



A Very Gentle Introduction to Multiparty Session Types

Nobuko Yoshida^(✉) and Lorenzo Gheri

Imperial College London, London, UK
{n.yoshida,l.gheri}@imperial.ac.uk

Abstract. Multiparty session types (MPST) are a formal specification and verification framework for message-passing protocols without central control: the desired interactions at the scale of the network itself are specified into a session (called *global type*). Global types are then projected onto *local types* (one for each participant), which describe the protocol from a local point of view. These local types are used to validate an application through type-checking, monitoring, and code generation. Theory of session types guarantees that local conformance of all participants induces *global conformance* of the network to the initial global type. This paper provides a very gentle introduction of the simplest version of multiparty session types for readers who are not familiar with session types nor process calculi.

Keywords: Multiparty session types · Process calculi · Distributed systems · Type safety · Progress

1 A Gentle Introduction to Multiparty Session Types

Backgrounds. Session types were introduced in a series of papers during the 1990s [8, 10, 18] in the context of pure concurrent processes and programming. Session types have since been studied in many contexts over the last decade—see the surveys of the field [4, 13].

A basic observation underlying session types is that a communication-based application often exhibits a highly structured sequence of interactions involving, for example, sequencing, choices and recursion, which as a whole form a natural unit of *session*. The structure of a session is abstracted as a *type* through an intuitive syntax, which is then used for validating programs through associated language primitives.

Multiparty session types generalise the binary session type theory to the multiparty case, preventing deadlock and communication errors in more sophisticated communication protocols involving any number (two or more) of participants. The central idea is to introduce global types, which describe multiparty conversations from a global perspective and provide a means to check

Work partially supported by EPSRC projects EP/K034413/1, EP/K011715/1, EP/L00058X/1, EP/N027833/1 and EP/N028201/1.

protocol compliance. The theory [11, 12] was born by transforming the industry idea developed during designing a protocol description language, Scribble [17], which was presented by Kohei Honda in [9] (see the historical background [20]). This original work was extended in various ways and applied to many different programming languages and tools, some of which use Scribble. The reader who is interested in programming languages can find the tools and papers at [Mobility Reading Group Home Page](#).

Binary Session Types. As a first example, consider a scenario where a customer tries to conclude a deal with a travel agency. We associate a process with the *customer* (**Customer**) and one with the *agency* (**Agency**). The task involves synchronisation in the communication between the customer and the agency. Synchronisation is obtained through the exchange of messages.

Specifically, the protocol works as follows.

1. The customer sends an order to the agency, namely the place they desire to visit in their next travel (let us say *Hawaii*). On receiving the request, the agency looks up the price relative to that travel and sends a quote (*quote*) back to the customer.
2. When the customer receives the *quote*, they make a decision (choice): either to *accept* or to *reject* the offer. The customer communicates their decision to the agency, which is waiting for such a message.
3. In case of acceptance, the agency waits that the customer communicates their address (*address*), then sends a confirmation date (*date*) for the trip back to the customer and the protocol terminates.
4. In case of rejection, the deal is not made and the protocol immediately terminates.

The multiparty session types methodology is as follows. First, define a *global type* that gives a shared contract of the allowed pattern of message exchanges in the system. Second, *project* the global type to each end-point participant to obtain a *local type*: an obligation on the message sends and receipts for each process that together ensure that the pattern of messages are allowed by the global type. Finally, check that the implementation of each process conforms to its local type.

In our protocol, from a global perspective, we expect to see the following pattern of message exchanges, encoded as a *global type* for the communication:

$$\begin{aligned}
 & \text{Customer} \rightarrow \text{Agency}: \text{Hawaii}(\text{bool}). \text{Agency} \rightarrow \text{Customer}: \text{quote}(\text{nat}). \\
 & \text{Customer} \rightarrow \text{Agency}: \{ \\
 & \quad \text{accept}(\text{bool}). \\
 & \quad \text{Customer} \rightarrow \text{Agency}: \text{address}(\text{nat}). \\
 & \quad \text{Agency} \rightarrow \text{Customer}: \text{date}(\text{nat}). \text{end}, \\
 & \quad \text{reject}(\text{bool}). \text{end} \\
 & \}
 \end{aligned} \tag{1}$$

The type describes the global pattern of communication between **Customer** and **Agency** using message exchanges, sequencing, and choice. The basic pattern

Customer \rightarrow **Agency**: $m(S)$ indicates a message with label m sent from the **Customer** to the **Agency**, together with an element of sort S . The communication starts with the customer sending the message *Hawaii* to the agency, then the agency sends the *quote* label together with a natural number and at this point the customer either sends an *accept* or a *reject* message. In case of *reject*, the protocol ends (type **end**); otherwise, the communication continues with the sequential exchange of an *address* (from the customer to the agency) and a *date* (from the agency to the customer). The operator “.” denotes sequencing, the “;” separates possible continuations for the protocol.

The global type states what are the valid message sequences allowed in the system. When we implement **Customer** and **Agency** separately, we would like to check that their composition conforms to the global type. Since there are only two participants, projecting to each participant is simple. From the perspective of the **Customer**, the communication can be described by the type:

$$\begin{aligned} & ?Hawaii(\text{bool}).!quote(\text{nat}). \\ & ((!accept(\text{bool}).!address(\text{nat}).?date(\text{nat}).\text{end}) \oplus (!reject(\text{bool}).\text{end})) \end{aligned} \quad (2)$$

where $!m$ denotes a message with label m sent to **Agency**, $?m$ denotes a message with label m received from **Agency**, and \oplus denotes an internal choice. Thus the type states that, **Customer** after sending the message with label *Hawaii*, waits for a natural number for the *quote* and, after receiving it, decides either to send $!accept(\text{bool}). !address(\text{nat})$, wait for a natural number for the *date* and then exit, or to send just $!reject(\text{bool})$ and exit.

From the viewpoint of **Agency**, the same global session is described by the dual type

$$\begin{aligned} & Hawaii(\text{bool}).!quote(\text{nat}). \\ & ((?accept(\text{bool}).?address(\text{nat}).!date(\text{nat}).\text{end}) \& (?reject(\text{bool}).\text{end})) \end{aligned} \quad (3)$$

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

We can now individually check that the implementations of the customer and the agency conform to these local types.

Multiparty Session Types. In the case of two parties, the safety can be checked given the local type (2) since its dual (3) is unique (up to unfolding recursions).

However, for applications involving *multiple parties*, the global type and its projection to each participant are essential to provide a shared contract among all participants.

For example, consider a simple ring protocol, where, after the customer’s acceptance, the agency needs to forward some details to the hotel, before a direct communication starts between the hotel and the customer. Namely, **Customer** sends a message *details* to **Agency**, which forwards the message to **Hotel**. After receiving the message, **Hotel** sends an acknowledgement *ok* to **Customer**. We start by specifying the global type as:

$$\begin{aligned} \text{Customer} \rightarrow \text{Agency}: details(\text{nat}).\text{Agency} \rightarrow \text{Hotel}: details(\text{nat}). \\ \text{Hotel} \rightarrow \text{Customer}: ok(\text{bool}).\text{end} \end{aligned} \quad (4)$$

As before, we want to check each process locally against a local type such that, if each process conforms to its local type, then the composition satisfies the global type. The global type in (4) is *projected* into the three local types:

Customer’s endpoint type: $\text{Agency!details}(\text{nat}).\text{Hotel?ok}(\text{bool}).\text{end}$
 Agency’s endpoint type: $\text{Customer?details}(\text{nat}).\text{Hotel!details}(\text{nat}).\text{end}$
 Hotel’s endpoint type: $\text{Agency?details}(\text{nat}).\text{Customer!ok}(\text{bool}).\text{end}$

where $\text{Agency!details}(\text{nat})$ means “send to **Agency** a *details* message,” and $\text{Hotel?ok}(\text{bool})$ means “receive from **Hotel** an *ok* message.” Then each process is type-checked against its own endpoint type. When the three processes are executed, their interactions automatically follow the stipulated scenario.

If instead we only used three separate binary session types to describe the message exchanges between **Customer** and **Agency**, between **Agency** and **Hotel**, and between **Hotel** and **Customer**, respectively, without using a global type, then we lose essential sequencing information in this interaction scenario. Consequently, we can no longer guarantee deadlock-freedom among these three parties. Since the three separate binary sessions can be interleaved freely, an implementation of the **Customer** that conforms to $\text{Hotel?ok}(\text{bool}).\text{Agency!details}(\text{nat}).\text{end}$ becomes typable. This causes the situation in which each of the three parties blocks indefinitely while waiting for a message to be delivered.

In what follows we will start from giving more examples in Sect. 2, which will be used as running examples throughout the whole paper. In Sect. 3 we will introduce a formal process calculus powerful enough to describe multiparty communication protocols, such as the ones from above. Finally in Sect. 4, we will introduce a type system; we will show how multiparty session types work by examples and present the key properties of typed processes.

2 Examples, Intuitively

Multiparty session types are aimed at enabling compositional reasoning about communication among different parties. In this section we present informally some essential communication protocols that can be handled with this methodology.

Let us for example consider the basic protocol from the previous section. This simple protocol for a travel agency is synthetically displayed by Fig. 1a. A more interesting protocol could be the one shown in Fig. 1b. Here the customer is allowed to try *again and again* to obtain the new quote, should they not like the first one. We will see that our calculus (Sect. 3) is endowed with recursion, which can model an indefinite number of iterations.

