



Reversible sessions with flexible choices

Ilaria Castellani¹ · Mariangiola Dezani-Ciancaglini² · Paola Giannini³ 

Received: 16 April 2018 / Accepted: 8 February 2019 / Published online: 1 March 2019
© Springer-Verlag GmbH Germany, part of Springer Nature 2019

Abstract

We propose a calculus for concurrent reversible multiparty sessions, equipped with a flexible choice operator allowing for different sets of participants in each branch. This operator is inspired by the notion of *connecting action* recently introduced by Hu and Yoshida to describe protocols with optional participants. We argue that this choice operator allows for a natural description of typical communication protocols. Our calculus also supports a compact representation of the history of processes and types, which facilitates the definition of rollback. Moreover, it implements a fine-tuned strategy for backward computation. We present a session type system for the calculus and show that it enforces the expected properties of session fidelity, forward progress and backward progress.

1 Introduction

Session types are a simple but expressive type formalism that specifies the structure of interactions. Traditionally, session types have been used to ensure safety properties of interactions, such as absence of communication errors, deadlock freedom and race freedom.

Reversibility is a means to improve system flexibility and reliability. Reversing a computation may be defined as the act of undoing some suffix of the computation, in order to return to a previously visited state. In a nondeterministic computation, this possibility may be used

Mariangiola Dezani-Ciancaglini: Partially supported by EU H2020-644235 Rephrase project, EU H2020-644298 HyVar project, IC1402 ARVI and Ateneo/CSP project RunVar.

Paola Giannini: This original research has the financial support of the Università del Piemonte Orientale. The authors acknowledge a partial support of COST Action IC1405 on Reversible Computation—extending horizons of computing.

✉ Paola Giannini
paola.giannini@uniupo.it

Ilaria Castellani
ilaria.castellani@inria.fr

Mariangiola Dezani-Ciancaglini
dezani@di.unito.it

¹ INRIA, Université Côte d’Azur, 2004 Route des Lucioles, 06902 Sophia Antipolis, France

² Dipartimento di Informatica, Università di Torino, Corso Svizzera 185, 10131 Turin, Italy

³ DiSIT, Università del Piemonte Orientale, Via Teresa Michel 11, 15121 Alessandria, Italy

to return to a previous branching point, in case the chosen branch has led to an unsuccessful state.

In the setting of structured communications, reversibility has been first studied for contracts [2,3] and transactions [11,12,21]. More recently, it has started to be investigated also for session calculi, both binary [24,34] and multiparty [14,26,28,35] (see Sect. 6 for more discussion on related work).

When reversing a structured interaction, one has to preserve the consistency of the global state: if one partner triggers a rollback, then all its communicating partners should roll back accordingly. Session types turn out to be very useful here, since they specify both the functionality of communications (sender, receiver and type of message), and the order in which they should occur.

We present a calculus for concurrent reversible multiparty sessions, whose distinctive feature is a flexible choice operator, allowing for different sets of participants in its branches. The only participant which is required to occur in all branches is the one which solves the choice, henceforth called the *choice leader*.

Our choice operator is inspired by the notion of *connecting action* recently introduced by Hu and Yoshida to describe protocols with optional participants [20]. The intuition behind connecting actions is that in some parts of the protocol, delimited by a choice construct, some participants are required to take part in the interaction while some others may be optional. Connecting actions are used to invite optional participants to join the interaction along some branches of the choice. For instance, in a PC meeting, it could happen that in case of divergent views about a paper, the PC chair launches a discussion among the concerned PC members, but also invites some additional PC members to join the discussion. These additional members are optional in the sense that they are not required to discuss on that particular paper, but they may be invited to do so in some cases.

Here we shall make a more permissive use of connecting actions than in [20].

The differences between our connecting actions and those in [20] are discussed in Sect. 6.

Compared to the standard choice operator of multiparty session calculi, our flexible choice allows for a more natural description of typical communication protocols, such as the one mentioned above. Another example is the vacation protocol discussed below, where Alice has to decide between two destinations, and depending on her decision she will wish to contact either an airline or a railway company but not both. This will be modelled by a choice between two connecting actions, one with the airline and one with the railway company.

Another notable feature of our calculus is that it gives a compact representation of the *history* of processes and types, which facilitates the definition of rollback. It also implements a fine-tuned strategy for backward computation, which is geared towards achieving compliance. In essence, a backward move can only return to a past choice point, and it can only be triggered by the leader of that choice; moreover, the past choice state is restored without the already explored branch, thus forcing the choice leader to engage into a different branch.

The main contributions of our paper may be summarised as follows:

- the introduction of a flexible choice operator based on *connecting actions* in a reversible multiparty session framework;
- a fine-tuned strategy for rollback to checkpointed choices, whereby rollback can only be triggered by choice leaders in predefined states of the computation, leading back to the choice state stripped off the unsuccessful path.

This work builds on our previous papers [8,14]. As regards the treatment of reversibility in multiparty sessions, the general principles and the use of checkpoints for return points are

taken from [14], while the formalisation of histories and the specific rollback mechanism are borrowed from [8]. The formalisation of histories is pushed a little further here than in [8], including notions of causality and conflict for communication occurrences. On the other hand, the use of connecting actions is new with respect to [8, 14], and so is the study of their interplay with reversible computations. This interplay is not entirely trivial, as it requires revisiting the definition of projection of a global type onto its participants. A relevant new notion to this purpose is that of *affecting choice* for a participant: intuitively, a choice specified by a global type affects a participant if the whole choice is needed to determine the projection on that participant. By definition, a choice affects all its required participants while it does not affect optional participants which appear in only one branch. If we ignore connecting actions, the expressive power of our calculus is comprised between that of [14] and that of [8]. Indeed, [8] allows for both parallel and sequential composition within processes and types, yielding a powerful but somewhat complex calculus. Here we decided to stick to a more standard syntax, in order to be able to focus on the main topic of the paper, namely backward computation in the presence of connecting actions.

Vacation protocol example To illustrate our approach, we present a simple protocol involving five parties, Alice, Bob, Carol, the airline company Alitalia and the railway company Trenitalia. This protocol, henceforth called the *vacation protocol*, will serve as our running example throughout the paper.

For Easter holidays, Alice receives two independent invitations from Bob and Carol. Bob proposes to Alice to visit him in Paris, while Carol offers to host her in her seaside house in Amalfi. Since Alice lives in Rome, depending on her decision she will need to book a flight to go to Paris or a train to reach Amalfi. To do so, Alice will either contact Alitalia or Trenitalia to buy a ticket, and then inform both Bob and Carol of her decision. In case of a plane or train strike, Alice is allowed to change her mind. Letting a, b, c, f, t denote respectively Alice, Bob, Carol, Alitalia and Trenitalia, a global type describing this communication protocol is:

$$b \xrightarrow{iP} a; c \xrightarrow{iA} a; (a \xleftrightarrow{ik} f; a \xrightarrow{yes} b; a \xrightarrow{no} c \text{ } \boxplus_C \text{ } a \xleftrightarrow{ik} t; a \xrightarrow{no} b; a \xrightarrow{yes} c)$$

where $\text{ } \boxplus_C \text{ }$ is a choice with checkpoint label C , meaning that the choice leader (i.e. the participant who makes the choice, Alice in this case) can rollback to take a different branch of the choice. The communications between Alice and Alitalia/Trenitalia are *connecting communications* (represented by messages of the form $\xleftrightarrow{\lambda}$). This means that there is no assurance for Alitalia and Trenitalia that they will receive a message from Alice. With standard global types, Alitalia and Trenitalia should receive a message from Alice in both branches of the choice, and this is clearly not realistic. A different way out would be to use three global types instead of one, respectively with sets of participants $\{Alice, Bob, Carol\}$, $\{Alice, Alitalia\}$ and $\{Alice, Trenitalia\}$. The drawback of the latter solution is that, in the presence of more than one global type, a simple session type system is not sufficient to ensure progress, and one needs to recourse to more refined type systems [29].

Thanks to connecting actions, the vacation protocol may be described by a single global type, without the constraints of standard global types.

Outline The rest of the paper is organised as follows. In Sect. 2 we define the syntax and operational semantics—both forward and backward—of our calculus. In Sect. 3 we introduce our syntax for global types and session types and we establish well-formedness conditions for global types. In Sect. 4 we present our type system and prove some preliminary properties. In Sect. 5 we prove the soundness of our type system, namely that it ensures the expected

semantic properties of session fidelity and forward and backward progress. We conclude in Sect. 6 with some discussion on related and future work. The Appendix proves the correctness of a possible implementation of backward reduction.

2 Calculus

We assume the following base sets: *simple messages*, ranged over by λ, λ', \dots and forming the set Msg ; *connecting messages*, ranged over by $\overset{\lambda}{\leftrightarrow}, \overset{\lambda'}{\leftrightarrow}, \dots$ and forming the set CMsg ; *checkpoint labels*, ranged over by C, C' and forming the set ChLa ; and *session participants*, ranged over by p, q, r and forming the set Part . We use Δ to range over both simple messages and connecting messages.

We use Δ to range over sets of checkpoint labels, and $\overline{\gamma}$ to stand for either the empty set or a singleton consisting of an overlined checkpoint label (curly brackets will be omitted around singletons):

$$\Delta ::= \emptyset \mid \Delta, C \qquad \overline{\gamma} ::= \emptyset \mid \overline{C}$$

Sets Δ and $\overline{\gamma}$ are associated with choices in processes. More precisely, a set Δ is associated with an external choice and said to be *passive*, while a set $\overline{\gamma}$ is associated with an internal choice and said to be *active*. Intuitively, an overlined label \overline{C} is the handle for a backward move: a participant who crossed an internal choice (henceforth called the *choice leader*) with checkpoint label \overline{C} , and then proceeded in the computation, may decide to return to that choice whenever she has the ability to send a message. Within a network, this backward move of the choice leader will have to be matched by backward moves of all the participants who did some action after crossing the matching external choice, whose checkpoint set contains label C . The asymmetry between Δ and $\overline{\gamma}$ is justified by the fact that there is only one choice leader who can send messages to various participants.

Let $\pi \in \{p? \Delta, p! \Delta \mid p \in \text{Part}, \Delta \in \text{Msg} \cup \text{CMsg}\}$ denote an *atomic action*, namely a simple input/output action or an input/output action establishing a connection. As in [20], connecting inputs may be dangling forever, whereas simple inputs will eventually take place. This gives a natural freedom in the definition of communication protocols as illustrated in the examples of the introduction. An atomic action can bear a hat, in which case it represents an already *executed action*. We use $\hat{\pi}$ to stand for either π or $\hat{\pi}$. External choices and internal choices are denoted by \sum and \oplus , respectively.

Definition 1 (*Processes*) Processes are defined by:

$$P ::= \Delta \sum_{i \in I} \widehat{p_i ? \Lambda_i}; P_i \mid \overline{\gamma} \oplus_{i \in I} \widehat{p_i ! \Lambda_i}; P_i \mid \mu X. P \mid X \mid \text{end}$$

where in both kinds of choice the pairs (p_i, Λ_i) are assumed to be all distinct. Moreover, for external choices we also assume that the Λ_i 's are either all simple or all connecting messages. This condition allows us to distinguish between simple and connecting external choices according to the kind of their inputs. This is essential for requiring that at least one of the inputs in a simple external choice will eventually be matched by a message. Instead all inputs in a connecting external choice can wait forever. Processes without hats are called *user processes* and the others are *runtime processes*.

We will omit empty sets of checkpoint labels, choice symbols in one-branch choices, and trailing end processes.

External and internal choices are assumed to be associative, commutative, and non-empty (except when combined with binary choices). A prefixed process may be either an *input process* or an *output process*. We require recursion to be guarded. Processes are treated equi-recursively, i.e. they are identified with their generated tree [33]. In other words, we consider processes up to the standard structural equivalence of processes.

The typing rules of Sect. 4 will ensure that in a choice, at most one of the first atomic actions bears a hat. This condition expresses the fact that executing a choice amounts to executing one of its branches.

In a full-fledged calculus, messages would carry values, namely they would be of the form $\Lambda(v)$. Here, for simplicity we consider only pure messages.

Networks are parallel compositions of pairs $p \llbracket P \rrbracket$, where participant p has behaviour P .

Definition 2 (*Networks*) Networks are defined by: $\mathbb{N} ::= p \llbracket P \rrbracket \mid \mathbb{N} \parallel \mathbb{N}$

The operator \parallel is associative and commutative, with neutral element $p \llbracket \text{end} \rrbracket$ for each p . These laws, together with the structural equivalence of processes, give the structural equivalence of networks.

Example 1 (*Networks*) A network for the vacation example of Sect. 1 is as follows:

$$\mathbb{N} = a \llbracket P^a \rrbracket \parallel b \llbracket P^b \rrbracket \parallel c \llbracket P^c \rrbracket \parallel f \llbracket P^f \rrbracket \parallel t \llbracket P^t \rrbracket$$

where

$$P^a = b?iP; c?iA; (P_1^a \overline{c} \oplus P_2^a) \text{ with } P_1^a = f! \overset{tk}{\leftarrow}; b!yes; c!no \text{ and } P_2^a = t! \overset{tk}{\leftarrow}; b!no; c!yes$$

$$P^b = a!iP; (a?yes \ c + \ a?no) \quad P^c = a!iA; (a?no \ c + \ a?yes) \quad P^f = P^t = c a? \overset{tk}{\leftarrow}$$

The operational semantics is given by two LTSs, one for processes and one for networks. In the LTS for processes, *forward transitions* have the form $P \xrightarrow{\pi} P'$ and *backward transitions* have the form $P \xrightarrow{\overline{C}} P'$ or $P \xrightarrow{C} P'$. We define $P \downarrow_{out}$ if $P \xrightarrow{p!\Lambda} P'$ for some p, Λ, P' . In the LTS for networks, forward and backward transitions have respectively the form $\mathbb{N} \xrightarrow{p!\Lambda} \mathbb{N}'$ and $\mathbb{N} \xrightarrow{C} \mathbb{N}'$.

The LTSs for processes and networks are given in Fig. 1. Rules [EXTCH] and [INTCH] allow an action to be extracted from one of the summands, as usual, but instead of discarding the other summands they record the fact that the choice has been crossed by marking the executed action with a hat. With this technique, inspired by [6] and already used for reversible computations in [8,26], all the dynamic operators are turned into static operators and nothing is lost of the original user process. Notice that when $I = \{j\}$ these rules become $\pi_j; P_j \xrightarrow{\pi_j} \widehat{\pi_j}; P_j$. Rule [BACKA] is the main backward rule: it applies to a past internal choice in which one branch has been partially executed, and it allows the process to roll back to the original choice where the executed branch is removed.¹ For this to be possible, the choice must have at least one alternative branch P_i and an overlined label \overline{C} , which will label the back transition. This is essential to ensure (by means of typing) that the choice leader will be the only participant habilitated to trigger a rollback to this choice. The condition $P \downarrow_{out}$

¹ We could also leave the executed branch, just erasing its hats. This change would preserve all the properties of the calculus. We prefer the current formulation, since it avoids the possibility of going back and fourth several times along the same branch (which could yield a livelock).

$$\begin{array}{c}
 \Delta \sum_{i \in I} \pi_i; P_i \xrightarrow{\pi_j} \Delta (\sum_{i \in I \setminus \{j\}} \pi_i; P_i + \widehat{\pi}_j; P_j) \quad j \in I \quad [\text{EXTCH}] \\
 \overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i \xrightarrow{\pi_j} \overline{\gamma} (\bigoplus_{i \in I \setminus \{j\}} \pi_i; P_i \oplus \widehat{\pi}_j; P_j) \quad j \in I \quad [\text{INTCH}] \\
 \frac{P \downarrow_{out} \quad I \neq \emptyset}{\overline{\mathcal{C}} (\bigoplus_{i \in I} \pi_i; P_i \oplus \widehat{\pi}; P) \xrightarrow{\overline{\mathcal{C}}} \overline{\mathcal{C}} \bigoplus_{i \in I} \pi_i; P_i} [\text{BACKA}] \\
 \frac{C \in \Delta}{\Delta \sum_{i \in I} \widehat{\pi}_i; P_i \xrightarrow{\mathcal{C}} \wr \Delta \sum_{i \in I} \widehat{\pi}_i; P_i \wr} [\text{BACKP}] \\
 \frac{P \xrightarrow{\pi} P'}{\mathcal{E}[P] \xrightarrow{\pi} \mathcal{E}[P']} [\text{CTFAT}] \quad \frac{P \xrightarrow{\overline{\mathcal{C}}} P' \quad \mathcal{E} \text{ ok for } \overline{\mathcal{C}}}{\mathcal{E}[P] \xrightarrow{\overline{\mathcal{C}}} \mathcal{E}[P']} [\text{CTBA}] \quad \frac{P \xrightarrow{\mathcal{C}} P' \quad \mathcal{E} \text{ ok for } C}{\mathcal{E}[P] \xrightarrow{\mathcal{C}} \mathcal{E}[P']} [\text{CTBP}] \\
 \frac{P \xrightarrow{q!A} P' \quad Q \xrightarrow{p?A} Q'}{\mathfrak{p}[\![P]\!] \parallel \mathfrak{q}[\![Q]\!] \parallel \mathbb{N} \xrightarrow{p!Aq} \mathfrak{p}[\![P']\!] \parallel \mathfrak{q}[\![Q']\!] \parallel \mathbb{N}} [\text{COM}] \\
 \frac{P \xrightarrow{\overline{\mathcal{C}}} P' \quad P_h \xrightarrow{\mathcal{C}} P'_h \quad h \in H \quad P_k \xrightarrow{\mathcal{C}} P'_k \quad k \in K}{\mathfrak{p}[\![P]\!] \parallel \Pi_{h \in H} \mathfrak{p}_h[\![P_h]\!] \parallel \Pi_{k \in K} \mathfrak{p}_k[\![P_k]\!] \xrightarrow{\mathcal{C}} \mathfrak{p}[\![P']\!] \parallel \Pi_{h \in H} \mathfrak{p}_h[\![P'_h]\!] \parallel \Pi_{k \in K} \mathfrak{p}_k[\![P'_k]\!]} [\text{BACK}]
 \end{array}$$

Fig. 1 LTS for processes and networks

means that in order to trigger a rollback, the process P must be “in lead”, namely able to do an output. In this way, rolling back acts as an alternative to one of its possible outputs.

