e., U is empty.

We now give a detailed description of each term in Figure 4.

Type $\mathtt{p} \to \mathtt{p}' : k\,\langle U \rangle.\ G'$ says that an interaction between two participants over a session channel with index k must take place. In an implementation of such a behaviour,

given a vector of session channels $\tilde{s}$, participant $\mathtt{p}$ would send some message of type U to participant $\mathtt{p}'$ over s_k and then the session would continue according to G'. The type U is called *carried type*. Note that the operator "." captures sequentiality. As an example, the global type

$$1 \to 3\colon k\,\langle \mathsf{int}\rangle.\ 3 \to 2\colon k'\,\langle \mathsf{bool}\rangle.\ \mathsf{end} \tag{3}$$

describes a protocol where, given a vector $\tilde{s}$, participant 1 sends an integer to participant 3 over session channel s_k and then, 3 sends a boolean to participant 2 over $s_{k'}$. This protocol can be written in the model introduced in the previous section as follows:

$$\underbrace{s_k!\langle 5\rangle;\mathbf{0}}_{1} \mid \underbrace{s_{k'}?(y);\mathbf{0}}_{2} \mid \underbrace{s_k?(x);s_{k'}!\langle x=5\rangle;\mathbf{0}}_{3}$$

According to the semantics given in Fig. 3, the processes above (where, for the sake of clarity, we have labelled each process with a number corresponding to a participant) will execute according to the specification given by the global type in (3). Obviously, the same channel can be used several times as in the following type:

$$1 \to 3\colon k\,\langle \mathsf{int}\rangle.\ 3 \to 2\colon k'\,\langle \mathsf{bool}\rangle.\ 2 \to 1\colon k\,\langle \mathsf{bool}\rangle. \tag{4}$$

A possible implementation respecting such a protocol is:

$$\underbrace{s_k!\langle 5\rangle;s_k?(z);\mathbf{0}}_{1} \mid \underbrace{s_{k'}?(y);s_k!\langle \mathsf{true}\rangle;\mathbf{0}}_{2} \mid \underbrace{s_k?(x);s_{k'}!\langle x=5\rangle;\mathbf{0}}_{3}$$

On the other hand, the following process would *not* satisfy the specification in (4):

$$\underbrace{s_k!\langle 5\rangle;s_k?(z);\mathbf{0}}_{1} \mid \underbrace{s_k!\langle \mathsf{true}\rangle;s_{k'}?(y);\mathbf{0}}_{2} \mid \underbrace{s_k?(x);s_{k'}!\langle x=5\rangle;\mathbf{0}}_{3}$$

Unfortunately, due to asynchrony, it is possible that participant 3 receives a boolean while participant 1 receives, later on, an integer causing a run-time error.

Type $\mathtt{p} \to \mathtt{p}'\colon k\,\{l_j\colon G_j\}_{j\in J}$ denotes branching of a session. Intuitively, participant $\mathtt{p}$ must send one of the labels in $\{l_j \mid j \in I\}$ on channel s_k to participant $\mathtt{p}'$. When l_i is sent, interactions described in G_i will take place. For example, the global type

$$1 \to 3\colon k\,\{\quad \mathsf{five}\colon 3 \to 2\colon k'\,\langle \mathsf{bool}\rangle.\ \mathsf{end},\quad \mathsf{notfive}\colon 3 \to 2\colon k'\,\langle \mathsf{bool}\rangle.\ \mathsf{end}\quad \}$$

could be implemented by the process

$$\underbrace{\mathsf{if}\ e\ \mathsf{then}\ \ s_k \lhd \mathsf{five}\ \ \mathsf{else}\ \ s_k \lhd \mathsf{notfive}}_{1} \mid \underbrace{s_{k'}?(y)}_{2} \mid \underbrace{s_{k'} \rhd \left\{ \begin{array}{rl} \mathsf{five}\colon & s_{k'}!\langle \mathsf{true}\rangle, \\ \mathsf{notfive}\colon & s_{k'}!\langle \mathsf{false}\rangle \end{array} \right\}}_{3}$$

Type $G \mid G'$ specifies the concurrent execution of the interactions in G and G'.

Type $\mu\mathbf{t}.G$ is a recursive type for recurring conversation structures, assuming type variables $(\mathbf{t}, \mathbf{t}', \dots)$ are guarded in the standard way, i.e. type variables only appear under the prefixes (hence contractive). We take an *equi-recursive* view, not distinguishing between $\mu\mathbf{t}.G$ and its unfolding $G[\mu\mathbf{t}.G/\mathbf{t}]$ [Pierce 2002]. We assume that $\langle G\rangle$ in the grammar of sorts is closed, i.e. without type variables.[3]

Type end represents the termination of the session and is often omitted. We identify "$G \mid \mathsf{end}$" and "$\mathsf{end} \mid G$" with G.

We conclude this subsection by giving the following definition:

[3]In the presence of the standard recursive sorts [Honda et al. 1998], which we omit for simpler presentation, we allow sort variables to occur in $\langle G\rangle$.

DEFINITION 3.1 (ACTION). We say that $\mathtt{p} \to \mathtt{p}' : k$ in $\mathtt{p} \to \mathtt{p}' \colon k\,\langle U \rangle.\ G'$ or $\mathtt{p} \to \mathtt{p}' \colon k\,\{l_j \colon G_j\}_{j \in J}$ is an *action from* $\mathtt{p}$ *to* $\mathtt{p}'$ *at* k.

3.2. Operational Semantics for Global Types

This subsection defines semantics of global types, introducing the labelled transition relation (LTS). The LTS is useful not only to give a clear justification for causal dependencies of global types defined in the next subsection, but also to prove the main theorems for the typing system later.

DEFINITION 3.2 (GLOBAL TYPE LABELLED TRANSITION RELATION). The syntax of labels ($\ell, \ell', ...$) of global types is defined as follows:

$$\ell \ ::= \ \mathtt{p} \to \mathtt{p}' \colon k\,\langle U \rangle \ \ \mid \ \ \mathtt{p} \to \mathtt{p}' \colon k\,\langle l \rangle$$

A label ℓ denotes a communication over a channel k of some type U or label l. Then the transition relation $G \xrightarrow{\ell} G'$ is defined by the following rules:

$$[\mathrm{GR1}] \quad \mathtt{p} \to \mathtt{q} \colon k\,\langle U \rangle.\ G \xrightarrow{\mathtt{p} \to \mathtt{q} \colon k\,\langle U \rangle} G \qquad [\mathrm{GR2}] \quad \mathtt{p} \to \mathtt{q} \colon k\,\{l_i \colon G_i\}_{i \in I} \xrightarrow{\mathtt{p} \to \mathtt{q} \colon k\,\langle l_j \rangle} G_j$$

$$[\mathrm{GR3}] \ \frac{G_1 \xrightarrow{\ell} G_2 \qquad \mathtt{q} \notin \ell}{\mathtt{p} \to \mathtt{q} \colon k\,\langle U \rangle.\ G_1 \xrightarrow{\ell} \mathtt{p} \to \mathtt{q} \colon k\,\langle U \rangle.\ G_2} \qquad [\mathrm{GR4}] \ \frac{\forall i \in I.\ G_i \xrightarrow{\ell} G_i' \quad \mathtt{q} \notin \ell}{\mathtt{p} \to \mathtt{q} \colon k\,\{l_i \colon G_i\}_{i \in I} \xrightarrow{\ell} \mathtt{p} \to \mathtt{q} \colon k\,\{l_i \colon G_i'\}_{i \in I}}$$

$$[\mathrm{GR5}] \ \frac{G_1 \xrightarrow{\ell} G_1'}{G_1 \mid G_2 \xrightarrow{\ell} G_1' \mid G_2} \qquad [\mathrm{GR6}] \ \frac{G_2 \xrightarrow{\ell} G_2'}{G_1 \mid G_2 \xrightarrow{\ell} G_1 \mid G_2'}$$

The rules allow to permute the order of two actions which are causally unrelated. This is defined by the condition $\mathtt{q} \notin \ell$ in [GR3,4]. Note that in [GR4], we require that each branch must be able to perform action ℓ.

As a simple example, consider $G = \mathtt{1} \to \mathtt{2} \colon k\,\langle \mathsf{int} \rangle.\ \mathtt{3} \to \mathtt{4} \colon k'\,\langle \mathsf{bool} \rangle.\ \mathsf{end}$ and let $\ell_1 = \mathtt{1} \to \mathtt{2} \colon k\,\langle \mathsf{int} \rangle$ and $\ell_2 = \mathtt{2} \to \mathtt{3} \colon k'\,\langle \mathsf{bool} \rangle$. Since the participants are pairwise distinct, we can perform the second action first. Hence, using [GR1] and [GR3] above, we have two possible transition relations from G as follows:

$$G \xrightarrow{\ell_1} \mathtt{3} \to \mathtt{4} \colon k'\,\langle \mathsf{bool} \rangle.\ \mathsf{end} \xrightarrow{\ell_2} \mathsf{end} \ \text{ and } \ G \xrightarrow{\ell_2} \mathtt{1} \to \mathtt{2} \colon k\,\langle \mathsf{int} \rangle.\ \mathsf{end} \xrightarrow{\ell_1} \mathsf{end}$$

Another interesting example is: $\mathtt{1} \to \mathtt{2} \colon k\,\langle \mathsf{int} \rangle.\ \mathtt{3} \to \mathtt{1} \colon k'\,\langle \mathsf{bool} \rangle.\ \mathsf{end}$. This global type means that the participant 1 is allowed to *receive* the message from the participant 3 before the message from 1 is received by 2 since they are delivered to the two different channels (i.e. queues). Thus with $\ell_3 = \mathtt{2} \to \mathtt{1} \colon k'\,\langle \mathsf{bool} \rangle$, we have:

$$G' \xrightarrow{\ell_1} \mathtt{3} \to \mathtt{1} \colon k'\,\langle \mathsf{bool} \rangle.\ \mathsf{end} \xrightarrow{\ell_3} \mathsf{end} \ \text{ and } \ G' \xrightarrow{\ell_3} \mathtt{1} \to \mathtt{2} \colon k\,\langle \mathsf{int} \rangle.\ \mathsf{end} \xrightarrow{\ell_1} \mathsf{end}$$

On the other hand, $G'' = \mathtt{3} \to \mathtt{1} \colon k'\,\langle \mathsf{bool} \rangle.\ \mathtt{1} \to \mathtt{2} \colon k\,\langle \mathsf{int} \rangle.\ \mathsf{end}$ has only one possible transition since the two inputs at the receiver $\mathtt{q}$ are ordered. Hence we only have the following one transition from G''.

$$G'' \xrightarrow{\ell_3} \mathtt{1} \to \mathtt{2} \colon k'\,\langle \mathsf{bool} \rangle.\ \mathsf{end} \xrightarrow{\ell_1} \mathsf{end}$$

This means the message from 1 is surely received at 2 *after* 1 received the message from 3 hence two actions are not permutable. The semantics of the permutation will be clearer when we introduce the causality relation between the actions in §3.5.

3.3. Action Ordering

Henceforth, we refer to the acyclic directed graph of a global type G as a standard regular tree representation [Pierce 2002]. In order to give a definition, we annotate the actions in $\mathtt{p} \to \mathtt{p}' : k\,\langle U \rangle.\, G'$ and $\mathtt{p} \to \mathtt{p}' : k\,\{l_j : G_j\}_{j \in J}$ by a node name $\mathtt{n}$.

DEFINITION 3.3 (REGULAR TREE REPRESENTATION). The *regular tree representation* $\mathsf{tree}(G)$ of a global type G is defined over the annotated unfolding of G such that

$\mathsf{tree}(\mathtt{n} : \mathtt{p} \to \mathtt{p}' : k\,\langle U \rangle.\, G')$ has root $\mathtt{n}$ with an edge to the root of $\mathsf{tree}(G')$
$\mathsf{tree}(\mathtt{n} : \mathtt{p} \to \mathtt{p}' : k\,\{l_j : G_j\}_{j \in J})$ has root $\mathtt{n}$ with edges to the roots of each $\mathsf{tree}(G_j)$ $(j \in J)$
$\mathsf{tree}(G_1 \mid G_2)$ has root $\mid$ with edges to the roots of each $\mathsf{tree}(G_i)$ $(i = 1, 2)$

Each node in $\mathsf{tree}(G)$ is labelled by the occurrence of its corresponding action or it has no label in the case of parallel. These node names are unique in the unfolding.

As an example, the global type

$$\mu \mathbf{t}.\ 1 \to 2 : k \left\{ \begin{array}{l} l_1 : 2 \to 1 : k\,\langle \mathsf{int} \rangle.\ \mathsf{end} \mid 2 \to 3 : k'\,\langle \mathsf{bool} \rangle.\ \mathsf{end} \\ l_2 : 2 \to 1 : k\,\langle \mathsf{nat} \rangle.\ \mathbf{t} \end{array} \right\}$$

has the following (infinite) regular tree representation:

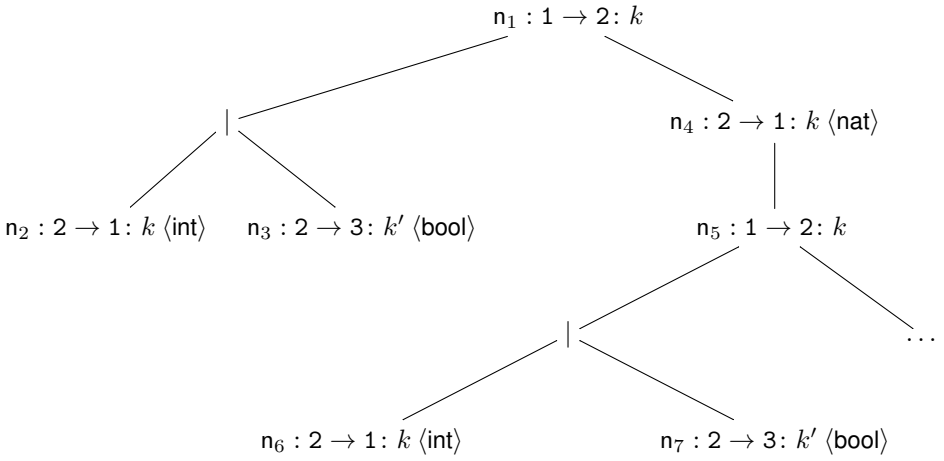

We now define:

DEFINITION 3.4. An action from $\mathtt{p}$ to $\mathtt{p}'$ at k is in a global type G, written $\mathtt{p} \to \mathtt{p}' : k \in G$, whenever, in the regular tree representation of G, there exists some node $\mathtt{n}$ with label $\mathtt{p} \to \mathtt{p}' : k$. We write $\mathtt{n} = \mathtt{p} \to \mathtt{p}' : k$ if $\mathtt{n}$ has label $\mathtt{p} \to \mathtt{p}' : k$.

DEFINITION 3.5. We denote

— $\mathsf{pid}(G)$ for the set of participants occurring in G (but not in any carried types).
— $\mathsf{sid}(G)$ for the number of the set of session channel indices in G (but not in any carried types).

For example, if $G = 1 \to 3 : k\,\langle \mathsf{int} \rangle.\ 3 \to 2 : k'\,\langle \mathsf{bool} \rangle.\ \mathsf{end}$, then $\mathsf{pid}(G) = \{1, 2, 3\}$ and $\mathsf{sid}(G) = 2$.

CONVENTION 3.6. We assume that in each action from $\mathtt{p}$ to $\mathtt{p}'$ we have $\mathtt{p} \neq \mathtt{p}'$, i.e. we prohibit reflexive interaction.

Below, we define the relation $\mathtt{n}_1 \prec \mathtt{n}_2 \in G$ which holds whenever $\mathtt{n}_1$ directly or indirectly occurs before $\mathtt{n}_2$ in the regular tree representation of G. For instance, in the global type $G = \mathtt{p}_1 \to \mathtt{p}_1' \colon k_1 \langle U_1 \rangle .\ \mathtt{p}_2 \to \mathtt{p}_2' \colon k_2 \langle U_2 \rangle .\ G_2'$ we have that $\mathtt{p}_1 \to \mathtt{p}_1' : k_1 \prec \mathtt{p}_2 \to \mathtt{p}_2' : k_2 \in G$.

DEFINITION 3.7 (ACTION ORDERING). We define $\prec$ as the least partial order such that:

(1)	$\mathtt{n}_1 \prec \mathtt{n}_2 \in \mathtt{p} \to \mathtt{p}' \colon k \langle U \rangle .\ G'$	if $\mathtt{n}_1 = \mathtt{p} \to \mathtt{p}' : k$ and $\mathtt{n}_2 \in G'$
(2)	$\mathtt{n}_1 \prec \mathtt{n}_2 \in \mathtt{p} \to \mathtt{p}' \colon k \{l_j \colon G_j\}_{j \in J}$	if $\mathtt{n}_1 = \mathtt{p} \to \mathtt{p}' : k$ and $\exists i \in J .\ \mathtt{n}_2 \in G_i$
(3)	$\mathtt{n}_1 \prec \mathtt{n}_2 \in \mathtt{p} \to \mathtt{p}' \colon k \langle U \rangle .\ G'$	if $\mathtt{n}_1 \prec \mathtt{n}_2 \in G'$
(4)	$\mathtt{n}_1 \prec \mathtt{n}_2 \in \mathtt{p}_1 \to \mathtt{p}_1' \colon k' \{l_j \colon G_j'\}_{j \in J}$	if $\mathtt{n}_1 \prec \mathtt{n}_2 \in G_i'$ for some $i \in J$
(5)	$\mathtt{n}_1 \prec \mathtt{n}_2 \in G_1 \mid G_2$	if $\mathtt{n}_1 \prec \mathtt{n}_2 \in G_i$ for some $i \in \{1, 2\}$

Above, (1,2) say that, in values and branching types, any nested action comes always after the top one for value and branch types. (3,4) say that if two actions are related by the action ordering in a subterm of some global type G then they are also related in G. Parallel composition and recursion are dealt with by (5,6).

The action ordering allows us to express intended causal dependencies in global types, which is subtle under asynchronous semantics. Consider the following simple global type:

$$G = \mathtt{A} \to \mathtt{B} \colon k \langle U \rangle .\ \ \mathtt{A} \to \mathtt{C} \colon k' \langle U' \rangle .\ \mathsf{end} \tag{5}$$

where A, B and C denote participants. We use this example to show an important difference between asynchronous communication and synchronous communication. In a "synchronous" interpretation of (5), the ordering would mean: "only after the first sending and receiving take place, the second sending and receiving take place". This is a suitable reading when sending and receiving constitute a single atomic action, as in synchronous languages, but *not* with asynchronous communication, where it is hard to impose such an ordering, since messages to distinct channels may not arrive in order e.g. C may receive the second output from A before its first message reaches B. This corresponds to [GR3] and [GR4] in Definition 3.2, where the action ℓ can be executed before the action in the prefix.

Thus, the present theory takes a more liberal interpretation of $\prec$, imposing sequencing *only on the actions of the same participant in ordered actions*. For example, in (5), A's two sending actions are ordered, but B's and C's receiving actions are not. This relation is explained in the next subsection with several examples.

3.4. Examples of Global Types

EXAMPLE 3.8 (TWO-BUYER PROTOCOL). The following is a global type of the two-buyer-protocol in §2.3. We write participants and channels with legible symbols though they are actually numbers (e.g., $\mathtt{Bi} = \mathtt{i}$, $\mathtt{S} = 3$, $b_1 = 1$, $b_2 = 2$, $b_2' = 3$ and $s = 4$):

1. $\mathtt{B1} \to \mathtt{S} :\ \ s\langle \mathsf{string} \rangle .$
2. $\mathtt{S} \to \mathtt{B1} :\ \ b_1\langle \mathsf{int} \rangle .$
3. $\mathtt{S} \to \mathtt{B2} :\ \ b_2\langle \mathsf{int} \rangle .$
4. $\mathtt{B1} \to \mathtt{B2} :\ b_2'\langle \mathsf{int} \rangle .$
5. $\mathtt{B2} \to \mathtt{S} :\ \ s \left\{ \begin{array}{l} \mathsf{ok} :\ \mathtt{B2} \to \mathtt{S} \colon s \langle \mathsf{string} \rangle .\ \mathtt{S} \to \mathtt{B2} \colon b_2 \langle \mathsf{date} \rangle .\ \mathsf{end}, \\ \mathsf{quit} :\ \mathsf{end} \end{array} \right\}$

The type gives a clear, abstract view of the whole conversation scenario. The following are several salient points in the asynchronous interpretation of this type:

— Consider Lines 3 and 4. Since they have different senders, the sending actions are unordered in spite of their $\prec$-ordering. *Hence if $b_2 = b_2'$ two messages have a conflict at s* (i.e. lose the ordering).
— Next, we consider the following causal chain from Line 1 to Line 3 to Line 5:

$$\texttt{B1} \;\rightarrow\; \texttt{S} \;\prec\; \texttt{S} \;\rightarrow\; \texttt{B2} \;\prec\; \texttt{B2} \;\rightarrow\; \texttt{S}$$

Above $\rightarrow$ can be interpreted as the ordering given by message delivery (see previous subsection), while $\prec$ is the action ordering. Note in particular two sending actions by B1 (Line 1) and by B2 (Line 5), both done at s, are causally ordered. By focussing on $\prec$ from the first S (of Line 1) to the last S (of Line 5), the receiving actions in Line 1 and the first $\texttt{B1} \rightarrow \texttt{S}$ in Line 5 are also ordered. Since the interaction in Line 1 will surely take place before the interaction in Line 5, no conflict occurs between these two communications in spite of their use of a common channel s.

EXAMPLE 3.9 (STREAMING PROTOCOL). We now present the global type of the simple streaming protocol in §2.3. Below we unfold its recursion once, and set: $d = 1$, $k = 2$, $c = 3$, $\texttt{K} = 1$, $\texttt{DP} = 2$, $\texttt{C} = 3$ and $\texttt{KP} = 4$.

$$\begin{array}{lll}
1. & \mu\mathbf{t}.\ \texttt{DP} \rightarrow \texttt{K}\colon d\,\langle \mathsf{bool}\rangle. & 4. \quad \texttt{DP} \rightarrow \texttt{K}\colon d\,\langle \mathsf{bool}\rangle. \\
2. & \quad\ \texttt{KP} \rightarrow \texttt{K}\colon k\,\langle \mathsf{bool}\rangle. & 5. \quad \texttt{KP} \rightarrow \texttt{K}\colon k\,\langle \mathsf{bool}\rangle. \\
3. & \quad\ \texttt{K} \rightarrow \texttt{C}\colon c\,\langle \mathsf{bool}\rangle. & 6. \quad \texttt{K} \rightarrow \texttt{C}\colon c\,\langle \mathsf{bool}\rangle.\mathbf{t}
\end{array}$$

The following arguments hold for any n-fold unfoldings.

— Lines 1 and 2 are temporally unordered in sending: but this does not cause conflict since channels d and k are distinct.
— Line 1 and its unfolding, Line 4, share d. But the two use the same sender and the same receiver, so each pair of actions are $\prec$-ordered, hence safe. Similarly for other unfolded actions.

EXAMPLE 3.10 (INSTRUMENT CONTROLLING). We now present another example from [OOI 2015] which focuses on the usage of an instrument through repeated commands, together with checking privileges initially and later reporting the status to the central operator in charge of the instrument. The global type description involves a user User, an operator Op and the instrument Instr and is given as follows.

$$\begin{array}{l}
\texttt{User} \rightarrow \texttt{Op}\colon 1\,\langle \mathsf{privilege}\rangle. \\
\texttt{Op} \rightarrow \texttt{User}\colon 2 \left\{ \begin{array}{l}
\mathsf{ok}:\ \mu\mathbf{t}. \\
\qquad \texttt{User} \rightarrow \texttt{Instr}\colon 3 \left\{ \begin{array}{l} \mathsf{move}:\ \mathbf{t} \\ \mathsf{photo}:\ \mathbf{t} \\ \mathsf{quit}:\ \texttt{Instr} \rightarrow \texttt{Op}\colon 4\,\langle \mathsf{string}\rangle.\ \mathsf{end} \end{array} \right\} \\
\\
\mathsf{no}:\ \mathsf{end}
\end{array} \right\}
\end{array}$$

Note that the protocol description given by the global type above can have several implementations. In particular, the instrument can be used with any combination of the operations move and photo. However, any sequence will be terminated by quit.

(II) Good
$\mathtt{A} \to \mathtt{B} : k$
$\mathtt{C} \to \mathtt{B} : k'$
$s_k! \mid (s_k?; s_{k'}?) \mid s_{k'}!$

(II) Bad
$\mathtt{A} \to \mathtt{B} : k$
$\mathtt{C} \to \mathtt{B} : k$
$s_k! \mid (s_k?; s_k?) \mid s_k!$

(IO) Good
$\mathtt{A} \to \mathtt{B} : k$
$\mathtt{B} \to \mathtt{C} : k'$
$s_k! \mid (s_k?; s_{k'}!) \mid s_{k'}?$

(IO) Bad
$\mathtt{A} \to \mathtt{B} : k$
$\mathtt{B} \to \mathtt{C} : k$
$s_k! \mid (s_k?; s_k!) \mid s_k?$

(OO, II) Good
$\mathtt{A} \to \mathtt{B} : k$
$\mathtt{A} \to \mathtt{B} : k$
$(s_k!; s_k!) \mid (s_k?; s_k?)$

(OI) Bad
$\mathtt{A} \to \mathtt{B} : k$
$\mathtt{C} \to \mathtt{A} : k$
$(s_k!; s_k?) \mid s_k? \mid s_k!$

(OO) Bad
$\mathtt{A} \to \mathtt{B} : k$
$\mathtt{A} \to \mathtt{C} : k$
$s_k? \mid (s_k!; s_k!) \mid s_k?$

Fig. 5. Causality Analysis

Below, we give a possible implementation of each participant:

$$\mathrm{User} \stackrel{\text{def}}{=} s_1!\langle high\rangle; s_2 \rhd \left\{ \begin{array}{l} \mathsf{ok}: \ s_3 \lhd \mathsf{move};\ s_3 \lhd \mathsf{photo};\ s_3 \lhd \mathsf{quit};\ \mathbf{0} \\ \mathsf{no}: \ \mathbf{0} \end{array} \right\}$$

$$\mathrm{Operator} \stackrel{\text{def}}{=} s_1?(x);\ \mathsf{if}\ f(x)\ \mathsf{then}\ \ s_2 \lhd \mathsf{ok};\ \mathbf{0}\ \ \mathsf{else}\ \ s_2 \lhd \mathsf{no};\ \mathbf{0}$$

$$\mathrm{Instrument} \stackrel{\text{def}}{=} \mu \mathbf{t}.\ s_3 \rhd \left\{ \begin{array}{l} \mathsf{move}: \ \mathbf{t} \\ \mathsf{photo}: \ \mathbf{t} \\ \mathsf{quit}: \ s_4!\langle report\rangle; \mathbf{0} \end{array} \right\}$$

3.5. A Safety Principle for Global Types: Linearity of Channels

For a conversation in a session to proceed properly, it is desirable that there is no conflict (racing) at session channels. The process $s_k!\langle \mathsf{true}\rangle \mid s_k!\langle 5\rangle \mid s_k?x;\ \mathsf{if}\ x\ \mathsf{then}\ P\ \mathsf{else}\ Q$ is a typical example of a race at channel s_k: if the second output synchronises with the first input we have a run-time error when evaluating the guard of the conditional. To ensure absence of such races, when a *common* channel is used in two communications, their sending actions and their receiving actions should (respectively) be ordered temporally (causality), so that no confusion arises at sending or receiving. If a global type satisfies this principle, then it specifies an ordering of interactions, and can be used as a basis of guaranteeing process behaviours through type checking. The correspondence between the linearity property and the LTS of the global types defined in Definition 3.2 will be used for proving the main theorems, Subject Reduction Theorem (Theorem 5.19) and Session Fidelity Theorem (Corollary 5.23).

Causality is induced in several ways in the present asynchronous model. We summarise all essential cases in Figure 5, with concrete process instances for illustration.

In the figure, IO indicates a causal ordering from input (receiving) to output (sending), similarly for II, OO and OI. In (II)-Bad, we demand $\mathtt{A} \neq \mathtt{C}$. We observe:

— The "good" and "bad" cases for II show that II alone is safe only when two channels differ. Similarly for IO.
— In OO,II, two outputs have the same sender and the same channel, so (by *message order-preservation*) outputs are ordered. Inputs are also ordered by $\prec$ hence they are safe.
— There is no ordering from output to input (due to asynchrony), so OI gives us no dependency.

These observations lead to the following causal relations on global types.