Let us consider a specification in Fig. 2, which describes a simple communication protocol among three parties. Bob is waiting for a message from Alice, *non-deterministically* allowing for messages with two different labels (*I1* and *I2*). Alice sends the message with label *I1* and the communication continues with Bob sending a message `message1` to Carol and finally Carol returning a message `answer` to Alice, which depends on the previous communication, including the choice of the label for the original message from Alice to Bob.

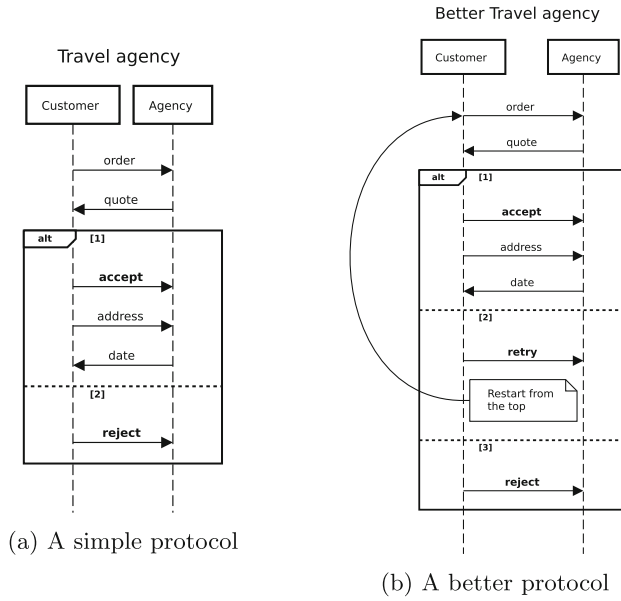


Fig. 1. Communication protocols for a travel agency

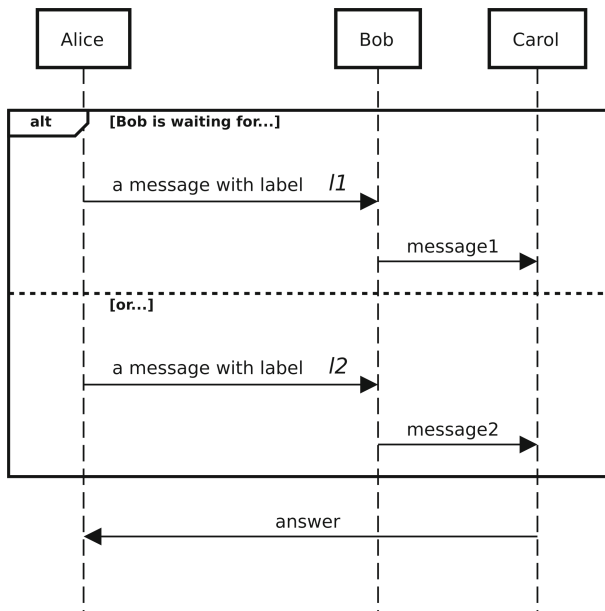


Fig. 2. Multiparty communication protocol

In the following sections, we model these protocols by a simple multiparty session calculus, and give types to processes modelling these protocols.

3 Synchronous Multiparty Session Calculus

This section introduces the syntax and semantics of a synchronous multiparty session calculus from [7], which simplifies the calculus in [14] by eliminating both shared channels for session initiations and session channels for communications inside sessions.

Notation 01 (Base sets). *We use the following base sets: values, ranged over by v, v', \dots ; expressions, ranged over by e, e', \dots ; expression variables, ranged over by x, y, z, \dots ; labels, ranged over by ℓ, ℓ', \dots ; session participants, ranged over by p, q, \dots ; process variables, ranged over by X, Y, \dots ; processes, ranged over by P, Q, \dots ; and multiparty sessions, ranged over by $\mathcal{M}, \mathcal{M}', \dots$.*

Syntax. A value v can be a natural number n , an integer i , or a boolean `true/false`. An expression e can be a variable, a value, or a term built from expressions by applying the operators `succ`, `neg`, `¬`, `⊕`, or the relation `>`. An *evaluation context* \mathcal{E} is an expression with exactly one hole. The operator `⊕` models non-determinism: $e_1 \oplus e_2$ is an expression that might yield either e_1 or e_2 .

The processes of the synchronous multiparty session calculus are defined by:

$$P ::= p!\ell\langle e \rangle.P \mid \sum_{i \in I} p?\ell_i(x_i).P_i \mid \text{if } e \text{ then } P \text{ else } Q \mid \mu X.P \mid X \mid \mathbf{0}$$

The output process $p!\ell\langle e \rangle.Q$ sends the value of expression e with label ℓ to participant p . The sum of input processes (external choice) $\sum_{i \in I} p?\ell_i(x_i).P_i$ is a process that can accept a value with label ℓ_i from participant p , for any $i \in I$. According to the label ℓ_i of the received value, the variable x_i is instantiated with the value in the continuation process P_i . We assume that the set I is always finite and non-empty.

The conditional process `if e then P else Q` represents the internal choice between processes P and Q . Which branch of the conditional process will be taken depends on the evaluation of the expression e . The process $\mu X.P$ is a recursive process. We assume that the recursive processes are *guarded*. For example, $\mu X.p?\ell(x).X$ is a valid process, while $\mu X.X$ is not. We often omit $\mathbf{0}$ from the tail of processes.

We define a *multiparty session* as a parallel composition of pairs (denoted by $p \triangleleft P$) of participants and processes:

$$\mathcal{M} ::= p \triangleleft P \mid \mathcal{M} \mid \mathcal{M}$$

with the intuition that process P plays the role of participant p , and can interact with other processes playing other roles in \mathcal{M} . A multiparty session is *well formed* if all its participants are different. We consider only well-formed multiparty sessions.

Example 1. We now show how to encode, in the above calculus, processes respecting the protocols informally presented in Fig. 1. Note that we picked one different label for each destination (that can be either Hawaii or France), as well as for each action (as “accept”, “reject”, ...).

Let us start with processes in Fig. 1a.

$$\begin{aligned}
 P'_{\text{Customer}} &= \text{Agency!Hawaii}\langle\text{true}\rangle.\text{Agency?quote}(x). \\
 &\quad \text{if } (x > 1000) \\
 &\quad \text{then Agency!reject}\langle\text{true}\rangle.\mathbf{0} \\
 &\quad \text{else Agency!accept}\langle\text{true}\rangle.\text{Agency!address}\langle 42\rangle.\text{Agency?date}(y).\mathbf{0} \\
 P'_{\text{Agency}} &= \text{Customer?Hawaii}(x).\text{Customer!quote}\langle 5000\rangle. \\
 &\quad (\text{Customer?accept}(y).\text{Customer?address}(z). \\
 &\quad \quad \text{Customer!date}\langle 25122019\rangle.\mathbf{0} \\
 &\quad + \text{Customer?reject}(y).\mathbf{0}) \\
 &+ \text{Customer?France}(x).\text{Customer!quote}\langle 1000\rangle. \\
 &\quad (\text{Customer?accept}(y).\text{Customer?address}(z). \\
 &\quad \quad \text{Customer!date}\langle 25122019\rangle.\mathbf{0} \\
 &\quad + \text{Customer?reject}(y).\mathbf{0})
 \end{aligned}$$

The customer here would ask for Hawaii as a destination, they will receive a quote that will be too expensive ($5000 > 1000$) and, thus, they will end up rejecting the offer (via sending `true` with label `reject`).

In what follows instead we will extend the above processes, giving to the customer the opportunity to retry, as suggested by the diagram in Fig. 1b.

$$\begin{aligned}
 P_{\text{Customer}} &= \text{Agency!Hawaii}\langle\text{true}\rangle.\text{Agency?quote}(x). \\
 &\quad \text{if } (x > 1000) \\
 &\quad \text{then Agency!retry}\langle\text{true}\rangle.\text{Agency!France}\langle\text{true}\rangle.\text{Agency?quote}(y). \\
 &\quad \quad \text{if } (y > 1000) \\
 &\quad \quad \text{then Agency!reject}\langle\text{true}\rangle.\mathbf{0} \\
 &\quad \quad \text{else Agency!accept}\langle\text{true}\rangle.\text{Agency!address}\langle 42\rangle.\text{Agency?date}(z).\mathbf{0} \\
 &\quad \text{else Agency!accept}\langle\text{true}\rangle.\text{Agency!address}\langle 42\rangle.\text{Agency?date}(y).\mathbf{0} \\
 P_{\text{Agency}} &= \mu X. \text{Customer?Hawaii}(x).\text{Customer!quote}\langle 5000\rangle. \\
 &\quad (\text{Customer?accept}(y).\text{Customer?address}(z). \\
 &\quad \quad \text{Customer!date}\langle 25122019\rangle.\mathbf{0} \\
 &\quad + \text{Customer?retry}(y).X + \text{Customer?reject}(y).\mathbf{0}) \\
 &+ \text{Customer?France}(x).\text{Customer!quote}\langle 1000\rangle. \\
 &\quad (\text{Customer?accept}(y).\text{Customer?address}(z). \\
 &\quad \quad \text{Customer!date}\langle 25122019\rangle.\mathbf{0} \\
 &\quad + \text{Customer?retry}(y).X + \text{Customer?reject}(y).\mathbf{0})
 \end{aligned}$$

In the example above, we have highlighted the syntactic construct allowing the customer to retry the purchase and obtain a new quote. In P_{Customer} the only occurrence of the label `retry` is within a deterministic choice (if - then - else - construct): if the quote is too high, the customer will communicate to the agency that they would like a different quote for a different destination. More interesting

Table 1. Expression evaluation.

$$\begin{array}{c}
\text{succ}(n) \downarrow (n + 1) \quad \text{neg}(i) \downarrow (-i) \quad \neg \text{true} \downarrow \text{false} \quad \neg \text{false} \downarrow \text{true} \quad v \downarrow v \\
(i_1 > i_2) \downarrow \left\{ \begin{array}{l} \text{true} \quad \text{if } i_1 > i_2, \\ \text{false} \quad \text{otherwise} \end{array} \right. \quad \frac{e_1 \downarrow v}{e_1 \oplus e_2 \downarrow v} \quad \frac{e_2 \downarrow v}{e_1 \oplus e_2 \downarrow v} \quad \frac{e \downarrow v \quad \mathcal{E}(v) \downarrow v'}{\mathcal{E}(e) \downarrow v'}
\end{array}$$

is how P_{Agency} handles the *retry* request, namely by recursion on X : the recursive call is activated each time the agency receives a *retry* request from the customer.

