



A theory of retractable and speculative contracts

Franco Barbanera^a, Ivan Lanese^{b,*}, Ugo de'Liguoro^c

^a Dipartimento di Matematica e Informatica, University of Catania, Italy

^b Dipartimento di Informatica – Scienza e Ingegneria, University of Bologna/INRIA, Italy

^c Dipartimento di Informatica, University of Torino, Italy



ARTICLE INFO

Article history:

Received 31 October 2017

Received in revised form 19 April 2018

Accepted 13 June 2018

Available online 1 August 2018

Keywords:

Behavioral contracts

Backtracking

Speculative execution

Compliance

Reversible computing

ABSTRACT

Behavioral contracts are abstract descriptions of expected communication patterns followed by either clients or servers during their interaction. Behavioral contracts come naturally equipped with a notion of *compliance*: when a client and a server follow compliant contracts, their interaction is guaranteed to progress or successfully complete. We study two extensions of behavioral contracts, *retractable contracts* dealing with *backtracking* and *speculative contracts* dealing with *speculative execution*. We show that the two extensions give rise to *the same notion of compliance*. As a consequence, they also give rise to the same *subcontract relation*, which determines when one server can be replaced by another preserving compliance. Moreover, compliance and subcontract relation are both decidable in quadratic time. Finally, we study the relationship between retractable contracts and calculi for reversible computing.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

Binary behavioral contracts [1–3] and binary session types [4] are abstractions of programs used to statically ensure that a client and a server successfully interact (see the survey in [5]). Along the years, the basic theory has been extended to deal with many features of clients and servers, such as exceptions [6], time [7], and so on. We consider here two new features: *backtracking*, allowing one to go back to previous stages of the interaction, and *speculative execution* [8], allowing one to try different alternatives concurrently. These two features have quite different origin and aims. Backtracking is used to avoid failures due to wrong past decisions in a wide range of settings, from the undo button in web browsers, to the execution model of Prolog, to techniques for rollback-recovery [9]. Speculative execution is used for efficiency reasons in different areas, from simulation [10], to thread-level optimization [11], to web services [12].

We present two extensions of binary contracts (Section 2): *retractable contracts* capturing backtracking, and *speculative contracts* capturing speculative execution. The two extensions are based on the same syntax, but naturally have different semantics. Essentially, they add to the session contracts of [13,14] (called first-order session behaviors in [13]) an operator of *external choice among output* operations. Classically, external means that the participant provides a set of alternatives, and the communication partner decides which one (s)he wants to take. This is opposed to internal choice, where the participant decides in isolation which alternative is taken. The setting in our extension is slightly more complex. The most interesting case is when an external choice among outputs and an external choice among inputs interact. In the retractable semantics, the client and the server agree on which option to explore, but they rollback and try a different possibility if

* Corresponding author.

E-mail addresses: barba@dmf.unict.it (F. Barbanera), ivan.lanese@gmail.com (I. Lanese), ugo.deliguoro@unito.it (U. de'Liguoro).

the computation gets stuck. In the speculative semantics, instead, all the possibilities are explored concurrently, and it is enough for one of them to succeed in order to guarantee the success of the whole computation.

This paper defines retractable and speculative contracts, and studies the related theory, considering the notions of *compliance* (Section 3), guaranteeing that the interaction progresses or successfully completes, *subcontract relation* (Section 4), determining when a server (resp. client) can be replaced by another server (resp. client) preserving compliance, and *dual contract* (Section 4), that is the most general contract (in terms of the subcontract relation) compliant with a given contract. Our analysis provides two main insights:

- Even if retractable contracts and speculative contracts have different semantics and give rise to different client–server interactions, the relations of compliance, subcontract and duality in the two settings do coincide. While surprising at first sight, this can be explained by noticing that in both the cases different alternatives are explored (sequentially for retractable contracts, in parallel for speculative contracts) and the success of one of them guarantees the success of the whole computation. In other terms, the two semantics provide different implementations of *angelic nondeterminism*, first described by Hoare [15].
- While retractable/speculative contracts are strictly more expressive than session contracts (indeed they are a conservative extension, see Section 3.1), their theory preserves the main good properties of the theory of session contracts. In particular, compliance and subcontract relations are both decidable (Section 3) in quadratic time (Section 5), and the dual of a contract always exists and has a simple syntactic characterization (Section 4).

A natural way to ensure the existence of the dual contract is to introduce an operator of internal choice among inputs. While this operator has limited practical impact, it makes the model more symmetric and the mathematical treatment simpler.

The results above make us confident in the fact that our semantics correctly captures the interaction patterns we are interested in. As further element supporting this, we show (Section 6) that the backtracking mechanism of retractable contracts can be seen as an application to behavioral contracts of the general theory proposed in [16] to define reversible extensions of process calculi.

This paper is an extended and revised version of [17] (a few preliminary results had been originally presented in a workshop paper [18]). Section 6, where the relation between retractable contracts and calculi for reversible computing is investigated, is completely new. Moreover, the analysis of the complexity of deciding compliance and subcontract relation has been refined, reducing the resulting complexity from a fifth power to a square. The paper also includes additional proofs and examples, and a more detailed discussion of related work. Finally, the whole presentation has been revised and improved.

Proofs omitted from the main part are collected in Appendix A.

2. Contracts for retractable and speculative interactions

We present below a uniform syntax for retractable and speculative contracts, with two semantics. It can be obtained from the syntax of (first-order) session contracts of [13,14] – dubbed SC in the present paper, and briefly recalled in Section 3.1.1 – by just adding external retractable/speculative choice among outputs and internal choice among inputs. As a matter of fact our contracts can also be seen as an extension of the retractable session contracts of [18], that we dub here **rc**, by simply adding internal choice among inputs.

Definition 1 (*Retractable/speculative contracts*). Let \mathcal{N} (*set of names*) be some countable set of symbols and let $\overline{\mathcal{N}}$ (*set of conames*) be $\{\bar{a} \mid a \in \mathcal{N}\}$, with $\mathcal{N} \cap \overline{\mathcal{N}} = \emptyset$. The set **rc** of *retractable/speculative contracts* is defined as the set of the *closed* expressions generated by the following grammar,

$\sigma, \rho :=$	1	SUCCESS
	$ \sum_{i \in I} a_i \cdot \sigma_i$	EXTERNAL INPUT CHOICE
	$ \sum_{i \in I} \bar{a}_i \cdot \sigma_i$	EXTERNAL OUTPUT CHOICE
	$ \bigoplus_{i \in I} a_i \cdot \sigma_i$	INTERNAL INPUT CHOICE
	$ \bigoplus_{i \in I} \bar{a}_i \cdot \sigma_i$	INTERNAL OUTPUT CHOICE
	$ x$	VARIABLE
	$ \text{rec } x \cdot \sigma$	RECURSION

where I is non-empty and finite, the names and the conames in choices are pairwise distinct and σ is not a variable in $\text{rec } x \cdot \sigma$. Recursion in **rc** is guarded and hence contractive in the usual sense. We take an equi-recursive view of recursion by equating $\text{rec } x \cdot \sigma$ with $\sigma[\text{rec } x \cdot \sigma / x]$.

Intuitively, a name a represents a communication channel, which can be used either in an input action (denoted a) or in an output action (denoted \bar{a}). The dot is used to denote precedence: to perform $a \cdot \sigma$ one first performs a and then continues as specified by σ . In the syntax above, branches $a_i \cdot \sigma_i$ and $\bar{a}_i \cdot \sigma_i$ can be composed either in internal choice or in

external choice. In internal choice the participant non-deterministically decides which branch he wants to take. In external input choice the participant presents a set of options and waits for his communication partner to decide which one to take. External output choice is the main novelty of this paper, and its behavior depends on whether the retractable or the speculative semantics is chosen. Hence, we will describe it in detail in the presentation of the two semantics.

We use α to range over $\mathcal{N} \cup \overline{\mathcal{N}}$, with the convention $\overline{\alpha} = \bar{a}$ if $\alpha = a$, and $\overline{\bar{a}} = a$ if $\alpha = \bar{a}$. We write $\alpha_1.\sigma_1 + \alpha_2.\sigma_2$ for binary external input/output choice and $\alpha_1.\sigma_1 \oplus \alpha_2.\sigma_2$ for binary internal input/output choice. They are both commutative by definition. Also, $\alpha.\sigma$ denotes both internal and external unary choice. This entails no ambiguity, since internal and external choices do coincide in the unary case. We also write $\alpha_k.\sigma_k + \sigma'$ for $\sum_{i \in I} \alpha_i.\sigma_i$ where $k \in I$ and $\sigma' = \sum_{i \in (I \setminus \{k\})} \alpha_i.\sigma_i$ (and similarly for internal choices). When no ambiguity can arise, we call just *contracts* the expressions in **rsC**. They are written by omitting all trailing **1**'s.

We discuss below the two interpretations and the two semantics for our contracts: the retractable one, and the speculative one.

2.1. Retractable semantics

The main novelty of the retractable semantics is that, when an external choice among outputs and an external choice among inputs interact, the client and the server agree on which option to explore *first*, but they rollback and try a different possibility if the computation gets stuck.

In order to deal with rollbacks, we decorate contracts with their *history*. A history H represents the alternatives that have been discharged in previous retractable choices (hence the branches that can be tried upon rollbacks due to synchronization failures). Histories are formalized as stacks, since we always rollback to the last retractable choice made. A current contract σ with history H is represented by $H \times \sigma$. We use the symbol ' \circ ' to stand for "no-remaining-alternatives".

Definition 2 (*Contracts with history*). Let *Histories* be the expressions generated by the grammar $H ::= () \mid H:\sigma$, where $\sigma \in \mathbf{rsC} \cup \{\circ\}$ and $\circ \notin \mathbf{rsC}$. Histories are hence stacks of contracts and \circ . Then the set of *contracts with history* is defined by:

$$\mathbf{rsCH} = \{H \times \sigma \mid H \in \text{Histories}, \sigma \in \mathbf{rsC} \cup \{\circ\}\}$$

We write just $\sigma_1:\dots:\sigma_k$ for the stack $(\dots((():\sigma_1):\dots):\sigma_k$. Moreover, we simply write σ for $():\sigma$ and $H_1:H_2$ for $(\dots(H_1:\sigma_1):\dots):\sigma_k$ where $H_2 = \sigma_1:\dots:\sigma_k$.

As standard for contracts, the definition of the retractable semantics is in two stages: we first define a labeled transition system (LTS) for contracts with history (Definition 3), and then we build on top of it a reduction semantics for pairs of contracts modeling one client and one server (Definition 4).

Definition 3 (*Semantics of contracts with history*).

$$\begin{array}{ll} (+) \ H \times \alpha.\sigma + \sigma' \xrightarrow{\alpha} H:\sigma' \times \sigma & (\oplus) \ H \times \alpha.\sigma \oplus \sigma' \xrightarrow{\tau} H \times \alpha.\sigma \\ (\alpha) \ H \times \alpha.\sigma \xrightarrow{\alpha} H:\circ \times \sigma & (\text{rb}) \ H:\sigma' \times \sigma \xrightarrow{\text{rb}} H \times \sigma' \end{array}$$

In the transition rule for external choice (+), the action α is executed, and the discharged branches in σ' are memorized on top of the history H . In internal choice (\oplus), instead, the selection of one branch is represented by a label τ , and the history H is unchanged. When a unary choice is executed (α), a ' \circ ' is added to the history, meaning that the only possible branch has been tried and no alternative is left. Rule (rb) pops the contract at the top of the stack, replacing the current one with it.

The client/server interaction is modeled by the reduction of their parallel composition, that can be either *forward*, consisting of CCS-style synchronizations 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 4 (*Semantics of retractable client/server pairs*). The following rules, plus the rule symmetric to (τ) w.r.t. \parallel , define the relation \longrightarrow over pairs of contracts with history:

$$\begin{array}{l} \text{(comm)} \\ \frac{H_1 \times \rho \xrightarrow{\alpha} H'_1 \times \rho' \quad H_2 \times \sigma \xrightarrow{\bar{\alpha}} H'_2 \times \sigma'}{H_1 \times \rho \parallel H_2 \times \sigma \longrightarrow H'_1 \times \rho' \parallel H'_2 \times \sigma'} \\ \text{(rbk)} \\ \frac{H_1 \times \rho \xrightarrow{\text{rb}} H'_1 \times \rho' \quad H_2 \times \sigma \xrightarrow{\text{rb}} H'_2 \times \sigma' \quad \rho \neq \mathbf{1}}{H_1 \times \rho \parallel H_2 \times \sigma \longrightarrow H'_1 \times \rho' \parallel H'_2 \times \sigma'} \end{array} \quad \begin{array}{l} \text{(\tau)} \\ \frac{H_1 \times \rho \xrightarrow{\tau} H_1 \times \rho'}{H_1 \times \rho \parallel H_2 \times \sigma \longrightarrow H_1 \times \rho' \parallel H_2 \times \sigma} \end{array}$$

Rule (rbk) applies only if neither (comm) nor (τ) do.

	$H_1 \times a.c + b.d$	\parallel	$H_2 \times \bar{a}.\bar{c} + \bar{b}.\bar{e}$	rollbackable interaction point
\longrightarrow	$H_1 : a.c \times d$	\parallel	$H_2 : \bar{a}.\bar{c} \times \bar{e}$	synchr. and state memorizations
\longrightarrow	$H_1 \times a.c$	\parallel	$H_2 \times \bar{a}.\bar{c}$	rollback (since no synchr. is possible)
\longrightarrow	...			etc.
	$H_1 \times a.c + b.d$	\parallel	$H_2 \times \bar{a}.\bar{c} \oplus \bar{b}.\bar{e}$	normal interaction point
\longrightarrow	$H_1 \times a.c + b.d$	\parallel	$H_2 \times \bar{b}.\bar{e}$	internal choice
\longrightarrow	$H_1 : a.c \times d$	\parallel	$H_2 : \circ \times \bar{e}$	synchr. and client-state memorization
\longrightarrow	$H_1 \times a.c$	\parallel	$H_2 \times \circ$	rollback (since no synchr. is possible)
\longrightarrow	...			rollback continues to an older past, if any (since no synchr. is possible for \circ)

Fig. 1. Rollback.

The *forward reduction* \longrightarrow_f is the relation generated by rules (τ) and (comm) . The *backward reduction* \longrightarrow_b is the relation generated by rule (rbk) .

Remark 1. The relation \longrightarrow is defined by means of a system of rules with priorities. Reduction relations with priority rules have been initially introduced in the context of term-rewriting [19] and of languages with pattern-matching [20]. Orderings on SOS rules were later proposed by Phillips and Ulidowski as an alternative to negative premises [21,22]. In fact alternative definitions would either consist in resorting to a “negated” reduction to specify that rollback can be applied when the forward semantics is stuck; or in introducing error states reached when a client and a server exhibit incompatible choices.

We opted for the present definition in order to get a simpler treatment of the reduction relation and hence simpler and more readable proofs.

Example 1. In order to get a better insight into the role of \circ in the rollback mechanism, observe that, for a client like $a.c + b.d$, rule $(+)$ in Definition 3 forces the memorization of the discharged branch, say $a.c$, independently from the shape of the server. If the server is, for instance, $\bar{a}.\bar{c} + \bar{b}.\bar{e}$ (see the first reduction sequence in Fig. 1) then the server memorizes a discarded branch too, $\bar{a}.\bar{c}$ in the example. Hence, after the synchronization failure of d and \bar{e} , the two discarded branches are recovered and executed. If the server is, instead, $\bar{a}.\bar{c} \oplus \bar{b}.\bar{e}$ (see the second reduction sequence in Fig. 1) then the server is not willing to rollback this choice, and it stores a \circ in the history. Upon synchronization failure, the discarded branch of the client and the \circ of the server are recovered, but, since \circ cannot reduce, this immediately causes a new synchronization failure and an *older* past (if any) is recovered. Summarizing, recovered branches are actually executed only if both the client and the server are willing to do so.

