A very gentle introduction for synchronous multiparty sessions

Nobuko Yoshida^a

^aImperial College London, UK

Keywords: Concurrency, Process Calculi, Multiparty Session Types

1 1. A Gentle Introduction to Multiparty Session Types

The aim of this section is to give a gentle introduction of multiparty session types for readers who are interested in programming but who are not familiar with session types nor process calculi.

Session types are introduced in a series of papers during the 1990s Honda (1993); Takeuchi et al. (1994);

⁵ Honda et al. (1998) in the context of pure concurrent processes and programming. Session types have since
⁶ been studied in many contexts over the last decade—see the surveys of the field Hüttel et al. (2016); Gay

 $_{7}$ and Ravera (2017).

We review multiparty session types, a methodology to enable compositional reasoning about communi cation.

As a simple example, consider a scenario in which a cart and arm assembly has to fetch objects. We associate a process with each physical component; thus, we model the scenario using a *cart* (Cart) and an *arm* (Arm) attached to the cart. The task involves synchronisation between the cart and the arm. Synchronisation is obtained through the exchange of messages.

¹⁴ Specifically, the protocol works as follows.

The cart sends the arm a fold command *fold*. On receiving the command, the arm folds itself. When
 the arm is completely folded, it sends back a message *ok* to the cart. On receipt of this message, the
 cart moves.

2. When the cart reaches the object, it stops and sends a *grab* message to the arm to grab the object.
 While the cart waits, the arm executes the grabbing operation, followed by a folding operation. Then
 the arm sends a message *ok* to the cart. This sequence may need to be repeated.

3. When all tasks are finished, the cart sends a message *done* to the arm, and the protocol 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 get 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:

$$\mu \mathbf{t}.\mathsf{Cart} \to \mathsf{Arm}: \{ fold.\mathsf{Arm} \to \mathsf{Cart}: ok.\mathsf{Cart} \to \mathsf{Arm}: grab.\mathsf{Arm} \to \mathsf{Cart}: ok.\mathbf{t}, done.\mathtt{end} \}$$
(1)

The type describes the global pattern of communication between Cart and Arm using message exchanges, sequencing, choice, and repetition. The basic pattern Cart \rightarrow Arm: *m* indicates a message *m* sent from the Cart to the Arm. The communication starts with the cart sending either a *fold* or a *done* command to the arm. In case of *done*, the protocol ends (type end); otherwise, the communication continues with the sequence *ok*. *grab. ok* followed by a repetition of the entire pattern. The operator "." denotes sequencing,

and the type $\mu \mathbf{t}.T$ denotes recursion of T.

The global type states what are the valid message sequences allowed in the system. When we implement Cart and Arm separately, we would like to check that their composition conforms to the global type. We can perform this check compositionally as follows.

Since there are only two participants, projecting to each participant is simple. From the perspective of the Cart, the communication can be described by the type:

$$\mu \mathbf{t}. ((!fold.?ok.!grab.?ok.\mathbf{t}) \oplus (!done. \mathsf{end}))$$
(2)

⁶ where !m denotes a message m sent (to the Arm) and ?m denotes a message m received from the Arm. and

 $_7 \oplus$ denotes an (internal) choice. Thus, the type states that Cart repeats actions ! fold. ? ok. ! grab. ? ok until

⁸ at some point it sends *done* and exits.

⁹ Dually, from the viewpoint of the Arm, the same global session is described by the dual type

$$\mu \mathbf{t}. \left(\left(?fold. !ok. ?grab. !ok. \mathbf{t} \right) \& \left(?done. \mathsf{end} \right) \right)$$

$$(3)$$

¹⁰ in which & means that a choice is offered externally.

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

The global type seems overkill if there are only two participants; indeed, the global type is uniquely determined given the local type (2) or its dual (3). 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.

¹⁷ For example, consider a simple ring protocol, where the Arm process above is divided into two parts,

¹⁸ Lower and Upper. Now, Cart sends a message *fold* to the lower arm Lower, which forwards the message to ¹⁹ Upper. After receiving the message, Upper sends an acknowledgement *ok* to Cart. We start by specifying

²⁰ the global type as:

2

$$Cart \rightarrow Lower: fold.Lower \rightarrow Upper: fold.Upper \rightarrow Cart: ok.end$$
 (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.