DEFINITION 3.11 (DEPENDENCY RELATIONS). Fix G. The relation $\prec_\phi$, with $\phi \in \{\mathsf{II}, \mathsf{IO}, \mathsf{OO}\}$, over actions is generated from:

$$\begin{array}{ll} \mathtt{n}_1 \prec_{\mathtt{II}} \mathtt{n}_2 & \text{if } \mathtt{n}_1 \prec \mathtt{n}_2 \in G \text{ and } \mathtt{n}_i = \mathtt{p}_\mathtt{i} \to \mathtt{p} : k_i \ (i = 1, 2) \\ \mathtt{n}_1 \prec_{\mathtt{IO}} \mathtt{n}_2 & \text{if } \mathtt{n}_1 \prec \mathtt{n}_2 \in G,\ \mathtt{n}_1 = \mathtt{p}_1 \to \mathtt{p} : k_1 \text{ and } \mathtt{n}_2 = \mathtt{p} \to \mathtt{p}_2 : k_2. \\ \mathtt{n}_1 \prec_{\mathtt{OO}} \mathtt{n}_2 & \text{if } \mathtt{n}_1 \prec \mathtt{n}_2 \in G,\ \mathtt{n}_i = \mathtt{p} \to \mathtt{p}_\mathtt{i} : k \ (i = 1, 2). \end{array}$$

— An *input dependency from* $\mathtt{n}_1$ *to* $\mathtt{n}_2$ is a chain of the form $\mathtt{n}_1 \prec_{\phi_1} \cdots \prec_{\phi_n} \mathtt{n}_2$ ($n \geq 0$) such that $\phi_i \in \{\mathsf{IO}\}$ for $1 \leq i \leq n-1$ and $\phi_n = \mathsf{II}$.
— An *output dependency from* $\mathtt{n}_1$ *to* $\mathtt{n}_2$ is a chain $\mathtt{n}_1 \prec_{\phi_1} \cdots \prec_{\phi_n} \mathtt{n}_2$ ($n \geq 1$) such that $\phi_i \in \{\mathsf{OO}, \mathsf{IO}\}$.

Note that, in the input dependency, the last II-ordering is necessary. In fact, if we allow the dependency to end with an IO-edge, an input at $\mathtt{n}_2$ is not checked. We can further clarify dependencies with the following graphical examples:

(a) (b)

In picture (a), there is an output dependency from the first to the third line which has been marked as OO, and an input dependency through all lines. However, in (b), we only have an output dependency. It is clear that, while (a) could be implemented in an asynchronous setting, the conversation in (b) would cause problems. In fact, the messages sent by A on k_1 could be delivered in the wrong order (first to C and then to B). The notion of linearity, hereby introduced, precisely captures such inconsistencies in global types.

DEFINITION 3.12 (LINEARITY). G is *linear* whenever, for all $\mathtt{n}_1 \prec \mathtt{n}_2 \in G$ such that $\mathtt{n}_i = \mathtt{p}_\mathtt{i} \to \mathtt{p}'_\mathtt{i} : k \ (i = 1, 2)$, both input and output dependencies from $\mathtt{n}_1$ to $\mathtt{n}_2$ exist. We inductively apply this constraint to all global types which G carries.

Observe that we do not require ordering between $\mathtt{n}_i \in G_k$ and $\mathtt{n}_j \in G_h$ in $\mathtt{p} \to \mathtt{p}' : k \, \{l_j : G_j\}_{j \in J}$ (for $h, k \in J$, $h \neq k$) since only one branch is performed. In fact, they cannot be related by $\prec$ according to Definition 3.7. We further clarify the condition on

branching with an example:

$$\mathtt{A} \to \mathtt{B} : t \left\{ \begin{array}{r} \mathsf{ok} : \ \mathtt{C} \to \mathtt{D} : s.\mathsf{end}, \\ \mathsf{quit} : \ \mathtt{C} \to \mathtt{D} : s.\mathsf{end} \end{array} \right\} \qquad \mathtt{A} \to \mathtt{B} : t. \left(\begin{array}{l} \mathtt{C} \to \mathtt{D} : s.\mathsf{end} \mid \\ \mathtt{C} \to \mathtt{D} : s.\mathsf{end} \end{array} \right)$$

(a) branching (b) parallel

The type (a) represents branching: since only one branch is selected, there is no conflict between the two actions $\mathtt{C} \to \mathtt{D} : s$. On the other hand, (b) denotes a concurrent execution of two independent $\mathtt{C} \to \mathtt{D} : s$, so an input conflict at $\mathtt{D}$ exists.

Linearity and its violation can be detected algorithmically, without infinite unfolding. First we observe we do need to unfold once.

$$\mu \mathbf{t}.(\mathtt{A} \to \mathtt{B} : s.\mathsf{end} \mid \ \mathtt{B} \to \mathtt{A} : t.\mathbf{t})$$

This is linear in its 0-th unfolding (i.e. we replace $\mathbf{t}$ with end): but when unfolded once, it becomes non-linear, as follows:

$$\mathtt{A} \to \mathtt{B} : s.\mathsf{end}, \ \mathtt{B} \to \mathtt{A} : t.\mu \mathbf{t}.(\mathtt{A} \to \mathtt{B} : s.\mathsf{end} \mid \ \mathtt{B} \to \mathtt{A} : t.\mathbf{t})$$

since the two actions $\mathtt{A} \to \mathtt{B} : s$ appear in parallel. This is witnessed by:

$$\mathsf{def}\ X(st) = ((s! \mid t?.s!.X\langle ts \rangle) \mid s?.t!)\ \mathsf{in}\ X\langle ts \rangle$$

where $(s! \mid t?.s!.X\langle ts \rangle)$ belongs to $\mathtt{A}$ and $s?.t!$ belongs to $\mathtt{B}$. Unfolding once is necessary also in global types that do not contain parallel global types. The example below shows a global type which satisfies the linearity condition:

$$\mu \mathbf{t}.\mathtt{A} \to \mathtt{B} : s.\mathtt{B} \to \mathtt{C} : s'.\mathtt{A} \to \mathtt{C} : s.\mathbf{t} \tag{6}$$

However, when unfolded once, it is no longer linear as:

$$\mathtt{A} \to \mathtt{B} : s.\mathtt{B} \to \mathtt{C} : s'.\mathtt{A} \to \mathtt{C} : s.\mu \mathbf{t}.\mathtt{A} \to \mathtt{B} : s.\mathtt{B} \to \mathtt{C} : s'.\mathtt{A} \to \mathtt{C} : s.\mathbf{t} \tag{7}$$

since there is no input and output dependencies between $\mathtt{A} \to \mathtt{C} : s$ and $\mathtt{A} \to \mathtt{B} : s$.

But in fact unfolding once turns out to be enough. Taking G as a syntax, let us call the *one-time unfolding of* G the result of unfolding once for each recursion in G (but never in carried types), and replacing the remaining variable with end. For example, the type in (6) would be first transformed into the type in (7) and finally become:

$$\mathtt{A} \to \mathtt{B} : s.\mathtt{B} \to \mathtt{C} : s'.\mathtt{A} \to \mathtt{C} : s.\mu \mathbf{t}.\mathtt{A} \to \mathtt{B} : s.\mathtt{B} \to \mathtt{C} : s'.\mathtt{A} \to \mathtt{C} : s.\mathsf{end}$$

PROPOSITION 3.13.
(1) *The one-time unfolding of a global type is linear if and only if its n-th unfolding is linear.*
(2) *The linearity of a global type is decidable.*

PROOF. For (1), the if-direction is obvious. The only if-direction is proved by induction on n. See Appendix A for the full proofs. (2) is an immediate corollary of (1). □

PROPOSITION 3.14. *Suppose G is linear and $G \xrightarrow{\ell} G'$. Then G' is linear.*

PROOF. By induction on the last LTS rule applied. The cases [GR1,GR2,GR5,GR6] are obvious. We prove the case [GR3]. The case [GR4] is similar. Suppose $G = \mathtt{p}_1 \to \mathtt{p}_2 \colon k\ \langle U \rangle$. $G_1 \xrightarrow{\ell} \mathtt{p}_1 \to \mathtt{p}_2 \colon k\ \langle U \rangle$. $G_2 = G'$ is derived by $G_1 \xrightarrow{\ell} G_2$ with $\mathtt{p}_2 \notin \ell$. Assume $\ell = \mathtt{q}_1 \to \mathtt{q}_2 \colon k'\ \langle U' \rangle$. Then we first prove if G satisfies the linearity condition, then $k \neq k'$. Suppose by contradiction, $k = k'$. Then there should be both output and input causalities from $\mathtt{n} = \mathtt{p}_1 \to \mathtt{p}_2 : k$ to $\mathtt{n}' = \mathtt{q}_1 \to \mathtt{q}_2 : k$. If there is the IO-causality from $\mathtt{n}$ to $\mathtt{n}'$ in G, then we cannot apply [GR3]. Hence there is only OO and II causalities from

$$
\begin{array}{llll}
\text{Value} & U & ::= & \tilde{S} \mid T@\mathtt{p} \\
\text{Sort} & S & ::= & \mathsf{bool} \mid \ldots \mid \langle G \rangle \\
\text{End-point} & T & ::= & k!\langle U \rangle;T & \text{send} \\
 & & \mid & k?\langle U \rangle;T & \text{receive} \\
 & & \mid & k \oplus \{l_i\colon T_i\}_{i\in I} & \text{selection} \\
 & & \mid & k\&\{l_i\colon T_i\}_{i\in I} & \text{branching} \\
 & & \mid & \mu \mathbf{t}.T \mid \mathbf{t} \mid \mathsf{end}
\end{array}
$$

Fig. 6. Syntax of End-point Session Types

$\mathtt{n}$ to $\mathtt{n}'$. In this case, we should have $\mathtt{p}_2 = \mathtt{q}_2$. This contradicts $\mathtt{p}_2 \notin \ell$ in [GR3]. Thus we assume $k \neq k'$. Then there are three cases.
Case (a) If $\mathtt{p}_1, \mathtt{p}_2, \mathtt{q}_1, \mathtt{q}_2$ are pairwise-distinct, then $\mathtt{p}_1 \to \mathtt{p}_2 : k \not\prec_\phi \mathtt{q}_1 \to \mathtt{q}_2 : k'$ in G. Thus no dependency relation from $\mathtt{p}_1 \to \mathtt{p}_2 : k$ to any action $\mathtt{n}' \neq \mathtt{q}_1 \to \mathtt{q}_2 : k'$ in G_1 is changed before and after the transition. Hence obviously $\mathtt{p}_1 \to \mathtt{p}_2 : k \prec_\phi \mathtt{n}' \in G$ implies $\mathtt{p}_1 \to \mathtt{p}_2 : k \prec_\phi \mathtt{n}' \in G'$ for all $\mathtt{n}'$ such that $\mathtt{n}' \neq \mathtt{q}_1 \to \mathtt{q}_2 : k'$. Hence G' is linear.
Case (b) Suppose $\mathtt{p}_1 = \mathtt{q}_2$ and others are pairwise distinct. Then $\mathtt{p}_1 \to \mathtt{p}_2 : k \not\prec_\phi \mathtt{q}_1 \to \mathtt{p}_1 : k'$ in G again. Hence by the same reasoning as above, $\mathtt{p}_1 \to \mathtt{p}_2 : k \prec_\phi \mathtt{n}' \in G$ implies $\mathtt{p}_1 \to \mathtt{p}_2 : k \prec_\phi \mathtt{n}' \in G'$ in $\mathtt{n}' \neq \mathtt{q}_1 \to \mathtt{q}_2 : k'$.
Case (c) Suppose $\mathtt{p}_1 = \mathtt{q}_1$ and others are pairwise distinct. Then again we have $\mathtt{p}_1 \to \mathtt{p}_2 : k \not\prec_\phi \mathtt{q}_1 \to \mathtt{q}_2 : k'$ in G. The rest is the same as the above cases. □

4. TYPE DISCIPLINE FOR MULTIPARTY SESSIONS

4.1. Programming Methodology for Multiparty Interactions

Once given global types as a description of global interactions among communicating processes, we can consider the following development steps for programs with multiparty sessions.

> *Step 1.* A programmer describes an intended interaction scenario as global type G, and checks that it is linear.
> *Step 2.* She develops code, one for the local behaviour of each participant, incrementally validating its conformance to the projection of G onto each participant by efficient type-checking.

The local behaviours might be developed by a team of programmers (who may as well be distributed geographically), in which case the use of a clear, precise global description is all the more essential. When programs are executed, their interactions are guaranteed to follow the stipulated scenario. Further, when transport issues interfere with communication, the global type gives a basic criteria by which communications are monitored and (in)validated at runtime. The type specification also serves as a basis for debugging, maintenance and upgrade.

For all these purposes, we need a type discipline which relates global types to communication behaviour of individual (end-point) programs, and guarantees key properties such as communication safety. This section introduces such a type discipline.

4.2. End-point Types

Syntax. End-point session types or *end-point types*, ranged over by $T, T', ..$, are types for the end-point behaviour of processes, acting as a link between the global types in Section 3, which give intended conversation structures of multiparty sessions, and processes in Section 2.1. The grammar is given in Figure 6 (the grammars for U and

S are repeated from Figure 4). All constructs come from binary session types [Honda et al. 1998] except for the following major changes for multiparty interactions.

— Since a process uses multiple channels for addressing multiple parties, a session type records the identity (number) of the session channel it uses at each action type.
— Since a type is used for type-checking each participant, we use a notation $T@\mathtt{p}$ (called *located type*) representing an end-point type T assigned to participant $\mathtt{p}$. A located type is also used for delegation.

The rest remains identical to the original session types [Honda et al. 1998]. Type $k?\,\langle U\rangle;T$ represents the behaviour of inputting values of type U at s_k (assume $s_1...s_n$ is shared at initialisation), then performing the actions represented by T. Similarly $k!\,\langle U\rangle;T$ is for sending.

Type $k\&\{l_i\colon T_i\}_{i\in I}$ describes a branching (external choice): it waits with n options at k, and then behaves as type T_i if the i-th label is selected; type $k\oplus\{l_i\colon T_i\}_{i\in I}$ represents the behaviour which selects one of the labels say l_i at k then behaves as T_i (internal choice). These four are *action prefixes* in end-point types. We call send and selection types *output types* and receive and branching *input types*. The rest is the same as the global types, demanding type variables occur guarded by a prefix and taking an equi-recursive approach for recursive types. We often omit end. Note that end-point types do not contain parallel composition, hence retaining simplicity.

Projection and Coherence. The following defines the projection of a global type to end-point types at each participant.

DEFINITION 4.1 (PROJECTION). The *projection of G onto* $\mathtt{p}$, written $G\upharpoonright\mathtt{p}$, is inductively given as:

$$-\quad(\mathtt{p}_1\to\mathtt{p}_2\colon k\,\langle U\rangle.\,G')\upharpoonright\mathtt{p}=\begin{cases}k!\,\langle U\rangle;(G'\upharpoonright\mathtt{p}) & \text{if } \mathtt{p}=\mathtt{p}_1\neq\mathtt{p}_2\\ k?\,\langle U\rangle;(G'\upharpoonright\mathtt{p}) & \text{if } \mathtt{p}=\mathtt{p}_2\neq\mathtt{p}_1\\ (G'\upharpoonright\mathtt{p}) & \text{if } \mathtt{p}\neq\mathtt{p}_2 \text{ and } \mathtt{p}\neq\mathtt{p}_1\end{cases}$$

$$-\quad(\mathtt{p}_1\to\mathtt{p}_2\colon k\,\{l_j\colon G_j\}_{j\in J})\upharpoonright\mathtt{p}=\begin{cases}k\oplus\{l_j\colon(G_j\upharpoonright\mathtt{p})\}_{j\in J} & \text{if } \mathtt{p}=\mathtt{p}_1\neq\mathtt{p}_2\\ k\&\{l_j\colon(G_j\upharpoonright\mathtt{p})\}_{j\in J} & \text{if } \mathtt{p}=\mathtt{p}_2\neq\mathtt{p}_1\\ (G_1\upharpoonright\mathtt{p}) & \text{if } \mathtt{p}\neq\mathtt{p}_2 \text{ and } \mathtt{p}\neq\mathtt{p}_1\\ & \text{and } \forall i,j\in J.G_i\upharpoonright\mathtt{p}=G_j\upharpoonright\mathtt{p}\end{cases}$$

$$-\quad(G_1\mid G_2)\upharpoonright\mathtt{p}=\begin{cases}G_i\upharpoonright\mathtt{p} & \text{if } \mathtt{p}\in G_i \text{ and } \mathtt{p}\notin G_j,\ i\neq j\in\{1,2\}\\ \mathsf{end} & \text{if } \mathtt{p}\notin G_1 \text{ and } \mathtt{p}\notin G_2\end{cases}$$

$$-\quad(\mu\mathbf{t}.G)\upharpoonright\mathtt{p}=\begin{cases}\mu\mathbf{t}.(G\upharpoonright\mathtt{p}) & \text{if } G\upharpoonright\mathtt{p}\neq\mathsf{end}\\ \mathsf{end}\end{cases}\qquad\mathbf{t}\upharpoonright\mathtt{p}=\mathbf{t}\qquad\mathsf{end}\upharpoonright\mathtt{p}=\mathsf{end}.$$

When none of the side conditions hold the map is undefined.

We regard the map to act on the syntax of global types. In the branching clause, all the projections of those participants whose behaviour does not depend on the branching should generate an identical end-point type (otherwise undefined); and in parallel composition, $\mathtt{p}$ should be contained in at most a single type, ensuring each type is single-threaded. Note that, for the sake of clarity, we forbid reflexive interactions directly in the definition of projection, making Convention 3.6 redundant. Below, in (2), the term $T_{\mathtt{p}}@\mathtt{p}$ was introduced at the beginning of §4.2.

DEFINITION 4.2 (COHERENCE).

(1) We say G is *coherent* if it is linear and $G \upharpoonright \mathtt{p}$ is well-defined for each $\mathtt{p} \in \mathsf{pid}(G)$, similarly for each carried global type inductively;
(2) $\{T_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I}$ is *coherent* if for some coherent G s.t. $I = \mathsf{pid}(G)$, we have $G \upharpoonright \mathtt{p} = T_{\mathtt{p}}$ for each $\mathtt{p} \in I$.

THEOREM 4.3. *Coherence of G is decidable.*

PROOF. By Proposition 3.13 (2), noting that the projection is only applied to a given global type without unfolding. A complexity analysis is given in [Deniélou and Yoshida 2010]. □

PROPOSITION 4.4. *Assume G is coherent and $G \xrightarrow{\ell} G'$. Then G' is coherent.*

PROOF. By Proposition 3.14, we only have to prove if G is projectable then G' is projectable. This can be done by induction on the last LTS rule applied. Cases [GR1,GR2] are straightforward by definition of projection, while [GR3,GR4,GR5,GR6] follow immediately by induction hypothesis. □

If the projection mapping is undefined, a global type is not coherent. Linearity guarantees linear channel usage including message-order preservation. The next examples demonstrate the need of these conditions.

4.3. Examples of Coherence

The following global type is linear but *not* coherent because the projection is undefined.

$$\mathtt{A} \to \mathtt{B} : k\{\mathsf{ok} : \mathtt{C} \to \mathtt{D} : k'\langle \mathsf{bool}\rangle,\ \mathsf{quit} : \mathtt{C} \to \mathtt{D} : k'\langle \mathsf{nat}\rangle\} \tag{8}$$

Intuitively, when we project this type onto C or D, regardless of the choice made by A, they should behave in the same way: participants C and D should be independent threads. If we change the above nat to bool as:

$$\mathtt{A} \to \mathtt{B} : k\{\mathsf{ok} : \mathtt{C} \to \mathtt{D} : k'\langle \mathsf{bool}\rangle, \mathsf{quit} : \mathtt{C} \to \mathtt{D} : k'\langle \mathsf{bool}\rangle\} \tag{9}$$

we can define the coherent projection as follows:

$$\begin{array}{c}\{\ k \oplus \{\mathsf{ok} : \mathsf{end},\ \mathsf{quit} : \mathsf{end}\}@\mathtt{A},\ k\&\{\mathsf{ok} : \mathsf{end},\ \mathsf{quit} : \mathsf{end}\}@\mathtt{B} \\ k'!\langle \mathsf{bool}\rangle@\mathtt{C},\ k'?\langle \mathsf{bool}\rangle@\mathtt{D}\ \}\end{array}$$

As examples of end-point types which are not coherent, consider processes in the second case of Figure 5:

$$(\mathsf{II})\ \ \mathrm{Bad}\ \ \{s!\langle\rangle@\mathtt{A},\ s?\langle\rangle; s?\langle\rangle@\mathtt{B},\ s!\langle\rangle@\mathtt{C}\}$$

This process is not coherent since the corresponding global type $\mathtt{A} \to \mathtt{B} : s.\mathtt{C} \to \mathtt{B} : s$ is not linear.

4.4. Typing System

The purpose of the typing system is to efficiently type behaviours *which are built by programmers* and hence which do not include runtime elements such as queues.

Environments and Type Algebra. The typing system uses a map from shared names to their sorts $(S, S', ..)$. As given in Figure 6, other than atomic types, a sort has the shape $\langle G\rangle$ assuming G is coherent. Using these sorts, we define:

$$\Gamma \ ::= \ \emptyset \ \mid \ \Gamma, u : S \ \mid \ \Gamma, X : \tilde{S}\tilde{T} \qquad\qquad \Delta \ ::= \ \emptyset \ \mid \ \Delta, \tilde{s} : \{T@\mathtt{p}\}_{\mathtt{p}\in I}$$

A *sorting* $(\Gamma, \Gamma', ..)$ is a finite map from names to sorts or from process variables to sequences of sorts and types. *Typing* $(\Delta, \Delta', ..)$ records linear usage of session channels. In binary sessions types, it assigned a type to a single channel; now it assigns a family of located types to a vector of session channels.

$$\Gamma, a : S \vdash a : S \qquad \Gamma \vdash \mathsf{true}, \mathsf{false} : \mathsf{bool} \qquad \frac{\Gamma \vdash e_i \triangleright \mathsf{bool}}{\Gamma \vdash e_1 \mathsf{or}\, e_2 : \mathsf{bool}} \qquad \text{[NAME], [BOOL], [OR]}$$

$$\frac{\Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash P \triangleright \Delta, \tilde{s} : (G \upharpoonright \mathtt{1})@\mathtt{1} \quad \{\mathtt{1}, \ldots, \mathtt{n}\} = \mathsf{pid}(G) \quad |\tilde{s}| = \mathsf{sid}(G)}{\Gamma \vdash \overline{a}_{[\mathtt{2..n}]}(\tilde{s}).P \triangleright \Delta} \qquad \text{[MCAST]}$$

$$\frac{\Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash P \triangleright \Delta, \tilde{s} : (G \upharpoonright \mathtt{p})@\mathtt{p} \quad \mathtt{p} \in \mathsf{pid}(G) \quad |\tilde{s}| = \mathsf{sid}(G)}{\Gamma \vdash a_{[\mathtt{p}]}(\tilde{s}).P \triangleright \Delta} \qquad \text{[MACC]}$$

$$\frac{\forall j.\, \Gamma \vdash e_j : S_j \qquad \Gamma \vdash P \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash s[k]!\langle \tilde{e} \rangle; P \triangleright \Delta, \tilde{s} : k!\, \langle \tilde{S} \rangle; T@\mathtt{p}} \qquad \text{[SEND]}$$

$$\frac{\Gamma, x : \tilde{S} \vdash P \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash s[k]?(\tilde{x}); P \triangleright \Delta, \tilde{s} : k?\, \langle \tilde{S} \rangle; T@\mathtt{p}} \qquad \text{[RCV]}$$

$$\frac{\Gamma \vdash P \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash s[k]!\langle\!\langle \tilde{t} \rangle\!\rangle; P \triangleright \Delta, \tilde{s} : k!\, \langle T'@\mathtt{p}' \rangle; T@\mathtt{p},\ \tilde{t} : T'@\mathtt{p}'} \qquad \text{[DELEG]}$$

$$\frac{\Gamma \vdash P \triangleright \Delta, \tilde{s} : T@\mathtt{p},\ \tilde{t} : T'@\mathtt{p}'}{\Gamma \vdash s[k]?((\tilde{t})); P \triangleright \Delta, \tilde{s} : k?\, \langle T'@\mathtt{p}' \rangle; T@\mathtt{p}} \qquad \text{[SREC]}$$

$$\frac{\Gamma \vdash P \triangleright \Delta, \tilde{s} : T_j@\mathtt{p} \qquad j \in I}{\Gamma \vdash s[k] \lhd l_j; P \triangleright \Delta, \tilde{s} : k \oplus \{l_i : T_i\}_{i \in I}@\mathtt{p}} \qquad \text{[SEL]}$$

$$\frac{\Gamma \vdash P_i \triangleright \Delta, \tilde{s} : T_i@\mathtt{p} \qquad \forall i \in I}{\Gamma \vdash\ s[k] \rhd \{l_i : P_i\}_{i \in I} \triangleright \Delta, \tilde{s} : k \,\&\{l_i : T_i\}_{i \in I}@\mathtt{p}} \qquad \text{[BRANCH]}$$

$$\frac{\Gamma \vdash P \triangleright \Delta \qquad \Gamma \vdash Q \triangleright \Delta'}{\Gamma \vdash P \mid Q \triangleright \Delta, \Delta'} \qquad \text{[CONC]}$$

$$\frac{\Gamma \vdash e \triangleright \mathsf{bool} \qquad \Gamma \vdash P \triangleright \Delta \qquad \Gamma \vdash Q \triangleright \Delta}{\Gamma \vdash \mathsf{if}\ e\ \mathsf{then}\ P\ \mathsf{else}\ Q \triangleright \Delta} \qquad \text{[IF]}$$

$$\frac{\Delta\ \mathsf{end\ only}}{\Gamma \vdash \mathbf{0} \triangleright \Delta} \qquad \frac{\Gamma, a : \langle G \rangle \vdash P \triangleright \Delta}{\Gamma \vdash (\boldsymbol{\nu}\, a)P \triangleright \Delta} \qquad \text{[INACT],[NRES]}$$

$$\frac{\Gamma \vdash \tilde{e} : \tilde{S} \qquad \Delta\ \mathsf{end\ only}}{\Gamma, X : \tilde{S}\tilde{T} \vdash X\langle \tilde{e}\tilde{s}_1..\tilde{s}_n \rangle \triangleright \Delta, \tilde{s}_1 : T_1@\mathtt{p}_1, .., \tilde{s}_n : T_n@\mathtt{p}_n} \qquad \text{[VAR]}$$

$$\frac{\Gamma, X : \tilde{S}\tilde{T}, \tilde{x} : \tilde{S} \vdash P \triangleright \tilde{s}_1 : T_1@\mathtt{p}_1..\tilde{s_n} : T_n@\mathtt{p}_n \quad \Gamma, X : \tilde{S}\tilde{T} \vdash Q \triangleright \Delta}{\Gamma \vdash \mathsf{def}\ X(\tilde{x}\tilde{s}_1..\tilde{s}_n) = P\ \mathsf{in}\ Q \triangleright \Delta} \qquad \text{[DEF]}$$

Fig. 7. Typing System for Expressions and Processes

NOTATION 4.5.

— We write Δ, Δ' to denote a typing made from the disjoint union of Δ and Δ' always assuming their domains contain disjoint sets of session channels.
— We write $\tilde{s} : T@\mathtt{p}$ for a singleton typing $\tilde{s} : \{T@\mathtt{p}\}$.

A family of located types is needed to link a set of session types to types of a set of processes created by the session initialisation.

Typing System. The type assignment system for processes is given in Figure 7. We use the following judgements for processes and expressions, respectively:

$$\Gamma \vdash P \triangleright \Delta \qquad\qquad \Gamma \vdash e : S$$

These read "under the environment Γ, process P has typing Δ" and "under the environment Γ, expression e has type S". If we set $|\tilde{s}| = 1$ and $\mathtt{n} = 2$, and delete $\mathtt{p}$ from located type, the rules are essentially identical with those for the original binary session [Yoshida and Vasconcelos 2007]. Below, we explain the key rules.

[NAME],[BOOL],[OR] are the rules for the expressions and identical with [Yoshida and Vasconcelos 2007].

[MCAST] is the rule for session request. The condition $\Gamma \vdash a : \langle G \rangle$ says that sessions established on shared channel a will execute according to global type G. Therefore, $\tilde{s}$ must be used in the body P as the *first* projection of G. Note how $\tilde{s}$ are bound in $\overline{a}_{[2..\mathtt{n}]}(\tilde{s}).P$ and therefore disappear from the typing. [MACC] is for the session accept, taking the $\mathtt{p}$-th projection. The end-point type $(G \upharpoonright \mathtt{p})@\mathtt{p}$ means that the participant $\mathtt{p}$ has $G \upharpoonright \mathtt{p}$, which is the projection of G onto $\mathtt{p}$, as its end-point type. In both rules, condition $|\tilde{s}| = \mathsf{sid}(G)$ (see Definition 3.5) ensures the number of session channels meets those in G. The typing $\tilde{s} : T@\mathtt{p}$ (stands for $\tilde{s} : \{T@\mathtt{p}\}$) means that each prefix does not contain parallel threads which share $\tilde{s}$.

