

Retractable Contracts*

Franco Barbanera

Dipartimento di Matematica e Informatica, University of Catania,
barba@dmi.unict.it

Mariangiola Dezani-Ciancaglini

Dipartimento di Informatica, University of Torino, †
dezani@di.unito.it

Ivan Lanese

Dipartimento di Informatica - Scienza e Ingegneria, University of Bologna/INRIA, ‡
ivan.lanese@gmail.com

Ugo de'Liguoro

Dipartimento di Informatica, University of Torino, §
deliguoro@di.unito.it

In calculi for modelling communication protocols, internal and external choices play dual roles. Two external choices can be viewed naturally as dual too, as they represent an agreement between the communicating parties. If the interaction fails, the past agreements are good candidates as points where to roll back, in order to take a different agreement. We propose a variant of contracts with synchronous rollbacks to agreement points in case of deadlock. The new calculus is equipped with a compliance relation which is shown to be decidable.

1 Introduction

In human as well as automatic negotiations, an interesting feature is the ability of rolling back to some previous point in case of failure, undoing previous choices and possibly trying a different path. *Rollbacks* are familiar to the users of web browsers, and so are also the troubles that these might cause during “undisciplined” interactions. Clicking the “back” button, or going to some previous point in the chronology when we are in the middle of a transaction, say the booking of a flight, can be as smart as dangerous. In any case, it is surely a behaviour that service programmers want to discipline. Also the converse has to be treated with care: a server discovering that an auxiliary service becomes available after having started a conversation could take advantage of it using some kind of rollback. However, such a server would be quite unfair if the rollback were completely hidden from the client.

Let us consider an example. A Buyer is looking for a bag ($\overline{\text{bag}}$) or a belt ($\overline{\text{belt}}$); she will decide how to pay, either by credit card ($\overline{\text{card}}$) or by cash ($\overline{\text{cash}}$), after knowing the price from a Seller. The Buyer behaviour can be described by the process:

$$\text{Buyer} = \overline{\text{bag}}.\text{price}.\overline{(\text{card} \oplus \text{cash})} \oplus \overline{\text{belt}}.\text{price}.\overline{(\text{card} \oplus \text{cash})}$$

*This work was partially supported by Italian MIUR PRIN Project CINA Prot. 2010LHT4KM and COST Action IC1201 BETTY.

†This author was partially supported by the Torino University/Compagnia San Paolo Project SALT.

‡This author was partially supported by the French ANR project REVER n. ANR 11 INSE 007 and COST Action IC1405.

§This author was partially supported by the Torino University/Compagnia San Paolo Project SALT.

where dot is sequential composition and \oplus is internal choice. The Seller does not accept credit card payments for items of low price, like belts, but only for more expensive ones, like bags:

$$\text{Seller} = \overline{\text{belt.price}}.\text{cash} + \text{bag}.\overline{\text{price}}.(\text{card} + \text{cash})$$

where $+$ is external choice. According to contract theory [6], Buyer is not compliant with Seller, since she can choose to pay the belt by card. Also, there is no obvious way to represent the buyer's will to be free in her decision about the payment and be compliant with a seller without asking the seller in advance. Nonetheless, when interacting with Seller, the buyer's decision is actually free at least in the case of purchase of a bag. For exploiting such a possibility the client (but also the server) should be able to tolerate a partial failure of her protocol, and to try a different path.

To this aim we add to (some) choices a possibility of rollback, in case the taken path fails to reach a success configuration. In this setting, choices among outputs are no more purely internal, since the environment may oblige to undo a wrong choice and choose a different alternative. For this reason, we denote choices between outputs which allow rollback as external, hence we use Buyer' below instead of Buyer:

$$\text{Buyer}' = \overline{\text{bag.price}}.(\overline{\text{card}} \oplus \overline{\text{cash}}) + \overline{\text{belt.price}}.(\overline{\text{card}} \oplus \overline{\text{cash}})$$

We thus explore a model of contract interaction in which synchronous rollback is triggered when client and server fail to reach an agreement.

In defining our model we build over some previous work reported in [2], where we have considered contracts with rollbacks. However, we depart from that model on three main aspects. First, in the present model rollback is used in a disciplined way to tolerate failures in the interaction, thus improving compatibility, while in [2] it is an internal decision of either client or server, which makes compatibility more difficult. Second, we embed checkpoints in the structure of contracts, avoiding explicit checkpoints. Third, we consider a stack of "pasts", called histories, instead of just one past for each participant, as in [2], thus allowing to undo many past choices looking for a successful alternative.

2 Contracts for retractable interactions

Our contracts can be obtained from the session behaviours of [1] or from the session contracts of [3] just adding external retractable choices between outputs.

Definition 2.1 (Retractable Contracts). *Let \mathcal{N} (set of names) be some countable set of symbols and $\overline{\mathcal{N}} = \{\overline{a} \mid a \in \mathcal{N}\}$ (set of conames), with $\mathcal{N} \cap \overline{\mathcal{N}} = \emptyset$. The set RC of **retractable contracts** is defined as the set of the **closed** expressions generated by the following grammar,*

$$\begin{array}{l} \sigma, \rho \quad := \quad | \mathbf{1} \quad \text{success} \\ \quad \quad \quad | \sum_{i \in I} a_i \cdot \sigma_i \quad \text{(retractable) input} \\ \quad \quad \quad | \sum_{i \in I} \overline{a}_i \cdot \sigma_i \quad \text{retractable output} \\ \quad \quad \quad | \bigoplus_{i \in I} \overline{a}_i \cdot \sigma_i \quad \text{unretractable output} \\ \quad \quad \quad | x \quad \text{variable} \\ \quad \quad \quad | \text{rec}.x.\sigma \quad \text{recursion} \end{array}$$

where I is non-empty and finite, the names and the conames in choices are pairwise distinct and σ is not a variable in $\text{rec}.x.\sigma$.