Rule [BACKP] is needed to allow the remaining participants to roll back. The mapping $\wr \wr$ erases hats from processes, yielding user processes, i.e.

$$\wr \Delta \sum_{i \in I} \widehat{\pi}_i; P_i \wr = \Delta \left(\sum_{i \in I} \pi_i; \wr P_i \wr \right) \quad \wr \overline{\gamma} \bigoplus_{i \in I} \widehat{\pi}_i; P_i \wr = \overline{\gamma} \left(\bigoplus_{i \in I} \pi_i; \wr P_i \wr \right)$$

and $\wr \wr$ acts homomorphically otherwise. Note that the rollback rules can only be applied to processes that are not user processes: in particular, one branch can be erased only if at least one of its actions has been executed.

Evaluation contexts are as expected.

Definition 3 (Evaluation contexts) Evaluation contexts \mathcal{E} are defined by:

$$\mathcal{E} ::= \Delta \left(\sum_{h \in H} \mathfrak{p}_h ? A_h; P_h + \widehat{\mathfrak{p}} ? A; \mathcal{E} \right) \mid \overline{\gamma} \left(\bigoplus_{h \in H} \mathfrak{p}_h ! A_h; P_h \oplus \widehat{\mathfrak{p}} ! A; \mathcal{E} \right) \mid []$$

An evaluation context \mathcal{E} is ok for C ($\overline{\mathcal{C}}$) if $C \notin \Delta$ ($\overline{\mathcal{C}} \neq \overline{\gamma}$) whenever \mathcal{E} has a sub-context of the shape $\Delta (\sum_{h \in H} P_h + \widehat{\pi}; \mathcal{E}')$ ($\overline{\gamma} (\bigoplus_{h \in H} P_h \oplus \widehat{\pi}; \mathcal{E}')$). We use this condition in rules [CTBA] and [CTBP] to assure that all participants involved in a recursion go back to the same checkpoint, namely to the outermost one, as in [28]. This is needed to assure subject reduction, see Example 6.

Rule [COM] is standard and deals with both simple and connecting messages. We write $P \not\stackrel{C}{\hookrightarrow}$ if Rule [CTBP] (with label C) cannot be applied to P . This means that C can only occur in user processes within P . In other words, P does not contain an executed $\Delta+$ with $C \in \Delta$. In a well-typed network, Rule [BACK] will make participant p roll back to an internal choice and moreover, all participants that can roll back to corresponding external choices will do so in the same step. This will be the basis for our soundness result in Sect. 5. Notice that the only requirements of rule [BACK] concern process P , since it is easy to verify that either $Q \stackrel{C}{\hookrightarrow} Q'$ or $Q \not\stackrel{C}{\hookrightarrow}$ for any other process Q . Note that a direct implementation of rule [BACK] would be unrealistic. We discuss a possible asynchronous implementation of this rule at the end of the section.

When the labels of transitions are not relevant, we write them simply as \longrightarrow and \curvearrowright . In this case, we use \longrightarrow^* to denote the reflexive and transitive closure of \longrightarrow and \curvearrowright^* to denote the reflexive and transitive closure of $\longrightarrow \cup \curvearrowright$.

Example 2 (Reduction of networks) We describe the evolution of the network N of Example 1. At each step we only show the participants that are modified by the reduction.

$$\begin{aligned}
 N &\xrightarrow{b!Pa} a\llbracket \widehat{b?iP}; c?iA; (P_1^a \bar{c} \oplus P_2^a) \rrbracket \parallel b\llbracket \widehat{aliP}; (a?yes_C + a?no) \rrbracket \parallel \dots \\
 &\xrightarrow{c!Aa} a\llbracket \widehat{b?iP}; c?iA; (P_1^a \bar{c} \oplus P_2^a) \rrbracket \parallel c\llbracket \widehat{a!iA}; (a?no_C + a?yes) \rrbracket \parallel \dots \\
 &\xrightarrow{a \stackrel{fk}{\leftrightarrow} f} a\llbracket \widehat{b?iP}; c?iA; (f! \stackrel{fk}{\leftrightarrow}; b!yes; c!no \bar{c} \oplus P_2^a) \rrbracket \parallel f\llbracket \widehat{cA?} \stackrel{fk}{\leftrightarrow} \rrbracket \parallel \dots \\
 &\xrightarrow{a?yes b} a\llbracket \widehat{b?iP}; c?iA; (f! \stackrel{fk}{\leftrightarrow}; b!yes; c!no \bar{c} \oplus P_2^a) \rrbracket \parallel b\llbracket \widehat{a!iP}; (a?yes_C + a?no) \rrbracket \parallel \dots \\
 &\xrightarrow{c} a\llbracket \widehat{b?iP}; c?iA; \bar{c}P_2^a \rrbracket \parallel b\llbracket \widehat{a!iP}; (a?yes_C + a?no) \rrbracket \parallel f\llbracket \widehat{cA?} \stackrel{fk}{\leftrightarrow} \rrbracket \parallel \dots
 \end{aligned}$$

The last rollback is triggered by a [BACKA] move of Alice synchronised with [BACKP] moves of Bob and Alitalia.

In the rest of this section we discuss a possible implementation of rule [BACK], which decomposes the simultaneous rollback of a set of participants into a sequence of backward moves, one for each participant in the set. The first move of the sequence is always that of the choice leader.

We introduce states to define systems, which are pairs of networks and states.

States are defined by $\sigma := \langle c, \mathbb{P}_1, \mathbb{P}_2 \rangle$, where c denotes either a checkpoint label C or ‘-’, and $\mathbb{P}_1, \mathbb{P}_2$ form a partition on the set of network participants. In a state of the form $\langle -, \mathbb{P}_1, \mathbb{P}_2 \rangle$, the set \mathbb{P}_2 is always \emptyset . A state is *ready* if it is of the shape $\langle -, \mathbb{P}, \emptyset \rangle$. A communication can only take place in a ready state, yielding a ready state again. Similarly, the starting move of a rollback can only take place in a ready state, but it gives rise to a state of the form $\langle C, \mathbb{P}_1, \mathbb{P}_2 \rangle$. In a state $\langle C, \mathbb{P}_1, \mathbb{P}_2 \rangle$, a rollback is underway: the participants in \mathbb{P}_1 have already contributed to the rollback, while the participants in \mathbb{P}_2 still need to respond to the rollback “call” from the choice leader.

Systems have the form $N \bowtie \langle c, \mathbb{P}_1, \mathbb{P}_2 \rangle$, where $\mathbb{P}_1 \cup \mathbb{P}_2$ is the participant set of N . A system $N \bowtie \sigma$ is *ready* if the state σ is ready. Figure 2 gives the LTS for systems. We use the labelled double arrows $\overset{\alpha}{\rightrightarrows}$ and the labelled curly arrows $\overset{sc}{\curvearrowright}, \overset{yc}{\curvearrowright}, \overset{nc}{\curvearrowright}, \overset{ec}{\curvearrowright}$ for communication and backward moves of systems, respectively. In this way the labelled arrows for networks and systems are disjoint.

The reverse rule [BACK] of Fig. 1 splits into four rules for systems:




Fig. 2 LTS for systems

- rule [BACKS] allows the choice leader p to start a rollback: the new state contains the checkpoint label and all participants but p are required to roll back, if possible;
- rule [BACKY] allows participant p to roll back to the checkpoint label memorised in the state: the new state records the rollback of p;
- rule [BACKN] modifies the state recording that participant p cannot roll back with the current checkpoint label;
- rule [BACKE] ends the rollback by erasing from the state the current checkpoint label when all participants have become aware of the rollback.

The participants different from the choice leader are split into those that may roll back and those that may not.

The first ones reduce by rule [BACKY] and the second ones reduce by rule [BACKN].

In “Appendix A” we prove the equivalence between the LTS of Fig. 1 and the one of Fig. 2.

Example 3 (Reduction of systems) We show now how the rules of Fig. 2 may be used to simulate the execution of the network N in Example 2. First observe that, since the initial state is $\sigma = \langle -, \{a, b, c, f, t\}, \emptyset \rangle$, the first four communications of the system can be derived by rule [COMS], leaving the state σ unchanged and yielding the same network as in Example 2 (fourth line), namely $N' = N_1 \parallel N_2$ where:

$$\begin{aligned}
 N_1 &= a \llbracket \widehat{b?iP}; \widehat{c?iA}; \widehat{(f! \overset{tk}{\leftrightarrow}); b!yes}; c!no \bar{c} \oplus P_2^a) \rrbracket \parallel b \llbracket \widehat{a!iP}; \widehat{(a?yes}_c + a?no) \rrbracket \\
 N_2 &= c \llbracket \widehat{a!iA}; \widehat{(a?no}_c + a?yes) \rrbracket \parallel f \llbracket \widehat{c a? \overset{tk}{\leftrightarrow}} \rrbracket \parallel t \llbracket \widehat{c a? \overset{tk}{\leftrightarrow}} \rrbracket
 \end{aligned}$$

Then, an application of rule [BACKS] starts the rollback and a possible evolution of the system $N' \checkmark \sigma$ is the following:

$$\begin{aligned}
 N' \checkmark \langle -, \{a, b, c, f, t\}, \emptyset \rangle & \stackrel{sc}{\hookrightarrow} a[\widehat{b?iP}; \widehat{c?iA}; \overline{c}P_2^a] \parallel \dots \checkmark \langle C, \{a\}, \{b, c, f, t\} \rangle & [\text{BACKS}] \\
 & \stackrel{yc}{\hookrightarrow} b[\widehat{a!iP}; (a?yes_{c+} a?no)] \parallel \dots \checkmark \langle C, \{a, b\}, \{c, f, t\} \rangle & [\text{BACKY}] \\
 & \stackrel{nc}{\hookrightarrow} c[\widehat{a!iA}; (a?no_{c+} a?yes)] \parallel \dots \checkmark \langle C, \{a, b, c\}, \{f, t\} \rangle & [\text{BACKN}] \\
 & \stackrel{yc}{\hookrightarrow} f[\widehat{c a?} \xleftrightarrow{tk}] \parallel \dots \checkmark \langle C, \{a, b, c, f\}, \{t\} \rangle & [\text{BACKY}] \\
 & \stackrel{nc}{\hookrightarrow} t[\widehat{c a?} \xleftrightarrow{tk}] \parallel \dots \checkmark \langle C, \{a, b, c, f, t\}, \emptyset \rangle & [\text{BACKN}] \\
 & \stackrel{ec}{\hookrightarrow} \dots \checkmark \langle -, \{a, b, c, f, t\}, \emptyset \rangle & [\text{BACKE}]
 \end{aligned}$$

The application of rule [BACKE] closes the rollback and brings the system back to the ready state σ . Note that the [BACKN] and [BACKY] moves can be performed in any order as long as they follow the [BACKS] move by Alice. The network N is typable in the system of Sect. 4 by the global type given in the Introduction. By the Subject Reduction Theorem (Theorem 1) also the network N' is typable. This example shows that the four back rules of Fig. 2 are all used for reducing typable networks.

3 Global types and session types

According to [18,19], a *multiparty session* is a series of communications among participants [20], which follows a predefined protocol specified by a *global type*. Global types are built from choices among communications with the same sender, possibly using recursion. The choices are required to be non-ambiguous, i.e. to have either different receivers or different messages. As in processes, choices in global types have *checkpoint labels* which mark them as return points for rollbacks, and the executed part is highlighted with hats. For consistency, when a communication has a hat, then also all the communications that cause it (see Definition 7) must have a hat. This condition is expressed by Definition 8.

The inputs/outputs of each participant are determined by her session type, which is obtained by projecting the global type of the whole conversation. To define session types we start with the syntax of session pre-types (Definition 9) and then we single out session types (Definition 11) using the projection of global types (Fig. 5). The crucial point in the definition of projection is to determine when a participant which is not the choice leader of a choice may be realised by an external choice of processes. To this aim we define a meet operation on session pre-types (Definition 10).

We use γ to denote either the empty set or a singleton made of a checkpoint label:

$$\gamma ::= \emptyset \mid C$$

Sets γ will be associated with choices in global types.

Let $\alpha^p \in \{p \xrightarrow{\Lambda} q \mid q \in \text{Part}, \Lambda \in \text{Msg} \cup \text{CMsg}\}$ denote an *atomic communication* with sender p . The communication $p \xrightarrow{\lambda} q$ is *simple*, while the communication $p \xrightarrow{\hat{\lambda}} q$ is *connecting*. An atomic communication can bear a hat, in notation $\widehat{\alpha}^p$, in which case it represents a past or *executed* communication. The symbol $\widetilde{\alpha}^p$ stands for either α^p or $\widehat{\alpha}^p$. We write $\widetilde{\alpha}$ instead of $\widetilde{\alpha}^p$ when the sender is not relevant. Given a communication $\alpha = p \xrightarrow{\Lambda} q$, we define the sender and receiver of α to be $\text{sender}(\alpha) = p$ and $\text{rcv}(\alpha) = q$. Moreover, we denote by $\text{part}(\alpha) = \{\text{rcv}(\alpha), \text{sender}(\alpha)\}$ the participants of the communication α .

A global type K specifies an interaction that is still to start. A general global type G specifies a partially executed interaction, whose parts that have been discarded in choices or remain to be executed are specified by subterms K . By $|I|$ we denote the cardinality of the set I .

$$\begin{aligned}
 \text{part}(\gamma \sqsupset_{i \in I} \widehat{\alpha}_i; \mathbf{G}_i) &= \bigcup_{i \in I} (\text{part}(\alpha_i) \cup \text{part}(\mathbf{G}_i)) \\
 \text{part}(\mu\mathbf{t}.\mathbf{K}) &= \text{part}(\mathbf{K}) \\
 \text{part}(\mathbf{t}) &= \text{part}(\text{End}) = \emptyset
 \end{aligned}$$

Fig. 3 Participants of global types

Definition 4 (*Global types*) Global types \mathbf{G} are defined by:

$$\begin{aligned}
 \mathbf{K} &::= \gamma \alpha; \mathbf{K} \mid \gamma \boxplus_{j \in J} \alpha_j^{\mathbf{P}}; \mathbf{K}_j \mid \mu\mathbf{t}.\mathbf{K} \mid \mathbf{t} \mid \text{End} \\
 \mathbf{G} &::= \mathbf{K} \mid \gamma \widetilde{\alpha}; \mathbf{G} \mid \gamma (\boxplus_{i \in I} \alpha_i^{\mathbf{P}}; \mathbf{K}_i \boxplus \widehat{\alpha}^{\mathbf{P}}; \mathbf{G})
 \end{aligned}$$

where $|J| > 1, |I| > 0$ and all atomic communications in a choice are different. I.e. $\alpha_h^{\mathbf{P}} \neq \alpha_k^{\mathbf{P}}$ for all $h \neq k \in J$ in a choice $\gamma \boxplus_{j \in J} \alpha_j^{\mathbf{P}}; \mathbf{K}_j$, and $\alpha_h^{\mathbf{P}} \neq \alpha_k^{\mathbf{P}}$ for all $h \neq k \in I$ and $\alpha_h^{\mathbf{P}} \neq \alpha^{\mathbf{P}}$ for all $h \in I$ in a partially executed choice $\gamma (\boxplus_{i \in I} \alpha_i^{\mathbf{P}}; \mathbf{K}_i \boxplus \widehat{\alpha}^{\mathbf{P}}; \mathbf{G})$.

The type $\gamma \widetilde{\alpha}; \mathbf{G}$ represents a prefixing communication. The choice operator \boxplus is n -ary ($n \geq 2$) and commutative, and “;” has higher precedence than \boxplus . The notation $\gamma (\boxplus_{i \in I} \alpha_i^{\mathbf{P}}; \mathbf{K}_i \boxplus \widehat{\alpha}^{\mathbf{P}}; \mathbf{G})$ stands for a choice where the branch initiating with $\widehat{\alpha}^{\mathbf{P}}$ has started to be executed, while the other branches have been discarded.

By abuse of notation, we will write $\gamma \boxplus_{j \in J} \alpha_j^{\mathbf{P}}; \mathbf{G}_j$ for either $\gamma \boxplus_{j \in J} \alpha_j^{\mathbf{P}}; \mathbf{G}_j$ or $\gamma (\boxplus_{j \in J \setminus \{k\}} \alpha_j^{\mathbf{P}}; \mathbf{G}_j \boxplus \widehat{\alpha}_k^{\mathbf{P}}; \mathbf{G}_k)$. In the examples we will write $\widehat{\alpha}_1^{\mathbf{P}}; \mathbf{G}_1 \boxplus \widehat{\alpha}_2^{\mathbf{P}}; \mathbf{G}_2$ and $\gamma (\widehat{\alpha}_1^{\mathbf{P}}; \mathbf{G}_1 \boxplus \dots \boxplus \widehat{\alpha}_n^{\mathbf{P}}; \mathbf{G}_n)$. (Note however that only one of the $\widehat{\alpha}_i^{\mathbf{P}}$ may have a hat.)

We use $\gamma \sqsupset_{i \in I} \alpha_i^{\mathbf{P}}; \mathbf{G}_i$ to stand for $\gamma \widehat{\alpha}^{\mathbf{P}}; \mathbf{G}$ when I is a singleton and for $\gamma \boxplus_{i \in I} \alpha_i^{\mathbf{P}}; \mathbf{G}_i$ otherwise. Moreover, $\gamma (\boxplus_{h \in H} \alpha_h^{\mathbf{P}}; \mathbf{G}_h \boxplus \widehat{\alpha}^{\mathbf{P}}; \mathbf{G})$ stands for $\gamma \widehat{\alpha}^{\mathbf{P}}; \mathbf{G}$ when $H = \emptyset$. We write $\gamma (\boxplus_{j \in J \setminus \{k\}} \alpha_j^{\mathbf{P}}; \mathbf{G}_j \boxplus \widehat{\alpha}_k^{\mathbf{P}}; \mathbf{G}_k)$ when we want to emphasise the branch $\widehat{\alpha}_k^{\mathbf{P}}; \mathbf{G}_k$.