Operational Semantics. *The value v of expression e* (notation $e \downarrow v$) is computed as expected, see Table 1. The successor operation **succ** is defined only on natural numbers, the negation **neg** is defined on integers, and \neg is defined only on boolean values. The internal choice $e_1 \oplus e_2$ evaluates either to the value of e_1 or to the value of e_2 .

The *computational rules of multiparty sessions* are given in Table 3. They are closed with respect to the structural congruence defined in Table 2. In rule [R-COMM], the participant q sends the value v choosing the label ℓ_j to participant p , who offers inputs on all labels ℓ_i with $i \in I$. In rules [T-CONDITIONAL] and [F-CONDITIONAL], the participant p chooses to continue as P if the condition e evaluates to **true** and as Q if e evaluates to **false**. Rule [R-STRUCT] states that

Table 2. Structural congruence.

$$\begin{array}{c}
\text{[S-REC]} \quad \mu X.P \equiv P\{\mu X.P/X\} \quad \text{[S-MULTI]} \quad P \equiv Q \Rightarrow p \triangleleft P \mid \mathcal{M} \equiv p \triangleleft Q \mid \mathcal{M} \\
\text{[S-PAR 1]} \quad p \triangleleft \mathbf{0} \mid \mathcal{M} \equiv \mathcal{M} \quad \text{[S-PAR 2]} \quad \mathcal{M} \mid \mathcal{M}' \equiv \mathcal{M}' \mid \mathcal{M} \quad \text{[S-PAR 3]} \quad (\mathcal{M} \mid \mathcal{M}') \mid \mathcal{M}'' \equiv \mathcal{M} \mid (\mathcal{M}' \mid \mathcal{M}'')
\end{array}$$

Table 3. Reduction rules.

$$\begin{array}{c}
\text{[R-COMM]} \quad \frac{j \in I \quad e \downarrow v}{p \triangleleft \sum_{i \in I} q? \ell_i(x).P_i \mid q \triangleleft p! \ell_j(e).Q \mid \mathcal{M} \longrightarrow p \triangleleft P_j\{v/x\} \mid q \triangleleft Q \mid \mathcal{M}} \\
\text{[T-CONDITIONAL]} \quad \frac{e \downarrow \text{true}}{p \triangleleft \text{if } e \text{ then } P \text{ else } Q \mid \mathcal{M} \longrightarrow p \triangleleft P \mid \mathcal{M}} \quad \text{[F-CONDITIONAL]} \quad \frac{e \downarrow \text{false}}{p \triangleleft \text{if } e \text{ then } P \text{ else } Q \mid \mathcal{M} \longrightarrow p \triangleleft Q \mid \mathcal{M}} \\
\text{[R-STRUCT]} \quad \frac{\mathcal{M}'_1 \equiv \mathcal{M}_1 \quad \mathcal{M}_1 \longrightarrow \mathcal{M}_2 \quad \mathcal{M}_2 \equiv \mathcal{M}'_2}{\mathcal{M}'_1 \longrightarrow \mathcal{M}'_2}
\end{array}$$

the reduction relation is closed with respect to structural congruence. We use \longrightarrow^* with the standard meaning.

We adopt some standard conventions regarding the syntax of processes and sessions. Namely, we will use $\prod_{i \in I} \mathbf{p}_i \triangleleft P_i$ as short for $\mathbf{p}_1 \triangleleft P_1 \mid \dots \mid \mathbf{p}_n \triangleleft P_n$, where $I = \{1, \dots, n\}$. We will sometimes use infix notation for external choice process. For example, instead of $\sum_{i \in \{1,2\}} \mathbf{p}^? \ell_i(x).P_i$, we will write $\mathbf{p}^? \ell_1(x).P_1 + \mathbf{p}^? \ell_2(x).P_2$.

Example 2. We now show the operational semantics in action. Consider the following multiparty session with three participants, **Alice**, **Bob** and **Carol** :

$$\mathcal{M} = \mathbf{Alice} \triangleleft P_{\mathbf{Alice}} \mid \mathbf{Bob} \triangleleft P_{\mathbf{Bob}} \mid \mathbf{Carol} \triangleleft P_{\mathbf{Carol}}$$

where

$$\begin{aligned} P_{\mathbf{Alice}} &= \mathbf{Bob}! \ell_1 \langle 50 \rangle. \mathbf{Carol}^? \ell_3(x). \mathbf{0} \\ P_{\mathbf{Bob}} &= \mathbf{Alice}^? \ell_1(x). \mathbf{Carol}! \ell_2 \langle 100 \rangle. \mathbf{0} + \mathbf{Alice}^? \ell_4(x). \mathbf{Carol}! \ell_2 \langle 2 \rangle. \mathbf{0} \\ P_{\mathbf{Carol}} &= \mathbf{Bob}^? \ell_2(x). \mathbf{Alice}! \ell_3 \langle \text{succ}(x) \rangle. \mathbf{0} \end{aligned}$$

This multiparty session reduces to

$$\mathbf{Alice} \triangleleft \mathbf{0} \mid \mathbf{Bob} \triangleleft \mathbf{0} \mid \mathbf{Carol} \triangleleft \mathbf{0}$$

after three communications occur. First, **Alice** sends to **Bob** natural number 50 with the label ℓ_1 . **Bob** is able to receive values with labels ℓ_1 and ℓ_4 . Next, the only possible communication is between **Bob** and **Carol**. So, **Carol** receives natural number 100 from **Bob**. The value 100 is substituted in the continuation process. Finally, since $\text{succ}(100) \downarrow 101$, **Carol** sends 101 to **Alice**. We can then reduce the session to, for example, $\mathbf{Alice} \triangleleft \mathbf{0}$, but not further.

Exercise 1. Prove the following:

1. $\mathbf{Customer} \triangleleft P'_{\mathbf{Customer}} \mid \mathbf{Agency} \triangleleft P'_{\mathbf{Agency}}$ reduces to $\mathbf{Customer} \triangleleft \mathbf{0} \mid \mathbf{Agency} \triangleleft \mathbf{0}$;
2. $\mathbf{Customer} \triangleleft P_{\mathbf{Customer}} \mid \mathbf{Agency} \triangleleft P_{\mathbf{Agency}}$ reduces to $\mathbf{Customer} \triangleleft \mathbf{0} \mid \mathbf{Agency} \triangleleft \mathbf{0}$.

From the end of Example 2, we can see that a session \mathcal{M} always has at least one participant, since we do not have neutral element for the parallel composition. In Sect. 4, we will introduce a type system ensuring that if a well-typed multiparty session has only one participant, then the corresponding process is $\mathbf{0}$ — hence, the participant's process has no inputs/outputs to perform.

The most crucial property is that when a multiparty session contains communications that will never be executed.

Definition 1. A multiparty session \mathcal{M} is stuck if $\mathcal{M} \not\equiv \mathbf{p} \triangleleft \mathbf{0}$ and there is no multiparty session \mathcal{M}' such that $\mathcal{M} \longrightarrow \mathcal{M}'$. A multiparty session \mathcal{M} gets stuck, notation $\text{stuck}(\mathcal{M})$, if it reduces to a stuck multiparty session.

The multiparty session \mathcal{M} in Example 2 does not get stuck. A similar multiparty session, where instead of $P_{\mathbf{Alice}}$ we take $P'_{\mathbf{Alice}} = \mathbf{Bob}! \ell_1 \langle 50 \rangle. \mathbf{Carol}^? \ell_5(x). \mathbf{0}$, gets stuck because of label mismatch.

4 Type System

This section introduces a type system for the calculus presented in Sect. 3 (the formulation is based on [7]). We formalise types and projections (Sect. 4.1), the subtyping relation (Sect. 4.2), and the typing rules and their properties (Sect. 4.3). All stated results in this paper are proved in [7].

4.1 Types and Projections

Global types provide global conversation scenarios of multiparty sessions, with a bird's eye view describing the message exchanges between pairs of participants.

