

Concurrent Reversible Sessions

Ilaria Castellani, Mariangiola Dezani-Ciancaglini, Paola Giannini

► **To cite this version:**

Ilaria Castellani, Mariangiola Dezani-Ciancaglini, Paola Giannini. Concurrent Reversible Sessions. CONCUR 2017 - 28th International Conference on Concurrency Theory , Sep 2017, Berlin, Germany. Leibniz International Proceedings in Informatics (LIPIcs), 85, pp.1-17, CONCUR 2017. <10.4230/LIPIcs.CONCUR.2017.30>. <hal-01639845>

HAL Id: hal-01639845

<https://hal.inria.fr/hal-01639845>

Submitted on 28 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Concurrent Reversible Sessions*

Ilaria Castellani¹, Mariangiola Dezani-Ciancaglini^{†2}, and
Paola Giannini^{‡3}

- 1 Université Côte d'Azur, INRIA,
2004 Route des Lucioles, 06902 Sophia Antipolis, France
ilaria.castellani@inria.fr
- 2 Dipartimento di Informatica, Università di Torino,
corso Svizzera 185, 10131 Torino, Italy
dezani@di.unito.it
- 3 DiSIT, Università del Piemonte Orientale,
Via Teresa Michel 11, 15121 Alessandria, Italy
paola.giannini@uniupo.it

Abstract

We present a calculus for concurrent reversible multiparty sessions, which improves on recent proposals in several respects: it allows for concurrent and sequential composition within processes and types, it gives a compact representation of the *past* of processes and types, which facilitates the definition of rollback, and it implements a fine-tuned strategy for backward computation. We propose a refined session type system for our calculus and show that it enforces the expected properties of session fidelity, forward and backward progress, as well as causal consistency. In conclusion, our calculus is a conservative extension of previous proposals, offering enhanced expressive power and refined analysis techniques.

1998 ACM Subject Classification F.1.2 Parallelism and Concurrency, F.3.3 Type Structure

Keywords and phrases Communication-centric Systems, Reversible Computation, Process Calculi, Multiparty Session Types

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2017.30

1 Introduction

Building on recent proposals, we argue that session types and reversibility can be fruitfully combined to yield a model for correct and reliable communication-centric systems. 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, on its side, may be viewed as 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. A condition which is usually required for undoing a computational step is that all its effects have been already undone [21, 22, 30, 33, 8, 16, 15]. This property is generally referred to as *causal consistency*.

* The authors acknowledge a partial support of COST Action IC1405 on Reversible Computation – extending horizons of computing.

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

‡ This original research has the financial support of the Università del Piemonte Orientale.



A stronger property, called *full reversibility* in [25], enables a system to restore exactly a past state, keeping no trace of the rollback [9, 20, 24]. Sometimes neither of these properties can be achieved, particularly in distributed systems, and one has to go for weaker properties. In some cases, it is even desirable that a restored state be not exactly the same as the original one: for instance, the restored state could keep some memory of the undone computational path, so as to avoid engaging in that path again in case it led to an unsuccessful state.

In the setting of structured communications, reversibility has been first studied for contracts [2, 3] and transactions [10, 11, 19]. Only recently has this issue started to be addressed for session calculi, both binary [33, 23] and multiparty [14, 34, 28, 24] (see Section 6 for more discussion on related work).

When reversing a structured interaction, one has to face the problem of preserving consistency of the global state: if one of the partners triggers a rollback, then all its communicating partners should roll back accordingly. This is where session types come to the rescue, with their precise specification of the functionality of communications (sender, receiver and message), and of the order in which they should occur.

We present a calculus for concurrent reversible multiparty sessions, which extends previous proposals in several ways. First of all, our protocols are more general than those specified by standard multiparty session types [18]: concurrent communications are allowed both between disjoint groups of participants, and, in a controlled way that excludes auto-concurrency, also within participants themselves. More specifically, we relax the linearity constraint of standard multiparty session types, to allow protocol participants to perform communications in parallel, as long as they do not generate races. To enable the control flows to join again after a bunch of parallel interactions, besides parallel composition we also introduce full sequential composition in the syntax of both types and processes.

Moreover, our calculus gives a compact representation of the *past* of processes and types, which facilitates the definition of rollback, and it implements a fine-tuned strategy for backward computation, which is geared towards achieving compliance.

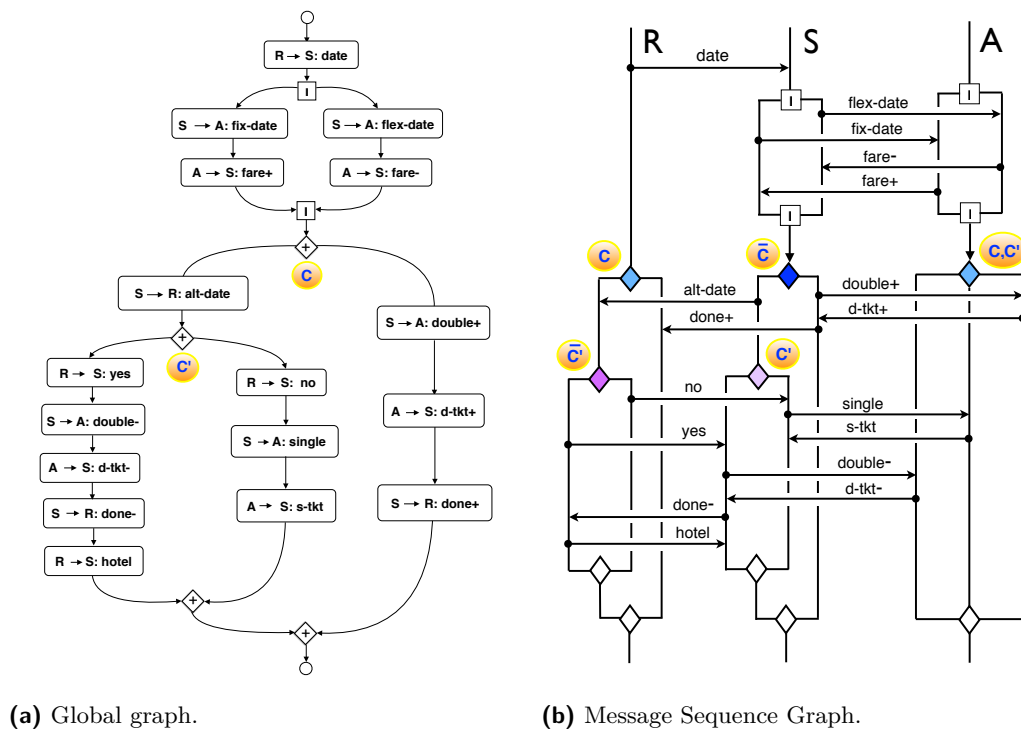
The main contributions of our paper may be summarised as follows:

- the introduction of *concurrent session types* in multiparty session frameworks;
- the proposal of a reversible session calculus which is more general than the existing ones, and which conservatively extends them, preserving the expected properties of session fidelity, forward and backward progress, and causal consistency;
- a fine-tuned strategy for rollback to checkpointed choices, which can only be triggered by choice leaders in predefined states of the computation, leading back to the choice state stripped off the unsuccessful path, so as redirect the computation towards alternative, potentially successful, paths.

The rest of the paper is organised as follows. In Section 2 we present our running example. In Section 3 we define the syntax and operational semantics – both forward and backward – of our calculus. In Section 4 we introduce our extended syntax for global types and session types and we establish well-formedness conditions for global types. Section 5 presents our type system and proves its soundness, namely that it ensures the expected semantic properties. We conclude in Section 6 with some discussion on related and future work.

2 Travel protocol example

To illustrate our approach, we present a simple travel protocol, involving three parties, a Researcher **R**, a Student **S** and an Airline company **A**. This protocol, henceforth called the *RSA protocol*, will serve as our running example throughout the paper.



■ **Figure 1** RSA protocol.

When graphically representing protocols, we shall use two kinds of graphs: the first one, called *global graph*, is borrowed from [13] and gives a global description of the protocol; the second one, which we call *message sequence graph*, is closer to message sequence charts and emphasises the control flows of the individual partners. In both cases, diamonds will be used as delimiters for the branching construct, and squares as delimiters for the parallel construct.

In the global graph, boxes represent communications and their whole functionality: sender, receiver and message. In the message sequence graph, the control flow of the protocol is split into the individual contributions of the partners: each partner is denoted by a vertical line; communications are represented by arrows from the sender to the receiver, labelled by messages; and branching and parallel constructs are replicated for each involved partner.

The RSA protocol is depicted in Figure 1a and Figure 1b according to these conventions. Let us now informally describe the protocol. Suppose that **R** and **S** wish to travel together to some conference, and that **R** has date constraints but is flexible about prices, while **S** has funding constraints but is flexible about dates. The protocol starts by **R** sending to **S** her chosen date. Then **S** sends two parallel flight requests to **A**: one for the date chosen by **R**, and one for more flexible dates. The airline **A** replies by sending in parallel the respective fares: a higher fare for the fixed date, and a lower fare for the flexible date. At this point, if **S** can afford the higher fare, she will book two tickets at this fare, then send a message `done+` to **R**, and the whole interaction will end here.