Remark 2. The semantics defined above for retractable client/server pairs can be seen as an instantiation on contracts of the standard reversible semantics for process calculi, see, e.g., [23,16,24,25]. In particular, the semantics would become a classic uncontrolled semantics (according to the terminology in [25]) by removing the four control mechanisms below:

1. the fact that only external choices are retractable;
2. the side condition $\rho \neq \mathbf{1}$ in rule (rbk) , which disallows backtrack after success;
3. the fact that rule (rbk) can be applied only if no other rule applies, ensuring that backtrack is enabled only when no forward reduction is possible;
4. the fact that in external choices the selected path is not stored in the history, so that each path can be tried at most once.

These mechanisms provide a semantic control of reversibility [25], specifying which rollback steps are allowed, and when. We formalize this intuition in Section 6. We discuss in Remark 3 the impact that removing the above control mechanisms would have on retractable contracts and on their theory.

Example 2. Retractable contracts allow one to first try a preferred alternative, but to accept also another alternative if the first one proves to be impossible to obtain. In cloud computing settings, companies may hire virtual machines and storing facilities from cloud providers with some agreed Quality of Service (QoS). A company is willing to hire at some medium or low price a certain amount of machines for online elaboration during the daytime, but, if the price is too high, it is also willing to switch to offline night elaboration. In this last case it is only willing to pay a low price.

A retractable contract with this behavior may be written as:

$$\text{cloudClient} = \overline{\text{QoSday}}.\overline{\text{priceMed.ok}} + \overline{\text{priceLow.ok}} + \overline{\text{QoSnight}}.\overline{\text{priceLow.ok}}$$

$$\begin{array}{l}
\overline{QoSday.(priceMed.ok} \\
\langle \rangle \times \frac{\overline{priceLow.ok} + \overline{QoSnight.priceLow.ok}}{\overline{priceLow.ok}} \parallel \langle \rangle \times \sum_{QoS} \overline{QoS.price_{QoS}.ok} \quad (1) \\
\rightarrow \langle \rangle : \overline{QoSnight.priceLow.ok} \\
\times \overline{priceMed.ok} + \overline{priceLow.ok} \parallel \langle \rangle : \sum_{QoS \neq QoSday} \overline{QoS.price_{QoS}.ok} \\
\times \overline{priceHigh.ok} \quad (2) \\
\rightarrow \langle \rangle \times \overline{QoSnight.priceLow.ok} \parallel \langle \rangle \times \sum_{QoS \neq QoSday} \overline{QoS.price_{QoS}.ok} \quad (3) \\
\rightarrow \langle \rangle : \circ \times \overline{priceLow.ok} \parallel \langle \rangle : \sum_{QoS \neq QoSday, QoSnight} \overline{QoS.price_{QoS}.ok} \\
\times \overline{priceLow.ok} \quad (4) \\
\rightarrow \langle \rangle : \circ : \circ \times \overline{ok} \parallel \langle \rangle : \sum_{QoS \neq QoSday, QoSnight} \overline{QoS.price_{QoS}.ok} : \circ : \circ \times \overline{ok} \quad (5) \\
\rightarrow \langle \rangle : \circ : \circ : \circ \times \mathbf{1} \parallel \langle \rangle : \sum_{QoS \neq QoSday, QoSnight} \overline{QoS.price_{QoS}.ok} : \circ : \circ \times \mathbf{1} \quad (6)
\end{array}$$

Fig. 2. An example of retractable interaction.

Notice that the contract does not specify which alternative the client prefers: this aspect of the client behavior is abstracted away.¹ A sample server is:

$$\text{cloudServer} = \sum_{QoS \in \{QoSday, QoSnight, \dots\}} \overline{QoS.price_{QoS}.ok}$$

A sample interaction is described in Fig. 2, where we assume that

$$\text{price}_{QoSday} = \text{priceHigh} \quad \text{and} \quad \text{price}_{QoSnight} = \text{priceLow}$$

In particular, the interaction starts from a rollbackable interaction point (1) and proceeds as follows: (2) synchronization on $QoSday$ and memorization of the discarded $QoSnight$ branch; (3) rollback, since no choice has been offered for the requested $\overline{priceHigh}$; (4) synchronization on $QoSnight$ and memorization of \circ (no further alternative branch is available); (5) synchronization on $\overline{priceLow}$; (6) synchronization on \overline{ok} and success state reached (client is satisfied).

2.2. Speculative semantics

The main idea of the speculative semantics is that in an external output choice all the options are tried concurrently: if at least one of them succeeds, then the whole computation succeeds. In order to represent concurrent trials we need runtime contracts featuring multiple *threads*. A thread is a contract preceded by zero or more prefixes of the form $\alpha @$, meaning that the action α has been performed in the past.

Definition 5 (*Contracts with threads*). *Contracts with threads* \mathbf{C} , used as runtime syntax for contracts, are parallel compositions of *threads* \mathbf{T} . Each thread is a contract prefixed by zero or more actions.

$$\mathbf{C} ::= \mathbf{T} \mid (\mathbf{C} \mid \mathbf{T}) \mid (\mathbf{T} \mid \mathbf{C}) \quad \mathbf{T} ::= \sigma \mid \alpha @ \mathbf{T}$$

We assume the operator ‘ \mid ’ to be associative and commutative.

We will always work with contracts with threads such that threads are uniquely identified by the sequence of actions prefixing them. Our prefixes may remind location prefixes from [27], however our prefixes are deterministically generated from performed actions, and they do not represent locations.

As for the retractable semantics, the definition of the speculative semantics is in two stages: we first define an LTS for contracts with threads (Definition 6), and then we use it to define a reduction semantics for pairs of contracts with threads representing one client and one server (Definition 7).

Definition 6 (*Semantics of contracts with threads*). The following rules define the LTS for contracts with threads. In the LTS, we use as labels: $\alpha ::= a \mid \bar{a}$ for actions, $\beta ::= \alpha \mid \alpha\beta$ for sequences of actions, and $\beta_\tau ::= \tau \mid \beta \mid \beta$, \mathbf{T} for general labels.

¹ As we will discuss in Section 7, it is possible to extend our contracts to express an order of preference among external output choices, as done in [26] in the context of session types.

$$\begin{array}{c}
\text{(Fork)} \quad \alpha.\sigma + \sigma' \xrightarrow{\alpha, \sigma'} \alpha @ \sigma \quad \text{(\oplus)} \quad \alpha.\sigma \oplus \sigma' \xrightarrow{\tau} \alpha.\sigma \quad \text{(\alpha)} \quad \alpha.\sigma \xrightarrow{\alpha} \alpha @ \sigma \\
\text{(@-\alpha)} \quad \frac{\mathbf{T} \xrightarrow{\beta} \mathbf{T}'}{\alpha @ \mathbf{T} \xrightarrow{\alpha\beta} \alpha @ \mathbf{T}'} \quad \text{(@-\alpha-T)} \quad \frac{\mathbf{T} \xrightarrow{\beta, \mathbf{T}'} \mathbf{T}'}{\alpha @ \mathbf{T} \xrightarrow{\alpha\beta, \alpha @ \mathbf{T}'} \alpha @ \mathbf{T}'} \quad \text{(@-\tau)} \quad \frac{\mathbf{T} \xrightarrow{\tau} \mathbf{T}'}{\alpha @ \mathbf{T} \xrightarrow{\tau} \alpha @ \mathbf{T}'} \quad \text{(ParL)} \quad \frac{\mathbf{T} \xrightarrow{\beta\tau} \mathbf{T}'}{\mathbf{T} | \mathbf{C} \xrightarrow{\beta\tau} \mathbf{T}' | \mathbf{C}}
\end{array}$$

In the rule for external choice (*Fork*), when an action α is executed, its continuation σ is prefixed by it. The other branches σ' need to be executed in a freshly spawned thread. Since such thread needs to be installed at top level, σ' is added to the label, and the actual installation is performed at the level of speculative client/server pairs (see rule (*comm*) in Definition 7). The rule for internal choice (\oplus) simply selects one of the available options. A unary choice (α) executes the action α and prefixes with it the continuation σ . Rules ($@-\alpha$), ($@-\alpha-T$), and ($@-\tau$) enable execution below an $@$ prefix. In particular, in rule ($@-\alpha$), the prefix itself is added to the label β . Prefixes uniquely identify threads, and ensure that each thread interacts only with the one with dual prefix which is running on the communication partner. This is specified in Definition 7 below. Rule ($@-\alpha-T$) is analogous to rule ($@-\alpha$), but the label also contains a thread \mathbf{T}' , and the prefix α is added to both β and \mathbf{T}' . No prefix is added to τ actions, propagated by rule ($@-\tau$). Rule (ParL) simply allows components of a parallel composition to execute (a symmetric rule is not needed thanks to the commutativity of $|$).

The interaction of a client with a server is modeled by the reduction of their parallel composition.

Definition 7 (*Semantics of speculative client/server pairs*). The following rules, plus the rule symmetric to (τ) w.r.t. \parallel , define the relation \longrightarrow over pairs of contracts with threads. In the LTS below, $?T$ denotes either the thread \mathbf{T} or nothing. Hence, $\beta, ?T$ and $\mathbf{C} | ?T$ are respectively β and \mathbf{C} if $?T$ is nothing, and β, \mathbf{T} and $\mathbf{C} | \mathbf{T}$ otherwise. Also, the duality operator extends from actions to sequences: $\overline{\alpha\beta} = \overline{\alpha}\overline{\beta}$.

$$\begin{array}{c}
\text{(comm)} \quad \frac{\mathbf{C} \xrightarrow{\beta, ?T} \mathbf{C}' \quad \overline{\mathbf{C}''} \xrightarrow{\overline{\beta}, ?T''} \overline{\mathbf{C}'''}}{\mathbf{C} \parallel \overline{\mathbf{C}''} \longrightarrow \mathbf{C}' | ?T \parallel \overline{\mathbf{C}'''} | ?T'''} \quad \text{(\tau)} \quad \frac{\mathbf{C} \xrightarrow{\tau} \mathbf{C}'}{\mathbf{C} \parallel \mathbf{C}'' \longrightarrow \mathbf{C}' \parallel \mathbf{C}''}
\end{array}$$

Rule (*comm*) allows threads performing dual sequences of actions to interact. This implies that both the performed actions and the prefixes of the threads executing them have to be dual. Threads in the labels, if present, are installed in parallel. Rule (τ) simply propagates the τ action.

Example 3. A server provides access to multiple algorithms for SAT solving [28]. A client first sends the problem instance to be solved, then selects the algorithm, and finally sends the relevant parameters. The server computes the solution according to the received commands, and sends it back. Since the most efficient technique depends on the problem instance [29], the server supports speculative execution to allow one to try different algorithms at the same time (this is called the portfolio approach). The server contract is described by:

$$\text{SATserver} = \text{inst}.\overline{\sum_i \text{alg}_i}.\overline{\sum_j \text{par}_j}.\overline{\text{sol}}$$

A simple client that tries both the DPLL approach and the walksat approach can be modeled as follows:

$$\text{SATclient} = \overline{\text{inst}}.(\overline{\text{DPLL}}.\overline{\text{par}}.\overline{\text{sol}} + \overline{\text{walksat}}.\overline{\text{par}}.\overline{\text{sol}})$$

A sample computation proceeds as described in Fig. 3, assuming that the server supports both DPLL and walksat. To keep the example simple we drop the choice of parameters. Let us see in more details how the creation of threads is managed. The first reduction in Fig. 3 is due to rule (*comm*), since

$$\overline{\text{inst}}.(\overline{\text{DPLL}}.\overline{\text{sol}} + \overline{\text{walksat}}.\overline{\text{sol}}) \xrightarrow{\overline{\text{inst}}} \overline{\text{inst}} @ (\overline{\text{DPLL}}.\overline{\text{sol}} + \overline{\text{walksat}}.\overline{\text{sol}})$$

and

$$\text{inst}.\overline{\sum_i \text{alg}_i}.\overline{\text{sol}} \xrightarrow{\text{inst}} \text{inst} @ \overline{\sum_i \text{alg}_i}.\overline{\text{sol}}.$$

The second reduction is also due to rule (*comm*), since, on the client side

$$\begin{array}{c}
\overline{\text{inst}} @ (\overline{\text{DPLL}}.\overline{\text{sol}} + \overline{\text{walksat}}.\overline{\text{sol}}) \xrightarrow{\overline{\text{inst}} @ \overline{\text{DPLL}}, \overline{\text{inst}} @ \overline{\text{walksat}}.\overline{\text{sol}}} \overline{\text{inst}} @ \overline{\text{DPLL}} @ \overline{\text{sol}} \quad \text{(Fork)} \\
\overline{\text{inst}} @ (\overline{\text{DPLL}}.\overline{\text{sol}} + \overline{\text{walksat}}.\overline{\text{sol}}) \xrightarrow{\overline{\text{inst}} @ \overline{\text{DPLL}}, \overline{\text{inst}} @ \overline{\text{walksat}}.\overline{\text{sol}}} \overline{\text{inst}} @ \overline{\text{DPLL}} @ \overline{\text{sol}} \quad \text{(@-\alpha-T)}
\end{array}$$

$$\begin{array}{l}
\overline{\text{inst.}(\overline{\text{DPLL}}.\text{sol} + \overline{\text{walksat}}.\text{sol})} \parallel \text{inst.} \sum_i \text{alg}_i.\overline{\text{sol}} \\
\rightarrow \overline{\text{inst}@\overline{\text{DPLL}}.\text{sol} + \overline{\text{walksat}}.\text{sol}} \parallel \text{inst}@\sum_i \text{alg}_i.\overline{\text{sol}} \\
\rightarrow \begin{array}{l} \overline{\text{inst}@\overline{\text{DPLL}}@\text{sol}} \\ | \overline{\text{inst}@\overline{\text{walksat}}.\text{sol}} \end{array} \parallel \begin{array}{l} \text{inst}@\overline{\text{DPLL}}@\overline{\text{sol}} \\ | \text{inst}@\sum_{\{i|A_i \neq \overline{\text{DPLL}}\}} \text{alg}_i.\overline{\text{sol}} \end{array} \\
\rightarrow \begin{array}{l} \overline{\text{inst}@\overline{\text{DPLL}}@\text{sol}} \\ | \overline{\text{inst}@\overline{\text{walksat}}@\text{sol}} \end{array} \parallel \begin{array}{l} \text{inst}@\overline{\text{DPLL}}@\overline{\text{sol}} \\ | \text{inst}@\overline{\text{walksat}}@\overline{\text{sol}} \\ | \text{inst}@\sum_{\{i|A_i \neq \overline{\text{DPLL}}, \overline{\text{walksat}}\}} \text{alg}_i.\overline{\text{sol}} \end{array} \\
\rightarrow \begin{array}{l} \overline{\text{inst}@\overline{\text{DPLL}}@\text{sol}} \\ | \overline{\text{inst}@\overline{\text{walksat}}@\text{sol}} @ \mathbf{1} \end{array} \parallel \begin{array}{l} \text{inst}@\overline{\text{DPLL}}@\overline{\text{sol}} \\ | \text{inst}@\overline{\text{walksat}}@\overline{\text{sol}} @ \mathbf{1} \\ | \text{inst}@\sum_{\{i|A_i \neq \overline{\text{DPLL}}, \overline{\text{walksat}}\}} \text{alg}_i.\overline{\text{sol}} \end{array}
\end{array}$$

Fig. 3. An example of speculative interaction.

whereas, on the server side,

$$\begin{array}{l}
\overline{\sum_i \text{alg}_i.\overline{\text{sol}}} \xrightarrow{\text{DPLL}, \sum_{\{i|A_i \neq \overline{\text{DPLL}}\}} \text{alg}_i.\overline{\text{sol}}} \overline{\text{DPLL}@\overline{\text{sol}}} \text{ (Fork)} \\
\overline{\text{inst}@\sum_i \text{alg}_i.\overline{\text{sol}}} \xrightarrow{\text{inst} \overline{\text{DPLL}}, \text{inst}@\sum_{\{i|A_i \neq \overline{\text{DPLL}}\}} \text{alg}_i.\overline{\text{sol}}} \overline{\text{inst}@\overline{\text{DPLL}}@\overline{\text{sol}}} \text{ (@-}\alpha\text{-T)}
\end{array}$$

3. Compliance

The compliance relation for session contracts [13,14] consists in requiring that, whenever no reduction is possible, all client's requests and offers have been satisfied, i.e. the client is in the success state $\mathbf{1}$. For retractable contracts, thanks to the retractable operational semantics taking care of forward and backward reductions, we can adopt the same definition. We use $\xrightarrow{*}$ to denote the reflexive and transitive closure of \rightarrow , and $\not\rightarrow$ to specify that no \rightarrow reduction exists.

