

Asynchronous Sessions with Input Races

Ilaria Castellani*

INRIA, Université Côte d’Azur, France

ilaria.castellani@inria.fr

Mariangiola Dezani-Ciancaglini

Dipartimento di Informatica, Università di Torino, Italy

dezani@di.unito.it

Paola Giannini[†]

DiSSTE, Università del Piemonte Orientale, Italy

paola.giannini@uniupo.it

We propose a calculus for asynchronous multiparty sessions where input choices with different senders are allowed in processes. We present a type system that accepts such input races provided they do not hinder lock-freedom.

1 Introduction

The foundational work on multiparty sessions [11] introduced the notion of *global type* for specifying the overall behaviour of multiparty protocols. The criterion for a session implementation to be correct with respect to its specification was formalised via the notion of projection: each process implementing the behaviour of a session participant was required to type-check against the local type obtained by projecting the global type on that participant.

However, the work [11] imposed strong restrictions on the syntax of global types, requiring all initial communications in the branches of a choice to have the same sender and the same receiver, and every third participant¹ to have the same behaviour in all branches. Although these were useful simplifying assumptions in order to achieve multiparty session correctness, they limited the expressiveness of global types, ruling out relevant protocols. For this reason, more permissive choice constructors were investigated in subsequent work [1, 2, 9, 16, 13, 3, 5, 14, 10, 17]. A widely adopted relaxation of the choice operator, originally proposed in [1], allows third participants to behave differently in different branches, provided they are notified of which branch has been chosen. Later proposals [9, 16, 13] accommodate processes with output choices among different receivers, for instance a client choosing among different servers. On the other hand *input races*, namely input choices among different senders, continued to be considered as problematic. As a consequence, common protocols such as a server shared by different clients could not be specified by means of global types.

Recent proposals introduce more flexibility in input choices for processes [3, 5, 14, 10, 17]. The work [10] defines the property of *race-freedom* for sessions as the absence, at any stage of computation, of a branching between communications from different senders towards the same receiver leading to distinct target states. A rather permissive type system is proposed, which is shown to be both sound and complete for a range of liveness properties when restricting attention to race-free sessions. The work [14] also addresses the input race problem, referred to as the “+–problem” there. While the proposed syntax for global and local types is completely free, two well-formedness conditions are imposed on types,

*This research has been supported by the ANR17-CE25-0014-01 CISC project.

[†]This original research has the financial support of the Università del Piemonte Orientale.

¹We call “third participant” any participant which is not involved in the first communication of a branch.

which are meant to prevent dangerous races. Sessions are synchronous in [10, 14] and asynchronous in [17], which is the work that is closest to ours. In that paper, input races are allowed under sophisticated conditions on projections of global types. These conditions track causalities between messages, and their soundness proof uses novel graph-theoretic techniques from the theory of message-sequence charts.

Consider for example the following session, where two participants p and q send concurrently a message to a third participant r , which is ready to receive both messages in any order:

Example 1.1 (Confluent input race) $p[[r!\ell]] \parallel q[[r!\ell']] \parallel r[[p?\ell; q?\ell' + q?\ell'; p?\ell]]$

No matter whether communication is synchronous or asynchronous, this session incurs a race². However, this race may be viewed as innocuous since after any branch is chosen, the input of the other branch is still available, leading to the same target state. Indeed, the race consists here of a choice between two different sequentialisations of concurrent inputs³. This kind of input race will be called *confluent*. The above session is well typed in [14], but not in [10, 17], where the syntax of global types forbids different senders in a choice. In [10], this session is also ruled out by the race-freedom condition.

By contrast, the following asynchronous session, where there is an apparent input race in the process of participant r , is actually race-free according to [10] because it cannot evolve to a state in which both inputs of r are simultaneously enabled. This kind of uneffective input race, which results from an agreement between the senders such that in every computation only one of them sends a message to the receiver, will be called *fake*.

Example 1.2 (Fake input race in asynchronous session)

$$p[[q!a; q?a; (q!\ell; r!b \oplus q!\ell')]] \parallel q[[p!a; p?a; (p?\ell + p?\ell'; r!b)]] \parallel r[[p?b + q?b]]$$

This session implements the following protocol between Alice, Bob and Carol, represented by participants p , q and r respectively:

- *Alice and Bob send each other the message “I arrived” and then they read their messages;*
- *Alice sends Bob either the message “I will tell Carol”, after which she sends Carol the message “We arrived”, or the message “Please tell Carol”;*
- *Bob reads either the message “I will tell Carol”, or the message “Please tell Carol” after which he sends Carol the message “We arrived”;*
- *Carol reads the message “We arrived” with sender either Alice or Bob.*

The session of Example 1.2 cannot be typed in [14, 10, 17] because the syntax of global types does not allow two participants to exchange messages by first performing both outputs and then both inputs. If we omit the initial exchange of messages between Alice and Bob, the resulting session can be typed in [10, 17] but not in [14].

Our goal is to devise a type system for asynchronous sessions that is permissive enough to accept the sessions of Example 1.1 and Example 1.2, while rejecting dangerous races that could lead to deadlock or starvation. In particular, we will not be able to type the sessions discussed in the introductions of [14, 17], since they both have a possibility of starvation.

For typing asynchronous sessions we use global types that split communications into outputs and inputs, following the approach advocated in [4, 7]. For instance, the session of Example 1.2 has the following global type: $p!q.a; q!p.a; p?q.a; q?p.a; p!\{q.\ell; G_1, q.\ell'; G_2\}$ where $G_1 = p!r.b; q?p.\ell; r?\{p.b, q.b\}$ and $G_2 = q?p.\ell'; q!r.b; r?\{p.b, q.b\}$. Here $p!q.a$ denotes a send from p to q of label a , $p?q.a$ denotes a receive by p from q of label a , $p!\{q.\ell; G_1, q.\ell'; G_2\}$ is an output choice with sender p and receiver q , and $r?\{p.b, q.b\}$ is an input choice with receiver r and senders p and q .

²Either initially, if communication is synchronous, or after performing both outputs, if communication is asynchronous.

³If we had a parallel construct $|$ for processes, this situation would be represented as $(p?\ell | q?\ell')$.