Suppose now that **S** cannot afford the higher fare. In this case, she will propose to **R** the alternative date corresponding to the cheaper fare. If **R** cannot travel at this date she replies `no`. Then **S** will book a single ticket for herself at the cheaper fare, and the interaction terminates. If instead this date is suitable for **R**, she replies `yes` to **S**. Now **S** books two tickets at the cheaper fare and sends a message `done-` to **R**. Then **R** informs **S** that the hotel

booking has been changed to fit the new date (supposing \mathbf{R} had previously booked the hotel for her chosen date), thereby closing the interaction.

So far we have described only the forward behaviour of the session. Suppose now that after proposing the alternative date to \mathbf{R} and receiving \mathbf{R} 's agreement, \mathbf{S} discovers that this date is not suitable anymore. Then \mathbf{S} will trigger a rollback to her internal choice with checkpoint label \overline{C} , and \mathbf{R} will have to roll back to her corresponding external choice with checkpoint label C . At this point, since the first branch of the choice has been already explored with no success, \mathbf{S} will have to choose the second branch and book the more expensive flight. On her side, \mathbf{R} could decide to roll back to her internal choice with label $\overline{C'}$ if she discovers that the hotel is fully booked at the alternative date, and then both \mathbf{S} and \mathbf{A} will have to roll back to their corresponding external choices with label C' .

In our calculus, only “the leader” of a choice, i.e., the participant who solved the choice by sending the first message, will be authorised to trigger a rollback. In the RSA protocol, the leader of the first choice is \mathbf{S} , while the leader of the second choice is \mathbf{R} . In the message sequence graph, the leader is distinguished by the fact that her choice has an overlined checkpoint. In the global graph, which collapses the distributed structure of the interaction, the choice leader is the sender in the first communication after the choice. In fact, the graph of Figure 1b can always be derived from that of Figure 1a through a projection operation.

3 Calculus

In this section, we introduce the syntax and semantics of our calculus. As usual in session calculi, we distinguish between user processes, which have not started to be executed yet, and runtime processes. The executed part of runtime processes will be marked with hats: this extension is instrumental to reversing computations. Processes may be decorated by checkpoint labels, marking them as possible rollback points. Our syntax generalises that of standard multiparty session calculi [18, 12], featuring the additional operators of parallel and sequential composition (the latter replacing the prefixing operator).

We assume the following base sets: *messages*, ranged over by λ, λ', \dots and forming the set Msg ; *checkpoint labels*, ranged over by C, C' and forming the set ChLa ; and *session participants*, ranged over by $\mathbf{p}, \mathbf{q}, \mathbf{r}$ and forming the set Part .

We use δ to range over general sets of checkpoint labels, and Δ to stand for either δ or δ extended by exactly one overlined checkpoint label:

$$\delta ::= \emptyset \quad | \quad \delta, C \quad \quad \Delta ::= \delta \quad | \quad \delta, \overline{C}$$

Sets of checkpoint labels are associated with choices. More precisely, an overlined checkpoint label can only be associated with an internal choice and is said to be *active*. A simple checkpoint label is *passive* and may be associated with both internal and external choices. Intuitively, an overlined checkpoint label is the handle of a backward move: a participant who crossed an internal choice (henceforth also called the *choice leader*) with an active checkpoint label, 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 that checkpoint label.

Let $\pi \in \{\mathbf{p}?\lambda, \mathbf{p}!\lambda \mid \mathbf{p} \in \text{Part}, \lambda \in \text{Msg}\}$ denote an *atomic action*, namely a directed input or output action. An atomic action can bear a hat, in which case it represents an already executed or *past* action. We use $\hat{\pi}$ to stand for either π or $\hat{\pi}$. External and internal choices can also bear hats, indicating that one branch has been chosen. We use $\hat{\Sigma}$ to stand for either Σ or $\hat{\Sigma}$, and similarly for $\hat{\oplus}$.

► **Definition 1.** *Processes* are defined by:

$$P ::= \widetilde{\sum}_{i \in I} \widetilde{\pi}_i; P_i \mid \widehat{\bigoplus}_{i \in I} \widetilde{\pi}_i; P_i \mid \Delta P \mid P \mid P \mid P; P \mid \mu X. P \mid X \mid \text{skip}$$

Processes without and with hats are called respectively *user processes* and *runtime processes*.

We will omit empty sets of checkpoint labels, choice symbols in one-branch choices, and trailing skip processes. External and internal choices are assumed to be associative, commutative, idempotent, and non-empty (except when combined with binary choices in evaluation contexts or finished processes, see below). A choice with a single branch may be either an *input process* or an *output process*. Parallel composition is associative and commutative, with neutral element **skip**. Sequential composition is associative, with neutral element **skip**. The operators have the following precedence: ‘;’, ‘+’, ‘ $\widehat{\bigoplus}$ ’, ‘|’.

Following [27], we require recursion to be:

- *guarded*, i.e. a recursion variable X can only appear free in the second argument of a sequential composition whose first argument is different from **skip**;
- *sequential*, i.e. a recursion variable X cannot occur free in any branch of a parallel composition.

Processes are treated equi-recursively, i.e. they are identified with their generated tree [31].

The typing rules of Section 5 will ensure the following well-formedness conditions for processes:

1. External choices are between input processes;
2. Internal choices are between output processes;
3. A choice without a hat has branches without hats;
4. A choice with a hat has exactly one branch with hats;
5. Only choices are decorated with sets of checkpoint labels;
6. Only internal choices are decorated with an active checkpoint label.

Conditions 1 and 2 reflect the active role of outputs and the passive role of inputs. Conditions 3 and 4 express the fact that executing a choice amounts to executing one of its branches. Conditions 5 and 6 specify that choices are the only return points for backward reductions. In the following we will take advantage of these restrictions to simplify definitions.

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 $\mathfrak{p}[P]$, where participant \mathfrak{p} has behaviour P .

► **Definition 2.** *Networks* are defined by: $\mathbb{N} ::= \mathfrak{p}[P] \mid \mathbb{N} \parallel \mathbb{N}$

The operator \parallel is associative and commutative, with neutral element $\mathfrak{p}[\text{skip}]$ for each \mathfrak{p} .

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 $P \xrightarrow{\mathfrak{p}\lambda} P'$ for some $\mathfrak{p}, \lambda, P'$.

In the LTS for networks, forward and backward transitions have respectively the form $N \xrightarrow{\mathfrak{p}\lambda\mathfrak{q}} N'$ and $N \xrightarrow{C} N'$.

► **Definition 3.** *Evaluation contexts* \mathcal{E} are defined by:

$$\mathcal{E} ::= \delta(\widehat{\sum}_{i \in I} \pi_i; P_i \widehat{+} \widehat{\pi}; \mathcal{E}) \mid \Delta(\widehat{\bigoplus}_{i \in I} \pi_i; P_i \widehat{\oplus} \widehat{\pi}; \mathcal{E}) \mid \mathcal{E} \mid P \mid \mathcal{E}; P \mid F; \mathcal{E} \mid []$$

The definition of evaluation context ensures that in a sequential composition the evaluation of the second component can only start when the first is a *finished process* F , defined by:

$$F ::= \delta(\widehat{\sum}_{i \in I} \pi_i; P_i \widehat{+} \widehat{\pi}; F) \mid \Delta(\widehat{\bigoplus}_{i \in I} \pi_i; P_i \widehat{\oplus} \widehat{\pi}; F) \mid F \mid F \mid F; F \mid \text{skip}$$

$$\begin{array}{c}
 \delta \sum_{i \in I} \pi_i; P_i \xrightarrow{\pi_j} \delta(\widehat{\sum}_{i \in I \setminus \{j\}} \pi_i; P_i \hat{+} \hat{\pi}_j; P_j) \quad j \in I \quad [\text{EXTCH}] \\
 \\
 \Delta \bigoplus_{i \in I} \pi_i; P_i \xrightarrow{\pi_j} \Delta(\widehat{\bigoplus}_{i \in I \setminus \{j\}} \pi_i; P_i \hat{\oplus} \hat{\pi}_j; P_j) \quad j \in I \quad [\text{INTCH}] \\
 \\
 \frac{P \downarrow_{out} \quad \overline{C} \in \Delta \quad I \neq \emptyset}{\Delta(\widehat{\bigoplus}_{i \in I} P_i \hat{\oplus} P) \xrightarrow{\overline{C}} \Delta \bigoplus_{i \in I} P_i} [\text{BACKCKT}] \quad \frac{C \in \Delta}{\Delta S \xrightarrow{C} \Delta \{ S \}} [\text{BACKP}] \\
 \\
 \frac{P \xrightarrow{\pi} P'}{\mathcal{E}[P] \xrightarrow{\pi} \mathcal{E}[P']} [\text{CTFAT}] \quad \frac{P \xrightarrow{\overline{C}} P' \quad \mathcal{E} \text{ ok for } \overline{C}}{\mathcal{E}[P] \xrightarrow{\overline{C}} \mathcal{E}[P']} [\text{CTBA}] \quad \frac{P \xrightarrow{C} P' \quad \mathcal{E} \text{ ok for } C}{\mathcal{E}[P] \xrightarrow{C} \mathcal{E}[P']} [\text{CTBP}] \\
 \\
 \frac{P_1 \xrightarrow{p^? \lambda} P'_1 \quad P_2 \xrightarrow{q^! \lambda} P'_2}{\mathfrak{q}[\![P_1]\!] \parallel \mathfrak{p}[\![P_2]\!] \parallel \mathbb{N} \xrightarrow{p \lambda q} \mathfrak{q}[\![P'_1]\!] \parallel \mathfrak{p}[\![P'_2]\!] \parallel \mathbb{N}} [\text{COM}] \\
 \\
 \frac{P \xrightarrow{\overline{C}} P' \quad P_i \xrightarrow{C} P'_i \quad i \in I \quad P_j \not\xrightarrow{C} \quad j \in J}{\mathfrak{p}[\![P]\!] \parallel \prod_{i \in I} \mathfrak{p}_i[\![P_i]\!] \parallel \prod_{j \in J} \mathfrak{p}_j[\![P_j]\!] \xrightarrow{C} \mathfrak{p}[\![P']\!] \parallel \prod_{i \in I} \mathfrak{p}_i[\![P'_i]\!] \parallel \prod_{j \in J} \mathfrak{p}_j[\![P_j]\!] } [\text{BACK}]
 \end{array}$$

■ **Figure 2** LTS for processes and networks.