²³ The global type in (4) is *projected* into the three endpoint session types:

	Cart 's endpoint type:	Lower! <i>fold</i> .Upper? <i>ok</i> .end
4	Lower's endpoint type:	${\sf Cart}? {\it fold}. {\sf Upper}! {\it fold}. {\tt end}$
	Upper's endpoint type:	Lower?fold.Cart!ok.end

where Lower! *fold* means "send to Lower a *fold* message," and Upper? *ok* means "receive from Upper 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 of a global type, we only used three separate binary session types to describe the mes-28 sage exchanges between Cart and Lower, between Lower and Upper, and between Upper and Cart, respec-29 tively, without using a global type, then we lose essential sequencing information in this interaction sce-30 nario. Consequently, we can no longer guarantee deadlock-freedom among these three parties. Since the 31 three separate binary sessions can be interleaved freely, an implementation of the Cart that conforms to 32 Upper? ok. Lower! fold. end becomes typable. This causes the situation that each of the three parties blocks 33 indefinitely while waiting for a message to be delivered. Thus, we shall use the power of multiparty session 34 types to ensure correct communication patterns. 35

³⁶ 2. Synchronous Multiparty Session Calculus

This section introduces the syntax and semantics of a synchronous multiparty session calculus. Since our focus is on subtyping, we simplify the calculus in (Kouzapas and Yoshida, 2013) eliminating both shared channels for session initiations and session channels for communications inside sessions — i.e, our calculus

⁴⁰ is akin to value-passing CCS (Milner, 1989, Chapter 2.8).

$$\begin{split} & \mathsf{succ}(\mathsf{n}) \downarrow (\mathsf{n}+1) \quad \mathsf{neg}(\mathsf{i}) \downarrow (-\mathsf{i}) \quad \neg \mathsf{true} \downarrow \mathsf{false} \quad \neg \mathsf{false} \downarrow \mathsf{true} \quad \mathsf{v} \downarrow \mathsf{v} \\ & (\mathsf{i}_1 > \mathsf{i}_2) \downarrow \begin{cases} \mathsf{true} \quad \mathrm{if} \; \mathsf{i}_1 > \mathsf{i}_2, & \\ \mathsf{false} \quad \mathrm{otherwise} \end{cases} \quad \frac{\mathsf{e}_1 \downarrow \mathsf{v}}{\mathsf{e}_1 \oplus \mathsf{e}_2 \downarrow \mathsf{v}} \quad \frac{\mathsf{e}_2 \downarrow \mathsf{v}}{\mathsf{e}_1 \oplus \mathsf{e}_2 \downarrow \mathsf{v}} \quad \frac{\mathsf{e} \downarrow \mathsf{v} \quad \mathcal{E}(\mathsf{v}) \downarrow \mathsf{v}'}{\mathcal{E}(\mathsf{e}) \downarrow \mathsf{v}'} \end{split}$$

Table 1: Expression evaluation.

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

Table 2: Structural congruence.

Notation 2.1 (Base sets). We use the following base sets: values, ranged over by v, v', \ldots ; expressions, ranged over by e, e', \ldots ; expression variables, ranged over by $x, y, z \ldots$; labels, ranged over by ℓ, ℓ', \ldots ; session participants, ranged over by p, q, \ldots ; process variables, ranged over by X, Y, \ldots ; processes, ranged over by P, Q, \ldots ; and multiparty sessions, ranged over by $\mathcal{M}, \mathcal{M}', \ldots$.

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 $\operatorname{succ}, \operatorname{neg}, \neg, \oplus$, or the relation > . An *evaluation context* \mathcal{E} is an expression with exactly one hole. The operator \oplus models non-determinism: $\mathbf{e}_1 \oplus \mathbf{e}_2$ is an expression that might yield either \mathbf{e}_1 or \mathbf{e}_2 .

⁹ The processes of the synchronous multiparty session calculus are defined by:

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

The output process $p!\ell(e).Q$ 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 **0** from the tail of processes.

¹⁸ We define a *multiparty session* as a parallel composition of pairs (denoted by $p \triangleleft P$) of participants and ¹⁹ processes:

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

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 multyparty session is *well formed* if all its participants are different. We consider only well-formed multiparty sessions.

²³ **Operational semantics.** The value v of expression e (notation $e \downarrow v$) is computed as expected, see Table 1. ²⁴ The successor operation succ is defined only on natural numbers, the negation neg is defined on integers, ²⁵ 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 ²⁶ the value of e_2 .
$$\begin{split} & [\operatorname{R-COMM}] & \underbrace{j \in I \quad \mathbf{e} \downarrow \mathbf{v}}_{\mathbf{p} \lhd \sum_{i \in I} \mathbf{q} ? \ell_i(x) . P_i \ | \ \mathbf{q} \lhd \mathbf{p} ! \ell_j(\mathbf{e}) . Q \ | \ \mathcal{M} \longrightarrow \mathbf{p} \lhd P_j \{ \mathbf{v} / x \} \ | \ \mathbf{q} \lhd Q \ | \ \mathcal{M}} \\ & [\operatorname{T-CONDITIONAL}] & [\operatorname{F-CONDITIONAL}] \\ & \underbrace{\mathbf{e} \downarrow \operatorname{true}}_{\mathbf{p} \lhd \operatorname{if} \mathbf{e} \operatorname{then} P \operatorname{else} Q \ | \ \mathcal{M} \longrightarrow \mathbf{p} \lhd P \ | \ \mathcal{M}} & \underbrace{\mathbf{p} \lhd \operatorname{if} \mathbf{e} \operatorname{then} P \operatorname{else} Q \ | \ \mathcal{M} \longrightarrow \mathbf{p} \lhd Q \ | \ \mathcal{M}}_{\mathbf{p} \lhd \operatorname{if} \mathbf{e} \operatorname{then} P \operatorname{else} Q \ | \ \mathcal{M} \longrightarrow \mathbf{p} \lhd Q \ | \ \mathcal{M}} \\ & \underbrace{\frac{[\operatorname{R-STRUCT}]}{\mathcal{M}_1' = \mathcal{M}_1 \ \mathcal{M}_1 \longrightarrow \mathcal{M}_2 \ \mathcal{M}_2 \equiv \mathcal{M}_2'}_{\mathcal{M}_1' \longrightarrow \mathcal{M}_2'}} \end{split}$$

Table 3: Reduction rules.