Definition 8 (Retractable compliance relation $\dashv\!\!\dashv^{\text{R}}$).

- i) The relation $\dashv\!\!\dashv^{\text{R}}$ on contracts with history is defined by:
 $H_1 \times \rho \dashv\!\!\dashv^{\text{R}} H_2 \times \sigma$ if, for any $H'_1, H'_2, \rho', \sigma'$ such that
 $H_1 \times \rho \parallel H_2 \times \sigma \xrightarrow{*} H'_1 \times \rho' \parallel H'_2 \times \sigma' \not\rightarrow$, we have $\rho' = \mathbf{1}$
- ii) The relation $\dashv\!\!\dashv^{\text{R}}$ on contracts is defined by: $\rho \dashv\!\!\dashv^{\text{R}} \sigma$ if $\langle \rangle \times \rho \dashv\!\!\dashv^{\text{R}} \langle \rangle \times \sigma$.

For speculative contracts we need to take into account the fact that the whole computation succeeds if at least one of its branches succeeds.

Definition 9 (Speculative compliance relation $\dashv\!\!\dashv^{\text{S}}$). The relation $\dashv\!\!\dashv^{\text{S}}$ on contracts is defined by:

$$\rho \dashv\!\!\dashv^{\text{S}} \sigma \text{ if for any } \mathbf{C}_\rho, \mathbf{C}_\sigma \text{ such that } \rho \parallel \sigma \xrightarrow{*} \mathbf{C}_\rho \parallel \mathbf{C}_\sigma \not\rightarrow \\
\text{there exist } \mathbf{C}, n, \alpha_1, \dots, \alpha_n \text{ such that } \mathbf{C}_\rho = \mathbf{C} | \alpha_1 @ \dots @ \alpha_n @ \mathbf{1}$$

We now provide a formal system characterizing compliance on both retractable and speculative contracts.

Definition 10 (Formal system for compliance \triangleright). Judgments in the formal system \triangleright are expressions of the form $\Gamma \triangleright \rho \curlywedge \sigma$, where the environment Γ is a finite set of expressions of the form $\delta \curlywedge \gamma$, with $\rho, \sigma, \delta, \gamma \in \mathbf{rsc}$. Axioms and rules are defined in Fig. 4.

Intuitively, the symbol \curlywedge is used as syntactical counterpart for both the relations $\dashv\!\!\dashv^{\text{R}}$ and $\dashv\!\!\dashv^{\text{S}}$ that, as we prove below, do coincide.

The only non standard rule of system \triangleright is $(+ \cdot +)$, which ensures compliance of two external choices when they contain respectively (at least) *one* α and the corresponding $\bar{\alpha}$, followed by compliant contracts. This contrasts with the rules $(\oplus \cdot +)$ and $(+ \cdot \oplus)$, where *each* α in an internal choice must have a corresponding $\bar{\alpha}$ in the external choice, followed by compliant contracts. No rule is provided for the case $(\oplus \cdot \oplus)$ since two internal choices are compliant only if both of them are unary choices (otherwise they may always get stuck by choosing incompatible actions). Since unary internal choice coincides with unary external choice, this case is taken into account by the rules we already have. Notice that rule

$$\begin{array}{c}
\text{(Ax)} \quad \Gamma \triangleright \mathbf{1} \dashv \sigma \qquad \text{(Hyp)} \quad \Gamma, \rho \dashv \sigma \triangleright \rho \dashv \sigma \qquad \text{(+}\cdot\text{+)} \quad \frac{\Gamma, \alpha.\rho + \rho' \dashv \bar{\alpha}.\sigma + \sigma' \triangleright \rho \dashv \sigma}{\Gamma \triangleright \alpha.\rho + \rho' \dashv \bar{\alpha}.\sigma + \sigma'} \\
\text{(\oplus}\cdot\text{+)} \quad \frac{\forall h \in I. \Gamma, \bigoplus_{i \in I} \bar{\alpha}_i.\rho_i \dashv \sum_{j \in I \cup J} \alpha_j.\sigma_j \triangleright \rho_h \dashv \sigma_h}{\Gamma \triangleright \bigoplus_{i \in I} \bar{\alpha}_i.\rho_i \dashv \sum_{j \in I \cup J} \alpha_j.\sigma_j} \\
\text{(+}\cdot\text{\oplus)} \quad \frac{\forall h \in I. \Gamma, \sum_{j \in I \cup J} \bar{\alpha}_j.\rho_j \dashv \bigoplus_{i \in I} \alpha_i.\sigma_i \triangleright \rho_h \dashv \sigma_h}{\Gamma \triangleright \sum_{j \in I \cup J} \bar{\alpha}_j.\rho_j \dashv \bigoplus_{i \in I} \alpha_i.\sigma_i}
\end{array}$$

Fig. 4. System \triangleright .

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

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

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

else fail

Fig. 5. The procedure **Prove**.

(+·+) implicitly represents the fact that, in the decision procedure for two contracts made of external choices, the possible synchronizing branches have to be tried, until either a successful one is found or all fail. Looking at a derivation bottom-up, at each application of a rule the considered pair of contracts is added to the environment Γ . In this way, if the same pair is reached again due to the equi-recursive view of contracts, the derivation can be closed using rule (Hyp). Rule (Ax) instead closes the derivation when the client reaches the success state $\mathbf{1}$. We write $\triangleright \rho \dashv \sigma$ instead of $\Gamma \triangleright \rho \dashv \sigma$ when Γ is empty.

Derivability in system \triangleright is decidable, since it is syntax-directed and proof reconstruction does terminate.

The procedure **Prove** in Fig. 5 clearly implements the formal system and can be used to prove decidability of \triangleright (see Appendix A.1).

Theorem 1. *Derivability in the formal system \triangleright is decidable.*

We can now prove (see Appendix A.1) the soundness and the completeness of the formal system \triangleright w.r.t. both the retractable and the speculative semantics.

Theorem 2 (Retractable soundness and completeness).

$$\triangleright \rho \dashv \sigma \text{ iff } \rho \dashv \|\text{R} \sigma$$

$$\begin{array}{c}
\frac{\overline{\gamma_1, \gamma_2, \gamma_3 \triangleright \mathbf{1} \wp \mathbf{1}}^{(Ax)}}{\overline{\gamma_1, \gamma_2 \triangleright \text{sol} \wp \overline{\text{sol}}}^{(+ \cdot \oplus)}} \\
\frac{\overline{\gamma_1 \triangleright \overline{\text{DPLL.sol} + \overline{\text{walksat.sol}}} \wp \sum_i \overline{\text{alg}_i.\overline{\text{sol}}}^{(+ \cdot +)}}{\overline{\triangleright \text{inst.}(\overline{\text{DPLL.sol} + \overline{\text{walksat.sol}}) \wp \text{inst.} \sum_i \overline{\text{alg}_i.\overline{\text{sol}}}^{(\oplus \cdot +)}}}
\end{array}$$

where $\gamma_1 = \overline{\text{inst.}(\overline{\text{DPLL.sol} + \overline{\text{walksat.sol}}) \wp \text{inst.} \sum_i \overline{\text{alg}_i.\overline{\text{sol}}}$
 $\gamma_2 = \overline{\text{DPLL.sol} + \overline{\text{walksat.sol}}} \wp \sum_i \overline{\text{alg}_i.\overline{\text{sol}}}$
 $\gamma_3 = \text{sol} \wp \overline{\text{sol}}$
and where, for some i , $\text{alg}_i = \text{walksat}$.

Fig. 6. A sample derivation in \triangleright .

Theorem 3 (Speculative soundness and completeness).

$$\triangleright \rho \wp \sigma \text{ iff } \rho \dashv\!\!\dashv^S \sigma$$

The client and the server of Example 3 can be now shown to be compliant, i.e. $\overline{\text{inst.}(\overline{\text{DPLL.sol} + \overline{\text{walksat.sol}}) \dashv\!\!\dashv^S \text{inst.} \sum_i \overline{\text{alg}_i.\overline{\text{sol}}}$, by providing a derivation for $\triangleright \overline{\text{inst.}(\overline{\text{DPLL.sol} + \overline{\text{walksat.sol}}) \wp \text{inst.} \sum_i \overline{\text{alg}_i.\overline{\text{sol}}}$, as shown in Fig. 6.

By the soundness and completeness of system \triangleright w.r.t. both the relations of retractable and speculative compliance, we immediately get that the two compliance relations do coincide.

Corollary 4 (Retractable and speculative compliances coincide).

$$\dashv\!\!\dashv^R = \dashv\!\!\dashv^S$$

By the above, from now on we write $\dashv\!\!\dashv$ instead of $\dashv\!\!\dashv^R$ or $\dashv\!\!\dashv^S$. So the following also easily follows.

Corollary 5 (Compliance decidability). *The relation $\dashv\!\!\dashv$ is decidable.*

Remark 3. We now discuss the impact on the compliance relation of the four mechanisms for controlling reversibility in the retractable semantics (see Remark 2). In particular, we analyze what would happen by dropping each one of them in isolation:

Dropping “Only external choices are retractable”:

each reduction could be undone. From the compliance point of view, all the choices would be retractable. Hence, retractable contracts would not be a conservative extension (see Section 3.1) of session contracts any more. The case we consider is strictly more general, since we allow for both retractable and unretractable choices, being our internal choices unretractable.

Dropping the side condition $\rho \neq \mathbf{1}$ in rule (rbk) of Definition 4:

any forward finite interaction would be followed by a rollback. In particular, most of the client/server pairs without recursion (except a few trivial ones, like $\langle \rangle \times \mathbf{1} \parallel \langle \rangle \times \sigma$) would end into $\langle \rangle \times \circ \parallel \langle \rangle \times \circ$. Thus all these pairs of contracts would not be compliant.

Dropping “rule (rbk) can be applied only if no other rule applies”:

interactions could rollback before succeeding. As in the case above, most client/server pairs (except a few trivial ones, but including recursive ones) could reduce to $\langle \rangle \times \circ \parallel \langle \rangle \times \circ$. Again all these pairs of contracts would not be compliant.

Dropping “in choices the selected path is not stored in the history”:

any client/server pair that would not normally succeed with at least one retractable choice could diverge by undoing and redoing the choice forever, thus trivially ensuring compliance.

None of the last three scenarios provides a reasonable setting. The first one would be reasonable, but too restrictive, since it would prevent to model all those irrevocable choices which are present in everyday interactions.

3.1. Conservativity

It is possible to show that all the relations on our retractable and speculative contracts (**rsC**) are conservative extensions of corresponding notions on (first-order) session contracts (SC) as defined in [13,14], and on the retractable session contracts (**rC**) as defined in [18].

$$\begin{array}{c}
\text{(Ax)} \quad \Gamma \triangleright_{\text{SC}} \mathbf{1} \dot{\vee} \sigma \qquad \text{(Hyp)} \quad \Gamma, \rho \dot{\vee} \sigma \triangleright_{\text{SC}} \rho \dot{\vee} \sigma \\
\text{(\oplus} \cdot \text{+)} \quad \frac{\forall h \in I. \Gamma, \bigoplus_{i \in I} \bar{a}_i \cdot \rho_i \dot{\vee} \sum_{j \in I \cup J} a_j \cdot \sigma_j \triangleright_{\text{SC}} \rho_h \dot{\vee} \sigma_h}{\Gamma \triangleright_{\text{SC}} \bigoplus_{i \in I} \bar{a}_i \cdot \rho_i \dot{\vee} \sum_{j \in I \cup J} a_j \cdot \sigma_j} \\
\text{(+} \cdot \text{\oplus)} \quad \frac{\forall h \in I. \Gamma, \sum_{j \in I \cup J} a_j \cdot \rho_j \dot{\vee} \bigoplus_{i \in I} \bar{a}_i \cdot \sigma_i \triangleright_{\text{SC}} \rho_h \dot{\vee} \sigma_h}{\Gamma \triangleright_{\text{SC}} \sum_{j \in I \cup J} a_j \cdot \rho_j \dot{\vee} \bigoplus_{i \in I} \bar{a}_i \cdot \sigma_i}
\end{array}$$

Fig. 7. System $\triangleright_{\text{SC}}$.

For the sake of completeness and readability, in the following subsection we briefly recall the formalism of session contracts together with some of their properties, while referring to [18] for retractable session contracts.

3.1.1. Session contracts

The set SC of session contracts (roughly interpreting session types [4] into contracts [1–3]) can be seen as the subset of elements in \mathbf{rSC} not containing external output choices and internal input choices, with the following operational semantics.

Definition 11 (*Semantics of session contracts*).

$$\bar{a} \cdot \sigma \oplus \sigma' \xrightarrow{\tau}_{\text{SC}} \bar{a} \cdot \sigma \qquad \alpha \cdot \sigma \xrightarrow{\alpha}_{\text{SC}} \sigma \qquad a \cdot \sigma + \sigma \xrightarrow{a}_{\text{SC}} \sigma$$

As done for \mathbf{rSC} (see Section 2), we write $\alpha_k \cdot \sigma_k + \sigma'$ for $\sum_{i \in I} \alpha_i \cdot \sigma_i$ where $k \in I$ and $\sigma' = \sum_{i \in (I \setminus \{k\})} \alpha_i \cdot \sigma_i$ (and similarly for internal choices). Moreover, we can look at session contracts up-to unfolding of recursion.

The next definitions introduce the LTS for client/server pairs of session contracts, and the corresponding compliance relation.

Definition 12 (*Semantics of client/server pairs of session contracts*).

$$\frac{\rho \xrightarrow{\tau}_{\text{SC}} \rho'}{\rho \parallel \sigma \rightarrow_{\text{SC}} \rho' \parallel \sigma} \qquad \frac{\sigma \xrightarrow{\tau}_{\text{SC}} \sigma'}{\rho \parallel \sigma \rightarrow_{\text{SC}} \rho \parallel \sigma'} \qquad \frac{\rho \xrightarrow{\alpha}_{\text{SC}} \rho' \quad \sigma \xrightarrow{\bar{\alpha}}_{\text{SC}} \sigma'}{\rho \parallel \sigma \rightarrow_{\text{SC}} \rho' \parallel \sigma'}$$

Definition 13 (*Compliance relation $\dashv\!\!\dashv_{\text{SC}}$ for session contracts*). The relation $\dashv\!\!\dashv_{\text{SC}} \subset \text{SC} \times \text{SC}$ is defined by:

$$\rho \dashv\!\!\dashv_{\text{SC}} \sigma \text{ if, for each } \rho', \sigma' \text{ such that} \\ \rho \parallel \sigma \xrightarrow{*}_{\text{SC}} \rho' \parallel \sigma' \dashv\!\!\dashv_{\text{SC}} \text{ we have } \rho' = \mathbf{1}.$$

The sound and complete formal system $\triangleright_{\text{SC}}$ for $\dashv\!\!\dashv_{\text{SC}}$ is recalled in Fig. 7.

Theorem 6. Let $\rho, \sigma \in \text{SC}$. $\rho \dashv\!\!\dashv_{\text{SC}} \sigma$ iff $\triangleright_{\text{SC}} \rho \dot{\vee} \sigma$.

Proof sketch. In [13] a formal system $\triangleright_{\text{H}}$ is devised which is sound and complete for compliance of higher-order session contracts (HSC), that is $\triangleright_{\text{H}} \rho \dot{\vee} \sigma$ iff $\rho \dashv\!\!\dashv_{\text{H}} \sigma$. The set SC can be looked at as the first-order restriction of HSC. It is easy to show that, for $\rho, \sigma \in \text{SC}$, $\triangleright_{\text{H}} \rho \dot{\vee} \sigma$ iff $\triangleright_{\text{SC}} \rho \dot{\vee} \sigma$. It is also not difficult to show that, for $\rho, \sigma \in \text{SC}$, $\rho \dashv\!\!\dashv_{\text{SC}} \sigma$ iff $\rho \dashv\!\!\dashv_{\text{H}} \sigma$.