Note that recursion in RC is guarded and hence contractive in the usual sense. We take an equi-recursive view of recursion by equating $\text{rec}.x.\sigma$ with $\sigma[\text{rec}.x.\sigma/x]$. We use α to range over $\mathcal{N} \cup \overline{\mathcal{N}}$, with the

$$\text{convention } \overline{\alpha} = \begin{cases} \overline{a} & \text{if } \alpha = a, \\ a & \text{if } \alpha = \overline{a}. \end{cases}$$

We write $\alpha_1.\sigma_1 + \alpha_2.\sigma_2$ for binary input/retractable output and $\overline{\alpha}_1.\sigma_1 \oplus \overline{\alpha}_2.\sigma_2$ for binary unretractable output. They are both commutative by definition. Also, $\overline{a}.\sigma$ may denote both unary retractable output and unary unretractable output. This is not a source of confusion since they have the same semantics.

From now on we call just *contracts* the expressions in RC. They are written by omitting all trailing **1**'s.

In order to deal with rollbacks we decorate contracts with histories, which memorise the alternatives in choices which have been discharged. We use 'o' as a placeholder for *no-remaining-alternatives*.

Definition 2.2 (Contracts with histories). *Let Histories be the expressions (referred to also as stacks) generated by the grammar:*

$$\gamma ::= [] \mid \gamma : \sigma$$

where $\sigma \in \text{RC} \cup \{\circ\}$ and $\circ \notin \text{RC}$. Then the set of contracts with histories is defined by:

$$\text{RCH} = \{ \gamma \prec \sigma \mid \gamma \in \text{Histories}, \sigma \in \text{RC} \cup \{\circ\} \}.$$

We write just $\sigma_1 : \dots : \sigma_k$ for the stack $(\dots ([:\sigma_1] : \dots) : \sigma_k)$. With a little abuse of notation we use ':' also to concatenate histories, and to add contracts in front of histories.

We can now discuss the operational semantics of our calculus (Definition 2.3). The reduction rule for the internal choice (\oplus) is standard, but for the presence of the $\gamma \prec \cdot$. Whereas, when reducing retractable choices ($+$), the discharged branches are memorised. When a single action is executed, the history is modified by adding a 'o', intuitively meaning that the only possible branch has been tried and no alternative is left. Rule (rb) recovers the contract on the top of the stack, replacing the current one with it.

Definition 2.3 (LTS of Contracts with Histories).

$$\begin{array}{ll} (+) \quad \gamma \prec \alpha.\sigma + \sigma' \xrightarrow{\alpha} \gamma : \sigma' \prec \sigma & (\oplus) \quad \gamma \prec \overline{a}.\sigma \oplus \sigma' \xrightarrow{\tau} \gamma \prec \overline{a}.\sigma \\ (\alpha) \quad \gamma \prec \alpha.\sigma \xrightarrow{\alpha} \gamma : \circ \prec \sigma & (\text{rb}) \quad \gamma : \sigma' \prec \sigma \xrightarrow{\text{rb}} \gamma \prec \sigma' \end{array}$$

The interaction of a client with a server is modelled by the reduction of their parallel composition, that can be either forward, consisting of CCS style synchronisations and single internal choices, or backward, only when there is no possible forward reduction, and the client is not satisfied, i.e. it is different from **1**.

Definition 2.4 (LTS of Client/Server Pairs). *We define the relation \longrightarrow over pairs of contracts with histories by the following rules:*

$$\begin{array}{c} \frac{\delta \prec \rho \xrightarrow{\alpha} \delta' \prec \rho' \quad \gamma \prec \sigma \xrightarrow{\overline{\alpha}} \gamma' \prec \sigma'}{\delta \prec \rho \parallel \gamma \prec \sigma \longrightarrow \delta' \prec \rho' \parallel \gamma' \prec \sigma'} \text{ (comm)} \\ \\ \frac{\delta \prec \rho \xrightarrow{\tau} \delta \prec \rho'}{\delta \prec \rho \parallel \gamma \prec \sigma \longrightarrow \delta \prec \rho' \parallel \gamma \prec \sigma} \text{ (\tau)} \\ \\ \frac{\gamma \prec \rho \xrightarrow{\text{rb}} \gamma' \prec \rho' \quad \delta \prec \sigma \xrightarrow{\text{rb}} \delta' \prec \sigma' \quad \rho \neq \mathbf{1}}{\gamma \prec \rho \parallel \delta \prec \sigma \longrightarrow \gamma' \prec \rho' \parallel \delta' \prec \sigma'} \text{ (rbk)} \end{array}$$

plus the rule symmetric to (τ) w.r.t. \parallel . Moreover, rule (rbk) applies only if neither (comm) nor (τ) do.

We will use $\xrightarrow{*}$ and $\not\rightarrow$ with the standard meanings.

Notice that, since ‘ \circ ’ cannot synchronise with anything, in case a partner rolls back to a ‘ \circ ’, it is forced to recover an *older* past (if any).