The rest of the paper is organised as follows. In Section 2 we introduce our calculus for asynchronous sessions. In Section 3 we define the syntax and semantics of global types. In Section 4 we present our type system, illustrate it with some examples, and establish the properties of Subject Reduction, Session Fidelity and Lock-freedom. We conclude in Section 5 with a discussion on future and related work.

2 Asynchronous Sessions

We assume the following base sets: *participants*, ranged over by p, q, r and forming the set Part , and *labels*, ranged over by ℓ, ℓ', \dots and forming the set Lab .

Definition 2.1 (Processes) Processes are defined by:

$$P ::=_{\rho} \mathbf{0} \mid \bigoplus_{i \in I} p_i ! \ell_i ; P_i \mid \bigoplus_{i \in I} p_i ? \ell_i ; P_i$$

where $I \neq \emptyset$ and $p_h ! \ell_h \neq p_k ! \ell_k$ and $p_h ? \ell_h \neq p_k ? \ell_k$ for $h, k \in I$ and $h \neq k$.

A process may be terminated, or it is an internal choice of outputs or an external choice of inputs. The symbol $::=_{\rho}$, in Definition 2.1 and in later definitions, indicates that the productions should be interpreted *coinductively* (they define possibly infinite processes) and that we focus on *regular* terms, namely, terms with finitely many distinct subterms. In this way, we only obtain processes which are solutions of finite sets of equations, see [6]. So, when writing processes, we shall use (mutually) recursive equations.

In the following, we will omit trailing $\mathbf{0}$'s when writing processes.

In a full-fledged calculus, processes would exchange labels of the form $\ell(v)$, where v is a value. For simplicity, we consider only pure labels here.

In our calculus, asynchronous communication is handled in the standard way, by storing sent labels in a queue together with sender and receiver names. Receivers may then fetch messages from the queue when required. We define *messages* to be triples $\langle p, \ell, q \rangle$, where p is the sender and q is the receiver, and *message queues* (or simply *queues*) to be possibly empty sequences of messages:

$$\mathcal{M} ::= \emptyset \mid \langle p, \ell, q \rangle \cdot \mathcal{M}$$

The order of messages in the queue is the order in which they will be read. Since the only reading order that matters is that between messages with the same sender and the same receiver, we consider message queues modulo the structural equivalence given by:

$$\mathcal{M} \cdot \langle p, \ell, q \rangle \cdot \langle r, \ell', s \rangle \cdot \mathcal{M}' \equiv \mathcal{M} \cdot \langle r, \ell', s \rangle \cdot \langle p, \ell, q \rangle \cdot \mathcal{M}' \text{ if } p \neq r \text{ or } q \neq s$$

Sessions are composed by a number of located processes of the form $p[[P]]$, each enclosed within a different participant p , and by a message queue.

Definition 2.2 (Networks and Sessions) Networks are defined by:

$$\mathbb{N} = p_1[[P_1]] \parallel \dots \parallel p_n[[P_n]] \text{ with } p_h \neq p_k \text{ for any } h \neq k$$

Sessions are defined by:

$$\mathbb{N} \parallel \mathcal{M}$$

where \mathbb{N} is a network, and \mathcal{M} is a message queue.

We assume the standard structural congruence \equiv on networks⁴, stating that parallel composition is associative and commutative and has neutral element $p[[\mathbf{0}]]$ for any fresh p .

If $P \neq \mathbf{0}$ we write $p[[P]] \in \mathbb{N}$ as short for $\mathbb{N} \equiv p[[P]] \parallel \mathbb{N}'$ for some \mathbb{N}' . This abbreviation is justified by the associativity and commutativity of parallel composition.

To define the *operational semantics* of sessions, we use an LTS whose transitions are decorated by outputs or inputs. Therefore, we define the set of *input/output communications* (communications for

⁴By abuse of notation, we use the same symbol as for structural equivalence on queues.

$$\begin{aligned}
& p[\oplus_{i \in I} q_i ! \ell_i; P_i] \parallel \mathbb{N} \parallel \mathcal{M} \xrightarrow{p!q_k \cdot \ell_k} p[[P_k]] \parallel \mathbb{N} \parallel \mathcal{M} \cdot \langle p, \ell_k, q_k \rangle \quad \text{where } k \in I \quad [\text{SEND}] \\
& q[\sum_{j \in J} p_j ? \ell_j; Q_j] \parallel \mathbb{N} \parallel \langle p_k, \ell_k, q \rangle \cdot \mathcal{M} \xrightarrow{q?p_k \cdot \ell_k} q[[Q_k]] \parallel \mathbb{N} \parallel \mathcal{M} \quad \text{where } k \in J \quad [\text{RCV}]
\end{aligned}$$

Figure 1: LTS for sessions.

short), ranged over by β, β' , to be $\{p!q.\ell, p?q.\ell \mid p, q \in \text{Part}, \ell \in \text{Lab}\}$, where $p!q.\ell$ represents the output of the label ℓ from participant p to participant q , and $p?q.\ell$ the input by participant p of the label ℓ sent by participant q .

The LTS semantics of networks, defined modulo \equiv , is specified by the two Rules [SEND] and [RCV] given in Figure 1. Rule [SEND] allows a participant p with an internal choice (a sender) to send to the participant q_k the label ℓ_k by adding it to the queue. Symmetrically, Rule [RCV] allows a participant q with an external choice (a receiver) to read the label ℓ_k sent by participant p_k , provided this label is among the ℓ_j 's specified in the choice. Thanks to structural equivalence, the first message from p_k to q that appears in the queue, if any, can always be moved to the top of the queue.

A key role in this paper is played by (possibly empty) sequences of communications. As usual we define them as traces.

Definition 2.3 (Traces) *(Finite) traces are defined by:*

$$\tau ::= \varepsilon \mid \beta \cdot \tau$$

We use $|\tau|$ to denote the length of the trace τ .

When $\tau = \beta_1 \cdot \dots \cdot \beta_n$ ($n \geq 1$) we write $\mathbb{N} \parallel \mathcal{M} \xrightarrow{\tau} \mathbb{N}' \parallel \mathcal{M}'$ as short for

$$\mathbb{N} \parallel \mathcal{M} \xrightarrow{\beta_1} \mathbb{N}_1 \parallel \mathcal{M}_1 \cdots \xrightarrow{\beta_n} \mathbb{N}_n \parallel \mathcal{M}_n = \mathbb{N}' \parallel \mathcal{M}'$$

We now introduce the notion of player, which is characteristic of asynchronous communication, where only one of the involved participants is active, namely the sender for an output communication and the receiver for an input communication. The player of a communication β is the participant who is active in β . The set of players of a trace is then obtained by collecting the players of all its communications.