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

⁷ We adopt some standard conventions regarding the syntax of processes and sessions. Namely, we will use ⁸ $\prod_{i \in I} \mathbf{p}_i \triangleleft P_i$ as short for $\mathbf{p}_1 \triangleleft P_1 \mid \ldots \mid \mathbf{p}_n \triangleleft P_n$, where $I = \{1, \ldots, n\}$. We will sometimes use infix notation for ⁹ 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 2.2. We now show the operational semantics in action. Consider the following multiparty session
 with three participants, Alice, Bob and Carol:

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

12 where

¹³ This mulitiparty session reduces to

 $\texttt{Alice} \triangleleft \mathbf{0} ~|~ \texttt{Bob} \triangleleft \mathbf{0} ~|~ \texttt{Carol} \triangleleft \mathbf{0}$

after three communications occur. First, Alice sends to Bob natural number 50 with the label ℓ_1 . Bob is

¹⁵ able to receive values with labels ℓ_1 and ℓ_4 . Next, the only possible communication is between Bob and Carol.

¹⁶ So, Carol receives natural number 100 from Bob. The value 100 is substituted in the continuation process. ¹⁷ Finally, since $succ(100) \downarrow 101$, Carol sends 101 to Alice. We can then reduce the session to, for example, ¹⁸ Alice $\downarrow 0$ but not further

¹⁸ Alice $\triangleleft \mathbf{0}$, but not further.

From the end of Example 2.2, 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 Section 3, 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 2.3. A multiparty session \mathcal{M} is stuck if $\mathcal{M} \neq \mathbf{p} \triangleleft \mathbf{0}$ and there is no multiparty session \mathcal{M}' such that $\mathcal{M} \longrightarrow \mathcal{M}'$. A multiparty session \mathcal{M} gets stuck, notation $\mathtt{stuck}(\mathcal{M})$, if it reduces to a stuck multiparty session.

E.g., the multiparty session \mathcal{M} in Example 2.2 is not stuck, and it does not get stuck. A similar multiparty session, where instead of P_{Alice} we take $P'_{\text{Alice}} = \text{Bob}!\ell_1(50).\text{Carol}?\ell_5(x).0$, gets stuck because of label mismatch.

4 3. Type System

This section introduces a type system for the calculus presented in Section 2. The formulation is based on Kouzapas and Yoshida (2013, 2015), with adaptations to account for our simplified calculus. We formalise types and projections (Section 3.1), the subtyping relation (Section 3.2), and the typing rules and their properties (Section 3.3).

9 3.1. 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.

¹² **Definition 3.1** (Sorts and global types). Sorts, ranged over by S, are defined as:

S ::= nat | int | bool

¹³ Global types, ranged over by \mathbb{G} , are terms generated by the following grammar:

$$\mathsf{G} \quad ::= \quad \mathsf{end} \quad | \quad \mu \mathbf{t}.\mathsf{G} \quad | \quad \mathbf{t} \quad | \quad \mathsf{p} \to \mathsf{q} : \{\ell_i(S_i).G_i\}_{i \in I}$$

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 3.1, the type $\mathbf{p} \to \mathbf{q} : \{\ell_i(S_i).\mathbf{G}_i\}_{i \in I}$ 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 \mathbf{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 $\mu \mathbf{t}.\mathbf{G}$, where recursion variable \mathbf{t} is bound and guarded in $\mathbf{G} - \text{e.g.}, \mu \mathbf{t}.\mathbf{p} \to \mathbf{q} : \ell(\text{nat}).\mathbf{t}$ is a valid global type, whereas $\mu \mathbf{t}.\mathbf{t}$ is not. We take the equi-recursive viewpoint, i.e. We identify $\mu \mathbf{t}.\mathbf{G}$ and $\mathbf{G}\{\mu \mathbf{t}.\mathbf{G}/\mathbf{t}\}$. We define the set of participants of a global type \mathbf{G} , by structural induction on \mathbf{G} , as follows:

$$\mathtt{pt}\{\mu\mathtt{t}.\mathsf{G}\}=\mathtt{pt}\{\mathsf{G}\}\qquad\mathtt{pt}\{\mathtt{end}\}=\mathtt{pt}\{\mathtt{t}\}=\emptyset\qquad\mathtt{pt}\{\mathtt{p}\to\mathtt{q}:\{\ell_i(S_i).G_i\}_{i\in I}\}=\{\mathtt{p},\mathtt{q}\}\cup\mathtt{pt}\{\mathsf{G}_i\}\;(i\in I)$$

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

- A local session type describes the behaviour of a single participant in a multiparty session.
- ²⁶ Definition 3.2 (Local Session Types). The grammar of session types, ranged over by T, is:

 $\mathsf{T} ::= \mathsf{end} | \bigwedge_{i \in I} \mathsf{p}?\ell_i(S_i).\mathsf{T}_i | \bigvee_{i \in I} \mathsf{q}!\ell_i(S_i).\mathsf{T}_i | \mu \mathsf{t}.\mathsf{T} | \mathsf{t}$

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, $p?\ell(int).end \land p?\ell(nat).end$ is not a type.