[SEND] and [RCV] are the rules for sending and receiving values. Since the k-th name $s[k]$ of $\tilde{s}$ is used as the subject, we record k in the type. Hence, vector $\tilde{s}$ has type $k!\langle \tilde{S} \rangle; T@\mathtt{p}$ in [SEND] and type $k?\langle \tilde{S} \rangle; T@\mathtt{p}$ in [RCV], under the assumption that it is used as $T@\mathtt{p}$ by the subterm P. Note how the relevant type prefixes ($k!\langle \tilde{S} \rangle$ for the output and $k?\langle \tilde{S} \rangle$ for the input) are composed. In both rules, "$\mathtt{p}$" in $T@\mathtt{p}$ ensures that P is (being inferred as) the behaviour for participant $\mathtt{p}$, and its domain should be $\tilde{s}$.

[DELEG] and [SREC] are the rules for delegation of a session and its dual. Delegation of a multiparty session passes the whole remaining capability to participate in a multiparty session: thus operationally we send the whole vector of session channels. The carried type T' is located, making sure that the behaviour by the receiver at the passed channels takes the role of a specific participant (here $\mathtt{p}'$) in the delegated multiparty session. The rest follows the standard delegation rule [Yoshida and Vasconcelos 2007], observing [DELEG] says that $\tilde{t} : T'@\mathtt{p}'$ does not appear in P, symmetrically to [SREC] which uses the channels in P.

[SEL] and [BRANCH], identical with [Yoshida and Vasconcelos 2007], are the rules for selection and branching.

[CONC] composes two processes if their end-point types are disjoint.

[IF], [INACT], [VAR], and [DEF] are standard. [NRES] is the restriction rule for shared name a.

In [INACT] and [VAR], "end only" means Δ only contains end as session types.

As for binary session types, the type checking problem for programs is decidable. Below we say Γ is *well-formed* when global types it uses are all coherent. A program (or a process) is *annotated* when each of its ν-bound shared names is annotated by a well-formed global type.

PROPOSITION 4.6. *Let Γ be well-formed and P be an annotated program. Then it is decidable whether $\Gamma \vdash P \triangleright \emptyset$ is derivable or not.*

PROOF. By annotation, in each typing rule in Figure 7, the conclusion uniquely determines its premise(s). Note also, by well-formedness, projection of global types P may use is always well-defined. □

In this article, we leave open the generalisation of the result to non-annotated programs, and the corresponding result for type inference.

4.5. Typing Examples

Two Buyer Protocol. Write $\overline{a}_{[2,3]}(b_1, b_2, b_2', s).Q_1$ and $a_{[2]}(b_1, b_2, b_2', s).Q_2$ for Buyer1 and Buyer 2 in §2.3. Then Q_1 and Q_2 have the following typing under $\Gamma = \{a : \langle G \rangle\}$ where G is given in the corresponding example in § 3.4, letting $\mathsf{Bi} = i$, $\mathsf{S} = 3$, $b_1 = 1$, $b_2 = 2$, $b_2' = 3$ and $s = 4$ and assuming P_1, P_2, Q are $\mathbf{0}$:

$$\Gamma \vdash Q_1 \triangleright \tilde{s} : s!\,\langle \mathsf{string} \rangle; b_1?\,\langle \mathsf{int} \rangle; b_2'!\,\langle \mathsf{int} \rangle @ \mathtt{B1}$$
$$\Gamma \vdash Q_2 \triangleright \tilde{s} : b_2?\,\langle \mathsf{int} \rangle; b_2'?\,\langle \mathsf{int} \rangle; s \oplus \{\mathsf{ok} : s!\,\langle \mathsf{string} \rangle; b_2?\,\langle \mathsf{date} \rangle; \mathsf{end},\ \mathsf{quit} : \mathsf{end}\} @ \mathtt{B2}$$

Similarly for Seller. After prefixing at a, we can compose all three by [CONC].

A Streaming Protocol. Let $\Gamma = \{a : \langle G' \rangle\}$ where G' is from the streaming example in § 3.4. Let $d = 1$, $k = 2$, $c = 3$, $\mathtt{K} = 1$, $\mathtt{DP} = 2$, $\mathtt{C} = 3$ and $\mathtt{KP} = 4$. Write R_1, R_2, R_3 and R_4 for the processes under the initial prefixes of Kernel, DataProducer, Consumer and KeyProducer, respectively. We can type them as:

$$\Gamma \vdash R_1 \triangleright dkc : \mu \mathbf{t}.d?\,\langle \mathsf{bool} \rangle; k?\,\langle \mathsf{bool} \rangle; c!\,\langle \mathsf{bool} \rangle; \mathbf{t} @ \mathtt{K}$$
$$\Gamma \vdash R_2 \triangleright dkc : \mu \mathbf{t}.d!\,\langle \mathsf{bool} \rangle; \mathbf{t} @ \mathtt{DP} \qquad \Gamma \vdash R_4 \triangleright dkc : \mu \mathbf{t}.c?\,\langle \mathsf{bool} \rangle; \mathbf{t} @ \mathtt{C}$$

(R_4 is similar as R_2). Note these types correspond to the projection of G' onto respective participants: thus Kernel, DataProducer, Consumer and KeyProducer are typable programs under Γ, which can be composed to make the initial configuration.

Delegation. One source of the expressiveness of the session types comes from a facility of *delegation* (often called *higher-order session passing*). We will type the example in § 3.4 and see the relationship with global and end-point types. Consider the following three participants:

$$\begin{array}{ll} \text{Alice} & \stackrel{\text{def}}{=} \overline{a}_{[2]}(t_1, t_2).\overline{b}_{[2,3]}(s_1, s_2).t_1!\langle\!\langle s_1, s_2 \rangle\!\rangle; \mathbf{0} \\ \text{Bob} & \stackrel{\text{def}}{=} a_{[2]}(t_1, t_2).b_{[1]}(s_1, s_2).t_1?((s_1, s_2)); s_1!\,\langle 1 \rangle; \mathbf{0} \\ \text{Carol} & \stackrel{\text{def}}{=} b_{[2]}(s_1, s_2).s_1?(x); P \end{array}$$

where Alice delegates its capability to Bob. Since there are two multicastings, there are two global specifications, one for a and another for b as follows:

$$\begin{array}{rcl} G_a & = & \mathtt{A} \to \mathtt{B} \colon t_1\ \langle s_1!\,\langle \mathsf{int} \rangle @ \mathtt{A} \rangle.\ \mathsf{end} \\ G_b & = & \mathtt{A} \to \mathtt{C} \colon s_1\ \langle \mathsf{int} \rangle.\ \mathsf{end} \end{array}$$

where the type $s_1!\,\langle \mathsf{int} \rangle @ \mathtt{A}$ means the capability to send an integer from participant $\mathtt{A}$ via channel s_1. This capability is passed to $\mathtt{B}$ so that $\mathtt{B}$ behaves as $\mathtt{A}$. However, since the two specifications are independent, $\mathtt{C}$ does not have to know who would pass the capability.

Let (Alice | Bob | Carol) $\to (\boldsymbol{\nu} \tilde{t}\tilde{s})(A \mid B \mid C \mid R)$ where A, B, C are the processes of Alice, Bob and Carol after the initial multicasting and R is the generated queues. Let $s_1 = 1, t_1 = 1, \mathtt{A} = 1, \mathtt{B} = 2, \mathtt{C} = 3$. We have the following typings under Γ with $P \equiv \mathbf{0}$:

$$\Gamma \vdash A \ \triangleright\ \tilde{t} : t_1!\,\langle s_1!\,\langle \mathsf{int} \rangle @ \mathtt{A} \rangle @ \mathtt{A},\ \tilde{s} : s_1!\,\langle \mathsf{int} \rangle @ \mathtt{A}$$
$$\Gamma \vdash B \ \triangleright\ \tilde{t} : t_1?\,\langle s_1!\,\langle \mathsf{int} \rangle @ \mathtt{A} \rangle @ \mathtt{B}$$
$$\Gamma \vdash C \ \triangleright\ \tilde{s} : s_1?\,\langle \mathsf{int} \rangle @ \mathtt{C}$$

where each end-point type reflects the original global specifications (e.g. Carol does not know Alice passed the capability to Bob and Bob behaves as Alice). These types give projections of G_a and G_b.

5. SAFETY AND PROGRESS

This section establishes the fundamental behavioural properties of typed processes. We follow three technical steps:

(1) We extend the typing rules to include those for runtime processes which involve message queues.
(2) We define reduction over session typings which eliminates a pair of minimal complementary actions from end-point types.
(3) We then relate the reduction of processes and that of typings: showing the latter follows the former gives us *subject reduction* (Theorem 5.19), *safety* (Theorem 5.22) and *session fidelity* (Corollary 5.23), while showing the former follows the latter under a certain condition gives us *progress* (Corollary 5.30).

By the correspondence between end-point types and global types, these results guarantee that interactions between typed processes exactly follow the conversation scenario specified in a global type.

Note that the typing system for runtime processes we shall introduce in this section is used solely for establishing the behavioural properties of typed processes, tracing how typability is preserved during reduction. This is in contrast to the simple typing system in § 4 which is for typing programs and program phrases.

5.1. Typing Runtime

How to Type a Queue. We first illustrate a key idea underlying our runtime typing using the following example.

$$\underbrace{s!\langle 3\rangle; s!\langle \mathsf{true}\rangle; \mathbf{0}}_{1} \;\mid\; s :: \emptyset \;\mid\; \underbrace{s?(x); s?(y); \mathbf{0}}_{2} \tag{10}$$

Above, process 1 sends an integer and a boolean to process 2 through queue $s :: \emptyset$. Process 1 can be typed with $s : 1!\langle \mathsf{nat}\rangle; 1!\langle \mathsf{bool}\rangle; \mathsf{end}@\mathtt{p}$ while process 2 by $s : 1?\langle \mathsf{nat}\rangle; 1?\langle \mathsf{bool}\rangle; \mathsf{end}@\mathtt{q}$. After a reduction, (10) changes into:

$$s!\langle \mathsf{true}\rangle; \mathbf{0} \;\mid\; s :: 3 \;\mid\; s?(x); s?(y); \mathbf{0} \tag{11}$$

Note that (11) is identical with (10) except that an output prefix in (10) changes its place to the queue. Thus we can go back from (11) to (10) by placing this message on the top of the process. A key idea in our runtime typing is *to carry out this "rollback of a message" in typing*, using an end-point type with a hole (a type context) for typing a queue. For example we type the queue in (11) as:

$$s : \{\; 1!\langle \mathsf{nat}\rangle; [\;]@\mathtt{p},\; [\;]@\mathtt{q} \;\} \tag{12}$$

where $[\,]$ indicates a hole (this will be formalised in Definition 5.2). Each of the holes above should be filled by the remaining end point type of s at $\mathtt{p}$ and $\mathtt{q}$. Hence, we cover the type $1!\langle \mathsf{bool}\rangle; \mathsf{end}$ with the type context for $\mathtt{p}$ given above, $1!\langle \mathsf{nat}\rangle; [\;]$, obtaining the type $1!\langle \mathsf{nat}\rangle; 1!\langle \mathsf{bool}\rangle; \mathsf{end}$ for $\mathtt{p}$, restoring the original typing.

Labels in a queue are also typed using a type context. For example $k : l_1 \cdot \mathsf{true} \cdot l_2$ can be typed with

$$k \oplus l_1 :\; k!\langle \mathsf{bool}\rangle;\; k \oplus l_2 : [\,], \tag{13}$$

omitting braces for a singleton selection. Now consider reduction

$$s_i \lhd \mathsf{ok}; P \mid s_i : \emptyset \quad \rightarrow \quad P \mid s_i : \mathsf{ok}. \tag{14}$$

Assume we type the left-hand side as

$$\tilde{s} : k \oplus \{\mathsf{ok} : T, \mathsf{quit} : T'\}@\mathtt{p}. \tag{15}$$

After the reduction, we obtain the type for P as

$$\tilde{s} : T@\mathrm{p}. \tag{16}$$

and the type for the queue as:

$$\tilde{s} : k \oplus \{\mathsf{ok} : [\ ]\}@\mathrm{p}. \tag{17}$$

By combining (16) and (17) as before, we obtain

$$\tilde{s} : k \oplus \{\mathsf{ok} : T\}@\mathrm{p}. \tag{18}$$

We now observe that the located type in (18) is a *subtype* of the located type in (15) in the standard session subtyping [Gay and Hole 2005; Carbone et al. 2007; 2012], which is formally defined as [Pierce and Sangiorgi 1996]:

DEFINITION 5.1. *The* subtyping *over end-point types, denoted $\leq_{\mathsf{sub}}$, is the maximal fixed point of function $\mathcal{F}$ that maps each binary relation $\mathcal{R}$ on end-point types as regular trees to $F(\mathcal{R})$ given as:*

— if $(T, T') \in \mathcal{R}$ then $(k!\langle U\rangle; T, k!\langle U\rangle; T') \in \mathcal{F}(\mathcal{R})$ and $(k?\langle U\rangle; T, k?\langle U\rangle; T') \in \mathcal{F}(\mathcal{R})$
— if $(T_i, T_i') \in \mathcal{R}$ for each $i \in I \subset J$ then $(\oplus\{l_i : T_i\}_{i\in I}, \oplus\{l_j : T_j'\}_{j\in J}) \in \mathcal{F}(\mathcal{R})$ and $(\&\{l_j : T_j\}_{i\in J}, \&\{l_i : T_i'\}_{i\in I}) \in \mathcal{F}(\mathcal{R})$.

If $T \leq_{\mathsf{sub}} T'$ then T is a subtype *of T' whereas T' is a* supertype *of T.*

Note that we do not have a subsumption rule for a program in Figure 7. On the other hand, we require a subtying relation between located types to type runtime processes.

Since $k \oplus \{\mathsf{ok} : T\} \leq_{\mathsf{sub}} k \oplus \{\mathsf{ok} : T, \mathsf{quit} : T'\}$, we can type the reductum in (14) using the located type given in (15), which is a supertype of the located type in (18), through the standard subsumption, achieving the required rollback.

Type Contexts. Below, we formalise the notion of type context used in the previous section.

DEFINITION 5.2. *The type contexts $(\mathcal{T}, \mathcal{T}', ...)$ and the extended session typing $(\Delta, \Delta', ...$ as before) are given as:*

$$\begin{array}{lcl} \mathcal{T} & ::= & [\ ] \mid k!\langle U\rangle; \mathcal{T} \mid k \oplus l_i : \mathcal{T} \\ H & ::= & T \mid \mathcal{T} \\ \Delta & ::= & \emptyset \mid \Delta, \tilde{s} : \{H_{\mathrm{p}}@\mathrm{p}\}_{\mathrm{p}\in I} \end{array}$$

Thus a type context represents a sequence of outputs and singleton selections which ends with a hole. As before, the notation "Δ, Δ'" denotes the union assuming the domains should not include a common channel name. The *isomorphism $\approx$ on type contexts* is generated from permutations given below:

DEFINITION 5.3 (PERMUTATION). In addition to the folding/unfolding of recursive types, we consider end-point types up to the following isomorphism (closed under all type constructors).

$$k!\langle U\rangle; k'!\langle U'\rangle; T \approx k'!\langle U'\rangle; k!\langle U\rangle; T \qquad (k \neq k') \tag{19}$$

$$k \oplus \{l_i : k' \oplus \{l_j' : T_{ij}\}_{j\in J}\}_{i\in I} \approx k' \oplus \{l_j' : k \oplus \{l_i : T_{ij}\}_{i\in I}\}_{j\in J} \qquad (k \neq k') \tag{20}$$

$$k \oplus \{l_i : k'!\langle U\rangle; T_i\}_{i\in I} \approx k'!\langle U\rangle; k \oplus \{l_i : T_i\}_{i\in I} \qquad (k \neq k') \tag{21}$$

The equations permute two consecutive outputs or selections with different subjects, capturing asynchrony in communication.

$$\frac{\Gamma \vdash P \triangleright_{\tilde{t}} \Delta \quad \Delta \leq \Delta'}{\Gamma \vdash P \triangleright_{\tilde{t}} \Delta'} \qquad \frac{\Delta \text{ end only}}{\Gamma \vdash s[k]:\emptyset \triangleright_{s[k]} \tilde{s}:\{[\,]@\mathtt{p}\}_{\mathtt{p}} \circ \Delta} \qquad \text{[SUBS],[QNIL]}$$

$$\frac{\Gamma \vdash v_i : S_i \quad \Gamma \vdash s[k]:\tilde{h} \triangleright_{s[k]} \Delta, \tilde{s}:(\{\mathcal{T}@\mathtt{q}\} \cup R) \qquad R = \{H_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I}}{\Gamma \vdash s[k] : \tilde{h} \cdot \tilde{v} \triangleright_{s[k]} \Delta, \tilde{s}:(\{\mathcal{T}[k!\langle \tilde{S}\rangle; [\,]]@\mathtt{q}\} \cup R)} \qquad \text{[QVAL]}$$

$$\frac{\Gamma \vdash s[k]:\tilde{h} \triangleright_{s[k]} \Delta, \tilde{s}:\{\mathcal{T}@\mathtt{q}\} \cup R \qquad R = \{H_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I}}{\Gamma \vdash s[k]:\tilde{h} \cdot \tilde{t}' \triangleright_{s[k]} \Delta, \tilde{s}:(\{\mathcal{T}[k!\langle T'@\mathtt{p}'\rangle; [\,]]@\mathtt{q}\} \cup R), \tilde{t}':T'@\mathtt{p}'} \qquad \text{[QSESS]}$$

$$\frac{\Gamma \vdash s[k] : \tilde{h} \triangleright_{s[k]} \Delta, \tilde{s} : \{\mathcal{T}@\mathtt{q}\} \cup R \qquad R = \{H_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I}}{\Gamma \vdash s[k] : \tilde{h} \cdot l \triangleright_{s[k]} \Delta, \tilde{s}:(\{\mathcal{T}[k \oplus l : [\,]]@\mathtt{q}\} \cup R)} \qquad \text{[QSEL]}$$

$$\frac{\Gamma \vdash P \triangleright_{\tilde{t}_1} \Delta \quad \Gamma \vdash Q \triangleright_{\tilde{t}_2} \Delta' \quad \tilde{t}_1 \cap \tilde{t}_2 = \emptyset \quad \Delta \asymp \Delta'}{\Gamma \vdash P \mid Q \triangleright_{\tilde{t}_1 \cdot \tilde{t}_2} \Delta \circ \Delta'} \qquad \text{[CONC]}$$

$$\frac{\Gamma \vdash P \triangleright_{\tilde{t}} \Delta, \tilde{s}:\{T_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I} \quad \tilde{s} \in \tilde{t} \quad \{T_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I} \text{ coherent}}{\Gamma \vdash (\boldsymbol{\nu}\, \tilde{s})P \triangleright_{\tilde{t}\setminus\tilde{s}} \Delta} \qquad \text{[CRES]}$$

Fig. 8. Selected Typing Rules for Runtime Processes

Assignments in Δ may contain both end-point types and type contexts. Below, we define the partial commutative algebra $\circ$ where $\mathsf{sid}(\mathcal{T})$ are the channel numbers in $\mathcal{T}$.

$$T \circ \mathcal{T} = \mathcal{T} \circ T = \mathcal{T}[T]$$
$$\mathcal{T} \circ \mathcal{T}' = \mathcal{T}[\mathcal{T}'] \quad (\mathsf{sid}(\mathcal{T}) \cap \mathsf{sid}(\mathcal{T}') = \emptyset)$$

In the first rule, we place the output types of message queues on that of a process. In the second, we compose the type contexts for two sets of messages from the mutually disjoint sets of queues. Note $\mathcal{T} \circ \mathcal{T}'$ is defined iff $\mathcal{T}' \circ \mathcal{T}$ is defined and in which case we have $\mathcal{T}[\mathcal{T}'] \approx \mathcal{T}'[\mathcal{T}]$. Note also $T \circ T'$ is never defined.

Below we define a simple algebra of environments for runtime processes.

DEFINITION 5.4 (TYPE ALGEBRA). A partial operator $\circ$ is defined as:

$$\{H_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I} \circ \{H'_{\mathtt{p}'}@\mathtt{p}'\}_{\mathtt{p}'\in J} \quad = \quad \{(H_{\mathtt{p}} \circ H'_{\mathtt{p}})@\mathtt{p}\}_{\mathtt{p}\in I\cap J} \cup \{H_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I\setminus J} \cup \{H'_{\mathtt{p}'}@\mathtt{p}'\}_{\mathtt{p}'\in J\setminus I}$$

assuming each $\circ$ on the right-hand side is defined. Otherwise the operation is undefined. Then we say Δ_1 and Δ_2 are *compatible*, written $\Delta_1 \asymp \Delta_2$, if for all $\tilde{s}_i \in \mathsf{dom}(\Delta_i)$ such that $\tilde{s}_1 \cap \tilde{s}_2 \neq \emptyset$, $\tilde{s} = \tilde{s}_1 = \tilde{s}_2$ and $\Delta_1(\tilde{s}) \circ \Delta_2(\tilde{s})$ is defined. When $\Delta_1 \asymp \Delta_2$, the *composition of* Δ_1 *and* Δ_2, written $\Delta_1 \circ \Delta_2$, is given as:

$$\Delta_1 \circ \Delta_2 \quad = \quad \{\, \Delta_1(\tilde{s}) \circ \Delta_2(\tilde{s}) \;\mid\; \tilde{s} \in \mathsf{dom}(\Delta_1) \cap \mathsf{dom}(\Delta_2) \,\} \cup \Delta_1 \setminus \mathsf{dom}(\Delta_2) \cup \Delta_2 \setminus \mathsf{dom}(\Delta_1)$$

The operation $\Delta \circ \Delta'$ is undefined if $\Delta \asymp \Delta'$ does not hold.

5.2. Typing Rules for Runtime

To guarantee that there is at most one queue for each channel, we use the typing judgement refined as:

$$\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$$

where $\tilde{s}$ (regarded as a set) records the session channels associated with the message queues. The typing rules for runtime are given in Figure 8. [SUBS] allows subsumption

($\leq_{\mathsf{sub}}$ is extended pointwise from types). [QNIL] starts from the empty hole for each participant, recording the session channel in the judgement. [QVAL] says when we enqueue $\tilde{v}$, the type for $\tilde{v}$ is added at the tail. [QSESS] and [QSEL] are the corresponding rules for delegated channels and a label.

[INACT] allows weakening for empty queue types, while [CONC] is refined to prohibit duplicated message queues. The rule does not use coherence (cf. Def.4.2 (2)) since coherence is meaningful only when all participants and queues are ready.

In [CRES], since we are hiding session channels, we now know no other participants can be added. Hence we check all message queues are composed and the given configuration at $\tilde{s}$ is coherent.

For the rest, we refine the original typing rules in Figure 7 not appearing in Figure 8 as follows (the full typing rules are listed in Appendix B).

— For [MCAST],[MACC],[RCV],[SREC],[BRANCH] and [DEF], we replace $\Gamma \vdash P \triangleright \Delta$ with $\Gamma \vdash P \triangleright_{\emptyset} \Delta$.
— [VAR] is similar to [INACT] (so that a queue can never occur in processes realising participants).
— For both [DEF] and [NRES], we replace $\Gamma \vdash P \triangleright \Delta$ by $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$.

Using these typing rules, we can check that the configurations at the beginning of this section, (10) and (11), are given an identical typing by "rolling back" the type of the message in the queues; similarly for the next redex and reductum pair in the same page, (15) and (16).

The typability in the original system in §4 and the one in this system coincide for processes without runtime elements.

PROPOSITION 5.5. *Let P be a program phrase and Δ be without a type context. Then $\Gamma \vdash P \triangleright \Delta$ in the typing system in § 4 iff $\Gamma \vdash P \triangleright_{\emptyset} \Delta$ is derived without using [*SUBS*] in the typing system in this section.*

PROOF. See Appendix B. □

PROPOSITION 5.6. *If $\Gamma \vdash P \triangleright_{s[1..m]} \Delta$ then P has a unique queue at $s[i]$ ($1 \leq i \leq m$), no other queue at a free channel occurs in P, and no queue in P is under any prefix.*

PROOF. It is routine by rule induction, see Appendix B.2. □

5.3. Type Reduction

Next we introduce a reduction relation over session typings, which abstractly represents interaction in processes at session channels. Below we assume well-formedness of types and typing.

DEFINITION 5.7 (TYPE REDUCTION). The syntax of labels ($\ell, \ell', ...$) of local types is defined as follows:

$$\ell \ ::= \ \mathtt{p} \to \mathtt{p}' \colon k\langle U \rangle \ \mid \ \mathtt{p} \to \mathtt{p}' \colon k\langle l \rangle \ \mid \ \mathtt{p} \to \mathtt{p}' \colon s[k]\langle U \rangle \ \mid \ \mathtt{p} \to \mathtt{p}' \colon s[k]\langle l \rangle$$

We generate $\Delta \xrightarrow{\ell} \Delta'$ by the following rule:

$$k!\langle U\rangle; H@\mathtt{p},\ k?\langle U\rangle; T@\mathtt{q} \xrightarrow{\mathtt{p}\to\mathtt{q}:\, k\,\langle U\rangle} H@\mathtt{p},\ T@\mathtt{q} \qquad \text{[TR-COM]}$$

$$k \oplus \{l : H, ...\}@\mathtt{p},\ k\&\{l : T, ...\}@\mathtt{q} \xrightarrow{\mathtt{p}\to\mathtt{q}:\, k\,\langle l\rangle} H@\mathtt{p},\ T@\mathtt{q} \qquad \text{[TR-BRA]}$$

$$\frac{H_1@\mathtt{p}_1,\ H_2@\mathtt{p}_2 \xrightarrow{\ell} H_1'@\mathtt{p}_1,\ H_2'@\mathtt{p}_2 \quad \mathtt{p}_1, \mathtt{p}_2 \in I \quad k \in \ell}{\tilde{s} : \{H_1@\mathtt{p}_1,\ H_2@\mathtt{p}_2, ...\}_{i\in I}, \Delta \xrightarrow{\ell[s[k]/k]} \tilde{s} : \{H_1'@\mathtt{p}_1,\ H_2'@\mathtt{p}_2, ...\}_{i\in I}, \Delta} \qquad \text{[TR-CONTEXT]}$$

$$\frac{\Delta \approx \Delta_0 \quad \Delta_0 \xrightarrow{\ell} \Delta_0' \quad \Delta_0' \approx \Delta'}{\Delta \xrightarrow{\ell} \Delta'} \qquad \text{[TR-ISO]}$$

In the sequel, we investigate the relationship between the LTS semantics of global and local types to prove the key properties for the main theorems.

DEFINITION 5.8 (FULL PROJECTION). *Assume G is coherent. Then the* full projection of G, *denoted by $[\![G]\!]$ is defined as the set $\{(G \upharpoonright \mathtt{p})@\mathtt{p} \mid \mathtt{p} \in \mathsf{pid}(G)\}$. We write $[\![G]\!] \xrightarrow{\ell} [\![G']\!]$ if $\tilde{s} : [\![G]\!] \xrightarrow{\ell[s[k]/k]} \tilde{s} : [\![G']\!]$.*

DEFINITION 5.9 (COHERENCE AND PARTIAL COHERENCE OF TYPINGS). (1) We say Δ *is coherent* if $\Delta(\tilde{s})$ is coherent for each $\tilde{s} \in \mathsf{dom}(\Delta)$. (2) Δ is *partially coherent* if for some Δ' we have $\Delta \asymp \Delta'$ and $\Delta \circ \Delta'$ is coherent.

The following lemma states that for any type reduction in the local types projected from G, its corresponding action $\mathtt{p}_i \to \mathtt{p}_j : k$ in G is the minimum with respect to $\prec_\phi$.