From the above statement, the thesis descends immediately. \square

3.1.2. Conservativity results

As previously said, it is not difficult to check that session contracts SC are a strict subset of retractable session contracts \mathbf{rC} , which, in turn, are a strict subset of the contracts \mathbf{rSC} we are presently investigating, namely: $\text{SC} \subsetneq \mathbf{rC} \subsetneq \mathbf{rSC}$. Obviously the strict inclusion $\text{SC} \subsetneq \mathbf{rSC}$ is not enough, by itself, to guarantee the retractable and speculative operational semantics for \mathbf{rSC} to be conservative extensions of the operational semantics of SC. We show that it is so in the following Proposition 7 (see Appendix A.1.1 for the proof). Informally, it states that both the forward retractable semantics \longrightarrow_f and the speculative semantics \longrightarrow of pairs of contracts in SC are annotated versions of their semantics as session contracts.

Proposition 7 (*Operational semantics conservativity*). Let $\rho, \sigma \in \text{SC}$.

- i) $\rho \parallel \sigma \xrightarrow{*}_{\text{SC}} \rho' \parallel \sigma'$ iff for all H_1, H_2 there exist H'_1, H'_2 such that $H_1 \times \rho \parallel H_2 \times \sigma \xrightarrow{*}_f H'_1 \times \rho' \parallel H'_2 \times \sigma'$

- ii) $\rho \parallel \sigma \xrightarrow{*}_{\text{SC}} \rho' \parallel \sigma'$ iff there exist $n, \alpha_1, \dots, \alpha_n, \mathbf{C}_\rho$ and \mathbf{C}_σ such that
- $$\rho \parallel \sigma \xrightarrow{*} \alpha_1 @ \dots @ \alpha_n @ \rho' \mid \mathbf{C}_\rho \parallel \overline{\alpha_1} @ \dots @ \overline{\alpha_n} @ \sigma' \mid \mathbf{C}_\sigma$$

where $\xrightarrow{*}_{\text{SC}}$ denotes the reduction relation on SC client/server pairs in the theory of session contracts.

We do not take into account conservativity of the retractable operational semantics for **rsC** over the one for **rC** because it is quite trivial, since the rules in the two semantics are essentially the same. A conservativity result of the speculative operational semantics for **rsC** over the one for **rC** would instead consist in a rather cumbersome and uninteresting statement.

The conservativity result for the operational semantics is not enough, in itself, to guarantee the theory of retractable compliance for **rsC** to be a conservative extension of both the theory of compliance for **rC** and for SC. Also in this case, however, we can prove it to be so, that is compliance for session contracts is the restriction of our relation $\dashv\!\!\dashv$ to elements in SC. Similarly, compliance for retractable session contracts is the restriction of $\dashv\!\!\dashv$ to elements in **rC**.

To prove such results, let $\dashv\!\!\dashv_{\text{rC}}$ and $\triangleright_{\text{rC}}$ be, respectively, the compliance relations on retractable session contracts and the formal system axiomatizing it (see [18]).

We first show that the logical theories of $\triangleright_{\text{SC}}$ and $\triangleright_{\text{rC}}$ are conservative extensions of the logical theory \triangleright .

Proposition 8 (Formal systems conservativity).

- i) Let $\rho, \sigma \in \text{SC}$: $\triangleright_{\text{SC}} \rho \dashv\!\!\dashv \sigma$ iff $\triangleright \rho \dashv\!\!\dashv \sigma$
 ii) Let $\rho, \sigma \in \text{rC}$: $\triangleright_{\text{rC}} \rho \dashv\!\!\dashv \sigma$ iff $\triangleright \rho \dashv\!\!\dashv \sigma$.

Proof sketch. By inspection of the rules of the formal systems, and by $\text{SC} \subsetneq \text{rC} \subsetneq \text{rsC}$. \square

From Proposition 8 and the soundness and completeness property of $\triangleright_{\text{SC}}$ (Theorem 6) and $\triangleright_{\text{rC}}$ [18, Theorems 3.9 and 3.12] we immediately get what follows.

Corollary 9 (Compliances conservativity).

- i) Let $\rho, \sigma \in \text{SC}$: $\rho \dashv\!\!\dashv_{\text{SC}} \sigma$ iff $\rho \dashv\!\!\dashv \sigma$
 ii) Let $\rho, \sigma \in \text{rC}$: $\rho \dashv\!\!\dashv_{\text{rC}} \sigma$ iff $\rho \dashv\!\!\dashv \sigma$.

4. Duality and the subcontract relation

Unlike the retractable session contracts of [18], in the present setting a natural notion of *duality* exists. The dual $\overline{\sigma}$ of an element σ of **rsC** is obtained, as for session contracts, by interchanging any name a with \overline{a} and $+$ with \oplus .

Formally, we first define duality for (possibly open) contracts, that we dub **rsCo**, and then we restrict such a definition to **rsC** (i.e., to closed expressions).

Definition 14 (Syntactic duality).

- i) Let $\sigma \in \text{rsCo}$. The syntactic dual $\overline{\sigma}$ of σ is defined by the following clauses:

$$\overline{\mathbf{1}} = \mathbf{1} \quad \overline{\overline{x}} = x \quad \overline{\text{rec } x. \sigma} = \text{rec } x. \overline{\sigma}$$

$$\overline{\sum_{i \in I} \alpha_i. \sigma_i} = \oplus_{i \in I} \overline{\alpha_i}. \overline{\sigma_i} \quad \overline{\oplus_{i \in I} \alpha_i. \sigma_i} = \sum_{i \in I} \overline{\alpha_i}. \overline{\sigma_i}$$

- ii) We define $\overline{(\cdot)} : \text{rsCo} \rightarrow \text{rsC}$ as the restriction to **rsC** of the duality function on **rsCo**, observing that $\overline{\sigma} \in \text{rsC}$ iff $\sigma \in \text{rsC}$.

From now on, in order to avoid too cumbersome definitions, any time an inductive definition on elements of **rsC** is provided, it will be tacitly assumed to be the restriction to **rsC** of the corresponding inductive definition on **rsCo**.

A first relevant property of duality is that any contract is compliant with its syntactic dual.

Proposition 10. For any $\sigma \in \text{rsC}$, $\sigma \dashv\!\!\dashv \overline{\sigma}$.

Proof. Since $\overline{\sigma}$ is obtained from σ by exchanging each α with $\overline{\alpha}$ and $+$ with \oplus , it is easy to get a derivation of $\triangleright \overline{\sigma} \dashv\!\!\dashv \sigma$. The thesis is then an immediate consequence of soundness and completeness of \triangleright . \square

The notion of dual contract allows one to combine pairs of contracts in the compliance relation using a sort of “transitive” property, provided that the two contracts “in the middle” are dual of each other (see Appendix A.2 for the proof). Operationally, one can think at the two contracts in the middle as the specification of a bidirectional forwarder, and the proposition means that if a client and a server can correctly interact via a forwarder then they can also correctly interact directly.

Proposition 11. For any $\rho, \sigma, \sigma' \in \mathbf{rsC}$, $\rho \dashv\!\!\dashv \sigma$ and $\bar{\sigma} \dashv\!\!\dashv \sigma'$ imply $\rho \dashv\!\!\dashv \sigma'$.

We will provide further properties of duality using the notion of subcontract relation. Indeed, the notion of compliance naturally induces a substitutability relation on servers, denoted \preceq_s , that we call *subcontract relation for servers*. Intuitively, larger elements in the relation are compliant with a superset of the clients of smaller elements. Such a relation may be used for implementing contract-based query engines (see [30] for a detailed discussion). An analogous subcontract relation, denoted \preceq_c , can be defined for clients.

Definition 15 (Subcontract relations for servers and for clients). Let $\sigma, \sigma' \in \mathbf{rsC}$. We define

- i) $\sigma \preceq_s \sigma' \triangleq \forall \rho \in \mathbf{rsC} [\rho \dashv\!\!\dashv \sigma \text{ implies } \rho \dashv\!\!\dashv \sigma']$
- ii) $\sigma \preceq_c \sigma' \triangleq \forall \rho \in \mathbf{rsC} [\sigma \dashv\!\!\dashv \rho \text{ implies } \sigma' \dashv\!\!\dashv \rho]$

Using Proposition 11 we can characterize both \preceq_s and \preceq_c in terms of duality and compliance, relate them and get their decidability (see Appendix A.2 for the proof).

Theorem 12. For any $\sigma, \sigma' \in \mathbf{rsC}$:

- i) $\sigma \preceq_s \sigma' \text{ iff } \bar{\sigma} \dashv\!\!\dashv \sigma'$
- ii) $\sigma \preceq_c \sigma' \text{ iff } \sigma' \dashv\!\!\dashv \bar{\sigma}$
- iii) $\sigma \preceq_s \sigma' \text{ iff } \bar{\sigma}' \preceq_c \bar{\sigma}$
- iv) $\sigma \preceq_s \sigma'$ and $\sigma \preceq_c \sigma'$ are decidable.

By item (iii) above, from now on we can concentrate on the relation \preceq_s .

We can now characterize duality in terms of the subcontract relation for servers: given a client ρ , its dual $\bar{\rho}$ is a least element among all its possible servers, namely it is a possible server, and it is smaller than all the other possible servers.

Proposition 13 (Dual as a least element w.r.t. \preceq_s). Let $\rho \in \mathbf{rsC}$. Then $\bar{\rho}$ is a server for ρ , namely $\rho \dashv\!\!\dashv \bar{\rho}$, and more precisely it is a least element in the set of the servers of ρ , that is,

$$\forall \sigma \in \mathbf{rsC} : \rho \dashv\!\!\dashv \sigma \text{ implies } \bar{\rho} \preceq_s \sigma$$

Proof. Suppose that $\rho \dashv\!\!\dashv \sigma$ and take any contract δ such that $\delta \dashv\!\!\dashv \bar{\rho}$. Since $\bar{\bar{\rho}} = \rho$, by Proposition 11 we know that $\delta \dashv\!\!\dashv \sigma$; hence $\bar{\rho} \preceq_s \sigma$ by definition. \square

Since we have not yet proved that the subcontract relation is a partial order, we do not know yet whether $\bar{\rho}$ is also a minimal, i.e. there is no smaller element, neither whether other least elements or minimal elements exist. These questions will be answered by Proposition 17.

As done for the compliance relation, we characterize now the subcontract relation for servers in terms of derivability in the following formal system, where the symbol \ll is used as syntactical counterpart of the relation \preceq_s .

Definition 16 (Formal system for subcontract \blacktriangleright). Judgments in the formal system \blacktriangleright are expressions of the form $\Gamma \blacktriangleright \rho \ll \sigma$, where the environment Γ is a finite set of expressions of the form $\delta \ll \gamma$, with $\rho, \sigma, \delta, \gamma \in \mathbf{rsC}$. Axioms and rules are defined in Fig. 8.

The rules in system \blacktriangleright can be read as a translation of the rules in system \triangleright (see Fig. 4) via Theorem 12(i). As for \triangleright , in $\Gamma \blacktriangleright \rho \ll \sigma$ we may drop Γ if empty.

Lemma 14. $\Gamma \blacktriangleright \sigma \ll \sigma' \text{ iff } \tilde{\Gamma} \triangleright \bar{\sigma} \dashv\!\!\dashv \sigma'$, where $\Gamma = \{\sigma_i \ll \sigma'_i\}_{i \in I}$ and $\tilde{\Gamma} = \{\bar{\sigma}_i \dashv\!\!\dashv \sigma'_i\}_{i \in I}$.

Proof. (\Rightarrow) By induction over the derivation of $\Gamma \blacktriangleright \sigma \ll \sigma'$.

(\Leftarrow) By induction over the derivation of $\tilde{\Gamma} \triangleright \bar{\sigma} \dashv\!\!\dashv \sigma'$. \square

System \blacktriangleright is sound and complete for the subcontract relation \preceq_s .

Theorem 15 (Soundness and completeness of \blacktriangleright). $\blacktriangleright \sigma \ll \sigma' \text{ iff } \sigma \preceq_s \sigma'$

Proof. (\Rightarrow) Let $\blacktriangleright \sigma \ll \sigma'$. By Lemma 14 we get $\triangleright \bar{\sigma} \dashv\!\!\dashv \sigma'$ and hence $\bar{\sigma} \dashv\!\!\dashv \sigma'$ by soundness of system \triangleright . The thesis now descends from Theorem 12.

$$\begin{array}{c}
\text{(AX-}\preceq_s\text{)} \qquad \qquad \qquad \text{(HYP-}\preceq_s\text{)} \\
\Gamma \triangleright \mathbf{1} \ll \sigma' \qquad \qquad \Gamma, \sigma \ll \sigma' \triangleright \sigma \ll \sigma' \\
\\
\text{(\oplus \cdot + -}\preceq_s\text{)} \\
\frac{\Gamma, \alpha.\sigma_1 \oplus \sigma_2 \ll \alpha.\sigma'_1 + \sigma'_2 \triangleright \sigma_1 \ll \sigma'_1}{\Gamma \triangleright \alpha.\sigma_1 \oplus \sigma_2 \ll \alpha.\sigma'_1 + \sigma'_2} \\
\\
\text{(+ \cdot + -}\preceq_s\text{)} \\
\frac{\forall h \in I. \Gamma, \sum_{i \in I} \alpha_i.\sigma_i \ll \sum_{j \in I \cup J} \alpha_j.\sigma'_j \triangleright \sigma_h \ll \sigma'_h}{\Gamma \triangleright \sum_{i \in I} \alpha_i.\sigma_i \ll \sum_{j \in I \cup J} \alpha_j.\sigma'_j} \\
\\
\text{(\oplus \cdot \oplus -}\preceq_s\text{)} \\
\frac{\forall h \in I. \Gamma, \bigoplus_{j \in I \cup J} \alpha_j.\sigma_j \ll \bigoplus_{i \in I} \alpha_i.\sigma'_i \triangleright \sigma_h \ll \sigma'_h}{\Gamma \triangleright \bigoplus_{j \in I \cup J} \alpha_j.\sigma_j \ll \bigoplus_{i \in I} \alpha_i.\sigma'_i}
\end{array}$$


Fig. 8. The formal system \triangleright .

Fig. 9. Subcontract preorder: a sample.

(\Leftarrow) Let $\sigma \preceq_s \sigma'$. By Theorem 12 we have that $\bar{\sigma} \dashv \sigma'$. By completeness of system \triangleright we get $\triangleright \bar{\sigma} \dashv \sigma'$. Now, by Lemma 14, we can obtain $\triangleright \sigma \ll \sigma'$. \square

System \triangleright can be used to show that \preceq_s is antisymmetric (see Appendix A.2) and hence that it is a partial order. Moreover $\bar{\rho}$ is also the minimum server of ρ : it is minimal, hence there is no smaller server, and there is a unique minimal.

Proposition 16. *The relation \preceq_s is antisymmetric. That is, for any $\sigma, \sigma' \in \mathbf{rsC}$,*

$$\sigma \preceq_s \sigma' \text{ and } \sigma' \preceq_s \sigma \text{ implies } \sigma = \sigma'$$

Proposition 17. *\preceq_s is a partial order $\wedge \forall \rho \in \mathbf{rsC}, \bar{\rho}$ is the minimum server of ρ .*

Proof sketch. We need to show \preceq_s to be reflexive, transitive and antisymmetric. Reflexivity and transitivity immediately descend from the definition of \preceq_s (Definition 15). Antisymmetry follows from Proposition 16.

For the second conjunct of the statement, suppose, towards a contradiction, that there exists $\sigma \neq \bar{\rho}$ such that $\rho \dashv \sigma$ and $\sigma \preceq_s \bar{\rho}$. By Proposition 13 we have $\bar{\rho} \preceq_s \sigma$. By antisymmetry we have $\sigma = \bar{\rho}$, against the hypothesis. \square

The structure of the partial order is shown in Fig. 9, where the relations between terms with a unique choice among actions $a, b, c, \bar{a}, \bar{b}$ and \bar{c} are pictured. Notably, duality can be seen graphically as central inversion.

Remark 4. Analogously to what done in Section 3.1, one can show the subcontract relation \preceq_s to be a conservative extension of the corresponding notion in SC. Moreover, the restriction of \preceq_s to \mathbf{rC} provides a suitable notion of subcontract for \mathbf{rC} (which has never been studied before).