Definition 2.4 (Players of communications and traces) *We denote by $\text{play}(\beta)$ the player of a communication β defined by*

$$\text{play}(p!q.\ell) = \text{play}(p?q.\ell) = p$$

We denote by $\text{Players}(\tau)$ the set of players of a trace τ defined by

$$\text{Players}(\varepsilon) = \emptyset \quad \text{Players}(\beta \cdot \tau) = \{\text{play}(\beta)\} \cup \text{Players}(\tau)$$

3 Global Types

As in [4, 7], our global types can be obtained from the standard ones [11, 12] by splitting output and input communications. The novelty is that we allow multiple receivers in output choices and multiple senders in input choices.

Definition 3.1 (Global types) *Global types G are defined by the following grammar:*

$$G ::=_{\rho} p!\{q_i.\ell_i; G_i\}_{i \in I} \mid p?\{q_i.\ell_i; G_i\}_{i \in I} \mid \text{End}$$

where $I \neq \emptyset$, $p \neq q_i$ for all $i \in I$ and $q_h.\ell_h \neq q_k.\ell_k$ for $h, k \in I$ and $h \neq k$.

As for processes, $::=_{\rho}$ indicates that global types are coinductively defined and *regular*.

The global type $p!\{q_i.l_i; G_i\}_{i \in I}$ specifies that player p sends the label ℓ_k with $k \in I$ to participant q_k and then the interaction described by the global type G_k takes place. The global type $p?\{q_i.l_i; G_i\}_{i \in I}$ specifies that player p receives the label ℓ_k with $k \in I$ from participant q_k and then the interaction described by the global type G_k takes place.

We define $\text{Players}(G)$ as the smallest set satisfying the following equations:

$$\text{Players}(\text{End}) = \emptyset \quad \text{Players}(p!\{q_i.l_i; G_i\}_{i \in I}) = \text{Players}(p?\{q_i.l_i; G_i\}_{i \in I}) = \{p\} \cup \bigcup_{i \in I} \text{Players}(G_i)$$

The regularity of global types ensures that the sets of players are finite. In Section 2 we used the same notation for the players of traces. In all cases, the context should make it easy to understand which function is in use.

To avoid starvation we require global types to satisfy a boundedness condition. To formalise boundedness we use ξ to denote a *path* in global type trees, i.e., a possibly infinite sequence of communications β . Note that a finite path is a trace in the sense of Definition 2.3. We extend the notation \cdot to denote also the concatenation of a finite sequence with a possibly infinite sequence. The function Paths gives the set of *paths* of a global type, which is the greatest set such that:

$$\begin{aligned} \text{Paths}(\text{End}) &= \{\varepsilon\} \\ \text{Paths}(p!\{q_i.l_i; G_i\}_{i \in I}) &= \bigcup_{i \in I} \{p!q_i.l_i \cdot \xi \mid \xi \in \text{Paths}(G_i)\} \\ \text{Paths}(p?\{q_i.l_i; G_i\}_{i \in I}) &= \bigcup_{i \in I} \{p?q_i.l_i \cdot \xi \mid \xi \in \text{Paths}(G_i)\} \end{aligned}$$

If $x \in \mathbf{N} \cup \{\omega\}$ is the length of ξ , we denote by $\xi[n]$ the n -th communication in the path ξ , where $1 \leq n < x$. It is handy to define the *depth* of a player p in a global type G , $\text{depth}(G, p)$.

Definition 3.2 (Depth of a player) *Let G be a global type. For $\xi \in \text{Paths}(G)$ set*

$$\text{depth}(\xi, p) = \inf\{n \mid \text{play}(\xi[n]) = p\}$$

and define $\text{depth}(G, p)$, the depth of p in G , as follows:

$$\text{depth}(G, p) = \begin{cases} \sup\{\text{depth}(\xi, p) \mid \xi \in \text{Paths}(G)\} & p \in \text{Players}(G) \\ 0 & \text{otherwise} \end{cases}$$

Note that $\text{depth}(G, p) = 0$ iff $p \notin \text{Players}(G)$. Moreover, if $p \neq \text{play}(\xi[n])$ for all $n \in \mathbf{N}$, then $\text{depth}(\xi, p) = \inf \emptyset = \infty$. Hence, if p is a player of a global type G and there is some path in G where p does not occur as a player, then $\text{depth}(G, p) = \infty$.

Definition 3.3 (Boundedness) *A global type G is bounded if $\text{depth}(G', p)$ is finite for all participants $p \in \text{Players}(G)$ and all types G' which occur in G .*

Example 3.4 *The following example shows the necessity of considering all types occurring in a global type for defining boundedness. Consider $G = r!q.l; q?r.l; G'$, where*

$$G' = p!\{q.l_1; q?p.l_1; q!r.l_3; r?q.l_3, q.l_2; q?p.l_2; G'\}$$

Then we have: $\text{depth}(G, p) = 3, \text{depth}(G, q) = 2, \text{depth}(G, r) = 1$, whereas $\text{depth}(G', p) = 1, \text{depth}(G', q) = 2, \text{depth}(G', r) = \infty$.

Since global types are regular the boundedness condition is decidable.