The session type end says that no further communication is possible and the protocol is completed. The external choice or branching type $\bigwedge_{i \in I} \mathsf{p}?\ell_i(S_i).\mathsf{T}_i$ requires to wait to receive a value of sort S_i (for some $i \in I$) from the participant 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 $\bigvee_{i \in I} \mathsf{q}!\ell_i(S_i).\mathsf{T}_i$ says that the participant implementing the type must choose a labelled message to send to q ; if the participant chooses the message ℓ_i , for some $i \in I$, it must include in the message to q a payload value of sort S_i , and continue as ¹ prescribed by T_i . Recursion is modelled by the session type $\mu t.T$. We adopt the following conventions: we

² do not write branch/selection symbols in case of a singleton choice, we do not write unnecessary parentheses,

 $_3$ and we often omit trailing ends.

10

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

$$\mathsf{pt}\{\bigwedge_{i\in I}\mathsf{p}?\ell_i(S_i).\mathsf{T}_i\}=\mathsf{pt}\{\bigvee_{i\in I}\mathsf{p}!\ell_i(S_i).\mathsf{T}_i\}=\{\mathsf{p}\}\cup\bigcup_{i\in I}\mathsf{pt}\{\mathsf{T}_i\}\quad\mathsf{pt}\{\mu\mathsf{t}.\mathsf{T}\}=\mathsf{pt}\{\mathsf{T}\}\quad\mathsf{pt}\{\mathsf{t}\}=\mathsf{pt}\{\mathsf{end}\}=\emptyset.$$

In Definition 3.3 below, we define the *global type projection* as a coinductive relation $G \upharpoonright_r T$ between global and session types. Our definition extends the one originally proposed by Honda et al. (2008, 2016),

⁷ along the lines of Yoshida et al. (2010) and Deniélou et al. (2012): i.e., it uses a merging operator \Box .

Definition 3.3. The projection of a global type onto a participant r is the largest relation \uparrow_r between global ypes and session types such that, whenever $G \uparrow_r T$:

•
$$r \notin pt\{G\}$$
 implies $T = end$; [PROJ-END]

•
$$G = p \rightarrow r : \{\ell_i(S_i).G_i\}_{i \in I} \text{ implies } T = \bigwedge_{i \in I} p?\ell_i(S_i).T_i, \text{ and } G_i \upharpoonright_r T_i, \forall i \in I;$$
 [PROJ-IN]

•
$$\mathbf{G} = \mathbf{r} \to \mathbf{q} : \{\ell_i(S_i).\mathbf{G}_i\}_{i \in I} \text{ implies } \mathbf{T} = \bigvee_{i \in I} \mathbf{q}! \ell_i(S_i).\mathbf{T}_i, \text{ and } \mathbf{G}_i \upharpoonright_{\mathbf{r}} \mathbf{T}_i, \forall i \in I;$$
 [PROJ-OUT]

•
$$G = p \rightarrow q : \{\ell_i(S_i), G_i\}_{i \in I} \text{ and } r \notin \{p, q\} \text{ implies that there are } \mathsf{T}_i, i \in I \text{ such that} \quad [PROJ-CONT]$$

 $\mathsf{T} = \prod_{i \in I} \mathsf{T}_i, \text{ and } \mathsf{G}_i \upharpoonright_r \mathsf{T}_i, \text{ for every } i \in I.$

¹¹ Above, \prod is the merging operator, that is a partial operation over session types defined as:

$$\mathsf{T}_{1} \sqcap \mathsf{T}_{2} = \begin{cases} \mathsf{T}_{1} & \text{if } \mathsf{T}_{1} = \mathsf{T}_{2} & [\text{MRG-ID}] \\ \\ \mathsf{T}_{3} & \text{if } \exists I, J : \begin{cases} \mathsf{T}_{1} = \bigwedge_{i \in I} \mathsf{p}'?\ell_{i}(S_{i}).\mathsf{T}_{i} & \text{and} \\ \mathsf{T}_{2} = \bigwedge_{j \in J} \mathsf{p}'?\ell_{j}(S_{j}).\mathsf{T}_{j} & \text{and} \\ \\ \mathsf{T}_{3} = \bigwedge_{k \in I \cup J} \mathsf{p}'?\ell_{k}(S_{k}).\mathsf{T}_{k} \end{cases} & [\text{MRG-BRA}] \\ \end{cases}$$

Proposition 3.4. The merging operation is associative, i.e.: $T \sqcap (T' \sqcap T'') = (T \sqcap T') \sqcap T''$.