The LTSs for processes and networks are given in Figure 2. 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 choice operator with a hat. With this technique, inspired by [6] and already used for reversible computations in [24], 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 $\pi_j; P_j \xrightarrow{\pi_j} \hat{\pi}_j; P_j$. Rule [BACKCKT] is the main backward rule: it applies to a past internal choice, where one branch has been partially executed, and it allows the process to roll back to the original choice where the executed branch is removed. For this to be possible, the choice must have at least one non executed branch P_i and a set of checkpoint labels containing 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 one who can decide to roll back to this choice. The condition $P \downarrow_{out}$ means that in order to trigger a rollback, P should be “in lead” again, namely able to do an output.

Rule [BACKP], where S denotes either $\widehat{\sum}_{i \in I} P_i$, or $\widehat{\bigoplus}_{i \in I} P_i$, is needed to allow the remaining participants to roll back. The mapping $\{ \}^C$ erases hats from processes, yielding user processes, i.e. $\{ \hat{\pi}; P \}^C = \pi; \{ P \}$ and $\{ \}^C$ 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. An evaluation context \mathcal{E} is ok for \overline{C} (C) if $\overline{C} \notin \Delta$ ($C \notin \Delta$) whenever \mathcal{E} is a sub-context of a context of the shape $\Delta(\widehat{\sum}_{i \in I} P_i \hat{+} \hat{\pi}; \mathcal{E}')$ or $\Delta(\widehat{\bigoplus}_{i \in I} P_i \hat{\oplus} \hat{\pi}; \mathcal{E}')$. 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 first one, as in [28]. This is needed to assure subject reduction, see Example 8.

Rule [COM] is standard. We write $P \not\xrightarrow{C}$ if neither Rule [BACKP] nor Rule [CTBP] (with label C) can be applied to P . This means that C can only occur in user processes within P .

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 this choice will do so in the same step. This will be the basis for our soundness result in Section 5. A direct implementation of this rule is clearly unrealistic. To that purpose, asynchronous communications including rollback messages should be used, as in [24]. As expected, sequences of network transitions generate *traces*, which are words over the infinite alphabet $\{p\lambda q \mid p, q \in \text{Part}, \lambda \in \text{Msg}\} \cup \text{ChLa}$.

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 $\longrightarrow \cup \curvearrowright$.

Note that our semantics does not enforce any scheduling policy. Hence, a network may generate an infinite trace that does not involve all participants. For instance, the network:

$$p[\mu X_1.q!\lambda_1;X_1 \mid \mu X_2.r!\lambda_2;X_2] \parallel q[\mu X_3.p?\lambda_1;X_3] \parallel r[\mu X_4.p?\lambda_2;X_4]$$

may generate the infinite trace $p\lambda_1q p\lambda_1q \dots$. In a more elaborate calculus, we could impose a *fair scheduling policy*, forcing p to communicate alternately with q and r , and in general, no communication to be iterated until each participant has communicated at least once.

A network for our RSA travel protocol in Section 2 is $\mathbb{N}_{in} = \mathbf{R}[\mathbf{P}^R] \parallel \mathbf{S}[\mathbf{P}^S] \parallel \mathbf{A}[\mathbf{P}^A]$, defined as follows (where messages are abbreviated in a programming language style):

- The process P^S is $\mathbf{R}?\text{dt};(P_1^S \mid P_2^S); (P_3^S_{\{\bar{C}\}} \oplus P_4^S)$ where

$$\begin{array}{ll} P_1^S = \mathbf{A}!\text{fxDt};\mathbf{A}?\text{frP1} & P_2^S = \mathbf{A}!\text{f1Dt};\mathbf{A}?\text{frMn} \\ P_3^S = \mathbf{R}!\text{a1Dt};(P_5^S_{\{C'\}} + P_6^S) & P_4^S = \mathbf{A}!\text{dbP1};\mathbf{A}?\text{dTkP1};\mathbf{R}!\text{dnP1} \\ P_5^S = \mathbf{R}?\text{yes};\mathbf{A}!\text{dbMn};P_7^S & P_6^S = \mathbf{A}?\text{dTkMn};\mathbf{R}!\text{dnMn};\mathbf{R}?\text{ht1} \\ P_7^S = \mathbf{A}?\text{dTkMn};\mathbf{R}!\text{dnMn};\mathbf{R}?\text{ht1} & P_8^S = \mathbf{R}?\text{no};\mathbf{A}!\text{sn};\mathbf{A}?\text{snTk} \end{array}$$

- The process P^R is $\mathbf{S}!\text{dt};(P_1^R_{\{C\}} + \mathbf{S}?\text{dnP1})$ where

$$P_1^R = \mathbf{S}?\text{a1Dt};(P_2^R_{\{\bar{C}'\}} \oplus \mathbf{S}!\text{no}) \quad P_2^R = \mathbf{S}!\text{yes};\mathbf{S}?\text{dnMn};\mathbf{S}!\text{ht1}$$

- The process P^A is $(P_1^A \mid P_2^A); \{C, C'\} \sum_{i \in \{3,4,5\}} P_i^A$ where

$$\begin{array}{lll} P_1^A = \mathbf{S}?\text{fxDt};\mathbf{S}!\text{frP1} & P_2^A = \mathbf{S}?\text{f1Dt};\mathbf{S}!\text{frMn} \\ P_3^A = \mathbf{S}?\text{dbMn};\mathbf{S}!\text{dTkMn} & P_4^A = \mathbf{S}?\text{sn};\mathbf{S}!\text{snTk} & P_5^A = \mathbf{S}?\text{dbP1};\mathbf{S}!\text{dTkP1} \end{array}$$

In the computation below, we denote with \widehat{P} the process P entirely marked with hats.

$$\begin{array}{l} \mathbb{N}_{in} \xrightarrow{\mathbf{R}\text{dtS}} \mathbf{R}[\widehat{\mathbf{S}!\text{dt}};(P_1^R_{\{C\}} + \mathbf{S}?\text{dnP1})] \parallel \mathbf{S}[\widehat{\mathbf{R}?\text{dt}};(P_1^S \mid P_2^S); (P_3^S_{\{\bar{C}\}} \oplus P_4^S)] \parallel \mathbf{A}[\widehat{\mathbf{P}^A}] \\ \longrightarrow^* \mathbf{R}[\widehat{\mathbf{S}!\text{dt}};(P_1^R_{\{C\}} + \mathbf{S}?\text{dnP1})] \parallel \mathbf{S}[\widehat{F^S}; (P_3^S_{\{\bar{C}\}} \oplus P_4^S)] \parallel \mathbf{A}[\widehat{Q^A}] \\ \text{where } F^S = \widehat{\mathbf{R}?\text{dt}};(\widehat{P_1^S} \mid \widehat{P_2^S}), Q^A = (\widehat{P_1^A} \mid \widehat{P_2^A}); \{C, C'\} \sum_{i \in \{3,4,5\}} \widehat{P_i^A} \\ \xrightarrow{\mathbf{S}\text{a1DtR}} \mathbf{R}[\widehat{\mathbf{S}!\text{dt}};(\widehat{\mathbf{S}?\text{a1Dt}};(\widehat{P_2^R}_{\{\bar{C}'\}} \oplus \widehat{\mathbf{S}!\text{no}})_{\{C\}} \widehat{\mathbf{S}?\text{dnP1}})] \parallel \\ \mathbf{S}[\widehat{F^S}; (\widehat{\mathbf{R}!\text{a1Dt}};(\widehat{P_5^S}_{\{C'\}} + \widehat{P_6^S})_{\{\bar{C}\}} \widehat{P_4^S})] \parallel \mathbf{A}[\widehat{Q^A}] \\ \longrightarrow^* \mathbf{R}[\widehat{\mathbf{S}!\text{dt}};(\widehat{\mathbf{S}?\text{a1Dt}};(\widehat{\mathbf{S}!\text{yes}};\widehat{\mathbf{S}?\text{dnMn}};\widehat{\mathbf{S}!\text{ht1}})_{\{\bar{C}'\}} \widehat{\mathbf{S}!\text{no}})_{\{C\}} \widehat{\mathbf{S}?\text{dnP1}})] \parallel \\ \mathbf{S}[\widehat{F^S}; (\widehat{\mathbf{R}!\text{a1Dt}};(\widehat{\mathbf{R}?\text{yes}};\widehat{\mathbf{A}!\text{dbMn}};P_7^S_{\{C'\}} \widehat{P_6^S})_{\{\bar{C}\}} \widehat{P_4^S})] \parallel \mathbf{A}[\widehat{Q^A}] \\ \curvearrowright \mathbf{R}[\widehat{\mathbf{S}!\text{dt}};(P_1^R_{\{C\}} + \mathbf{S}?\text{dnP1})] \parallel \mathbf{S}[\widehat{F^S}; P_4^S] \parallel \mathbf{A}[\widehat{Q^A}] \end{array}$$

Note that the last network differs from that in Line 2 only by the absence of process P_3^S .