When the checkpoint set γ associated with a choice is empty, we omit it. We call *rooted interaction* a subterm $\widetilde{\alpha}; \mathbf{G}$. So $\widetilde{\alpha}; \mathbf{G}$ can denote either a rooted interaction or a prefixing communication (the context will disambiguate if needed). A type $\gamma \boxplus_{j \in J} \alpha_j^{\mathbf{P}}; \mathbf{G}_j$ is a choice among more than one rooted interaction with the same sender, at most one of which bears a hat. Recursion must be guarded and it is treated equi-recursively.

We use $\text{part}(\mathbf{G})$ to denote the set of participants of \mathbf{G} , as defined in Fig. 3.

We represent global types as trees, with $\gamma \boxplus$ or γ ; on internal nodes, atomic communications on the branches, and End on the leaves. For simplicity we assume that each index set I, J, H is of the form $\{1, \dots, n\}$, where n is its cardinality. The j th branch of a $\gamma \boxplus$ node is labelled by $\widetilde{\alpha}_j$, and the unique branch of a γ ; node, called 0th branch to distinguish it from proper choice branches, is labelled by the corresponding communicating communication. In case the global type has some recursive subtype, the tree is an infinite (regular) tree.

Example 4 (*Trees*) Figure 4 shows the tree representing the global type


$$\mathbf{G} = \widehat{\text{p}} \xrightarrow{\lambda_0} \mathbf{q}; \mu\mathbf{t}.\mathbf{q} \xleftrightarrow{\lambda_1} \mathbf{r}; \mathbf{r} \xrightarrow{\lambda_2} \mathbf{p}; \mathbf{t}_C \boxplus \mathbf{q} \xrightarrow{\lambda_3} \mathbf{p}; \mathbf{p} \xrightarrow{\lambda_4} \mathbf{q}$$

where α_i denotes the communication with message $\Lambda_i = \lambda_i$ or $\Lambda_i = \xleftrightarrow{\lambda_i}$ for $i = 0, \dots, 4$.

We introduce now a notation to distinguish different occurrences of the same communication $\widetilde{\alpha}$ within a type \mathbf{G} .

Communication occurrences in \mathbf{G} , denoted by ξ, ξ' , have the form $\sigma \widetilde{\alpha}$, where $\widetilde{\alpha}$ is a possibly hatted communication and σ is a finite string over Nat, which represents the *path*

Fig. 4 Tree representation of $G = \widehat{\alpha}_0; \mu t.(\alpha_1; \alpha_2; t_C \boxplus \alpha_3; \alpha_4)$



leading to that particular occurrence of α in the tree of G . Formally, the path σ records the branch chosen at each $\gamma \boxplus$ or $\gamma;$ node in G . In particular, the last element of σ specifies the branch labelled by the communication $\tilde{\alpha}$. (Note that in $\sigma \tilde{\alpha}$ the path σ is not empty since it contains at least the index of the branch labelled by $\tilde{\alpha}$.) We use \sqsubseteq (\sqsubset) to denote the prefix ordering (strict prefix ordering) between strings of naturals.

For instance, the two shown occurrences of α_2 in the tree of Fig. 4 are identified by $010\alpha_2$ and $01010\alpha_2$.

Definition 5 (Occurrences) The set of *communication occurrences* of G , written $Occ(G)$, is defined by:

- $Occ(\text{End}) = \emptyset$ $Occ(\mu t.K) = Occ(K\{\mu t.K/t\})$
- $Occ(\gamma \tilde{\alpha}; G) = \{0\tilde{\alpha}\} \cup \{0\xi \mid \xi \in Occ(G)\}$
- $Occ(\gamma \boxplus_{j \in J} \tilde{\alpha}_j; G_j) = \{j\tilde{\alpha}_j \mid j \in J\} \cup \{j\xi \mid j \in J \text{ and } \xi \in Occ(G_j)\}$

In the following, communication occurrences will be simply referred to as *occurrences*. Given an occurrence $\xi = \sigma \tilde{\alpha}$ of G , we define the path of ξ in G to be $\text{path}(\xi) = \sigma$, the communication of ξ to be $\text{comm}(\xi) = \tilde{\alpha}$ and the participants of ξ to be $\text{part}(\xi) = \text{part}(\alpha)$.

An occurrence ξ is *executed* in G if $\text{comm}(\xi)$ bears a hat. We define $ExOcc(G)$, the *set of executed occurrences of G* , as follows:

$$ExOcc(G) = \{\xi \mid \xi \in Occ(G) \text{ and } \text{comm}(\xi) = \widehat{\alpha} \text{ for some } \alpha\}$$

The relation of conflict on occurrences is handy. Intuitively, two occurrences are in conflict if they mutually exclude each other in any computation.

Definition 6 (Conflict) The *conflict relation* in G , denoted by $\#_G$, is the symmetric irreflexive relation on $Occ(G)$ defined by: $\xi \#_G \xi'$ if there exist $\sigma_0, \sigma_1, \sigma_2$ such that $\text{path}(\xi) = \sigma_0 i \sigma_1$ and $\text{path}(\xi') = \sigma_0 j \sigma_2$ with $i \neq j$.

Two conflicting occurrences have paths that diverge at some choice node of the tree. Since the syntax of global types does not allow hats on more than one branch of a choice, this implies that two conflicting occurrences can never be both executed.

We now formalise the notion of causality on occurrences by defining the set of causes of a given occurrence. Intuitively, the set of causes of ξ in G is the set of occurrences in G on which ξ depends, namely those that need to be executed before ξ in any computation allowed by G .

Definition 7 (Causes)

Given an occurrence $\xi \in \text{Occ}(G)$, the *set of causes of ξ in G* , written $\text{Causes}(G, \xi)$, is the smallest set that contains an occurrence ξ' if:

- either $\text{path}(\xi') \sqsubset \text{path}(\xi)$ and $\text{part}(\xi') \cap \text{part}(\xi) \neq \emptyset$;
- or $\xi' \in \text{Causes}(G, \xi'')$ and $\xi'' \in \text{Causes}(G, \xi)$ for some ξ'' .

Our notions of conflict and causality are similar to those defined in [37] for Prime Event Structures (although simpler, since we only deal with trees here).

We require global types to respect causality of occurrences, as formalised in the following definition.

Definition 8 (Causally correct global type) A global type G is *causally correct* if, whenever an occurrence ξ is executed in G , then every occurrence in $\text{Causes}(G, \xi)$ is executed in G , namely if $\xi \in \text{ExOcc}(G)$ implies $\text{Causes}(G, \xi) \subseteq \text{ExOcc}(G)$.

The conflict relation is *hereditary*, namely if $\xi \#_G \xi'$ and $\xi' \in \text{Causes}(G, \xi'')$, then $\xi \#_G \xi''$. To show this property (which is an axiom in Prime Event Structures), it is enough to observe that if $\text{path}(\xi) = \sigma_0 i \sigma_1$ and $\text{path}(\xi') = \sigma_0 j \sigma_2$, then $\sigma_0 j \sigma_2 \sqsubset \text{path}(\xi'')$ implies $\text{path}(\xi'') = \sigma_0 j \sigma_2 \sigma_3$ for some σ_3 .

Session types are projections of global types onto participants. They represent the contributions of individual participants to the session. The projection of a choice yields a union for the choice leader and intersections for the receivers. Checkpoint labels of global types are preserved by the projection onto session types, and the checkpoint label of the choice leader is distinguished by overlining it.

We now define *session pre-types*, which are a superset of session types. Session types will be session pre-types which are projections of global types. Session pre-types are obtained from processes by replacing external and internal choices with intersections and unions, X with \mathbf{t} and end with End , with similar conventions.

Definition 9 (Session pre-types) *Session pre-types* are defined by:

$$\mathsf{T} ::= \Delta \bigwedge_{i \in I} \widetilde{\mathbf{p}_i ? \Lambda_i}; T_i \mid \overline{\mathbf{y}} \bigvee_{i \in I} \widetilde{\mathbf{p}_i ! \Lambda_i}; T_i \mid \mu \mathbf{t}. \mathsf{T} \mid \mathbf{t} \mid \text{End}$$

As for processes, we assume that intersections and unions are not ambiguous, i.e. that the pairs $(\mathbf{p}_i, \Lambda_i)$ ($i \in I$) are all distinct and the Λ_i 's in intersections are either all simple or all connecting messages.

We want projection to ensure that in a global choice, the choice leader makes the decision and all the other participants act accordingly. We do so by requiring that, for any participant except the choice leader, the set of projections of the choice branches on that participant, say $\{T_i \mid i \in I\}$, be consistent, i.e., the *meet of the types in the set*, $\prod_{i \in I} T_i$, be defined.

The meet $\prod_{i \in I} T_i$ is a partial operator, which checks that the T_i 's are compatible behaviours and then combines them into a single session type. Intuitively, $\prod_{i \in I} T_i$ is defined if the

concerned participant receives a message that “notifies” her about the chosen T_i . If one of the T_i ’s is an intersection of simple inputs, then so must be all the other T_i ’s. Instead, intersections of connecting inputs can be combined with End.

To build the meet of intersection types, we define an auxiliary operator \mathbb{m} , which takes an intersection $\bigwedge_{i \in I} \widetilde{p_i?A_i}; T_i$ and an input type (after removing checkpoint labels in both of them) and combines them, if possible:

$$\begin{aligned} \left(\bigwedge_{i \in I} \widetilde{p_i?A_i}; T_i \right) \mathbb{m} \widetilde{p?A}; T &= \bigwedge_{i \in I} \widetilde{p_i?A_i}; T_i \wedge \widetilde{p?A}; T \\ &\text{if the resulting intersection is not ambiguous} \\ \left(\bigwedge_{i \in I} \widetilde{p_i?A_i}; T_i \right) \mathbb{m} \widetilde{p?A}; T &= \bigwedge_{i \in I} \widetilde{p_i?A_i}; T_i \\ &\text{if } p = p_j \text{ and } A = A_j \text{ and } T = T_j \text{ for some } j \in I \\ \left(\bigwedge_{i \in I} \widetilde{p_i? \overset{\lambda_i}{\leftrightarrow}}; T_i \right) \mathbb{m} \text{End} &= \bigwedge_{i \in I} \widetilde{p_i? \overset{\lambda_i}{\leftrightarrow}}; T_i \end{aligned}$$

To define \mathbb{m} on two intersection types, we just iterate the above definition on the members of one of the intersections. In a similar way, we can extend \mathbb{m} to a set of types. Using \mathbb{m} we are able to build meets.

Definition 10 (Meet) The meet of a set of session pre-types is defined by:

$$\bigsqcap_{i \in I} T_i = \begin{cases} \Delta \mathbb{m}_{i \in I} T'_i & \text{if } T_i = \Delta_i T'_i \text{ for all } i \in I \text{ and } \Delta = \bigcup_{i \in I} \Delta_i \\ \text{End} & \text{if } T_i = \text{End for all } i \in I \end{cases}$$

If possible we combine the T_i with \mathbb{m} as explained above. The set of checkpoint labels of the resulting intersection is the union of the corresponding sets for the T_i ’s. If some T_i is End, it means that the participant terminates in the branch T_i . Then it must terminate or be connecting in all the other branches.

We do not need to define the meet when the arguments are recursive types, since we consider recursion equi-recursively. The meet is not defined in all other cases, i.e. between intersections and unions, between unions etc.

To define projection we need one last auxiliary operator, the *labelling* of a session pre-type T by a set γ of at most one checkpoint label (notation $\gamma[T]$), defined by:

$$\gamma[\Delta \bigwedge_{i \in I} \widetilde{\pi_i}; T_i] = \gamma \cup \Delta \bigwedge_{i \in I} \widetilde{\pi_i}; T_i \quad \gamma[\text{End}] = \text{End}$$

Labelling adds γ to the first (possibly empty) set Δ found in T . Intuitively, the checkpoint labels that are spread along successive choices in the global type may get grouped together on a single local choice represented by an intersection when projected on participants. Labelling of union and recursive types is not defined since this operator is only used with types resulting from the application of the meet.

The projection of global types uses the projections of rooted interactions, see the first line of Fig. 5, where the projection operator “ $\bar{\cdot}$ ” has higher precedence than “ \cdot ”. The projection of a choice is a union for the sending participant, and it is otherwise computed as the meet of the projections on the branches. By abuse of notation, in Fig. 5 $\bar{\gamma}$ means \bar{C} if $\gamma = C$ and \emptyset if $\gamma = \emptyset$. The meet operation can be undefined, and therefore also the projection of global types can be undefined. The definition of projection does not ensure that all the branches of a

$$\begin{aligned}
 (\widehat{p \xrightarrow{A} q}; G) \upharpoonright p &= \widehat{q!A}; G \upharpoonright p & (\widehat{p \xrightarrow{A} q}; G) \upharpoonright q &= \widehat{p?A}; G \upharpoonright q & (\widehat{p \xrightarrow{A} q}; G) \upharpoonright r &= G \upharpoonright r, \text{ if } r \neq p, q \\
 (\gamma \widetilde{\alpha^p}; G) \upharpoonright r &= \begin{cases} \overline{\gamma}(\widetilde{\alpha^p}; G) \upharpoonright r & \text{if } r = p \\ \gamma(\widetilde{\alpha^p}; G) \upharpoonright r & \text{otherwise} \end{cases} \\
 (\gamma \boxplus_{i \in I} \widetilde{\alpha_i^p}; G_i) \upharpoonright r &= \begin{cases} \overline{\gamma} \bigvee_{i \in I} (\widetilde{\alpha_i^p}; G_i) \upharpoonright r & \text{if } r = p \\ \gamma \llbracket \bigwedge_{i \in I} (\widetilde{\alpha_i^p}; G_i) \upharpoonright r \rrbracket & \text{otherwise} \end{cases} \\
 (\mu t. K) \upharpoonright p &= \begin{cases} \mu t. K \upharpoonright p & \text{if } p \in \text{part}(K) \\ \text{End} & \text{otherwise} \end{cases} & \mathbf{t} \upharpoonright p &= \mathbf{t} & \text{End} \upharpoonright p &= \text{End}
 \end{aligned}$$

Fig. 5 Projection of global types onto participants

choice $\gamma \boxplus_{i \in I} \widetilde{\alpha_i^p}; G_i$ have the same participants. They can differ for participants whose first communication is a connecting input. Notice that projection respects hats and checkpoint labels (transforming C to \overline{C} for the leader of a C -labelled choice).

Example 5 (Projection) Let us see our notion of projection at work on a few examples.

1. If G is the global type shown in Example 4 we get

$$\begin{aligned}
 G \upharpoonright p &= \widehat{q! \lambda_0}; \mu t. (r? \lambda_2; \mathbf{t}_C \wedge q? \lambda_3; q! \lambda_4) & G \upharpoonright r &= \mu t. C q? \xleftrightarrow{\lambda_1}; p! \lambda_2; \mathbf{t} \\
 G \upharpoonright q &= \widehat{p? \lambda_0}; \mu t. (r! \xleftrightarrow{\lambda_1}; \mathbf{t}_{\overline{C}} \vee p! \lambda_3; p? \lambda_4)
 \end{aligned}$$

2. An example showing how projections can decorate intersections by more checkpoint labels is $G = p \xrightarrow{\lambda_1} q; (p \xrightarrow{\lambda_2} r_{C_1} \boxplus p \xrightarrow{\lambda_3} s) \boxplus p \xrightarrow{\lambda_4} q; (q \xrightarrow{\lambda_5} r_{C_2} \boxplus q \xrightarrow{\lambda_6} s)$ where $G \upharpoonright r = p? \xleftrightarrow{\lambda_2}_{\{C_1, C_2\}} \wedge q? \xleftrightarrow{\lambda_5}$.
3. Let $G = \mu t. (p \xrightarrow{\lambda_1} q; \mathbf{t} \boxplus p \xrightarrow{\lambda_2} q)$ and $G' = p \xrightarrow{\lambda_1} q; G \boxplus p \xrightarrow{\lambda_2} q$. Then $G \upharpoonright q = \mu t. (p? \lambda_1; \mathbf{t} \wedge p? \lambda_2)$ and $G' \upharpoonright q = p? \lambda_1; G \upharpoonright q \wedge p? \lambda_2$.

In the following we will consider only *causally correct and projectable global types*. The Subject Reduction Theorem (Theorem 1) shows that these conditions are preserved by reducing networks that can be typed using the rules introduced in the next section.

We end this section by defining session types.

Definition 11 (Session types) A session pre-type T is a *session type* if $T = G \upharpoonright p$ for some global type G and some participant p .

4 Type system

In this section we define our type system and prove some initial results about it. The shape of *typing judgements* is $\Gamma \vdash P : T$, where the environment Γ associates process variables with session types: $\Gamma ::= \emptyset \mid \Gamma, X : T$. Process typing exploits the correspondence between external choices and intersections, internal choices and unions. The typing rules for both processes and networks are given in Fig. 6. Typing respects hats: in rules [T-EXTCH], and [T-INTCH], the hats in processes and types are exactly on the same actions.