5. Complexity issues

The decision algorithm **Prove** for compliance of Fig. 5 is simple, but, as it is, its complexity is exponential. The example below shows how an exponential number of recursive calls can be actually reached. It is an adaptation of the example presented in [31](§11) for the subtyping relation for recursive arrow and product types.

$$\begin{array}{c}
\text{(A}\infty\text{)} \\
\triangleright \mathbf{1} \dashv \sigma \\
\\
\frac{(\oplus \cdot +\infty)}{(\forall h \in I) \rho_h \dashv \sigma_h} \\
\oplus_{i \in I} \bar{\alpha}_i \cdot \rho_i \dashv \sum_{j \in I \cup J} \alpha_j \cdot \sigma_j \\
\\
\text{(A}\infty\text{)} \\
\frac{(\cdot +\infty)}{\rho \dashv \sigma} \\
\alpha \cdot \rho + \rho' \dashv \bar{\alpha} \cdot \sigma + \sigma' \\
\\
\frac{(\oplus \cdot +\infty)}{(\forall h \in I) \rho_h \dashv \sigma_h} \\
\sum_{j \in I \cup J} \bar{\alpha}_j \cdot \rho_j \dashv \oplus_{i \in I} \alpha_i \cdot \sigma_i
\end{array}$$

Fig. 10. The non-well founded system \triangleright_∞ .

For each $n \in \mathbb{N}$ we define two contracts ρ_n and σ_n by induction, as follows.

$$\begin{array}{ll}
\rho_0 = \text{rec } x.(a.x + b.x) & \rho_{n+1} = \text{rec } x.(a.x + b.\rho_n) \\
\sigma_0 = \text{rec } x.(\bar{a}.x \oplus \bar{b}.x) & \sigma_{n+1} = \text{rec } x.(\bar{a}.\sigma_n \oplus \bar{b}.x)
\end{array}$$

As for the example in [31], the size of ρ_n and σ_n is linear in n , since ρ_n and σ_n appear just once in the definitions of ρ_{n+1} and σ_{n+1} , respectively.

By complete induction over n it is possible to prove that

Fact 18. For any $n \in \mathbb{N}$, $\rho_n \dashv \sigma_n$.

In particular, the computation of $\text{Prove}(\emptyset \triangleright \rho_n \dashv \sigma_n)$ builds a derivation for $\triangleright \rho_n \dashv \sigma_n$ in an actual exponential number of calls. In order to show that, let us first notice that any call of the shape $\text{Prove}(\Gamma \triangleright \rho_k \dashv \sigma_k)$, with $k \neq 0$ and $\rho_k \dashv \sigma_k \notin \Gamma$, produces two immediate calls: $\text{Prove}(\Gamma_1 \triangleright \rho_k \dashv \sigma_{k-1})$ and $\text{Prove}(\Gamma_1 \triangleright \rho_{k-1} \dashv \sigma_k)$. These latter two calls, produce, besides others, a call of $\text{Prove}(\Gamma_2 \triangleright \rho_{k-1} \dashv \sigma_{k-1})$ and one of $\text{Prove}(\Gamma_3 \triangleright \rho_{k-1} \dashv \sigma_{k-1})$, with $\Gamma_2 \neq \Gamma_3$, as it is shown in the following recursive-call tree (where we abbreviate “Prove” by “Pr”).

$$\begin{array}{c}
\text{Pr}(\Gamma \triangleright \rho_k \dashv \sigma_k) \\
\text{Pr}(\Gamma_1 \triangleright \rho_k \dashv \sigma_{k-1}) \quad \text{Pr}(\Gamma_1 \triangleright \rho_{k-1} \dashv \sigma_k) \\
\text{Pr}(\Gamma_2 \triangleright \rho_{k-1} \dashv \sigma_{k-2}) \quad \text{Pr}(\Gamma_2 \triangleright \rho_{k-1} \dashv \sigma_{k-1}) \quad \text{Pr}(\Gamma_3 \triangleright \rho_{k-1} \dashv \sigma_{k-1}) \quad \text{Pr}(\Gamma_3 \triangleright \rho_{k-2} \dashv \sigma_k) \\
\dots \dots \text{ etc.}
\end{array}$$

where $\Gamma_1 = \{\Gamma, \rho_k \dashv \sigma_k\}$

and $\Gamma_2 = \{\Gamma, \rho_k \dashv \sigma_k, \rho_k \dashv \sigma_{k-1}\} \neq \{\Gamma, \rho_k \dashv \sigma_k, \rho_{k-1} \dashv \sigma_k\} = \Gamma_3$.

So, given n , the computation of $\text{Pr}(\emptyset \triangleright \rho_n \dashv \sigma_n)$ results in an overall number of *at least* 2^n recursive calls.

Nonetheless, the complexity of the compliance decision procedure can be drastically reduced down to a polynomial – actually quadratic – complexity, as shown below.

A quadratic decision algorithm.

We first define a non-well founded version of system \triangleright , that we dub \triangleright_∞ .

Definition 17 (The non-well founded system \triangleright_∞). We write $\triangleright_\infty \rho \dashv \sigma$ whenever there exists a finite or infinite derivation tree formed by the rules in Fig. 10 having $\rho \dashv \sigma$ as conclusion, and such that each finite branch ends with an instance of axiom (A ∞).

Systems \triangleright and \triangleright_∞ are equivalent (see Appendix A.3 for the proof).

Lemma 19. $\triangleright \rho \dashv \sigma$ iff $\triangleright_\infty \rho \dashv \sigma$.

In Fig. 11 we present a decision algorithm Decide_{\dashv} , based on the procedures \mathbf{P} and \mathbf{P}^+ . Whereas the procedure Prove returns, if any, a derivation in \triangleright , Decide_{\dashv} just checks for the existence of a derivation in \triangleright_∞ (and hence in \triangleright by Lemma 19). Its execution resembles that of an alternating Turing machine, where nodes corresponding to rules $(\oplus \cdot +\infty)$ and $(\cdot +\infty)$ are universal and nodes corresponding to $(+\cdot +\infty)$ are existential; $\mathbf{P}(A, F, L, b)$ attempts to prove all statements in its goal list L , while $\mathbf{P}^+(A, F, L, b)$ succeeds if at least one goal in L is satisfiable.

The procedure \mathbf{P} is an adaptation of the concrete subtyping algorithm for recursive arrow and product types of [31](§10) to the present, more complex context, where \sum and \oplus can be looked at as constructors of arbitrary arity. It consists of a proof reconstruction procedure for \triangleright_∞ using a *depth-first* technique. \mathbf{P} accumulates in its first argument A all the judgments it encounters during the search, in order to avoid looping over the same judgments (a role similar to Γ in system \triangleright). With

```

Decide-||( $\rho \dashv \sigma$ ) = let (A, F, b) = P( $\emptyset, \emptyset, [\rho \dashv \sigma], \mathbf{ok}$ )
in b = ok

where

P(A, F, [ ], b) = (A, F, b)
P(A, F, ( $\rho \dashv \sigma$ ) : xs, b) =
-1- if  $\rho = \mathbf{1}$  then P(A, F, xs, b)
-2- else if  $\rho \dashv \sigma \in A$  then P(A, F, xs, b)
-3- else if  $\rho \dashv \sigma \in F$  then (A, F, fail)
-4- else if  $\rho = \sum_{i \in I} \alpha_i . \rho_i$  and  $\sigma = \sum_{j \in J} \bar{\alpha}_j . \sigma_j$  and  $I \cap J = \{i_1, \dots, i_n\}$ 
-5- then let (A0, F0, b0) = P+(A ∪ { $\rho \dashv \sigma$ }, F, [ $\rho_{i_1} \dashv \sigma_{i_1} \dots \rho_{i_n} \dashv \sigma_{i_n}$ ], b)
-6- in if b0 = fail then (A0, F0, fail)
-7- else P(A0, F0, xs, b0)
-8- else if  $\rho = \bigoplus_{i \in I} \bar{\alpha}_i . \rho_i$  and  $\sigma = \sum_{j \in I \cup J} \alpha_j . \sigma_j$  and  $I = \{i_1, \dots, i_n\}$ 
-9- then let (A0, F0, b0) = P(A ∪ { $\rho \dashv \sigma$ }, F, [ $\rho_{i_1} \dashv \sigma_{i_1} \dots \rho_{i_n} \dashv \sigma_{i_n}$ ], b)
-10- in if b0 = fail then (A0, F0, fail)
-11- else P(A0, F0, xs, b0)
-12- else if  $\rho = \sum_{j \in I \cup J} \bar{\alpha}_j . \rho_j$  and  $\sigma = \bigoplus_{i \in I} \alpha_i . \sigma_i$  and  $I = \{i_1, \dots, i_n\}$ 
-13- then let (A0, F0, b0) = P(A ∪ { $\rho \dashv \sigma$ }, F, [ $\rho_{i_1} \dashv \sigma_{i_1} \dots \rho_{i_n} \dashv \sigma_{i_n}$ ], b)
-14- in if b0 = fail then (A0, F0, fail)
-15- else P(A0, F0, xs, b0)
-16- else if  $\rho = \text{rec } x . \rho'$  then P(A, F, ({ $\text{rec } x . \rho' / x$ }) $\rho' \dashv \sigma$ ) : xs, b)
-17- else if  $\sigma = \text{rec } x . \sigma'$  then P(A, F, ( $\rho \dashv \{\text{rec } x . \sigma' / x\} \sigma'$ ) : xs, b)
-18- else (A, F ∪ { $\rho \dashv \sigma$ }, fail)

and where

P+(A, F, [ $\rho \dashv \sigma$ ], b) = P(A, F, [ $\rho \dashv \sigma$ ], b)
P+(A, F, ( $\rho \dashv \sigma$ ) : xs, b) =
-19- let (A0, F0, b0) = P(A, F, [ $\rho \dashv \sigma$ ], b) in
-20- if b0 = fail then P+(A ∪ A0, F ∪ F0, xs, ok)
-21- else (A0, F0, b0)

```

Fig. 11. The quadratic decision procedure for compliance.

respect to the algorithm in [31](§10) we have two further parameters, F and b. The argument F accumulates the judgments for which it has been found that no derivation exists. When a rule ($+ \cdot +_{\infty}$) is encountered, the algorithm proceeds by calling the procedure **P**⁺ which, in case a premise is unprovable, goes on checking the other premises. The negative information inferred about unprovable judgments is stored in F and it is carried along by the procedure **P**⁺ (as well as the positive information stored in A) in order not to duplicate work. The argument b, that can be either **ok** or **fail**, is used to record whether the last call was successful or not, and it is used by **P**⁺ to know whether it has to stop with success, or to check a new premise.

Let us note that, contrary to the previous treatment, while studying the algorithm **Decide**_{-||}, we abandon the equi-recursive view of recursion, and we represent a contract by a particular explicit and (possibly) recursive expression.

We now give an upper bound on the complexity of the algorithm **Decide**_{-||} (see Appendix A.3 for the proof), based on the notion of size of a contract below.

Definition 18. The size $\text{size}(\sigma)$ of a contract σ is inductively defined as follows:

$$\begin{aligned}
 \text{size}(\mathbf{1}) &= 1 \\
 \text{size}(\sum_{i \in I} \alpha_i . \sigma_i) &= \sum_{i \in I} (1 + \text{size}(\sigma_i)) & \text{size}(x) &= 1 \\
 \text{size}(\bigoplus_{i \in I} \alpha_i . \sigma_i) &= \sum_{i \in I} (1 + \text{size}(\sigma_i)) & \text{size}(\text{rec } x . \sigma) &= 1 + \text{size}(\sigma)
 \end{aligned}$$

Proposition 20 (Complexity of **Decide**_{-||}). The complexity of the algorithm **Decide**_{-||} is $\mathcal{O}(n^2)$, where, for an argument $\rho \dashv \sigma$, n is the maximum between $\text{size}(\rho)$ and $\text{size}(\sigma)$.

We remark that for obtaining such a complexity it is fundamental to avoid to recompute information, positive or negative, hence the need for sets A and F. It is also fundamental to be able to perform insertion and membership check in those sets

in constant time. This can be done by implementing them as Boolean vectors indexed by all the pairs of subterms of the original ρ and σ .

From Proposition 20, using Theorem 12, it is immediate to get a decision procedure for the subcontract relation out of that for compliance:

$$\mathbf{Decide}_{\preceq_s}(\rho \preceq_s \sigma) = \mathbf{Decide}_{\perp}(\bar{\rho} \dashv \sigma)$$

Corollary 21 (Complexity of $\mathbf{Decide}_{\preceq_s}$). *The complexity of the algorithm $\mathbf{Decide}_{\preceq_s}$ is $\mathcal{O}(n^2)$, where, for an argument $\rho \preceq_s \sigma$, n is the maximum between $\text{size}(\rho)$ and $\text{size}(\sigma)$.*

Proof. Immediate by noticing that building the dual of a given contract takes linear time. \square

Remark 5. The polynomial decision procedure \mathbf{Decide}_{\perp} applies also to the formalism of retractable session contracts of [18]. In fact, the sound and complete formal system $\triangleright_{\mathbf{rC}}$ (and the corresponding decision procedure) is the restriction to elements of \mathbf{rC} of the system in Fig. 4. Obviously, when applied to elements of \mathbf{rC} , the clauses -8- and -12- of \mathbf{Decide}_{\perp} do not need to take into account the possibility of internal input choices.

6. Retractable contracts vs reversible computing

In this section we explore the relations between our retractable contracts and calculi for reversible computing (see [25] for an overview). In [16], Phillips and Ulidowski provide an automatic technique to derive, from the forward semantics of a given calculus, its reversible semantics. In principle, we would like to apply the technique in [16] to the forward calculus underlying our retractable contracts, which is presented below. It is built by equipping retractable contracts with the semantics obtained by replacing both a and \bar{a} with α in the semantics of session contracts (recalled in Definition 12 in Section 3.1.1).

Definition 19 (Retractable contracts underlying semantics). *The following rules define the LTS for retractable contracts and the reduction semantics for client/server pairs.*

$$\frac{\alpha.\sigma \oplus \sigma' \xrightarrow{\tau}_U \alpha.\sigma}{\rho \parallel \sigma \rightarrow_U \rho' \parallel \sigma} \quad \frac{\alpha.\sigma \xrightarrow{\alpha}_U \sigma}{\rho \parallel \sigma \rightarrow_U \rho \parallel \sigma'} \quad \frac{\alpha.\sigma + \sigma' \xrightarrow{\alpha}_U \sigma}{\rho \parallel \sigma \rightarrow_U \rho' \parallel \sigma'} \quad \frac{\rho \xrightarrow{\alpha}_U \rho' \quad \sigma \xrightarrow{\bar{\alpha}}_U \sigma'}{\rho \parallel \sigma \rightarrow_U \rho' \parallel \sigma'}$$

However, the direct application of the technique in [16] requires the forward semantics to be specified as an LTS and to satisfy a number of conditions. More precisely, the semantics should be specified by rules in the simple path format [16, Def. 2.1] which must be either static rules [16, Def. 2.2], choice rules [16, Def. 2.3] or choice axioms [16, Def. 2.4]. The semantics in Definition 19 is a two-level semantics, where the lower level is an LTS and the top level a reduction semantics. Furthermore, some of the rules do not satisfy the required conditions. For instance, the rule for internal choice does not belong to any of the classes of rules above.

Thus, in order to directly apply the technique, we transform the syntax and the semantics of our forward calculus as follows:

- we merge the two levels of syntax (contracts and client/server pairs) into one, and specify its semantics as an LTS by considering reductions as transitions with label τ ;
- we transform internal choice into τ -prefixed external choice;
- we separate action prefixing from internal/external choice.

The resulting calculus is defined as follows.

Definition 20 (Transformed contracts). *The set TC of transformed contracts is the set of closed expressions generated by the following grammar.*

$$\sigma := \alpha_{\tau}.\sigma \mid \sum_{i \in I} \sigma_i \mid x \mid \text{rec } x.\sigma \mid \sigma \parallel \sigma' \mid \mathbf{1}$$

where α_{τ} denotes a , \bar{a} or τ .