The following examples show the different behaviours of retractable and unretractable outputs. We decorate arrows with the name of the used reduction rule. As a first example we consider a possible reduction of the process discussed in the Introduction.

Example 2.5. As in the Introduction, let $\text{Buyer}' = \overline{\text{bag}}.\text{price}.\overline{(\text{card} \oplus \text{cash})} + \overline{\text{belt}}.\text{price}.\overline{(\text{card} \oplus \text{cash})}$ be a client and $\text{Seller} = \overline{\text{belt}}.\text{price}.\text{cash} + \overline{\text{bag}}.\text{price}.\overline{(\text{card} + \text{cash})}$ a server; then

$$\begin{array}{l}
\begin{array}{c}
\text{comm} \\
\text{comm} \\
\tau \\
\text{rbk} \\
\text{rbk} \\
\text{comm} \\
\text{comm} \\
\tau \\
\text{comm} \\
\not\rightarrow
\end{array}
\begin{array}{c}
\overline{\text{bag}}.\text{price}.\overline{(\text{card} \oplus \text{cash})} \prec \text{price}.\overline{(\text{card} \oplus \text{cash})} \\
\overline{\text{bag}}.\text{price}.\overline{(\text{card} \oplus \text{cash})} : \circ \prec \overline{(\text{card} \oplus \text{cash})} \\
\overline{\text{bag}}.\text{price}.\overline{(\text{card} \oplus \text{cash})} : \circ \prec \overline{\text{card}} \\
\overline{\text{bag}}.\text{price}.\overline{(\text{card} \oplus \text{cash})} \prec \circ \\
[] \prec \overline{\text{bag}}.\text{price}.\overline{(\text{card} \oplus \text{cash})} \\
\circ \prec \text{price}.\overline{(\text{card} \oplus \text{cash})} \\
\circ : \circ \prec \overline{(\text{card} \oplus \text{cash})} \\
\circ : \circ \prec \overline{\text{card}} \\
\circ : \circ : \circ \prec \mathbf{1}
\end{array}
\parallel
\begin{array}{c}
[] \prec \text{Seller} \\
\overline{\text{bag}}.\text{price}.\overline{(\text{card} + \text{cash})} \prec \overline{\text{price}}.\text{cash} \\
\overline{\text{bag}}.\text{price}.\overline{(\text{card} + \text{cash})} : \circ \prec \text{cash} \\
\overline{\text{bag}}.\text{price}.\overline{(\text{card} + \text{cash})} : \circ \prec \text{cash} \\
\overline{\text{bag}}.\text{price}.\overline{(\text{card} + \text{cash})} \prec \circ \\
[] \prec \overline{\text{bag}}.\text{price}.\overline{(\text{card} + \text{cash})} \\
\circ \prec \overline{\text{price}}.\overline{(\text{card} + \text{cash})} \\
\circ : \circ \prec \overline{(\text{card} + \text{cash})} \\
\circ : \circ \prec \overline{(\text{card} + \text{cash})} \\
\circ : \circ : \text{cash} \prec \mathbf{1}
\end{array}
\end{array}$$

Example 2.6. Let $\rho = \text{rec } x.(\overline{b}.x \oplus \overline{a}.c.x)$ and $\sigma = \text{rec } x.(b.x + a.\overline{e}.x)$. The following reduction sequence leads the parallel composition of these contracts to a deadlock.

$$\begin{array}{l}
\rho \parallel \sigma \xrightarrow{\tau} [] \prec \overline{a}.c.\rho \parallel [] \prec \text{rec } x.(b.x + a.\overline{e}.x) \\
\text{comm} \xrightarrow{\quad} \circ \prec c.\rho \parallel b.\sigma \prec \overline{e}.\sigma \\
\text{rbk} \xrightarrow{\quad} [] \prec \circ \parallel [] \prec b.\sigma \\
\not\rightarrow
\end{array}$$

Example 2.7. Let us now modify the above example by using retractable outputs in the client, so making the two contracts in parallel always reducible. The following reduction shows that there can be an infinite number of rollbacks in a sequence, even if it is not possible to have an infinite reduction containing only rollbacks. Notice how the stack keeps growing indefinitely.

Notice that rule $(+, +)$ implicitly represents the fact that, in the decision procedure for two contracts made of retractable choices, the possible synchronising branches have to be tried, until either a successful one is found or all fail.

Example 3.3. Let us formally show that, for the Buyer' and Seller of the Introduction, we have Buyer' \dashv Seller.

For the sake of readability, let

$\Gamma' = \text{Buyer}' \dashv \text{Seller}$, $\text{price}.\overline{\text{card}} \oplus \overline{\text{cash}} \dashv \text{price}.\overline{\text{card}} + \text{price}.\overline{\text{cash}}$ and $\Gamma'' = \Gamma'$, $\overline{\text{card}} \oplus \overline{\text{cash}} \dashv \text{card} + \text{cash}$

$$\frac{\frac{\frac{\Gamma'' \triangleright \mathbf{1} \dashv \mathbf{1}}{\Gamma' \triangleright \overline{\text{card}} \oplus \overline{\text{cash}} \dashv \text{card} + \text{cash}} \text{(Ax)} \quad \frac{\Gamma'' \triangleright \mathbf{1} \dashv \mathbf{1}}{\Gamma' \triangleright \overline{\text{card}} \oplus \overline{\text{cash}} \dashv \text{card} + \text{cash}} \text{(Ax)}}{\Gamma' \triangleright \overline{\text{card}} \oplus \overline{\text{cash}} \dashv \text{card} + \text{cash}} \text{(\oplus, +)}}{\text{Buyer}' \dashv \text{Seller} \triangleright \text{price}.\overline{\text{card}} \oplus \overline{\text{cash}} \dashv \text{price}.\overline{\text{card}} + \text{price}.\overline{\text{cash}} \text{(\oplus, +)}} \text{(\oplus, +)}$$

$$\frac{\text{Buyer}' \dashv \text{Seller} \triangleright \text{price}.\overline{\text{card}} \oplus \overline{\text{cash}} \dashv \text{price}.\overline{\text{card}} + \text{price}.\overline{\text{cash}} \text{(\oplus, +)}}{\triangleright \text{Buyer}' \dashv \text{Seller}} \text{(\oplus, +)}$$