$\frac{\Gamma \vdash P_i : \mathbb{T}_i \ (i \in I)}{\Gamma \vdash \Delta \sum_{i \in I} \tilde{\pi}_i; P_i : \Delta \bigwedge_{i \in I} \tilde{\pi}_i; \mathbb{T}_i} \text{ [T-EXTCH]}$	$\frac{\Gamma \vdash P_i : \mathbb{T}_i \ (i \in I)}{\Gamma \vdash \overline{\gamma} \bigoplus_{i \in I} \tilde{\pi}_i; P_i : \overline{\gamma} \bigvee_{i \in I} \tilde{\pi}_i; \mathbb{T}_i} \text{ [T-INTCH]}$
$\Gamma \vdash \text{end} : \text{End} \text{ [T-end]}$	$\frac{\Gamma, X : \mathbf{t} \vdash P : \mathbb{T}}{\Gamma \vdash \mu X.P : \mu \mathbf{t}.\mathbb{T}} \text{ [T-REC]}$
$\frac{\Gamma, X : \mathbf{t} \vdash X : \mathbf{t}}{\Gamma, X : \mathbf{t} \vdash X : \mathbf{t}} \text{ [T-VAR]}$	
$\frac{\vdash P_i : \mathbb{T}_i \quad \mathbb{T}_i \leq G \upharpoonright \mathbf{p}_i \quad (1 \leq i \leq n) \quad \text{part}(G) \subseteq \{\mathbf{p}_1, \dots, \mathbf{p}_n\}}{\vdash \mathbf{p}_1 \llbracket P_1 \rrbracket \parallel \dots \parallel \mathbf{p}_n \llbracket P_n \rrbracket : G} \text{ [T-NET]}$	

Fig. 6 Typing rules for processes and networks

$\frac{\forall i \in I : \mathbb{T}_i \leq \mathbb{T}'_i}{\Delta \bigwedge_{i \in I \cup I'} \tilde{\pi}_i; \mathbb{T}_i \leq \Delta \bigwedge_{i \in I} \tilde{\pi}_i; \mathbb{T}'_i} \text{ [SUB-IN]}$	$\frac{\forall i \in I : \mathbb{T}_i \leq \mathbb{T}'_i}{\overline{\gamma} \bigvee_{i \in I} \tilde{\pi}_i; \mathbb{T}_i \leq \overline{\gamma} \bigvee_{i \in I} \tilde{\pi}_i; \mathbb{T}'_i} \text{ [SUB-OUT]}$
$\frac{\forall i \in I : \pi_i \text{ is a connecting input}}{\Delta \bigwedge_{i \in I} \tilde{\pi}_i; \mathbb{T}_i \leq \text{End}} \text{ [SUB-IN-SKIP]}$	$\frac{}{\text{End} \leq \text{End}} \text{ [SUB-END]}$

Fig. 7 Subtyping rules

Figure 7 gives the subtyping rules, where the double line indicates that the rules are interpreted *coinductively* [33] (Chapter 21). Subtyping takes into account the rules for intersection and union and preserves hats. Rule [SUB-IN-SKIP] reflects the fact that connecting inputs can be added without causing problems. Rule [T-NET] is the only rule for typing networks: it requires that the types of all processes be subtypes of the projections of a unique global type. The condition $\text{part}(G) \subseteq \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ ensures the presence of all session participants and allows the typing of sessions containing $\mathbf{p} \llbracket \text{end} \rrbracket$ for any \mathbf{p} , a property needed to guarantee invariance of types under structural equivalence of networks. Clearly, typing imposes constraints on the way hats and checkpoint labels are placed within processes.

Example 6 (*ok condition on evaluation contexts*) The following example shows the need of the *ok* condition in reduction rules to assure subject reduction. Let $P = \mu X.(q!\lambda_1; r!\lambda_2; X \ C \oplus q!\lambda_3; r!\lambda_4; X)$ and $Q = \mu Y.(p?\lambda_1; Y \ C + p?\lambda_3; Y)$ and $R = \mu Z.(p?\lambda_2; Z \ C + p?\lambda_4; Z)$. The network $\mathbf{p} \llbracket P \rrbracket \parallel \mathbf{q} \llbracket Q \rrbracket \parallel \mathbf{r} \llbracket R \rrbracket$ reduces by forward reductions to $\mathbf{p} \llbracket P' \rrbracket \parallel \mathbf{q} \llbracket Q' \rrbracket \parallel \mathbf{r} \llbracket R' \rrbracket$ where

$P' = \widehat{q!\lambda_1}; \widehat{r!\lambda_2}; (q!\lambda_1; r!\lambda_2; P \ C \oplus q!\lambda_3; r!\lambda_4; P) \ C \oplus q!\lambda_3; r!\lambda_4; P$ and $Q' = \widehat{p?\lambda_1}; (p?\lambda_1; Q \ C + \widehat{p?\lambda_3}; Q) \ C + p?\lambda_3; Q$ and $R' = \widehat{p?\lambda_2}; R \ C + p?\lambda_4; R$. By rule [BACK], $\mathbf{p} \llbracket P' \rrbracket \parallel \mathbf{q} \llbracket Q' \rrbracket \parallel \mathbf{r} \llbracket R' \rrbracket \xrightarrow{C} \mathbf{p} \llbracket P'' \rrbracket \parallel \mathbf{q} \llbracket Q \rrbracket \parallel \mathbf{r} \llbracket R \rrbracket$ where $P'' = q!\lambda_3; r!\lambda_4; P$. Without the condition \mathcal{E} *ok* for \overline{C} on rule [CTBA], we could have also the backward move:

$\mathbf{p} \llbracket P' \rrbracket \parallel \mathbf{q} \llbracket Q' \rrbracket \parallel \mathbf{r} \llbracket R' \rrbracket \xrightarrow{C} \mathbf{p} \llbracket P''' \rrbracket \parallel \mathbf{q} \llbracket Q'' \rrbracket \parallel \mathbf{r} \llbracket R \rrbracket$, where $P''' = (\widehat{q!\lambda_1}; \widehat{r!\lambda_2}; q!\lambda_1; r!\lambda_2; P) \ C \oplus q!\lambda_3; r!\lambda_4; P$ and $Q'' = \widehat{p?\lambda_1}; Q \ C + p?\lambda_3; Q$. Then Subject Reduction would fail, since the network $\mathbf{p} \llbracket P \rrbracket \parallel \mathbf{q} \llbracket Q \rrbracket \parallel \mathbf{r} \llbracket R \rrbracket$ is typable with the global type $\mu \mathbf{t}.\mathbf{p} \xrightarrow{\lambda_1} \mathbf{q}; \mathbf{p} \xrightarrow{\lambda_2} \mathbf{r}; \mathbf{t} \ C \boxplus \mathbf{p} \xrightarrow{\lambda_3} \mathbf{q}; \mathbf{p} \xrightarrow{\lambda_4} \mathbf{r}; \mathbf{t}$, while $\mathbf{p} \llbracket P''' \rrbracket \parallel \mathbf{q} \llbracket Q'' \rrbracket \parallel \mathbf{r} \llbracket R \rrbracket$ is not typable. In fact the output $r!\lambda_2$ has a hat in the session type of P''' , while the corresponding input $p?\lambda_2$ does not have a hat in the session type of R . This example shows also

why it would not be possible to roll back to the innermost checkpoints in recursive processes, since they could be different for different participants.

In the remainder of this section we prove some properties of our type system which will be used to show the soundness results in the next section. The reader not interested in proofs can go directly to Sect. 5. We start with the classical lemmas of inversion and canonical forms.

Lemma 1 (Inversion Lemma)

1. If $\Gamma \vdash \Delta \sum_{i \in I} \tilde{\pi}_i; P_i : \top$, then $\top = \Delta \bigwedge_{i \in I} \tilde{\pi}_i; \top_i$ and $\Gamma \vdash P_i : \top_i$ for $i \in I$.
2. If $\Gamma \vdash \bar{\gamma} \bigoplus_{i \in I} \tilde{\pi}_i; P_i : \top$, then $\top = \bar{\gamma} \bigvee_{i \in I} \tilde{\pi}_i; \top_i$ and $\Gamma \vdash P_i : \top_i$ for $i \in I$.
3. If $\Gamma \vdash \mu X. P : \top$, then $\top = \mu t. \top'$ and $\Gamma, X : t \vdash P : \top'$.
4. If $\Gamma \vdash X : \top$, then $\top = t$ and $\Gamma = \Gamma', X : t$.
5. If $\Gamma \vdash \text{end} : \top$, then $\top = \text{End}$.
6. If $\vdash \mathfrak{p}_1 \llbracket P_1 \rrbracket \parallel \dots \parallel \mathfrak{p}_n \llbracket P_n \rrbracket : \mathbb{G}$, then $\vdash P_i : \top_i$ and $\top_i \leq \mathbb{G} \upharpoonright \mathfrak{p}_i$ for $1 \leq i \leq n$ and $\text{part}(\mathbb{G}) \subseteq \{\mathfrak{p}_1, \dots, \mathfrak{p}_n\}$.

Lemma 2 (Canonical Form Lemma)

1. If $\Gamma \vdash P : \Delta \bigwedge_{i \in I} \tilde{\pi}_i; \top_i$, then $P = \Delta \sum_{i \in I} \tilde{\pi}_i; P_i$ and $\Gamma \vdash P_i : \top_i$ for $i \in I$.
2. If $\Gamma \vdash P : \bar{\gamma} \bigvee_{i \in I} \tilde{\pi}_i; \top_i$, then $P = \bar{\gamma} \bigoplus_{i \in I} \tilde{\pi}_i; P_i$ and $\Gamma \vdash P_i : \top_i$ for $i \in I$.
3. If $\Gamma \vdash P : \mu t. \top$, then $P = \mu X. Q$ and $\Gamma, X : t \vdash Q : \top$.
4. If $\Gamma \vdash P : t$, then $P = X$ and $\Gamma = \Gamma', X : t$.
5. If $\Gamma \vdash P : \text{End}$, then $P = \text{end}$.
6. If $\vdash \mathbb{N} : \mathbb{G}$ and $\text{part}(\mathbb{G}) = \{\mathfrak{p}_1, \dots, \mathfrak{p}_n\}$, then $\mathbb{N} \equiv \mathfrak{p}_1 \llbracket P_1 \rrbracket \parallel \dots \parallel \mathfrak{p}_n \llbracket P_n \rrbracket$ and $\vdash P_i : \top_i$ and $\top_i \leq \mathbb{G} \upharpoonright \mathfrak{p}_i$ for $1 \leq i \leq n$.

The mapping $\wr \wr$ defined at page 7 for processes may be extended in the obvious way to session types. As expected, this mapping preserves typing.

Lemma 3 If $\vdash P : \top$, then $\vdash \wr P \wr : \wr \top \wr$.

Another classical lemma we need for proving Subject Reduction is the following, which retrieves the shape of processes and networks from their labelled transitions.

- Lemma 4**
1. If $P \xrightarrow{\mathfrak{p}^? \Lambda} P'$, then $P = \mathcal{E}[\Delta \sum_{i \in I} \pi_i; P_i]$, where $\pi_j = \mathfrak{p}^? \Lambda$ for some $j \in I$, and $P' = \mathcal{E}[\Delta (\sum_{i \in I \setminus \{j\}} \pi_i; P_i + \hat{\pi}_j; P_j)]$.
 2. If $P \xrightarrow{\mathfrak{p}! \Lambda} P'$, then $P = \mathcal{E}[\bar{\gamma} \bigoplus_{i \in I} \pi_i; P_i]$, where $\pi_j = \mathfrak{p}! \Lambda$ for some $j \in I$, and $P' = \mathcal{E}[\bar{\gamma} (\bigoplus_{i \in I \setminus \{j\}} \pi_i; P_i \oplus \hat{\pi}_j; P_j)]$.
 3. If $P \xrightarrow{\bar{C}} P'$, then $P = \mathcal{E}[\bar{C} (\bigoplus_{i \in I} \pi_i; Q_i \oplus \hat{\pi}; Q)]$ and $Q \downarrow_{\text{out}}$ and $I \neq \emptyset$ and $\mathcal{E} \text{ ok}$ for \bar{C} and $P' = \mathcal{E}[\bar{C} \bigoplus_{i \in I} \pi_i; Q_i]$.
 4. If $P \xrightarrow{C} P'$, then $P = \mathcal{E}[\Delta \sum_{i \in I} P_i]$ and $C \in \Delta$ and $\mathcal{E} \text{ ok}$ for C and $P' = \mathcal{E}[\wr \Delta \sum_{i \in I} P_i \wr]$.
 5. If $P \not\xrightarrow{C}$, then P does not contain an executed external choice with checkpoint label Δ and $C \in \Delta$.
 6. If $\mathbb{N} \xrightarrow{\mathfrak{p}! \Lambda \mathfrak{q}} \mathbb{N}'$, then $\mathbb{N} = \mathfrak{p} \llbracket P \rrbracket \parallel \mathfrak{q} \llbracket Q \rrbracket \parallel \mathbb{N}_0$ and $P \xrightarrow{\mathfrak{q}! \Lambda} P'$ and $Q \xrightarrow{\mathfrak{p}^? \Lambda} Q'$ and $\mathbb{N}' = \mathfrak{p} \llbracket P' \rrbracket \parallel \mathfrak{q} \llbracket Q' \rrbracket \parallel \mathbb{N}_0$.
 7. If $\mathbb{N} \xrightarrow{C} \mathbb{N}'$, then $\mathbb{N} = \mathfrak{p} \llbracket P \rrbracket \parallel \prod_{h \in H} \mathfrak{p}_h \llbracket P_h \rrbracket \parallel \prod_{k \in K} \mathfrak{p}_k \llbracket P_k \rrbracket$ and $P \xrightarrow{C} P'$ and $P_h \xrightarrow{C} P'_h$ for all $h \in H$ and $P_k \not\xrightarrow{C}$ for all $k \in K$ and

$$\mathbb{N}' = \mathfrak{p} \llbracket P' \rrbracket \parallel \prod_{h \in H} \mathfrak{p}_h \llbracket P'_h \rrbracket \parallel \prod_{k \in K} \mathfrak{p}_k \llbracket P_k \rrbracket$$

We now introduce some notions that are specific to the present type system, and prove some results about them. More precisely:

- we define session contexts (Definition 12) and we show that they correspond to evaluation contexts (Lemma 5);
- by extending the definition of meet to session contexts (Definition 14), we show how session contexts may be retrieved as projections of global contexts (Definition 13);
- we show that the maximal hatted path in the tree of a global type splits the global type into a global context and the global type filling the hole (Definition 15);
- the main result is Lemma 6: it devises the shape of global types starting from one of its projections. It uses the notion of “global type affecting a participant” (Definition 16).

We start by defining session contexts, which mirror process evaluation contexts (Definition 3):

Definition 12 (*Session contexts*) *Session contexts* are defined by:

$$\mathcal{T} ::= \Delta(\bigwedge_{h \in H} p_h ? A_h; T_h \wedge \widehat{p} ? \Lambda; \mathcal{T}) \mid \overline{\gamma}(\bigvee_{h \in H} p_h ! A_h; T_h \vee \widehat{p} ! \Lambda; \mathcal{T}) \mid []$$

We generalise typing, subtyping, and the definitions of ok for C (\overline{C}) to session contexts. We omit these definitions, which are trivial due to the correspondence between process evaluation contexts and session contexts.

The following lemma gives easy relations between contexts and typing/subtyping.

- Lemma 5**
1. If $\mathcal{T}[T] \leq T'$ and $T' \neq \text{End}$, then $T' = T'[T'']$ and $\mathcal{T} \leq T'$ and $T \leq T''$. Moreover if T is ok for C (\overline{C}), then T' is ok for C (\overline{C}).
 2. If $T \leq \mathcal{T}[T']$, then $T = T'[T'']$ and $T' \leq \mathcal{T}$ and $T'' \leq T'$. Moreover if T is ok for C (\overline{C}), then T' is ok for C (\overline{C}).
 3. If $\Gamma \vdash \mathcal{E}[P] : T$, then $T = \mathcal{T}[T']$ and $\Gamma \vdash \mathcal{E} : \mathcal{T}$ and $\Gamma \vdash P : T'$. Moreover if \mathcal{E} is ok for C (\overline{C}), then \mathcal{T} is ok for C (\overline{C}).
 4. If $\Gamma \vdash \mathcal{E} : \mathcal{T}$ and $\Gamma \vdash P : T$ and $\mathcal{T}[T]$ is a type, then $\Gamma \vdash \mathcal{E}[P] : \mathcal{T}[T]$.
 5. If $\Gamma \vdash P : \mathcal{T}[T]$, then $P = \mathcal{E}[Q]$ and $\Gamma \vdash \mathcal{E} : \mathcal{T}$ and $\Gamma \vdash Q : T$. Moreover if T is ok for C (\overline{C}), then \mathcal{E} is ok for C (\overline{C}).

Proof By induction on contexts.

We define now global contexts, which identify subtypes of global types which do not occur in discarded branches.

Definition 13 (*Global contexts*) *Global contexts* are defined by:

$$\mathcal{G} ::= \gamma \tilde{\alpha}; \mathcal{G} \mid \gamma(\boxplus_{i \in I} \alpha_i; K_i \boxplus \tilde{\alpha}; \mathcal{G}) \mid []$$

We extend to session contexts the labelling defined at page 14 by $\gamma[] = []$. This allows us to define projections of global contexts into session contexts using $[] \upharpoonright p = []$ and generalising the definition of the meet operator as follows:

Definition 14 (*Meet with contexts*) The *meet of session contexts and session types* is defined by:

$$\begin{aligned} \Delta(\bigwedge_{h \in H} T_h \wedge \widehat{\pi}; \mathcal{T}) \sqcap \Delta T &= \Delta \cup \Delta((\bigwedge_{h \in H} T_h \sqcap T) \wedge \widehat{\pi}; \mathcal{T}) && \text{if the resulting} \\ &&& \text{intersection is a} \\ &&& \text{session context;} \\ \Delta(\bigwedge_{h \in H} T_h \wedge \widehat{\pi}; \mathcal{T}) \sqcap \text{End} &= \Delta(\bigwedge_{h \in H} T_h \wedge \widehat{\pi}; \mathcal{T}) && \text{if } T_h \text{ for } h \in H \text{ and} \\ &&& \widehat{\pi} \text{ are all connecting inputs.} \end{aligned}$$

Projections of global context define session contexts only for a subset of participants. A simple example is $\mathcal{G} = p \xrightarrow{\lambda} q; []$, whose projections on all $r \neq p, q$ is $\mathcal{G} \upharpoonright r = []$. However, $\mathcal{G} \upharpoonright p = q! \lambda; []$ and $\mathcal{G} \upharpoonright q = p? \lambda; []$ are not session contexts because the actions before the hole are not hatted. A more interesting example is the ternary choice:

$$\mathcal{G}' = p \xleftrightarrow{\lambda_1} q \boxplus p \xleftrightarrow{\lambda_2} r \boxplus p \xleftrightarrow{\lambda_3} q; []$$

In fact $\mathcal{G}' \upharpoonright p = q! \xleftrightarrow{\lambda_1} \vee r! \xleftrightarrow{\lambda_2} \vee q! \xleftrightarrow{\lambda_3}; []$ and $\mathcal{G}' \upharpoonright q = p? \xleftrightarrow{\lambda_1} \wedge p? \xleftrightarrow{\lambda_3}; []$ are session contexts, while $\mathcal{G}' \upharpoonright r = p? \xleftrightarrow{\lambda_2} \wedge []$ is not. Notice that without connecting communications such kind of example would be longer, since the participants q and r would have to occur in all branches of the choice.