LEMMA 5.10 (PROJECTION AND CAUSALITY). *Assume $[\![G]\!] = \{T_i@\mathtt{p}_i\}_{i\in I}$ and there exists $i, j \in I$ such that $T_i@\mathtt{p}_i, T_j@\mathtt{p}_j \xrightarrow{\ell} T_i'@\mathtt{p}_i, T_j'@\mathtt{p}_j$ with $k \in \ell$. Then there is no action $\mathtt{q} \to \mathtt{q}' : k' \in G$ such that $(\mathtt{q} \to \mathtt{q}' : k') \prec_\phi (\mathtt{p}_i \to \mathtt{p}_j : k) \in G$ with either (i) $\phi \in \{\mathsf{II}, \mathsf{IO}\}$ or (ii) $\phi = \mathsf{OO}$ and $k = k'$.*

PROOF. By the linearity of G, if T_i is the output type at k, then there is no output type at k except T_i. Similarly if T_j is an input type at k, there is no input type at k except T_j in $[\![G]\!]$. Then it is obvious by the definition of $\prec$. □

The key lemma which states the one-to-one correspondence between the semantics of a global type and the semantics of its projected local types follows.

LEMMA 5.11 (GLOBAL AND LOCAL TYPES). *Suppose G is coherent. Then $G \xrightarrow{\ell} G'$ iff $[\![G]\!] \xrightarrow{\ell} [\![G']\!]$.*

PROOF. The only-if direction is straightforward by definition of $[\![G]\!]$. We prove the if direction by induction on the derivation of $[\![G]\!] \xrightarrow{\ell} [\![G']\!]$.

Let us first analyse the case where either [TR-COM] or [TR-BRA] are in the premise of [TR-CONTEXT]. If $\mathtt{p}_1$ and $\mathtt{p}_2$ are toplevel in G then we can straightforwardly use [GR1] and [GR2] from Definition 3.2. Otherwise, if [TR-BRA] is in the premise, then it must be the case that both $\mathtt{p}_1$ and $\mathtt{p}_2$ are not top level in G. This follows by the definition of projection and applicability of [TR-CONTEXT] with [TR-ISO] not in the premise. In such a case, by definition of projection, all roles different from $\mathtt{p}_1$ and $\mathtt{p}_2$ must behave the same on each branch. Hence, the precondition of [GR4] is satisfied and we can apply such rule. Note that the global type obtained after reduction can be projected to the reductum of [TR-CONTEXT] as expected. The case where [T-COM] is used in the premise of [TR-CONTEXT] is similar.

If [TR-ISO] is in the premise of [TR-CONTEXT] then we must have done a permutation of some outputs/selections. We show that such a behaviour can be emulated by the global type semantics. Suppose that

$$\begin{array}{rcl} \tilde{s} : [\![G]\!] & = & \tilde{s} : \{k!\langle U\rangle; k'!\langle U'\rangle; T_1@\mathtt{1}, k?\langle U\rangle; T_2@\mathtt{2}, k'?\langle U'\rangle; T_3@\mathtt{3}\} \\ \stackrel{\ell_0}{\to}\ \tilde{s} : [\![G']\!] & = & \tilde{s} : \{k'!\langle U'\rangle; T_1@\mathtt{1}, T_2@\mathtt{2}, k'?\langle U'\rangle; T_3@\mathtt{3}\} \end{array}$$

where $G = \mathtt{1} \to \mathtt{2} : k\langle U\rangle.\mathtt{1} \to \mathtt{3} : k'\langle U'\rangle.G_1$ and $G' = \mathtt{1} \to \mathtt{3} : k'\langle U'\rangle.G_1$ with $\ell_0[s[k]/k] = \ell$.

Now, suppose

$$\begin{array}{rcl} \tilde{s} : [\![G]\!] & \approx & \tilde{s} : \{k'!\langle U'\rangle; k!\langle U\rangle; T_1@\mathtt{1}, k?\langle U\rangle; T_2@\mathtt{2}, k'?\langle U'\rangle; T_3@\mathtt{3}\} \\ \stackrel{\ell'_0}{\to}\ \tilde{s} : [\![G_0]\!] & = & \tilde{s} : \{k!\langle U\rangle; T_1@\mathtt{1}, k?\langle U\rangle; T_2@\mathtt{2}, T_3@\mathtt{3}\} \end{array}$$

by (19) in Definition 5.3 where $G_0 = \mathtt{1} \to \mathtt{2} : k\langle U\rangle.G_1$. By the definition of LTS ([GR3] in Definition 3.2), we can obtain $G_1 \stackrel{\ell'}{\to} G_0$ with $\ell'_0[s[k]/k] = \ell'$, as required. Other cases are similar. □

The following lemma states that the transitions from global types and projected local types are deterministic.

LEMMA 5.12 (DETERMINACY).

(1) *Suppose G is coherent. Then $G \stackrel{\ell}{\to} G_1$ and $G \stackrel{\ell}{\to} G_2$ imply $G_1 \approx G_2$.*
(2) *Suppose Δ is coherent. Then $\Delta \stackrel{\ell}{\to} \Delta_1$ and $\Delta \stackrel{\ell}{\to} \Delta_2$ imply $\Delta_1 \approx \Delta_2$.*
(3) *Suppose G is coherent, and $G \stackrel{\ell_1}{\to} G_1$ and $G \stackrel{\ell_2}{\to} G_2$ with $k \in \ell_1, \ell_2$. Then $\ell_1 = \ell_2$*
(4) *Suppose Δ is coherent, and $\Delta \stackrel{\ell_1}{\to} \Delta_1$ and $\Delta \stackrel{\ell_2}{\to} \Delta_2$ with $s[k] \in \ell_1, \ell_2$. Then $\ell_1 = \ell_2$.*

PROOF. (1) is immediate noting that if $G_1 \mid G_2 \stackrel{\ell}{\to} G'_1 \mid G_2$ then $G_2 \stackrel{\ell}{\not\to}$ (since the participants are disjoint between G_1 and G_2). (2) is by (1) and Lemma 5.10. (3) and (4) are similar with (1) and (2), respectively. □

The following proposition states that (1) transitions of Δ is closed under $\asymp$; (2,3) Δ is invaliant w.r.t. partial and coherence; and (4) the transition of a global type and its mapping have exact correspondence.

PROPOSITION 5.13.

(1) *$\Delta_1 \stackrel{\ell}{\to} \Delta'_1$ and $\Delta_1 \asymp \Delta_2$ imply $\Delta'_1 \asymp \Delta_2$ and $\Delta_1 \circ \Delta_2 \stackrel{\ell}{\to} \Delta'_1 \circ \Delta_2$.*
(2) *Let Δ be coherent. Then $\Delta \stackrel{\ell}{\to} \Delta'$ implies Δ' is coherent.*
(3) *Let Δ be partial coherent. Then $\Delta \stackrel{\ell}{\to} \Delta'$ implies Δ' is partial coherent.*
(4) *Let Δ be coherent and $\Delta(\tilde{s}) = [\![G]\!]$. Then $\Delta \stackrel{\ell}{\to} \Delta'$ with $s[k] \in \ell$ iff $G \stackrel{\ell[k/s[k]]}{\to} G'$ with $\Delta'(\tilde{s}) = [\![G']\!]$.*

PROOF. For (1) suppose $\Delta_1 \stackrel{\ell}{\to} \Delta'_1$ with $s[k] \in \ell$ and $\Delta_1 \asymp \Delta_2$. Note by definition of $\Delta_1 \asymp \Delta_2$, each pair of vectors of channels from $\Delta_{1,2}$ either coincide or are disjoint, i.e. (a) $s[k] \in \tilde{s} \in \mathsf{dom}(\Delta_1) \cap \mathsf{dom}(\Delta_2)$ or (b) $s[k] \in \tilde{s} \in \mathsf{dom}(\Delta_i)$ and $\tilde{s} \notin \mathsf{dom}(\Delta_j)$ with $i \neq j$. For case (a), since the typed reduction only erases the top input and output pair in Δ_1, we have $\Delta'_1 \asymp \Delta_2$ by the inductive hypothesis and Lemma 5.12 (2). Then $\Delta_1 \circ \Delta_2 \stackrel{\ell}{\to} \Delta'_1 \circ \Delta_2$ is by definition.

Case (b) is vacuous since the reduction does not relate to the domain of Δ_2. Hence $\Delta'_1 \asymp \Delta_2$.

For (2), suppose Δ is coherent and $\Delta \xrightarrow{\ell} \Delta'$. Suppose the associated redex is in $\Delta(\tilde{s})$. By coherence we can write $\Delta(\tilde{s})$ as $[\![G]\!]$ for some coherent G. By Lemma 5.11, there exists G' such $G \xrightarrow{\ell'} G'$ such that $\ell'[s[k]/k] = \ell$ and $[\![G']\!] = \Delta'(\tilde{s})$. Then by Proposition 4.4, G' is coherent. Hence $[\![G']\!]$ and $\Delta'(\tilde{s})$ are both coherent.

Implication (3) is immediate from (1) and (2).

Finally the only if-direction of (4) follows directly from Definition 5.8, while the if direction is immediate by Lemma 5.11. □

5.4. Subject Reduction and Communication Safety

For subject reduction we use the following lemmas. In the first lemma below, we say that two typings, Δ_1 and Δ_2, *share a common target channel in their type contexts* when, for some $\tilde{s}$ and k, we have: (1) $\mathcal{T}_1@\mathtt{p} \in \Delta_1(\tilde{s})$ and $\mathcal{T}_2@\mathtt{p} \in \Delta_2(\tilde{s})$; and (2) $k!\langle U\rangle$ or $k \oplus l$ occurs in $\mathcal{T}_1$ and $k!\langle U'\rangle$ or $k \oplus l'$ occurs in $\mathcal{T}_2$ (i.e. they have an output/selection type at a shared channel).[4]

LEMMA 5.14 (PARTIAL COMMUTATIVITY AND ASSOCIATIVITY OF $\circ$). *$\circ$ on typings is partially commutative and associative with identity $\emptyset$ under the condition that, whenever we compose two typings, they never share a target channel in their type contexts (in the above sense).*

PROOF. See Appendix B.3. □

LEMMA 5.15. *Assume $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$. Then all free names and free variables in P occur in Γ and all free channels in P occur in Δ.*

Below a *derivation of* $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ is a derivation tree of the typing rules for runtime processes (fully listed in Appendix B) whose conclusion is $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$.

LEMMA 5.16 (PERMUTATION). (1) *Assume given a derivation of $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ which uses [*SUBS*] at its last two steps. Then $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ has a derivation identical with the original one except its last two steps are replaced by a single application of [*SUBS*].* (2) *Assume given a derivation of $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ which uses [*SUBS*] as its last rule and another rule which is not one of [*SUBS*], [*SEL*] and [*BRANCH*]. Then $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ has a derivation which is identical with the original one except that the last two rules used are permuted.*

PROOF. (1) is immediate from the transitivity of [SUBS]. (2) is routine. □

LEMMA 5.17 (QUEUE). *The following rules are admissible in the typing system for runtime processes. Below, let $\tilde{s} = s[1..k..n]$ and let us assume that occurrences of $\circ$ in the premise of each rule are well-defined.*

$$\frac{\Gamma \vdash s[k] :: \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}@\mathtt{p}\} \quad \Gamma \vdash \tilde{v} \triangleright \tilde{S}}{\Gamma \vdash s[k] :: \tilde{h} \cdot \tilde{v} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}[k!\langle \tilde{S}\rangle; [\,]]@\mathtt{p}\}} \quad [\mathbf{Q}\textsc{val}]$$

$$\frac{\Gamma \vdash s[k] :: \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}@\mathtt{p}\} \quad \{\tilde{t}\}\ \textit{fresh}}{\Gamma \vdash s :: \tilde{h} \cdot \tilde{t} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}[k!\langle T@\mathtt{p}'\rangle; [\,]]@\mathtt{p}\},\ \tilde{t} : \{T@\mathtt{p}'\}} \quad [\mathbf{Q}\textsc{sess}]$$

[4] Whenever we compose two processes, their typings never share a common target channel in their type contexts in this sense because, by the disjointness of mentioned channels for queues, target channels in type contexts can never coincide.

$$\frac{\Gamma \vdash s :: \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}@\mathtt{p}\}}{\Gamma \vdash s :: \tilde{h} \cdot l \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}[k \oplus \{.., l : [\,], ..\}]@\mathtt{p}\}} \quad [\mathrm{Q}\textsc{sel}]$$

$$\frac{\Gamma \vdash s :: \tilde{v} \cdot \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{k!\,\langle \tilde{S}\rangle; \mathcal{T}@\mathtt{p}\}@\mathtt{p}}{\Gamma \vdash s :: \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}@\mathtt{p}\}} \quad [\mathrm{Q}\textsc{valdq}]$$

$$\frac{\Gamma \vdash s :: \tilde{t} \cdot \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{k!\,\langle T@\mathtt{p}'\rangle; \mathcal{T}@\mathtt{p}\},\ \tilde{t} : \{T@\mathtt{p}'\}@\mathtt{p}'}{\Gamma \vdash s :: \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}@\mathtt{p}\}} \quad [\mathrm{Q}\textsc{sessdq}]$$

$$\frac{\Gamma \vdash s :: l \cdot \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{k \oplus l : \mathcal{T}@\mathtt{p}\}}{\Gamma \vdash s :: \tilde{h} \triangleright_{\tilde{s}'} \Delta \circ \tilde{s} : \{\mathcal{T}@\mathtt{p}\}} \quad [\mathrm{Q}\textsc{seldq}]$$

PROOF. See Appendix B.2. □

Below we do not require the substitution lemmas for session channels and process variables, cf. [Yoshida and Vasconcelos 2007].

LEMMA 5.18 (SUBSTITUTION AND WEAKENING). (1) *(substitution)* $\Gamma, x : S \vdash P \triangleright_{\tilde{s}} \Delta$ *and* $\Gamma \vdash v : S$ *imply* $\Gamma \vdash P[v/x] \triangleright_{\tilde{s}} \Delta$. (2) *(weakening) Whenever* $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ *is derivable then its weakening,* $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta, \Delta'$ *for disjoint* Δ' *where* Δ' *contains only empty type contexts and for types* end*, is also derivable.*

PROOF. Standard, see [Yoshida and Vasconcelos 2007]. □

Among the lemmas above, the lemmas for queues are needed for treating reduction involving queues in the present asynchronous operational semantics. We can now establish subject reduction.

Subject Reduction, Communication Safety and Session Fidelity. By the above proposition and the substitution lemma, we obtain:

THEOREM 5.19 (SUBJECT CONGRUENCE AND REDUCTION).

(*1*) $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ *and* $P \equiv P'$ *imply* $\Gamma \vdash P' \triangleright_{\tilde{s}} \Delta$.
(*2*) $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ *such that* Δ *is coherent and* $P \to P'$ *imply* $\Gamma \vdash P' \triangleright_{\tilde{s}} \Delta'$ *where* $\Delta = \Delta'$ *or* $\Delta \xrightarrow{\ell} \Delta'$ *for some* ℓ.
(*3*) $\Gamma \vdash P \triangleright_{\emptyset} \emptyset$ *and* $P \to P'$ *imply* $\Gamma \vdash P' \triangleright_{\emptyset} \emptyset$.

PROOF. See Appendix B.4. □

REMARK 5.20. Theorem 5.19 (3) and the subsequent results (in particular Theorem 5.22 and Corollary 5.30 below) tell us, through Proposition 5.5, that the typing system in § 4, which is for programs and program phrases, guarantees type safety and other significant behavioural properties for typable programs, noting typability of (annotated) programs is decidable by Proposition 4.6.

Theorem 5.19 immediately entails the lack of the standard type errors in expressions (such as $\mathsf{true} + 3$). The type discipline also satisfies, as in the preceding session type disciplines [Honda et al. 1998], communication error freedom, including linear usage of channels. We first introduce the reduction context $\mathcal{E}$ as follows:

$$\mathcal{E} \quad ::= \quad \mathcal{E} \,|\, P \quad | \quad P \,|\, \mathcal{E} \quad | \quad (\boldsymbol{\nu} n)\mathcal{E} \quad | \quad \mathsf{def}\ D\ \mathsf{in}\ \mathcal{E}$$

We also say and write:

— A prefix is *at* s (resp. *at* a) if its subject (i.e. its initial channel) is s (resp. a). Further a prefix is *emitting* if it is request, output, delegation or selection, otherwise it is *receiving*.
— A prefix is *active* if it is not under a prefix or an `if` branch, after any unfoldings by [DEF]. We write $P\langle\!\langle s \rangle\!\rangle$ if P contains an active subject at s after applying [DEF], and $P\langle\!\langle s! \rangle\!\rangle$ (resp. $P\langle\!\langle s? \rangle\!\rangle$) if P contains an emitting (resp. receiving) active prefix at s.
— P has a *redex at* s if it has an active prefix at s among its redexes.

Below and henceforth we safely confuse a channel (as a number) in a typing and the corresponding free session channel of a process.

LEMMA 5.21. *Assume* $\Gamma \vdash P \triangleright_{\emptyset} \Delta$ *s.t.* $\Delta \circ \Delta_0$ *is coherent for some* Δ_0.

(*1*) *If* $P\langle\!\langle s \rangle\!\rangle$ *then* P *contains either a unique active prefix at* s *or a unique active emitting prefix and a unique active receiving prefix at* s.
(*2*) *If* P *contains an active emitting (resp. receiving) prefix at* s *then* Δ *contains an emitting (resp. receiving) minimal prefix at* s.

PROOF. By easy rule induction, see Appendix B.6. □

The following result adapts the standard properties for synchronous session types [Takeuchi et al. 1994; Honda et al. 1998; Yoshida and Vasconcelos 2007] to multiparty asynchronous session types. Note that reductions may go wrong for several reasons. Traditional problems include non-boolean values in a conditional, as in $\mathsf{if}\ a\ \mathsf{then}\ P\ \mathsf{else}\ Q$, and arity mismatch for process definitions such as as in $\mathsf{def}\ X(yx) = P\ \mathsf{in}\ X\langle\mathsf{true}\rangle$. Here, we are instead interested in *communication safety*, which ensures there is no error when participants interact with each other. Since interactions always happen at session channels, we focus on the linearity property (no races) and the interactions between processes and their corresponding queue. Below we assume the standard bound name convention.

THEOREM 5.22 (COMMUNICATION SAFETY). *Suppose* $\Gamma \vdash P \triangleright_{\tilde{t}} \Delta$ *s.t.* Δ *is coherent and* P *has a redex at free* s. *Then:*

(*1*) *(linearity)* $P \equiv \mathcal{E}[s :: \tilde{h}]$ *such that either*
  (*a*) $P\langle\!\langle s? \rangle\!\rangle$, s *occurs exactly once in* $\mathcal{E}$ *and* $\tilde{h} \neq \emptyset$; *or*
  (*b*) $P\langle\!\langle s! \rangle\!\rangle$ *and* s *occurs exactly once in* $\mathcal{E}$; *or*
  (*c*) $P\langle\!\langle s? \rangle\!\rangle$, $P\langle\!\langle s! \rangle\!\rangle$, *and* s *occurs exactly twice in* $\mathcal{E}$.
(*2*) *(error-freedom) if* $P \equiv \mathcal{E}[R]$ *with* $R\langle\!\langle s? \rangle\!\rangle$ *being a redex:*
  (*a*) *If* $R \equiv s?(\tilde{y}); Q$ *then* $P \equiv \mathcal{E}'[s : \tilde{v} \cdot \tilde{h}]$ *for some* $\mathcal{E}'$ *and* $|\tilde{v}| = |\tilde{y}|$.
  (*b*) *If* $R \equiv s?((\tilde{s}\,)); Q$ *then* $P \equiv \mathcal{E}'[s : \tilde{t} \cdot \tilde{h}]$ *for some* $\mathcal{E}'$ *and* $|\tilde{s}| = |\tilde{t}|$.
  (*c*) *If* $R \equiv s \triangleright \{l_i : Q_i\}_{i \in I}$ *then* $P \equiv \mathcal{E}'[s : l_j \cdot \tilde{h}]$ *for some* $\mathcal{E}'$ *and* $j \in I$.

PROOF. For (1), let $P \equiv (\boldsymbol{\nu}\tilde{n})(P_0 | s : \tilde{h} | Q)$ where P_0 does not contain a queue and Q only contains queues (by Proposition 5.6). By Lemma 5.21 we know P_0 has either a single active prefix or a pair of a receiving active prefix and an emitting active prefix. So we have three cases:

— $P_0\langle\!\langle s? \rangle\!\rangle$ and there is no other active prefixes at s: if so because there is a redex in P the queue cannot be empty.
— $P_0\langle\!\langle s! \rangle\!\rangle$ and there is no other active prefixes at s: then this gives us a redex.
— $P_0\langle\!\langle s! \rangle\!\rangle$ and $P_0\langle\!\langle s? \rangle\!\rangle$. Then at least the former gives a redex but the latter can also give a redex.

Hence as required.

For (2), if P satisfies the stated condition then we can write $P \equiv \mathcal{E}'[s : \tilde{h}|R]$ and $S \stackrel{\text{def}}{=} s : \tilde{h}|R$ form a redex, with the same typing by Theorem 5.7 (1). Since this should have a partially coherent typing it in particular means the pair of active prefixes at s in the tying of S should be complementary. The rest is by the direct correspondence between the type constructors and the prefixes. □

By Theorems 5.19 and 5.22, a typed process "never goes wrong" in the sense that its interaction at a multiparty session channel is always one-to-one and that each delivered value matches the receiving prefix.

By Lemma 5.21 (2) and by the typing of the associated queue, this delivery precisely corresponds to a redex in the session typing.

As the corollary of Theorem 5.19(2) and Proposition 5.13(4), we obtain *session fidelity*: the interactions of a typable process exactly follow the specification described by its global type.

COROLLARY 5.23 (SESSION FIDELITY). *Assume* $\Gamma \vdash P \triangleright_{\tilde{t}} \Delta$ *such that* Δ *is coherent and* $\Delta(\tilde{s}) = [\![G]\!]$. *If*

(*1*) $P\langle\!\langle s[k]?\rangle\!\rangle \to P'$ *at the redex of* $s[k]$, *then* $\Gamma \vdash P' \triangleright_{\tilde{t}} \Delta'$ *with* $G \stackrel{\ell}{\to} G'$ *with* $k \in \ell$ *and* $[\![G']\!] = \Delta'(\tilde{s})$, *or*
(*2*) $P\langle\!\langle s[k]!\rangle\!\rangle \to P'$ *at the redex of* $s[k]$, *then* $\Gamma \vdash P' \triangleright_{\tilde{t}} \Delta$.

PROOF. In (1), the conclusion $\Gamma \vdash P' \triangleright_{\tilde{t}} \Delta'$ where $\Delta = \Delta'$ or $\Delta \stackrel{\ell}{\to} \Delta'$ follows directly from Theorem 5.19(2). The second conclusion $G \stackrel{\ell}{\to} G'$ with $k \in \ell$ and $[\![G']\!] = \Delta'(\tilde{s})$ follow directly from Lemma 5.12 (3) and Proposition 5.13 (4). If not, a sender puts some value in the queue. Hence (2) obviously holds. □

5.5. Progress

Communication safety says that if a process ever does a reduction, it conforms to the typing and it is linear. If interactions within a session are not hindered by initialisation and communication of *different* sessions, then the converse holds: the reduction predicted by the typing surely takes place, that is the standard progress property in binary session types [Dezani-Ciancaglini et al. 2006; Honda et al. 1998]. First we define:

DEFINITION 5.24. Let $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$. Then P is *queue-full* when $\{\tilde{s}\}$ coincide with the set of session channels occurring in Δ.

A process is queue-full when it has a queue for each session channel. The following precludes interleaving of other sessions (including initialisations and communications) which can introduce deadlock. For example, two session initialisations $a[2](s).b[2](t).s?;t!$ and $\overline{a}[2](s).\overline{b}[2](t).t?;s!$ cause deadlock. Observe, because we have multiparty sessions, there is less need to use interleaved sessions.

DEFINITION 5.25 (SIMPLE). A process P is *simple* when it is typable with a type derivation where the session typing in the premise and the conclusion of each prefix rule in Figure 7 is restricted to at most a singleton. I.e. *(1)* Δ of [MCAST], [MACC], [SEND], [RCV], [BRANCH] and [VAR] are empty; *(2)* Neither [RCV] nor [DELEG] is used; *(3)* Δ of [IF], [INACT], [NRES] and [DEF] contains at most a singleton; and in [CONC], either Δ, Δ' contains at most a singleton.

Thus each prefixed subterm in a simple process has only a unique session.

PROPOSITION 5.26. *Let P_0 be simple and $P_0 \to^* P$. Then no delegation prefix (input or output) occurs in P and for each prefix with a shared name in P, say $a[\mathtt{i}](\tilde{s}).P'$ or $\overline{a}[2..\mathtt{n}](\tilde{s}).P'$, there is no free session channels in P' except $\tilde{s}$.*

PROOF. See Appendix B.7. □

Another element which can hinder progress is when interactions at shared names cannot proceed.

DEFINITION 5.27 (WELL-LINKED). We say P is *well-linked* when for each $P \to^* Q$, whenever Q has an active prefix whose subject is a (free or bound) shared name, then it is always part of a redex.

Thus, in a simple well-linked P, each session is never hindered by other sessions nor by a name prefixing. The key lemma for simple processes follows. Below we safely confuse a channel in a typing and the corresponding free session channel of a process.

LEMMA 5.28. *Let $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ and P is simple. If there is an active receiving (resp. active emitting) prefix in Δ at s and none of prefixes at s in P is under a prefix at a shared name or under an* `if`*-branch, then $P \langle\!\langle s? \rangle\!\rangle$ (resp. either $P \langle\!\langle s! \rangle\!\rangle$ or the queue at s is not empty).*

PROOF. By rule induction using Proposition 5.26, see Appendix B.8. □

PROPOSITION 5.29. *Let $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$, Δ is coherent, P is simple, well-linked and queue-full. Then:*

(*1*) *If $P \not\equiv \mathbf{0}$ then $P \to P'$ for some P'.*
(*2*) *If $\Delta(\tilde{t}) = [\![G]\!]$ and $G \xrightarrow{\ell} G'$ with $k \in \ell$, then $P \to^+ P'$ at the redex at t_k s. t. $\Gamma \vdash P' \triangleright_{\tilde{s}} \Delta'$ with $\Delta'(\tilde{t}) = [\![G']\!]$.*

PROOF. Let P be simple, queue-full and well-linked, and $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ such that Δ is coherent. Without loss of generality we can assume P does not have hidings (we can just take off and the result is again simple, queue-full, well-linked and coherent). Since Δ is coherent, if Δ contains any prefix then, by Proposition 5.26, it should form a redex (together with another prefix to form the image of a identical set). By Lemma 5.28 and Theorem 5.22 (1,2) and by the well-linkedness, either there is an if-branch above the prefix or P has an active prefix (or prefixes) at s in P. For the former, this if-branch itself cannot be under any prefix since that violates the activeness at s in Δ. So this if-branch can reduce; hence, we conclude the case.

If not then by Lemma 5.28 there are the following cases:

(a) $P \equiv \mathcal{E}[Q\langle\!\langle s! \rangle\!\rangle | s : \tilde{h} | R\langle\!\langle s? \rangle\!\rangle]$, in which case there is at least one redex in P between the emitting prefix and the queue.
(b) $P \equiv \mathcal{E}[s : \tilde{h} | R\langle\!\langle s? \rangle\!\rangle]$ with $\tilde{h}$ non-empty, in which case there is a redex between the non-empty queue and the receiving redex.
(c) $P \equiv \mathcal{E}[Q\langle\!\langle s! \rangle\!\rangle | s : \tilde{h}]$, in which case there is a redex as in (a).

In each case there is a reduction hence done. □

(2) above gives the converse of Corollary 5.23: if the global type has a reduction, then the process can always realise it.

COROLLARY 5.30 (PROGRESS). Let P be a simple and well-linked program. Then P has the *progress property* in the sense that $P \to^* P'$ implies either $P' \equiv \mathbf{0}$ or $P' \to P''$ for some P''.

PROOF. Immediate from Proposition 5.26, Lemma 5.28 and Theorem 5.29. □

A simple application of Theorems 5.19 (3) and 5.22 and Corollary 5.30 for processes from §2.3 follow. Below *communication mismatch* stands for the violation of the conditions given in Theorem 5.22 (2).

PROPOSITION 5.31 (PROPERTIES OF TWO PROTOCOLS).

(1) *Let* $Buyer1|Buyer2|Seller \rightarrow^* P$. *Then* P *is well-typed, simple and well-linked,* P *has no communication mismatch, and either* $P \equiv \mathbf{0}$ *or* $P \rightarrow P'$ *for some* P'.
(2) *Similarly for* $DataProducer|KeyProducer|Kernel|Consumer$.

PROOF. Immediate from Corollary 5.30 because these two configurations are typable programs each of which loses its shared name in the initial reduction (at which point all the occurrences of the shared name are used). □

The significance of the progress result under these constraints is that, if a typable program ever gets stuck during reduction, then its causes are other than the structure of individual typed conversations: thus we are ensured that the causes of deadlock (if any) in typed interactions do not lie in each conversation structure itself, allowing their well-articulated analysis.