Definition 2 (Sorts and global types). Sorts, ranged over by S , are defined as:

$$S ::= \text{nat} \mid \text{int} \mid \text{bool}$$

Global types, ranged over by G , are terms generated by the following grammar:

$$G ::= \text{end} \mid \mu\mathbf{t}.G \mid \mathbf{t} \mid \mathbf{p} \rightarrow \mathbf{q} : \{\ell_i(S_i).G_i\}_{i \in I}$$

We require that $\mathbf{p} \neq \mathbf{q}$, $I \neq \emptyset$, and $\ell_i \neq \ell_j$ whenever $i \neq j$, for all $i, j \in I$. We postulate that recursion is guarded. Unless otherwise noted, global types are closed: a recursion variable \mathbf{t} only occurs bounded by $\mu\mathbf{t} \dots$

In Definition 2, the type $\mathbf{p} \rightarrow \mathbf{q} : \{\ell_i(S_i).G_i\}_{i \in I}$ formalises a protocol where participant \mathbf{p} must send to \mathbf{q} one message with label ℓ_i and a value of type S_i as payload, for some $i \in I$; then, depending on which ℓ_i was sent by \mathbf{p} , the protocol continues as G_i . Value types are restricted to sorts, that can be natural (nat), integer (int) and boolean (bool). The type end represents a terminated protocol. Recursive protocol is modelled as $\mu\mathbf{t}.G$, where recursion variable \mathbf{t} is bound and guarded in G — e.g., $\mu\mathbf{t}.\mathbf{p} \rightarrow \mathbf{q} : \ell(\text{nat}).\mathbf{t}$ is a valid global type, whereas $\mu\mathbf{t}.\mathbf{t}$ is not. We take the equi-recursive viewpoint, i.e. we identify $\mu\mathbf{t}.G$ and $G\{\mu\mathbf{t}.G/\mathbf{t}\}$.

We define the *set of participants of a global type* G , by structural induction on G , as follows:

$$\begin{aligned} \text{pt}\{\mu\mathbf{t}.G\} &= \text{pt}\{G\} & \text{pt}\{\text{end}\} &= \text{pt}\{\mathbf{t}\} = \emptyset \\ \text{pt}\{\mathbf{p} \rightarrow \mathbf{q} : \{\ell_i(S_i).G_i\}_{i \in I}\} &= \{\mathbf{p}, \mathbf{q}\} \cup \text{pt}\{G_i\} \quad (i \in I) \end{aligned}$$

We will often write $\mathbf{p} \in G$ instead of $\mathbf{p} \in \text{pt}\{G\}$.

A *local session type* describes the behaviour of a single participant in a multiparty session.

Definition 3 (Local Session Types). The grammar of session types, ranged over by T , is:

$$T ::= \text{end} \mid \&_{i \in I} \mathbf{p} ? \ell_i(S_i).T_i \mid \oplus_{i \in I} \mathbf{q} ! \ell_i(S_i).T_i \mid \mu\mathbf{t}.T \mid \mathbf{t}$$

We require that $\ell_i \neq \ell_j$ whenever $i \neq j$, for all $i, j \in I$. We postulate that recursion is always guarded. Unless otherwise noted, session types are closed.

Note that, according to the previous definition, labels in a type need to be pairwise different. For example, $\mathbf{p}?\ell(\mathbf{int}).\mathbf{end}\&\mathbf{p}?\ell(\mathbf{nat}).\mathbf{end}$ is not a type.

The session type \mathbf{end} says that no further communication is possible and the protocol is completed. The *external choice* or *branching type* $\&_{i \in I} \mathbf{p}?\ell_i(S_i).T_i$ requires to wait to receive a value of sort S_i (for some $i \in I$) from the participant \mathbf{p} , via a message with label ℓ_i ; if the received message has label ℓ_i , the protocol will continue as prescribed by T_i . The *internal choice* or *selection type* $\oplus_{i \in I} \mathbf{q}!\ell_i(S_i).T_i$ says that the participant implementing the type must choose a labelled message to send to \mathbf{q} ; if the participant chooses the message ℓ_i , for some $i \in I$, it must include in the message to \mathbf{q} a payload value of sort S_i , and continue as prescribed by T_i . Recursion is modelled by the session type $\mu\mathbf{t}.T$. We adopt the following conventions: we do not write branch/selection symbols in case of a singleton choice, we do not write unnecessary parentheses, and we often omit trailing \mathbf{ends} .

The set $\mathbf{pt}\{T\}$ of participants of a session type T is defined inductively as follows

$$\begin{aligned} \mathbf{pt}\{\&_{i \in I} \mathbf{p}?\ell_i(S_i).T_i\} &= \mathbf{pt}\{\oplus_{i \in I} \mathbf{p}!\ell_i(S_i).T_i\} = \{\mathbf{p}\} \cup \bigcup_{i \in I} \mathbf{pt}\{T_i\} \\ \mathbf{pt}\{\mu\mathbf{t}.T\} &= \mathbf{pt}\{T\} \quad \mathbf{pt}\{\mathbf{t}\} = \mathbf{pt}\{\mathbf{end}\} = \emptyset. \end{aligned}$$

In what follows we introduce the concept of *projection* of a global type onto a participant, didactically into two steps: first we give the simpler version of the projection from [11, 12], then we extend it with the *merging* operation. The second definition will extend the domain of the projection as a partial function on global types, i.e., the second version of the projection will be well-defined on a wider range of global types.

Definition 4. *The projection (no-merging version) of a global type onto a participant r is defined by recursion on G :*

- $\mathbf{end}|r = \mathbf{end}$; [PROJ-END]
- $\mathbf{t}|r = \mathbf{t}$; [PROJ-VAR]
- $(\mu\mathbf{t}.G)|r = \mu\mathbf{t}.(G|r)$ if $r \in \mathbf{pt}\{G\}$; [PROJ-REC-1]
- $(\mu\mathbf{t}.G)|r = \mathbf{end}$ if $r \notin \mathbf{pt}\{G\}$; [PROJ-REC-2]
- $\mathbf{p} \rightarrow r : \{\ell_i(S_i).G_i\}_{i \in I}|r = \&_{i \in I} \mathbf{p}?\ell_i(S_i).G_i|r$; [PROJ-IN]
- $r \rightarrow \mathbf{q} : \{\ell_i(S_i).G_i\}_{i \in I}|r = \oplus_{i \in I} \mathbf{q}!\ell_i(S_i).G_i|r$; [PROJ-OUT]
- $\mathbf{p} \rightarrow \mathbf{q} : \{\ell_i(S_i).G_i\}_{i \in I}|r = G_{i_0}|r$ where $i_0 \in I$, [PROJ-CONT']
if $r \notin \{\mathbf{p}, \mathbf{q}\}$ and, for all $i, j \in I$, $G_i|r = G_j|r$;
- *undefined otherwise.*

We describe the clauses of Definition 4:

[PROJ-END, PROJ-VAR] give the behaviour of projections on \mathbf{end} and type variables;

- [PROJ-REC-1, PROJ-REC-2] give the behaviour of projections on recursive types; in particular [PROJ-REC-2] is needed: if we applied only [PROJ-REC-1] to any recursive type, we will obtain $(\mu \mathbf{t}. \mathbf{p} \rightarrow \mathbf{q} : \ell(\mathbf{nat}). \mathbf{t}) \upharpoonright r = \mu \mathbf{t}. \mathbf{t}$, namely we will allow for unguarded occurrences of \mathbf{t} , which we do not accept as valid;
- [PROJ-IN] (**resp.** [PROJ-OUT]) states that a global type \mathbf{G} starting with a communication from \mathbf{p} to \mathbf{r} (resp. from \mathbf{r} to \mathbf{q}) projects onto an external (resp. internal) choice $\&_{i \in I} \mathbf{p} ? \ell_i(S_i). \mathbf{G}_i \upharpoonright r$ (resp. $\oplus_{i \in I} \mathbf{q} ! \ell_i(S_i). \mathbf{G}_i \upharpoonright r$), provided that the continuations of $\&_{i \in I} \mathbf{p} ? \ell_i(S_i). \mathbf{G}_i \upharpoonright r$ (resp. $\oplus_{i \in I} \mathbf{q} ! \ell_i(S_i). \mathbf{G}_i \upharpoonright r$) are also projections of the corresponding global type continuations $\mathbf{G}_i \upharpoonright r$.
- [PROJ-CONT'] states that if \mathbf{G} starts with a communication between \mathbf{p} and \mathbf{q} , and we are projecting \mathbf{G} onto a third participant \mathbf{r} , then we need to make sure that continuation is *the same* on all branches; just below we will see how this restriction can be relaxed.