Global types in parallel with queues, dubbed *type configurations*, are given semantics by means of the LTS in Figure 2. The first two rules allow top level outputs and inputs to be performed in the standard way. The remaining two rules allow communications to be performed inside output and input choices. These inside rules are needed to enable interleaving between independent communications despite the sequential structure of global types. For example, we want to allow $p!q.l; r!s.l' \parallel \emptyset \xrightarrow{r!s.l'} p!q.l \parallel \langle r, l', s \rangle$ when $p \neq r$, because, intuitively, outputs performed by different players should be independent. This

$$\begin{array}{c}
\text{[TOP-OUT]} \frac{}{p!\{q_i.l_i; G_i\}_{i \in I} \parallel \mathcal{M} \xrightarrow{p!q_h.l_h} G_h \parallel \mathcal{M} \cdot \langle p, \ell_h, q_h \rangle} h \in I \\
\text{[TOP-IN]} \frac{}{p?\{q_i.l_i; G_i\}_{i \in I} \parallel \langle q_h, \ell_h, p \rangle \cdot \mathcal{M} \xrightarrow{q_h?p.l_h} G_h \parallel \mathcal{M}} h \in I \\
\text{[INSIDE-OUT]} \frac{G_i \parallel \mathcal{M} \cdot \langle p, \ell_i, q_i \rangle \xrightarrow{\beta} G'_i \parallel \mathcal{M}' \cdot \langle p, \ell_i, q_i \rangle \quad \forall i \in I}{p!\{q_i.l_i; G_i\}_{i \in I} \parallel \mathcal{M} \xrightarrow{\beta} p!\{q_i.l_i; G'_i\}_{i \in I} \parallel \mathcal{M}'} p \neq \text{play}(\beta) \\
\text{[INSIDE-IN]} \frac{G_j \parallel \mathcal{M} \xrightarrow{\beta} G'_j \parallel \mathcal{M}' \quad \forall j \in J}{p?\{q_i.l_i; G_i\}_{i \in I} \parallel \mathcal{M} \xrightarrow{\beta} p?\{q_i.l_i; G'_i\}_{i \in I} \parallel \mathcal{M}'} \begin{array}{l} J = \text{rm}(\{\langle q_i, \ell_i, p \rangle\}_{i \in I}, \mathcal{M}) \neq \emptyset \\ p \neq \text{play}(\beta) \quad \beta \neq q_l!p.l_l \\ G'_i = G_k \quad k \in J \quad \forall l \in I \setminus J \end{array}
\end{array}$$

Figure 2: LTS for type configurations.

justifies the condition $p \neq \text{play}(\beta)$ in Rules [INSIDE-OUT] and [INSIDE-IN]. In Rule [INSIDE-OUT] we require all branches to be able to perform the β transition. This avoids for example:

$$\begin{array}{c}
p!\{q.l; q?p.l; r!p.l; p?r.l, q.l'; q?p.l'; r!p.l'; p?r.l'\} \parallel \emptyset \xrightarrow{r!p.l} \\
p!\{q.l; q?p.l; p?r.l, q.l'; q?p.l'; r!p.l'; p?r.l'\} \parallel \langle r, \ell, p \rangle
\end{array}$$

which, in case we choose the right branch, leads to the configuration $p?r.l' \parallel \langle r, \ell, p \rangle \cdot \langle r, \ell', p \rangle$.

The shapes of the queues appearing in the premise of Rule [INSIDE-OUT] ensure that β is not the matching input for any output in the choice. In Rule [INSIDE-IN], we consider only the branches with corresponding messages on top of the queue (called *live* branches), using the index set of ready messages $\text{rm}(\{\langle q_i, \ell_i, p \rangle\}_{i \in I}, \mathcal{M})$ defined as follows, where m ranges over messages.

Definition 3.5 *Given a set of messages $\{m_i\}_{i \in I}$ and a queue \mathcal{M} , the index set of the “ready messages” in this set is defined by: $\text{rm}(\{m_i\}_{i \in I}, \mathcal{M}) = \{i \in I \mid \mathcal{M} \equiv m_i \cdot \mathcal{M}_i\}$.*

The mapping rm plays a crucial role also in the typing rule for input choices, as we will see in Section 4. The condition $J \neq \emptyset$ means that there is at least one live branch. The condition $\beta \neq q_l!p.l_l$ for all $l \in I \setminus J$ ensures that the occurrence of β does not generate a message that would “awaken” some dead, i.e. not live, branch of the choice. In the resulting choice, the dead branches become an arbitrary live branch (condition $G'_i = G_k$ for some $k \in J$ and for all $l \in I \setminus J$). In fact, such branches could also be omitted as they can never be awakened.

It is easy to check that the LTS of type configurations preserves boundedness of global types. Therefore, from now on we will assume all our global types to be bounded.

4 Type System

Global types are an abstraction of sessions. Usually, global types are projected to participants, yielding local types which are assigned to processes. The simplicity of our calculus and the flexibility of our global types allow us to formulate a type system where global types are directly derived for sessions, using judgements of the form $G \vdash \mathbb{N} \parallel \mathcal{M}$. The typing rules are given in Figure 3.

Rules [OUT] and [IN] just add simultaneously outputs and inputs to global types and to the corresponding processes inside networks. The condition $\text{Players}(G_i) \setminus \{p\} = \text{Players}(\mathbb{N})$ for all $i \in I$ ensures that all players in \mathbb{N} are also players in G . For example, this condition prevents the derivation of

$$\begin{array}{c}
\text{[END]} \frac{}{\text{End} \vdash p[\mathbf{0}] \parallel \mathbf{0}} \\
\text{[OUT]} \frac{G_i \vdash p[P_i] \parallel N \parallel \mathcal{M} \cdot \langle p, \ell_i, q_i \rangle \quad \text{Players}(G_i) \setminus \{p\} = \text{Players}(N) \quad \forall i \in I}{p! \{q_i, \ell_i; G_i\}_{i \in I} \vdash p[\bigoplus_{i \in I} q_i! \ell_i; P_i] \parallel N \parallel \mathcal{M}} \\
\text{[IN]} \frac{G_j \vdash p[P_j] \parallel N \parallel \mathcal{M}_j \quad \mathbf{l}(G_j, \mathbb{M}) \quad \forall j \in J \quad \text{Players}(G_j) \setminus \{p\} = \text{Players}(N) \quad \forall i \in I}{p? \{q_i, \ell_i; G_i\}_{i \in I} \vdash p[\sum_{h \in H} q_h? \ell_h; P_h] \parallel N \parallel \mathcal{M}}
\end{array}$$

$$\begin{array}{l}
J = \text{rm}(\{ \langle q_i, \ell_i, p \rangle \}_{i \in I}, \mathcal{M}) \\
= \text{rm}(\{ \langle q_h, \ell_h, p \rangle \}_{h \in H}, \mathcal{M}) \neq \emptyset \\
\mathcal{M} \equiv \langle q_j, \ell_j, p \rangle \cdot \mathcal{M}_j \quad \forall j \in J \\
\mathbb{M} = \{ \langle q_l, \ell_l, p \rangle \mid l \in (I \cup H) \setminus J \text{ \& } q_l \neq q_j \quad \forall j \in J \}
\end{array}$$

Figure 3: Typing rules for sessions.