The paths that are not on discarded branches in the tree representation of a global type G may be used to split G between a global context and the subtype filling the hole. This is formalised in the following definition.

Definition 15 (*Contexts and subtypes determined by paths*) Let σ be a path in a global type G .

1. The *context determined by σ* , $Ctx(G, \sigma)$, is defined by:

- $Ctx(G, \epsilon) = []$
- $Ctx(\gamma \tilde{\alpha}; G', 0\sigma) = \gamma \tilde{\alpha}; Ctx(G', \sigma)$
- $Ctx(\gamma (\boxplus_{j \in J \setminus \{k\}} \alpha_j; K_j \boxplus \tilde{\alpha}_k; G_k), k\sigma) = \gamma (\boxplus_{j \in J \setminus \{k\}} \alpha_j; K_j \boxplus \tilde{\alpha}_k; Ctx(G_k, \sigma))$
- $Ctx(\mu t.K, \sigma) = Ctx(K\{\mu t.K/t\}, \sigma)$

2. The *subtype determined by σ* , $SubT(G, \sigma)$, is defined by:

- $SubT(G, \epsilon) = G$
- $SubT(\gamma \tilde{\alpha}; G', 0\sigma) = SubT(G', \sigma)$
- $SubT(\gamma (\boxplus_{j \in J \setminus \{k\}} \alpha_j; K_j \boxplus \tilde{\alpha}_k; G_k), k\sigma) = SubT(G_k, \sigma)$
- $SubT(\mu t.K, \sigma) = SubT(K\{\mu t.K/t\}, \sigma)$

Note that $Ctx(G, \sigma)$ and $SubT(G, \sigma)$ are defined if and only if the occurrences on σ are not in conflict with executed occurrences, in other words if all $\xi \in ExOcc(G)$ are such that $path(\xi) \sqsubseteq \sigma$ or $\sigma \sqsubset path(\xi)$.

It is easy to verify that $G = Ctx(G, \sigma)[SubT(G, \sigma)]$ whenever $Ctx(G, \sigma)$ is defined.

The properties of our calculus mainly depend on the possibility of deriving information on the shape of a global type from its projections on participants. A useful notion is that of global type ‘affecting’ a participant, which essentially means that the whole type is needed in order to obtain the projection on that participant. A global type G affects p if:

- either there is one branch of G whose first communication has participant p
- or there are two branches of G whose projections on p are both different from End .

Definition 16 (*Affecting global types*) A global type $\gamma \boxplus_{i \in I} \tilde{\alpha}_i; G_i$ affects participant p if one of the following holds:

- $p \in part(\alpha_i)$ for some $i \in I$;
- $|\{i \in I \mid \tilde{\alpha}_i; G_i \upharpoonright p \neq End\}| > 1$.

For instance, if G is the global type of Example 4, then G affects participants p and q but it does not affect participant r , as shown by their projections given in Example 5(1).

In the projection of a global type G on p , the intersections have sets of checkpoint labels. We know that at most one of these checkpoint labels is the label of the subtype of G which affects p . But which one? The projection does not give us enough information to decide. We overcome this problem by means of an equality relation on session types which disregards the sets of checkpoint labels decorating intersections.

We say that two session types T and T' are *equal up to intersection labelling*, dubbed $T \doteq T'$, if they are equal ($T = T'$) or they are the same intersection with different sets of checkpoint labels ($T = \Delta T''$, $T' = \Delta' T''$ for some T'').

We have now enough machinery to prove that if the projection of G on p is split into a session context and a session type T which is either a union or an intersection of simple inputs, then this is mirrored by a splitting of G into a global context and a subtype which affects p . Moreover, if the session context is a hole, then all the choices along the path leading to the hole in G must be unary choices. Distinguishing the two shapes of T we can also show that:

- if T is a union type, then the subtype of G is a choice with leader p ;
- if T is an intersection of simple inputs, then each “long enough” path in G leads to a communication that projects onto one of these inputs.

Lemma 6 1. *If $G \upharpoonright p = \mathcal{T}[\overline{\gamma} \bigvee_{i \in I} T_i]$, then there is a unique path σ in G such that $SubT(G, \sigma) = \gamma \sqsupset_{i \in I} \widetilde{\alpha}_i^p$; G_i affects p and*

$$Ctx(G, \sigma) \upharpoonright p = \mathcal{T} \text{ and } SubT(G, \sigma) \upharpoonright p = \overline{\gamma} \bigvee_{i \in I} T_i.$$

2. *If $G \upharpoonright p = \mathcal{T}[\Delta \bigwedge_{i \in I} \widetilde{q}_i ? \lambda_i; T_i]$, then there is a unique path σ in G such that $SubT(G, \sigma)$ affects p and*

$$Ctx(G, \sigma) \upharpoonright p = \mathcal{T} \text{ and } SubT(G, \sigma) \upharpoonright p \doteq \Delta \bigwedge_{i \in I} \widetilde{q}_i ? \lambda_i; T_i.$$

Moreover for each path σ' in $SubT(G, \sigma)$ there are $j \in I$ and σ_j such that

$$\sigma \sigma_j \xrightarrow{\lambda_j} p \in Occ(G) \text{ and either } \sigma_j \sqsubseteq \sigma' \text{ or } \sigma' \sqsubset \sigma_j.$$

In both cases if \mathcal{T} is the hole, then σ is a possibly empty string of 0's.

Proof (1) By induction on session type contexts. If $\mathcal{T} = []$, then $G \upharpoonright p = \overline{\gamma} \bigvee_{i \in I} T_i$ and by definition of projection $G = \gamma \sqsupset_{i \in I} \widetilde{\alpha}_i^p; G_i$, so G affects p . In this case $\sigma = \epsilon$ is the required path, since $SubT(G, \epsilon) = G$ affects p and $Ctx(G, \epsilon) \upharpoonright p = [] \upharpoonright p = []$ and $SubT(G, \epsilon) \upharpoonright p = G \upharpoonright p = \overline{\gamma} \bigvee_{i \in I} T_i$.

Moreover, σ is unique because for any path $\sigma' \neq \epsilon$ the first communication along σ' would be an output from p , which would imply $Ctx(G, \sigma') \upharpoonright p \neq []$ whenever $Ctx(G, \sigma') \upharpoonright p$ is defined.

If $\mathcal{T} = \Delta (\bigwedge_{i \in I} T_i \widehat{q} ? \Lambda; T')$, then $\widehat{q} \xrightarrow{\Lambda} p$ occurs in G by definition of projection. This means that there exists $\xi \in Occ(G)$ such that $\xi = \sigma' q \xrightarrow{\Lambda} p$ and there does not exist $\xi' \in Occ(G)$ such that $path(\xi') \sqsubset \sigma'$ and $p \in part(\xi')$. The definition of meet implies that there is no $j \in I$ such that T_j starts by $q ? \Lambda$. Let $G' = SubT(G, \sigma')$. By construction $G' \upharpoonright p = T'[\overline{\gamma} \bigvee_{i \in I} T_i]$. By induction hypothesis on T' there is a unique path σ'' such that $SubT(G', \sigma'')$ affects p and $Ctx(G', \sigma'') \upharpoonright p = T'$ and $SubT(G', \sigma'') \upharpoonright p = \overline{\gamma} \bigvee_{i \in I} T_i$. We can then choose $\sigma = \sigma' \sigma''$. The proof of the other inductive case is similar.

(2) By induction on session type contexts. Let $T = \Delta \bigwedge_{i \in I} \widetilde{q_i ? \lambda_i}; T_i$. If $T = []$ the proof proceeds by an inner induction on the path $\text{ap}(p, G)$ defined by:

$$\text{ap}(p, G) = \begin{cases} 0 \text{ ap}(p, G') & \text{if } G = \tilde{\alpha}; G' \text{ and } p \notin \text{part}(\alpha) \\ \epsilon & \text{otherwise} \end{cases}$$

Intuitively, $\text{ap}(p, G)$ represents the access path to the first choice affecting p in G . Let $\sigma = \text{ap}(p, G)$. If $\sigma = \epsilon$, then $G = \gamma \sqsupset_{h \in H} \tilde{\alpha}_h; G_h$ must satisfy either $|H| > 1$ or $p \in \text{part}(\alpha_j)$ for the unique $j \in H$. In the latter case, G trivially affects p . If $|H| > 1$, then by definition of projection and the fact that all inputs are simple, there are two branches whose projections on p are different from End . Therefore G affects p . In both cases we have $\text{Ctx}(G, \epsilon) \upharpoonright p = [] \upharpoonright p = []$ and $\text{SubT}(G, \epsilon) \upharpoonright p = G \upharpoonright p = T$.

We show that σ is unique. Let $\sigma' = j \sigma''$ for some $j \in H$ and suppose that $\text{Ctx}(G, \sigma') \upharpoonright p$ is defined. If $p \in \text{part}(\alpha_j)$, then clearly $\text{Ctx}(G, \sigma') \upharpoonright p \neq []$. If $|H| > 1$, then $\text{Ctx}(G, \sigma') \upharpoonright p$ contains the projection on p of all the paths of G starting with a branch different from the j th one. Since p appears as receiver along all these paths as shown below, this again implies $\text{Ctx}(G, \sigma') \upharpoonright p \neq []$.

If $\sigma = 0 \sigma'$ where $\sigma' = \text{ap}(p, G')$ and $G = \tilde{\alpha}; G'$, then by inductive hypothesis on σ' we have that $\text{SubT}(G', \sigma')$ affects p and $\text{Ctx}(G', \sigma') \upharpoonright p = []$ and $\text{SubT}(G', \sigma') \upharpoonright p = T$. Therefore, $\text{SubT}(G, 0 \sigma')$ affects p and $\text{Ctx}(G, 0 \sigma') \upharpoonright p = []$ and $\text{SubT}(G, 0 \sigma') \upharpoonright p = T$. By induction hypothesis σ' is unique and so is σ .

The proof for $T \neq []$ is as in (1).

In the remaining we prove that for each path σ' in $\text{SubT}(G, \sigma)$ there are σ'' and $j \in I$

such that $\sigma \sigma'' q_j \xrightarrow{\lambda_j} p \in \text{Occ}(G)$ and either $\sigma'' \sqsubseteq \sigma'$ or $\sigma' \sqsubset \sigma''$. The path σ' crosses different choices. By the assumption that all inputs are simple and by the definition of \mathfrak{m} , the projections of the branches on p must be intersections of types starting with inputs with different senders or messages, namely $\widetilde{q_j ? \lambda_j}$ ($j \in I' \subseteq I$). Moreover, the

corresponding communications $q_j \xrightarrow{\lambda_j} p$ ($j \in I'$) must be the first communications involving p in paths of $\text{SubT}(G, \sigma)$. Therefore, given σ' in $\text{SubT}(G, \sigma)$ there are σ'' and α such that $\xi = \sigma \sigma'' \tilde{\alpha} \in \text{Occ}(G)$ and ξ is the first occurrence with $p \in \text{part}(\xi)$ and either $\sigma'' \sqsubseteq \sigma'$ or $\sigma' \sqsubset \sigma''$. The definition of projection implies $\tilde{\alpha} = q_j \xrightarrow{\lambda_j} p$ for some $j \in I'$. □

5 Soundness

This section is devoted to the formulation and the proof of the properties of our calculus, i.e. subject reduction, session fidelity and progress. These proofs are not trivial due to the presence of both reverse computations and connecting communications.

Lemma 7 deals with subtypes of global contexts. It uses the notions of “alive” and “enabled” for communication occurrences. An occurrence of a communication is alive if it can be executed in the future, and it is enabled if it can be immediately executed. We also define the notions of “alive” and “enabled” for checkpoint labels as they are needed for session fidelity. A checkpoint label is alive if it can be the target of a rollback in the future, and it is enabled if the next action of the choice leader is an output, possibly triggering a backward reduction of the network.

Definition 17 (*Alive and enabled*)

1. An occurrence ξ is alive in G if $\xi \in Occ(G) \setminus ExOcc(G)$ and there is no $\xi' \in ExOcc(G)$ such that $\xi' \#_G \xi$.
2. An occurrence ξ is enabled in G if it is alive in G and $Causes(G, \xi) \subseteq ExOcc(G)$.
3. A path σ is ok for the checkpoint label C in G if C does not occur along the path σ in the tree of G .
4. A checkpoint label C is alive in G if for some path σ ok for C (in G)

$$SubT(G, \sigma) = C(\boxplus_{j \in J \setminus \{k\}} \alpha_j^P; K_j \boxplus \widehat{\alpha}_k^P; G_k)$$

and $|J| > 1$ and there are σ' and α^P such that $\sigma k \sigma' \alpha^P \in Occ(G)$ is alive in G .

5. A checkpoint label C is enabled in G if it is alive in G and $\sigma k \sigma' \alpha^P$ is enabled in G , where σ, k, σ' and α^P are as in (3).

For instance in the global type of Example 4 the occurrence $01\alpha_1$ is enabled and the occurrence $0101\alpha_1$ is alive but not enabled. The checkpoint label C is enabled for the path ϵ in

the global type $p \xrightarrow{\lambda_1} q_C \boxplus p \xrightarrow{\lambda_2} q; p \xleftrightarrow{\lambda_3} r$, since the occurrence $20p \xrightarrow{\lambda_3} r$ is enabled in G . The checkpoint label C is alive but not enabled for the path ϵ in the global type $p \xrightarrow{\lambda_1} q_C \boxplus p \xrightarrow{\lambda_2} q; q \xrightarrow{\lambda_3} p; p \xleftrightarrow{\lambda_4} r$, since the occurrence $200p \xleftrightarrow{\lambda_4} r$ is alive, but not enabled in G .

- Lemma 7**
1. If $\sigma j \alpha$ is enabled in G and $p \in part(\alpha)$, then $Ctx(G, \sigma) \upharpoonright p$ is a session context.
 2. If $SubT(G, \sigma) = \gamma \boxplus_{i \in I} \alpha_i; G_i$, then $\sigma i \alpha_i$ is alive in G for all $i \in I$.
 3. If $SubT(G, \sigma) = \gamma (\boxplus_{i \in I \setminus \{j\}} \alpha_i^P; K_i \boxplus \alpha_j^P; G_j)$ and $recv(\alpha_j^P) = q$ and both $Ctx(G, \sigma) \upharpoonright p$ and $Ctx(G, \sigma) \upharpoonright q$ are session contexts, then $\sigma j \alpha_j^P$ is enabled in G .

Proof (1) If $\sigma j \alpha$ is enabled in G , then $Causes(G, \sigma j \alpha) \subseteq ExOcc(G)$ and $\xi \in ExOcc(G)$ implies $path(\xi) \sqsubseteq \sigma$. Hence $Ctx(G, \sigma)$ is defined. If $p \in part(\alpha)$, define $Causes_p(G, \sigma j \alpha) = Causes(G, \sigma j \alpha) \cap \{\xi \in Occ(G) \mid p \in part(\xi)\}$. Since the projection of $Ctx(G, \sigma)$ on p retains only $Causes_p(G, \sigma j \alpha)$ and forgets everything else on the path σ , and moreover it preserves and reflects hats, $Ctx(G, \sigma) \upharpoonright p$ is a session context.

(2) From $SubT(G, \sigma) = \gamma \boxplus_{i \in I} \alpha_i; G_i$ it follows that $\sigma i \alpha_i \in Occ(G) \setminus ExOcc(G)$. Moreover, since $SubT(G, \sigma)$ is defined, the path σ must follow the executed branches of G , if any, thus $\xi \in Causes(G, \sigma i \alpha_i)$ implies $path(\xi) \sqsubseteq \sigma$. Hence $\sigma i \alpha_i$ is alive in G .

(3) By (2) $\sigma j \alpha_j^P$ is alive in G . We prove that $Causes(G, \sigma j \alpha_j^P) \subseteq ExOcc(G)$ by induction on the definition of $Causes$. If $\xi \in Causes(G, \sigma j \alpha_j^P)$ because $part(\xi) \cap \{p, q\} \neq \emptyset$, then it must be $\xi \in ExOcc(G)$, otherwise the projection of the communication of ξ would appear unhatted in $Ctx(G, \sigma) \upharpoonright p$ or in $Ctx(G, \sigma) \upharpoonright q$. Suppose now that $\xi \in Causes(G, \sigma j \alpha_j^P)$ because $\xi \in Causes(G, \xi')$ and $\xi' \in Causes(G, \sigma j \alpha_j^P)$. By induction hypothesis $\xi' \in ExOcc(G)$, hence by causal correctness of G also $\xi \in ExOcc(G)$. □

Theorem 1 (Subject reduction) *If $\vdash N : G$ and $N \xrightarrow{\ast} N'$, then $\vdash N' : G'$ for some G' .*

Proof By case analysis on the reduction rules for networks. Since the case of structural equivalence is trivial, there are only two rules to consider, Rule [COM] and Rule [BACK]. Let the applied rule be [COM], then

$$\frac{P \xrightarrow{q!A} P' \quad Q \xrightarrow{p?A} Q'}{p \llbracket P \rrbracket \parallel q \llbracket Q \rrbracket \parallel N \xrightarrow{p!Aq} p \llbracket P' \rrbracket \parallel q \llbracket Q' \rrbracket \parallel N}$$

By Lemma 4(2) $P = \mathcal{E}[\overline{c}(\bigoplus_{i \in I} \pi_i; P_i)]$, where $\pi_j = p!A$ for some $j \in I$, and $P' = \mathcal{E}[\overline{c}(\bigoplus_{i \in I \setminus \{j\}} \pi_i; P_i \oplus \widehat{\pi}_j; P_j)]$. By Lemma 4(1) $Q = \mathcal{E}'[\Delta(\sum_{h \in H} \pi'_h; Q_h)]$, where $\pi'_k = p?A$ for some $k \in H$, and $Q' = \mathcal{E}'[\Delta(\sum_{h \in H \setminus \{k\}} \pi'_h; Q_h + \widehat{\pi}'_k; Q_k)]$.