4 Global Types and Session Types

A *multiparty session* is a series of communications among a fixed number of participants [18], which follows a predefined protocol specified by a *global type*.

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

$$\gamma ::= \emptyset \quad | \quad \{C\}$$

Sets γ will be associated with choices in global types.

Let $\alpha^p \in \{p \xrightarrow{\lambda} q \mid q \in \text{Part}, \lambda \in \text{Msg}\}$ denote an *atomic communication* with sender p . An atomic communication can bear a hat, in notation $\widehat{\alpha}^p$, in which case it represents an executed or *past* communication. The symbol $\widetilde{\alpha}^p$ stands for either α^p or $\widehat{\alpha}^p$.

Global types G are built from choices among atomic communications with the same sender, with the constructs of parallel composition, sequential composition, and recursion. A global type K specifies an interaction that is still to start. On the opposite, a global type H specifies a completely executed interaction, in which there is one path entirely marked with hats. A general global type G specifies a partially executed interaction, whose executed part (history) is specified by subterms H , while the parts that have been discarded in choices or remain to be executed are specified by subterms K .

► **Definition 4.** *Global types* G are defined by:

$$\begin{aligned} K & ::= \gamma \boxplus_{i \in I} \alpha_i^p; K_i \quad | \quad K \parallel K \quad | \quad K; K \quad | \quad \mu t. K \quad | \quad t \quad | \quad \text{Skip} \\ H & ::= \gamma(\widehat{\boxplus}_{i \in I} \alpha_i^p; K_i \widehat{\boxplus} \widehat{\alpha}^p; H) \quad | \quad H \parallel H \quad | \quad H; H \quad | \quad \text{Skip} \\ G & ::= K \quad | \quad \gamma(\widehat{\boxplus}_{i \in I} \alpha_i^p; K_i \widehat{\boxplus} \widehat{\alpha}^p; G) \quad | \quad G \parallel G \quad | \quad G; K \quad | \quad H; G \end{aligned}$$

The choice operator \boxplus is n -ary, commutative and idempotent. We use $\widehat{\boxplus}$ for either \boxplus or $\widehat{\boxplus}$. In the binary case we write $\widehat{\alpha}_1^p; G \gamma \widehat{\boxplus} \alpha_2^p; G'$. The notation $\gamma(\widehat{\boxplus}_{i \in I} \alpha_i^p; K_i \widehat{\boxplus} \widehat{\alpha}^p; G)$ stands for a choice where the branch initiating with $\widehat{\alpha}^p$ has started to be executed. By abuse of notation, we will write $\gamma \widehat{\boxplus}_{i \in I} \alpha_i^p; G_i$ for either $\gamma \boxplus_{i \in I} \alpha_i^p; G_i$ or $\gamma(\widehat{\boxplus}_{i \in I \setminus \{j\}} \alpha_i^p; G_i \widehat{\boxplus} \alpha_j^p; G_j)$. When the checkpoint set γ associated with a choice is empty, we omit it. We call *rooted interaction* a subterm $\widetilde{\alpha}^p; G$. Hence a type $\gamma \widehat{\boxplus}_{i \in I} \alpha_i^p; G_i$ is a choice among a number of rooted interactions with the same sender, at most one of which bears a hat. If I is a singleton we write $\widetilde{\alpha}^p; G$. So $\widetilde{\alpha}^p; G$ can denote either a rooted interaction or a choice with a single branch (the context will disambiguate if needed). We use $\text{pa}(G)$ to denote *the set of participants of* G , namely all p, q such that $p \xrightarrow{\lambda} q$ occurs in G .

As for processes, parallel and sequential composition are associative, with neutral element **Skip**, parallel composition is commutative, and recursion is guarded, sequential, and treated equi-recursively. The operators have the following precedence: ‘;’, ‘ \boxplus ’, ‘ \parallel ’.

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. In this paper, session types will only be defined as projections of well-formed global types, see Definition 7. (There is no circularity since global type well-formedness does not depend on session types.) Session pre-types are obtained from processes by replacing external and internal choices with intersections and unions, X with t and skip with **Skip**, with similar conventions.

► **Definition 5.** *Session pre-types* are defined by:

$$T ::= \widetilde{\bigwedge}_{i \in I} \widetilde{\pi}_i; T_i \quad | \quad \widetilde{\bigvee}_{i \in I} \widetilde{\pi}_i; T_i \quad | \quad \Delta T \quad | \quad T | T \quad | \quad T; T \quad | \quad \mu t. T \quad | \quad t \quad | \quad \text{Skip}$$

We call an intersection $\widetilde{\bigwedge}_{i \in I} \widetilde{p_i ? \lambda_i}; T_i$ *non-ambiguous* if its initial inputs are all distinct, namely if $i \neq j$ implies either $p_i \neq p_j$ or $\lambda_i \neq \lambda_j$. Projectability will require non-ambiguity of intersections. Our definition of projection (see Figure 3) is more liberal than in other session calculi, since we have an extended syntax for global types and we want to maximise the set of global types that are projectable. In particular, we allow equal messages between the same pair of participants in choices (this is usually forbidden [12, 28]). In any case, 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 *join of the types in the set*, $\bigsqcup_{i \in I} T_i$, be defined.

The join $\bigsqcup_{i \in I} T_i$ is a partial operator, which checks that the T_i 's are compatible behaviours and then combines them into a single session type. (In the definition of join we may assume that all the T_i 's are projections of global types.) Intuitively, $\bigsqcup_{i \in I} T_i$ is defined if the concerned participant either has the same behaviour in all T_i , or, if this is not the case, if she receives a message that “notifies” her about the chosen T_i before she starts differentiating her behaviour.

If one of the T_i 's is an intersection of input behaviours, then so must be all the other T_i 's.

To join intersection types, we define an auxiliary operator \uplus , which takes a non-ambiguous intersection $\widetilde{\bigwedge}_{i \in I} \widetilde{p_i ? \lambda_i}; T_i$ and a session type $\widehat{p ? \lambda}; T$, and combines them, if possible:

$$\left(\widetilde{\bigwedge}_{i \in I} \widetilde{p_i ? \lambda_i}; T_i\right) \uplus \widehat{p ? \lambda}; T = \begin{cases} \widetilde{\bigwedge}_{i \in I} \widetilde{p_i ? \lambda_i}; T_i \widetilde{\wedge} \widehat{p ? \lambda}; T & \text{if } p_i \neq p \text{ or } \lambda_i \neq \lambda \text{ for all } i \in I \\ \widehat{\bigwedge}_{i \in I \setminus \{j\}} \widetilde{p_i ? \lambda_i}; T_i \widetilde{\wedge} (\widehat{p_j ? \lambda_j} \widetilde{\cup} \widehat{p ? \lambda}); T \sqcup T_j & \text{if } p_j = p \\ & \text{and } \lambda_j = \lambda \text{ (} j \in I \text{)} \end{cases}$$

where $\widehat{p ? \lambda} \widetilde{\cup} p ? \lambda = p ? \lambda \widetilde{\cup} \widehat{p ? \lambda} = \widehat{p ? \lambda}$ and $p ? \lambda \widetilde{\cup} p ? \lambda = p ? \lambda$ and the obtained intersections have hats if one of the input behaviours has a hat. If the prefix $p ? \lambda$ is different from all the $p_i ? \lambda_i$, then \uplus produces the intersection of the two types. In this case, receiving message λ from p amounts to being notified of the choice of branch T . If instead $p ? \lambda = p_j ? \lambda_j$ for some $j \in I$, then we try to combine the types $\widehat{p ? \lambda}; T$ and $\widehat{p_j ? \lambda_j}; T_j$ by factoring out their common prefix and producing $\widehat{p ? \lambda}; T \sqcup T_j$ (where $p ? \lambda$ has a hat if one of the two building prefixes has a hat), so that the resulting intersection is again non-ambiguous. If $T \sqcup T_j$ is defined, then T and T_j will be discriminated later if they are different, see G_1 in Example 6.

To define \uplus 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 \uplus to a set of intersection types.

We may now formally define the join of a set of session types.

$$\bigsqcup_{i \in I} T_i = \begin{cases} \Delta \uplus_{i \in I, j \in J_i} \widetilde{p_{i,j} ? \lambda_{i,j}}; T_{i,j} & \text{if } T_i = \Delta_i \widetilde{\bigwedge}_{j \in J_i} \widetilde{p_{i,j} ? \lambda_{i,j}}; T_{i,j} \text{ for all } i \in I \\ & \text{and } \Delta = \bigcup_{i \in I} \Delta_i \\ T; \bigsqcup_{i \in I} T'_i & \text{if } T_i = T; T'_i \text{ for all } i \in I \\ \text{Skip} & \text{if } T_i = \text{Skip for all } i \in I \end{cases}$$

The definitions of $\bigsqcup_{i \in I} T_i$ and \uplus are mutually recursive. Let us examine the three clauses. If all T_i 's are intersections of session types, then we combine them with \uplus as explained above. The checkpoint set of the resulting intersection is the union of the checkpoint sets of the T_i 's.

If all T_i 's start with the same interaction T , and their continuations T'_i 's are consistent, then we factor out T and we produce the session type T followed by the join of the T'_i 's.