Example 3.4. The contracts of Example 2.7 can be formally proved to be compliant by means of the following derivation in our formal system. Actually such a derivation can be looked at as the result of the decision procedure implicitly described by the formal system.

$$\frac{\frac{\overline{b}.\rho + \overline{a}.c.\rho \dashv b.\sigma + a.\overline{e}.\sigma \triangleright \rho \dashv \sigma}{\triangleright \overline{b}.\rho + \overline{a}.c.\rho \dashv b.\sigma + a.\overline{e}.\sigma} \text{(HYP)}}{\triangleright \overline{b}.\rho + \overline{a}.c.\rho \dashv b.\sigma + a.\overline{e}.\sigma} \text{(\oplus, +)}$$

In applying the rules we exploit the fact that we consider contracts modulo recursion fold/unfold.

We can show that derivability in this formal system is decidable, since it is syntax directed and it does not admit infinite derivations.

We denote by \mathcal{D} a derivation in the system of Definition 3.2. The procedure **Prove** in Figure 1 clearly implements the formal system, that is it is straightforward to check the following

Fact 3.5.

- i) **Prove**($\Gamma \triangleright \rho \dashv \sigma$) \neq **fail** iff $\Gamma \triangleright \rho \dashv \sigma$.
- ii) **Prove**($\Gamma \triangleright \rho \dashv \sigma$) = $\mathcal{D} \neq$ **fail** implies $\Gamma \triangleright \rho \dashv \sigma$

Theorem 3.6. *Derivability in the formal system is decidable.*

Proof. By Fact 3.5, we only need to show that the procedure **Prove** always terminates. Notice that in all recursive calls **Prove**($\Gamma, \rho \dashv \sigma \triangleright \rho_k \dashv \sigma_k$) inside **Prove**($\Gamma \triangleright \rho \dashv \sigma$) the expressions ρ_k and σ_k are subexpressions of ρ and σ respectively (because of unfolding of recursion they can also be ρ and σ). Since contract expressions generate regular trees, there are only finitely many such subexpressions. This implies that the number of different calls of procedure **Prove** is always finite. \square

In the remaining of this section we will show the soundness and the completeness of the formal system using some auxiliary lemmas.

Prove($\Gamma \triangleright \rho \dashv \sigma$)

if $\rho = \mathbf{1}$ **then** $\frac{}{\Gamma \triangleright \mathbf{1} \dashv \sigma}$ (Ax)

else if $\rho \dashv \sigma \in \Gamma$ **then** $\frac{}{\Gamma, \rho \dashv \sigma \triangleright \rho \dashv \sigma}$ (HYP)

else if $\rho = \sum_{i \in I} \alpha_i \cdot \rho_i$ **and** $\sigma = \sum_{j \in J} \bar{\alpha}_j \cdot \sigma_j$
and exists $k \in I \cap J$ **s.t.** $\mathcal{D} = \mathbf{Prove}(\Gamma, \rho \dashv \sigma \triangleright \rho_k \dashv \sigma_k) \neq \mathbf{fail}$
then $\frac{\mathcal{D}}{\Gamma \triangleright \rho \dashv \sigma}$ (+, +) **else fail**

else if $\rho = \bigoplus_{i \in I} \bar{\alpha}_i \cdot \rho_i$ **and** $\sigma = \sum_{j \in J} \alpha_j \cdot \sigma_j$ **and** $I \subseteq J$
and for all $k \in I$ $\mathcal{D}_k = \mathbf{Prove}(\Gamma, \rho \dashv \sigma \triangleright \rho_k \dashv \sigma_k) \neq \mathbf{fail}$
then $\frac{\forall k \in I \mathcal{D}_k}{\Gamma \triangleright \rho \dashv \sigma}$ (\oplus , +)

else if $\rho = \sum_{i \in I} \alpha_i \cdot \rho_i$ **and** $\sigma = \bigoplus_{j \in J} \bar{\alpha}_j \cdot \sigma_j$ **and** $I \supseteq J$
and for all $k \in J$ $\mathcal{D}_k = \mathbf{Prove}(\Gamma, \rho \dashv \sigma \triangleright \rho_k \dashv \sigma_k) \neq \mathbf{fail}$
then $\frac{\forall k \in J \mathcal{D}_k}{\Gamma \triangleright \rho \dashv \sigma}$ (+, \oplus) **else fail**

else fail

Figure 1: The procedure **Prove**.

Soundness It is useful to show that if a configuration is stuck, then both histories are empty. This is a consequence of the fact that the property “the histories of client and server have the same length” is preserved by reductions.