$G \vdash p[P] \parallel q[Q]$ with $G = p!q.\ell; G$ and $P = q!\ell; P$ and Q arbitrary.

Rule [OUT] considers all branches of the global type, since the choice of the sent message is arbitrary. This rule requires that the session resulting from the output of a branch be typed with the corresponding branch of the global type.

Rule [IN] requires that the global type and the process read the same messages on the queue. To this end, it uses the index set of ready messages defined in Definition 3.5, collecting the indices of the live branches of the global type and of the input process⁵, and asking them to be equal (condition $\text{rm}(\{ \langle q_i, \ell_i, p \rangle \}_{i \in I}, \mathcal{M}) = \text{rm}(\{ \langle q_h, \ell_h, p \rangle \}_{h \in H}, \mathcal{M})$). This set of indices must not be empty (condition $J \neq \emptyset$). Only the branches of the global type and of the input process thus selected are compared in the premises of Rule [IN]. Note that in this way we allow more freedom than in the synchronous subtyping for session types [8]. In Rule [IN], in order to ensure the condition $\beta \neq q_l!p.\ell_l$ for all $l \in I \setminus J$ required by the transition Rule [INSIDE-IN], we want to prevent the enqueueing of messages that would transform a dead branch of the process or of the global type into a live branch. To this end, we introduce a predicate which forbids a global type to generate such messages. Let \mathbb{M} range over sets of messages.

Definition 4.1 *The type G is inactive for the set of messages \mathbb{M} , if $\mathbf{l}(G, \mathbb{M})$ holds, where:*

$$\begin{array}{l}
\mathbf{l}(\text{End}, \mathbb{M}) \quad \mathbf{l}(p? \{q_i, \ell_i; G_i\}_{i \in I}, \mathbb{M}) \quad \text{if } \mathbf{l}(G_i, \mathbb{M}) \quad \forall i \in I \\
\mathbf{l}(p! \{q_i, \ell_i; G_i\}_{i \in I}, \mathbb{M}) \quad \text{if } \langle p, \ell_i, q_i \rangle \notin \mathbb{M} \text{ and } \mathbf{l}(G_i, \mathbb{M}) \quad \forall i \in I
\end{array}$$

The predicate $\mathbf{l}(G, \mathbb{M})$ looks for outputs in G which produce messages in \mathbb{M} . The regularity of global types guarantees the computability of this predicate. Notice that $\mathbf{l}(G, \mathbb{M})$ also ensures that the network cannot produce messages in \mathbb{M} . This is due to the typing Rule [OUT] prescribing that messages put on the queue by the global type be the same as the ones of the network.

For example consider the following sequence of transitions

$$\begin{array}{l}
q[r!\ell'] \parallel r[p? \ell; q? \ell' + q? \ell'; p? \ell'] \parallel \langle p, \ell, r \rangle \xrightarrow{q!r.\ell'} r[p? \ell; q? \ell' + q? \ell'; p? \ell'] \parallel \langle p, \ell, r \rangle \cdot \langle q, \ell', r \rangle \\
 \xrightarrow{r?q.\ell'} r[p? \ell'] \parallel \langle p, \ell, r \rangle
\end{array}$$

Since the input and the message in $r[p? \ell'] \parallel \langle p, \ell, r \rangle$ do not match, this session cannot be typed and therefore also the session $q[r!\ell'] \parallel r[p? \ell; q? \ell' + q? \ell'; p? \ell'] \parallel \langle p, \ell, r \rangle$ should not be typable. Without checking the inactivity predicate we can type this session by the global type

$$(*) \quad r? \{p.\ell; q!r.\ell'; r?q.\ell', q.\ell'; q!r.\ell'; r?p.\ell'\}$$

⁵As for global types, a branch of the input process is live if it has a corresponding message on top of the queue, and dead otherwise.

as follows:

$$\frac{\frac{\frac{\text{End} \vdash r[\mathbf{0}] \parallel \emptyset}{r?q.\ell' \vdash r[q?\ell'] \parallel \langle q, \ell', r \rangle}}{q!r.\ell'; r?q.\ell' \vdash q[r!\ell'] \parallel r[q?\ell'] \parallel \emptyset}}{r?\{p.\ell; q!r.\ell'; r?q.\ell', q.\ell'; q!r.\ell'; r?p.\ell'\} \vdash q[r!\ell'] \parallel r[p?\ell; q?\ell' + q?\ell'; p?\ell'] \parallel \langle p, \ell, r \rangle}}$$

The problem here is that Rule [IN] does not check the dead branches of the global type. The role of the inactivity predicate is just to ensure that the transitions will not awake dead branches. This is done by checking the outputs in the live branches. In this example the output $q!r.\ell'$ is in the branch starting with the input $r?p.\ell$ and the queue contains $\langle p, \ell, r \rangle$. So the typing Rule [IN] cannot be applied since $\mathbf{I}(q!r.\ell'; r?q.\ell', \{\langle q, \ell', r \rangle\})$ does not hold.

Notice that the session in Example 1.1 has the transition

$p[r!\ell] \parallel q[r!\ell'] \parallel r[p?\ell; q?\ell' + q?\ell'; p?\ell] \parallel \emptyset \xrightarrow{p!r.\ell} q[r!\ell'] \parallel r[p?\ell; q?\ell' + q?\ell'; p?\ell] \parallel \langle p, \ell, r \rangle$
and the resulting session differs from that of the previous example only for the label of the last input. Correspondingly, the global types $r?\{p.\ell; q!r.\ell'; r?q.\ell', q.\ell'; r?p.\ell\}$ and $r?\{p.\ell; q!r.\ell'; r?q.\ell', q.\ell'; r?p.\ell'\}$ only differ for the labels of the underlined inputs. Therefore $r?\{p.\ell; q!r.\ell'; r?q.\ell', q.\ell'; r?p.\ell\}$ cannot be derived for the session $q[r!\ell'] \parallel r[p?\ell; q?\ell' + q?\ell'; p?\ell] \parallel \langle p, \ell, r \rangle$, since as we just saw the predicate $\mathbf{I}(q!r.\ell'; r?q.\ell', \{\langle q, \ell', r \rangle\})$ does not hold. In fact, this is expected since the session can do a transition $\xrightarrow{q!r.\ell'}$ that the type configuration cannot mimic. On the other hand, this session can be typed by the global type