If some T_i is *Skip*, it means that the participant terminates in the branch T_i . Then it must terminate in all branches, so all T_i 's must be *Skip*.

The join is not defined between intersections and unions, nor between unions with different prefixes. Our join extends that of [12], to deal with parallel and sequential composition.

$$\begin{aligned}
& \widetilde{(\mathbf{p} \xrightarrow{\lambda} \mathbf{q}; \mathbf{G})} \upharpoonright \mathbf{p} = \widetilde{\mathbf{q}! \lambda; \mathbf{G}} \upharpoonright \mathbf{p} \quad \widetilde{(\mathbf{p} \xrightarrow{\lambda} \mathbf{q}; \mathbf{G})} \upharpoonright \mathbf{q} = \widetilde{\mathbf{p} ? \lambda; \mathbf{G}} \upharpoonright \mathbf{q} \quad \widetilde{(\mathbf{p} \xrightarrow{\lambda} \mathbf{q}; \mathbf{G})} \upharpoonright \mathbf{r} = \mathbf{G} \upharpoonright \mathbf{r}, \text{ if } \mathbf{r} \neq \mathbf{p}, \mathbf{q} \\
& (\gamma \boxplus_{i \in I} \widetilde{\alpha_i^p}; \mathbf{G}_i) \upharpoonright \mathbf{r} = \begin{cases} \widetilde{\gamma \bigvee_{i \in I} (\alpha_i^p; \mathbf{G}_i)} \upharpoonright \mathbf{r} & \text{if } \mathbf{r} = \mathbf{p} \\ \gamma[\bigwedge_{i \in I} (\alpha_i^p; \mathbf{G}_i)] \upharpoonright \mathbf{r} & \text{otherwise} \end{cases}
\end{aligned}$$

■ **Figure 3** Projection of global types onto participants.

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

$$\begin{aligned}
\gamma[\bigwedge_{i \in I} \widetilde{\pi}_i; \mathbf{T}_i] &= \gamma \bigwedge_{i \in I} \widetilde{\pi}_i; \mathbf{T}_i & \gamma[\bigvee_{i \in I} \widetilde{\pi}_i; \mathbf{T}_i] &= \gamma \bigvee_{i \in I} \widetilde{\pi}_i; \mathbf{T}_i & \gamma[\Delta \mathbf{T}] &= \gamma \cup \Delta \mathbf{T} \text{ if } \gamma \cap \Delta = \emptyset \\
\gamma[\mathbf{T}_0 \mid \mathbf{T}_1] &= \gamma[\mathbf{T}_0] \mid \gamma[\mathbf{T}_1] & \gamma[\mathbf{T}_0; \mathbf{T}_1] &= \gamma[\mathbf{T}_0]; \mathbf{T}_1 & \gamma[\mu \mathbf{t}. \mathbf{T}'] &= \mu \mathbf{t}. \gamma[\mathbf{T}'] & \gamma[\text{Skip}] &= \text{Skip}
\end{aligned}$$

Prefixing adds γ to the first (possibly empty) set Δ found in \mathbf{T} . Intuitively, the checkpoint labels that are spread along successive choices in the global type may get grouped together on a single local choice when projected on participants. We do not need to define $\gamma[\mathbf{t}]$ since recursion is guarded. The condition in the clause for $\gamma[\Delta \mathbf{T}]$ is needed to avoid checkpoint labels which can never be used. This condition rules out for instance the global type:

$$(\mathbf{p} \xrightarrow{\lambda_1} \mathbf{q}; \mathbf{p} \xrightarrow{\lambda'_1} \mathbf{r}; (\mathbf{p} \xrightarrow{\lambda_2} \mathbf{q}; \mathbf{p} \xrightarrow{\lambda'_2} \mathbf{r} \{C\} \boxplus \mathbf{p} \xrightarrow{\lambda_3} \mathbf{q}; \mathbf{p} \xrightarrow{\lambda'_3} \mathbf{r})) \{C\} \boxplus \mathbf{p} \xrightarrow{\lambda_4} \mathbf{q}; \mathbf{p} \xrightarrow{\lambda'_4} \mathbf{r}$$

in which no backward reduction could return to the innermost occurrence of C . This is due to the condition on evaluation contexts in rules [CTBA] and [CTBP].

The projection of global types uses the projections of rooted interactions, see the first line of Figure 3. The projection of a choice is a union for the sending participant, and it is otherwise computed as the join of the projections on the branches. With $\overline{\gamma}$ we denote $\{\overline{C}\}$ if $\gamma = \{C\}$ and \emptyset if $\gamma = \emptyset$. The join operation can be undefined, and therefore also the projection of global types can be undefined. Since $\text{Skip} \sqcup \mathbf{T}$ is defined only if $\mathbf{T} = \text{Skip}$, the definition of projection ensures that all the branches of a choice $\gamma \boxplus_{i \in I} \widetilde{\alpha_i^p}; \mathbf{G}_i$ have the same participants, i.e., $\text{pa}(\widetilde{\alpha_i^p}; \mathbf{G}_i) = \text{pa}(\widetilde{\alpha_j^p}; \mathbf{G}_j)$ for all $i, j \in I$. The omitted cases in Figure 3 are either standard (recursion and Skip) or homomorphic projections (parallel and sequential composition). Notice that projection respects hats and checkpoint labels.

► **Example 6.** Let $\mathbf{G}_1 = (\widehat{\mathbf{p} \xrightarrow{\lambda} \mathbf{q}; \mathbf{p} \xrightarrow{\lambda_1} \mathbf{q}; \mathbf{q} \xrightarrow{\lambda_2} \mathbf{r}}) \widehat{\boxplus} (\widehat{\mathbf{p} \xrightarrow{\lambda} \mathbf{q}; \mathbf{p} \xrightarrow{\lambda_3} \mathbf{q}; \mathbf{q} \xrightarrow{\lambda_4} \mathbf{r}})$. The projections of \mathbf{G}_1 onto its participants are: $\mathbf{G}_1 \upharpoonright \mathbf{p} = (\widehat{\mathbf{q}! \lambda; \mathbf{q}! \lambda_1}) \widehat{\vee} (\widehat{\mathbf{q}! \lambda; \mathbf{q}! \lambda_3})$, $\mathbf{G}_1 \upharpoonright \mathbf{q} = \widehat{\mathbf{p} ? \lambda; ((\widehat{\mathbf{p} ? \lambda_1; \mathbf{r}! \lambda_2}) \widehat{\wedge} (\widehat{\mathbf{p} ? \lambda_3; \mathbf{r}! \lambda_4}))}$ and $\mathbf{G}_1 \upharpoonright \mathbf{r} = \widehat{\mathbf{q} ? \lambda_2 \wedge \mathbf{q} ? \lambda_4}$. Even though \mathbf{q} receives the same message λ from \mathbf{p} in the two branches of the choice, before sending two different messages to \mathbf{r} it receives a second message from \mathbf{p} , which identifies the chosen branch. Also the global type $\mathbf{G}_2 = (\mathbf{p} \xrightarrow{\lambda_1} \mathbf{q}; \mathbf{r} \xrightarrow{\lambda_2} \mathbf{q}) \boxplus (\mathbf{p} \xrightarrow{\lambda_3} \mathbf{q}; \mathbf{r} \xrightarrow{\lambda_2} \mathbf{q})$ is projectable, since \mathbf{r} sends the same message λ_2 in both branches.

Instead, the global type $\mathbf{G}_3 = (\widehat{\mathbf{p} \xrightarrow{\lambda} \mathbf{q}; \mathbf{q} \xrightarrow{\lambda_1} \mathbf{r}}) \widehat{\boxplus} (\widehat{\mathbf{p} \xrightarrow{\lambda} \mathbf{q}; \mathbf{q} \xrightarrow{\lambda_2} \mathbf{r}})$ is not projectable on the participant \mathbf{q} : in fact $(\widehat{\mathbf{p} ? \lambda; \mathbf{r}! \lambda_1}) \sqcup (\widehat{\mathbf{p} ? \lambda; \mathbf{r}! \lambda_2})$ is not defined, since it yields $\widehat{\mathbf{p} ? \lambda; (\mathbf{r}! \lambda_1 \sqcup \mathbf{r}! \lambda_2)}$ and $\mathbf{r}! \lambda_1 \sqcup \mathbf{r}! \lambda_2$ is not defined. This is justified since \mathbf{q} should send two different messages to \mathbf{r} in the two branches, while no previous communication identifies the chosen branch.

The next example features a checkpoint label set containing both a simple label and an overlined one. Let $\mathbf{G}_4 = \mathbf{p} \xrightarrow{\lambda_1} \mathbf{q}; \mathbf{p} \xrightarrow{\overline{\lambda_1}} \mathbf{r}; \mathbf{G}_5 \{C'\} \boxplus \mathbf{p} \xrightarrow{\lambda_2} \mathbf{q}; \mathbf{p} \xrightarrow{\overline{\lambda_2}} \mathbf{r}; \mathbf{G}_5$ where $\mathbf{G}_5 =$

$s \xrightarrow{\lambda_3} p; s \xrightarrow{\lambda'_3} r_{\{C\}} \boxplus s \xrightarrow{\lambda_4} p; s \xrightarrow{\lambda'_4} r$. Then the projection of G_4 onto s is $G_4 \upharpoonright s = p! \lambda_3 r! \lambda'_3 \{C', \bar{C}\}^\forall p! \lambda_4 r! \lambda'_4$.