Example 3. We now show an example of some projections of global types. Consider the global type (where \mathbf{p} , \mathbf{q} and \mathbf{r} are pairwise distinct):

$$\mathbf{G} = \mathbf{p} \rightarrow \mathbf{q} : \{\ell_1(\mathbf{nat}). \mathbf{G}_1, \ell_2(\mathbf{bool}). \mathbf{G}_1\} \text{ where } \mathbf{G}_1 = \mathbf{q} \rightarrow \mathbf{r} : \{\ell_3(\mathbf{int}), \ell_4(\mathbf{nat})\}$$

We have:

$$\begin{aligned} \mathbf{G} \upharpoonright \mathbf{p} &= \mathbf{q} ! \ell_1(\mathbf{nat}). (\mathbf{G}_1 \upharpoonright \mathbf{p}) \oplus \mathbf{q} ! \ell_2(\mathbf{bool}). (\mathbf{G}_1 \upharpoonright \mathbf{p}) \\ &= \mathbf{q} ! \ell_1(\mathbf{nat}). \mathbf{end} \oplus \mathbf{q} ! \ell_2(\mathbf{bool}). \mathbf{end} \end{aligned}$$

$$\begin{aligned} \mathbf{G} \upharpoonright \mathbf{q} &= \mathbf{p} ? \ell_1(\mathbf{nat}). (\mathbf{G}_1 \upharpoonright \mathbf{q}) \& \mathbf{p} ? \ell_2(\mathbf{bool}). (\mathbf{G}_1 \upharpoonright \mathbf{q}) \\ &= \mathbf{p} ? \ell_1(\mathbf{nat}). (\mathbf{r} ! \ell_3(\mathbf{int}) \oplus \mathbf{r} ! \ell_4(\mathbf{nat})) \& \mathbf{p} ? \ell_2(\mathbf{bool}). (\mathbf{r} ! \ell_3(\mathbf{int}) \oplus \mathbf{r} ! \ell_4(\mathbf{nat})) \end{aligned}$$

$$\mathbf{G} \upharpoonright \mathbf{r} = \mathbf{G}_1 \upharpoonright \mathbf{r} = (\mathbf{q} ? \ell_3(\mathbf{int}) \& \mathbf{q} ? \ell_4(\mathbf{nat}))$$

Note that $\mathbf{G} \upharpoonright \mathbf{r}$ is well defined only because the continuation of the communication after the exchange $\mathbf{p} \rightarrow \mathbf{q}$ is equal, namely \mathbf{G}_1 , in both branches.

In the following Definition 5, we give a more permissive definition of projection, that handles also cases in which the continuation is not the same in all branches, but the types are somehow “compatible”, namely they can be merged. Our definition follows [7], which extends [11, 12], along the lines of [19] and [2]. i.e., it uses a *merging operator* \sqcap .

Definition 5. *The projection of a global type onto a participant \mathbf{r} is defined by recursion on \mathbf{G} :*

- $\mathbf{end} \upharpoonright \mathbf{r} = \mathbf{end}$; [PROJ-END]
- $\mathbf{t} \upharpoonright \mathbf{r} = \mathbf{t}$; [PROJ-VAR]
- $(\mu \mathbf{t}). \mathbf{G} \upharpoonright \mathbf{r} = \mu \mathbf{t}. (\mathbf{G} \upharpoonright \mathbf{r})$ if $\mathbf{r} \in \mathbf{pt}\{\mathbf{G}\}$; [PROJ-REC-1]
- $(\mu \mathbf{t}). \mathbf{G} \upharpoonright \mathbf{r} = \mathbf{end}$ if $\mathbf{r} \notin \mathbf{pt}\{\mathbf{G}\}$; [PROJ-REC-2]
- $\mathbf{p} \rightarrow \mathbf{r} : \{\ell_i(S_i). \mathbf{G}_i\}_{i \in I} \upharpoonright \mathbf{r} = \&_{i \in I} \mathbf{p} ? \ell_i(S_i). \mathbf{G}_i \upharpoonright \mathbf{r}$; [PROJ-IN]
- $\mathbf{r} \rightarrow \mathbf{q} : \{\ell_i(S_i). \mathbf{G}_i\}_{i \in I} \upharpoonright \mathbf{r} = \oplus_{i \in I} \mathbf{q} ! \ell_i(S_i). \mathbf{G}_i \upharpoonright \mathbf{r}$; [PROJ-OUT]
- $\mathbf{p} \rightarrow \mathbf{q} : \{\ell_i(S_i). \mathbf{G}_i\}_{i \in I} \upharpoonright \mathbf{r} = \sqcap_{i \in I} (\mathbf{G}_i \upharpoonright \mathbf{r})$ if $\mathbf{r} \notin \{\mathbf{p}, \mathbf{q}\}$; [PROJ-CONT]
- *undefined otherwise.*

Above, \sqcap is the merging operator, that is a partial operation over session types defined as:

$$\mathsf{T}_1 \sqcap \mathsf{T}_2 = \begin{cases} \mathsf{T}_1 & \text{if } \mathsf{T}_1 = \mathsf{T}_2 & \text{[MRG-ID]} \\ \mathsf{T}_3 & \text{if } \exists I, J : \begin{cases} \mathsf{T}_1 = \&x_{i \in I} \mathsf{p}'? \ell_i(S_i). \mathsf{T}_i & \text{and} \\ \mathsf{T}_2 = \&x_{j \in J} \mathsf{p}''? \ell_j(S_j). \mathsf{T}_j & \text{and} \\ \mathsf{T}_3 = \&x_{k \in I \cup J} \mathsf{p}'? \ell_k(S_k). \mathsf{T}_k \end{cases} & \text{[MRG-BRA]} \\ \text{undefined otherwise.} \end{cases}$$

Proposition 1. *The merging operation is associative, i.e.: $\mathsf{T} \sqcap (\mathsf{T}' \sqcap \mathsf{T}'') = (\mathsf{T} \sqcap \mathsf{T}') \sqcap \mathsf{T}''$.*

By Definition 5, merging a type with itself results in itself (rule [MRG-ID]). Moreover, Definition 5 allows to combine different external choices (rule [MRG-ID]) *if and only if* common labels have identical sorts and identical continuations, as formalised in Proposition 2 below and illustrated in Examples 4 to 6.

Proposition 2. *For two types $\mathsf{T}' = \&x_{i \in I} \mathsf{p}'? \ell_i(S_i). \mathsf{T}_i$ and $\mathsf{T}'' = \&x_{j \in J} \mathsf{p}''? \ell_j(S_j). \mathsf{T}_j$, we have that $\mathsf{T}' \sqcap \mathsf{T}''$ is defined if and only if $\mathsf{p}' = \mathsf{p}''$ and, whenever $\ell_i = \ell_j$ (for some $i \in I$ and $j \in J$), $S_i = S_j$ and $\mathsf{T}_i = \mathsf{T}_j$.*

Example 4. We now give some small examples that illustrate the definition of the merging operator (here, $i \neq j$ implies $\ell_i \neq \ell_j$):

$$\begin{aligned} \mathsf{q}! \ell(\mathsf{nat}) \sqcap \mathsf{q}! \ell(\mathsf{nat}) &= \mathsf{q}! \ell(\mathsf{nat}) \\ \mathsf{p}! \ell(\mathsf{nat}) \sqcap \mathsf{q}! \ell(\mathsf{nat}) &\quad \text{undefined: outputs to different participants} \\ \mathsf{q}! \ell_3(\mathsf{nat}) \sqcap \mathsf{q}! \ell_4(\mathsf{nat}) &\quad \text{undefined: outputs with different labels} \\ (\mathsf{q}? \ell_3(\mathsf{int}) \&\mathsf{q}? \ell_5(\mathsf{nat})) \sqcap (\mathsf{q}? \ell_4(\mathsf{int}) \&\mathsf{q}? \ell_5(\mathsf{nat})) \\ &= \mathsf{q}? \ell_3(\mathsf{int}) \&\mathsf{q}? \ell_4(\mathsf{int}) \&\mathsf{q}? \ell_5(\mathsf{nat}) \\ \mathsf{q}? \ell_3(\mathsf{nat}) \sqcap \mathsf{q}? \ell_3(\mathsf{nat}). \mathsf{q}? \ell_3(\mathsf{nat}) & \\ &\quad \text{undefined: same prefixes, but different continuations} \\ \mathsf{q}? \ell(\mathsf{nat}) \sqcap \mathsf{q}? \ell(\mathsf{int}) &\quad \text{undefined: the payload sorts do not match} \end{aligned}$$

Now we understand better how clause [PROJ-CONT] works: it states that if G starts with a communication between p and q , and we are projecting G onto a third participant r , then we need to (1) skip the initial communication, (2) project all the continuations onto r , and (3) *merge* the resulting session types, using the *merging operator* \sqcap .

As a result, clause [PROJ-CONT] of Definition 5 allows participant r to receive different messages (from a same participant p') in different branches of a global type, as shown in Example 5 below.

Example 5. We demonstrate interesting points of Definition 5. First, we show some projections of global types. Consider the global type:

$$\mathsf{G} = \mathsf{p} \rightarrow \mathsf{q} : \{\ell_1(\mathsf{nat}). \mathsf{G}_1, \ell_2(\mathsf{bool}). \mathsf{G}_2\} \quad \text{where} \quad \begin{cases} \mathsf{G}_1 = \mathsf{q} \rightarrow \mathsf{r} : \{\ell_3(\mathsf{int}), \ell_5(\mathsf{nat})\} \\ \mathsf{G}_2 = \mathsf{q} \rightarrow \mathsf{r} : \{\ell_4(\mathsf{int}), \ell_5(\mathsf{nat})\} \\ \mathsf{r} \neq \mathsf{p} \end{cases}$$

We have:

$$\begin{aligned}
G \upharpoonright p &= q!l_1(\text{nat}).(G_1 \upharpoonright p) \oplus q!l_2(\text{bool}).(G_2 \upharpoonright p) \\
&= q!l_1(\text{nat}).\text{end} \oplus q!l_2(\text{bool}).\text{end} \\
G \upharpoonright q &= p?l_1(\text{nat}).(G_1 \upharpoonright q) \& p?l_2(\text{bool}).(G_2 \upharpoonright q) \\
&= p?l_1(\text{nat}).(r!l_3(\text{int}) \oplus r!l_5(\text{nat})) \& p?l_2(\text{bool}).(r!l_4(\text{int}) \oplus r!l_5(\text{nat})) \\
G \upharpoonright r &= G_1 \upharpoonright r \sqcap G_2 \upharpoonright r = (q?l_3(\text{int})\&q?l_5(\text{nat})) \sqcap (q?l_4(\text{int})\&q?l_5(\text{nat})) \\
&= q?l_3(\text{int})\&q?l_4(\text{int})\&q?l_5(\text{nat})
\end{aligned}$$

Note that in G , q could output different messages towards r , depending on whether p sends l_1 or l_2 to q ; therefore, in $G \upharpoonright r$, the possible inputs of r in G_1 and G_2 are merged into a larger external choice that supports all possible outputs of q .