We define below a translation function $\llbracket \cdot \rrbracket$ from either a retractable contract σ or a client/server pair $\sigma \parallel \rho$ into transformed contracts.

Definition 21 (*Translation function*). The translation function $\llbracket \cdot \rrbracket : \mathbf{rsCo} \cup (\mathbf{rsCo} \times \mathbf{rsCo}) \longrightarrow \mathbf{TC}$ is defined inductively as follows:

$$\begin{aligned} \llbracket \sum_{i \in I} \alpha_i . \sigma_i \rrbracket &= \sum_{i \in I} \alpha_i . \llbracket \sigma_i \rrbracket & \llbracket x \rrbracket &= x & \llbracket \text{rec } x . \sigma \rrbracket &= \text{rec } x . \llbracket \sigma \rrbracket \\ \llbracket \bigoplus_{i \in I} \alpha_i . \sigma_i \rrbracket &= \sum_{i \in I} \tau . \alpha_i . \llbracket \sigma_i \rrbracket & \llbracket \mathbf{1} \rrbracket &= \mathbf{1} & \llbracket \sigma \parallel \rho \rrbracket &= \llbracket \sigma \rrbracket \parallel \llbracket \rho \rrbracket \end{aligned}$$

Transformed contracts allow for general parallel composition and mixed choice. However, the restriction of transformed contracts to the image of function $\llbracket \cdot \rrbracket$ is closed under reduction, as shown by the semantics below. Thus, from now on, we consider only the transformed contracts in the image of $\llbracket \cdot \rrbracket$.

Definition 22 (*Semantics of transformed contracts*). The rules below, plus their symmetric, define the semantics of transformed contracts.

$$\begin{array}{c} \alpha_\tau . \sigma \xrightarrow{\alpha_\tau} \sigma \\ \rho \xrightarrow{\tau} \rho' \\ \hline \rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma \end{array} \quad \begin{array}{c} \sigma \xrightarrow{\alpha_\tau} \sigma' \\ \sigma + \rho \xrightarrow{\alpha_\tau} \sigma' \\ \rho \xrightarrow{\alpha} \rho' \quad \sigma \xrightarrow{\bar{\alpha}} \sigma' \\ \hline \rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma' \end{array}$$

It is easy to check that the LTS for transformed contracts and the LTS underlying retractable contracts model the same client/server interactions.

Proposition 22. Let $\sigma, \rho, \sigma', \rho' \in \mathbf{rsC}$.

$$\sigma \parallel \rho \longrightarrow_U \sigma' \parallel \rho' \quad \text{iff} \quad \llbracket \sigma \parallel \rho \rrbracket \xrightarrow{\tau} \llbracket \sigma' \parallel \rho' \rrbracket$$

Proof. By inspection of the rules. \square

One can apply to the LTS in Definition 22 the technique in [16], obtaining the LTS below. In order to simplify the treatment, we replaced the keys used in [16] to annotate actions with an underline. While this is not correct in general, it is correct in the image of our client/server pairs, since keys are used to distinguish interactions with different communication partners, but in our case for each action there is at most one possible partner.

Definition 23 (*Reversible transformed contracts*). The rules below, plus their symmetric, define the *forward semantics* of reversible transformed contracts.

$$\begin{array}{c} \text{std}(X) \\ \hline \alpha_\tau . X \xrightarrow{\alpha_\tau} \underline{\alpha_\tau} . X \\ X \xrightarrow{\tau} X' \\ \hline X \parallel Y \xrightarrow{\tau} X' \parallel Y \end{array} \quad \begin{array}{c} X \xrightarrow{\alpha_\tau} X' \\ \hline \underline{\alpha_\tau} . X \xrightarrow{\alpha_\tau} \underline{\alpha_\tau} . X' \\ X \xrightarrow{\alpha} X' \quad Y \xrightarrow{\bar{\alpha}} Y' \\ \hline X \parallel Y \xrightarrow{\tau} X' \parallel Y' \end{array} \quad \begin{array}{c} X \xrightarrow{\alpha_\tau} X' \quad \text{std}(Y) \\ \hline X + Y \xrightarrow{\alpha_\tau} X' + Y \end{array}$$

In the rules, $\text{std}(X)$ holds if X does not contain underlined prefixes.

The *backward semantics* of reversible transformed contracts, denoted by arrow \rightsquigarrow , is obtained by changing the direction of the arrows in the rules above.

Fig. 12 shows how a sample reduction sequence for a retractable client/server pair corresponds to a reduction sequence for the corresponding transformed contract. We formalize the correspondence hinted at in the example by defining a notion of *barbed simulation* between retractable client/server pairs and transformed contracts, and showing that each client/server pair (with empty histories) is simulated by its translation (see Appendix A.4 for the proof). We start by defining barbs for both retractable client–server pairs and reversible transformed contracts.

Definition 24 (*Barbs*). Given a retractable client–server pair $H_1 \times \rho \parallel H_2 \times \sigma$ and a name or coname α , we say that $H_1 \times \rho \parallel H_2 \times \sigma$ has a barb at α , written $H_1 \times \rho \parallel H_2 \times \sigma \downarrow_\alpha$, when defined by the following structural induction on $H_1 \times \rho \parallel H_2 \times \sigma$:

$$\begin{aligned} H_1 \times \rho \parallel H_2 \times \sigma \downarrow_\alpha &\text{ iff } H_1 \times \rho \downarrow_\alpha \text{ or } H_2 \times \sigma \downarrow_\alpha \\ H_1 \times \rho \downarrow_\alpha &\text{ iff } \rho \downarrow_\alpha \\ \sum_{i \in I} \alpha_i . \sigma_i \downarrow_\alpha & \\ \bigoplus_{i \in I} \alpha_i . \sigma_i \downarrow_\alpha & \\ \text{rec } x . \sigma \downarrow_\alpha &\text{ iff } \sigma \downarrow_\alpha \end{aligned}$$

$\llbracket a.c + b.d \rrbracket = a.c + b.d$	$\llbracket \bar{a}.\bar{c} + \bar{b}.(e \oplus f) \rrbracket = \bar{a}.\bar{c} + \bar{b}.(e + f)$
$\langle \rangle \times a.c + b.d \parallel \langle \rangle \times \bar{a}.\bar{c} + \bar{b}.(e \oplus f)$	$a.c + b.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f)$
$\rightarrow_f a.c \times d \parallel \bar{a}.\bar{c} \times e \oplus f$	$\rightarrow a.c + \underline{b}.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f)$
$\rightarrow_f a.c \times d \parallel \bar{a}.\bar{c} \times e$	$\rightarrow a.c + \underline{b}.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f)$
$\rightarrow_b \langle \rangle \times a.c \parallel \langle \rangle \times \bar{a}.\bar{c}$	$\rightsquigarrow a.c + \underline{b}.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f)$
$\rightarrow_f \circ \times c \parallel \circ \times \bar{c}$	$\rightarrow \underline{a}.c + b.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f)$
$\rightarrow_f \circ : \circ \times \mathbf{1} \parallel \circ : \circ \times \mathbf{1}$	$\rightarrow \underline{a}.\underline{c} + b.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f)$

Fig. 12. Retractable client/server pairs vs transformed contracts: an example.

Similarly, given a reversible transformed contract X and a name or coname α , we say that X has a barb at α , written $X \downarrow_\alpha$, when defined by the following structural induction on X :

$$\begin{aligned}
& \alpha.X \downarrow_\alpha \\
& \tau.X \downarrow_\alpha \text{ iff } X \downarrow_\alpha \\
& \underline{\beta}.X \downarrow_\alpha \text{ iff } X \downarrow_\alpha \\
& \underline{\tau}.X \downarrow_\alpha \text{ iff } X \downarrow_\alpha \\
& \sum_{i \in I} X_i \downarrow_\alpha \text{ iff } (\text{std}(\sum_{i \in I} X_i) \wedge \exists i. X_i \downarrow_\alpha) \text{ or } (\exists i. X_i = \underline{\alpha} \wedge X_i \downarrow_\alpha) \\
& X \parallel X' \downarrow_\alpha \text{ iff } X \downarrow_\alpha \text{ or } X' \downarrow_\alpha \\
& \text{rec } x.X \downarrow_\alpha \text{ iff } X \downarrow_\alpha
\end{aligned}$$

While the definition of barbs for retractable client–server pairs is standard, the one for transformed contracts has a few peculiarities, that we now discuss. The main point is that underlined actions are actions that have already been executed, hence they do not show barbs, but propagate barbs from their continuation. This is particularly tricky in the case of choice: if the transformed contract is not standard (that is, it has some underlined action), then only the selected branch (that is, the one with an underlined action) needs to be inspected for barbs. For instance, $a.c + b.d \downarrow_a$ and $a.c + b.d \downarrow_b$, while $\underline{a}.c + b.d \downarrow_c$ and has no other barb. We also notice that τ actions propagate barbs from their continuation.

Definition 25 (Barbed simulation). Let R be a relation between retractable client/server pairs and reversible transformed contracts. R is a barbed simulation iff for each $(H_1 \times \rho \parallel H_2 \times \sigma, X) \in R$:

- if $H_1 \times \rho \parallel H_2 \times \sigma \rightarrow_f H'_1 \times \rho' \parallel H'_2 \times \sigma'$ then $X \xrightarrow{\tau} X'$ and $(H'_1 \times \rho' \parallel H'_2 \times \sigma', X') \in R$;
- if $H_1 \times \rho \parallel H_2 \times \sigma \rightarrow_b H'_1 \times \rho' \parallel H'_2 \times \sigma'$ then $X \xrightarrow{\tau^+} X'$ and $(H'_1 \times \rho' \parallel H'_2 \times \sigma', X') \in R$;
- if $H_1 \times \rho \parallel H_2 \times \sigma \downarrow_\alpha$ then $X \downarrow_\alpha$;

where $\xrightarrow{\tau^+}$ is the transitive closure of $\xrightarrow{\tau}$.

Theorem 23. For any $\rho, \sigma \in \text{rsC}$ there is a simulation R such that $(\langle \rangle \times \rho \parallel \langle \rangle \times \sigma, \llbracket \rho \parallel \sigma \rrbracket) \in R$.

Note that the opposite of Theorem 23 cannot hold because of the mechanisms to control reversibility discussed in Remark 2. Indeed, the technique in [16] generates an uncontrolled semantics. A sample difference is that in transformed contracts we can have an infinite reduction sequence persistently choosing the right branch after the backward reduction, as follows.

$$\begin{aligned}
& a.c + b.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f) \xrightarrow{+} a.c + \underline{b}.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f) \\
& \rightsquigarrow^+ a.c + b.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f) \xrightarrow{+} a.c + \underline{b}.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f) \\
& \rightsquigarrow^+ a.c + b.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + f) \xrightarrow{+} \text{etc.}
\end{aligned}$$

It is easy to check, instead, that $\langle \rangle \times a.c + b.d \parallel \langle \rangle \times \bar{a}.\bar{c} + \bar{b}.(e \oplus f)$ can perform no infinite reduction sequence since the chosen branch is discarded upon rollback.

We remark that the present investigation just relates the retractable semantics of contracts and reversible calculi, showing that the retractable semantics correctly models a form of reversibility. In itself, this provides little justification to the speculative semantics. However, since the retractable and the speculative semantics give rise to the same notion of compliance, if the retractable semantics is correct than also the speculative semantics needs to be correct, at least from the abstract point of view determined by considering compliance.

7. Related work and conclusion

We have presented two conservative extensions of the session contracts of [32,13,14], a formalism interpreting session types [4] into a subset of contracts [1–3]. One extension deals with backtracking and one with speculative execution. We have shown that they both give rise to *the same* compliance relation, and, as a consequence, to the same subcontract (both for servers and for clients) and duality relations. For each of these relations we provided syntactic characterizations of the semantic concepts, allowing for efficient ways of checking them.

We discussed in the Introduction the improvements with respect to the previous work presented in [18,17]. Another closely related work is [33,34], where a different form of contracts with rollback is presented. Our retractable contracts depart from that model on three main aspects: (1) we use rollback in a disciplined way to tolerate failures in the interaction, thus making it easier for contracts to be compliant, while in [33,34] the decision to rollback is internal to each participant, and the partner needs to be compliant with both forward behavior and all the behaviors emerging from each possible rollback; (2) we embed checkpoints in the structure of contracts, avoiding explicit checkpoints; (3) we keep a stack of “pasts”, instead of just a single past as in [33,34].

Reversibility, generalizing backtracking by allowing one to go back to any past state, has also been studied in the setting of binary session types [35–37]. There however the emphasis is on defining the reversible engine, based on causal-consistent reversibility [25], and not on studying compliance or subtyping (which would correspond to our subcontract relation).

Similarly to our retractable contracts, long running transactions with compensations, and in particular interacting transactions [38], allow one to undo past agreements. In interacting transactions, however, abort (which corresponds to our backtracking) can occur at any time, not only when an agreement cannot be found as in our case. Also, each transaction offers just two possibilities, and they are sorted: first the normal execution, then the compensation. Finally, compliance of interacting transactions has never been studied.

In [39] a game-theoretical interpretation of the retractable session contracts of [18] has been provided. Such an interpretation is likely to extend to the retractable contracts presented here.

It would also be worth to investigate whether our approach can be extended to multi-party sessions [40]. Uncontrolled reversibility (according to the terminology in [25]) in multi-party sessions has been studied in [41]. An investigation of multi-party sessions with rollbacks and named checkpoints has also been already undertaken, in [42,43]. In such a paper, however, the cause of a rollback is not a synchronization failure, but it is completely transparent to the calculus. Moreover, chosen branches are not discarded and can be retried upon rollback.

Because of the relevance of higher-order features in type systems, and of session delegation in type systems with sessions in particular, also higher-order session contracts, i.e. session contracts with delegation, have been investigated [13,44]. It is hence worth studying the integration of backtracking (or speculative execution) and session delegation.

As mentioned in Example 2, any possible preference among alternatives in external output choices is abstracted away in our contracts. Actually, it would be possible to model preference by dropping commutativity of $+$ and thus considering an order for alternatives in choices, with the highest priority option selected always first. Such a possibility has been explored in [26], in the context of orchestrated session types. There the information obtained when compliance is decided is embodied in an orchestrator enabling to drive the choices of processes. In presence of a choice which can possibly lead to a failure, the orchestrator forces the processes to follow one of the *safe* branches, so avoiding the need for mechanisms such as backtracking or speculative execution. Instead, a setting combining rollback and preference among alternatives has been studied from an operational point of view in [45].

A last line of future work concerns the applicability of our results to real world applications. In general, contracts can be extracted from the specification or the code of clients and servers (see for instance [46], where a form of contracts is extracted from Java programs and exploited for deadlock analysis), and then used to check compliance, thus guaranteeing successful interaction. Of course, our contracts come handy when clients and servers use either backtracking, as provided, e.g., by the CBack library [47] for C, or speculative parallelism, as provided, e.g., by the C# library described in [8]. How to actually extract from client/server software based on such or similar libraries the corresponding retractable or speculative contracts in a rigorous and possibly automatic way is still to be seen.

Acknowledgements

This work was partially supported by the COST Action IC1405 on “Reversible computation – extending horizons of computing”. The first and third authors were partially supported also by the COST Action EUTYPES CA-15123 and, respectively, Project “Chance” of the University of Catania and Project FORMS 2015 of University of Turin. The second author is member of the INdAM Research group GNCS. We thank Mariangiola Dezani-Ciancaglini for interesting discussions and helpful suggestions. We also thank the anonymous reviewers for their useful comments and remarks.

Appendix A. Proofs

A.1. Proofs of Section 3

Proof of Theorem 1. It is straightforward to check the following

Fact 24.

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 \mathcal{D} has conclusion $\Gamma \triangleright \rho \dashv \sigma$.

By Fact 24, we only need to show that the procedure **Prove** always terminates. Note 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, respectively, ρ and σ (because of the equi-recursive view 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

Soundness and Completeness Proofs (Theorems 2 and 3)

We begin with the proof of Soundness and Completeness of system \triangleright with respect to the retractable compliance.

Retractable soundness and completeness. 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 25. If $\langle \rangle \times \rho \parallel \langle \rangle \times \sigma \xrightarrow{*} H_1 \times \rho' \parallel H_2 \times \sigma' \dashv$, then $H_1 = H_2 = \langle \rangle$.