Let us now look back at the RSA protocol. Its global type G , which may be seen as a syntactic description of the graph in Figure 1a, is the following:

$$\begin{aligned} G &= \mathbf{R} \xrightarrow{dt} \mathbf{S}; (G_1 \parallel G_2); (G_3 \{C\} \boxplus G_4) \text{ where} \\ G_1 &= \mathbf{S} \xrightarrow{fxDt} \mathbf{A}; \mathbf{A} \xrightarrow{frP1} \mathbf{S} & G_2 &= \mathbf{S} \xrightarrow{flDt} \mathbf{A}; \mathbf{A} \xrightarrow{frMn} \mathbf{S} \\ G_3 &= \mathbf{S} \xrightarrow{alDt} \mathbf{R}; (G_5 \{C'\} \boxplus G_6) & G_4 &= \mathbf{S} \xrightarrow{dbP1} \mathbf{A}; \mathbf{A} \xrightarrow{dTkp1} \mathbf{S}; \mathbf{S} \xrightarrow{dnP1} \mathbf{R} \\ G_5 &= \mathbf{R} \xrightarrow{yes} \mathbf{S}; \mathbf{S} \xrightarrow{dbMn} \mathbf{A}; \mathbf{A} \xrightarrow{dTkmn} \mathbf{S}; \mathbf{S} \xrightarrow{dnMn} \mathbf{R}; \mathbf{R} \xrightarrow{ht1} \mathbf{S} & G_6 &= \mathbf{R} \xrightarrow{no} \mathbf{S}; \mathbf{S} \xrightarrow{sn} \mathbf{A}; \mathbf{A} \xrightarrow{snTk} \mathbf{S} \end{aligned}$$

The projection of G on the participant \mathbf{S} (T^S) has the same structure as G , since \mathbf{S} is involved in all communications of the protocol:

$$\begin{aligned} T^S &= \mathbf{R}?dt; (T_1^S \mid T_2^S); (T_3^S \{C'\}^\forall T_4^S) \text{ where} \\ T_1^S &= \mathbf{A}!fxDt; \mathbf{A}?frP1 & T_2^S &= \mathbf{A}!flDt; \mathbf{A}?frMn \\ T_3^S &= \mathbf{R}!alDt; (T_5^S \{C'\} \wedge T_6^S) & T_4^S &= \mathbf{A}!dbP1; \mathbf{A}?dTkp1; \mathbf{R}!dnP1 \\ T_5^S &= \mathbf{R}?yes; \mathbf{A}!dbMn; \mathbf{A}?dTkmn; \mathbf{R}!dnMn; \mathbf{R}?ht1 & T_6^S &= \mathbf{R}?no; \mathbf{A}!sn; \mathbf{A}?snTk \end{aligned}$$

On the other hand, the projections of G on \mathbf{A} and \mathbf{R} are simpler, and are not given here.

In the presence of sequentialisation and recursion, projectability of global types is not enough to ensure *starvation freedom* in communication protocols. For example, the global type $\mu t. p \xrightarrow{\lambda} q; t; p \xrightarrow{\lambda'} r$ exhibits starvation, since the participant r will indefinitely wait for message λ' from participant p . This problem can be avoided by requiring that the communications which follow a loop involve only participants who are also communicating in the body of the loop. We formalise this by the predicate *sequentially well-formed* on global types. This predicate applied to $G; G'$ requires that the participants of G' be contained in the set of *allowed followers* of G , denoted by $\text{af}(G)$ and defined by

$$\begin{aligned} \text{af}(\gamma \boxplus_{i \in I} \tilde{\alpha}_i; G_i) &= \bigcap_{i \in I} \text{af}(G_i) & \text{af}(G \parallel G') &= \text{af}(G) \cup \text{af}(G') & \text{af}(G; G') &= \text{af}(G) \cap \text{af}(G') \\ \text{af}(\mu t. G) &= \begin{cases} \text{pa}(G) & \text{if } t \in G \\ \text{af}(G) & \text{otherwise} \end{cases} & \text{af}(\text{Skip}) &= \text{Part} \end{aligned}$$

In the definition of af for parallel composition of global types, the union is justified by the example $(\mu t. p \xrightarrow{\lambda_1} q; t \parallel \mu t. r \xrightarrow{\lambda_2} s; t); p \xrightarrow{\lambda_3} r$.

In order to avoid deadlocks, we need to impose some restrictions on the use of checkpoint labels and of messages in parallel communications. The same checkpoint label should not occur in two parallel global types. For suppose we allowed the global type:

$$(p \xrightarrow{\lambda_1} q; p \xrightarrow{\lambda'_1} r_{\{C\}} \boxplus p \xrightarrow{\lambda_2} q; p \xrightarrow{\lambda'_2} r) \parallel (p \xrightarrow{\lambda_3} q; p \xrightarrow{\lambda'_3} r_{\{C\}} \boxplus p \xrightarrow{\lambda_4} q; p \xrightarrow{\lambda'_4} r)$$

Here there are two parallel interactions, both involving participants p, q and r . Assume that message λ_1 has been exchanged in the first interaction, and message λ_3 has been exchanged in the second. Now participant p may want to roll back to the choice labelled by $\{C\}$ in the first interaction. Then participant q should roll back to the same choice, while the above type would allow her to roll back to the choice labelled by $\{C\}$ in the second interaction.

As regards messages, we already observed that we allow different branches of a choice to have equal senders, equal messages and equal receivers. Our definition of projection ensures that the behaviour of a participant is either independent of the chosen branch or identifiable through some communications. For parallel global types, instead, we require that different branches do not have equal senders, equal messages and equal receivers. Indeed, if we allowed

$$\begin{array}{c}
\frac{\Gamma \vdash P_i : \mathbb{T}_i \ (i \in I)}{\Gamma \vdash \delta \sum_{i \in I} \pi_i; P_i : \delta \bigwedge_{i \in I} \pi_i; \mathbb{T}_i} \text{[T-EXTCH]} \quad \frac{\Gamma \vdash P_i : \mathbb{T}_i \ (i \in I)}{\Gamma \vdash \Delta \bigoplus_{i \in I} \pi_i; P_i : \Delta \bigvee_{i \in I} \pi_i; \mathbb{T}_i} \text{[T-INTCH]} \\
\\
\frac{\Gamma \vdash P_i : \mathbb{T}_i \ (i \in I) \quad \Gamma \vdash P : \mathbb{T}}{\Gamma \vdash \delta \widehat{\sum}_{i \in I} \pi_i; P_i \widehat{+} \widehat{\pi}; P : \delta \widehat{\bigwedge}_{i \in I} \pi_i; \mathbb{T}_i \widehat{\wedge} \widehat{\pi}; \mathbb{T}} \text{[T-EXTCH-RT]} \\
\\
\frac{\Gamma \vdash P_i : \mathbb{T}_i \ (i \in I) \quad \Gamma \vdash P : \mathbb{T}}{\Gamma \vdash \Delta \widehat{\bigoplus}_{i \in I} \pi_i; P_i \widehat{\oplus} \widehat{\pi}; P : \Delta \widehat{\bigvee}_{i \in I} \pi_i; \mathbb{T}_i \widehat{\vee} \widehat{\pi}; \mathbb{T}} \text{[T-INTCH-RT]} \\
\\
\frac{\vdash P_i : \mathbf{G} \upharpoonright \mathfrak{p}_i \quad (1 \leq i \leq n) \quad \mathfrak{pa}(\mathbf{G}) \subseteq \{\mathfrak{p}_1, \dots, \mathfrak{p}_n\}}{\vdash \mathfrak{p}_1 \llbracket P_1 \rrbracket \parallel \dots \parallel \mathfrak{p}_n \llbracket P_n \rrbracket : \mathbf{G}} \text{[T-NET]}
\end{array}$$

■ **Figure 4** Main typing rules for processes and networks.

two parallel communications with equal messages between the same pair of participants, then we could have a global type $\mathfrak{p} \xrightarrow{\lambda} \mathfrak{q}; \mathbf{G}_1 \parallel \mathfrak{p} \xrightarrow{\lambda} \mathfrak{q}; \mathbf{G}_2$, and processes implementing participants \mathfrak{p} and \mathfrak{q} could “cross” their communications.

A global type meeting the above two requirements on messages and checkpoint labels is called respectively *message well formed* and *checkpoint label well formed*.

Last but not least, following [7, 26] we require that the order of communications prescribed by a global type be witnessed by at least one of the session participants. This excludes, for instance, the global type $\mathfrak{p} \xrightarrow{\lambda} \mathfrak{q}; \mathfrak{r} \xrightarrow{\lambda'} \mathfrak{s}$, since the two communications have disjoint sets of participants, and therefore they should be independent. The correct global type representing this situation is $\mathfrak{p} \xrightarrow{\lambda} \mathfrak{q} \parallel \mathfrak{r} \xrightarrow{\lambda'} \mathfrak{s}$. This condition is necessary for a global type to be *implementable* by a collection of processes, and we refer to [7] for further justifications of this choice. It is easy to formalise this requirement as a well-formedness condition on the set of traces generated by global types, defined in the standard way.

To sum up, we define well-formedness of global types as follows.

► **Definition 7** (Well-formed global types). *A global type is well formed* when:

1. Its projections are defined on all participants;
2. It is sequentially well formed, message well formed and checkpoint label well formed;
3. Its set of traces is well formed.

In the following we will consider only well-formed global types.