$$q!r.\ell'; r?\{p.\ell; r?q.\ell', q.\ell'; r?p.\ell\}$$

as follows:

$$\frac{\frac{\frac{\text{End} \vdash r[\mathbf{0}] \parallel \emptyset}{r?q.\ell' \vdash r[q?\ell'] \parallel \langle q, \ell', r \rangle} \quad \frac{\text{End} \vdash r[\mathbf{0}] \parallel \emptyset}{r?p.\ell \vdash r[p?\ell] \parallel \langle p, \ell, r \rangle}}{r?\{p.\ell; r?q.\ell', q.\ell'; r?p.\ell\} \vdash r[p?\ell; q?\ell' + q?\ell'; p?\ell] \parallel \langle p, \ell, r \rangle \cdot \langle q, \ell', r \rangle}}{q!r.\ell'; r?\{p.\ell; r?q.\ell', q.\ell'; r?p.\ell\} \vdash q[r!\ell'] \parallel r[p?\ell; q?\ell' + q?\ell'; p?\ell] \parallel \langle p, \ell, r \rangle}}$$

In this derivation Rule [OUT] is applied first, and thus Rule [IN] is applied only when the queue contains the matching messages for both branches. Therefore Rule [IN] checks the continuations of both branches, and the inactivity predicate holds trivially for each of them.

Notice that bringing forward the output from q to r in the global type $(*)$ does not enable us to type:

$$q[r!\ell'] \parallel r[p?\ell; q?\ell' + q?\ell'; p?\ell] \parallel \langle p, \ell, r \rangle$$

In fact, we cannot complete the derivation:

$$\frac{\frac{\frac{\text{End} \vdash r[\mathbf{0}] \parallel \emptyset}{r?q.\ell' \vdash r[q?\ell'] \parallel \langle q, \ell', r \rangle} \quad r?p.\ell' \vdash r[p?\ell'] \parallel \langle p, \ell, r \rangle}}{r?\{p.\ell; r?q.\ell', q.\ell'; r?p.\ell'\} \vdash r[p?\ell; q?\ell' + q?\ell'; p?\ell] \parallel \langle p, \ell, r \rangle \cdot \langle q, \ell', r \rangle}}{q!r.\ell'; r?\{p.\ell; r?q.\ell', q.\ell'; r?p.\ell'\} \vdash q[r!\ell'] \parallel r[p?\ell; q?\ell' + q?\ell'; p?\ell] \parallel \langle p, \ell, r \rangle}}$$

Indeed, we cannot apply Rule [IN] to derive the top right judgement $r?p.\ell' \vdash r[p?\ell'] \parallel \langle p, \ell, r \rangle$, since the only input does not match the message in the queue.

We can also type the following recursive version of Example 1.1:

$$p[P] \parallel q[Q] \parallel r[R]$$

where $P = r!\ell; P$, $Q = r!\ell'; Q$ and $R = p?\ell; q?\ell'; R + q?\ell'; p?\ell; R$. A suitable global type is

$$G = p!r.\ell; q!r.\ell'; r?\{p.\ell; r?q.\ell'; G, q.\ell'; r?p.\ell; G\}$$

as shown by the following derivation:

$$\begin{array}{c}
\vdots \\
\hline \hline
G \vdash p[[P]] \parallel q[[Q]] \parallel r[[R]] \parallel \emptyset \\
\hline \hline
r?q.\ell'; G \vdash p[[P]] \parallel q[[Q]] \parallel r[q?\ell'; R] \parallel \langle q, \ell', r \rangle \\
\hline \hline
r?\{p.\ell; r?q.\ell'; G, q.\ell'; r?p.\ell; G\} \vdash p[[P]] \parallel q[[Q]] \parallel r[[R]] \parallel \langle p, \ell, r \rangle \cdot \langle q, \ell', r \rangle \\
\hline \hline
q!r.\ell'; r?\{p.\ell; r?q.\ell'; G, q.\ell'; r?p.\ell; G\} \vdash p[[P]] \parallel q[[Q]] \parallel r[[R]] \parallel \langle p, \ell, r \rangle \\
\hline \hline
G \vdash p[[P]] \parallel q[[Q]] \parallel r[[R]] \parallel \emptyset
\end{array}$$

Our type system enjoys the properties of Session Fidelity and Subject Reduction. Moreover, it ensures the semantic property of Lock-freedom. Since every participant can freely perform outputs, to prove this property we only have to show that all inputs can be enabled. For lack of space we only give the most interesting case in the proof of Subject Reduction.

Theorem 4.2 (Session Fidelity) *If $G \vdash \mathbb{N} \parallel \mathcal{M}$ and $G \parallel \mathcal{M} \xrightarrow{\beta} G' \parallel \mathcal{M}'$, then $\mathbb{N} \parallel \mathcal{M} \xrightarrow{\beta} \mathbb{N}' \parallel \mathcal{M}'$ and $G' \vdash \mathbb{N}' \parallel \mathcal{M}'$.*

Theorem 4.3 (Subject Reduction) *If $G \vdash \mathbb{N} \parallel \mathcal{M}$ and $\mathbb{N} \parallel \mathcal{M} \xrightarrow{\beta} \mathbb{N}' \parallel \mathcal{M}'$, then $G \parallel \mathcal{M} \xrightarrow{\beta} G' \parallel \mathcal{M}'$ and $G' \vdash \mathbb{N}' \parallel \mathcal{M}'$.*

Proof. The proof is by induction on $d = \text{depth}(G, p)$ where $p = \text{play}(\beta)$. Notice that $\mathbb{N} \parallel \mathcal{M} \xrightarrow{\beta} \mathbb{N}' \parallel \mathcal{M}'$ implies $p \in \text{Players}(\mathbb{N})$, which together with $G \vdash \mathbb{N} \parallel \mathcal{M}$ implies $p \in \text{Players}(G)$. Then $d > 0$. Moreover d is finite since G is bounded.