By Definition 3.3, merging a type with itself results in itself (rule [MRG-ID]). Moreover, Definition 3.3 allows to combine different external choices (rule [MRG-BRA]) *if and only if* common labels have identical sorts and identical continuations, as formalised in Prop. 3.5 below and illustrated in Examples 3.6, 3.8 and 3.9.

Proposition 3.5. For two types $\mathsf{T}' = \bigwedge_{i \in I} \mathsf{p}'?\ell_i(S_i).\mathsf{T}_i$ and $\mathsf{T}'' = \bigwedge_{j \in J} \mathsf{p}''?\ell_j(S_j).\mathsf{T}_j$, 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 3.6. We now give some small examples that illustrate the definition of the merging operator (here, $i \neq j$ implies $\ell_i \neq \ell_j$):

$q!\ell(\texttt{nat}) \sqcap q!\ell(\texttt{nat})$	$= q! \ell(nat)$
$p!\ell(\mathtt{nat}) \sqcap q!\ell(\mathtt{nat})$	undefined: outputs to different participants
$q!\ell_3(\texttt{nat}) \sqcap q!\ell_4(\texttt{nat})$	undefined: outputs with different labels
$ \begin{array}{l} \left(q?\ell_3(\texttt{int}) \land q?\ell_5(\texttt{nat})\right) \sqcap \left(q?\ell_4(\texttt{int}) \land q?\ell_5(\texttt{nat})\right) \\ \qquad $	$= q?\ell_3(int) \wedge q?\ell_4(int) \wedge q?\ell_5(nat)$ undefined: same prefixes, but different continuations undefined: the payload sorts do not match

²⁰ **Proposition 3.7.** The projection relation $|_r$ is a partial function.

²¹ We now describe the clauses of Definition 3.3:

- clause [PROJ-END] states that when a global type G is projected onto a participant r who does not appear in G (e.g., because G = end), then the result is the inactive session type end;
- clause [PROJ-IN] (resp. [PROJ-OUT]) states that a global type G starting with a communication from p to
 r (resp. from r to q) projects onto an external (resp. internal) choice T, provided that the continuations
 of T are also projections of the corresponding global type continuations.
- clause [PROJ-CONT] 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 \Box .

As a result, clause [PROJ-CONT] of Definition 3.3 allows participant r to receive different messages (from a same participant p') in different branches of a global type, as shown in Example 3.8 below.

Example 3.8. We demonstrate interesting points of Definition 3.3. First, we show some projections of global types. Consider the global type:

$$\mathsf{G} \ = \ \mathsf{p} \rightarrow \mathsf{q} : \{\ell_1(\texttt{nat}).\mathsf{G}_1, \ell_2(\texttt{bool}).\mathsf{G}_2\} \qquad \textit{where } \left\{ \begin{array}{l} \mathsf{G}_1 \ = \ \mathsf{q} \rightarrow \mathsf{r} : \{\ell_3(\texttt{int}), \ell_5(\texttt{nat})\} \\ \mathsf{G}_2 \ = \ \mathsf{q} \rightarrow \mathsf{r} : \{\ell_4(\texttt{int}), \ell_5(\texttt{nat})\} \\ \mathsf{r} \neq \mathsf{p} \end{array} \right.$$

13 We have:

$$\mathsf{G}\restriction \mathsf{p} \ = \ \mathsf{q}!\ell_1(\texttt{nat}).(\mathsf{G}_1\restriction \mathsf{p}) \ \lor \ \mathsf{q}!\ell_2(\texttt{bool}).(\mathsf{G}_2\restriction \mathsf{p}) \ = \ \mathsf{q}!\ell_1(\texttt{nat}).\texttt{end} \ \lor \ \mathsf{q}!\ell_2(\texttt{bool}).\texttt{end}$$

$$\begin{split} \mathsf{G} &\models \mathsf{q} := \mathsf{p}?\ell_1(\mathtt{nat}).(\mathsf{G}_1 \models \mathsf{q}) \land \mathsf{p}?\ell_2(\mathtt{bool}).(\mathsf{G}_2 \models \mathsf{q}) \\ &= \mathsf{p}?\ell_1(\mathtt{nat}).(\mathsf{r}!\ell_3(\mathtt{int}) \lor \mathsf{r}!\ell_5(\mathtt{nat})) \land \mathsf{p}?\ell_2(\mathtt{bool}).(\mathsf{r}!\ell_4(\mathtt{int}) \lor \mathsf{r}!\ell_5(\mathtt{nat})) \\ \mathsf{G} \models \mathsf{r} := \mathsf{G}_1 \models \mathsf{r} \sqcap \mathsf{G}_2 \models \mathsf{r} := (\mathsf{q}?\ell_3(\mathtt{int}) \land \mathsf{q}?\ell_5(\mathtt{nat})) \sqcap (\mathsf{q}?\ell_4(\mathtt{int}) \land \mathsf{q}?\ell_5(\mathtt{nat})) \\ &= \mathsf{q}?\ell_3(\mathtt{int}) \land \mathsf{q}?\ell_4(\mathtt{int}) \land \mathsf{q}?\ell_5(\mathtt{nat}) \end{split}$$

¹⁴ Note that in G, q could output different messages towards r, depending on whether p sends ℓ_1 or ℓ_2 to q; ¹⁵ therefore, in G|r, the possible inputs of r in G₁ and G₂ are merged into a larger external choice that supports ¹⁶ all possible outputs of q.