6. EXTENSIONS AND RELATED WORK

We outline applications and several possible extensions of the presented framework, then discuss related works. We also sumarise recent results and applications of multiparty session types after the publication of the extended abstract [Honda et al. 2008a].

6.1. Applications and Extensions

Applications. As we have already discussed (cf. §1 and §4.1), the type discipline we have explored in the present paper is intended to be used as a typed foundation for the development of communication-centred software in various ways and at different development stages. Types will also serve as a core specification upon which other formal specifications and techniques such as program analyses and assertions may be built.

A global type serves as an agreement of a protocol following which each end-point program will execute its communication. An automatic method to check well-formedness of the global types (linearity, by Proposition 3.13, and coherence, by Proposition 4.3) guarantees the behaviours specified by the global specification. Development of individual programs for end-point communications, which materialise a global conversation, is assisted in several ways: first, the projection of a global type to each participant (well-defined by coherence) directly suggests the possible shape of end-point interactional behaviour. Second, during development, a programmer can check whether her program conforms to the agreed global type through type-checking the program against an appropriate projection of the global type (Proposition 4.6). The global type and its projections may also be used as a basis of the debugging/testing process, including automatic generation of test suites.

Once the development of all programs is complete, their typability ensures, in the absence of systems errors (such as transport-level failure), that the runtime behaviour of the deployed programs satisfy the key properties including communication safety, session fidelity and progress, through the theorems in §5 (cf. Theorems 5.19 and 5.22 and Corollary 5.30). Since global types and their projections specify possible legitimate interaction sequences of the deployed programs, they can be used for runtime monitoring, flagging those communications which go out of expected conversation sequences and thus signalling the existence of system-level errors (which is another direct consequence of the theorems in §5), thus helping locating the cause of such errors. These

static and dynamic validations of programs may add further precision by using refined specifications, such as logical assertions following (global and local) session type structures as we have recently done in [Bocchi et al. 2010].

The effectiveness of these applications hinges on the exactitude with which global types and the associated type discipline can assure basic properties of programs, and thus is underpinned by the formal results discussed in the present paper. At the same time, in order to put these ideas into practise, the presented framework may need various extensions as well as engineering experiments. Some possible extensions of the presented type discipline are discussed in the following.

Existing Extensions of Binary Session Types. In the literature, several extensions of binary session type disciplines have been proposed, including subtyping [Gay and Hole 2005], bounded polymorphism [Gay 2008], integration with security annotations to guarantee authentication properties [Bonelli et al. 2005], and integration with higher-order π-calculus [Mostrous and Yoshida 2007; 2009]. We believe that integrations with these extensions should be possible and will enrich the expressive power and applicability of the theory.

Multithreaded Participants. Another straightforward extension is to allow a multithreaded participant, so that a single participant can perform parallel conversations with others during a session. For this extension, we need to augment end-point types with the parallel composition $T_1 \mid T_2$, equipped with the following isomorphism (using type contexts in §5): $\mathcal{T}[T_1] \mid T_2 \approx \mathcal{T}[T_1 \mid T_2]$ if for no k there is an output at k in both $\mathcal{T}$ and T_2 (such a prefix adds false OO-dependency), as well as commutativity and associativity. Linearity between T_1 and T_2 in T_1, T_2 is given by coherence via projection. This extension has been recently studied with more advanced dynamic roles in [Deniélou and Yoshida 2011].

Graph-Based Global Types and Type Inference. The syntax of global types uses the standard abstract syntax tree. We can further generalise this tree-based syntax to graph structures to obtain a strictly more expressive type language, enlarging typability. Consider the two end-point processes $P \equiv s!\langle t\rangle?$ and $Q \equiv t!\langle s\rangle?$: their parallel composition does not introduce conflict hence it is linear and safe. This situation however cannot be represented in the current global types since two "prefixes" criss-cross each other. Interestingly, our linearity conditions in §3.5, based on input/output dependencies, can directly capture the safety of this configuration. All we need to do is to take the graphs of prefixes and II, IO and OO-edges (cf. Figure 5) under the linearity condition (precisely following §3.5) as global types, augmented with an acyclicity condition on chains of these causal edges. All other definitions and results stay the same.

Our recent work [Mostrous et al. 2009] studies a generation of graph-based global types from end-point types, where we also use such graph-based types for solving the type inference problem for (the generalised version of) the presented type discipline. This is further extended in [Deniélou and Yoshida 2012; Lange et al. 2015] making a connection with Communicating Automata. See §6.3.

Synchrony and Asynchrony. Most of the session types currently studied are binary and synchronous [Honda et al. 1998]. In some computing environments (e.g. tightly coupled SMP), synchrony would be more suitable. Adding synchrony means we have more causality: OO-dependency between different names as well as the OI-dependency (i.e. the dependency from output to input, cf. Figure 5), which in asynchrony never arises §3.4. Our subsequent work [Bejleri and Yoshida 2009] studies a synchronous multiparty session type.

A different direction is to consider asynchronous message passing without order-preservation [Honda and Tokoro 1991] which are also used in some computing envi-

ronments (though in many environments we have efficient order-preserving transport such as TCP). Again we can use our modular articulation, by taking off OO-edges to obtain a consistent theory for pure asynchrony.

Multicast Primitives for Sessions Communication. The two-buyer protocol uses a multicasting prefix notation $s, t!\langle V\rangle$. The present work treats it as a macro for $s!\langle V\rangle; t!\langle W\rangle$ which has an essentially identical abstract semantics. Having proper multicasting primitives for session communication is however useful especially in the case of sessions involving a large number of participants, using multicast protocols such as IP-multicast through APIs. It also enriches the type structures: we extend $\mathtt{p} \to \mathtt{p}' : k$ in the prefix of global types to $\mathtt{p} \to \mathtt{p}_1, .., \mathtt{p}_n : \{k_1, ..., k_n\}$ (with a practical adaptation such as group addressing), representing the multicast of a message to $\mathtt{p}_1, .., \mathtt{p}_n$ via channels $k_1, ..., k_n$ by participant $\mathtt{p}$, similarly we extend end-point session types to $\tilde{k}!\langle U\rangle$ from $k!\langle U\rangle$. Causality analysis remains the same by decomposing each multicasting prefix into its unicasting elements and considering causality for each of them. Our subsequent work [Bettini et al. 2008; Coppo et al. 2015b; Coppo et al. 2015a] uses multicasting and proves the progress properties in asynchronous multiparty sessions.

6.2. Related Work

There is a large literature on session types for both process calculi (in particular π-calculi) and programming languages. Below we discuss some of the most closely related works.

Asynchronous Session Types. Multiparty session types are based on message-order preserving asynchronous communication. Operational semantics of binary sessions based on asynchronous communication was first considered by [Neubauer and Thiemann 2004b]. Recently, [Gay and Vasconcelos 2009] studies the asynchronous version of binary sessions for an ML-like language [Vasconcelos et al. 2006]. In [Gay and Vasconcelos 2009], message queues are given two endpoint channels and a direction.

[Coppo et al. 2007] study asynchronous binary session types for Java, extending the previous work [Dezani-Ciancaglini et al. 2006], and prove progress by introducing an effect system. The resulting system does not allow interleaving sessions so that interactions involving more than two parties such as our examples in §2.3 cannot be represented. Our theorem establishes the progress property on multiple session channels, significantly enlarging the framework in [Coppo et al. 2007]. Recently, [Dezani-Ciancaglini et al. 2007] propose a typing system for progress in binary synchronous interleaving sessions. There, typable processes obey the partial orders of shared and session channels inferred during type-checking. Because of the use of global types, processes typed by our multiparty session typing do not have to follow such ordering; on the other hand, the system in [Dezani-Ciancaglini et al. 2007] does not require the simpleness condition (Definition 5.25). In [Dezani-Ciancaglini et al. 2007], a progress property is defined as follows: a typable process never reduces to a process which contains open sessions (this amounts to containing session channels) and which is irreducible in any inactive context (represented by another inactive process running in parallel). A combination of this and our multiparty session typing systems will enlarge typability, guaranteeing progress in many situations. See also §6.3.

The concurrent work done by [Bonelli and Compagnoni 2007], which is independently conceived and developed, studies a multiparty session typing for asynchronous communication. While treating the common topic, the technical direction of their work is different from that of the present work. Instead of global types, they solely use what we call (recursion-free) end-point types. In type checking, end-point types are projected to each binary session, so that type safety can be ensured using duality. Since we lose

sequencing information in this way, the progress property is not guaranteed. The use of global types in the present work leads to transparent treatment of type structures such as recursion, the guarantee of stronger behavioural properties such as progress, and (arguably) more intelligible description of multiparty interaction structures.

Global Description of Session Types. There are two recent works which studied global descriptions of sessions in the context of Web services and business protocols, by the present authors [Carbone et al. 2007; 2012] and by [Bhargavan et al. 2009]. Our work [Carbone et al. 2007; 2012] presented an *executable language* for directly describing Web interactions from a global viewpoint and provided the framework for projecting a description in the language to end-point processes. The use of global description for *types* and its associated theories have not been developed in [Carbone et al. 2007]. The type disciplines for the two (global and end-point) calculi studied in [Carbone et al. 2007] are based on binary synchronous session types, hence safety and progress for multiparty interactions are not considered. See also §6.3 for further extensions of [Carbone et al. 2007; 2012].

The work [Bhargavan et al. 2009] investigates approaches to cryptographically protecting session execution from both external attackers in networks and malicious session principals. Their session specification models an interaction sequence between two or more constituent *roles*, an abstraction of network peers. The description is given as a graph whose node represents a specific state of a role in a session, and whose edge denotes a dyadic communication and control flow. The purpose of the message flow graphs in [Bhargavan et al. 2009] is more to serve as a model for systems and programs than to offer a type discipline for programming languages.

First their work does not (aim to) present compositional typing rules for processes. Secondly their flow graphs do not (try to) represent such elements as local control flow (e.g. prefixing), channel-based communication and delegation. Third their operational structures may not be oriented towards type abstractions: for example their choice structures are based on transitions of flow graphs than additive structures realisable by branching and selection.

Integration of their and our approaches is an interesting further topic: for example, we may consider developing a runtime validation method for multiparty sessions using flow models induced by our global types.

With a similar intent to address secure implementation of multiparty sessions, the works in [Carbone and Guttman 2009b; 2009a] provide an abstract semantics for global types without parallel composition and recursion into the Strand Spaces model [Thayer et al. 1999]. The semantic function exploits a projection similar to ours.

Semantics of Delegation. For a simpler presentation, we used the operational semantics of delegation from [Honda et al. 1998] which demands that delegated channels do not occur in the receiver. This prevents a process from acting as two or more participants in the same session, which usually leads to a deadlock. The duplication check is easily implementable in a way analogous to the standard mechanism of firewalls. The more generous rule [Gay and Hole 2005; Yoshida and Vasconcelos 2007] allows substitution of session channels as in [RECV], which also satisfies type safety and progress through annotations on channels and types. This annotation extends the method in [Gay and Hole 2005; Yoshida and Vasconcelos 2007]: instead of polarities we use indices by participants to annotate each usage of channels. With this change the whole theories remain intact with exactly the same operational semantics and typing for programs. We study this delegation in [Bettini et al. 2008; Bejleri and Yoshida 2009].

Linear and Behavioural Types for Mobile Processes. Among many works on types for mobile processes, session type disciplines in general and the present work in par-

ticular are most closely related with linear/IO-typed π-calculi with causality information. The session type disciplines are related with linear and IO-typed π-calculi with causality information. The causality analysis in global types is partly inspired by the graph-based linear types developed in [Yoshida 1996; Yoshida et al. 2004] where ordering among multiple linear names (which correspond to session channels) guarantees deadlock-freedom of typed processes. Several works [Kobayashi 2006; Igarashi and Kobayashi 2004] study generalised forms of linear typing for guaranteeing different kinds of deadlock-freedom, incorporating synchronisations and locking.

A main difference of session type disciplines from these and other preceding works in this field is a notion of rigorously structured conversations and their direct type abstraction. See [Acciai and Boreale 2008; Dezani-Ciancaglini et al. 2007] for detailed discussions, including comparisons between the session-based and the behavioural-based ones [Yoshida 1996; Yoshida et al. 2004; Kobayashi 2006]; in [Acciai and Boreale 2008; Dezani-Ciancaglini et al. 2007; Bettini et al. 2008], structured session primitives help to give simpler typing systems for progress for binary sessions.

By raising the level of abstraction through the use of structured primitives such as separate session initiation, branching and recursion, session types can describe complex interaction structures more intelligibly and enable efficient type checking. These features would have direct applicability for the design of programming languages with communication [Hu et al. 2008; Carbone et al. 2007; 2012; Honda et al. 2007; Sackman and Eisenbach 2008; Pucella and Tov 2008; Scribble 2008].

One of the novelties of the present work is the introduction of global description as types and a use of their projection for type-checking. They offer a modular and systematic causality analysis rather than directly working on individual syntax and operational semantics, with adaptations to asynchronous and synchronous communications. Composability of multiple programs is transparent through projection of a common global type while complex syntax of types and typing are required in the traditional approach. To our knowledge, this method has not been investigated so far in the types of mobile processes.

Advanced Process Calculi and Types. Several process calculi for broadcasting have been investigated to model and analyse broadcasting networks including (recently) mobile ad-hoc networks, starting from Prasad's thesis [Prasad 2001]. Recent works focus on behavioural equivalences with lts [Merro 2007; Mezzetti and Sangiorgi 2006; Prasad 2006] and static analysis [Nanz et al. 2007] to investigate a number of different broadcasting. None of them studied the typing system and provided a strong progress guarantee as ensured by our session types. Our session types use a static participant information in the syntax and types. Recent advanced typing systems for location-based distributed processes [Hennessy 2007] use the similar notion for types $T@\mathtt{p}$, allowing dynamically instantiate locations into the capabilities using dependent type techniques. Since our aim is to prove the simplest extension of the original session types to multiparty, the static participants are enough even for delegations. It is a valuable further study to investigate a dynamic change of participant numbers when session initialisation (without explicitly declaring $\mathtt{p}$ in the syntax) by using channel dependent types [Mostrous and Yoshida 2007] or polymorphism.

Other Recent Service-Oriented Calculi. A vast amount of formal work for Service-Orientation has been done using process calculi and session types. The reader can refer two recent surveys [Dezani-Ciancaglini and de' Liguoro 2010; Castagna et al. 2011] for more comparisons. We focus on the most related recent work. Different approaches to the description of service-oriented multiparty communications are taken in [Bravetti and Zavattaro 2007; Bruni et al. 2008]. In [Bravetti and Zavattaro 2007], the global and local views of protocols are described using a synchronous CCS-based

calculus as a contract language, and testing-preorders to check subcontract compliance; [Bruni et al. 2008] proposes a distributed calculus which provides communications either inside sessions or inside locations, modelling merging running sessions. *Contracts* [Castagna and Padovani 2009] use a process-based specification of protocols, where conformance means Must-preorder (so that we can ensure liveness). The system in [Castagna and Padovani 2009] can type more processes than session types, thanks to the flexibility of process syntax for describing protocols. However, typable processes themselves in [Castagna and Padovani 2009] may not always satisfy the properties of session types such as progress: it is proved later by checking whether the type meets a certain form. Hence proving progress with contracts effectively requires an exploration of all possible paths (interleavings, choices) of a protocol.

The work [Caires and Vieira 2010] proposes a proof system which builds a well-founded ordering on events to enforce progress for processes of the Conversation Calculus [Vieira et al. 2008] where dynamic join and leave of participants are treated. These recent works do not treat a prescription of protocols given by the global types, with the efficient projection and type-checking, which can ensure strong safety properties. Our recent work [Deniélou and Yoshida 2011] extends a dynamic join and leaving mechanism based on the multiparty session types introducing a notion of *roles* which represent a unit of local behaviours.

6.3. Recent Works based on Multiparty Session Types

This subsection summarises works based on Multiparty Session Types published after the extended abstract [Honda et al. 2008a] of this article.

Theoretical Studies on Multiparty Session Types. Extensions of the original multiparty session types [Honda et al. 2008a] has been proposed, often motivated by use cases resulting from industry applications. Such extensions include: a subtyping for asynchronous multiparty session types enhancing efficiency [Mostrous et al. 2009], motivated by financial protocols and multicore algorithms; parametrised global types for parallel programming and Web service descriptions [Deniélou et al. 2012]; communication buffered analysis [Deniélou and Yoshida 2010]; extensions to the sumtype and its encoding [Nielsen et al. 2010] for describing Healthcare workflows; and exception handling for multiparty conversations [Capecchi et al. 2016] for Web services and financial protocols; a liveness-preserving refinement for multiparty session types [Padovani 2014b].

Multiparty session types can be extended with logical assertions following the design by contract framework [Bocchi et al. 2010]. This framework is enriched in [Bocchi et al. 2012] to handle stateful logical assertions, while [Chen and Honda 2012] offers more fine-grained property analysis for multiparty session types with these stateful assertions.

In [Deniélou and Yoshida 2011] roles are inhabited by an arbitrary number of participants which can dynamically join and leave a session. The paper [Swamy et al. 2011] shows that the multirole session types [Deniélou and Yoshida 2011] can be naturally represented in a dependent-typed language.

To enhance expressivity and flexibility of multiparty session types, the work [Demangeon and Honda 2012] proposes nested, higher-order multiparty session types and the work [Castagna et al. 2012] studies a generalisation of choices and parallelism. The paper [Carbone and Montesi 2013] directly types a global description language [Carbone et al. 2012] by multiparty session types without using local types. This direct approach can type processes which are untypable in the original multiparty session typing (i.e. the communication type system in this article). The paper [Montesi and

Yoshida 2013] extends the work in [Carbone and Montesi 2013] to compositional global description languages.

As another line of the study, we extend the multiparty session types to express temporal properties [Bocchi et al. 2014b]. In this work, the global times are enriched with time constraints, in a way similar to timed automata.

A type system enforcing a stronger correspondence between nondeterministic choices expressed in multiparty session types and the behaviour of processes involved in multiparty sessions has been investigated in [Bocchi et al. 2014a].

Progress and Session Interleaving. Multiparty session types are a convenient methodology for ensuring progress of systems of communicating processes. However, progress is only guaranteed within a *single* session [Honda et al. 2008a; Dezani-Ciancaglini and de' Liguoro 2010; Deniélou and Yoshida 2011], but not when multiple sessions are interleaved. The first papers considering progress for interleaved sessions required the nesting of sessions in Java [Dezani-Ciancaglini et al. 2006; Coppo et al. 2007]. These systems can guarantee progress for only one single active binary session. The work [Coppo et al. 2015b] develops a static interaction type system for global progress in dynamically interleaved and interfered multiparty sessions. A type inference algorithm for this system has been studied in [Coppo et al. 2013], although for finite types only. The work [Padovani 2014a, technical report] presents a type system for the linear π-calculus that can ensure progress even in presence of session interleaving, exploiting an encoding similar to that described in [Dardha et al. 2012] of sessions into the linear π-calculus. However, not *all* multiparty sessions can be encoded into well-typed linear π-calculus processes. In this respect, the richer structure of multiparty session types increases the range of systems for which non-trivial properties such as progress can be guaranteed.

Security. Enforcement of *integrity* properties in multiparty sessions, using session types, has been studied in [Bhargavan et al. 2009; Planul et al. 2009]. These papers propose a compiler which, given a multiparty session description, implements cryptographic protocols that guarantee session execution integrity.

The work [Capecchi et al. 2010] and in its extended version [Capecchi et al. 2014] propose a session type system for a calculus of multiparty sessions enriched with security levels, adding access control and secure information flow requirements in the typing rules, and show that this type system guarantees preservation of data confidentiality during session execution. In [Capecchi et al. 2015] this calculus is equipped with a monitored semantics, which blocks the execution of processes as soon as they attempt to leak information, raising an error.

Behavioural Semantics. Typed behavioural theory has been one of the central topics in the study of the π-calculus throughout its history, for example, reasoning about various encodings into the typed π-calculi [Pierce and Sangiorgi 1996; Yoshida 1996; Kouzapas et al. 2016]. In the context of typed bisimulations and reduction-closed theories, the work [Kouzapas and Yoshida 2014] shows that unique behavioural theories can be constructed based on the multiparty session types. The behavioural theory in [Kouzapas and Yoshida 2014] treats the mutual effects of multiple choreographic sessions which are shared among distributed participants as their common knowledge or agreements, reflecting the origin of choreographic frameworks [WS-CDL 2003]. These features related to multiparty session type discipline make the theory distinct from any type-based bisimulations in the literature and also applicable to a real choreographic usecase from a large-scale distributed system. This bisimulation is called *globally governed*, since it uses global multiparty specifications to regulate the conversational behaviour of distributed processes. It is an interesting future work to extend

this work towards more scalable session bisimulations for the eventful session types and the higher-order π-calculus studied in [Kouzapas et al. 2015].

Runtime Monitoring and Adaptations. Multiparty session types were originally developed to be used for static type checking of communicating processes. Via collaborations with Ocean Observatories Initiative [OOI 2015], it was discovered that the framework of multiparty session types can be naturally extended to runtime type checking (monitoring). A formulation of the runtime monitoring (dynamic or runtime type checking) is firstly proposed in [Chen et al. 2012]. Later the work [Bocchi et al. 2013] has formally proved its correctness and properties guaranteed by the runtime monitoring based on multiparty session types.

Works addressing adaptation for multiparty communications include [Dalla Preda et al. 2014] and [Coppo et al. 2014]. The paper [Dalla Preda et al. 2014] proposes a choreographic language for distributed applications. Adaptation follows a rule-based approach, in which all interactions, under all possible changes produced by the adaptation rules, proceed as prescribed by an abstract model. In [Coppo et al. 2014] a calculus based on global types, monitors and processes is introduced and adaptation is triggered after the execution of the communications prescribed by a global type, in reaction to changes of the global state.

Linkages with Other Frameworks. The work [Deniélou and Yoshida 2012] gives a linkage between communicating automata [Brand and Zafiropulo 1983] and a general graphical version of multiparty session types, proving a correspondence between the safety properties of communicating automata and multiparty session types. This work [Deniélou and Yoshida 2012] studies more detailed semantics for global and local types, relating with other frameworks such as model checking and logical verification for contracts [Villard 2011; Basu et al. 2012] (see [Deniélou and Yoshida 2012, §5] for detailed comparisons).

The paper [Deniélou and Yoshida 2013] studies the sound and complete characterisation of the multiparty session types in communicating automata (called *multiparty compatibility*) and applies the result to the synthesis of the multiparty session types. The inference of global types from a set of local types is also studied in [Lange and Tuosto 2012]. The techniques developed in [Deniélou and Yoshida 2013; Lange and Tuosto 2012] are extended to a synthesis of general graphical multiparty session types in [Lange et al. 2015]. This connection is extended to timed communicating automata [Krcál and Yi 2006]: the work [Bocchi et al. 2015] proposes general conditions of progress and non-zero properties of timed communicating automata at the top of multiparty compatibility.

The work [Fossati et al. 2014] studies the relationship of multiparty session types with Petri Nets. It proposes a conformance relation between global session nets and endpoint programs, and proves its safety.

A recent work [Carbone et al. 2015] studies a relationship with Linear Logic and multiparty session types along the line of [Wadler 2012; Caires and Pfenning 2010].

Implementations based on Multiparty Session Types. We are currently designing and implementing a modelling and specification language with multiparty session types [SAVARA 2010; Scribble 2008] in collaboration with some industrial partners [Honda et al. 2011; Honda et al. 2014]. This protocol language is called Scribble. An article [Yoshida et al. 2013] also explains the origin and recent development on Scribble.

Java protocol optimisation [Sivaramakrishnan et al. 2010] based on multiparty session types and generation of multiparty cryptographic protocols [Bhargavan et al. 2009] are also studied. The multiparty session type theory is applied to Healthcare

workflows [Henriksen et al. 2013]. Its prototype implementation (the multiparty session π-processes with sumtypes) is available from [Apims 2014].

Based on the runtime type checking theory, we are implementing a runtime monitoring [Demangeon et al. 2015; Hu et al. 2013; Neykova et al. 2013] under collaborations with Ocean Observatories Initiative [OOI 2015]. The work [Demangeon et al. 2015; Hu et al. 2013] allows interruptions in Scribble and proves the correctness of this extension. Further we generalise the Python implementation to the Actor framework [Neykova and Yoshida 2014]. In order to express temporal properties studied in timed multiparty session types [Bocchi et al. 2014b], the work [Neykova et al. 2014] extends Scribble with timed constrains and implements the runtime monitoring in Python.

We also apply the multiparty session types to high-performance parallel programming in C [Ng et al. 2012; Ng et al. 2012] and MPI [Ng and Yoshida 2014]. A parametrised version of Scribble [Ng and Yoshida 2014; Ng et al. 2013] based on the theory of parametrised multiparty session types [Deniélou et al. 2012] is developed. This extension, called Pabble, is used for automatically generating MPI parallel programs from sequential C code in [Ng et al. 2015].

7. CONCLUSION

One of the main open problems of the session types is whether binary sessions can be extended to n-party sessions and, if they can, what is their additional expressiveness and benefits. This paper answers the question affirmatively. The present theory can guarantee stronger conformance to stipulated conversation structures than binary sessions when a protocol involves more than two parties. We proposed a new efficient type checking system and proved type safety and progress, extended to multiparty interactions. The central technical underpinning of the present work is the introduction of global types, which offer an intuitive syntax for describing multiparty conversation structures from a global viewpoint; and the use of their projection for efficient typechecking, proposing a new effective methodology for programming multiparty interactions in distributed environments. Global types also offer a basis of a clean modular causal analysis systematically applicable to both synchronous and asynchronous communications, ensuring the progress and session fidelity.

There are several significant future topics on the theory and applications of the proposed theory. We are currently starting to use this generalised session type structure as one of the formal foundations for the following applications: for the next version of a web service description language (based on an idea from [WS-CDL 2003]) developed in Scribble from JBoss Red Hat [Scribble 2008], a message scheme for financial protocols, for a testable architecture, SAVARA from JBoss Red Hat [SAVARA 2010], for a specification for message middleware from AMQP [AMQP 2015] for a specification for large distributed systems from Ocean Observatories Initiative [OOI 2015], for software development life cycle from Zero Diviation Life Cycle (ZDLC) [zdl 2015]. In particular, we are currently designing and implementing a modelling and specification language with multiparty session types [Scribble 2008] for these standards with our industry collaborators. This consists of three layers: the first layer is a global type which corresponds to a signature of class models in UML; the second one is for conversation models where signatures and variables for multiple conversations are integrated; and the third layer includes extensions of the existing languages (such as Java [Hu et al. 2008; Hu et al. 2010; Ng et al. 2011]) which implement conversation models. Other future topics include tools assistance for the design and elaboration of global types; incorporation of typed exceptions to sessions; and integration of the type discipline with diverse specification concerns including security and monitoring for distributed messages by the assertional methods [Bocchi et al. 2010].

ACKNOWLEDGMENTS

We would like to thank Andi Bejleri for an early collaboration on this work.

REFERENCES

2015. Zero Deviation Lifecycle. http://www.zdlc.co. (2015).

Lucia Acciai and Michele Boreale. 2008. A Type System for Client Progress in a Service-Oriented Calculus. In *Concurrency, Graphs and Models (LNCS)*, Vol. 5065. Springer, Pisa, Italy, 642–658.

AMQP 2015. Advanced Message Queuing Protocol. http://www.iona.com/opensource/amqp/. (2015).

Apims 2014. Apims. (2014). http://thelas.dk/index.php?title=Apims.

Samik Basu, Tevfik Bultan, and Meriem Ouederni. 2012. Deciding Choreography Realizability. In *Symposium on Principles of Programming Languages*. ACM, Philadelphia, USA, 191–202.

Andi Bejleri and Nobuko Yoshida. 2009. Synchronous Multiparty Session Types. In *In Proceedings of Programming Languages Approaches to Concurrency and Communication-Centric Software (PLACES'08) (ENTCS)*, Vol. 241. Elsevier, Oslo, Norway, 3–33.

Lorenzo Bettini, Mario Coppo, Loris D'Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. 2008. Global Progress in Dynamically Interleaved Multiparty Sessions. In *International Conference on Concurrency Theory (LNCS)*, Vol. 5201. Springer, Toronto, Canada, 418–433.

Karthikeyan Bhargavan, Ricardo Corin, Pierre-Malo Deniélou, Cédric Fournet, and James Leifer. 2009. Cryptographic Protocol Synthesis and Verification for Multiparty Sessions. In *Computer Security Foundations Symposium*. IEEE, New York, USA, 124–140.

Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. 2013. Monitoring Networks through Multiparty Session Types. In *IFIP Joint International Conference on Formal Techniques for Distributed Systems (LNCS)*, Vol. 7892. Springer, Florence, Italy, 50–65.

Laura Bocchi, Romain Demangeon, and Nobuko Yoshida. 2012. A Multiparty Multi-Session Logic. In *7th International Symposium on Trustworthy Global Computing (LNCS)*, Vol. 8191. Springer, Newcastle upon Tyne, UK, 111–97.

Laura Bocchi, Kohei Honda, Emilio Tuosto, and Nobuko Yoshida. 2010. A theory of design-by-contract for distributed multiparty interactions. In *International Conference on Concurrency Theory (LNCS)*, Vol. 6269. Springer, Paris, France, 162–176.

Laura Bocchi, Julien Lange, and Nobuko Yoshida. 2015. Meeting Deadlines Together. In *International Conference on Concurrency Theory (LIPIcs)*, Vol. 42. Schloss Dagstuhl, Madrid, Spain, 283–296.

Laura Bocchi, Hernán C. Melgratti, and Emilio Tuosto. 2014a. Resolving Non-determinism in Choreographies. In *European Symposium on Programming (LNCS)*, Vol. 8410. Springer, Grenoble, France, 493–512.

Laura Bocchi, Weizhen Yang, and Nobuko Yoshida. 2014b. Timed Multiparty Session Types. In *International Conference on Concurrency Theory (LNCS)*, Vol. 8704. Springer, Rome, Italy, 419–434.

Eduardo Bonelli, Adriana Compagnoni, and Elsa Gunter. 2005. Correspondence Assertions for Process Synchronization in Concurrent Communications. *Journal of Functional Programming* 15, 2 (2005), 219–248.

Eduardo Bonelli and Adriana B. Compagnoni. 2007. Multipoint Session Types for a Distributed Calculus. In *Trustworthy Global Computing (LNCS)*, Vol. 4912. Springer, Sophia-Antipolis, France, 240–256.

BPMNC 2012. Business Process Model and Notation 2.0 Choreography. (2012). http://en.bpmn-community.org/tutorials/34/.

Daniel Brand and Pitro Zafiropulo. 1983. On Communicating Finite-State Machines. *Journal of ACM* 30 (April 1983), 323–342. Issue 2.

Mario Bravetti and Gianluigi Zavattaro. 2007. Towards a Unifying Theory for Choreography Conformance and Contract Compliance. In *Software Composition (LNCS)*, Vol. 4829. Springer, Braga, Portugal, 34–50.

Roberto Bruni, Ivan Lanese, Hernan Melgratti, and Emilio Tuosto. 2008. Multiparty Sessions in SOC. In *Coordination Models and Languages (LNCS)*, Vol. 5052. Springer, Oslo, Norway, 67–82.

Luís Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In *International Conference on Concurrency Theory (LNCS)*, Vol. 6269. Springer, Paris, France, 222–236.

Luís Caires and Hugo Torres Vieira. 2010. Conversation types. *Theoretical Computer Science* 411, 51-52 (2010), 4399–4440.

Sara Capecchi, Ilaria Castellani, and Mariangiola Dezani-Ciancaglini. 2014. Typing Access Control and Secure Information Flow in Sessions. *Information and Computation* 238 (2014), 68–105.

Sara Capecchi, Ilaria Castellani, and Mariangiola Dezani-Ciancaglini. 2015. Information Flow Safety in Multiparty Sessions. (2015). To appear.

Sara Capecchi, Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and Tamara Rezk. 2010. Session Types for Access and Information Flow Control. In *International Conference on Concurrency Theory (LNCS)*, Vol. 6269. Springer, Paris, France, 237–252.

Sara Capecchi, Elena Giachino, and Nobuko Yoshida. 2016. Global Escape in Multiparty Sessions. *Mathematical Structures in Computer Science* 26, 2 (2016), 156–205.

Marco Carbone and Joshua Guttman. 2009a. Choreographies with Secure Boxes and Compromised Principals. In *Interaction and Concurrency Experience - Structured Interactions (EPTCS)*, Vol. 12. Bologna, Italy, 1–16.

Marco Carbone and Joshua Guttman. 2009b. Execution Models for Choreographies and Cryptoprotocols. In *Workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software (EPTCS)*, Vol. 17. York, UK, 31–42.

Marco Carbone, Kohei Honda, and Nobuko Yoshida. 2007. Structured Communication-Centred Programming for Web Services. In *European Symposium on Programming (LNCS)*, Vol. 4421. Springer, Braga, Portugal, 2–17.

Marco Carbone, Kohei Honda, and Nobuko Yoshida. 2008. Structured Interactional Exceptions in Session Types. In *International Conference on Concurrency Theory (LNCS)*, Vol. 5201. Springer, Toronto, Canada, 402–417.

Marco Carbone, Kohei Honda, and Nobuko Yoshida. 2012. Structured Communication-Centered Programming for Web Services. *ACM Transactions on Programming Languages and Systems* 34, 2 (2012), 8.

Marco Carbone, Kohei Honda, Nobuko Yoshida, Robin Milner, Gary Brown, and Steve Ross-Talbot. 2006. A Theoretical Basis of Communication-Centred Concurrent Programming. http://www.w3.org/2002/ws/chor/. (2006).

Marco Carbone and Fabrizio Montesi. 2013. Deadlock-freedom-by-design: Multiparty Asynchronous Global Programming. In *Symposium on Principles of Programming Languages*. ACM, Rome, Italy, 263–274.

Marco Carbone, Fabrizio Montesi, Carsten Schrmann, and Nobuko Yoshida. 2015. Multiparty Session Types as Coherence Proofs. In *International Conference on Concurrency Theory (LIPIcs)*, Vol. 42. Schloss Dagstuhl, Madrid, Spain, 412–426.

Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, and Luca Padovani. 2011. On Global Types and Multiparty Sessions. In *International Conference on Formal Methods for Open Object-based Distributed Systems (LNCS)*, Vol. 6722. Springer, Reykjavik, Iceland, 1–28.

Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, and Luca Padovani. 2012. On Global Types and Multi-Party Session. *Logical Methods in Computer Science* 8, 1 (2012), 24.

Giuseppe Castagna and Luca Padovani. 2009. Contracts for Mobile Processes. In *International Conference on Concurrency Theory (LNCS)*. Springer, Bologna, Italy, 211–228.

Tzu-Chun Chen, Laura Bocchi, Pierre-Malo Deniélou, Kohei Honda, and Nobuko Yoshida. 2012. Asynchronous Distributed Monitoring for Multiparty Session Enforcement. In *Trustworthy Global Computing (LNCS)*, Vol. 7173. Springer, Newcastle upon Tyne, UK, 25–45.

Tzu-Chun Chen and Kohei Honda. 2012. Specifying Stateful Asynchronous Properties for Distributed Programs. In *International Conference on Concurrency Theory (LNCS)*, Vol. 7454. Springer, Newcastle upon Tyne, UK, 209–224.

Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani, and Nobuko Yoshida. 2013. Inference of Global Progress Properties for Dynamically Interleaved Multiparty Sessions. In *Coordination Models and Languages (LNCS)*, Vol. 7890. Springer, Florence, Italy, 45–59.

Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani, and Nobuko Yoshida. 2015a. A Gentle Introduction to Multiparty Asynchronous Session Types. In *SFM-15:MP (LNCS)*, Vol. 9104. Springer, Bertinoro, Italy, 146–178.

Mario Coppo, Mariangiola Dezani-Ciancaglini, and Betti Venneri. 2014. Self-Adaptive Multiparty Sessions. *Service Oriented Computing and Applications* 9, 3-4 (2014), 249–268.

Mario Coppo, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. 2007. Asynchronous Session Types and Progress for Object-Oriented Languages. In *IFIP International Conference on Formal Methods for Open Object-based Distributed Systems (LNCS)*, Vol. 4468. Springer, Paphos, Cyprus, 1–31.

Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. 2015b. Global Progress for Dynamically Interleaved Multiparty Sessions. *Mathematical Structures in Computer Science* 26, 2 (2015), 238–302.

Mila Dalla Preda, Saverio Giallorenzo, Ivan Lanese, Jacopo Mauro, and Maurizio Gabbrielli. 2014. AIOCJ: A Choreographic Framework for Safe Adaptive Distributed Applications. In *International Conference on Software Language Engineering (LNCS)*, Vol. 8706. Springer, Västerås, Sweden, 161–170.

Ornela Dardha, Elena Giachino, and Davide Sangiorgi. 2012. Session Types Revisited. In *International Symposium on Principles and Practice of Declarative Programming*. ACM Press, Leuven, Belgium, 139–150.

Romain Demangeon and Kohei Honda. 2012. Nested Protocols in Session Types. In *International Conference on Concurrency Theory (LNCS)*, Vol. 7454. Springer, Newcastle upon Tyne, UK, 272–286.

Romain Demangeon, Kohei Honda, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. 2015. Practical Interruptible Conversations: Distributed Dynamic Verification with Multiparty Session Types and Python. *Formal Methods in System Design* 46, 3 (2015), 197–225.

Pierre-Malo Deniélou and Nobuko Yoshida. 2010. Buffered Communication Analysis in Distributed Multiparty Sessions. In *International Conference on Concurrency Theory (LNCS)*, Vol. 6269. Springer, Paris, France, 343–357.

Pierre-Malo Deniélou and Nobuko Yoshida. 2011. Dynamic Multirole Session Types. In *Symposium on Principles of Programming Languages*. ACM, Austin, USA, 435–446.

Pierre-Malo Deniélou and Nobuko Yoshida. 2012. Multiparty Session Types Meet Communicating Automata. In *European Symposium on Programming (LNCS)*, Helmut Seidl (Ed.), Vol. 7211. Springer, Tallin, Estonia, 194–213.

Pierre-Malo Deniélou and Nobuko Yoshida. 2013. Multiparty Compatibility in Communicating Automata: Characterisation and Synthesis of Global Session Types. In *International Colloquium on Automata, Languages and Programming (LNCS)*, Vol. 7966. Springer, Riga, Latvia, 174–186.

Pierre-Malo Deniélou, Nobuko Yoshida, Andi Bejleri, and Raymond Hu. 2012. Parameterised Multiparty Session Types. *Logical Methods in Computer Science* 8, 4 (2012).

Mariangiola Dezani-Ciancaglini and Ugo de' Liguoro. 2010. Sessions and Session Types: an Overview. In *International Workshop on Web Services and Formal Methods (LNCS)*, Vol. 6194. Springer, Bologna, Italy, 1–28.

Mariangiola Dezani-Ciancaglini, Ugo de'Liguoro, and Nobuko Yoshida. 2007. On Progress for Structured Communications. In *Trustworthy Global Computing (LNCS)*, Vol. 4912. Springer, Sophia-Antipolis, France, 257–275.

Mariangiola Dezani-Ciancaglini, Sophia Drossopoulou, Dimitris Mostrous, and Nobuko Yoshida. 2009. Objects and session types. *Information and Computation* 207, 5 (2009), 595–641.

Mariangiola Dezani-Ciancaglini, Dimitris Mostrous, Nobuko Yoshida, and Sophia Drossopoulou. 2006. Session Types for Object-Oriented Languages. In *European Conference on Object-Oriented Programming (LNCS)*, Vol. 4067. Springer, Nantes, France, 328–352.

Manuel Fähndrich, Mark Aiken, Chris Hawblitzel, Orion Hodson, Galen C. Hunt, James R. Larus, , and Steven Levi. 2006. Language Support for Fast and Reliable Message-based Communication in Singularity OS. In *EuroSys2006 (ACM SIGOPS)*. ACM Press, Leuven, Belgium, 177–190.

Luca Fossati, Raymond Hu, and Nobuko Yoshida. 2014. Multiparty Session Nets. In *Trustworthy Global Computing (LNCS)*, Vol. 8902. Springer, Rome, Italy, 112–127.

Pablo Garralda, Adriana Compagnoni, and Mariangiola Dezani-Ciancaglini. 2006. BASS: Boxed Ambients with Safe Sessions. In *International Symposium on Principles and Practice of Declarative Programming*. ACM Press, Venice, Italy, 61–72.

Simon Gay. 2008. Bounded Polymorphism in Session Types. *MSCS* 18 (2008), 895–930.

Simon Gay and Malcolm Hole. 2005. Subtyping for Session Types in the Pi-Calculus. *Acta Informatica* 42, 2/3 (2005), 191–225.

Simon Gay and Vasco T. Vasconcelos. 2009. Linear Type Theory for Asynchronous Session Types. *Journal of Functional Programming* (2009).

Simon Gay, Vasco T. Vasconcelos, António Ravara, Nils Gesbert, and Alexandre Z. Caldeira. 2010. Modular Session Types for Distributed Object-Oriented Programming. In *Symposium on Principles of Programming Languages*. ACM, Madrid, Spain, 299–312.

Jean-Yves Girard. 1987. Linear Logic. *Theoretical Computer Science* 50 (1987), 1–102.

Matthew Hennessy. 2007. *A Distributed Pi-Calculus*. Cambridge University Press.

Anders Henriksen, Lasse Nielsen, Thomas Hildebrandt, Nobuko Yoshida, , and Fritz Henglein. 2013. Trustworthy Pervasive Healthcare Services via Multi-party Session Type. In *Foundations of Health Information Engineering and Systems (LNCS)*, Vol. 7789. Paris, France, 124–141.

Kohei Honda. 1993. Types for Dyadic Interaction. In *International Conference on Concurrency Theory (LNCS)*, Eike Best (Ed.), Vol. 715. Springer, Hildesheim, Germany, 509–523.

Kohei Honda, Raymond Hu, Rumyana Neykova, Tzu-Chun Chen, Romain Demangeon, Pierre-Malo Deniélou, and Nobuko Yoshida. 2014. Structuring Communication with Session Types. In *Concurrent Objects and Beyond (LNCS)*, Vol. 8665. Springer, 105–127.

Kohei Honda, Aybek Mukhamedov, Gary Brown, Tzu-Chun Chen, and Nobuko Yoshida. 2011. Scribbling Interactions with a Formal Foundation. In *International Conference on Distributed Computing and Internet Technology (LNCS)*, Vol. 6536. Springer, Bhubaneswar, India, 55–75.

Kohei Honda and Mario Tokoro. 1991. An Object Calculus for Asynchronous Communication. In *European Conference on Object-Oriented Programming*, Vol. 512. Geneva, Switzerland, 133–147.

Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Disciplines for Structured Communication-based Programming. In *European Symposium on Programming (LNCS)*, Vol. 1381. Springer-Verlag, Lisbon, Portugal, 22–138.

Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2007. Web Services, Mobile Processes and Types. *The Bulletin of the European Association for Theoretical Computer Science* February, 91 (2007), 165–185.

Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008a. Multiparty Asynchronous Session Types. In *Symposium on Principles of Programming Languages*. ACM, San Francisco, USA, 273–284.

Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008b. Multiparty Asynchronous Session Types. (2008). Web page. http://www.doc.ic.ac.uk/~yoshida/multiparty.

Raymond Hu, Dimitrious Kouzapas, Oliver Pernet, Nobuko Yoshida, and Kohei Honda. 2010. Type-Safe Eventful Sessions in Java. In *European Conference on Object-Oriented Programming (LNCS)*, Vol. 6183. Springer, Maribor, Slovenia, 329–353.

Raymond Hu, Rumyana Neykova, Nobuko Yoshida, and Romain Demangeon. 2013. Practical Interruptible Conversations: Distributed Dynamic Verification with Session Types and Python. In *Runtime Verification (LNCS)*, Vol. 8174. Springer, Rennes, France, 148–130.

Raymond Hu, Nobuko Yoshida, and Kohei Honda. 2008. Session-Based Distributed Programming in Java. In *European Conference on Object-Oriented Programming*, Vol. 5142. Springer, Paphos, Cyprus, 516–541.

Atsushi Igarashi and Naoki Kobayashi. 2004. A Generic Type System for the Pi-Calculus. *Theoretical Computer Science* 311, 1-3 (2004), 121–163.

International Telecommunication Union. 1996. Recommendation Z.120: Message Sequence Chart. (1996).

Naoki Kobayashi. 2006. A New Type System for Deadlock-Free Processes. In *International Conference on Concurrency Theory (LNCS)*, Vol. 4137. Bonn, Germany, 233–247.

Dimitrios Kouzapas, Jorge A. Perez, and Nobuko Yoshida. 2015. Characteristic Bisimulations for Higher-Order Session Processes. In *International Conference on Concurrency Theory (LIPIcs)*, Vol. 42. Schloss Dagstuhl, Madrid, Spain, 398–411.

Dimitrios Kouzapas and Nobuko Yoshida. 2014. Globally Governed Session Semantics. *Logical Methods in Computer Science* 10, 4 (2014).

Dimitrios Kouzapas, Nobuko Yoshida, Raymond Hu, and Kohei Honda. 2016. On Asynchronous Eventful Session Semantics. *Mathematical Structures in Computer Science* 26, 2 (2016), 303–364.

Pavel Krcál and Wang Yi. 2006. Communicating Timed Automata: The More Synchronous, the More Difficult to Verify. In *Computer Aided Verification (LNCS)*. Springer, Seattle, USA, 249–262.

Leslie Lamport. 1978. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM* 21, 7 (July 1978), 558–564.

Julien Lange and Emilio Tuosto. 2012. Synthesising Choreographies from Local Session Types. In *International Conference on Concurrency Theory (LNCS)*, Vol. 7454. Springer, Newcastle upon Tyne, UK, 225–239.

Julien Lange, Emilio Tuosto, and Nobuko Yoshida. 2015. From Communicating Machines to Graphical Choreographies. In *Symposium on Principles of Programming Languages*. ACM Press, Mumbai, India, 221–232.

Massimo Merro. 2007. An Observational Theory for Mobile Ad Hoc Networks. In *Electronic Notes in Theoretical Computer Science*, Vol. 172. Elsevier, 275–293.

Nicola Mezzetti and Davide Sangiorgi. 2006. Towards a Calculus For Wireless Systems. In *Electronic Notes in Theoretical Computer Science*, Vol. 158. Elsevier, 331–353.

Leonardo Gaetano Mezzina. 2008. How to Infer Finite Session Types in a Calculus of Services and Sessions. In *Coordination Models and Languages (LNCS)*, Vol. 5052. Springer, Oslo, Norway, 216–231.

Fabrizio Montesi and Nobuko Yoshida. 2013. Compositional Choreographies. In *International Conference on Concurrency Theory (LNCS)*, Vol. 8052. Springer, Buenos Aires, Argentina, 439–425.

Dimitris Mostrous and Nobuko Yoshida. 2007. Two Session Typing Systems for Higher-Order Mobile Processes. In *Typed Lambda Calculi and Applications (LNCS)*, Vol. 4583. Springer, Paris, France, 321–335.

Dimitris Mostrous and Nobuko Yoshida. 2009. Session-Based Communication Optimisation for Higher-Order Mobile Processes. In *Typed Lambda Calculi and Applications (LNCS)*, Pierre-Louis Curien (Ed.), Vol. 5608. Springer, Brasilia, Brazil, 203–218.

Dimitris Mostrous, Nobuko Yoshida, and Kohei Honda. 2009. Global Principal Typing in Partially Commutative Asynchronous Sessions. In *European Symposium on Programming (LNCS)*, Vol. 5502. Springer, York, UK, 316–332.

Sebastian Nanz, Flemming Nielson, and Hanne Riis Nielson. 2007. Topology-Dependent Abstractions of Broadcast Networks. In *International Conference on Concurrency Theory*. Lisbon, Portugal, 226–240.

Matthias Neubauer and Peter Thiemann. 2004a. An Implementation of Session Types. In *Practical Aspects of Declarative Languages (LNCS)*, Vol. 3057. Springer, Dallas, USA, 56–70.

Matthias Neubauer and Peter Thiemann. 2004b. Session Types for Asynchronous Communication. (2004). Universität Freiburg.

Rumyana Neykova, Laura Bocchi, and Nobuko Yoshida. 2014. Timed Runtime Monitoring for Multiparty Conversations. In *Workshop on Behavioural Types (EPTCS)*, Vol. 162. Rome, Italy, 19–26.

Rumyana Neykova and Nobuko Yoshida. 2014. Multiparty Session Actors. In *Coordination Models and Languages (LNCS)*, Vol. 8459. Springer, Berlin, Germany, 131–146.

Rumyana Neykova, Nobuko Yoshida, and Raymond Hu. 2013. SPY: Local Verification of Global Protocols. In *Runtime Verification (LNCS)*, Vol. 8174. Springer, Rennes, France, 363–358.

Nicholas Ng, Jose G.F. Coutinho, and Nobuko Yoshida. 2015. Protocols by Default: Safe MPI Code Generation based on Session Types. In *Compiler Construction (LNCS)*. Springer, London, UK, 212–232.

Nicholas Ng and Nobuko Yoshida. 2014. Pabble: Parameterised Scribble. *Service Oriented Computing and Applications* 9, 3-4 (2014), 1–16.

Nicholas Ng, Nobuko Yoshida, and Kohei Honda. 2012. Multiparty Session C: Safe Parallel Programming with Message Optimisation. In *TOOLS (LNCS)*, Vol. 7304. Springer, Prague, Czech Republic, 202–218.

Nicholas Ng, Nobuko Yoshida, and Wayne Luk. 2013. Scalable Session Programming for Heterogeneous High-Performance Systems. In *International Conference on Software Engineering and Formal Methods (LNCS)*, Vol. 8368. Springer, Madrid, Spain, 82–98.

Nicholas Ng, Nobuko Yoshida, Xin Yu Niu, Kuen Hung Tsoi, and Wayne Luk. 2012. Session Types: Towards Safe and Fast Reconfigurable Programming. *SIGARCH CAN* 40 (2012), 22–27. Issue 5.

Nicholas Ng, Nobuko Yoshida, Olivier Pernet, Raymond Hu, and Yiannos Kryftis. 2011. Safe Parallel Programming with Session Java. In *Coordination Models and Languages (LNCS)*, Vol. 6721. Springer, Reykjavik, Iceland, 110–126.

Lasse Nielsen, Nobuko Yoshida, and Kohei Honda. 2010. Multiparty Symmetric Sum Types. In *Expressiveness in Concurrency (EPTCS)*, Vol. 41. Paris, France, 121–135.

OOI 2015. Ocean Observatories Initiative. http://www.oceanleadership.org/programs-and-partnerships/ocean-observing/ooi/. (2015).

Luca Padovani. 2014a. Deadlock and Lock Freedom in the Linear π-Calculus. In *Computer Science Logic and Logic in Computer Science*. ACM Press, Vienna, Austria, 72:1–72:10.

Luca Padovani. 2014b. Fair Subtyping for Multi-Party Session Types. *Mathematical Structures in Computer Science* (2014), 1–41.

B. Pierce and D. Sangiorgi. 1996. Typing and Subtyping for Mobile Processes. *Mathematical Structures in Computer Science* 6, 5 (1996), 409–454.

Benjamin C. Pierce. 2002. *Types and Programming Languages*. MIT Press.

Jérémy Planul, Ricardo Corin, and Cédric Fournet. 2009. Secure Enforcement for Global Process Specifications. In *International Conference on Concurrency Theory (LNCS)*, Vol. 5710. Springer, Bologna, Italy, 511–526.

K.V.S. Prasad. 2001. Broadcast Calculus Interpreted in CCS upto Bisimulation. In *Electronic Notes in Theoretical Computer Science*, Vol. 52. Elsevier, 83–100. Issue 1.

K.V.S. Prasad. 2006. A Prospectus for Mobile Broadcasting Systems. In *Electronic Notes in Theoretical Computer Science*, Vol. 162. Elsevier, 295–300.

Riccardo Pucella and Jesse Tov. 2008. Haskell Session Types with (Almost) No Class. In *Haskell Symposium*. ACM SIGPLAN, Victoria, Canada.

Matthew Sackman and Susan Eisenbach. 2008. Session Types in Haskell. (2008). draft.

SAVARA 2010. SAVARA JBoss Project. http://www.jboss.org/savara. (2010).

Bruce Schneier. 1993. *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. John Wiley & Sons, Inc.

Scribble. 2008. Scribble Project. (2008). www.scribble.org.

K. C. Sivaramakrishnan, Karthik Nagaraj, Lukasz Ziarek, and Patrick Eugster. 2010. Efficient Session Type Guided Distributed Interaction. In *Coordination Models and Languages (LNCS)*, Vol. 6116. Springer, Amsterdam, Holland, 152–167.

Stephen Sparkes. 2006. Conversation with Steve Ross-Talbot. *ACM Queue* 4, 2 (March 2006).

Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. 2011. Secure distributed programming with value-dependent types. In *International Conference on Functional programming*. IEEE, Tokyo, Japan, 266–278.

Kaku Takeuchi, Kohei Honda, and Makoto Kubo. 1994. An Interaction-based Language and its Typing System. In *Parallel Architectures and Languages Europe (LNCS)*, Vol. 817. Springer-Verlag, Athens, Greece, 398–413.

F. Javier Thayer, Jonathan C. Herzog, and Joshua D. Guttman. 1999. Strand Spaces: Proving Security Protocols Correct. *Journal of Computer Security* 7, 2/3 (1999), 191–230.

Vasco T. Vasconcelos, Simon Gay, and António Ravara. 2006. Typechecking a Multithreaded Functional Language with Session Types. *Theoretical Computer Science* 368, 1–2 (2006), 64–87.

Hugo Torres Vieira, Luís Caires, and João Costa Seco. 2008. The Conversation Calculus: A Model of Service-Oriented Computation. In *European Symposium on Programming (LNCS)*, Vol. 4960. Springer, Budapest, Hungary, 269–283.

Jules Villard. 2011. *Heaps and Hops*. Ph.D. Dissertation. ENS Cachan.

Phil Wadler. 2012. Proposition as Sessions. In *International Conference on Functional Programming*. IEEE, Copenhagen, Denmark, 273–286.

WS-CDL. 2003. Web Services Choreography Working Group. http://www.w3.org/2002/ws/chor/. (2003).

Nobuko Yoshida. 1996. Graph Types for Monadic Mobile Processes.. In *Foundations of Software Technology and Theoretical Computer Science (LNCS)*, Vol. 1180. Springer, Hyderabad, India, 371–386.

Nobuko Yoshida, Martin Berger, and Kohei Honda. 2004. Strong Normalisation in the π-Calculus. *Information and Computation* 191(2004) (2004), 145–202.

Nobuko Yoshida, Raymond Hu, Rumyana Neykova, and Nicholas Ng. 2013. The Scribble Protocol Language. In *Trustworthy Global Computing (LNCS)*, Vol. 8358. Springer, Buenos Aires, Argentina, 22–41.

Nobuko Yoshida and Vasco Thudichum Vasconcelos. 2007. Language Primitives and Type Discipline for Structured Communication-Based Programming Revisited: Two Systems for Higher-Order Session Communication. *Electronic Notes Theoretical Computer Science* 171, 4 (2007), 73–93.

Nobuko Yoshida, Vasco Thudichum Vasconcelos, Hervé Paulino, and Kohei Honda. 2008. Session-Based Compilation Framework for Multicore Programming. In *International Symposium on Formal Methods for Components and Objects (LNCS)*, Vol. 5751. Springer, Sophia Antipolis, France, 226–246.

Received January 2009; revised February 2013 and August 2015; accepted September 2015

A. PROOF OF PROPOSITION 3.13

Below the proofs of both (1) and (2) induce concrete algorithms. Global types are generally treated as regular trees (except e.g. when we consider substitution). We first introduce the following notation.

NOTATION A.1.

(i) In the following we write $G(0)$, $G(1)$, ..., $G(n)$, ... for the result of n-times unfolding of each recursion in G. For example if G is $\mu \mathbf{t}.G'$ and this is the only recursion in G, then $G(0)$ is given as $G'[\mathsf{end}/\mathbf{t}]$, $G(1)$ is given as $G'[G(0)/\mathbf{t}]$ and, for each n, $G(n+1)$ is given as $G'[G(n)/\mathbf{t}]$. If G contains more than one recursion we perform the unfolding of each of its recursions. For convenience we set $G(-1)$ to be the empty graph.

(ii) Observing each $G(n+1)$ is the result of adding zero or more unfoldings to $G(n)$, so that $G(n+1)$ contains the exact copy of $G(n)$, we write $G(n+1)\backslash G(n)$ to denote the newly added (unfolded) part of $G(n+1)$.

(iii) Given a node n in $G(m+1)\backslash G(m)$, we can jump back from n once to reach its "original" in $G(m)\backslash G(m-1)$ (which is $G(0)$ if $m=0$). This exact copy of n which was created "one unfolding ago", is called the *one-time folding of* n, or simply the *folding of* n. In

the same way we define *the i-th folding of* n which is in $G(m-i+1)\backslash G(m-i)$ (which is $G(0)$ if $i = m+1$). Note there are $m+1$ such "foldings" of n in $G(m+1)\backslash G(m)$.