By Lemma 1(6) $\vdash P : T$ and $\vdash Q : S$ and $T \leq G \upharpoonright p$ and $S \leq G \upharpoonright q$. By Lemma 5(3) $T = T[T']$ and $\vdash \mathcal{E} : T$ and $\vdash \overline{c}(\bigoplus_{i \in I} \pi_i; P_i : T'$ and $S = T'[S']$ and $\vdash \mathcal{E}' : T'$ and $\vdash \Delta(\sum_{h \in H} \pi'_h; Q_h : S'$. By Lemma 1(2) and (1) $T' = \overline{c}(\bigvee_{i \in I} \pi_i; T_i$ and $S' = \Delta(\bigwedge_{h \in H} \pi'_h; S_h$.

By Lemma 5(1) $G \upharpoonright p = T_0[T']$ and $T \leq T_0$ and $\overline{c}(\bigvee_{i \in I} \pi_i; T_i \leq T''$ and $G \upharpoonright q = T'_0[S']$ and $T' \leq T'_0$ and $\Delta(\bigwedge_{h \in H} \pi'_h; S_h \leq S''$. By definition of \leq we get $T'' = \overline{c}(\bigvee_{i \in I} \pi_i; T'_i$ with $T_i \leq T'_i$ for $i \in I$ and $S'' = \Delta(\bigwedge_{h \in H'} \pi'_h; S'_h$ with $H \supseteq H'$ and $S_h \leq S'_h$ for $h \in H'$.

From $G \upharpoonright p = T_0[\overline{c}(\bigvee_{i \in I} \pi_i; T'_i)]$ it follows by Lemma 6(1) that there is a unique path σ of G such that $SubT(G, \sigma) = \overline{c}(\bigvee_{i \in I} \alpha_i^p; G_i$ affects p and $Ctx(G, \sigma) \upharpoonright p = T_0$ and $SubT(G, \sigma) \upharpoonright q = \overline{c}(\bigvee_{i \in I} \pi_i; T'_i$. Since $\pi_j = q!A$, we have $\alpha_j^p = p \xrightarrow{A} q$. Consider now the projection $SubT(G, \sigma) \upharpoonright q = (\overline{c}(\bigvee_{i \in I} \alpha_i^p; G_i) \upharpoonright q$. It must be $SubT(G, \sigma) \upharpoonright q = \overline{c}(\bigvee_{i \in I} \alpha_i^p; G_i) \upharpoonright q \doteq \Delta(\bigwedge_{h \in H' \setminus \{k\}} \pi'_h; S'_h \wedge \pi'_k; S'_k) = S''$, where $\gamma \in \Delta$ and $\pi'_k = p?A$.

By definition we have $G = Ctx(G, \sigma)[SubT(G, \sigma)]$. Then we can choose $G' = Ctx(G, \sigma)[\overline{c}(\bigvee_{i \in I \setminus \{j\}} \alpha_i^p; K_i \square \alpha_j^p; K_j)]$. Indeed $G' \upharpoonright p = T_0[T_0]$ and $G' \upharpoonright q = T'_0[S_0]$, where $T_0 = \overline{c}(\bigvee_{i \in I \setminus \{j\}} \pi_i; T_i \vee \widehat{\pi}_j; T_j)$ and $S_0 = \Delta(\bigwedge_{h \in H' \setminus \{k\}} \pi_h; S_h \wedge \widehat{\pi}'_k; S'_k)$, while $G' \upharpoonright r = G \upharpoonright r$ for $r \neq p, q$. By Lemma 5(4) $\vdash P' : T[T_0]$ and $\vdash Q' : T'[S_0]$. Since $T[T_0] \leq G' \upharpoonright p$ and $T'[S_0] \leq G' \upharpoonright q$ we can derive $\vdash p \ll P' \parallel \parallel q \ll Q' \parallel \parallel N : G'$. By Lemma 7(3) $\sigma j \alpha_j^p$ is enabled in G , being both $Ctx(G, \sigma) \upharpoonright p$ and $Ctx(G, \sigma) \upharpoonright q$ session contexts. This implies the causal correctness of G' . The projectability of G' is immediate.

Let now the applied reduction rule be [BACK], then:

$$P \xrightarrow{\overline{C}} P' \quad P_h \xrightarrow{C} P'_h \quad h \in H \quad P_k \not\xrightarrow{C} \quad k \in K$$

$$p \ll P \parallel \parallel \Pi_{h \in H} p_h \ll P_h \parallel \parallel \Pi_{k \in K} p_k \ll P_k \parallel \parallel \xrightarrow{\overline{C}} p \ll P' \parallel \parallel \Pi_{h \in H} p_h \ll P'_h \parallel \parallel \Pi_{k \in K} p_k \ll P_k \parallel \parallel$$

By Lemma 4(3) $P = \mathcal{E}[\overline{c}(\bigoplus_{i \in I} \pi_i; Q_i \oplus \widehat{\pi}; Q)]$ and $Q \downarrow_{out}$ and $I \neq \emptyset$ and \mathcal{E} ok for \overline{C} and $P' = \mathcal{E}[\overline{c}(\bigoplus_{i \in I} \pi_i; Q_i)]$. By Lemma 4(4) $P_h = \mathcal{E}_h[S_h]$, and $P'_h = \mathcal{E}_h[\overline{c}S_h]$, where $S_h = \Delta_h \sum_{\ell \in L_h} R_\ell$ and $C \in \Delta_h$ and \mathcal{E}_h ok for C with $h \in H$. By Lemma 4(5), in a P_k with $k \in K$ there cannot be executed external choices whose checkpoint labels contain C .

By Lemma 1(6) $\vdash P : T$ for some $T \leq G \upharpoonright p$, and $\vdash P_j : T_j$ for some $T_j \leq G \upharpoonright p_j$ for each $j \in H \cup K$. By Lemma 1(1), in a T_k with $k \in K$ there cannot be executed intersections whose checkpoint labels contain C .

By Lemma 5(3) $T = T[T']$, where $\vdash \mathcal{E} : T$ and T ok for \overline{C} and

$$\vdash \overline{c}(\bigoplus_{i \in I} \pi_i; Q_i \oplus \widehat{\pi}; Q) : T'$$

and $T_h = T_h[S_h]$, where $\vdash \mathcal{E}_h : T_h$ and T_h ok for C and $\vdash S_h : S_h$ for each $h \in H$. By Lemma 1(2) $T' = \overline{c}(\bigvee_{i \in I} \pi_i; T_i \vee \widehat{\pi}; T_Q)$ and by Lemma 1(1) S_h is an intersection with Δ_h as set of checkpoint labels ($h \in H$).

By Lemma 5(1) $G \upharpoonright p = T'[T']$ and $T \leq T'$ and T' ok for \overline{C} and

$$\overline{c} \left(\bigvee_{i \in I} \pi_i; T_i \vee \widehat{\pi}; T_Q \right) \leq T''$$

and $G \upharpoonright p_h = T'_h[S'_h]$ and $T_h \leq T'_h$ and T'_h ok for C and $S_h \leq S'_h$ for $h \in H$. By definition of \leq we get $T'' = \overline{c}(\bigvee_{i \in I} \pi_i; T'_i \vee \widehat{\pi}; T'_Q)$ with $T_i \leq T'_i$ for $i \in I$ and $T_Q \leq T'_Q$ and S'_h is an intersection with Δ_h as set of checkpoint labels ($h \in H$).

By Lemma 6(1), from $G \upharpoonright p = T'[\overline{c}(\bigvee_{i \in I} \pi_i; T'_i \widehat{\vee} \widehat{\pi}; T'_Q)]$ it follows that there is a unique path σ of G such that $SubT(G, \sigma) = C(\boxplus_{i \in I} \alpha_i^p; K_i \boxplus \widehat{\alpha}^p; G_Q)$ affects p and $Ctx(G, \sigma) \upharpoonright p = T'$ and $SubT(G, \sigma) \upharpoonright p = \overline{c}(\bigvee_{i \in I} \pi_i; T'_i \widehat{\vee} \widehat{\pi}; T'_Q)$. The conditions T' ok for \overline{C} and T_h ok for C ($h \in H$) and T_k without executed intersections whose checkpoint labels contain C ($k \in K$) assure that the path σ is ok for C , i.e. that $Ctx(G, \sigma)$ does not contain occurrences of C . Therefore the occurrence of C in $SubT(G, \sigma)$ is the outermost one and it is projected on the checkpoint labels of T' and S_h for $h \in H$. Then we can choose $G' = Ctx(G, \sigma)[C(\boxplus_{i \in I} \alpha_i^p; K_i)]$. In fact $G' \upharpoonright p = T'[\overline{c}(\bigvee_{i \in I} \pi_i; T'_i)]$ and $G' \upharpoonright p_h = T'_h[\widehat{S}'_h]$ for $h \in H$ and $G' \upharpoonright p_k = G \upharpoonright p_k$ for $k \in K$. We can derive $\vdash P' : T[\overline{c}(\bigvee_{i \in I} \pi_i; T_i)]$ by Lemma 5(4) and $\vdash P'_h : T_h[\widehat{S}'_h]$ for $h \in H$ by Lemma 3. Since $T[\overline{c}(\bigvee_{i \in I} \pi_i; T_i)] \leq G' \upharpoonright p$ and $T_h[\widehat{S}'_h] \leq G' \upharpoonright p_h$ for $h \in H$ we may conclude that $\vdash p \llbracket P' \rrbracket \parallel \Pi_{h \in H} p_h \llbracket P'_h \rrbracket \parallel \Pi_{k \in K} p_k \llbracket P_k \rrbracket : G'$. The causal correctness of G implies the causal correctness of G' since G' is obtained from G by erasing one branch in a choice. The projectability of G' is immediate. \square

A standard property enforced by session types is session fidelity: all communications occur as specified by global types. In our case this applies also to backward reductions. The proof relies on the proof of subject reduction.

Theorem 2 (Session fidelity) *Let $\vdash N : G$.*

1. *If $\sigma p \xrightarrow{\Delta} q$ is enabled in G , then $N \xrightarrow{p \Delta q} N'$.*
2. *If $N \xrightarrow{p \Delta q} N'$, then $\sigma p \xrightarrow{\Delta} q$ is enabled in G for some σ .*
3. *If C is enabled in G , then $N \overset{C}{\curvearrowright} N'$.*
4. *If $N \overset{C}{\curvearrowright} N'$, then C is enabled in G .*

Proof (1) Let $\sigma = \sigma' j$. Since $\sigma' j p \xrightarrow{\Delta} q$ is enabled in G , from Lemma 7(1) for some session contexts T and T' we have $T = Ctx(G, \sigma') \upharpoonright p$ and $T' = Ctx(G, \sigma') \upharpoonright q$. Therefore the projections of G on p and q may be written as $G \upharpoonright p = T[\overline{y}(\bigvee_{i \in I \setminus \{j\}} \pi_i; T_i \vee q! \Delta; T_j)]$ and $G \upharpoonright q = T'[\Delta(\bigwedge_{h \in H \setminus \{k\}} \pi'_h; S_h \wedge p? \Delta; S_k)]$.

By Lemma 2(6) $N = p \llbracket P \rrbracket \parallel q \llbracket Q \rrbracket \parallel N''$ and $\vdash P : T$ and $\vdash Q : S$ and $T \leq G \upharpoonright p$ and $S \leq G \upharpoonright q$. By Lemma 5(2) and the definition of \leq we have :

- $T = T_1[\overline{y}(\bigvee_{i \in I} \pi_i; T'_i)]$ where $\pi_j = q! \Delta$ with $j \in I$ and $T'_i \leq T_i$ ($i \in I$) and $T \leq T_1$
- $S = T_2[\Delta(\bigwedge_{h \in H'} \pi'_h; S'_h)]$ where $\pi_k = p? \Delta$ with $k \in H$ and $H' \supseteq H$ and $S'_h \leq S_h$ ($h \in H$) and $T' \leq T_2$.

By Lemma 5(5) we get $P = \mathcal{E}[P']$ and $\vdash \mathcal{E} : T_1$ and $\vdash P' : \overline{y}(\bigvee_{i \in I} \pi_i; T'_i)$ and $Q = \mathcal{E}'[Q']$ and $\vdash \mathcal{E}' : T_2$ and $\vdash Q' : \Delta(\bigwedge_{h \in H'} \pi'_h; S'_h)$. By Lemma 2(2) $P' = \overline{y}(\bigoplus_{i \in I} \pi_i; P_i)$ and by Lemma 2(1) $Q' = \Delta(\sum_{h \in H'} \pi'_h; Q_h)$. Therefore

$$P \xrightarrow{q! \Delta} \mathcal{E}[P''] \text{ where } P'' = \overline{y}(\bigoplus_{i \in I \setminus \{j\}} \pi_i; P_i \oplus \widehat{q! \Delta}; P_j)$$

and

$$Q \xrightarrow{p? \Delta} \mathcal{E}'[Q''] \text{ where } Q'' = \Delta(\sum_{h \in H' \setminus \{k\}} \pi'_h; Q_h + \widehat{p? \Delta}; Q_k),$$

which imply $N \xrightarrow{p \Delta q} p \llbracket \mathcal{E}[P''] \rrbracket \parallel q \llbracket \mathcal{E}'[Q''] \rrbracket \parallel N''$.

(2) By Lemma 4(6) $\mathbb{N} = \mathbf{p}[\![P]\!] \parallel \mathbf{q}[\![Q]\!] \parallel \mathbb{N}_0$ and $P \xrightarrow{\mathbf{q}!\Delta} P'$ and $Q \xrightarrow{\mathbf{p}!\Delta} Q'$. Therefore rule [COM] has been applied and from the proof of Theorem 1 we know that $\sigma \mathbf{p} \xrightarrow{\Delta} \mathbf{q}$ is enabled in \mathbb{G} for some σ .

(3) From Definition 17(4) we know that there exists a path σ ok for C such that $SubT(\mathbb{G}, \sigma) = \mathcal{C}(\boxplus_{j \in J \setminus \{k\}} \alpha_j^{\mathbf{p}}; K_j \boxplus \alpha_k^{\mathbf{p}}; \mathbb{G}_k)$. Since $\mathbb{G} = Ctx(\mathbb{G}, \sigma)[SubT(\mathbb{G}, \sigma)]$ is causally correct all the communications involving \mathbf{p} in $Ctx(\mathbb{G}, \sigma)$ are executed communications. This implies that $\mathcal{T} = Ctx(\mathbb{G}, \sigma) \upharpoonright \mathbf{p}$ is a session context and thus the projection of \mathbb{G} on \mathbf{p} may be written as $\mathbb{G} \upharpoonright \mathbf{p} = \mathcal{T}[\overline{\mathcal{V}}(\bigvee_{j \in J \setminus \{k\}} \pi_j; T_j \vee \widehat{\pi}_k; \mathbb{T})]$. Therefore \mathcal{T} is ok for C and $(\alpha_j^{\mathbf{p}}; K_j) \upharpoonright \mathbf{p} = \pi_j; T_j$ for $j \in J \setminus \{k\}$ and $(\alpha_k^{\mathbf{p}}; \mathbb{G}_k) \upharpoonright \mathbf{p} = \widehat{\pi}_k; \mathbb{T}$. By Lemma 2(6) $\mathbb{N} = \mathbf{p}[\![P]\!] \parallel \mathbb{N}'$ and $\vdash P : \mathcal{T}'$ and $\mathcal{T}' \leq \mathbb{G} \upharpoonright \mathbf{p}$. By Lemma 5(2) $\mathcal{T}' = \mathcal{T}'[\overline{\mathcal{V}}(\bigvee_{j \in J \setminus \{k\}} S_j \vee \widehat{\pi}_k; \mathbb{S})]$ and $\mathcal{T}' \leq \mathcal{T}$ and $\pi_j; S_j \leq \pi_j; T_j$ for $j \in J \setminus \{k\}$ and $\widehat{\pi}_k; \mathbb{S} \leq \widehat{\pi}_k; \mathbb{T}$.

By Lemma 5(5) $P = \mathcal{E}[P']$ and $\vdash \mathcal{E} : \mathcal{T}'$ and $\vdash P' : \overline{\mathcal{V}}(\bigvee_{j \in J \setminus \{k\}} \pi_j; S_j \vee \widehat{\pi}_k; \mathbb{S})$. By Lemma 2(2) $P' = \mathcal{E}[\overline{\mathcal{V}}(\bigoplus_{j \in J \setminus \{k\}} \pi_j; Q_j \oplus \widehat{\pi}_k; \mathbb{Q})]$ and $\vdash \widehat{\pi}_k; \mathbb{Q} : \widehat{\pi}_k; \mathbb{S}$. Let σ' and h be such that $\sigma k \sigma' h \alpha^{\mathbf{p}} \in Occ(\mathbb{G})$ is enabled in \mathbb{G} . Then $SubT(\mathbb{G}, \sigma k \sigma') = \overline{\mathcal{V}}(\boxplus_{i \in I} \alpha_i^{\mathbf{p}}; K_i$ with $h \in I$ and $\alpha_h^{\mathbf{p}} = \alpha^{\mathbf{p}}$ and for some \mathcal{T}'' we have that $\mathbb{T} = \mathcal{T}''[\overline{\mathcal{V}}(\bigvee_{i \in I} T'_i)]$. By Lemma 5(2) $\mathbb{S} = \mathcal{T}'''[\overline{\mathcal{V}}(\bigvee_{i \in I} S'_i)]$ for some \mathcal{T}''' . By Lemma 2(2) $Q = \mathcal{E}'[\overline{\mathcal{V}}(\bigoplus_{i \in I} Q'_i)]$, which implies $Q \downarrow_{out}$. Then $P \xrightarrow{\overline{\mathcal{C}}} P'$ and $\mathbb{N} \xrightarrow{\overline{\mathcal{C}}} \mathbb{N}'$, since all processes in \mathbb{N} may contain or may not contain executed internal choices labelled C .