Importantly, by Definition 5, there exist global types that *cannot* be projected onto all their participants. This is because G might describe meaningless protocols, that cause the merging operation \sqcap in clause [PROJ-CONT] to be undefined, as shown in Example 6 below.

Example 6. We show two global types that cannot be projected according to the Definition 5. Consider the global type $G = p \rightarrow q : \{l_1(\text{nat}).G_1, l_2(\text{bool}).G_2\}$, with $G_1 = r \rightarrow q : l_3(\text{nat})$ and $G_2 = r \rightarrow q : l_4(\text{nat})$. Then,

$$\begin{aligned}
G \upharpoonright p &= q!l_1(\text{nat}) \oplus q!l_2(\text{bool}) \\
G \upharpoonright q &= p?l_1(\text{nat}).r?l_3(\text{nat}) \& p?l_2(\text{bool}).r?l_4(\text{nat}) \\
G \upharpoonright r &= q!l_3(\text{nat}) \sqcap q!l_4(\text{nat}) \quad (\text{undefined if } l_3 \neq l_4)
\end{aligned}$$

Intuitively, when $l_3 \neq l_4$, $G \upharpoonright r$ is undefined because in G , depending on whether p and q exchange l_1 or l_2 , r is supposed to send either l_3 or l_4 to q ; however, r is not privy to the interactions between p and q , and thus, G provides an invalid specification for r . Instead, if $l_3 = l_4$, then by Definition 5 we have $G \upharpoonright r = q!l_3(\text{nat}) \sqcap q!l_3(\text{nat}) = q!l_3(\text{nat})$.

Now, consider the global type $G' = p \rightarrow q : \{l_1(\text{nat}).G'_1, l_2(\text{bool}).G'_2\}$, with $G'_1 = q \rightarrow r : l_3(\text{nat})$ and $G'_2 = q \rightarrow r : l_3(\text{nat}).q \rightarrow r : l_3(\text{nat})$. Then,

$$\begin{aligned}
G' \upharpoonright p &= q!l_1(\text{nat}) \oplus q!l_2(\text{bool}) \\
G' \upharpoonright q &= p?l_1(\text{nat}).r!l_3(\text{nat}) \& p?l_2(\text{bool}).r!l_3(\text{nat}).r!l_3(\text{nat}) \\
G' \upharpoonright r &= q?l_3(\text{nat}) \sqcap q?l_3(\text{nat}).q?l_3(\text{nat}) \quad (\text{undefined})
\end{aligned}$$

Here, $G' \upharpoonright r$ is undefined because in G' , depending on whether p and q exchange l_1 or l_2 , r is supposed to receive either one or two messages l_3 from q ; however, as in the previous example, r is not aware of the interactions between p and q , and thus, G provides an invalid specification for r . This example could be fixed, e.g., by replacing l_3 with $l' \neq l_3$ in G'_2 , or by letting $G'_1 = G'_2$: both fixes would make $G' \upharpoonright r$ defined, similarly to Example 5.

4.2 Subtyping

The subtyping relation \leq is used to augment the flexibility of the type system (introduced in Sect. 4.3): by determining when a type T is “smaller” than T' , it allows to use a process typed by the former whenever a process typed by the latter is required.

Definition 6 (Subsorting and subtyping). Subsorting \leq : is the least reflexive binary relation such that $\mathbf{nat} \leq \mathbf{int}$.

Subtyping \leq is the largest relation between session types coinductively defined by the following rules:

$$\begin{array}{c}
\text{[SUB-END]} \quad \mathbf{end} \leq \mathbf{end} \\
\\
\begin{array}{c}
\text{SUB-IN} \\
\forall i \in I : S'_i \leq S_i \quad T_i \leq T'_i \\
\hline
&_{i \in I \cup \mathbf{p}} ?\ell_i(S_i).T_i \leq &_{i \in I \mathbf{p}} ?\ell_i(S'_i).T'_i
\end{array}
\quad
\begin{array}{c}
\text{SUB-OUT} \\
\forall i \in I : S_i \leq S'_i \quad T_i \leq T'_i \\
\hline
&_{i \in I \mathbf{p}} !\ell_i(S_i).T_i \leq &_{i \in I \cup \mathbf{p}} !\ell_i(S'_i).T'_i
\end{array}
\end{array}$$

Intuitively, the session subtyping \leq in Definition 6 says that T is smaller than T' when T is “less liberal” than T' — i.e., when T allows for less internal choices, and demands to handle more external choices.¹ A peculiarity of the relation is that, apart from a pair of inactive session types, only inputs and outputs from/to a same participant can be related (with additional conditions to be satisfied). Note that the double line in the subtyping rules indicates that the rules are interpreted *coinductively* [15, Chapter 21]

[SUB-END] says that \mathbf{end} is only subtype of itself.

[SUB-IN] relates external choices from the same participant \mathbf{p} : the subtype must support all the choices of the supertype, and for each common message label, the continuations must be related, too; note that the carried sorts are contravariant: e.g., if the supertype requires to receive a message $\ell_i(\mathbf{nat})$ (for some $i \in I$), then the subtype can support $\ell_i(\mathbf{int})$ or $\ell_i(\mathbf{nat})$, since $\mathbf{nat} \leq \mathbf{int}$ and $\mathbf{nat} \leq \mathbf{nat}$.

[SUB-OUT] relates internal choices towards the same participant \mathbf{p} : the subtype must offer a subset of the choices of the supertype, and for each common message label, the continuations must be related, too; note that the carried sorts are covariant: e.g., if the supertype allows to send a message $\ell_i(\mathbf{int})$ (for some $i \in I$), then the subtype can allow to send $\ell_i(\mathbf{int})$ or $\ell_i(\mathbf{nat})$, since $\mathbf{int} \leq \mathbf{int}$ and $\mathbf{nat} \leq \mathbf{int}$.

Lemma 1. *The subtyping relation \leq is reflexive and transitive.*

¹ Readers familiar with the theory of session types might notice that our subtyping relation is inverted w.r.t. the original binary session subtyping, introduced in the works of [3, 6]. In such works, smaller types have less internal choices, and more external choices: this is because they formalise a “channel-oriented” notion of subtyping, while we adopt a “process-oriented” view. For a thorough analysis and comparison of the two approaches, see [5].

4.3 Type System

We now introduce a type system for the multiparty session calculus presented in Sect. 3. We distinguish three kinds of typing judgments:

$$\Gamma \vdash e : S \qquad \Gamma \vdash P : \mathbb{T} \qquad \vdash \mathcal{M} : \mathbb{G}$$

where Γ is the *typing environment*:

$$\Gamma ::= \emptyset \mid \Gamma, x : S \mid \Gamma, X : \mathbb{T}$$

i.e., a mapping that associates expression variables with sorts, and process variables with session types.

We say that *a multiparty session \mathcal{M} is well typed* if there is a global type \mathbb{G} such that $\vdash \mathcal{M} : \mathbb{G}$. If a multiparty session is well typed, we will sometimes write just $\vdash \mathcal{M}$.

Table 4. Typing rules for expressions.

$\Gamma \vdash n : \mathbf{nat}$	$\Gamma \vdash i : \mathbf{int}$	$\Gamma \vdash \mathbf{true} : \mathbf{bool}$	$\Gamma \vdash \mathbf{false} : \mathbf{bool}$	$\Gamma, x : S \vdash x : S$
$\frac{\Gamma \vdash e : \mathbf{nat}}{\Gamma \vdash \mathbf{succ}(e) : \mathbf{nat}}$	$\frac{\Gamma \vdash e : \mathbf{int}}{\Gamma \vdash \mathbf{neg}(e) : \mathbf{int}}$	$\frac{\Gamma \vdash e : \mathbf{bool}}{\Gamma \vdash \mathbf{neg}(e) : \mathbf{bool}}$		
$\frac{\Gamma \vdash e_1 : S \quad \Gamma \vdash e_2 : S}{\Gamma \vdash e_1 \oplus e_2 : S}$	$\frac{\Gamma \vdash e_1 : \mathbf{int} \quad \Gamma \vdash e_2 : \mathbf{int}}{\Gamma \vdash e_1 > e_2 : \mathbf{bool}}$	$\frac{\Gamma \vdash e : S \quad S \leq S'}{\Gamma \vdash e : S'}$		

The typing rules for expressions are given in Table 4, and are self-explanatory. The typing rules for processes and multiparty sessions are content of Table 5:

[T-SUB] is the *subsumption rule*: a process with type \mathbb{T} is also typed by the supertype \mathbb{T}' ;

[T-0] says that a terminated process implements the terminated session type;

[T-REC] types a recursive process $\mu X.P$ with \mathbb{T} if P can be typed as \mathbb{T} , too, by extending the typing environment with the assumption that X has type \mathbb{T} ;

[T-VAR] uses the typing environment assumption that process X has type \mathbb{T} ;

[T-INPUT-CHOICE] types a summation of input prefixes as a branching type. It requires that each input prefix targets the same participant q , and that, for all $i \in I$, each continuation process P_i is typed by the continuation type \mathbb{T}_i , having the bound variable x_i in the typing environment with sort S_i . Note that the rule implicitly requires the process labels ℓ_i to be pairwise distinct (as per Definition 3);

[T-OUT] types an output prefix with a singleton selection type, provided that the expression in the message payload has the correct sort S , and the process continuation matches the type continuation;

[T-CHOICE] types a conditional process with \mathbb{T} if its sub-processes can be typed by \mathbb{T} and expression e is boolean.