Importantly, by Definition 3.3, 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 \square in clause [PROJ-CONT] to be undefined, as shown in Example 3.9 below.

Example 3.9. We show two global types that cannot be projected according to the Definition 3.3. Consider the global type $G = p \rightarrow q : \{\ell_1(nat).G_1, \ell_2(bool).G_2\}, with G_1 = r \rightarrow q : \ell_3(nat) \text{ and } G_2 = r \rightarrow q : \ell_4(nat).$ Then,

Intuitively, when $\ell_3 \neq \ell_4$, $G \upharpoonright r$ 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 $\ell_3 = \ell_4$, then by Definition 3.3 we have $G \upharpoonright r = q!\ell_3(nat) \sqcap q!\ell_3(nat) = q!\ell_3(nat).$

Now, consider the global type $G' = p \rightarrow q : \{\ell_1(nat).G'_1, \ell_2(bool).G'_2\}$, with $G'_1 = q \rightarrow r : \ell_3(nat)$ and $G'_2 = q \rightarrow r : \ell_3(nat).q \rightarrow r : \ell_3(nat).$ Then,

$$\begin{array}{lll} \mathsf{G}'\!\!\upharpoonright\!\!\mathsf{p} &=& \mathsf{q}!\ell_1(\texttt{nat})\vee\mathsf{q}!\ell_2(\texttt{bool})\\ \mathsf{G}'\!\!\upharpoonright\!\!\mathsf{q} &=& \mathsf{p}?\ell_1(\texttt{nat}).\mathsf{r}!\ell_3(\texttt{nat})\wedge\mathsf{p}?\ell_2(\texttt{bool}).\mathsf{r}!\ell_3(\texttt{nat}).\mathsf{r}!\ell_3(\texttt{nat})\\ \mathsf{G}'\!\!\upharpoonright\!\!\mathsf{r} &=& \mathsf{q}?\ell_3(\texttt{nat})\sqcap\mathsf{q}?\ell_3(\texttt{nat}).\mathsf{q}?\ell_3(\texttt{nat}) \ \textit{(undefined)} \end{array}$$

¹ 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 ² receive either one or two messages ℓ_3 from q; however, as in the previous example, r is not aware of the

³ interactions between **p** and **q**, and thus, **G** provides an invalid specification for **r**. This example could be fixed,

4 e.g., by replacing ℓ_3 with $\ell' \neq \ell_3$ in G'_2 , or by letting $G'_1 = G'_2$: both fixes would make $G' \upharpoonright r$ defined, similarly

5 to Example 3.8.

6 3.2. Subtyping

The subtyping relation \leq is used to augment the flexibility of the type system (introduced in Section 3.3):

 $_{\rm s}~$ by determining when a type T is "smaller" than $T^\prime,$ it allows to use a process typed by the former whenever

⁹ a process typed by the latter is required.

Definition 3.10 (Subsorting and subtyping). Subsorting \leq : is the least reflexive binary relation such that nat \leq : int.

¹² Subtyping \leq is the largest relation between session types coinductively defined by the following rules:

$$\stackrel{[\text{SUB-END}]}{\text{end} \leqslant \text{end}} \quad \frac{ \begin{bmatrix} \text{SUB-IN} \end{bmatrix}}{\bigwedge_{i \in I \cup J} \mathbf{p}?\ell_i(S_i).\mathsf{T}_i \leqslant \bigwedge_{i \in I} \mathbf{p}?\ell_i(S_i').\mathsf{T}_i'} \quad \begin{array}{c} \begin{bmatrix} \text{SUB-OUT} \end{bmatrix} \\ \forall i \in I: \quad S_i \leq :S_i' \quad \mathsf{T}_i \leqslant \mathsf{T}_i' \\ \hline \bigvee_{i \in I} \mathbf{p}!\ell_i(S_i).\mathsf{T}_i \leqslant \bigvee_{i \in I \cup J} \mathbf{p}!\ell_i(S_i').\mathsf{T}_i' \\ \hline \end{array}$$

Intuitively, the session subtyping \leq in Definition 3.10 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* (Pierce, 2002, Chapter 21).

• Rule [SUB-END] says that end is only subtype of itself.

• Rule [SUB-IN] relates external choices from the same participant **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 $\ell_i(nat)$ (for some $i \in I$), then the subtype can support $\ell_i(int)$ or $\ell_i(nat)$, since $nat \leq$: int and $nat \leq$: nat.