5 Type System and Soundness

The shape of *typing judgements* is $\Gamma \vdash P : \mathbb{T}$, where the environment Γ associates process variables with session types: $\Gamma ::= \emptyset \mid \Gamma, X : \mathbb{T}$. Process typing exploits the correspondence between external choices and intersections, internal choices and unions. Figure 4 shows the interesting rules. Typing respects hats, in rules [T-EXTCH], [T-INTCH], [T-EXTCH-RT], [T-INTCH-RT] intersections and unions have hats if choices have hats. Subtyping \leq on session types takes into account the standard rules for intersection and union and preserves hats, while checkpoint sets can decrease with deletion of types in intersections and increase with addition of types in unions. Subtyping is used in a standard subsumption rule (omitted). The other omitted rules for processes are just homomorphisms as expected from the syntax of processes and session types. Rule [T-NET] is the only rule for typing networks: it

requires that the types of all processes be projections of a unique global type. The condition $\text{pa}(\mathbf{G}) \subseteq \{\mathfrak{p}_1, \dots, \mathfrak{p}_n\}$ ensures the presence of all session participants and allows the typing of sessions containing $\mathfrak{p}[\text{skip}]$ for any \mathfrak{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 8.** Let $P = \mu X. \mathfrak{q}!\lambda_1; r!\lambda_2; X_{\{C\} \oplus \mathfrak{q}!\lambda_3; r!\lambda_4; X}$ and $Q = \mu Y. \mathfrak{p}?\lambda_1; Y_{\{C\} + \mathfrak{p}?\lambda_3; Y}$ and $R = \mu Z. \mathfrak{p}?\lambda_2; Z_{\{C\} + \mathfrak{p}?\lambda_4; Z}$. The network $\mathfrak{p}[P] \parallel \mathfrak{q}[Q] \parallel r[R]$ reduces by forward reductions to $\mathfrak{p}[P'] \parallel \mathfrak{q}[Q'] \parallel r[R']$ where $P' = \widehat{\mathfrak{q}!\lambda_1}; \widehat{r!\lambda_2}; (\mathfrak{q}!\lambda_1; r!\lambda_2; P_{\{C\} \oplus \mathfrak{q}!\lambda_3; r!\lambda_4; P})_{\{C\} \oplus \mathfrak{q}!\lambda_3; r!\lambda_4; P}$ and $Q' = \widehat{\mathfrak{p}?\lambda_1}; (\mathfrak{p}?\lambda_1; Q_{\{C\} + \mathfrak{p}?\lambda_3; Q})_{\{C\} + \mathfrak{p}?\lambda_3; Q}$ and $R' = \widehat{\mathfrak{p}?\lambda_2}; R_{\{C\} + \mathfrak{p}?\lambda_4; R}$. By Rule [BACK], $\mathfrak{p}[P'] \parallel \mathfrak{q}[Q'] \parallel r[R'] \xrightarrow{C} \mathfrak{p}[P''] \parallel \mathfrak{q}[Q] \parallel r[R]$ where $P'' = \mathfrak{q}!\lambda_3; r!\lambda_4; P$. Without the condition $\mathcal{E} \text{ ok}$ for \overline{C} on rule [CTBA], we could have also the backward move: $\mathfrak{p}[P'] \parallel \mathfrak{q}[Q'] \parallel r[R'] \xrightarrow{C} \mathfrak{p}[P'''] \parallel \mathfrak{q}[Q''] \parallel r[R]$, where $P''' = (\widehat{\mathfrak{q}!\lambda_1}; \widehat{r!\lambda_2}; \mathfrak{q}!\lambda_1; r!\lambda_2; P)_{\{C\} \oplus \mathfrak{q}!\lambda_3; r!\lambda_4; P}$ and $Q'' = \widehat{\mathfrak{p}?\lambda_1}; Q_{\{C\} + \mathfrak{p}?\lambda_3; Q}$. Then Subject Reduction would fail, since the network $\mathfrak{p}[P] \parallel \mathfrak{q}[Q] \parallel r[R]$ is typable, while $\mathfrak{p}[P'''] \parallel \mathfrak{q}[Q''] \parallel r[R]$ is not typable. In fact the output $r!\lambda_2$ has a hat in the session type of P''' , while the corresponding input $\mathfrak{p}?\lambda_2$ 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 last checkpoints in recursive processes, since they could be different for different participants.

It is easy to verify that the type \mathbf{T}^S given at page 11 can be derived for the process P^S given at the end of Section 3.

The present type system is not informative enough: $\vdash \mathbf{N} : \mathbf{G}$ does not imply that a communication in \mathbf{G} can be done in a forward computation of \mathbf{N} . For example, consider $\vdash \mathbf{N}_0 : \mathbf{G}_0$ where $\mathbf{N}_0 = \mathfrak{p}[\mathfrak{q}!\lambda] \parallel \mathfrak{q}[\mathfrak{p}?\lambda + \mathfrak{p}?\lambda']$ and $\mathbf{G}_0 = \mathfrak{p} \xrightarrow{\lambda} \mathfrak{q} \boxplus \mathfrak{p} \xrightarrow{\lambda'} \mathfrak{q}$. The communication $\mathfrak{p} \xrightarrow{\lambda'} \mathfrak{q}$ cannot occur in \mathbf{N}_0 . To discuss properties, we introduce a type system which represents more closely the evolution of networks.

We write $\Gamma \vdash^* P : \mathbf{T}$ if this judgement can be derived using the typing rules of Figure 4 without the help of the subtyping rule. Let \leq^* be the subtyping relation between session types obtained from \leq by requiring that unions have the same number of disjuncts.

We write $\vdash^* \mathbf{N} : \mathbf{G}$ if this judgement can be derived using the typing rule:

$$\frac{\vdash^* P_i : \mathbf{T}_i \quad \mathbf{T}_i \leq^* \mathbf{G} \upharpoonright \mathfrak{p}_i \quad (1 \leq i \leq n) \quad \text{pa}(\mathbf{G}) \subseteq \{\mathfrak{p}_1, \dots, \mathfrak{p}_n\}}{\vdash \mathfrak{p}_1[P_1] \parallel \dots \parallel \mathfrak{p}_n[P_n] : \mathbf{G}} \text{ [T-NET}^* \text{]}$$

The relation between the type systems \vdash^* and \vdash is expressed by Theorem 9, whose proof uses standard Inversion Lemmas. The property of Subject Reduction for \vdash^* is then established by Theorem 10. Its proof relies on the close correspondence between the evolution of well-typed nets and the evolution of their types along forward and backward computations.

► **Theorem 9.** *If $\vdash^* \mathbf{N} : \mathbf{G}$, then $\vdash \mathbf{N} : \mathbf{G}$. If $\vdash \mathbf{N} : \mathbf{G}$, then $\vdash^* \mathbf{N} : \mathbf{G}'$ for some \mathbf{G}' .*

► **Theorem 10 (Subject Reduction).** *$\vdash^* \mathbf{N} : \mathbf{G}$ and $\mathbf{N} \xrightarrow{\ast} \mathbf{N}'$ imply $\vdash^* \mathbf{N}' : \mathbf{G}'$ for some \mathbf{G}' .*

Notice that Theorems 9 and 10 imply Subject Reduction for \vdash .

A standard property enforced by session types is Session Fidelity: all communications occur as specified by global types. To deal with backward reductions we introduce a mapping $\llbracket \cdot \rrbracket$ on global types, which erases all communications with hats and all discarded branches of choices. A checkpoint label C is *alive* in \mathbf{G} if $\llbracket \mathbf{G} \rrbracket$ contains a choice $\{C\} \boxplus_{i \in I} \mathfrak{p} \xrightarrow{\lambda_i} \mathfrak{q}_i; \mathbf{G}_i$ with more than one branch, such that for some $i \in I$, λ, \mathfrak{q} , the global type \mathbf{G}_i contains $\mathfrak{p} \xrightarrow{\lambda} \mathfrak{q}$.

► **Theorem 11** (Session Fidelity). *If $\vdash \mathbb{N} : G$, then all traces in forward computations of \mathbb{N} are suffixes of traces of G . Moreover if $\mathbb{N} \longrightarrow^* \mathbb{N}' \overset{C}{\curvearrowright} \mathbb{N}''$, then C is alive in G .*

We discuss now forward and backward progress.

► **Theorem 12** (Forward Progress). *If $\vdash p[\mathcal{E}[P]] \parallel \mathbb{N} : G$ with P user process, $P \neq \text{skip}$, then there are P', \mathbb{N}' such that $p[\mathcal{E}[P]] \parallel \mathbb{N} \longrightarrow^* p[\mathcal{E}[P']] \parallel \mathbb{N}'$ and P' has one hat.*

Notice that the standard formulation of progress [12], which requires that each input/output that is persistently offered be eventually consumed, is an easy consequence of this theorem.

► **Theorem 13** (Backward Progress). *If $\vdash^* \mathbb{N} : G$ and C is alive in G , then there is \mathbb{N}' such that $\mathbb{N} \longrightarrow^* \mathbb{N}'$ and $\mathbb{N}' \overset{C}{\curvearrowright} \mathbb{N}''$.*

We end this section with a remark on causal consistency and full reversibility. Clearly, our calculus enjoys causal consistency since a rollback to a checkpointed choice cancels all communications done after that checkpoint. Instead, full reversibility does not hold, since rollbacks erase explored branches in internal choices.