Lemma 3.7. *If $\delta \prec \rho' \parallel \gamma \prec \sigma' \not\rightarrow$, then $\delta = \gamma = []$.*

Proof. Clearly $\delta \prec \rho' \parallel \gamma \prec \sigma' \not\rightarrow$ implies either $\delta = []$ or $\gamma = []$. Observe that:

- rule (comm) adds one element to both stacks;
- rule (τ) does not modify both stacks;
- rule (rbk) removes one element from both stacks.

Then starting from two stacks containing the same number of elements, the reduction always produces two stacks containing the same number of elements. So $\delta = []$ implies $\gamma = []$ and vice versa. \square

Next we state a lemma relating the logical rules of the formal system for compliance to the reduction rules; notice that the formers do not mention stacks in their judgments.

Lemma 3.8. *If the following is an instance of rule (+, +), or (\oplus , +), or (+, \oplus):*

$$\frac{\Gamma, \rho \dashv \sigma \triangleright \rho_1 \dashv \sigma_1 \quad \dots \quad \Gamma, \rho \dashv \sigma \triangleright \rho_n \dashv \sigma_n}{\Gamma \triangleright \rho \dashv \sigma}$$

then for all δ, γ and for all $i = 1, \dots, n$ there exist δ_i, γ_i such that

$$\delta \prec \rho \parallel \gamma \prec \sigma \xrightarrow{*} \delta_i \prec \rho_i \parallel \gamma_i \prec \sigma_i$$

and rule (rbk) is not used, namely no rollback occurs.

Proof. By inspection of the deduction and reduction rules. \square

Theorem 3.9 (Soundness). *If $\triangleright \rho \dashv \sigma$, then $\rho \dashv \sigma$.*

Proof. The proof is by contradiction. Assume $\rho \not\dashv \sigma$. Then there is a reduction

$$[] \prec \rho \parallel [] \prec \sigma \xrightarrow{*} [] \prec \rho' \parallel [] \prec \sigma' \not\rightarrow$$

with $\rho' \neq \mathbf{1}$. Note that both the histories are empty thanks to Lemma 3.7. We proceed by induction on the number n of steps in the reduction.

Let us consider the base case ($n = 0$). In this case $\rho \neq \mathbf{1}$ and there is no possible synchronization. Rule Ax is not applicable since $\rho \neq \mathbf{1}$. Rule HYP is not applicable since Γ is empty. The other rules are not applicable otherwise we would have a possible synchronization.

Let us consider the inductive case. We have a case analysis on the topmost operators in ρ and σ . Let us start with the case where both topmost operators are retractable sums, i.e., $\rho = a.\rho_k + \rho''$ and $\sigma = \bar{a}.\sigma_k + \sigma''$. Thus,

$$[] \prec a.\rho_k + \rho'' \parallel [] \prec \bar{a}.\sigma_k + \sigma'' \xrightarrow{*} [] \prec \rho'' \parallel [] \prec \sigma'' \xrightarrow{*} [] \prec \rho' \parallel [] \prec \sigma' \not\rightarrow.$$

By definition $\rho'' \not\dashv \sigma''$, and since this requires a reduction of length $< n$, by inductive hypothesis $\not\rightarrow \rho'' \dashv \sigma''$. Also the above reduction begins by:

$$[] \prec a.\rho_k + \rho'' \parallel [] \prec \bar{a}.\sigma_k + \sigma'' \longrightarrow \rho'' \prec \rho_k \parallel \sigma'' \prec \sigma_k \xrightarrow{*} \rho'' \prec \rho'_k \parallel \sigma'' \prec \sigma'_k \longrightarrow [] \prec \rho'' \parallel [] \prec \sigma''$$

for some ρ'_k, σ'_k . This implies, by the conditions on rule (rbk) and by Lemma 3.7:

$$[] \prec \rho_k \parallel [] \prec \sigma_k \xrightarrow{*} [] \prec \rho'_k \parallel [] \prec \sigma'_k \not\rightarrow.$$

It follows that $\rho_k \not\dashv \sigma_k$, and by inductive hypothesis $\not\rightarrow \rho_k \dashv \sigma_k$. We now claim that:

$$\not\rightarrow \rho'' \dashv \sigma'' \text{ and } \not\rightarrow \rho_k \dashv \sigma_k \text{ imply } \not\rightarrow a.\rho_k + \rho'' \dashv \bar{a}.\sigma_k + \sigma''$$

Toward a contradiction let us assume that $\triangleright a.\rho_k + \rho'' \dashv \bar{a}.\sigma_k + \sigma''$; then the only applicable rule is $(+, +)$, which requires both

$$a.\rho_k + \rho'' \dashv \bar{a}.\sigma_k + \sigma'' \triangleright \rho'' \dashv \sigma'' \text{ and } a.\rho_k + \rho'' \dashv \bar{a}.\sigma_k + \sigma'' \triangleright \rho_k \dashv \sigma_k.$$

Because the only difference between these statements and $\triangleright \rho'' \dashv \sigma''$ and $\triangleright \rho_k \dashv \sigma_k$ is the assumption $a.\rho_k + \rho'' \dashv \bar{a}.\sigma_k + \sigma''$, which is used only in rule HYP , there must be at least one branch of the derivation tree of $\triangleright a.\rho_k + \rho'' \dashv \bar{a}.\sigma_k + \sigma''$ ending by such a rule. By Lemma 3.8 this implies that $[] \prec a.\rho_k + \rho'' \parallel [] \prec \bar{a}.\sigma_k + \sigma'' \xrightarrow{*} [] \prec a.\rho_k + \rho'' \parallel [] \prec \bar{a}.\sigma_k + \sigma''$ by a reduction never using rule (rbk). By definition this implies that $[] \prec a.\rho_k + \rho'' \dashv [] \prec \bar{a}.\sigma_k + \sigma''$, contradicting the hypothesis.