Proof. Clearly $H_1 \times \rho' \parallel H_2 \times \sigma' \dashv$ implies either $H_1 = \langle \rangle$ or $H_2 = \langle \rangle$. Observe that: rule (comm) adds one element to both stacks; rule (τ) does not modify any of the 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 $H_1 = \langle \rangle$ implies $H_2 = \langle \rangle$ and vice versa. \square

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

Lemma 26. If $H_1 \times \rho \dashv^R H_2 \times \sigma$, then $H'_1 : H_1 \times \rho \dashv^R H'_2 : H_2 \times \sigma$ for all H'_1, H'_2 .

Proof sketch. It suffices to show that if in a sequence of reductions H'_1 or H'_2 are used, then $H_1 \times \rho \dashv^R H_2 \times \sigma$. \square

The following lemma gives all possible shapes of compliant contracts.

Lemma 27. We have $\rho \dashv^R \sigma$ iff 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 $\exists k \in I \cap J$. $\rho_k \dashv^R \sigma_k$;
3. $\rho = \bigoplus_{i \in I} \bar{\alpha}_i \cdot \rho_i$, $\sigma = \sum_{j \in J} \alpha_j \cdot \sigma_j$, $I \subseteq J$ and $\forall k \in I$. $\rho_k \dashv^R \sigma_k$;
4. $\rho = \sum_{i \in I} \bar{\alpha}_i \cdot \rho_i$, $\sigma = \bigoplus_{j \in J} \alpha_j \cdot \sigma_j$, $I \supseteq J$ and $\forall k \in J$. $\rho_k \dashv^R \sigma_k$.

Proof. The if part is immediate. We prove the only if part by contraposition and by cases on 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 $\rho_{k_i} \dashv^R \sigma_{k_i}$ for $1 \leq i \leq n$. Then we get

$$\langle \rangle \times \rho_{k_i} \parallel \langle \rangle \times \sigma_{k_i} \xrightarrow{*} H_i \times \rho'_i \parallel H'_i \times \sigma'_i \dashv$$

for $1 \leq i \leq n$, where $\rho'_i \neq \mathbf{1}$ and $H_i = H'_i = \langle \rangle$ by Lemma 25. This implies

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

by Lemma 26. Let $I' = I \setminus J$ and $J' = J \setminus I$. We can reduce $\langle \rangle \times \rho \parallel \langle \rangle \times \sigma$ only as follows:

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

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

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

$$\langle \rangle \times \rho \parallel \langle \rangle \times \sigma \longrightarrow \langle \rangle \times \bar{\alpha}_k \cdot \rho_k \parallel \langle \rangle \times \sigma \not\rightarrow$$

Otherwise $I \subseteq J$ and $\rho_k \not\stackrel{R}{\sim} \sigma_k$ for some $k \in I$. By reasoning as above we have

$$\langle \rangle \times \rho_k \parallel \langle \rangle \times \sigma_k \xrightarrow{*} \langle \rangle \times \rho' \parallel \langle \rangle \times \sigma' \not\rightarrow$$

$$\text{and } \circ \times \rho_k \parallel \sum_{j \in J \setminus \{k\}} \alpha_j \cdot \sigma_j \times \sigma_k \xrightarrow{*} \circ \times \rho' \parallel \sum_{j \in J \setminus \{k\}} \alpha_j \cdot \sigma_j \times \sigma'$$

which imply

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

In both the cases we conclude that $\rho \not\stackrel{R}{\sim} \sigma$.

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

Lemma 27 suggests that $\dashv\!\!\dashv^R$ can be coinductively defined, or equivalently that $\dashv\!\!\dashv^R = \bigcap_n \dashv\!\!\dashv_n^R$ where $\dashv\!\!\dashv_0^R$ is the trivial relation $\mathbf{rsC} \times \mathbf{rsC}$, and for all $n > 0$, $\dashv\!\!\dashv_n^R \sigma$ and we have $\rho \dashv\!\!\dashv_n^R \sigma$ if one of the following holds:

1. $\rho = \sum_{i \in I} \alpha_i \cdot \rho_i$, $\sigma = \sum_{j \in J} \bar{\alpha}_j \cdot \sigma_j$ and $\exists k \in I \cap J$. $\rho_k \dashv\!\!\dashv_{n-1}^R \sigma_k$;
2. $\rho = \bigoplus_{i \in I} \bar{\alpha}_i \cdot \rho_i$, $\sigma = \sum_{j \in J} \alpha_j \cdot \sigma_j$, $I \subseteq J$ and $\forall k \in I$. $\rho_k \dashv\!\!\dashv_{n-1}^R \sigma_k$;
3. $\rho = \sum_{i \in I} \bar{\alpha}_i \cdot \rho_i$, $\sigma = \bigoplus_{j \in J} \alpha_j \cdot \sigma_j$, $I \supseteq J$ and $\forall k \in J$. $\rho_k \dashv\!\!\dashv_{n-1}^R \sigma_k$.

We write:

1. $\models^R \Gamma$ if for all $\rho' \dashv\!\!\dashv^R \sigma' \in \Gamma$ we have $\rho' \dashv\!\!\dashv^R \sigma'$
2. $\Gamma \models^R \rho \dashv\!\!\dashv \sigma$ if $\models^R \Gamma$ implies $\rho \dashv\!\!\dashv \sigma$.

We also write $\Gamma \models_n^R \rho \dashv\!\!\dashv \sigma$ if $\dashv\!\!\dashv^R$ is replaced by $\dashv\!\!\dashv_n^R$ in the above.

Observing that $\dashv\!\!\dashv_{n+1}^R \subseteq \dashv\!\!\dashv_n^R$, we have that $\models_{n+1}^R \Gamma$ implies $\models_n^R \Gamma$. Also it is immediate to verify that the following holds:

Fact 28. If $\Gamma \models_n^R \rho \dashv\!\!\dashv \sigma$ for all n , then $\Gamma \models^R \rho \dashv\!\!\dashv \sigma$.

We can now prove Theorem 2, which can be restated as follows.

Theorem 2 (Retractable soundness and completeness).

$$\triangleright \rho \dashv\!\!\dashv \sigma \text{ iff } \models^R \rho \dashv\!\!\dashv \sigma$$

Proof. (\Rightarrow) For this direction we can actually prove a stronger statement, namely $\Gamma \triangleright \rho \dashv\!\!\dashv \sigma \Rightarrow \Gamma \models^R \rho \dashv\!\!\dashv \sigma$

By Fact 28 it suffices to prove that if $\Gamma \triangleright \rho \dashv\!\!\dashv \sigma$ then $\Gamma \models_n^R \rho \dashv\!\!\dashv \sigma$ for all n , which we establish by simultaneous induction over n and over the derivation \mathcal{D} of $\Gamma \triangleright \rho \dashv\!\!\dashv \sigma$.

If \mathcal{D} ends by either Ax or HYP then the thesis trivially holds. If \mathcal{D} ends by:

$$\frac{(+ \cdot +)}{\Gamma, \alpha \cdot \rho + \rho' \dashv\!\!\dashv \bar{\alpha} \cdot \sigma + \sigma' \triangleright \rho \dashv\!\!\dashv \sigma} \Gamma \triangleright \alpha \cdot \rho + \rho' \dashv\!\!\dashv \bar{\alpha} \cdot \sigma + \sigma'$$

then we have to show that $\models_n^R \Gamma$ implies $\alpha \cdot \rho + \rho' \dashv\!\!\dashv_n^R \bar{\alpha} \cdot \sigma + \sigma'$. By induction over n we know that $\Gamma \models_{n-1}^R \alpha \cdot \rho + \rho' \dashv\!\!\dashv \bar{\alpha} \cdot \sigma + \sigma'$; from this and the fact that $\models_n^R \Gamma$ implies $\models_{n-1}^R \Gamma$, we obtain that $\alpha \cdot \rho + \rho' \dashv\!\!\dashv_{n-1}^R \bar{\alpha} \cdot \sigma + \sigma'$, and hence that $\models_{n-1}^R \Gamma, (\alpha \cdot \rho + \rho' \dashv\!\!\dashv \bar{\alpha} \cdot \sigma + \sigma')$. By induction over \mathcal{D} it follows that $\rho \dashv\!\!\dashv_{n-1}^R \sigma$, which implies $\alpha \cdot \rho + \rho' \dashv\!\!\dashv_n^R \bar{\alpha} \cdot \sigma + \sigma'$ by definition of $\dashv\!\!\dashv_n^R$, as desired.

The cases in which \mathcal{D} ends by either $(\oplus \cdot +)$ or $(+ \cdot \oplus)$ are similar.

(\Leftarrow) By Theorem 1 each computation of **Prove**($\triangleright \rho \dashv\!\!\dashv \sigma$) always terminates. By Lemma 27 and Fact 24, $\rho \dashv\!\!\dashv^R \sigma$ implies that **Prove**($\triangleright \rho \dashv\!\!\dashv \sigma$) \neq **fail**, and hence $\triangleright \rho \dashv\!\!\dashv \sigma$. \square

We proceed now with the proof of Soundness and Completeness of system \triangleright with respect to the speculative compliance.

Speculative soundness and completeness. A lemma similar to Lemma 27 holds for speculative compliance as well.

Lemma 29. *We have $\rho \dashv^S \sigma$ if and only if 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 $\exists k \in I \cap J$. $\rho_k \dashv^S \sigma_k$;
3. $\rho = \bigoplus_{i \in I} \bar{\alpha}_i \cdot \rho_i$, $\sigma = \sum_{j \in J} \alpha_j \cdot \sigma_j$, $I \subseteq J$ and $\forall k \in I$. $\rho_k \dashv^S \sigma_k$;
4. $\rho = \sum_{i \in I} \bar{\alpha}_i \cdot \rho_i$, $\sigma = \bigoplus_{j \in J} \alpha_j \cdot \sigma_j$, $I \supseteq J$ and $\forall k \in J$. $\rho_k \dashv^S \sigma_k$.

Proof. The if part is immediate. We prove the only if part by contraposition and by cases on 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 $\rho_{k_i} \not\vdash^S \sigma_{k_i}$ for $1 \leq i \leq n$. Then we get $\rho_{k_i} \parallel \sigma_{k_i} \xrightarrow{*} \mathbf{C}_i^\rho \parallel \mathbf{C}_i^\sigma \not\vdash$ for $1 \leq i \leq n$, where $\mathbf{C}_i^\rho \neq \mathbf{C} \mid \alpha_1 @ \dots @ \alpha_n @ \mathbf{1}$.

This implies

$$\begin{aligned} & \sum_{i \in I} \alpha_i \cdot \rho_i \parallel \sum_{j \in J} \bar{\alpha}_j \cdot \sigma_j \\ \xrightarrow{*} & (\prod_{i \in I \cap J} \alpha_i @ \rho_i) \mid \sum_{i \in I \setminus J} \alpha_i \cdot \rho_i \parallel (\prod_{j \in I \cap J} \bar{\alpha}_j @ \sigma_j) \mid \sum_{j \in J \setminus I} \bar{\alpha}_j \cdot \sigma_j \end{aligned}$$

where $\prod_{i \in I} \mathbf{C}_i$ denotes the parallel composition of \mathbf{C}_i for each $i \in I$.

Terms $\sum_{i \in I \setminus J} \alpha_i \cdot \rho_i$ and $\sum_{j \in J \setminus I} \bar{\alpha}_j \cdot \sigma_j$ cannot interact with other terms. Instead, term $\prod_{i \in I \cap J} \alpha_i @ \rho_i$ can interact only with term $\prod_{j \in I \cap J} \bar{\alpha}_j @ \sigma_j$ and vice versa. The interaction follows the computations $\rho_{k_i} \parallel \sigma_{k_i} \xrightarrow{*} \mathbf{C}_i^\rho \parallel \mathbf{C}_i^\sigma \not\vdash$, with the added prefixes α_i and $\bar{\alpha}_i$. However, none of these computations produces a thread of the form $\alpha'_1 @ \dots @ \alpha'_n @ \mathbf{1}$, hence $\rho \not\vdash^S \sigma$.

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

$$\rho \parallel \sigma \longrightarrow \bar{\alpha}_k \cdot \rho_k \parallel \sigma \not\vdash$$

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

$$\begin{aligned} \rho_k \parallel \sigma_k & \xrightarrow{*} \mathbf{C}^\rho \parallel \mathbf{C}^\sigma \not\vdash \quad \text{and} \\ & \rho \parallel \sigma \\ & \longrightarrow \bar{\alpha}_k \cdot \rho_k \parallel \sigma \\ & \longrightarrow \bar{\alpha}_k @ \rho_k \parallel \alpha_k @ \sigma_k \mid \sum_{j \in J \setminus \{k\}} \alpha_j \cdot \sigma_j \\ & \xrightarrow{*} \bar{\alpha}_k @ \mathbf{C}^\rho \parallel \alpha_k @ \mathbf{C}^\sigma \mid \sum_{j \in J \setminus \{k\}} \alpha_j \cdot \sigma_j \\ & \not\vdash \end{aligned}$$

where $\alpha @ \mathbf{C}$ denotes $\prod_{i \in I} \alpha @ \mathbf{T}_i$ if $\mathbf{C} = \prod_{i \in I} \mathbf{T}_i$.

In both cases we conclude that $\rho \not\vdash^S \sigma$.

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

As with \dashv^R , we define the family of relations \dashv_n^S on the basis of Lemma 29, which are such that $\dashv^S = \bigcap_n \dashv_n^S$; similarly we define the respective notions $\Gamma \models^S \rho \dashv \sigma$ and $\Gamma \models^{S_n} \rho \dashv \sigma$. Speculative soundness and completeness can hence be restated as follows.

Theorem 3 (*Speculative soundness and completeness*).

$$\triangleright \rho \dashv \sigma \text{ iff } \models^S \rho \dashv \sigma$$

Proof. (\Rightarrow) This implication can be proved in the same way as Theorem 2.

(\Leftarrow) By Theorem 1 each computation of **Prove**($\triangleright \rho \dashv \sigma$) always terminates. By Lemma 27 and Fact 24, $\rho \dashv^R \sigma$ implies that **Prove**($\triangleright \rho \dashv \sigma$) \neq **fail**, and hence $\triangleright \rho \dashv \sigma$. \square

A.1.1. Proofs of Section 3.1

Before proving Proposition 7 of Section 3.1, we need a simple technical lemma and a fact.

Lemma 30. *Let $\rho \in \text{SC}$. Then either $\rho = \bar{a} \cdot \rho_1 \oplus \rho_2$, or $\rho = \alpha \cdot \rho'$, or $\rho = a \cdot \rho_1 + \rho_2$, or $\rho = \mathbf{1}$. Moreover, for any H,*

- i) $\bar{a}. \rho_1 \oplus \rho_2 \xrightarrow{\tau} \bar{a}. \rho_1$ if and only if $H \times \bar{a}. \rho_1 \oplus \rho_2 \xrightarrow{\tau} H \times \bar{a}. \rho_1$;
- ii) $\alpha. \rho' \xrightarrow{\alpha} \rho'$ if and only if $H \times \alpha. \rho' \xrightarrow{\alpha} H : \circ \times \rho'$;
- iii) $a. \rho_1 + \rho_2 \xrightarrow{a} \rho_1$ if and only if $H \times a. \rho_1 + \rho_2 \xrightarrow{a} H : \rho_2 \times \rho_1$
- iv) $\rho = \mathbf{1}$ if and only if $(H \times \rho \not\xrightarrow{\alpha} \text{ and } H \times \rho \not\xrightarrow{\tau})$.

Proof. Easy, by definition of session contract and by Definitions 4 and 23. \square

Fact 31. Let $\rho, \sigma \in \text{SC}$. $H_1 \times \rho \parallel H_2 \times \sigma \xrightarrow{f} H'_1 \times \rho' \parallel H'_2 \times \sigma'$ implies $\rho', \sigma' \in \text{SC}$

We can now proceed with the proof of Proposition 7.

Proof of Proposition 7. (i) The inclusion $\text{SC} \subseteq \text{rsC}$ holds by definition. Hence, given $\rho, \sigma \in \text{SC}$ we have $\rho, \sigma \in \text{rsC}$ and $H_1 \times \rho, H_2 \times \sigma \in \text{rsCH}$.