Let $d > 1$ and $G = r?\{q_i.\ell_i; G_i\}_{i \in I}$ with $r \neq p$. Since $G \vdash \mathbb{N} \parallel \mathcal{M}$ must be derived using Rule [IN], we get:

$\mathbb{N} \equiv r[[\Sigma_{h \in H} q_h ? \ell_h; R_h]] \parallel \mathbb{N}_0 \quad G_j \vdash r[[R_j]] \parallel \mathbb{N}_0 \parallel \mathcal{M}_j$ for all $j \in J \quad \mathbf{l}(G_j, \mathbb{M})$ for all $j \in J$
where $J = \text{rm}(\{\langle q_i, \ell_i, p \rangle\}_{i \in I}, \mathcal{M}) = \text{rm}(\{\langle q_h, \ell_h, p \rangle\}_{h \in H}, \mathcal{M})$ and $\mathcal{M} \equiv \langle q_j, \ell_j, r \rangle \cdot \mathcal{M}_j$ for all $j \in J$ and $\mathbb{M} = \{\langle q_l, \ell_l, r \rangle \mid l \in (I \cup H) \setminus J \ \& \ q_l \neq q_j \ \forall j \in J\}$. The condition $r \neq p$ ensures that the transition $\mathbb{N} \parallel \mathcal{M} \xrightarrow{\beta} \mathbb{N}' \parallel \mathcal{M}'$ does not modify the process of participant r and does not dequeue any message with receiver r from \mathcal{M} . Therefore we get $\mathbb{N}' \equiv r[[\Sigma_{h \in H} q_h ? \ell_h; R_h]] \parallel \mathbb{N}'_0$ and $\mathcal{M}' \equiv \langle q_j, \ell_j, r \rangle \cdot \mathcal{M}'_j$ for all $j \in J$. Moreover the transition can be done also if the process $\Sigma_{h \in H} q_h ? \ell_h; R_h$ is replaced by an arbitrary process and top messages with receiver r are dequeued. Therefore

$$r[[R_j]] \parallel \mathbb{N}_0 \parallel \mathcal{M}_j \xrightarrow{\beta} r[[R_j]] \parallel \mathbb{N}'_0 \parallel \mathcal{M}'_j \text{ for all } j \in J$$

It is easy to verify that $\text{depth}(G_j, p) < \text{depth}(G, p)$. Then by induction we get $G_j \parallel \mathcal{M}_j \xrightarrow{\beta} G'_j \parallel \mathcal{M}'_j$ and $G'_j \vdash r[[R_j]] \parallel \mathbb{N}'_0 \parallel \mathcal{M}'_j$ for all $j \in J$. Let $G' = r?\{q_i.\ell_i; G'_i\}_{i \in I}$ where $G'_i = G_{j_0}$ for some $j_0 \in J$ and all $i \in I \setminus J$. From $G_j \parallel \mathcal{M}_j \xrightarrow{\beta} G'_j \parallel \mathcal{M}'_j$ we get $G_j \parallel \mathcal{M} \xrightarrow{\beta} G'_j \parallel \mathcal{M}'$ for all $j \in J$. Then we derive $G \parallel \mathcal{M} \xrightarrow{\beta} G' \parallel \mathcal{M}'$ by Rule [INSIDE-IN]. The condition $\mathbf{l}(G_j, \mathbb{M})$ for all $j \in J$ ensures that \mathcal{M}' cannot contain a message $\langle q_k, \ell_k, r \rangle$ with $k \in (I \cup H) \setminus J$ and $q_k \neq q_j$ for all $j \in J$. The condition $r \neq p$ ensures that the transition $\xrightarrow{\beta}$ cannot dequeue a message with r as receiver. Hence $\text{rm}(\{\langle q_h, \ell_h, p \rangle\}_{h \in H}, \mathcal{M}') = J$. It is easy to verify that $\mathbf{l}(G_j, \mathbb{M})$ implies $\mathbf{l}(G'_j, \mathbb{M})$ for all $j \in J$. From $G'_j \vdash r[[R_j]] \parallel \mathbb{N}'_0 \parallel \mathcal{M}'_j$ for all $j \in J$ we get $\text{Players}(G'_i) \setminus \{r\} = \text{Players}(\mathbb{N}'_0)$ for all $i \in I$. Then all premises of Rule [IN] hold and we can derive $G' \vdash \mathbb{N}' \parallel \mathcal{M}'$. \square

Theorem 4.4 (Lock-freedom) *If $G \vdash \mathbb{N} \parallel \mathcal{M}$ and $p[[P]] \in \mathbb{N}$, then $\mathbb{N} \parallel \mathcal{M} \xrightarrow{\tau\beta}$ with $\text{play}(\beta) = p$ for some τ, β .*

As expected, since queues in type configurations can arbitrarily grow, our type system is undecidable.

Theorem 4.5 (Undecidability) *Typing is undecidable.*

In order to recover from this undecidability result, we can define an inductive version of typing, thus obtaining a sound algorithm. This inductive definition follows the standard pattern to deal with regular structures for global types, and it requires the same queue at the beginning and at the end of each cycle.

5 Related Work and Conclusion

We proposed flexible choice operators for an asynchronous multiparty session calculus, in order to ensure the classical session correctness properties for a larger class of protocols than is usually done. Several other proposals for relaxing the constraints of the original choice operator of [11] were already mentioned in Section 1. We now discuss some of them in more detail.

In [3], which builds on [13], we pushed this flexibility even further by allowing input choices with different senders in processes, without restrictions. The same approach was followed in [5]. However, this liberal approach turned out to be incorrect, as pointed out in [10], as it allows the following (synchronous) network to be typed, while it is not deadlock-free. Indeed, this network can reach a deadlock if p chooses its second branch, leading both s and t to choose their second branch too. Then, if r chooses its first branch, it will be unable to complete it.

$$\begin{aligned} & p[(s!a; t!a; r!d) \oplus (s!b; t!b)] \parallel r[(s?c; t?e; p?d) + (t?e; s?c)] \parallel \\ & s[(p?a; r!c) + (p?b; r!c)] \parallel t[(p?a; r!e) + (p?b; r!e)] \end{aligned}$$

In fact, this session is not race-free according to the *race-freedom* condition proposed in [10]. Note that the asynchronous session obtained by composing this network with the empty queue is not typable in our type system. Indeed, its typability would contradict Subject Reduction or Lock-freedom, since it has the derivative $r[p?d] \parallel \emptyset$ which is stuck. This session cannot be typed in [14] either, since participant r does not satisfy the required well-formedness conditions. Also the type system of [17] rejects this session, the reason being that participant r can read the same message in more than one branch. More precisely, participant r can read the message c from participant s and the message e from participant t in both branches. This control is realised by annotating projections with the set of available messages. We take advantage of queues in type configurations for a similar but less refined control, which uses the predicate ensuring that a global type is inactive for a given set of messages.