All other cases are similar. \square

Completeness The following lemma proves that compliance is preserved by the concatenation of histories to the left of the current histories.

Lemma 3.10. *If $\delta \prec \rho \dashv \gamma \prec \sigma$, then $\delta' : \delta \prec \rho \dashv \gamma' : \gamma \prec \sigma$ for all δ', γ' .*

Proof. It suffices to show that

$$\delta \prec \rho \Vdash \gamma \prec \sigma \text{ implies } \rho' : \delta \prec \rho \Vdash \gamma \prec \sigma \text{ and } \delta \prec \rho \Vdash \sigma' : \gamma \prec \sigma$$

which we prove by contraposition.

Suppose that $\rho' : \delta \prec \rho \not\Vdash \gamma \prec \sigma$; then

$$\rho' : \delta \prec \rho \parallel \gamma \prec \sigma \xrightarrow{*} \delta' \prec \rho'' \parallel \gamma' \prec \sigma'' \not\rightarrow \text{ and } \rho'' \neq \mathbf{1}$$

If ρ' is never used, then $\delta' = \rho' : \delta''$ and $\gamma' = []$, so that we get

$$\delta \prec \rho \parallel \gamma \prec \sigma \xrightarrow{*} \delta'' \prec \rho'' \parallel [] \prec \sigma'' \not\rightarrow$$

Otherwise we have that

$$\rho' : \delta \prec \rho \parallel \gamma \prec \sigma \xrightarrow{*} \rho' \prec \rho'' \parallel \gamma' \prec \sigma'' \rightarrow [] \prec \rho' \parallel \gamma'' \prec \sigma'''$$

and we assume that $\xrightarrow{*}$ is the shortest such reduction. It follows that $\rho'' \neq \mathbf{1}$. By the minimality assumption about the length of $\xrightarrow{*}$ we know that ρ' neither has been restored by some previous application of rule (rbk), nor pushed back into the stack before. We get

$$\delta \prec \rho \parallel \gamma \prec \sigma \xrightarrow{*} [] \prec \rho'' \parallel \gamma'' \prec \sigma'' \not\rightarrow$$

In both cases we conclude that $\delta \prec \rho \not\Vdash \gamma \prec \sigma$ as desired.

Similarly we can show that $\delta \prec \rho \not\Vdash \sigma' : \gamma \prec \sigma$ implies $\delta \prec \rho \not\Vdash \gamma \prec \sigma$. \square

The following lemma gives all possible shapes of compliant contracts. It is the key lemma for the proof of completeness.

Lemma 3.11. *If $\rho \Vdash \sigma$, then one of the following conditions holds:*

1. $\rho = \mathbf{1}$;
2. $\rho = \sum_{i \in I} \alpha_i \cdot \rho_i$, $\sigma = \sum_{j \in J} \bar{\alpha}_j \cdot \sigma_j$ and $\exists k \in I \cap J$. $\rho_k \Vdash \sigma_k$;
3. $\rho = \bigoplus_{i \in I} \bar{a}_i \cdot \rho_i$, $\sigma = \sum_{j \in J} a_j \cdot \sigma_j$, $I \subseteq J$ and $\forall k \in I$. $\rho_k \Vdash \sigma_k$;
4. $\rho = \sum_{i \in I} a_i \cdot \rho_i$, $\sigma = \bigoplus_{j \in J} \bar{a}_j \cdot \sigma_j$, $I \supseteq J$ and $\forall k \in J$. $\rho_k \Vdash \sigma_k$.

Proof. By contraposition and by cases of the possible shapes of ρ and σ .

Suppose $\rho = \sum_{i \in I} \alpha_i \cdot \rho_i$, $\sigma = \sum_{j \in J} \bar{\alpha}_j \cdot \sigma_j$, $I \cap J = \{k_1, \dots, k_n\}$ and $\rho_{k_i} \not\Vdash \sigma_{k_i}$ for $1 \leq i \leq n$. Then we get

$$[] \prec \rho_{k_i} \parallel [] \prec \sigma_{k_i} \xrightarrow{*} \delta_i \prec \rho'_i \parallel \gamma_i \prec \sigma'_i \not\rightarrow$$

for $1 \leq i \leq n$, where $\rho'_i \neq \mathbf{1}$ and $\delta_i = \gamma_i = []$ by Lemma 3.7. This implies

$$\sum_{i \in I \setminus \{k_1\}} \alpha_i \cdot \rho_i \prec \rho_{k_1} \parallel \sum_{j \in J \setminus \{k_1\}} \bar{\alpha}_j \cdot \sigma_j \prec \sigma_{k_1} \xrightarrow{*} \sum_{i \in I \setminus \{k_1\}} \alpha_i \cdot \rho_i \prec \rho'_1 \parallel \sum_{j \in J \setminus \{k_1\}} \bar{\alpha}_j \cdot \sigma_j \prec \sigma'_1$$

by Lemma 3.10. Let $I' = I \setminus J$ and $J' = J \setminus I$. We can reduce $[] \prec \rho \parallel [] \prec \sigma$ only as follows:

$$\begin{aligned}
[] \prec \rho \parallel [] \prec \sigma &\longrightarrow \sum_{i \in I \setminus \{k_1\}} \alpha_i \cdot \rho_i \prec \rho_{k_1} \parallel \sum_{j \in J \setminus \{k_1\}} \bar{\alpha}_j \cdot \sigma_j \prec \sigma_{k_1} && \text{by (comm)} \\
&\xrightarrow{*} \sum_{i \in I \setminus \{k_1\}} \alpha_i \cdot \rho_i \prec \rho'_1 \parallel \sum_{j \in J \setminus \{k_1\}} \bar{\alpha}_j \cdot \sigma_j \prec \sigma'_1 \\
&\longrightarrow [] \prec \sum_{i \in I \setminus \{k_1\}} \alpha_i \cdot \rho_i \parallel [] \prec \sum_{j \in J \setminus \{k_1\}} \bar{\alpha}_j \cdot \sigma_j && \text{by (rbk)} \\
&\quad \vdots \qquad \qquad \qquad \quad \vdots \\
&\xrightarrow{*} \sum_{i \in I'} \alpha_i \cdot \rho_i \prec \rho'_n \parallel \sum_{j \in J'} \bar{\alpha}_j \cdot \sigma_j \prec \sigma'_n \\
&\longrightarrow [] \prec \sum_{i \in I'} \alpha_i \cdot \rho_i \parallel [] \prec \sum_{j \in J'} \bar{\alpha}_j \cdot \sigma_j && \text{by (rbk)}
\end{aligned}$$

and $[] \prec \sum_{i \in I'} \alpha_i \cdot \rho_i \parallel [] \prec \sum_{j \in J'} \bar{\alpha}_j \cdot \sigma_j$ is stuck since $I' \cap J' = \emptyset$.

Suppose $\rho = \bigoplus_{i \in I} \bar{\alpha}_i \cdot \rho_i$ and $\sigma = \sum_{j \in J} a_j \cdot \sigma_j$. If $I \not\subseteq J$ let $k \in I \setminus J$; then we get

$$\begin{aligned}
[] \prec \rho \parallel [] \prec \sigma &\longrightarrow [] \prec \bar{\alpha}_k \cdot \rho_k \parallel [] \prec \sigma && \text{by } (\tau) \\
&\not\rightarrow
\end{aligned}$$

Otherwise $I \subseteq J$ and $\rho_k \not\# \sigma_k$ for some $k \in I$. By reasoning as above we have

$$[] \prec \rho_k \parallel [] \prec \sigma_k \xrightarrow{*} [] \prec \rho' \parallel [] \prec \sigma' \not\rightarrow$$

and

$$\circ \prec \rho_k \parallel \sum_{j \in J \setminus \{k\}} a_j \cdot \sigma_j \prec \sigma_k \xrightarrow{*} \circ \prec \rho' \parallel \sum_{j \in J \setminus \{k\}} a_j \cdot \sigma_j \prec \sigma'$$

which imply

$$\begin{aligned}
[] \prec \rho \parallel [] \prec \sigma &\longrightarrow [] \prec \bar{\alpha}_k \cdot \rho_k \parallel [] \prec \sigma && \text{by } (\tau) \\
&\longrightarrow \circ \prec \rho_k \parallel \sum_{j \in J \setminus \{k\}} a_j \cdot \sigma_j \prec \sigma_k && \text{by (comm)} \\
&\xrightarrow{*} \circ \prec \rho' \parallel \sum_{j \in J \setminus \{k\}} a_j \cdot \sigma_j \prec \sigma' \\
&\longrightarrow [] \prec \circ \parallel [] \prec \sum_{j \in J \setminus \{k\}} a_j \cdot \sigma_j && \text{by (rbk)} \\
&\not\rightarrow
\end{aligned}$$

In both cases we conclude that $\rho \not\# \sigma$.

The proof for the case $\rho = \sum_{i \in I} a_i \cdot \rho_i$, $\sigma = \bigoplus_{j \in J} \bar{\alpha}_j \cdot \sigma_j$ is similar. \square

Theorem 3.12 (Completeness). *If $\rho \# \sigma$, then $\triangleright \rho \dashv \sigma$.*

Proof. By Theorem 3.6 each computation of **Prove**($\triangleright \rho \dashv \sigma$) always terminates. By Lemma 3.11 and Fact 3.5, $\rho \# \sigma$ implies that **Prove**($\triangleright \rho \dashv \sigma$) \neq **fail**, and hence $\triangleright \rho \dashv \sigma$. \square

4 Related work and conclusions

Since the pioneering work by Danos and Krivine [7], reversible concurrent computations have been widely studied. A main point is that understanding which actions can be reversed is not trivial in a concurrent setting, since there is no unique “last” action. Since [7], the most common notion of reversibility in concurrency is *causal-consistent* reversibility: any action can be undone if no other action depending on it has been executed (and not yet undone). The name highlights the relation with causality, which

makes the approach applicable even in settings where there is no unique notion of time, but makes it quite complex.

The first calculus for which a causal-consistent reversible extension has been defined is CCS in [7], using a stack of memories for each thread. Later, causal-consistent reversible extensions have been defined by Phillips and Ulidowski [13] for calculi definable by SOS rules in a general format (without mobility), using keys to bind synchronised actions together, and by Lanese et al. [10] for the higher-order π -calculus, using explicit memory processes to store history information and tags to track causality. A survey of causal-consistent reversibility can be found in [11].