(4) By Lemma 4(7) $\mathbb{N} = \mathbf{p}[\![P]\!] \parallel \prod_{h \in H} \mathbf{p}_h[\![P_h]\!] \parallel \prod_{k \in K} \mathbf{p}_k[\![P_k]\!] \parallel P \xrightarrow{\overline{\mathcal{C}}} P'$ and $P_h \xrightarrow{\overline{\mathcal{C}}} P'_h$ for all $h \in H$ and $P_k \not\xrightarrow{\overline{\mathcal{C}}}$ for all $k \in K$ and

$$\mathbb{N}' = \mathbf{p}[\![P']\!] \parallel \prod_{h \in H} \mathbf{p}_h[\![P'_h]\!] \parallel \prod_{k \in K} \mathbf{p}_k[\![P_k]\!]$$

Therefore rule [BACK] has been applied and from the proof of Theorem 1 we know that:

- $P = \mathcal{E}[\overline{\mathcal{C}}(\bigoplus_{i \in I} \pi_i; Q_i \oplus \widehat{\pi}; \mathbb{Q})]$ and $Q \downarrow_{out}$ and $I \neq \emptyset$ and \mathcal{E} ok for C
- $SubT(\mathbb{G}, \sigma) = \mathcal{C}(\boxplus_{i \in I} \alpha_i^{\mathbf{p}}; K_i \boxplus \alpha^{\mathbf{p}}; \mathbb{G}_Q)$ and σ is ok for C
- $\vdash P : \mathcal{T}$ and $\mathcal{T} \leq \mathbb{G} \upharpoonright \mathbf{p}$ and $\vdash Q : \mathcal{T}_Q$ and $\mathcal{T}_Q \leq \mathbb{G}_Q \upharpoonright \mathbf{p}$
- $\mathbb{G} \upharpoonright \mathbf{p} = \mathcal{T}'[\overline{\mathcal{C}}(\bigvee_{i \in I} \pi_i; T'_i \vee \widehat{\pi}; T'_Q)]$ and $\mathcal{T}_Q \leq \mathcal{T}'_Q$.

Let $\alpha^{\mathbf{p}}; \mathbb{G}_Q$ be the ℓ th branch in $SubT(\mathbb{G}, \sigma)$. By definition $Q \downarrow_{out}$ implies $Q = \mathcal{E}'[\overline{\mathcal{V}}(\bigoplus_{j \in I'} Q'_j)]$.

From $\vdash Q : \mathcal{T}_Q$ and Lemma 5(3) $\mathcal{T}_Q = \mathcal{T}_Q[S_Q]$ and $\vdash \overline{\mathcal{V}}(\bigoplus_{j \in I'} Q'_j : S_Q)$. By Lemma 1(2) $S_Q = \overline{\mathcal{V}}(\bigvee_{j \in I'} S_j)$, which together with $\mathcal{T}_Q \leq \mathcal{T}'_Q$ imply $\mathcal{T}'_Q = \mathcal{T}'_Q[\overline{\mathcal{V}}(\bigvee_{j \in I'} S'_j)]$ by Lemma 5(1). Let $\mathcal{T}'' = \mathcal{T}'[\overline{\mathcal{C}}(\bigvee_{i \in I} \pi_i; T'_i \vee \widehat{\pi}; T'_Q)]$, then $\mathbb{G} \upharpoonright \mathbf{p} = \mathcal{T}''[\overline{\mathcal{V}}(\bigvee_{j \in I'} S'_j)]$. By Lemma 6(1) there is σ' such that $SubT(\mathbb{G}, \sigma \ell \sigma') = \overline{\mathcal{V}}(\boxplus_{j \in I'} \beta_j^{\mathbf{p}}; K_j$ affects \mathbf{p} and $Ctx(\mathbb{G}, \sigma \ell \sigma') \upharpoonright \mathbf{p} = \mathcal{T}''$ and $SubT(\mathbb{G}, \sigma \ell \sigma') \upharpoonright \mathbf{p} = \overline{\mathcal{V}}(\bigvee_{j \in I'} S'_j)$. From Lemma 7(2) $\sigma \ell \sigma' j \beta_j^{\mathbf{p}}$ is alive in \mathbb{G} for all $j \in I'$ and then C is alive in \mathbb{G} . We finally conclude that C is enabled in \mathbb{G} , since $Ctx(\mathbb{G}, \sigma \ell \sigma') \upharpoonright \mathbf{p}$ is a session context, and this implies that $\sigma \ell \sigma' j \beta_j^{\mathbf{p}}$ is enabled in \mathbb{G} for all $j \in I'$. □

The remainder of this section is devoted to the proof of forward and backward progress.

For forward progress we first show that a communication which is alive in a global type has at least one enabled cause, and will eventually become enabled.

- Lemma 8** 1. If $\xi \in \text{Occ}(\mathbb{G})$ is alive in \mathbb{G} , then either ξ is enabled in \mathbb{G} or there is some occurrence in $\text{Causes}(\mathbb{G}, \xi)$ which is enabled in \mathbb{G} .
2. If $\vdash \mathbb{N} : \mathbb{G}$ and $\xi \in \text{Occ}(\mathbb{G})$ is alive in \mathbb{G} , then $\mathbb{N} \longrightarrow^* \mathbb{N}'$ with $\vdash \mathbb{N}' : \mathbb{G}'$ and ξ is enabled in \mathbb{G}' .

Proof (1) By induction on the cardinality of the set of non executed occurrences in $\text{Causes}(\mathbb{G}, \xi)$, namely on $n = |\text{Causes}(\mathbb{G}, \xi) \setminus \text{ExOcc}(\mathbb{G})|$. If $n = 0$, then ξ is enabled in \mathbb{G} by definition. Let $n > 0$. Note that every $\xi' \in \text{Causes}(\mathbb{G}, \xi) \setminus \text{ExOcc}(\mathbb{G})$ must be alive, because otherwise there would exist an executed ξ'' such that $\xi'' \#_{\mathbb{G}} \xi'$, which by conflict heredity would imply $\xi'' \#_{\mathbb{G}} \xi$, contradicting the assumption that ξ is alive. So, let us choose a non executed $\xi' \in \text{Causes}(\mathbb{G}, \xi)$. If ξ' is enabled we are done, otherwise, since ξ' is alive and $\text{Causes}(\mathbb{G}, \xi') \subset \text{Causes}(\mathbb{G}, \xi)$, by inductive hypothesis there is ξ'' in $\text{Causes}(\mathbb{G}, \xi')$ which is enabled in \mathbb{G} .

(2) Again, we proceed by induction on $n = |\text{Causes}(\mathbb{G}, \xi) \setminus \text{ExOcc}(\mathbb{G})|$. If $n = 0$, then ξ is enabled in \mathbb{G} and the result is immediate by Theorem 2(1). Let now $n > 0$. By (1) there exists ξ' in $\text{Causes}(\mathbb{G}, \xi)$ which is enabled in \mathbb{G} . If $\xi' = \sigma \mathfrak{p} \xrightarrow{\Delta} \mathfrak{q}$, by Theorem 2(1) we have $\mathbb{N} \xrightarrow{\mathfrak{p}\Delta\mathfrak{q}} \mathbb{N}''$ and then by Theorem 1 we have $\vdash \mathbb{N}'' : \mathbb{G}''$. Let $\xi'' = \sigma \mathfrak{p} \xrightarrow{\Delta} \mathfrak{q}$. Since $\text{Causes}(\mathbb{G}'', \xi) = \text{Causes}(\mathbb{G}, \xi) \setminus \{\xi'\} \cup \{\xi''\}$ and $\text{ExOcc}(\mathbb{G}'') = \text{ExOcc}(\mathbb{G}) \cup \xi''$, we have $\text{Causes}(\mathbb{G}'', \xi) \setminus \text{ExOcc}(\mathbb{G}'') \subset \text{Causes}(\mathbb{G}, \xi) \setminus \text{ExOcc}(\mathbb{G})$. Then by induction $\mathbb{N}'' \longrightarrow^* \mathbb{N}'$ (whence $\mathbb{N} \longrightarrow^* \mathbb{N}'$) with $\vdash \mathbb{N}' : \mathbb{G}'$ and ξ is enabled in \mathbb{G}' . \square

Connecting communications cause the failure of the standard progress property, since a participant offering an external choice between connecting communications can wait forever. Instead, processes offering outputs or simple inputs can always communicate.

Theorem 3 (Forward progress)

- If $\vdash \mathfrak{p} \llbracket \mathcal{E}[\overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i] \rrbracket \parallel \mathbb{N} : \mathbb{G}$, then for every $j \in I$ there exists an \mathbb{N}_j such that $\mathfrak{p} \llbracket \mathcal{E}[\overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i] \rrbracket \parallel \mathbb{N} \longrightarrow^* \mathfrak{p} \llbracket \mathcal{E}[\overline{\gamma} (\bigoplus_{i \in I \setminus \{j\}} \pi_i; P_i \oplus \widehat{\pi}_j; P_j)] \rrbracket \parallel \mathbb{N}_j$.
- If $\vdash \mathfrak{p} \llbracket \mathcal{E}[\Delta \sum_{i \in I} \pi_i; P_i] \rrbracket \parallel \mathbb{N} : \mathbb{G}$ and all π_i are simple inputs, then there is \mathbb{N}' such that $\mathfrak{p} \llbracket \mathcal{E}[\Delta \sum_{i \in I} \pi_i; P_i] \rrbracket \parallel \mathbb{N} \longrightarrow^* \mathfrak{p} \llbracket \mathcal{E}[\Delta (\sum_{i \in I \setminus \{j\}} \pi_i; P_i + \widehat{\pi}_j; P_j)] \rrbracket \parallel \mathbb{N}'$ for some $j \in I$.

Proof (1) Let $P = \mathcal{E}[\overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i]$, by Lemma 1(6) $\vdash P : \mathbb{T}$ and $\mathbb{T} \leq \mathbb{G} \upharpoonright \mathfrak{p}$. By Lemma 5(3) $\mathbb{T} = \mathbb{T}[\mathbb{T}']$ and $\vdash \mathcal{E} : \mathbb{T}$ and $\vdash \overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i : \mathbb{T}'$. By Lemma 1(2) $\mathbb{T}' = \overline{\gamma} \bigvee_{i \in I} \pi_i; T_i$. By Lemma 5(1) $\mathbb{G} \upharpoonright \mathfrak{p} = \mathbb{T}'[\mathbb{T}'']$ and $\mathbb{T} \leq \mathbb{T}'$ and $\overline{\gamma} \bigvee_{i \in I} \pi_i; T_i \leq \mathbb{T}''$. By definition of \leq we get $\mathbb{T}'' = \overline{\gamma} \bigvee_{i \in I} \pi_i; T'_i$ with $T_i \leq T'_i$ for $i \in I$. Therefore $\mathbb{G} \upharpoonright \mathfrak{p} = \mathbb{T}'[\overline{\gamma} \bigvee_{i \in I} \pi_i; T'_i]$.

From Lemma 6(1), for some σ we have that $\text{SubT}(\mathbb{G}, \sigma) = \gamma \boxtimes_{i \in I} \alpha_i^{\mathfrak{p}}; K_i$ and $\text{Ctx}(\mathbb{G}, \sigma) \upharpoonright \mathfrak{p} = \mathbb{T}'$ and $\text{SubT}(\mathbb{G}, \sigma) \upharpoonright \mathfrak{p} = \overline{\gamma} \bigvee_{i \in I} \pi_i; T'_i$. This implies, by Lemma 7(2), that $\sigma j \alpha_j^{\mathfrak{p}}$ is alive in \mathbb{G} . Therefore from Lemma 8(2) we get that

$$\mathfrak{p} \llbracket \mathcal{E}[\overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i] \rrbracket \parallel \mathbb{N} \longrightarrow^* \mathfrak{p} \llbracket \mathcal{E}[\overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i] \rrbracket \parallel \mathbb{N}'_j$$

with $\vdash \mathfrak{p} \llbracket \mathcal{E}[\overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i] \rrbracket \parallel \mathbb{N}'_j : \mathbb{G}'$ and $\sigma j \alpha_j^{\mathfrak{p}}$ enabled in \mathbb{G}' . From Theorem 2(1) we conclude $\mathfrak{p} \llbracket \mathcal{E}[\overline{\gamma} \bigoplus_{i \in I} \pi_i; P_i] \rrbracket \parallel \mathbb{N}'_j \xrightarrow{\alpha_j^{\mathfrak{p}}} \mathfrak{p} \llbracket \mathcal{E}[\overline{\gamma} (\bigoplus_{i \in I \setminus \{j\}} \pi_i; P_i \oplus \widehat{\pi}_j; P_j)] \rrbracket \parallel \mathbb{N}_j$.

(2) Let $P = \mathcal{E}[\Delta \sum_{i \in I} \pi_i; P_i]$, by Lemma 1(6) $\vdash P : \mathbb{T}$ and $\mathbb{T} \leq \mathbb{G} \upharpoonright \mathfrak{p}$. By Lemma 5(3) $\mathbb{T} = \mathbb{T}[\mathbb{T}']$ and $\vdash \mathcal{E} : \mathbb{T}$ and $\vdash \Delta \sum_{i \in I} \pi_i; P_i : \mathbb{T}'$. By Lemma 1(1) $\mathbb{T}' = \Delta \bigwedge_{i \in I} \pi_i; T_i$. By Lemma 5(1) $\mathbb{G} \upharpoonright \mathfrak{p} = \mathbb{T}'[\mathbb{T}'']$ and $\mathbb{T} \leq \mathbb{T}'$ and $\Delta \bigwedge_{i \in I} \pi_i; T_i \leq \mathbb{T}''$. By definition

of \leq we get and $T'' = \Delta \bigwedge_{i \in I'} \pi_i; T'_i$ with $I \supseteq I'$ and $T_i \leq T'_i$ for $i \in I'$. Therefore $G \upharpoonright p = T'[\Delta \bigwedge_{i \in I'} \pi_i; T'_i]$. From Lemma 6(2) we have that there is σ such that $SubT(G, \sigma)$ affects p and $Ctx(G, \sigma) \upharpoonright p = T'$ and $SubT(G, \sigma) \upharpoonright p \doteq \Delta \bigwedge_{i \in I'} \pi_i; T'_i$. For each $i \in I'$, the action π_i must have the form $\pi_i = q_i ? \lambda_i$ and be obtained projecting a communication $\alpha_i = q_i \xrightarrow{\lambda_i} p$. We consider two cases: either $SubT(G, \sigma)$ is an executed choice, or it is either a non executed choice or a sequence starting with a non executed communication.

- If $SubT(G, \sigma) = \gamma(\Box_{h \in H \setminus \{k\}} \beta_h; K_h \Box \widehat{\beta}_k; G_k)$, then let σ_k be the longest path in $SubT(G, \sigma)$ such that $k \sqsubseteq \sigma_k$ and $\sigma \sigma_k = \text{path}(\xi)$ for some $\xi \in ExOcc(G)$. By Lemma 6(2) there are σ' and $j \in I'$ such that $\sigma \sigma' \alpha_j \in Occ(G)$ and either $\sigma' \sqsubseteq \sigma_k$ or $\sigma_k \sqsubset \sigma'$. Since $\xi' \in ExOcc(G)$ implies $\text{path}(\xi') \sqsubseteq \sigma \sigma_k$, then ξ' is not in conflict with $\sigma \sigma' \alpha_j$. Therefore $\sigma \sigma' \alpha_j$ is alive in G .
- If $SubT(G, \sigma) = \gamma \Box_{h \in H} \beta_h; G_h$, then $\sigma h \beta_h$ is alive in G for all $h \in H$ by Lemma 7(2). Then by Lemma 6(2) for any $h \in H$ we can find σ_h such that $\sigma \sigma_h \alpha_j$ is alive in G for some $j \in I'$.

In both cases there is σ_p such that $\sigma_p \alpha_j$ is alive in G for some $j \in I'$. Lemma 8(2) implies

$$p \llbracket \mathcal{E}[\Delta \sum_{i \in I} \pi_i; P_i] \rrbracket \llbracket N \longrightarrow^* p \llbracket \mathcal{E}[\Delta \sum_{i \in I} \pi_i; P_i] \rrbracket \llbracket N'_j \rrbracket$$

and $\vdash p \llbracket \mathcal{E}[\Delta \sum_{i \in I} \pi_i; P_i] \rrbracket \llbracket N'_j \rrbracket : G'$ with $\sigma_p \alpha_j$ enabled in G' . From Theorem 2(1) we conclude $p \llbracket \mathcal{E}[\Delta \sum_{i \in I} \pi_i; P_i] \rrbracket \llbracket N'_j \rrbracket \xrightarrow{\alpha_j} p \llbracket \mathcal{E}[\Delta (\sum_{i \in I \setminus \{j\}} \pi_i; P_i \oplus \widehat{\pi}_j; P_j)] \rrbracket \llbracket N'_j \rrbracket$. \square

Notice that the standard formulation of progress [13], which requires that each *simple* input and each output that is persistently offered be eventually consumed, is an easy consequence of this theorem.

Theorem 4 (Backward progress) *If $\vdash N : G$ and C is alive in G , then there is N' such that $N \longrightarrow^* N'$ and $N' \overset{C}{\curvearrowright} N''$.*

Proof By definition $G = Ctx(G, \sigma)[C(\boxplus_{j \in J \setminus \{k\}} \alpha_j^P; K_j \boxplus \widehat{\alpha}_k^P; G_k)]$ and $|J| > 1$ and σ is ok for C (in G) and there are σ' and α^P such that $\sigma k \sigma' \alpha^P \in Occ(G)$ is alive in G . Lemma 8(2) implies that $N \longrightarrow^* N'$ and $\vdash N : G'$ with $\sigma k \sigma' \alpha^P$ enabled in G' .

This implies that C is enabled in G . We conclude $N' \overset{C}{\curvearrowright} N''$ by Theorem 2(3). \square

To sum up, our calculus enjoys:

- subject reduction for both forward and backward computations;
- session fidelity for enabled communications and for rollbacks to enabled checkpoint labels;
- forward progress, assuring the absence of dangling actions but for connecting inputs;
- backward progress for alive checkpoint labels.