As future work, we plan to investigate three different variations of our typing. The first one would be a weakening of condition $\mathbf{I}(G, \mathbb{M})$ in Rule [IN], taking into account the order of sent messages and the expected inputs. The second one would be a strengthening of our typing in order to forbid orphan messages. The last one would be a weakening of our typing in order to allow optional participants in the branches of choices, possibly using connecting communications as in [13, 3].

Another direction that would be worth investigating is the relationship between our approach and input races for sessions based on Classical Linear Logic, see [15, 18].

To make our type system more efficient we will design two algorithms, taking inspiration from [7], one for inferring global types for networks and the other one for checking the correctness of global types for queues, allowing also cycles in which queues increase.

Acknowledgments This paper came into being thanks to Ross Horne, who pointed out to us the example reported in the Conclusion. We are indebted to him for many interesting discussions on this subject and for suggestions on a previous version of this paper. We also thank the anonymous referees for helpful comments.

References

- [1] Marco Carbone, Nobuko Yoshida & Kohei Honda (2009): *Asynchronous Session Types: Exceptions and Multiparty Interactions*. In Marco Bernardo, Luca Padovani & Gianluigi Zavattaro, editors: *Formal Methods for Web Services*, LNCS 5569, Springer, pp. 187–212, doi:10.1007/978-3-642-01918-0_5.
- [2] Giuseppe Castagna, Mariangiola Dezani-Ciancaglini & Luca Padovani (2012): *On Global Types and Multiparty Sessions*. *Logical Methods in Computer Science* 8, pp. 1–45, doi:10.2168/LMCS-8(1:24)2012.
- [3] Iliaria Castellani, Mariangiola Dezani-Ciancaglini & Paola Giannini (2019): *Reversible Sessions with Flexible Choices*. *Acta Informatica* 56(7), pp. 553–583, doi:10.1007/s00236-019-00332-y.
- [4] Iliaria Castellani, Mariangiola Dezani-Ciancaglini & Paola Giannini (2021): *Global types and event structure semantics for asynchronous multiparty sessions*. *CoRR* abs/2102.00865, doi:10.48550/arXiv.2102.00865.
- [5] Iliaria Castellani, Mariangiola Dezani-Ciancaglini, Paola Giannini & Ross Horne (2020): *Global Types with Internal Delegation*. *Theoretical Computer Science* 807, pp. 128–153, doi:10.1016/j.tcs.2019.09.027.
- [6] Bruno Courcelle (1983): *Fundamental Properties of Infinite Trees*. *Theoretical Computer Science* 25, pp. 95–169, doi:10.1016/0304-3975(83)90059-2.
- [7] Francesco Dagnino, Paola Giannini & Mariangiola Dezani-Ciancaglini (2021): *Deconfined Global Types for Asynchronous Sessions*. *CoRR* abs/2111.11984, doi:10.48550/arXiv.2111.11984.
- [8] Romain Demangeon & Kohei Honda (2012): *Nested Protocols in Session Types*. In Maciej Koutny & Irek Ulidowski, editors: *CONCUR*, LNCS 7454, Springer, pp. 272–286, doi:10.1007/978-3-642-32940-1_20.
- [9] Pierre-Malo Deniérou & Nobuko Yoshida (2012): *Multiparty Session Types Meet Communicating Automata*. In Helmut Seidl, editor: *ESOP*, LNCS 7211, Springer, pp. 194–213, doi:10.1007/978-3-642-28869-2_10.
- [10] Rob van Glabbeek, Peter Höfner & Ross Horne (2021): *Assuming Just Enough Fairness to make Session Types Complete for Lock-freedom*. In Leonid Libkin, editor: *LICS*, ACM Press, pp. 1–13, doi:10.1109/LICS52264.2021.9470531.
- [11] Kohei Honda, Nobuko Yoshida & Marco Carbone (2008): *Multiparty Asynchronous Session Types*. In George C. Necula & Philip Wadler, editors: *POPL*, ACM Press, pp. 273–284, doi:10.1145/1328438.1328472.
- [12] Kohei Honda, Nobuko Yoshida & Marco Carbone (2016): *Multiparty Asynchronous Session Types*. *Journal of ACM* 63(1), pp. 9:1–9:67, doi:10.1145/2827695.
- [13] Raymond Hu & Nobuko Yoshida (2017): *Explicit Connection Actions in Multiparty Session Types*. In Marieke Huisman & Julia Rubin, editors: *FASE*, LNCS 10202, Springer, pp. 116–133, doi:10.1007/978-3-662-54494-5_7.
- [14] Sung-Shik Jongmans & Nobuko Yoshida (2020): *Exploring Type-Level Bisimilarity towards More Expressive Multiparty Session Types*. In Peter Müller, editor: *ESOP*, LNCS 12075, Springer, pp. 251–279, doi:10.1007/978-3-030-44914-8_10.
- [15] Wen Kokke, J. Garrett Morris & Philip Wadler (2020): *Towards Races in Linear Logic*. *Logical Methods in Computer Science* 16(4), doi:10.23638/LMCS-16(4:15)2020.
- [16] Julien Lange, Emilio Tuosto & Nobuko Yoshida (2015): *From Communicating Machines to Graphical Choreographies*. In Sriram K. Rajamani & David Walker, editors: *POPL*, ACM Press, pp. 221–232, doi:10.1145/2676726.2676964.
- [17] Rupak Majumdar, Madhavan Mukund, Felix Stutz & Damien Zufferey (2021): *Generalising Projection in Asynchronous Multiparty Session Types*. In Serge Haddad & Daniele Varacca, editors: *CONCUR*, LIPIcs 203, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 35:1–35:24, doi:10.4230/LIPIcs.CONCUR.2021.35.

- [18] Zesen Qian, G. A. Kavvos & Lars Birkedal (2021): *Client-server sessions in linear logic*. *Proc. ACM Program. Lang.* 5(ICFP), pp. 1–31, doi:10.1145/3473567.