In [9], Lanese et al. enrich the calculus of [10] with a fine-grained rollback primitive, showing the subtleties of defining a rollback operator in a concurrent setting. The first papers exploring reversibility in a context of sessions (see, e.g., [12] for a comparison between session types and contracts) are [14, 15], by Tiezzi and Yoshida. These papers define the semantics for reversible sessions by adapting the approach in [10], but do not consider compliance. Compliance has been first studied in [2]. We already discussed the differences between the present work and [2] in the Introduction.

A main point of our approach is that it exploits the fact that contracts describe sequential interactions (in a concurrent setting) to avoid the complexity of causal-consistent reversibility, allowing for a simpler semantics (compared, e.g., to the one of [9]).

Similarly to our approach, long running transactions with compensations, and in particular interacting transactions [16], allow to undo past agreements. In interacting transactions, however, a new possibility is tried when an exception is raised, not when an agreement cannot be found as in our case. Also, the possible options are sorted: first the normal execution, then the compensation. Finally, compliance of interacting transactions has never been studied. In the field of sessions, the most related works are probably the ones studying exceptions in binary sessions [5] and in multi-party sessions [4]. As for transactions, they aim at dealing with exceptions more than at avoiding to get stuck since an agreement cannot be found.

We plan to investigate whether our approach can be extended to multi-party sessions [8], the rationale being that parallelism is controlled by the global type, hence possibly part of the complexity due to concurrency can be avoided. The sub-behaviour relation induced by our notion of compliance is also worth being thoroughly studied.

Acknowledgments. We are grateful to the anonymous reviewers for their useful remarks.

References

- [1] Franco Barbanera & Ugo de'Liguoro (2010): *Two Notions of Sub-behaviour for Session-based Client/Server Systems*. In: *PPDP*, ACM Press, pp. 155–164, doi:10.1145/1836089.1836109.
- [2] Franco Barbanera, Mariangiola Dezani-Ciancaglini & Ugo de'Liguoro (2014): *Compliance for Reversible Client/Server Interactions*. In: *BEAT, EPTCS 162*, pp. 35–42, doi:10.4204/EPTCS.162.5.
- [3] Giovanni Bernardi & Matthew Hennessy (2014): *Modelling Session Types using Contracts*. *Math. Struct. in Comp. Science*, doi:10.1017/S0960129514000243. To appear.
- [4] Sara Capecchi, Elena Giachino & Nobuko Yoshida (2010): *Global Escape in Multiparty Sessions*. In: *FSTTCS, LIPIcs 8*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 338–351, doi:10.4230/LIPIcs.FSTTCS.2010.338.
- [5] Marco Carbone, Kohei Honda & Nobuko Yoshida (2008): *Structured Interactional Exceptions in Session Types*. In: *CONCUR, LNCS 5201*, Springer, pp. 402–417, doi:10.1007/978-3-540-85361-9_32.

- [6] Giuseppe Castagna, Nils Gesbert & Luca Padovani (2009): *A Theory of Contracts for Web Services*. *ACM Trans. on Prog. Lang. and Sys.* 31(5), pp. 19:1–19:61, doi:10.1145/1538917.1538920.
- [7] Vincent Danos & Jean Krivine (2004): *Reversible Communicating Systems*. In: *CONCUR, LNCS 3170*, Springer, pp. 292–307, doi:10.1007/978-3-540-28644-8_19.
- [8] Kohei Honda, Nobuko Yoshida & Marco Carbone (2008): *Multiparty Asynchronous Session Types*. In: *POPL*, ACM Press, pp. 273–284, doi:10.1145/1328897.1328472.
- [9] Ivan Lanese, Claudio Antares Mezzina, Alan Schmitt & Jean-Bernard Stefani (2011): *Controlling Reversibility in Higher-Order Pi*. In: *CONCUR, LNCS 6901*, Springer, pp. 297–311, doi:10.1007/978-3-642-23217-6_20.
- [10] Ivan Lanese, Claudio Antares Mezzina & Jean-Bernard Stefani (2010): *Reversing Higher-Order Pi*. In: *CONCUR, LNCS 6269*, Springer, pp. 478–493, doi:10.1007/978-3-642-15375-4_33.
- [11] Ivan Lanese, Claudio Antares Mezzina & Francesco Tiezzi (2014): *Causal-Consistent Reversibility*. *Bulletin of the EATCS* 114.
- [12] Cosimo Laneve & Luca Padovani (2008): *The Pairing of Contracts and Session Types*. In: *Concurrency, Graphs and Models, LNCS 5065*, pp. 681–700, doi:10.1007/978-3-540-68679-8_42.
- [13] Iain C. C. Phillips & Irek Ulidowski (2007): *Reversing Algebraic Process Calculi*. *J. of Logic and Alg. Progr.* 73(1-2), pp. 70–96, doi:10.1016/j.jlap.2006.11.002.
- [14] Francesco Tiezzi & Nobuko Yoshida (2014): *Towards Reversible Sessions*. In: *PLACES, EPTCS 155*, pp. 17–24, doi:10.4204/EPTCS.155.3.
- [15] Francesco Tiezzi & Nobuko Yoshida (2015): *Reversible session-based pi-calculus*. *J. Log. Algebr. Meth. Program.* 84(5), pp. 684–707, doi:10.1016/j.jlamp.2015.03.004.
- [16] Edsko de Vries, Vasileios Koutavas & Matthew Hennessy (2010): *Communicating Transactions - (Extended Abstract)*. In: *CONCUR, LNCS 6269*, Springer, pp. 569–583, doi:10.1007/978-3-642-15375-4_39.