• Rule [SUB-OUT] relates internal choices towards the same participant 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 $\ell_i(int)$ (for some $i \in I$), then the subtype can allow to send $\ell_i(int)$ or $\ell_i(nat)$, since $int \leq :$ int and $nat \leq :$ int.

²⁸ Lemma 3.11. The subtyping relation \leq is reflexive and transitive.

²⁹ 3.3. Type system

We now introduce a type system for the multiparty session calculus presented in Section 2. We distinguish three kinds of typing judgments:

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

¹Readers 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 Gay and Hole (1999, 2005). 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 Gay (2016).

 $\Gamma \vdash \mathsf{n}: \texttt{nat} \qquad \Gamma \vdash \mathsf{i}: \texttt{int} \qquad \Gamma \vdash \mathsf{true}: \texttt{bool} \qquad \Gamma \vdash \mathsf{false}: \texttt{bool} \qquad \Gamma, x: S \vdash x: S$

$\Gamma \vdash \texttt{e}: \texttt{nat}$	$\Gamma \vdash \texttt{e}:\texttt{int}$	$\Gamma \vdash \texttt{e}:\texttt{bool}$
$\Gamma \vdash \texttt{succ}(e):\texttt{na}$	t $\Gamma \vdash \texttt{neg}(\texttt{e}) : \texttt{int}$	$\overline{\Gamma \vdash \neg \texttt{e}:\texttt{bool}}$
$\Gamma \vdash e_1: S \Gamma \vdash e_2: S$	$\Gamma \vdash e_1: \mathtt{int} \Gamma \vdash e_2: \mathtt{int}$	$\Gamma \vdash \mathbf{e}: S S \leq: S'$
$\Gamma \vdash e_1 \oplus \overline{e_2 : S}$	$\Gamma \vdash e_1 > e_2:\texttt{bool}$	$\Gamma \vdash e:S'$

Table 4: Typing rules for expressions.

¹ where Γ is the *typing environment*:

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

² i.e., a mapping that associates expression variables with sorts, and process variables with session types.

³ We say that a multiparty session \mathcal{M} is well typed if there is a global type G such that $\vdash \mathcal{M} : \mathsf{G}$. If a ⁴ multiparty session is well typed, we will sometimes write just $\vdash \mathcal{M}$.

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 T is also typed by the supertype T';

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

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

• [T-INPUT-CHOICE] types a summation of input prefixes as a branching type. It requires that each input prefix targets the same participant \mathbf{q} , and that, for all $i \in I$, each continuation process P_i is typed by the continuation type 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 3.2);

• [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 T if its sub-processes can be typed by 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 G. The condition $pt{G} \subseteq {p_i \mid i \in I}$ allows to also type sessions containing $p \triangleleft 0$: this is needed to assure invariance of typing.

Example 3.12. We show that the multiparty session M from Example 2.2 is well typed. Consider the following global type:

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

 $\vdash P_{\texttt{Alice}}: \texttt{T}_{\texttt{Alice}} \qquad \vdash P_{\texttt{Bob}}: \texttt{T}_{\texttt{Bob}} \qquad \vdash P_{\texttt{Carol}}: \texttt{T}_{\texttt{Carol}}$

Table 5: Typing rules for processes and sessions.

1 where:

² Now, let:

$$\mathsf{T}'_{\texttt{Alice}} = \mathsf{Bob}! \ell_1(\texttt{nat}).\mathsf{Carol}? \ell_3(\texttt{nat}).\mathsf{end} \lor \mathsf{Bob}! \ell_4(\texttt{nat}).\mathsf{Carol}? \ell_3(\texttt{nat}).\mathsf{end}$$

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

 $\mathsf{G}{\upharpoonright} \texttt{Alice} = \mathsf{T}'_{\texttt{Alice}} \qquad \mathsf{G}{\upharpoonright} \texttt{Bob} = \mathsf{T}_{\texttt{Bob}} \qquad \mathsf{G}{\upharpoonright} \texttt{Carol} = \mathsf{T}_{\texttt{Carol}}$

4 we conclude:

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

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.

¹⁰ **Definition 3.13** (Global types consumption and reduction). The consumption of the communication $\mathbf{p} \xrightarrow{\ell} \mathbf{q}$ ¹¹ for the global type \mathbf{G} (notation $\mathbf{G} \setminus \mathbf{p} \xrightarrow{\ell} \mathbf{q}$) is the global type coinductively defined as follows:

¹² The reduction of global types is the smallest pre-order relation closed under the rule: $G \Longrightarrow G \setminus p \xrightarrow{\ell} q$

13 Example 3.14. We show that a projection of a global type before the consumption might require to support

 $_{14}$ more external choices than the projection after the consumption. Take G, its subterm G₁, from Example 3.8,

 $_{15}$ $\,$ and their types denoted as ${\sf G}$ and ${\sf G}_1,$ respectively. Also take the projection:

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

and recall the explanation on how G[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:

$$\begin{split} \mathsf{G} \setminus \mathsf{p} \xrightarrow{\ell_1} \mathsf{q} &= \mathsf{G}_1 &= \mathsf{q} \to \mathsf{r} : \{\ell_3(\texttt{int}), \ell_5(\texttt{nat})\} \\ (\mathsf{G} \setminus \mathsf{p} \xrightarrow{\ell_1} \mathsf{q}) \restriction \mathsf{r} &= \mathsf{G}_1 \restriction \mathsf{r} &= \mathsf{q}?\ell_3(\texttt{int}) \land \mathsf{q}?\ell_5(\texttt{nat}) \end{split}$$

and we obtain $G \upharpoonright r \leq (G \setminus p \xrightarrow{\ell_1} q) \upharpoonright r$. 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 \mathbf{r} in \mathbf{G} , and is therefore typed by $\mathbf{G} \upharpoonright \mathbf{r}$, has to support the input of ℓ_4 from \mathbf{q} , by rule [T-INPUT-CHOICE] in Table 5. After the transition from \mathbf{G} to \mathbf{G}_1 , the same process is also τ_7 typed by $\mathbf{G}_1 \upharpoonright \mathbf{r}$, by rule [T-SUB] — but will never receive a message ℓ_4 from \mathbf{q} .

⁸ We can now prove subject reduction.

⁹ Theorem 3.15 (Subject Reduction). Let $\vdash \mathcal{M} : \mathsf{G}$. For all \mathcal{M}' , if $\mathcal{M} \longrightarrow \mathcal{M}'$, then $\vdash \mathcal{M}' : \mathsf{G}'$ for some G' ¹⁰ such that $\mathsf{G} \Longrightarrow \mathsf{G}'$.

Theorem 3.16 (Progress). If $\vdash \mathcal{M} : \mathsf{G}$, then either $\mathcal{M} \equiv \mathsf{p} \triangleleft \mathsf{0}$ 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 3.17 (Type Safety). If $\vdash M : G$, then it does not hold stuck(M).

¹⁵ *Proof.* Direct consequence of Theorem 3.15, Theorem 3.16, and Definition 2.3.

16 References

- Deniélou, P., Yoshida, N., Bejleri, A., Hu, R., 2012. Parameterised multiparty session types. Logical Methods in Computer
 Science 8. doi:10.2168/LMCS-8(4:6)2012.
- Gay, S., Hole, M., 2005. Subtyping for session types in the pi calculus. Acta Informatica 42, 191–225. doi:10.1007/
 s00236-005-0177-z.
- 21 Gay, S., Ravera, A. (Eds.), 2017. Behavioural Types: from Theory to Tools. River Publishers.
- Gay, S.J., 2016. Subtyping supports safe session substitution, in: A List of Successes That Can Change the World: Essays
 Dedicated to Philip Wadler on the Occasion of His 60th Birthday, Springer International Publishing. pp. 95–108. doi:10.
 1007/978-3-319-30936-1_5.
- Gay, S.J., Hole, M., 1999. Types and subtypes for client-server interactions, in: ESOP, pp. 74–90. doi:10.1007/3-540-49099-X\
 _6.
- ²⁷ Honda, K., 1993. Types for dyadic interaction, in: CONCUR'93, pp. 509-523.
- Honda, K., Vasconcelos, V.T., Kubo, M., 1998. Language primitives and type disciplines for structured communication-based
 programming, in: ESOP, Springer. pp. 22–138. doi:10.1007/BFb0053567.
- Honda, K., Yoshida, N., Carbone, M., 2008. Multiparty asynchronous session types, in: POPL, ACM Press. pp. 273–284.
 doi:10.1145/1328438.1328472.

32 Honda, K., Yoshida, N., Carbone, M., 2016. Multiparty asynchronous session types. Journal of ACM 63, 1–67.

- 33 Hüttel, H., Lanese, I., Vasconcelos, V.T., Caires, L., Carbone, M., Deniélou, P.M., Mostrous, D., Padovani, L., Ravara, A.,
- Tuosto, E., Vieira, H.T., Zavattaro, G., 2016. Foundations of session types and behavioural contracts. ACM Comput. Surv. 49. doi:10.1145/2873052.
- Kouzapas, D., Yoshida, N., 2013. Globally governed session semantics, in: CONCUR, Springer. pp. 395–409. doi:10.1145/
 1328438.1328472.
- 38 Kouzapas, D., Yoshida, N., 2015. Globally governed session semantics. Logical Methods in Computer Science 10.
- ³⁹ Milner, R., 1989. Communication and Concurrency. Prentice-Hall, Inc.
- 40 Pierce, B.C., 2002. Types and Programming Languages. MIT Press.
- Takeuchi, K., Honda, K., Kubo, M., 1994. An Interaction-based Language and its Typing System, in: PARLE'94, pp. 398–413.
 doi:10.1007/3-540-58184-7_118.
- Yoshida, N., Deniélou, P., Bejleri, A., Hu, R., 2010. Parameterised multiparty session types, in: FOSSACS, Springer. pp.
 128–145. doi:10.1007/978-3-642-12032-9_10.