6 Related Work and Conclusion

Since the seminal work by Danos and Krivine on reversible CCS [9], reversible computation has been widely studied in process calculi. In [29], 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 [21], Lanese et al. defined 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 [20]. In [8], Cristescu et al. proposed a causal semantic model for the reversible π -calculus.

Reversibility for structured communications was first studied in [10, 11, 19], 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] and [3], Barbanera et al. 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.

Our work is most closely related to the recent proposals [33, 34, 23, 22, 24, 28, 26, 14]. Tiezzi and Yoshida [33] use tags and memories to allow full reversibility of binary sessions with delegation. In [34], 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 [23, 22] 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 [24], 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 [26], 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 [26] is that inputs are potentially irreversible actions, unless they appear within a loop. Finally, our paper builds on [14] and improves it in several respects: it has a more liberal syntax for types and processes, gives a more compact representation of past communications and implements a fine-tuned strategy for backward moves, geared towards achieving compliance.

It should be noted that in all the above-mentioned proposals, the syntax of global types (if any) does not include recursion, parallel composition nor sequential composition. Thus, some of our examples cannot be expressed in these models. On the other hand, the context-free session types of [32], recently proposed to capture the type-safe serialisation of recursive data types, include sequential composition and recursion. Our session types generalise them by offering in addition parallel composition, and by handling multiparty sessions.

The properties of our calculus are standard for session types, but their proofs require some ingenuity due to the generality of our syntax and to the specificity of our rollback mechanism. A limitation of our work is that the starting points of rollbacks are statically determined. By contrast, these points are determined dynamically in [26], offering a more realistic solution. We plan to introduce similar runtime conditions for rollback in our calculus.

Since our session types and global types are “truly concurrent”, we would also like to study their interpretation into some non-interleaving semantic model. A natural candidate that comes to mind is the model of Event Structures, for which reversible variants have already been proposed [30, 16]. 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 are grateful to the anonymous reviewers for their useful suggestions, which led to substantial improvements. We also thank Claudio Antares Mezzina for pointing out a technical problem in an earlier version of the paper.

References

- 1 Franco Barbanera and Ugo de’ Liguoro. Sub-behaviour relations for session-based client/server systems. *Mathematical Structures in Computer Science*, 25(6):1339–1381, 2015. doi:10.1017/S096012951400005X.
- 2 Franco Barbanera, Mariangiola Dezani-Ciancaglini, and Ugo de’Liguoro. Reversible client/server interactions. *Formal Aspects of Computing*, 28(4):697–722, 2016. doi:10.1007/s00165-016-0358-2.
- 3 Franco Barbanera, Mariangiola Dezani-Ciancaglini, Ivan Lanese, and Ugo de’ Liguoro. Retractable contracts. In *PLACES*, volume 203 of *EPTCS*, pages 61–72, 2016. doi:10.4204/EPTCS.203.
- 4 Giovanni Bernardi and Matthew Hennessy. Modelling session types using contracts. *Mathematical Structures in Computer Science*, 26(3):510–560, 2016. doi:10.1017/S0960129514000243.
- 5 Gérard Boudol and Ilaria Castellani. Permutation of transitions: an event structure semantics for CCS and SCCS. In *REX*, volume 354 of *LNCS*, pages 411–427. Springer, 1988. doi:10.1007/BFb0013028.
- 6 Gérard Boudol and Ilaria Castellani. Flow models of distributed computations: three equivalent semantics for CCS. *Information and Computation*, 114(2):247–314, 1994. doi:10.1006/inco.1994.1088.
- 7 Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, and Luca Padovani. On Global Types and Multi-Party Sessions. *Logical Methods in Computer Science*, 8:1–45, 2012. doi:10.2168/LMCS-8(1:24)2012.
- 8 Ioana Cristescu, Jean Krivine, and Daniele Varacca. Rigid families for the reversible π -calculus. In *RC*, volume 9720 of *LNCS*, pages 3–19. Springer, 2016. doi:10.1007/978-3-319-40578-0_1.
- 9 Vincent Danos and Jean Krivine. Reversible communicating systems. In *CONCUR*, volume 3170 of *LNCS*, pages 292–307. Springer, 2004. doi:10.1007/978-3-540-28644-8_19.

- 10 Edsko de Vries, Vasileios Koutavas, and Matthew Hennessy. Communicating transactions - (extended abstract). In *CONCUR*, volume 6269 of *LNCS*, pages 569–583. Springer, 2010. doi:10.1007/978-3-642-15375-4_39.
- 11 Edsko de Vries, Vasileios Koutavas, and Matthew Hennessy. Liveness of communicating transactions - (extended abstract). In *APLAS*, volume 6461 of *LNCS*, pages 392–407. Springer, 2010. doi:10.1007/978-3-642-17164-2_27.
- 12 Pierre-Malo Deniérou and Nobuko Yoshida. Dynamic multirole session types. In *POPL*, pages 435–446. ACM Press, 2011. doi:10.1145/1926385.1926435.
- 13 Pierre-Malo Deniérou and Nobuko Yoshida. Multiparty session types meet communicating automata. In *ESOP*, volume 7211 of *LNCS*, pages 194–213. Springer, 2012. doi:10.1007/978-3-642-28869-2_10.
- 14 Mariangiola Dezani-Ciancaglini and Paola Giannini. Reversible multiparty sessions with checkpoints. In *EXPRESS/SOS*, volume 222 of *EPTCS*, pages 60–74, 2016. doi:10.4204/EPTCS.222.5.
- 15 Elena Giachino, Ivan Lanese, Claudio Antares Mezzina, and Francesco Tiezzi. Causal-consistent rollback in a tuple-based language. *Journal of Logic and Algebraic Methods in Programming*, 88:99–120, 2017. URL: <https://doi.org/10.1016/j.jlamp.2016.09.003>, doi:10.1016/j.jlamp.2016.09.003.
- 16 Eva Graversen, Iain Phillips, and Nobuko Yoshida. Towards a categorical representation of reversible event structures. In *PLACES*, volume 246 of *EPTCS*, pages 49–60, 2017. doi:10.4204/EPTCS.246.9.
- 17 Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP*, volume 1381 of *LNCS*, pages 22–138. Springer, 1998. doi:10.1007/BFb0053567.
- 18 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In *POPL*, pages 273–284. ACM Press, 2008. doi:10.1145/1328897.1328472.
- 19 Vasileios Koutavas, Carlo Spaccasassi, and Matthew Hennessy. Bisimulations for communicating transactions - (extended abstract). In *FOSSACS*, volume 8412 of *LNCS*, pages 320–334. Springer, 2014. doi:10.1007/978-3-642-54830-7_21.
- 20 Ivan Lanese, Claudio Antares Mezzina, Alan Schmitt, and Jean-Bernard Stefani. Controlling reversibility in higher-order pi. In *CONCUR*, volume 6901 of *LNCS*, pages 297–311. Springer, 2011. doi:10.1007/978-3-642-23217-6_20.
- 21 Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Reversing higher-order pi. In *CONCUR*, volume 6269 of *LNCS*, pages 478–493. Springer, 2010. doi:10.1007/978-3-642-15375-4_33.
- 22 Claudio Antares Mezzina and Jorge A. Pérez. Reversible semantics in session-based concurrency. In *ICTCS*, volume 1720 of *CEUR Workshop Proceedings*, pages 221–226. CEUR-WS.org, 2016.
- 23 Claudio Antares Mezzina and Jorge A. Pérez. Reversible sessions using monitors. In *PLACES*, volume 211 of *EPTCS*, pages 56–64, 2016. doi:10.4204/EPTCS.211.6.
- 24 Claudio Antares Mezzina and Jorge A. Pérez. Causally consistent reversible choreographies. *CoRR*, abs/1703.06021, 2017.
- 25 Claudio Antares Mezzina, Rudolf Schlatte, and al. COST Action on Reversible Computation, State of the Art Report, Working Group 2 on Software and Systems, 2012. URL: http://www.informatik.uni-bremen.de/ictcost/wg2_soar.pdf.
- 26 Claudio Antares Mezzina and Emilio Tuosto. Choreographies for automatic recovery. *CoRR*, abs/1705.09525, 2017.
- 27 Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *LNCS*. Springer, 1980. doi:10.1007/3-540-10235-3.

- 28 Rumyana Neykova and Nobuko Yoshida. Let it recover: multiparty protocol-induced recovery. In *CC*, pages 98–108. ACM Press, 2017. doi:10.1145/3033019.
- 29 Iain Phillips and Irek Ulidowski. Reversing algebraic process calculi. *Journal of Logic and Algebraic Methods in Programming*, 73(1-2):70–96, 2007. doi:10.1016/j.jlap.2006.11.002.
- 30 Iain Phillips and Irek Ulidowski. Reversibility and asymmetric conflict in event structures. *Journal of Logic and Algebraic Methods in Programming*, 84(6):781–805, 2015. doi:10.1016/j.jlamp.2015.07.004.
- 31 Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- 32 Peter Thiemann and Vasco T. Vasconcelos. Context-free session types. In *ICFP*, pages 462–475. ACM Press, 2016. doi:10.1145/2951913.2951926.
- 33 Francesco Tiezzi and Nobuko Yoshida. Reversible session-based pi-calculus. *Journal of Logical and Algebraic Methods in Programming*, 84(5):684–707, 2015. doi:10.1016/j.jlamp.2015.03.004.
- 34 Francesco Tiezzi and Nobuko Yoshida. Reversing single sessions. In *RC*, volume 9720 of *LNCS*, pages 52–69. Springer, 2016. doi:10.1007/978-3-319-40578-0_4.