Proof of (1). Below we say there are input/output dependencies from n_1 to n_2 when there is an input dependency *and* an output dependency from n_1 to n_2.

Claim. (A) Suppose $\mathsf{n}_{1,2}$ and their respective i-th foldings $\mathsf{n}'_{1,2}$ are in $G(m)$. Then there are both input/output dependencies from n_1 to n_2 iff there are both input/output dependencies from n'_1 to n'_2. **(B)** Let n' be the folding of n. Then there is always both input and output dependencies from n' to n.

PROOF OF CLAIM. (A) is immediate since the graph structure of the foldings is identical to that of the originals (i.e. we can simply "fold" the original two onto their foldings and all prefix relations coincide). (B) is obvious since there always exist both II and OO dependencies by the definition of linearity. □

We now prove the statement. Fix a global type G and assume $G(1)$ is linear. We show by induction on n ($n \geq 1$) that each $G(n)$ is linear. Henceforth we ignore nodes in carried types.

Base step. This is linearity of $G(1)$ which is the assumption itself.

Induction Step. Suppose $G(n)$ is linear. Then take two nodes n_1 and n_2 in $G(n+1)$ (but not in carried types) which happen to share a common channel. We show there are input/output dependencies from n_1 to n_2, or the same holds in the reverse direction. We say such $\mathsf{n}_{1,2}$ are *conflict-free* for brevity. We do case analysis depending on the position of these nodes in $G(m+1)$.

(i) If $\mathsf{n}_{1,2}$ are in $G(n)$ then they already have input/output dependencies by induction hypothesis.

(ii) If n_1 is in $G(n)\backslash G(n-1)$ and n_2 is in $G(n+1)\backslash G(n)$ then take their two foldings say n'_1 and n'_2 respectively. By induction hypothesis they are conflict-free by a pair of dependency chains. By Claim A we are done.

(iii) If n_1 is in $G(n-i)$ ($i \geq 1$) and n_2 is in $G(n+1)\backslash G(n)$ then take the folding of n_2 say n'_2 which is in $G(n)$. By induction we know n_1 and n'_2 are conflict-free.

By Claim B, there are both input and output dependencies from n_2 to n'_2. Thus we have both input and output dependencies from n_1 to n'_2 and n'_2 to n_2 (hence n_1 to n_2). Now we connect these chains and we are done. □

B. FULL TYPING RULES FOR RUNTIME PROCESSES

This appendix first presents the full typing rules except those for expressions.

$$\frac{\Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : (G \upharpoonright 1)@1 \quad \{1, \ldots, n\} = \mathsf{pid}(G) \quad |\tilde{s}| = \mathsf{sid}(G)}{\Gamma \vdash_\emptyset \overline{a}_{[2..n]}(\tilde{s}).P \triangleright \Delta} \quad [\textsc{Mcast}]$$

$$\frac{\Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : (G \upharpoonright \mathtt{p})@\mathtt{p} \quad \mathtt{p} \in \mathsf{pid}(G) \quad |\tilde{s}| = \mathsf{sid}(G)}{\Gamma \vdash_\emptyset a_{[\mathtt{p}]}(\tilde{s}).P \triangleright \Delta} \quad [\textsc{Macc}]$$

$$\frac{\Gamma \vdash e_j : S_j \qquad \Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash_\emptyset s[k]!\langle \tilde{e} \rangle; P \triangleright \Delta, \tilde{s} : k!\langle \tilde{S} \rangle; T@\mathtt{p}} \qquad \frac{\Gamma, \tilde{x} : \tilde{S} \vdash P_\emptyset \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash_\emptyset s[k]?(\tilde{x}); P \triangleright \Delta, \tilde{s} : k?\langle \tilde{S} \rangle; T@\mathtt{p}} \quad [\textsc{Send}], [\textsc{Rcv}]$$

$$\frac{\Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash_\emptyset s[k]!\langle\langle \tilde{t} \rangle\rangle; P \triangleright \Delta, \tilde{s} : k!\langle T'@\mathtt{p}' \rangle; T@\mathtt{p}, \ \tilde{t} : T'@\mathtt{p}'} \qquad \frac{\Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : T@\mathtt{p}, \ \tilde{t} : T'@\mathtt{p}'}{\Gamma \vdash_\emptyset s[k]?((\tilde{t})); P \triangleright \Delta, \tilde{s} : k?\langle T'@\mathtt{p}' \rangle; T@\mathtt{p}} \quad [\textsc{Deleg}],[\textsc{SRec}]$$

$$\frac{\Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : T_j@\mathtt{p} \qquad j \in I}{\Gamma \vdash_\emptyset s[k] \lhd l; P \triangleright \Delta, \tilde{s} : k \oplus \{l_i : T_i\}_{i \in I}@\mathtt{p}} \qquad \frac{\Gamma \vdash_\emptyset P_i \triangleright \Delta, \tilde{s} : T_i@\mathtt{p} \qquad \forall i \in I}{\Gamma \vdash_\emptyset \ s[k] \rhd \{l_i : P_i\}_{i \in I} \triangleright \Delta, \tilde{s} : k \,\&\{l_i : T_i\}_{i \in I}@\mathtt{p}} \quad [\textsc{Sel}],[\textsc{Branch}]$$

$$\frac{\Gamma \vdash_\emptyset e \triangleright \mathsf{bool} \qquad \Gamma \vdash P \triangleright \Delta \qquad \Gamma \vdash Q \triangleright \Delta}{\Gamma \vdash_\emptyset \mathsf{if}\ e\ \mathsf{then}\ P\ \mathsf{else}\ Q \triangleright \Delta} \quad [\textsc{If}]$$

$$\frac{\Gamma \vdash P \triangleright_{\tilde{t}_1} \Delta \quad \Gamma \vdash_{\tilde{t}_2} Q \triangleright \Delta' \quad \tilde{t}_1 \cap \tilde{t}_2 = \emptyset \quad \Delta \asymp \Delta'}{\Gamma \vdash_{\tilde{t}_1 \cdot \tilde{t}_2} P \mid Q \triangleright_{\tilde{t}_1 \cdot \tilde{t}_2} \Delta \circ \Delta'} \quad [\textsc{Conc}]$$

$$\frac{\Delta \ \mathsf{end\ only} \quad \Delta' \ [\,] \ \mathsf{only}}{\Gamma \vdash \mathbf{0} \triangleright_\emptyset \Delta, \Delta'} \qquad \frac{\Gamma \vdash P \triangleright_{\tilde{t}} \Delta \quad \Delta \leq \Delta'}{\Gamma \vdash P \triangleright_{\tilde{t}} \Delta'} \quad [\textsc{Inact}],[\textsc{Subs}]$$

$$\frac{\Gamma, a : \langle G \rangle \vdash_{\tilde{t}} P \triangleright \Delta}{\Gamma \vdash_{\tilde{t}} (\boldsymbol{\nu}\, a)P \triangleright \Delta} \qquad \frac{\Gamma \vdash P \triangleright_{\tilde{t}} \Delta, \tilde{s} : \{T_\mathtt{p}@\mathtt{p}\}_{\mathtt{p} \in I} \quad \tilde{s} \in \tilde{t} \quad \{T_\mathtt{p}@\mathtt{p}\}_{\mathtt{p} \in I} \ \mathsf{coherent}}{\Gamma \vdash_{\tilde{t} \setminus \tilde{s}} \ (\boldsymbol{\nu}\, \tilde{s})P \triangleright \Delta}$$

$[\textsc{NRes}],[\textsc{CRes}]$

$$\frac{\Gamma \vdash \tilde{e} : \tilde{S} \qquad \Delta \ \mathsf{end\ only}}{\Gamma, X : \tilde{S}\tilde{T} \vdash_\emptyset X\langle \tilde{e}\tilde{s}_1..\tilde{s}_n \rangle \triangleright \Delta, \tilde{s}_1 : T_1@\mathtt{p}_1, .., \tilde{s}_n : T_n@\mathtt{p}_n} \quad [\textsc{Var}]$$

$$\frac{\Gamma, X : \tilde{S}\tilde{T}, \tilde{x} : \tilde{S} \vdash_\emptyset P \triangleright \tilde{s}_1 : T_1@\mathtt{p}_1..\tilde{s_n} : T_n@\mathtt{p}_n \quad \Gamma, X : \tilde{S}\tilde{T} \vdash_{\tilde{t}} Q \triangleright \Delta}{\Gamma \vdash_{\tilde{t}} \mathsf{def}\ X(\tilde{x}\tilde{s}_1..\tilde{s}_n) = P\ \mathsf{in}\ Q \triangleright \Delta} \quad [\textsc{Def}]$$

The typing rules for queues are from Figure 8.

B.1. Proof of Proposition 5.5

Suppose P is a program phrase. By definition, P is without queues and without bound channels. We show two implications.

(1) $\Gamma \vdash P \triangleright \Delta$ implies $\Gamma \vdash P \triangleright_\emptyset \Delta$: Suppose P is typable in the original typing rules (for program phrases). Since the typing rules for runtime processes subsume the original rules, they can type P with the same derivation.

(2) $\Gamma \vdash P \triangleright_\emptyset \Delta$ without [Subs] implies $\Gamma \vdash P \triangleright \Delta$: Suppose P is typable in the refined system as $\Gamma \vdash P \triangleright_\emptyset \Delta$ without type contexts in Δ and without using [Subs]. By the

lack of [SUBS] in the derivation, the derivation precisely follows the structure of P. We inspect the potential differences between the original rules and the refined rules.

— (Use of Type Contexts in Derivation) Suppose the derivation uses a type context. The only place it can be taken off is [CONC]. Since there is no queue in P this means the type context has been empty as the result of weakening by [INACT]. Hence its use can be taken off from the derivations.
— (Use of Refined Constraints on Queue Channels in Judgements) Since the only rule which decreases the number of mentioned queue channels in the judgement (as in $\triangleright_{\tilde{s}}$) is [CRES] we know each judgement in the derivation has the $\emptyset$ as its mentioned queue channels. Hence the constraint on queue channels in [CONC] and other rules are never used.

Thus this derivation for P in the refined rules offers the derivation in the original rules as is, hence done. □

B.2. Proof of Proposition 5.6

Assume $\Gamma \vdash P \triangleright_{s[1..m]} \Delta$. We call $s_1...s_m$ in $\Gamma \vdash P \triangleright_{s[1..m]} \Delta$, the judgement's *mentioned queue channels* or simply *queue channels*.

We first show there is one-to-one correspondence between the free queues in P and the mentioned queue channels by inspecting each rule.

Case [INACT]: Zero queue channel to zero queue.

Case [QNIL]: It connects precisely one channel to one queue.

Case [QVAL], [QSESS], [QSEL]: These "enqueue" rules leave the number of channels one assigned to the unique queue channel.

Case [MCAST], [MACC], [SEND], [RCV], [DELEG], [SREC], [SEL] and [BRANCH], [IF], [VAR], [DEF]: Each of these process construction rules leaves the queue channels unchanged (empty).

Case [CONC]: in the premise, assume $\Gamma \vdash P \triangleright_{\tilde{t}_1} \Delta$ and $\Gamma \vdash_{\tilde{t}_2} Q \triangleright \Delta'$ the free queues in P have channels $\tilde{t}_1$ while the free queues in Q have channels $\tilde{t}_2$. Since we assume $\tilde{t}_1 \cap \tilde{t}_2 = \emptyset$ and $P|Q$ have exactly the sum of their respective queues.

Case [NRES]: The rule leaves both the queues and the queue channels unchanged.

Case [CRES]: The rule precisely takes off those channels whose channels become bound.

Case [SUBS]: No change in the process and no change in the queue. This exhausts all cases.

By the case analysis above, we conclude that free queues and mentioned queue channels precisely correspond to each other. Further the case analysis also shows that each prefix rule assumes the process has no free queue before prefixing (in the premise). Further a program phrase cannot have channel restriction so that all of its existing queues should be recorded in queue channels. We can now conclude that no queue can be under a prefix. □

B.3. Proof of Lemma 5.14

By the definition of $\circ$ on Δ, it suffices to show the commutativity and associativity at the level of types and type contexts, assuming that combined type contexts never share a target channel (in the sense defined just before Lemma 5.14, page 33).

We first show the commutativity. We write $H_1 \asymp H_2$ (which we read: "H_1 and H_2 are coherent") when $H_1 \circ H_2$ is defined. Note $H_1 \circ H_2$ means either:

— both of $H_{1,2}$ are type contexts and they do not share a target channel; or
— one of $H_{1,2}$ is a type context and the other is a type.

Below the designation "context-context" below means the case when we compose two contexts, similarly for others.

Case Context-Context: We consider the composition of $\mathcal{T}_{1,2}$ which are disjoint in targets (by our assumption). Then we always have:

$$\mathcal{T}_1 \asymp \mathcal{T}_2 \tag{22}$$
$$\mathcal{T}_1 \circ \mathcal{T}_2 = \mathcal{T}_1[\mathcal{T}_2] \tag{23}$$

By the symmetry of $\asymp$ (or equivalently by the assumption on target channels) we have:

$$\mathcal{T}_2 \asymp \mathcal{T}_1 \tag{24}$$
$$\mathcal{T}_2 \circ \mathcal{T}_1 = \mathcal{T}_2[\mathcal{T}_1] \tag{25}$$

Because of the isomorphism by the permutation equivalence for target-disjoint type contexts (cf. Section 5.1, paragraph **Type contexts**: recall $\approx$ is extended to type contexts unlike $\leq_{\mathsf{sub}}$) we have $\mathcal{T}_1[\mathcal{T}_2] \approx \mathcal{T}_2[\mathcal{T}_1]$ hence we are done.

Case Type-Context: Immediate since, by definition, $\mathcal{T} \asymp T$ and $T \asymp \mathcal{T}$ always and $\mathcal{T} \circ T = T \circ \mathcal{T} = \mathcal{T}[T]$.

Case Context-Type: Symmetric to the case above.

Case Type-Type: Never defined hence vacuous.

This exhausts all cases.

Next we show associativity.

Case Context-Context-Context: We consider the composition of $\mathcal{T}_{1,2,3}$, showing $(\mathcal{T}_1 \circ \mathcal{T}_2) \circ \mathcal{T}_3$ and $\mathcal{T}_1 \circ (\mathcal{T}_2 \circ \mathcal{T}_3)$ coincide in definedness and their resulting values. Assume $\mathcal{T}_{1,2}$ are mutually disjoint in target channels, similarly for $\mathcal{T}_1[\mathcal{T}_2]$ and $\mathcal{T}_3$. Then automatically:

$$\mathcal{T}_1 \asymp \mathcal{T}_2 \tag{26}$$
$$\mathcal{T}_1[\mathcal{T}_2] \asymp \mathcal{T}_3 \tag{27}$$
$$\mathcal{T}_1[\mathcal{T}_2] \circ \mathcal{T}_3 = \mathcal{T}_1[\mathcal{T}_2][\mathcal{T}_3] \tag{28}$$

By (27) we have:

$$\mathcal{T}_2 \asymp \mathcal{T}_3 \tag{29}$$
$$\mathcal{T}_1 \asymp \mathcal{T}_2[\mathcal{T}_3] \tag{30}$$
$$\mathcal{T}_1 \circ \mathcal{T}_2[\mathcal{T}_3] = \mathcal{T}_1[\mathcal{T}_2[\mathcal{T}_3]] \tag{31}$$

Since $\mathcal{T}_1[\mathcal{T}_2][\mathcal{T}_3] = \mathcal{T}_1[\mathcal{T}_2[\mathcal{T}_3]]$ we are done. The other direction is symmetric.

Case Context-Context-Type: We consider the composition of $\mathcal{T}_{1,2}$ and T, showing that the definedness and the resulting value of $(\mathcal{T}_1 \circ \mathcal{T}_2) \circ T$ and $\mathcal{T}_1 \circ (\mathcal{T}_2 \circ T)$ coincide. This case is not symmetric hence we show both directions. First if $\mathcal{T}_{1,2}$ are disjoint then automatically:

$$\mathcal{T}_1 \asymp \mathcal{T}_2 \tag{32}$$
$$\mathcal{T}_1[\mathcal{T}_2] \asymp T \tag{33}$$
$$\mathcal{T}_1[\mathcal{T}_2] \circ T = \mathcal{T}_1[\mathcal{T}_2][T] \tag{34}$$

We also always have:

$$\mathcal{T}_2 \asymp T \tag{35}$$
$$\mathcal{T}_1 \asymp \mathcal{T}_2[T] \tag{36}$$
$$\mathcal{T}_1 \circ \mathcal{T}_2[T] = \mathcal{T}_1[\mathcal{T}_2[T]] \tag{37}$$

Since $\mathcal{T}_1[\mathcal{T}_2][T] = \mathcal{T}_1[\mathcal{T}_2[T]]$ we are done. For the other direction, we first compose $\mathcal{T}_2$ and T then compose $\mathcal{T}_1$. As noted we always have

$$\mathcal{T}_2 \asymp T \tag{38}$$
$$\mathcal{T}_1 \asymp \mathcal{T}_2[T] \tag{39}$$
$$\mathcal{T}_1 \circ \mathcal{T}_2[T] = \mathcal{T}_1[\mathcal{T}_2[T]] \tag{40}$$

By our assumption $\mathcal{T}_1$ and $\mathcal{T}_2$ do not share a target channel. Hence:

$$\mathcal{T}_1 \asymp \mathcal{T}_2 \tag{41}$$
$$\mathcal{T}_1[\mathcal{T}_2] \asymp T \tag{42}$$
$$\mathcal{T}_1[\mathcal{T}_2] \circ T = \mathcal{T}_1[\mathcal{T}_2][T] \tag{43}$$

Again we note $\mathcal{T}_1[\mathcal{T}_2[T]] = \mathcal{T}_1[\mathcal{T}_2][T]$; hence we are done.

Case Type-Context-Context, Context-Type-Context: By the case Context-Context-Type above and commutativity.

Since we can never combine two types this exhausts all cases. □

B.4. Proof of Subject Reduction Theorem (Theorem 5.19)

(1) is by rule induction on $\equiv$ showing, in both ways, that if one side has a typing then the other side has the same typing. In the following we safely ignore uninteresting (permutable) final applications of [SUBS] in derivations by way of Lemma 5.16.

Case $P \mid \mathbf{0} \equiv P$: First assume $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$. By $\Gamma \vdash \mathbf{0} \triangleright_{\emptyset} \emptyset$ and by applying [CONC] to these two sequents we immediately obtain $\Gamma \vdash P|\mathbf{0} \triangleright_{\tilde{s}} \Delta$, as required. For the converse direction assume $\Gamma \vdash P|\mathbf{0} \triangleright_{\tilde{s}} \Delta$. We can safely assume (via Lemma 5.16) that the last rule applied is [CONC]. Thus we can set $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta_1$ and $\Gamma \vdash \mathbf{0} \triangleright_{\emptyset} \Delta_2$ such that $\Delta_1 \circ \Delta_2 = \Delta$. Note we can safely regard $\Gamma \vdash \mathbf{0} \triangleright \Delta_2$ as being inferred by the axiom [INACT] since applying [SUBS] to empty types and empty type contexts again lead to the empty types and empty type contexts: thus Δ_2 consists of only empty types and empty type contexts. Thus, in the composition $\Delta_1 \circ \Delta_2$, the empty types and some of the empty type contexts from Δ_2 are added to Δ_1 to generate Δ. Let this added part be Δ_2'. Since we can weaken Δ_1 in the first sequent with Δ_2' using Lemma 5.18 (2) we are done.

Case $P \mid Q \equiv Q \mid P$: By symmetry of the rule we have only to show one direction. Suppose $\Gamma \vdash P|Q \triangleright_{\tilde{s}} \Delta$. We can safely assume the last rule applied is [CONC]. We can thus set $\Gamma \vdash P \triangleright_{\tilde{t}} \Delta_1$ and $\Gamma \vdash Q \triangleright_{\tilde{r}} \Delta_2$ such that $\Delta_1 \asymp \Delta_2$, $\Delta_1 \circ \Delta_2 = \Delta$ and $\tilde{t} \uplus \tilde{r} = \tilde{s}$. By Lemma 5.14 we know $\Delta_2 \asymp \Delta_1$ and $\Delta_2 \circ \Delta_1 = \Delta$ hence by applying [CONC] with the premises reversed we are done.

Case $(P \mid Q) \mid R \equiv P \mid (Q \mid R)$: By the establishment of the previous case again we have only to show one direction. Suppose $\Gamma \vdash (P \mid Q) \mid R \triangleright_{\tilde{s}} \Delta$. We can safely assume: $\Gamma \vdash P \triangleright_{\tilde{t}} \Delta_1$, $\Gamma \vdash P \triangleright_{\tilde{r}} \Delta_2$ and $\Gamma \vdash P \triangleright_{\tilde{q}} \Delta_3$ such that $\Delta_1 \asymp \Delta_2$, $(\Delta_1 \circ \Delta_2) \asymp \Delta_3$ and $(\Delta_1 \circ \Delta_2) \circ \Delta_3 = \Delta$, as well as $\tilde{t} \uplus \tilde{r} \uplus \tilde{q} = \tilde{s}$. By the last condition, no two of Δ_1, Δ_2 and Δ_3 share a common target channel in their type contexts (in the sense given just before Lemma 5.14, page 33) because if the queue for a certain channel does not

exist in a sequent then it cannot be used as a target channel in a type context in its typing. Thus we can apply Lemma 5.14 to know $\Delta_2 \asymp \Delta_3$, $\Delta_1 \asymp (\Delta_2 \circ \Delta_3)$ and $\Delta_1 \circ (\Delta_2 \circ \Delta_3) = \Delta$. By applying [CONC] in an appropriate order we are done.

The remaining rules are reasoned exactly as in [Yoshida and Vasconcelos 2007] (note the arguments for congruence rules are direct from the compositionality of the typing rules). This concludes the proof of (1).

For (2), we establish the following stronger claim by rule induction.

Claim. Suppose $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta$ and Δ is partially coherent (cf. Definition 5.9). Then $P \to P'$ implies $\Gamma \vdash P \triangleright_{\tilde{s}} \Delta'$ such that either $\Delta \xrightarrow{\ell} \Delta'$ or $\Delta = \Delta'$.

All results on reduction on coherent typing is immediately applicable to partially coherent typing by Proposition 5.13 (1). Further by Proposition 5.13 (3), Δ' above is again partially coherent. Below we again ignore irrelevant final application of [SUBS] through Lemma 5.16. All rule names are those of the typing rules .

Case [LINK]: Let $R \stackrel{\text{def}}{=} \overline{a}_{[2..n]}(\tilde{s}).P_1 \mid a_{[2]}(\tilde{s}).P_2 \mid \cdots \mid a_{[n]}(\tilde{s}).P_n$ which is a redex of [LINK]. We write R_1 for $\overline{a}_{[2..n]}(\tilde{s}).P_1$ and R_i for $a_{[i]}(\tilde{s}).P_i$ $(2 \leq i \leq n)$. Assume:

$$\Gamma \vdash R \triangleright \Delta \tag{44}$$

By Lemma 5.15 we know $a \in \mathsf{dom}(\Gamma)$. Let $\Gamma(a) = G$. Since (44) can only be inferred by the sequence of [CONC] (up to permutable [SUBS], similarly in the following), we know $\Gamma \vdash R_i \triangleright \Delta_i$ $(1 \leq i \leq n)$ such that $\Delta_1 \circ ... \circ \Delta_n = \Delta$. By [MCAST] and [MACC] this means:

$$\Gamma \vdash P_i \triangleright \Delta_i, \tilde{s} : \{(G \upharpoonright i)@i\} \tag{45}$$

for each $1 \leq i \leq n$. Hence by the successive applications of [CONC] we reach:

$$\Gamma \vdash (\Pi_i P_i) \mid (\Pi_i s_i :: \emptyset) \triangleright_{\tilde{s}} \Delta, \tilde{s} : \{(G \upharpoonright i)@i\}_{1 \leq i \leq n} \tag{46}$$

Since $\{(G \upharpoonright i)@i\}_i$ collects all projections of G we can apply [CRES] to obtain:

$$\Gamma \vdash (\nu \tilde{s})(\ (\Pi_i P_i) \mid (\Pi_i s_i :: \emptyset)\ ) \triangleright \Delta \tag{47}$$

for a reductum of [LINK]. Note the typing does not change.

Case [SEND]: We use the first rule of Lemma 5.17 for "rolling back" a message. Suppose we have:

$$\Gamma \vdash s!\langle \tilde{e} \rangle; P \mid s :: \tilde{h} \triangleright_s \Delta \tag{48}$$

Since [CONC] is the only rule to derive this process we can set

$$\Gamma \vdash s!\langle \tilde{e} \rangle; P \triangleright_{\emptyset} \Delta_1 \tag{49}$$

and

$$\Gamma \vdash s :: \tilde{h} \triangleright_s \Delta_2 \tag{50}$$

such that $\Delta_1 \circ \Delta_2 = \Delta$. Since (49) can only be inferred from [SEND] we know, first:

$$\Gamma \vdash e_j : S_j \tag{51}$$

for each e_j in $\tilde{e}$; and, second, for some p and for some $\tilde{s}$ which includes s,

$$\Delta_1 \ = \ \Delta_1' \circ \tilde{s} : k!\,\langle \tilde{S} \rangle; T@\mathtt{p} \tag{52}$$

and moreover

$$\Gamma \vdash P \triangleright_{\emptyset} \Delta_1' \circ \tilde{s} : T@\mathtt{p}. \tag{53}$$

On the other hand by $\Delta_1 \asymp \Delta_2$ and (50) we know:

$$\Delta_2 \;=\; \Delta_2' \;\circ\; \tilde{s} : \mathcal{T}@\mathtt{p} \tag{54}$$

Now assume $\tilde{e} \downarrow \tilde{v}$. Notice by (51) we have $\Gamma \vdash v_j : S_j$ for each v_j in $\tilde{v}$. Thus by Lemma 5.17, [QVAL], we infer:

$$\Gamma \;\vdash\; s :: \tilde{h} \cdot \tilde{v} \;\triangleright\; \Delta_2' \;\circ\; \tilde{s} : \mathcal{T}[k!\,\langle \tilde{S} \rangle; [\,]]@\mathtt{p}. \tag{55}$$

By the algebra of located types and type contexts:

$$\begin{array}{rl} (\Delta_1' \circ \tilde{s} : T@\mathtt{p}) \;\circ\; (\Delta_2' \circ \tilde{s} : \mathcal{T}[k!\,\langle \tilde{S} \rangle; [\,]]@\mathtt{p}) & \\ = & (\Delta_1' \circ \tilde{s} : k!\,\langle \tilde{S} \rangle; T@\mathtt{p}) \circ (\Delta_2' \circ \tilde{s} : \mathcal{T}[\,]@\mathtt{p}) \\ = & \Delta \end{array}$$

Thus by applying [CONC] to (49) and (50) we obtain:

$$\Gamma \;\vdash\; P \mid s :: \tilde{h} \cdot \tilde{v} \;\triangleright\; \Delta \tag{56}$$

which gives the expected typing for the reductum of [SEND], with no type change.

Case [DELEG]: Similar to [SEND] using the second rule of Lemma 5.17, see Appendix B.5.