[T-SESS] types multiparty sessions, by associating typed processes to participants. It requires that the processes being composed in parallel can play as participants of a global communication protocol: hence, their types must be projections of a single global type \mathbf{G} . The condition $\text{pt}\{\mathbf{G}\} \subseteq \{\mathbf{p}_i \mid i \in I\}$ allows to also type sessions containing $\mathbf{p} \triangleleft \mathbf{0}$: this is needed to assure invariance of typing.

Example 7. We show that the multiparty session \mathcal{M} from Example 2 is well typed. Consider the following global type:

$$\begin{aligned} \mathbf{G} = \text{Alice} \rightarrow \text{Bob} : \\ \{ \ell_1(\text{nat}).\text{Bob} \rightarrow \text{Carol} : \ell_2(\text{nat}).\text{Carol} \rightarrow \text{Alice} : \ell_3(\text{nat}).\text{end}, \\ \ell_4(\text{nat}).\text{Bob} \rightarrow \text{Carol} : \ell_2(\text{nat}).\text{Carol} \rightarrow \text{Alice} : \ell_3(\text{nat}).\text{end} \}. \end{aligned}$$

We show that participants `Alice`, `Bob` and `Carol` respect the prescribed protocol \mathbf{G} , by showing that they participate in a well-typed multiparty session. Applying rules from Table 5, we derive

$$\vdash P_{\text{Alice}} : \mathbb{T}_{\text{Alice}} \quad \vdash P_{\text{Bob}} : \mathbb{T}_{\text{Bob}} \quad \vdash P_{\text{Carol}} : \mathbb{T}_{\text{Carol}}$$

where:

$$\begin{aligned} \mathbb{T}_{\text{Alice}} &= \text{Bob}!\ell_1(\text{nat}).\text{Carol}?\ell_3(\text{nat}).\text{end} \\ \mathbb{T}_{\text{Bob}} &= \text{Alice}?\ell_1(\text{nat}).\text{Carol}!\ell_2(\text{nat}).\text{end} \ \& \ \text{Alice}?\ell_4(\text{nat}).\text{Carol}!\ell_2(\text{nat}).\text{end} \\ \mathbb{T}_{\text{Carol}} &= \text{Bob}?\ell_2(\text{nat}).\text{Alice}!\ell_3(\text{nat}).\text{end} \end{aligned}$$

Now, let:

$$\mathbb{T}'_{\text{Alice}} = \text{Bob}!\ell_1(\text{nat}).\text{Carol}?\ell_3(\text{nat}).\text{end} \oplus \text{Bob}!\ell_4(\text{nat}).\text{Carol}?\ell_3(\text{nat}).\text{end}$$

Since it holds that $\mathbb{T}_{\text{Alice}} \leq \mathbb{T}'_{\text{Alice}}$, and the projections of \mathbf{G} to the participants are

$$\mathbf{G}|_{\text{Alice}} = \mathbb{T}'_{\text{Alice}} \quad \mathbf{G}|_{\text{Bob}} = \mathbb{T}_{\text{Bob}} \quad \mathbf{G}|_{\text{Carol}} = \mathbb{T}_{\text{Carol}}$$

Table 5. Typing rules for processes and sessions.

$\frac{[\text{T-SUB}]}{\Gamma \vdash P : \mathbb{T} \quad \mathbb{T} \leq \mathbb{T}'}{\Gamma \vdash P : \mathbb{T}'}$	$[\text{T-0}] \quad \Gamma \vdash \mathbf{0} : \text{end}$	$\frac{[\text{T-REC}]}{\Gamma, X : \mathbb{T} \vdash P : \mathbb{T}}{\Gamma \vdash \mu X.P : \mathbb{T}}$	$[\text{T-VAR}] \quad \Gamma, X : \mathbb{T} \vdash X : \mathbb{T}$
$\frac{[\text{T-INPUT-CHOICE}]}{\Gamma \vdash \sum_{i \in I} \mathbf{q}?\ell_i(x_i).P_i : \&_{i \in I} \mathbf{q}?\ell_i(S_i).\mathbb{T}_i} \quad \forall i \in I \quad \Gamma, x_i : S_i \vdash P_i : \mathbb{T}_i$		$\frac{[\text{T-OUT}]}{\Gamma \vdash \mathbf{q}!\ell(\mathbf{e}).P : \mathbf{q}!\ell(S).\mathbb{T}} \quad \Gamma \vdash \mathbf{e} : S \quad \Gamma \vdash P : \mathbb{T}$	
$\frac{[\text{T-CHOICE}]}{\Gamma \vdash \text{if } \mathbf{e} \text{ then } P_1 \text{ else } P_2 : \mathbb{T}} \quad \Gamma \vdash \mathbf{e} : \text{bool} \quad \Gamma \vdash P_1 : \mathbb{T} \quad \Gamma \vdash P_2 : \mathbb{T}$	$\frac{[\text{T-SESS}]}{\Gamma \vdash \prod_{i \in I} \mathbf{p}_i \triangleleft P_i : \mathbf{G}} \quad \forall i \in I \quad \vdash P_i : \mathbf{G} _{\mathbf{p}_i} \quad \text{pt}\{\mathbf{G}\} \subseteq \{\mathbf{p}_i \mid i \in I\}$		

we conclude:

$$\vdash \text{Alice} \triangleleft P_{\text{Alice}} \mid \text{Bob} \triangleleft P_{\text{Bob}} \mid \text{Carol} \triangleleft P_{\text{Carol}} : G.$$

Example 8. Let us consider processes P_{Customer} and P_{Agency} from Example 1. As a suitable global type for the session

$$\text{Customer} \triangleleft P_{\text{Customer}} \mid \text{Agency} \triangleleft P_{\text{Agency}}$$

we pick the following.

$$\begin{aligned} G &= \mu t. \text{Customer} \rightarrow \text{Agency} : \{ \text{Hawaii}(\text{bool}).G_1, \text{France}(\text{bool}).G_1 \} \\ &\text{where} \\ G_1 &= \text{Agency} \rightarrow \text{Customer} : \{ \text{quote}(\text{nat}).\text{Customer} \rightarrow \text{Agency} : \{ \\ &\quad \text{accept}(\text{bool}).\text{Customer} \rightarrow \text{Agency} : \text{address}(\text{nat}). \\ &\quad \quad \quad \text{Agency} \rightarrow \text{Customer} : \text{date}(\text{nat}).\text{end}, \\ &\quad \text{retry}(\text{bool}).t, \\ &\quad \text{reject}(\text{bool}).\text{end} \} \} \end{aligned}$$

Observe the recursive behaviour over t . Now, let us project G onto its participant Agency (there is no need to merge types here, namely we can use Definition 4):

$$\begin{aligned} G \upharpoonright \text{Agency} &= \mu t. (\text{Customer}? \text{Hawaii}(\text{bool}).G_1 \upharpoonright \text{Agency} \\ &\quad \oplus \text{Customer}? \text{France}(\text{bool}).G_1 \upharpoonright \text{Agency}) \\ &\text{and} \\ G_1 \upharpoonright \text{Agency} &= \text{Customer}! \text{quote}(\text{nat}).(\\ &\quad (\text{Customer}? \text{accept}(\text{bool}).\text{Customer}? \text{address}(\text{nat}). \\ &\quad \quad \quad \text{Customer}! \text{date}(\text{nat}).\text{end}) \\ &\quad \& (\text{Customer}? \text{retry}(\text{bool}).t) \\ &\quad \& (\text{Customer}? \text{reject}(\text{bool}).\text{end})) \end{aligned}$$

Also this projection presents a recursive behaviour as expected. The reader will now be able to derive:

$$\vdash \text{Customer} \triangleleft P_{\text{Customer}} \mid \text{Agency} \triangleleft P_{\text{Agency}} : G.$$

Exercise 2. This exercise is intended to guide the reader to obtain the final result (using the same notation as in Examples 1 and 8):

$$\vdash \text{Customer} \triangleleft P_{\text{Customer}} \mid \text{Agency} \triangleleft P_{\text{Agency}} : G$$

Let us proceed step by step following Example 7.

1. Given the global type G and its projection $G \upharpoonright \text{Agency}$, derive $G \upharpoonright \text{Customer}$.
2. Then derive $\vdash P_{\text{Customer}} : G \upharpoonright \text{Customer}$ and $\vdash P_{\text{Agency}} : G \upharpoonright \text{Agency}$.

Hint 1. The reader might want to use the identification of $\mu t.G$ and $G\{\mu t.G/t\}$ (this holds both for global and session types), namely the possibility to unfold any recursive construct for types.

Hint 2. In order to prove points 2. and 3. the reader might want to proceed in two steps: first, finding appropriate session types $\mathsf{T}_{\text{Customer}}$ and $\mathsf{T}_{\text{Agency}}$ for processes P_{Customer} and P_{Agency} respectively, and second, proving $\mathsf{T}_{\text{Customer}} \leq \mathsf{G} \upharpoonright \text{Customer}$ and $\mathsf{T}_{\text{Agency}} \leq \mathsf{G} \upharpoonright \text{Agency}$.