We end this section with a remark on causal consistency and the loop lemma [10]. Causal consistency states that a cause cannot be reversed without first reversing its effects. Clearly, our calculus enjoys causal consistency since a rollback to a checkpointed choice removes all communications that were done after that checkpoint. The loop lemma prescribes that each transition has an inverse. The loop lemma does not hold for our calculus, since a reverse computation can only go back to a checkpointed choice whose leader currently offers an output.

6 Related work and conclusion

Since the seminal work by Danos and Krivine on reversible CCS [10], reversible computation has been widely studied in process calculi. In [31], Phillips and Ulidowski proposed a method for reversing process operators defined in a general SOS format, and noted that thread identifiers and histories were needed to record the past of computations. In [30] the authors use a cursor to mark past actions, thus avoiding the use of extra memory information as in [10].

In [23], Lanese et al. extended the approach of Danos and Krivine by defining a reversible variant of the higher-order π -calculus, using tags to identify threads, and explicit memory processes. This calculus was enriched with a fine-grained rollback primitive in [22]. In [9], Cristescu et al. proposed a causal semantic model for the reversible π -calculus.

Reversibility for structured communications was first studied in [11,12,21], where *transactions* with rollback and coordinated checkpoints were modelled in an extended CCS. More recently, reversibility has been incorporated into *contracts* [1,4] and *session calculi* [17,18]. In [2,3], the authors investigated the notions of compliance and sub-behaviour for contracts with checkpoints. While rollbacks are forgetful in [2], in [3] they are used as a strategy to achieve compliance: in this case, after a rollback a process cannot engage again in the previously explored branch, presumably unsuccessful. We used a similar idea for defining our rollback here.

Our backward mechanism is most closely related to the recent proposals [24–26,34,35]. Tiezzi and Yoshida [34] use tags and memories to allow full reversibility of binary sessions with delegation. In [35], two forms of reversibility are considered: either a session is completely reversed in a single backward step, or any intermediate state is restored. Mezzina and Pérez [24,25] use monitors as memories for reversing binary sessions. A key novelty of this work is the use of session types with present and past. In [26], this approach is generalised to multiparty sessions, asynchronous higher-order communications, and decoupled rollbacks.

In [28], Neykova and Yoshida provide an algorithm to analyse and extract causal dependencies from a given multiparty global type, and use it to ensure that communicating processes are safely recovered from consistent states in the presence of a failure. In [27], Mezzina and Tuosto propose a semantic control of reversibility: a computation along a branch is reversed according to the guards on the current configuration. A feature of [27] is that inputs are potentially irreversible actions, unless they appear within a loop.

Our work builds on our previous papers [8,14]. In [14] we introduced a multiparty session calculus in which choices could be labelled with checkpoints. Participants could revert to one of these checkpointed choices in order to make a different choice. Global types were used to control reversibility, and shown to enforce the properties of fidelity and progress (both forward and backward). In [8] we enriched the syntax for types and processes with parallel and sequential composition, and used a more compact representation for past communications and a more refined strategy for backward moves. The current work borrows ideas and techniques from both papers while sticking to a simpler syntax than [8], and improves on both of them by allowing connecting communications.

As regards dynamic participants, in [13,15] global types prescribe the behaviours of fixed roles, and each role includes an arbitrary number of participants, which can dynamically join and leave roles. Instead, in the Conversation Calculus [7,36] the protocols describe communications between participants which can dynamically join and leave conversations. The model based on conversation contexts and labelled message-passing primitives is quite different from multiparty session types and this makes it difficult to adapt their approach to our setting.

The paper [20] has been our inspiration for connecting messages, but our calculus is more permissive than the original one. In fact, in [20] a connecting message can be exchanged only between two participants that did not communicate before, unless they have been disconnected by an ad hoc primitive. This means for example that $p \xrightarrow{\lambda_1} q; q \xrightarrow{\lambda_2} p$ is not allowed. Another difference of [20] with respect to our calculus is that in a choice, the first message received by a participant must be a connecting message, which forbids for instance $p \xrightarrow{\lambda_1} q \boxplus p \xrightarrow{\lambda_2} r$. A further restriction is that all the inputs in an external choice must have the same sender. A final major difference is that communication in [20] is asynchronous.

The properties of our calculus are standard for reversible session calculi, but their proofs require some ingenuity due to the presence of connecting communications and to the specificity of our rollback mechanism. As argued already in Sect. 3, connecting communications may be used to avoid the use of multiple global types. They could also be useful for incorporating delegation into global types. For example, a seller could delegate a bank to receive a credit card number only when the client wants to buy an item, otherwise the bank is not involved in the transaction (so delegation would appear in one branch of a choice but not in the others). This is a topic we plan to investigate. As regards the conditions for reversing a computation, a current limitation of our work is that the starting points of rollbacks are statically determined. By contrast, these points are determined dynamically in [27], offering a more realistic solution. We plan to introduce similar runtime conditions for rollback in our calculus.

Finally, we would like to study the interpretation of global types into a model of Event Structures, for which reversible variants have already been proposed [16,32]. We plan to explore a reversible variant of *Flow Event Structures* [5], a model that has already been used to interpret CCS processes with past in [6].

Acknowledgements We would like to thank the anonymous referees for their helpful comments.

A: Reductions of networks and systems

It is easy to define a bijection between networks and initial systems:

$$S(\mathbb{N}) = \mathbb{N} \check{\boxtimes} \langle -, \mathbb{P}, \emptyset \rangle \quad \mathcal{N}(\mathbb{N} \check{\boxtimes} \langle -, \mathbb{P}, \emptyset \rangle) = \mathbb{N}$$

where \mathbb{P} is the set of participants in \mathbb{N} .

We establish now an operational correspondence between networks and their associated systems. More precisely we prove the following:

- *Communication preservation and reflection*: every network communication is simulated by a communication in the associated system and vice-versa;
- *Rollback preservation*: every network rollback is simulated by a sequence of backward moves in the associated system, initiating with a starting backward move and terminating with an ending backward move;
- *Rollback reflection*: every starting backward move in a system can be extended to a complete sequence of backward moves such that the resulting system is the image of a rollback in the source network.

In the following theorem we denote by $\overset{C}{\curvearrowright}$ the transitive closure of $\overset{y_C}{\curvearrowright}$ and $\overset{n_C}{\curvearrowright}$.

Theorem 5 1. *If $\mathbb{N} \xrightarrow{p \Delta q} \mathbb{N}'$, then $S(\mathbb{N}) \xrightarrow{p \Delta q} S(\mathbb{N}')$.*

2. If $\mathcal{S}(\mathbb{N}) \xrightarrow{p\Delta q} \mathcal{S}(\mathbb{N}')$, then $\mathbb{N} \xrightarrow{p\Delta q} \mathbb{N}'$.
3. If $\mathbb{N} \xrightarrow{C} \mathbb{N}'$, then $\mathcal{S}(\mathbb{N}) \xrightarrow{\text{sc}, \text{cc}, \text{ec}} \mathcal{S}(\mathbb{N}')$.
4. If $\mathbb{N} \checkmark \langle -, \mathbb{P}, \emptyset \rangle \xrightarrow{\text{sc}} \mathbb{N}' \checkmark \langle C, \mathbb{P}_1, \mathbb{P}_2 \rangle$, then $\mathbb{N}' \checkmark \langle C, \mathbb{P}_1, \mathbb{P}_2 \rangle \xrightarrow{\text{cc}, \text{ec}} \mathbb{N}'' \checkmark \langle -, \mathbb{P}, \emptyset \rangle$ and $\mathbb{N} \xrightarrow{C} \mathbb{N}''$.

Proof (1) and (2). The result is immediate since Rule [COM] of Fig. 1 and Rule [COMS] of Fig. 2 have the same antecedents.

(3). If $\mathbb{N} \xrightarrow{C} \mathbb{N}'$, then the applied rule is

$$\frac{P \xrightarrow{\bar{C}} P' \quad P_i \xrightarrow{C} P'_i \quad 1 \leq i \leq m \quad P_i \not\xrightarrow{C} \quad m+1 \leq i \leq n}{\mathbb{N} \xrightarrow{C} \mathbb{N}'} \text{[BACK]}$$

where

$$\mathbb{N} = p[[P]] \parallel \prod_{1 \leq i \leq n} p_i[[P_i]] \quad \text{and} \quad \mathbb{N}' = p[[P']] \parallel \prod_{1 \leq i \leq m} p_i[[P'_i]] \parallel \prod_{m+1 \leq i \leq n} p_i[[P_i]]$$

Then $\mathcal{S}(\mathbb{N}) = \mathbb{N} \checkmark \sigma$ and $\mathcal{S}(\mathbb{N}') = \mathbb{N}' \checkmark \sigma$ with $\sigma = \langle -, \{p\} \cup \{p_i \mid 1 \leq i \leq n\}, \emptyset \rangle$. We get

$$\begin{aligned} \mathbb{N} \checkmark \sigma &\xrightarrow{\text{sc}} p[[P']] \parallel \prod_{1 \leq i \leq n} p_i[[P_i]] \checkmark \langle C, \{p\}, \{p_i \mid 1 \leq i \leq n\} \rangle \\ &\xrightarrow{\text{cc}} p[[P']] \parallel p_1[[P'_1]] \parallel \prod_{2 \leq i \leq n} p_i[[P_i]] \checkmark \langle C, \{p, p_1\}, \{p_i \mid 2 \leq i \leq n\} \rangle \\ &\xrightarrow{\text{cc}} \dots \\ &\xrightarrow{\text{cc}} \mathbb{N}' \checkmark \langle C, \{p\} \cup \{p_i \mid 1 \leq i \leq m\}, \{p_i \mid m+1 \leq i \leq n\} \rangle \\ &\xrightarrow{\text{nc}} \mathbb{N}' \checkmark \langle C, \{p\} \cup \{p_i \mid 1 \leq i \leq m+1\}, \{p_i \mid m+2 \leq i \leq n\} \rangle \\ &\xrightarrow{\text{nc}} \dots \\ &\xrightarrow{\text{nc}} \mathbb{N}' \checkmark \langle C, \{p\} \cup \{p_i \mid 1 \leq i \leq n\}, \emptyset \rangle \\ &\xrightarrow{\text{ec}} \mathbb{N}' \checkmark \sigma \end{aligned}$$

(4). If $\mathbb{N} \checkmark \langle -, \mathbb{P}, \emptyset \rangle \xrightarrow{\text{sc}} \mathbb{N}' \checkmark \langle C, \mathbb{P}_1, \mathbb{P}_2 \rangle$, then Rule [BACKS] has been applied, so

$$P \xrightarrow{\bar{C}} P' \quad \mathbb{N} = p[[P]] \parallel \mathbb{N}_1 \quad \mathbb{N}' = p[[P']] \parallel \mathbb{N}_1 \quad \mathbb{P}_1 = \{p\} \quad \mathbb{P}_2 = \mathbb{P} \setminus \{p\}$$

Since \mathbb{P}_2 can be split in two subsets of participants, according to whether the associated processes satisfy the premise of Rule [BACKY] or of Rule [BACKN] in Fig. 2, we may assume without loss of generality that

$$\mathbb{N}_1 = \prod_{1 \leq i \leq m} p_i[[P_i]] \parallel \prod_{m+1 \leq i \leq n} p_i[[P_i]]$$

where for all i , $1 \leq i \leq n$, $P_i \xrightarrow{C} P'_i$ if $i \leq m$ and $P_i \not\xrightarrow{C}$ otherwise. With a sequence of reductions as in the proof of point (3), we then obtain for some network \mathbb{N}'' :

$$\mathbb{N}' \checkmark \langle C, \{p\}, \mathbb{P} \setminus \{p\} \rangle \xrightarrow{\text{cc}, \text{ec}} \mathbb{N}'' \checkmark \langle -, \mathbb{P}, \emptyset \rangle.$$

Since P and all the P_i for $1 \leq i \leq n$ satisfy the premises of Rule [BACK] in Fig. 1, we may apply this rule to \mathbb{N} to conclude $\mathbb{N} \xrightarrow{C} \mathbb{N}''$, as required. □

References

1. Barbanera, F., de' Liguoro, U.: Sub-behaviour relations for session-based client/server systems. *Math. Struct. Comput. Sci.* **25**(6), 1339–1381 (2015)
2. Barbanera, F., Dezani-Ciancaglini, M., de' Liguoro, U.: Reversible client/server interactions. *Form. Asp. Comput.* **28**(4), 697–722 (2016)
3. Barbanera, F., Dezani-Ciancaglini, M., Lanese, I., de' Liguoro, U.: Retractable contracts. In: PLACES, volume 203 of EPTCS, pp. 61–72 (2016)
4. Bernardi, G., Hennessy, M.: Modelling session types using contracts. *Math. Struct. Comput. Sci.* **26**(3), 510–560 (2016)
5. Boudol, G., Castellani, I.: Permutation of transitions: an event structure semantics for CCS and SCCS. In: REX School, volume 354 of LNCS, pp. 411–427. Springer (1988)
6. Boudol, G., Castellani, I.: Flow models of distributed computations: three equivalent semantics for CCS. *Inf. Comput.* **114**(2), 247–314 (1994)
7. Caires, L., Vieira, H.T.: Analysis of service oriented software systems with the conversation calculus. In: FACS, volume 6921 of LNCS, pp. 6–33. Springer (2010)
8. Castellani, I., Dezani-Ciancaglini, M., Giannini, P.: Concurrent reversible sessions. In: CONCUR, volume 85 of LIPIcs, pp. 30:1–30:17. Schloss Dagstuhl (2017)
9. Cristescu, I., Krivine, J., Varacca, D.: Rigid families for the reversible π -calculus. In: RC, volume 9720 of LNCS, pp. 3–19. Springer (2016)
10. Danos, V., Krivine, J.: Reversible communicating systems. In: CONCUR, volume 3170 of LNCS, pp. 292–307. Springer (2004)
11. de Vries, E., Koutavas, V., Hennessy, M.: Communicating transactions—(extended abstract). In: CONCUR, volume 6269 of LNCS, pp. 569–583. Springer (2010)
12. de Vries, E., Koutavas, V., Hennessy, M.: Liveness of communicating transactions—(extended abstract). In: APLAS, volume 6461 of LNCS, pp. 392–407. Springer (2010)
13. Deniérou, P.-M., Yoshida, N.: Dynamic multirole session types. In: POPL, pp. 435–446. ACM Press (2011)
14. Dezani-Ciancaglini, M., Giannini, P.: Reversible multiparty sessions with checkpoints. In: EXPRESS/SOS, volume 222 of EPTCS, pp. 60–74 (2016)
15. Giachino, E., Sackman, M., Drossopoulou, S., Eisenbach, S.: Softly safely spoken: Role playing for session types. In: PLACES (2009)
16. Graversen, E., Phillips, I., Yoshida, N.: Towards a categorical representation of reversible event structures. In: PLACES, volume 246 of EPTCS, pp. 49–60 (2017)
17. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type disciplines for structured communication-based programming. In: ESOP, volume 1381 of LNCS, pp. 22–138. Springer (1998)
18. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: POPL, pp. 273–284. ACM Press (2008)
19. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. *J. ACM* **63**(1), 9 (2016)
20. Hu, R., Yoshida, N.: Explicit connection actions in multiparty session types. In: Fundamental Approaches to Software Engineering, volume 10202 of LNCS, pp. 116–133. Springer (2017)
21. Koutavas, V., Spaccasassi, C., Hennessy, M.: Bisimulations for communicating transactions—(extended abstract). In: FOSSACS, volume 8412 of LNCS, pp. 320–334. Springer (2014)
22. Lanese, I., Mezzina, C.A., Schmitt, A., Stefani, J.-B.: Controlling reversibility in higher-order pi. In: CONCUR, volume 6901 of LNCS, pp. 297–311. Springer (2011)
23. Lanese, I., Mezzina, C.A., Stefani, J.-B.: Reversing higher-order pi. In: CONCUR, volume 6269 of LNCS, pp. 478–493. Springer (2010)
24. Mezzina, C.A., Pérez, J.A.: Reversible semantics in session-based concurrency. In: ICTCS, volume 1720 of CEUR, pp. 221–226 (2016). CEUR-WS.org
25. Mezzina, C.A., Pérez, J.A.: Reversible sessions using monitors. In: PLACES, volume 211 of EPTCS, pp. 56–64 (2016)
26. Mezzina, C.A., Pérez, J.A.: Causally consistent reversible choreographies: a monitors-as-memories approach. In: PPDP, pp. 127–138. ACM Press (2017)
27. Mezzina, C.A., Tuosto, E.: Choreographies for automatic recovery (2017). CoRR, [arXiv:1705.09525](https://arxiv.org/abs/1705.09525)
28. Neykova, R., Yoshida, N.: Let it recover: multiparty protocol-induced recovery. In: CC, pp. 98–108. ACM Press (2017)
29. Padovani, L.: Type reconstruction for the linear π -calculus with composite regular types. *Log. Methods Comput. Sci.* **11**(4), 1–23 (2015)
30. Phillips, I., Ulidowski, I.: Operational semantics of reversibility in process algebra. In: APC, volume 162 of ENTCS, pp. 281–286 (2006)

31. Phillips, I., Ulidowski, I.: Reversing algebraic process calculi. *J. Log. Algebr. Methods Program.* **73**(1–2), 70–96 (2007)
32. Phillips, I., Ulidowski, I.: Reversibility and asymmetric conflict in event structures. *J. Log. Algebr. Methods Program.* **84**(6), 781–805 (2015)
33. Pierce, B.C.: *Types and Programming Languages*. MIT Press, Cambridge (2002)
34. Tiezzi, F., Yoshida, N.: Reversible session-based pi-calculus. *J. Log. Algebr. Methods Program.* **84**(5), 684–707 (2015)
35. Tiezzi, F., Yoshida, N.: Reversing single sessions. In: *RC*, volume 9720 of LNCS, pp. 52–69. Springer (2016)
36. Vieira, H.T., Caires, L., Seco, J.C.: The conversation calculus: a model of service-oriented computation. In: *ESOP*, volume 4960 of LNCS, pp. 269–283. Springer (2008)
37. Winskel, G.: *Events in Computation*. PhD thesis, Department of Computer Science, University of Edinburgh (1980)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.