Case [LABEL]: We use the third rule of Lemma 5.17 together with the subtyping $\leq_{\mathsf{sub}}$. Suppose we have:

$$\Gamma \;\vdash\; s \lhd l; P \mid s :: \tilde{h} \;\triangleright_s\; \Delta \tag{57}$$

which is the redex of [LABEL]. Since [CONC] is the only rule to derive this process we can set, without loss of generality:

$$\Gamma \;\vdash\; s \lhd l; P \;\triangleright_\emptyset\; \Delta_1 \tag{58}$$

and

$$\Gamma \;\vdash\; s :: \tilde{h} \;\triangleright_s\; \Delta_2 \tag{59}$$

such that $\Delta_1 \circ \Delta_2 = \Delta$. Since (58) can only be inferred from [SEL] as the last rule (up to permutable applications of [SUBS]), we know, for some $\mathtt{p}$ and for some $\tilde{s}$ which includes s and for some $\{l_i\}$ which includes l,

$$\Delta_1 \;=\; \Delta_1' \circ \tilde{s} : k \oplus \{l_i : T_i\}_{i \in I}@\mathtt{p} \tag{60}$$

and moreover

$$\Gamma \;\vdash\; P \;\triangleright_\emptyset\; \Delta_1' \circ \tilde{s} : T_i@\mathtt{p}, \quad \text{for } i \in I. \tag{61}$$

On the other hand we can write:

$$\Delta_2 \;=\; \Delta_2' \;\circ\; \tilde{s} : \mathcal{T}@\mathtt{p} \tag{62}$$

By (59), (62) and Lemma 5.17, [QSEL], we infer:

$$\Gamma \;\vdash\; s :: \tilde{h} \cdot l \;\triangleright\; \Delta_2' \;\circ\; \tilde{s} : \mathcal{T}[k \oplus l : [\,]]@\mathtt{p}. \tag{63}$$

By the algebra of located types and type contexts together with subtyping:

$$\begin{array}{rl} (\Delta_1' \circ \tilde{s} : T_i@\mathtt{p}) \;\circ\; (\Delta_2' \circ \tilde{s} : \mathcal{T}[k \oplus l : [\,]]@\mathtt{p}) & \\ = & \Delta_1' \circ \Delta_2' \circ \tilde{s} : \mathcal{T}[k \oplus l : T_i]@\mathtt{p} \\ \leq_{\mathsf{sub}} & \Delta_1' \circ \Delta_2' \circ \tilde{s} : \mathcal{T}[k \oplus \{l_i : T_i\}_{i \in I}]@\mathtt{p} \\ = & (\Delta_1' \circ \tilde{s} : k \oplus \{l_i : T_i\}_{i \in I}@\mathtt{p}) \circ (\Delta_2' \circ \tilde{s} : \mathcal{T}@\mathtt{p}) \\ = & \Delta \end{array}$$

Thus we obtain, by applying [CONC] to (61), (63) then applying [SUBS] (the subsumption rule):

$$\Gamma \vdash P \mid s :: \tilde{h} \cdot l \triangleright \Delta \tag{64}$$

which gives the expected typing for the reductum of [SEND], with no type change.

Case [RECV]: By the first of the latter three rules of Lemma 5.17 together with Lemma 5.18. Suppose

$$\Gamma \vdash s?(\tilde{x}); P \mid s :: \tilde{v} \cdot \tilde{h} \triangleright_s \Delta \tag{65}$$

Since [CONC] is the only possible last rule (up to permutable [SUBS]) we can set

$$\Gamma \vdash s?(\tilde{x}); P \triangleright_\emptyset \Delta_1 \tag{66}$$

and

$$\Gamma \vdash s :: \tilde{v} \cdot \tilde{h} \triangleright_s \Delta_2 \tag{67}$$

such that $\Delta_1 \circ \Delta_2 = \Delta$. Since (66) can only be inferred from [RCV] we know, for some $\mathtt{p}$ and for some $\tilde{s}$ which includes s,

$$\Delta_1 = \Delta_1' \circ \tilde{s} : k? \langle \tilde{S} \rangle ; T@\mathtt{p} \tag{68}$$

and moreover

$$\Gamma, \tilde{x} : \tilde{S} \vdash P \triangleright_\emptyset \Delta_1' \circ \tilde{s} : T@\mathtt{p}. \tag{69}$$

By Lemma 5.18, we obtain:

$$\Gamma \vdash P[\tilde{v}/\tilde{x}] \triangleright_\emptyset \Delta_1' \circ \tilde{s} : T@\mathtt{p}. \tag{70}$$

Further by $\Delta_1 \asymp \Delta_2$ and (67) we know:

$$\Delta_2 = \Delta_2' \circ \tilde{s} : k! \langle \tilde{S} \rangle . \mathcal{T}@\mathtt{p} \tag{71}$$

By Lemma 5.17, [QVALDQ], we infer:

$$\Gamma \vdash s :: \tilde{h} \triangleright \Delta_2' \circ \tilde{s} : \mathcal{T}@\mathtt{p}. \tag{72}$$

Then we obtain:

$$\begin{aligned} \Delta &\stackrel{\text{def}}{=} (\Delta_1' \circ \tilde{s} : k? \langle \tilde{S} \rangle . T@\mathtt{p}) \circ (\Delta_2' \circ \tilde{s} : k! \langle \tilde{S} \rangle . \mathcal{T}@\mathtt{p}) \\ &\stackrel{\ell}{\to} \quad (\Delta_1', \tilde{s} : T@\mathtt{p}) \circ (\Delta_2' \circ \tilde{s} : \mathcal{T}@\mathtt{p}) \quad (\stackrel{\text{def}}{=} \Delta') \end{aligned}$$

Thus by applying [CONC] to (66) and (67) we obtain:

$$\Gamma \vdash P[\tilde{v}/\tilde{x}] \mid s :: \tilde{h} \triangleright \Delta' \tag{73}$$

such that $\Delta \stackrel{\ell}{\to} \Delta'$, as required. Note this case demands reduction of typings.

Case [SREC], [BRANCH]: Similar to [RECV], using the latter two rules of Lemma 5.17, see Appendix B.5.

Case [IFT], [IFF], [DEF], [DEFIN]: Standard, cf. [Yoshida and Vasconcelos 2007]. No difference in the typing.

Case [SCOP]: When a shared name is hidden, assume

$$\Gamma \vdash (\boldsymbol{\nu} a) P \triangleright_{\tilde{s}} \Delta \tag{74}$$

and $P \to P'$. Then we can set

$$\Gamma, a : \langle G \rangle \ \vdash \ P \ \triangleright_{\tilde{s}} \ \Delta. \tag{75}$$

By induction hypothesis we know

$$\Gamma, a : \langle G \rangle \ \vdash \ P' \ \triangleright_{\tilde{s}} \ \Delta' \tag{76}$$

such that either $\Delta \stackrel{\ell}{\to}^{0,1} \Delta'$. Hence by [NRES] we have

$$\Gamma \ \vdash \ (\boldsymbol{\nu} a)P' \ \triangleright_{\tilde{s}} \ \Delta' \tag{77}$$

as required. When session channels are hidden, suppose

$$\Gamma \ \vdash \ (\boldsymbol{\nu} \tilde{s})P \ \triangleright_{\tilde{t} \setminus \tilde{s}} \ \Delta \tag{78}$$

and $P \to P'$. We can set:

$$\Gamma \ \vdash \ P \ \triangleright_{\tilde{t}} \ \Delta, \tilde{s} : \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \tag{79}$$

where $\{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I}$ is coherent. By induction hypothesis

$$\Gamma \ \vdash \ P' \ \triangleright_{\tilde{t}} \ \Delta', \tilde{s} : \{T'_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \tag{80}$$

where either $\Delta \stackrel{\ell}{\to}^{0,1} \Delta'$ or $\{s\} : \{T_p @ p\}_{p \in I} \to^{0,1} \{s\} : \{T'_p @ p\}_{p \in I}$. By Proposition 5.13 (2) $\{T'_p @ p\}_{p \in I}$ is again coherent. Hence by [CRES] we obtain

$$\Gamma \ \vdash \ (\boldsymbol{\nu} \tilde{s})P' \ \triangleright_{\tilde{t} \setminus \tilde{s}} \ \Delta' \tag{81}$$

as required.

Case [PAR]: Suppose we have $\Gamma \ \vdash \ P|Q \ \triangleright_{\tilde{t}_1 \cdot \tilde{t}_2} \ \Delta$ and $P \to P'$. By [CONC] we have $\Gamma \ \vdash \ P \ \triangleright_{\tilde{t}_1} \ \Delta_1$ and $\Gamma \ \vdash \ Q \ \triangleright_{\tilde{t}_2} \ \Delta_2$ such that $\Delta_1 \circ \Delta_2 = \Delta$. By induction hypothesis we have $\Gamma \ \vdash \ P' \ \triangleright_{\tilde{t}_1} \ \Delta'_1$ such that $\Delta_1 \to^{0,1} \Delta'_1$. By Proposition 5.13 (1) we have $\Delta'_1 \asymp \Delta_2$ hence $\Gamma \ \vdash \ P'|Q \ \triangleright_{\tilde{t}_1 \cdot \tilde{t}_2} \ \Delta'_1 \circ \Delta_2$. Noting Proposition 5.13 (1) also says that $(\Delta_1 \circ \Delta_2) \to^{0,1} (\Delta'_1 \circ \Delta_2)$ we are done.

Case [STR]: Immediate from Subject Congruence (the first clause of this theorem). This exhausts all cases for (2).

(3) is because the empty typing $\emptyset$ is always coherent. □

B.5. Remaining Cases of Theorem 5.19

Case [DELEG]: We use the second rule of Lemma 5.17. Suppose we have:

$$\Gamma \ \vdash \ s!\langle\!\langle \tilde{t} \rangle\!\rangle; P \mid s :: \tilde{h} \ \triangleright_s \ \Delta \tag{82}$$

Since [CONC] is the only rule to derive this process we can set

$$\Gamma \ \vdash \ s!\langle\!\langle \tilde{t} \rangle\!\rangle; P \ \triangleright_{\emptyset} \ \Delta_1 \tag{83}$$

and

$$\Gamma \ \vdash \ s :: \tilde{h} \ \triangleright_s \ \Delta_2 \tag{84}$$

such that $\Delta_1 \circ \Delta_2 = \Delta$. Since (83) can only be inferred from [DELEG] we know, for some $\mathtt{p}$ and for some $\tilde{s}$ which includes s,

$$\Delta_1 \ = \ \Delta'_1 \circ (\tilde{s} : k! \, \langle T' @ \mathtt{p}' \rangle . T @ \mathtt{p}, \ \tilde{t} : T' @ \mathtt{p}') \tag{85}$$

and moreover

$$\Gamma \ \vdash \ P \ \triangleright_{\emptyset} \ \Delta'_1, \ \tilde{s} : T @ \mathtt{p}. \tag{86}$$

On the other hand by $\Delta_1 \asymp \Delta_2$ and (50) we know:

$$\Delta_2 \;=\; \Delta_2' \;\circ\; \tilde{s} : \mathcal{T}[\,]@\mathtt{p} \tag{87}$$

By Lemma 5.17, [QSESS], we infer:

$$\Gamma \;\vdash\; s::\tilde{h}\cdot\tilde{t} \;\triangleright_{\tilde{s}'}\; \Delta_2' \circ \tilde{s} : \{\mathcal{T}[k!\,\langle T@\mathtt{p}'\rangle.[\,]]@\mathtt{p}\},\; \tilde{t} : \{T@\mathtt{p}'\}. \tag{88}$$

By the algebra of located types and type contexts:

$$\begin{array}{rcl}(\Delta_1',\; \tilde{s}{:}T@\mathtt{p}) \;\circ\; (\Delta_2' \circ \tilde{s} : \{\mathcal{T}[k!\,\langle T@\mathtt{p}'\rangle.[\,]]@\mathtt{p}\},\; \tilde{t} : \{T@\mathtt{p}'\}) & & \\ \quad = \quad (\Delta_1' \circ (\tilde{s} : k!\,\langle T'@\mathtt{p}'\rangle.T@\mathtt{p},\; \tilde{t}{:}T'@\mathtt{p}')) \;\circ\; (\Delta_2' \;\circ\; \tilde{s} : \mathcal{T}[\,]@\mathtt{p}) & & \\ \quad = \quad \Delta & & \end{array}$$

Thus by applying [CONC] to (83) and (84) we obtain:

$$\Gamma \;\vdash\; P \mid s::\tilde{h}\cdot\tilde{t} \;\triangleright\; \Delta \tag{89}$$

which gives the expected typing for the reductum of [DELEG], with no type change.

Case [SREC]: By the second to the last rule of Lemma 5.17. Suppose

$$\Gamma \;\vdash\; s?((\tilde{t}\,)); P \mid s::\tilde{t}\cdot\tilde{h} \;\triangleright_s\; \Delta \tag{90}$$

Since [CONC] is the only possible last rule (up to permutable [SUBS]) we can set

$$\Gamma \;\vdash\; s?((\tilde{t}\,)); P \;\triangleright_\emptyset\; \Delta_1 \tag{91}$$

and

$$\Gamma \;\vdash\; s::\tilde{t}\cdot\tilde{h} \;\triangleright_s\; \Delta_2 \tag{92}$$

such that $\Delta_1 \circ \Delta_2 = \Delta$. Since (91) can only be inferred from [SREC] we know, for some p and for some $\tilde{s}$ which includes s,

$$\Delta_1 \;=\; \Delta_1' \circ \tilde{s} : k?\,\langle T'@\mathtt{p}'\rangle.T@\mathtt{p} \tag{93}$$

and moreover

$$\Gamma \;\vdash\; P \;\triangleright_\emptyset\; \Delta_1' \circ \tilde{s}{:}T@\mathtt{p},\; \tilde{t}{:}T'@\mathtt{p}' \tag{94}$$

By $\Delta_1 \asymp \Delta_2$ and (92) we know:

$$\Delta_2 \;=\; \Delta_2' \;\circ\; \tilde{s} : k!\,\langle T'@\mathtt{p}'\rangle.\mathcal{T}@\mathtt{p},\; \tilde{t} : T'@\mathtt{p}' \tag{95}$$

By Lemma 5.17, [QSESSDQ], we infer:

$$\Gamma \;\vdash\; s::\tilde{h} \;\triangleright\; \Delta_2' \;\circ\; \tilde{s} : \mathcal{T}@\mathtt{p}. \tag{96}$$

Then we obtain:

$$\begin{array}{rcl}\Delta \;\stackrel{\text{def}}{=}\; (\Delta_1' \circ \tilde{s} : k?\,\langle T'@\mathtt{p}'\rangle.T@\mathtt{p}) \circ (\Delta_2' \;\circ\; \tilde{s} : k!\,\langle T'@\mathtt{p}'\rangle.\mathcal{T}@\mathtt{p},\; \tilde{t} : T'@\mathtt{p}') & & \\ \quad \stackrel{\ell}{\to} \quad (\Delta_1' \;\circ\; \tilde{s}{:}T@\mathtt{p},\; \tilde{t} : T'@\mathtt{p}') \circ (\Delta_2' \;\circ\; \tilde{s} : \mathcal{T}@\mathtt{p}) \quad (\stackrel{\text{def}}{=} \Delta') & & \end{array}$$

Thus by applying [CONC] to (91) and (92) we obtain:

$$\Gamma \;\vdash\; P \mid s::\tilde{h} \;\triangleright\; \Delta' \tag{97}$$

such that $\Delta \stackrel{\ell}{\to} \Delta'$, as required. Note this case again demands reduction of typings.

Case [BRANCH]: By the last rule of Lemma 5.17. Suppose

$$\Gamma \;\vdash\; s \rhd \{l_i : P_i\}_{i\in I} \mid s::l_j\cdot\tilde{h} \;\triangleright_s\; \Delta \tag{98}$$

where we assume $j \in I$. Since [CONC] is the only possible last rule (up to permutable [SUBS]) we can set

$$\Gamma \vdash s \triangleright \{l_i\colon P_i\}_{i\in I} \triangleright_\emptyset \Delta_1 \tag{99}$$

and

$$\Gamma \vdash s::l_j \cdot \tilde{h} \triangleright_s \Delta_2 \tag{100}$$

such that $\Delta_1 \circ \Delta_2 = \Delta$. First for Δ_2 we know, for some p and for some $\tilde{s}$ which includes s:

$$\Delta_2 = \Delta_2' \circ \tilde{s} : k \oplus l_j : \mathcal{T}@\mathtt{p} \tag{101}$$

where by assumption we have $j \in I$. Since (99) can only be inferred from [BRANCH] and by $\Delta_1 \asymp \Delta_2$, we also know:

$$\Delta_1 = \Delta_1' \circ \tilde{s}\colon k \,\&l_j : T_j@\mathtt{p} \tag{102}$$

(where $\&l_j : T_j$ is the singleton notation as in selection) and moreover

$$\Gamma \vdash P_i \triangleright_\emptyset \Delta_1' \circ \tilde{s}\colon T_i@\mathtt{p} \tag{103}$$

for each $i \in I$ (so (102) is inferred using [SUBS]). By Lemma 5.17, [QSELDQ], we infer:

$$\Gamma \vdash s::\tilde{h} \triangleright \Delta_2' \circ \tilde{s} : \mathcal{T}@\mathtt{p}. \tag{104}$$

Then we obtain:

$$\begin{aligned}\Delta &\stackrel{\text{def}}{=} (\Delta_1' \circ \tilde{s}\colon k \,\&l_j : T_j@\mathtt{p}) \circ (\Delta_2' \circ \tilde{s} : k \oplus l_j : \mathcal{T}@\mathtt{p}) \\ &\stackrel{\ell}{\to} \quad (\Delta_1',\ \tilde{s}\colon T_j@\mathtt{p}) \circ (\Delta_2' \circ \tilde{s} : \mathcal{T}@\mathtt{p}) \quad (\stackrel{\text{def}}{=} \Delta')\end{aligned}$$

Thus by applying [CONC] to (99) and (100) we obtain:

$$\Gamma \vdash P \mid s::\tilde{h} \triangleright \Delta' \tag{105}$$

such that $\Delta \stackrel{\ell}{\to} \Delta'$, as required. Again we need a reduction of typings. □

B.6. Proof of Lemma 5.21

Proof of (1) and (2). We prove the following claim which implies both (1) and (2) by rule induction on the typing rules. Below and henceforth we are confusing a free session channel and its numeric representation in the typing. Recall Δ is partially coherent when for some Δ_0 we have $\Delta \asymp \Delta_0$ and $\Delta \circ \Delta_0$ is coherent.

Claim. Assume $\Gamma \vdash P \triangleright_{\tilde{t}} \Delta$ s.t. Δ is partially coherent and there is no queue at s. Assume $P\langle\!\langle s \rangle\!\rangle$. Then one of the following conditions holds.

(a) P contains a unique active receiving (resp. emitting) prefix at s and Δ contains a unique minimal receiving (resp. emitting) prefix at s (Δ may contain another minimal prefix at s).
(b) P contains a unique minimal receiving prefix at s and a unique minimal emitting prefix at s. Moreover Δ contains a unique minimal receiving prefix at s and a unique minimal emitting prefix at s.

Case [MCAST], [MACC]: Vacuous since in this case the unique active prefix in the process is at a shared name.

Case [SEND], [RCV], [DELEG], [SREC], [SEL] and [BRANCH]: Immediate since there can only be a unique active channel name which is by the given prefixing.

Case [INACT], [IF], [VAR], [DEF], [QNIL], [QVAL], [QSESS], [QSEL]: Vacuous.

Case [CONC]: Suppose

$$\Gamma \vdash P \triangleright_{\tilde{t}_1} \Delta, \qquad \Gamma \vdash_{\tilde{t}_2} Q \triangleright \Delta' \tag{106}$$

such that $\tilde{t}_1 \cap \tilde{t}_2 = \emptyset$ and $\Delta \asymp \Delta'$. Observe if $\Delta \circ \Delta'$ is partially coherent then Δ and Δ' respectively are partially coherent by definition. By induction hypothesis we can assume P and Q satisfy the required condition.

(1) If only one party has an active prefix at s there is nothing to prove.
(2) If both are active at s, suppose both processes, hence Δ and Δ', have receiving active prefixes at s. Then this cannot be partially coherent since if so then the assumed completion of $\Delta \circ \Delta'$ to a coherent typing should also contain two minimal receiving prefixes which are impossible by the definition of $\circ$. Similarly when two include active emitting prefixes at s, hence as required.

Note that this pair may *not* be a redex: we do not (have to) validate coherence until we hide channels, however it is important that there is one output and one input since if not there will be a conflict at s.

Case [NRES]: Vacuous since there is no change either in the process nor in the typing.

Case [CRES]: Vacuous since there is no difference in the typing for s nor in the activeness in prefixes.

Case [SUBS]: Vacuous again. □

B.7. Proof of Proposition 5.26

We show the following logically equivalent result:

Claim. (1) If P is simple then

(1-a) no delegation prefix (input or output) occurs in P and
(1-b) for each prefix with a shared name in P, say $a[\mathtt{i}](\tilde{s}).P'$ or $\overline{a}[2..\mathtt{n}](\tilde{s}).P'$, there is no free session channels in P' except $\tilde{s}$.

(2) If P is simple and $P \to P'$ then P' is again simple.

We first show (1) by rule induction on typing rules.

Case [MCAST]: The rule reads:

$$\frac{\Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash_{\emptyset} P \triangleright \Delta, \tilde{s} : (G \upharpoonright \mathtt{1})@\mathtt{1} \quad |\tilde{s}| = \mathsf{sid}(G)}{\Gamma \vdash_{\emptyset} \overline{a}_{[2..\mathtt{n}]}(\tilde{s}).P \triangleright \Delta}$$

First by simplicity we know $\Delta = \emptyset$ (since if not the premise has at least a doubleton typing). (1-a) is immediate from the induction hypothesis since the rule does not add a delegation prefix: For (1-b) if P' in $a[i](\tilde{\mathtt{s}}).P'$ (resp. $\overline{a}[2..\mathtt{n}](\tilde{s}).P'$) has free session channels then we cannot have $\Delta = \emptyset$, violating simplicity.

Case [MACC]: The rule reads:

$$\frac{\Gamma \vdash a : \langle G \rangle \quad \Gamma \vdash_{\emptyset} P \triangleright \Delta, \tilde{s} : (G \upharpoonright \mathtt{p})@\mathtt{p} \quad |\tilde{s}| = \mathsf{sid}(G)}{\Gamma \vdash_{\emptyset} a_{[\mathtt{p}]}(\tilde{s}).P \triangleright \Delta}$$

Again $\Delta = \emptyset$, and the remaining reasoning is precisely the same as [MCAST].

Case [SEND]: The rule reads:

$$\frac{\Gamma \vdash e_j : S_j \qquad \Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash_\emptyset s[k]!\langle \tilde{e} \rangle; P \triangleright \Delta, \tilde{s} : k! \langle \tilde{S} \rangle; T@\mathtt{p}}$$

Again $\Delta = \emptyset$. (1-a) is immediate from the induction hypothesis since the rule does not add any delegation prefix. (1-b) is again immediate from the induction hypothesis since the rule doe not add a shared-name prefix.

Case [RCV]: The rule reads:

$$\frac{\Gamma, \tilde{x} : \tilde{S} \vdash P_\emptyset \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash_\emptyset s[k]?(\tilde{x}); P \triangleright \Delta, \tilde{s} : k? \langle \tilde{S} \rangle; T@\mathtt{p}}$$

Precisely the same as in [SEND] above.

Case [DELEG]: The rule reads:

$$\frac{\Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash_\emptyset s[k]!\langle\!\langle \tilde{t} \rangle\!\rangle; P \triangleright \Delta, \tilde{s} : k! \langle T'@\mathtt{p}' \rangle; T@\mathtt{p}, \ \tilde{t} : T'@\mathtt{p}'}$$

Even if $\Delta = \emptyset$ the conclusion's typing becomes a doubleton hence this rule cannot be applied.

Case [SREC]: The rule reads:

$$\frac{\Gamma \vdash_\emptyset P \triangleright \Delta, \tilde{s} : T@\mathtt{p}, \ \tilde{t} : T'@\mathtt{p}'}{\Gamma \vdash_\emptyset s[k]?((\tilde{t})); P \triangleright \Delta, \tilde{s} : k? \langle T'@\mathtt{p}' \rangle; T@\mathtt{p}}$$

which is again impossible to apply (the premise's typing becomes a doubleton).

Case [SEL],[BRANCH]: Similar with [SEND] and [RCV].

Case [IF], [CONC], [CRES], [NRES], [SUBS], [DEF]: By the shape of these rules, in each rule, there is no addition or removal of a prefix from the premise to the conclusion. Hence both (1-a/b) are immediate from the induction hypothesis.

Case [INACT], [VAR], [QNIL], [QVAL], [QSESS], [QSEL]: Vacuous since no prefixes are involved.

Hence as required.

For (2) suppose a derivation of P is simple. By the proof of Theorem 5.19, if $P \to P'$ then we have essentially the same derivation for both P and P' except:

— taking off the lost pair of prefixes from that of P (three pair of prefix rules);
— one of the branches is chosen (conditional)
— copying some part from the derivation for P to that of P' (for recursion)

In each case clearly the simplicity of the derivation for P implies that of P', as required. □

B.8. Proof of Lemma 5.28

Suppose:

(C1). $\Gamma \vdash P \triangleright \Delta$.
(C2). P is simple.
(C3). Δ has a minimal receiving (resp. emitting) prefix at s.
(C4). none of the prefixes at s in P are under a shared name.
(C5). none of the prefixes at s in P are under a conditional branch.

Under these conditions, we show that P has an active receiving prefix (resp. has an active emitting prefix or a non-empty queue). We use rule induction on typing rules.

Case [MCAST], [MACC]: By Proposition 5.26 there can be no free session channels hence vacuous (since (C3) is not satisfied).

Case [SEND]: The "simple" rule reads:

$$\frac{\Gamma \vdash e_j : S_j \qquad \Gamma \vdash_\emptyset P \triangleright \tilde{s} : T@\mathtt{p}}{\Gamma \vdash_\emptyset s[k]!\langle \tilde{e} \rangle; P \triangleright \tilde{s} : k! \langle \tilde{S} \rangle; T@\mathtt{p}}$$

Observe that there can be no other minimal prefix in the typing in the conclusion than the newly introduced prefix itself: this corresponds to the unique minimal prefix in the typing.

Case [RCV]: The "simple" rule reads:

$$\frac{\Gamma, \tilde{x} : \tilde{S} \vdash P_\emptyset \triangleright \Delta, \tilde{s} : T@\mathtt{p}}{\Gamma \vdash_\emptyset s[k]?(\tilde{x}); P \triangleright \Delta, \tilde{s} : k? \langle \tilde{S} \rangle; T@\mathtt{p}}$$

Same as [SEND].

Case [SREC], [DELEG]: By Proposition 5.26 these rules are not used in derivation of a simple process, hence vacuous.

Case [SEL],[BRANCH]: Similar with [SEND],[RCV].

Case [IF]: Vacuous since (C5) does not hold.

Case [CONC]: The rule reads:

$$\frac{\Gamma \vdash P \triangleright_{\tilde{t}_1} \Delta \quad \Gamma \vdash_{\tilde{t}_2} Q \triangleright \Delta' \quad \tilde{t}_1 \cap \tilde{t}_2 = \emptyset \quad \Delta \asymp \Delta'}{\Gamma \vdash_{\tilde{t}_1 \cdot \tilde{t}_2} P \mid Q \triangleright_{\tilde{t}_1 \cdot \tilde{t}_2} \Delta \circ \Delta'}$$

We first observe:

Claim A1. If the result of the operation $\circ$ on typings (when defined) has a minimal input prefix then one of the original typings also has the same.

This is because, direct from the definition of $\circ$, if $\circ$ results in an input minimal input prefix then it cannot come from a type context (which contains only an output prefix) hence it can come only from the same in the premise. Further:

Claim A2. If the result of the operation $\circ$ on typings (when defined) has a minimal output prefix then one of the premises also has the same in the form of either the corresponding non-empty type context or the corresponding type ("corresponding" means that the minimal prefix coincides).

Above the details of the shape of a typing is in fact unnecessary.
Claim B. The composition $\mid$ preserves activeness of each prefix.

This is immediate from the definition.

Now we reason by induction. In the case of an input prefix in the typing, by Claim A1 we know one of the premises also contains an input prefix in the typing. Hence the corresponding process has an active input prefix by induction hypothesis. By Claim B we are done.

On the other hand in the case of an output prefix in the typing, by Claim A2 we know one of the premises also contains the same (either as the corresponding type context or the corresponding output prefix) in the typing. Hence by induction hypothesis the corresponding process has an active output prefix or a non-empty queue. Hence by induction hypothesis we are done. By Claim B we are done.

Case [INACT], [VAR]: Vacuous since in this case the typing does not contain any active channel hence violating (C3).

Case [SUBS], : The subsumption does not add any new active prefix in the typing hence by induction hypothesis we are done.

Case [DEF]: As [SUBS] above.

Case [QVAL], [QSESS], [QSEL]: In these cases we have a minimal emitting prefix in the typing; and we have a corresponding non-empty queue, as required.

Case [QNIL]: Vacuous since (C3) is violated.

Case [NRES]: This reads:

$$\frac{\Gamma, a : \langle G \rangle \vdash_{\tilde{t}} P \triangleright \Delta}{\Gamma \vdash_{\tilde{t}} (\boldsymbol{\nu}\, a) P \triangleright \Delta}$$

which shows there is no change in the typing and in the process with respect to (free) active/minimal prefixes hence immediate by induction hypothesis.

Case [CRES]: This reads:

$$\frac{\Gamma \vdash P \triangleright_{\tilde{t}} \Delta, \tilde{s} : \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \quad \tilde{s} \in \tilde{t} \quad \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \text{ coherent}}{\Gamma \vdash_{\tilde{t} \setminus \tilde{s}} \ (\boldsymbol{\nu}\, \tilde{s}) P \triangleright \Delta}$$

Suppose in the conclusion there is a minimal prefix at s in Δ. Then it is also minimal in the premise hence by induction hypothesis we are done.

This exhausts all cases. □