Exercise 3. Prove that: $\vdash \text{Customer} \triangleleft P'_{\text{Customer}} \mid \text{Agency} \triangleleft P'_{\text{Agency}} : \mathsf{G}$.

The proposed type system for multiparty sessions enjoys two fundamental properties: typed sessions only reduce to typed sessions (subject reduction), and typed sessions never get stuck. The remaining of this section is devoted to the proof of these properties.

In order to state subject reduction, we need to formalise how global types are modified when multiparty sessions reduce and evolve.

Definition 7 (Global types consumption and reduction). *The consumption of the communication $\mathsf{p} \xrightarrow{\ell} \mathsf{q}$ for the global type G (notation $\mathsf{G} \setminus \mathsf{p} \xrightarrow{\ell} \mathsf{q}$) is the global type coinductively defined as follows:*

$$\begin{aligned} (\mathsf{p} \rightarrow \mathsf{q} : \{\ell_i(S_i).G_i\}_{i \in I}) \setminus \mathsf{p} \xrightarrow{\ell} \mathsf{q} &= G_k && \text{if } \exists k \in I : \ell = \ell_k \\ (\mathsf{r} \rightarrow \mathsf{s} : \{\ell_i(S_i).G_i\}_{i \in I}) \setminus \mathsf{p} \xrightarrow{\ell} \mathsf{q} &= \mathsf{r} \rightarrow \mathsf{s} : \{\ell_i(S_i).G_i \setminus \mathsf{p} \xrightarrow{\ell} \mathsf{q}\}_{i \in I} \\ &&& \text{if } \begin{cases} \{\mathsf{r}, \mathsf{s}\} \cap \{\mathsf{p}, \mathsf{q}\} = \emptyset \\ \forall i \in I : \{\mathsf{p}, \mathsf{q}\} \subseteq G_i \end{cases} \end{aligned}$$

The reduction of global types is the smallest pre-order relation closed under the rule: $\mathsf{G} \Longrightarrow \mathsf{G} \setminus \mathsf{p} \xrightarrow{\ell} \mathsf{q}$

Example 9. We show that a projection of a global type before the consumption might require to support more external choices than the projection after the consumption. Take G , its subterm G_1 , from Example 5, and their types denoted as G and G_1 , respectively. Also take the projection:

$$\mathsf{G} \upharpoonright \mathsf{r} = \mathsf{q}?\ell_3(\text{int})\&\mathsf{q}?\ell_4(\text{int})\&\mathsf{q}?\ell_5(\text{nat})$$

and recall the explanation on how $\mathsf{G} \upharpoonright \mathsf{r}$ above merges all the possible inputs that r might receive from q , depending on whether p first sends ℓ_1 or ℓ_2 to q .

We have:

$$\begin{aligned} \mathsf{G} \setminus \mathsf{p} \xrightarrow{\ell_1} \mathsf{q} &= G_1 = \mathsf{q} \rightarrow \mathsf{r} : \{\ell_3(\text{int}), \ell_5(\text{nat})\} \\ (\mathsf{G} \setminus \mathsf{p} \xrightarrow{\ell_1} \mathsf{q}) \upharpoonright \mathsf{r} &= G_1 \upharpoonright \mathsf{r} = \mathsf{q}?\ell_3(\text{int})\&\mathsf{q}?\ell_5(\text{nat}) \end{aligned}$$

and we obtain $\mathsf{G} \upharpoonright \mathsf{r} \leq (\mathsf{G} \setminus \mathsf{p} \xrightarrow{\ell_1} \mathsf{q}) \upharpoonright \mathsf{r}$. The reason is that, after the transition from G to G_1 , there is no possibility for q to send ℓ_4 to r , hence r does not need to support such a message in its projection.

Note that a process that plays the role of r in G , and is therefore typed by $\mathsf{G} \upharpoonright \mathsf{r}$, has to support the input of ℓ_4 from q , by rule [T-INPUT-CHOICE] in Table 5. After the transition from G to G_1 , the same process is also typed by $G_1 \upharpoonright \mathsf{r}$, by rule [T-SUB] — but will never receive a message ℓ_4 from q .

We can now prove subject reduction.

Theorem 1 (Subject Reduction). *Let $\vdash \mathcal{M} : \mathbf{G}$. For all \mathcal{M}' , if $\mathcal{M} \longrightarrow \mathcal{M}'$, then $\vdash \mathcal{M}' : \mathbf{G}'$ for some \mathbf{G}' such that $\mathbf{G} \Longrightarrow \mathbf{G}'$.*

Theorem 2 (Progress). *If $\vdash \mathcal{M} : \mathbf{G}$, then either $\mathcal{M} \equiv \mathbf{p}\triangleleft\mathbf{0}$ or there is \mathcal{M}' such that $\mathcal{M} \longrightarrow \mathcal{M}'$.*

As a consequence of subject reduction and progress, we get the safety property stating that a typed multiparty session will never get stuck.

Theorem 3 (Type Safety). *If $\vdash \mathcal{M} : \mathbf{G}$, then it does not hold $\text{stuck}(\mathcal{M})$.*

Proof. Direct consequence of Theorems 1, 2, and Definition 1.

Finally the reader wishes to learn more about the full asynchronous multiparty session types (which includes channel passing, asynchrony (FIFO queues), shared names and parameterised recursion) can proceed to [1]. The article [16] also discusses type soundness of various MPST calculi.

Acknowledgements. We would like to thank Fangyi Zhou for their careful proof-reading and detailed comments on the paper.

References

1. Coppo, M., Dezani-Ciancaglini, M., Padovani, L., Yoshida, N.: A gentle introduction to multiparty asynchronous session types. In: Bernardo, M., Johnsen, E.B. (eds.) SFM 2015. LNCS, vol. 9104, pp. 146–178. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-18941-3_4
2. Denielou, P., Yoshida, N., Bejleri, A., Hu, R.: Parameterised multiparty session types. *Log. Methods Comput. Sci.* **8**(4) (2012). [https://doi.org/10.2168/LMCS-8\(4:6\)2012](https://doi.org/10.2168/LMCS-8(4:6)2012)
3. Gay, S., Hole, M.: Subtyping for session types in the pi calculus. *Acta Informatica* **42**(2/3), 191–225 (2005). <https://doi.org/10.1007/s00236-005-0177-z>
4. Gay, S., Ravera, A. (eds.): *Behavioural Types: From Theory to Tools*. River Publishers (2017)
5. Gay, S.J.: Subtyping supports safe session substitution. In: Lindley, S., McBride, C., Trinder, P., Sannella, D. (eds.) *A List of Successes That Can Change the World*. LNCS, vol. 9600, pp. 95–108. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-30936-1_5
6. Gay, S., Hole, M.: Types and subtypes for client-server interactions. In: Swierstra, S.D. (ed.) *ESOP 1999*. LNCS, vol. 1576, pp. 74–90. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-49099-X_6
7. Ghilezan, S., Jaksic, S., Pantovic, J., Scalas, A., Yoshida, N.: Precise subtyping for synchronous multiparty sessions. *J. Log. Algebr. Meth. Program.* **104**, 127–173 (2019). <https://doi.org/10.1016/j.jlamp.2018.12.002>
8. Honda, K.: Types for dyadic interaction. In: Best, E. (ed.) *CONCUR 1993*. LNCS, vol. 715, pp. 509–523. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-57208-2_35

9. Honda, K., Mukhamedov, A., Brown, G., Chen, T.-C., Yoshida, N.: Scribbling interactions with a formal foundation. In: Natarajan, R., Ojo, A. (eds.) ICDCIT 2011. LNCS, vol. 6536, pp. 55–75. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19056-8_4
10. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0053567>
11. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: POPL, pp. 273–284. ACM Press (2008). <https://doi.org/10.1145/1328438.1328472>
12. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. *J. ACM* **63**, 1–67 (2016)
13. Hüttel, H., et al.: Foundations of session types and behavioural contracts. *ACM Comput. Surv.* **49**(1) (2016). <https://doi.org/10.1145/2873052>
14. Kouzapas, D., Yoshida, N.: Globally governed session semantics. In: D’Argenio, P.R., Melgratti, H. (eds.) CONCUR 2013. LNCS, vol. 8052, pp. 395–409. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40184-8_28
15. Pierce, B.C.: *Types and Programming Languages*. MIT Press (2002)
16. Scalas, A., Yoshida, N.: Less is more: multiparty session types revisited. *ACM Program. Lang. POPL* (2019). <https://doi.org/10.1145/3290343>
17. Scribble home page. <http://www.scribble.org>
18. Takeuchi, K., Honda, K., Kubo, M.: An interaction-based language and its typing system. In: Halatsis, C., Maritsas, D., Philokyprou, G., Theodoridis, S. (eds.) PARLE 1994. LNCS, vol. 817, pp. 398–413. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-58184-7_118
19. Yoshida, N., Deniérou, P.-M., Bejleri, A., Hu, R.: Parameterised multiparty session types. In: Ong, L. (ed.) FoSSaCS 2010. LNCS, vol. 6014, pp. 128–145. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12032-9_10
20. Yoshida, N., Hu, R., Neykova, R., Ng, N.: The scribble protocol language. In: Abadi, M., Lluch Lafuente, A. (eds.) TGC 2013. LNCS, vol. 8358, pp. 22–41. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-05119-2_3