- (\Rightarrow) By induction on the length of the reduction sequence $\xrightarrow{*}_{\text{SC}}$, using Definitions 4 and 12 and Lemma 30 to check all the possible cases for the reductions of client/server pairs.
- (\Leftarrow) Using Fact 31 we first show that no pair of the form $H_1 \times \alpha. \rho_1 + \rho_2 \parallel H_2 \times \bar{\alpha}. \sigma_1 + \sigma_2$ can ever appear inside the reduction sequence $\xrightarrow{*}_f$. Then we proceed by induction on the length of the reduction sequence $\xrightarrow{*}_f$ using Definitions 4 and 12 and Lemma 30 to check all the possible cases for the reductions of client/server pairs with histories.

(ii) By induction on the length of the derivation. The base case is trivial. Let us consider the inductive case.

- (\Rightarrow) the actions performed by σ' and ρ' can be performed as well by $\alpha_1 @ \dots @ \alpha_n @ \rho'$ and $\bar{\alpha}_1 @ \dots @ \bar{\alpha}_n @ \sigma'$, with the only possible side effects of adding complementary actions to the prefixes, and of spawning further parallel threads. Since both the effects are compatible with the thesis, we are done.
- (\Leftarrow) if the action is a τ action, then the corresponding session contract can perform it as well. If it is a synchronization, by definition of the semantics it involves two threads with complementary prefixes $\alpha_1 @ \dots @ \alpha_n @ \rho'$ and $\bar{\alpha}_1 @ \dots @ \bar{\alpha}_n @ \sigma'$. By inductive hypothesis $\rho \parallel \sigma \xrightarrow{*}_{\text{SC}} \rho' \parallel \sigma'$, hence the thesis follows since ρ' and σ' can match the synchronization. \square

A.2. Proofs of Section 4

Proof of Proposition 11. Like \dashv^{R} , \dashv can be coinductively defined, since, by Corollary 4, $\dashv^{\text{R}} = \dashv^{\text{S}} = \dashv$. So $\dashv = \bigcap_n \dashv_n$ where $\dashv_n = \dashv_n^{\text{R}}$.

For the definition of \dashv_n see the discussion after Lemma 27.

We shall prove that

$$\forall n. [(\rho \dashv_n \sigma \text{ and } \bar{\sigma} \dashv_n \sigma') \text{ implies } \rho \dashv_n \sigma'] \quad (\text{A.1})$$

So, from $\rho \dashv \sigma$ and $\bar{\sigma} \dashv \sigma'$ we have that $\forall n. \rho \dashv_n \sigma$ and $\forall n. \bar{\sigma} \dashv_n \sigma'$ and hence, by (A.1) we can get $\forall n. \rho \dashv_n \sigma'$, that is $\rho \dashv \sigma'$.

We show now (A.1) by induction on n .

The base case is trivial. Let then $n > 0$ with $\rho \dashv_n \sigma$ and $\bar{\sigma} \dashv_n \sigma'$. We proceed by cases according to the possible shapes of ρ and σ in the definition of \dashv_n .

$\rho = \mathbf{1}$ Immediate.

$$\rho = \sum_{i \in I} \alpha_i. \rho_i, \sigma = \sum_{j \in J} \bar{\alpha}_j. \sigma_j \text{ and } \exists k \in I \cap J. \rho_k \dashv_{n-1} \sigma_k$$

We have then that $\bar{\sigma} = \bigoplus_{j \in J} \alpha_j. \bar{\sigma}_j$. So, by $\bar{\sigma} \dashv_n \sigma'$ and by definition of \dashv_n we have that $\sigma' = \sum_{h \in H} \bar{\alpha}_h. \sigma'_h$, $J \subseteq H$ and $\forall j \in J. \bar{\sigma}_j \dashv_{n-1} \sigma'_j$. By the induction hypothesis we can hence get that $\exists k \in (I \cap J) \subseteq (I \cap H)$ such that $\rho_k \dashv_{n-1} \sigma'_k$, that means, by definition, that $\rho \dashv_n \sigma'$.

$$\rho = \bigoplus_{i \in I} \bar{\alpha}_i. \rho_i, \sigma = \sum_{j \in J} \alpha_j. \sigma_j, I \subseteq J \text{ and } \forall k \in I. \rho_k \dashv_{n-1} \sigma_k$$

We have then that $\bar{\sigma} = \bigoplus_{j \in J} \bar{\alpha}_j. \bar{\sigma}_j$. So, by $\bar{\sigma} \dashv_n \sigma'$ and by definition of \dashv_n we have that $\sigma' = \sum_{h \in H} \alpha_h. \sigma'_h$, $J \subseteq H$ and $\forall j \in J. \bar{\sigma}_j \dashv_{n-1} \sigma'_j$. By the induction hypothesis we can hence get that $\forall k \in I \subseteq J \subseteq H. \rho_k \dashv_{n-1} \sigma'_k$, that means, by definition, that $\rho \dashv_n \sigma'$.

$$\rho = \sum_{i \in I} \bar{\alpha}_i. \rho_i, \sigma = \bigoplus_{j \in J} \alpha_j. \sigma_j, I \supseteq J \text{ and } \forall k \in J. \rho_k \dashv_{n-1} \sigma_k$$

We have then that $\bar{\sigma} = \sum_{j \in J} \bar{\alpha}_j. \bar{\sigma}_j$. So, by $\bar{\sigma} \dashv_n \sigma'$ and by definition of \dashv_n we have to take into account two cases.

$$\sigma' = \sum_{h \in H} \alpha_h \cdot \sigma'_h \text{ and } \exists k \in J \cap H. \sigma_k \dashv_{n-1} \sigma'_k$$

By the induction hypothesis we can get that $\exists k \in (J \cap H) \subseteq (I \cap H)$ such that $\rho_k \dashv_{n-1} \sigma'_k$, that means, by definition, that $\rho \dashv_n \sigma'$.

$$\sigma' = \bigoplus_{h \in H} \alpha_h \cdot \sigma'_h, J \supseteq H \text{ and } \forall h \in H. \sigma_h \dashv_{n-1} \sigma'_h$$

By the induction hypothesis we can get that $\forall h \in H \subseteq J \subseteq I. \rho_h \dashv_{n-1} \sigma'_h$, that means, by definition, that $\rho \dashv_n \sigma'$. \square

Proof of Theorem 12. (i) (\Rightarrow) By contraposition, assume that $\bar{\sigma} \not\vdash \sigma'$. Since $\bar{\sigma} \dashv \sigma$ by Proposition 10, then by definition of \preceq_s we have $\sigma \not\preceq_s \sigma'$.

(\Leftarrow) Let $\bar{\sigma} \dashv \sigma'$. If $\rho \dashv \sigma$, then from $\bar{\sigma} \dashv \sigma'$, we get $\rho \dashv \sigma'$ by Proposition 11, and therefore $\sigma \preceq_s \sigma'$ by definition.

(ii) (\Rightarrow) By contraposition, assume that $\sigma' \not\vdash \bar{\sigma}$. Since $\sigma \dashv \bar{\sigma}$ by Proposition 10, then by definition of \preceq_c we have $\sigma \not\preceq_c \sigma'$.

(\Leftarrow) Let $\sigma' \dashv \bar{\sigma}$. If $\sigma \dashv \rho$, then from $\sigma' \dashv \bar{\sigma}$, we get $\sigma' \dashv \rho$ by Proposition 11, and therefore $\sigma \preceq_c \sigma'$ by definition.

(iii) From Item (i) we have $\sigma \preceq_s \sigma'$ iff $\bar{\sigma} \dashv \sigma'$. From Item (ii) we have $\bar{\sigma}' \preceq_c \bar{\sigma}$ iff $\bar{\sigma} \dashv \bar{\sigma}'$. The thesis follows since $\bar{\sigma}' = \bar{\sigma}$.

(iv) From Items (i) and (ii) and decidability of $\bar{\sigma} \dashv \sigma'$. \square

Proof sketch of Proposition 16. For the proof of the antisymmetric property we cannot rely directly on the definition of \preceq_s , since from $\sigma \preceq_s \sigma'$ and $\sigma' \preceq_s \sigma$ we can only infer that σ and σ' have the same set of clients. This does not trivially imply that $\sigma = \sigma'$.

We can proceed, instead, roughly as follows. Let $\sigma, \sigma' \in \mathbf{rsC}$ be such that $\sigma \preceq_s \sigma'$ and $\sigma' \preceq_s \sigma$. By completeness of \blacktriangleright we get $\blacktriangleright \sigma \ll \sigma'$ and $\blacktriangleright \sigma' \ll \sigma$. By having such derivations, we can infer that in each of them no rule $(\oplus \cdot + \cdot \preceq_s)$ can be present. Moreover, in each application of rule $(+ \cdot + \cdot \preceq_s)$ or rule $(\oplus \cdot \oplus \cdot \preceq_s)$, we have necessarily that $J = \emptyset$. Out of that we can infer $\sigma = \sigma'$. \square

A.3. Proofs of Section 5

Proof sketch of Lemma 19. (\Rightarrow) A derivation of $\triangleright \rho \dashv \sigma$ containing no occurrence of rule (HYP) can be turned in a derivation of $\triangleright_{\infty} \rho \dashv \sigma$ by erasing all the environments.

Otherwise, we observe that if an occurrence $\Gamma', \rho' \dashv \sigma' \triangleright \rho' \dashv \sigma'$ of rule (HYP) occurs in a given derivation of $\triangleright \rho \dashv \sigma$, then the expression $\rho' \dashv \sigma'$ is discharged by a rule (namely, it occurs in the environment of the premise(s) but not in the one of the conclusion; we dub the occurrence of the rule a *discharging occurrence*) present in the path from the root to (HYP). Hence a derivation of $\triangleright_{\infty} \rho \dashv \sigma$ can be built by performing infinitely many times the replacement of all the occurrences of rule (HYP) with the subderivations having their corresponding discharging occurrences as roots. The infinite tree obtained by applying such a procedure and then erasing all the environments is a derivation of $\triangleright_{\infty} \rho \dashv \sigma$.

(\Leftarrow) Take an infinite derivation of $\triangleright_{\infty} \rho \dashv \sigma$ (for a finite one the result is immediate). Since all contracts in the premises of a rule are subexpressions of those in the conclusion, and contracts are regular expressions, then in an infinite branch there must be at least a judgment occurring infinitely many times. It is not difficult hence to get a derivation of $\triangleright \rho \dashv \sigma$ by replacing all the expressions of the form $\rho' \dashv \sigma'$ occurring for the second time (starting from the root) in a branch, by rules (HYP), and by inserting environments accordingly. \square

Proof of Lemma 20. First observe that:

- (1) The recursive calls in lines -1-, -2-, 19- do leave unaltered the arguments A and F; so do those in lines -7-, -11-, -15-, -20- when the results A_0 and F_0 of the previous call are equal, respectively, to its first and second argument. In the depth-first search for a proof in \triangleright_{∞} all these calls do correspond to trying and finding a subproof for a further pair of subterms of the initial ρ and σ once the search for a proof for the current pair has been (successfully or unsuccessfully) completed. This means that the overall number of such calls during the execution of the algorithm cannot exceed the number of pairs of subterms of the initial ρ and σ , so it is $\mathcal{O}(n^2)$;
- (2) Also the recursive calls in lines -16-, -17- do leave unaltered the arguments A and F, but there cannot be more than two consecutive calls of this sort. This implies that the overall number of such calls cannot exceed two times the number of all the other calls;
- (3) in all the other recursive calls the cardinality of $A \cup F$ strictly increases.

Hence, by the above and by the termination conditions of the algorithm, the overall number of calls cannot exceed $(n^2 + n^2) + (2 \times (n^2 + n^2))$.

We now observe that both the element-insertion and the membership-check operations for the sets A and F can be performed in $\mathcal{O}(1)$. It is in fact enough to build, before running the algorithm, two boolean vectors indexed by all the pairs of subterms of the initial ρ and σ . Building such vectors takes $\mathcal{O}(n^2)$ time.

Before concluding, we also observe – for what concerns the membership checks – that without loss of generality we can assume the initial ρ and σ to be length-minimal (where a contract γ is *length-minimal* whenever, for any contract δ such that γ and δ represent the same regular tree, the length of γ is not greater than that of δ). It can hence be checked that length-minimality is maintained by the argument of any recursive call but those in lines -16- and -17-, which are however performed after the membership check.

The complexity of the algorithm is hence $\mathcal{O}(n^2)$. \square

A.4. Proofs of Section 6

Proof sketch of Theorem 23. The definition of the relation R is quite convoluted, hence we will not spell out it, but we present the main ideas below. Essentially, one starts by adding to R all pairs of the form $(\langle \rangle \times \rho \parallel \langle \rangle \times \sigma, \llbracket \rho \parallel \sigma \rrbracket)$. This takes into account client/server pairs that have not performed any action. It is easy to check by induction on the definition of the translation function that the condition on bars is satisfied.

In order to satisfy the other conditions, we need to add pairs obtained by performing actions. On the side of transformed contracts, performing actions just means underlining prefixes (there is no need to explicitly unfold recursion since we are using an equi-recursive interpretation of recursion), with the constraints that a prefix of an underlined prefix is underlined, and that in choices at most one alternative can be underlined.

A given transformed contract $X \parallel X'$ with underlined prefixes is in the relation R with multiple retractable client/server pairs, each of the form $H \times \rho \parallel H' \times \rho'$. The possible contracts with history $H \times \rho$ (resp. $H' \times \rho'$) are derived from X (resp. X') by applying the following procedure recursively to the top element of X , until all the underlined prefixes have been removed.

- if $X = \underline{\alpha}.X_1$ then a \circ is pushed on the history, and the procedure applied to X_1 ;
- if $X = \underline{\tau}.X_1 + \sigma$ then the procedure is applied to X_1 ;
- if $X = \underline{\alpha}.X_1 + \sigma'$ then the procedure is applied to X_1 . Note that here $\sigma' = \sum_{i \in I} \pi_i \cdot \sigma_i$, where each π_i is an action α . In fact, the conditions on underlines described above forbid to have underlined actions, while the form of terms produced by the translation function forbids to have τ s. Nevertheless, there are multiple possibilities for the term to push in the history, one for each subset of the alternatives in I . Formally, we have a possibility for each term of the form $\sum_{i \in J} \pi_i \cdot \sigma_i$ with $J \subseteq I$, where the choice on no elements denotes \circ . Branches which are removed (that is, in $I \setminus J$) correspond to alternatives that have already been tried and discarded. The actual term to push is the only contract $\hat{\sigma}$ such that $\llbracket \hat{\sigma} \rrbracket = \sum_{i \in J} \pi_i \cdot \sigma_i$. Essentially, it is obtained by translating back τ -prefixed choices into internal choices.

It is easy to check that the resulting relation R is a simulation. \square

References

- [1] S. Carpineti, G. Castagna, C. Laneve, L. Padovani, A formal account of contracts for web services, in: WS-FM, in: LNCS, vol. 4184, Springer, 2006, pp. 148–162.
- [2] C. Laneve, L. Padovani, The must preorder revisited: an algebraic theory for web services contracts, in: CONCUR, in: LNCS, vol. 4703, Springer, 2007, pp. 212–225.
- [3] G. Castagna, N. Gesbert, L. Padovani, A theory of contracts for web services, ACM Trans. Program. Lang. Syst. 31 (5) (2009) 19:1–19:61.
- [4] K. Honda, V.T. Vasconcelos, M. Kubo, Language primitives and type disciplines for structured communication-based programming, in: ESOP, in: LNCS, vol. 1381, Springer, 1998, pp. 22–138.
- [5] H. Hüttel, I. Lanese, V.T. Vasconcelos, L. Caires, M. Carbone, P.-M. Deniérou, D. Mostrous, L. Padovani, A. Ravara, E. Tuosto, H.T. Vieira, G. Zavattaro, Foundations of session types and behavioural contracts, ACM Comput. Surv. 49 (1) (2016) 3:1–3:36.
- [6] M. Carbone, K. Honda, N. Yoshida, Structured interactional exceptions in session types, in: CONCUR, in: LNCS, vol. 5201, Springer, 2008, pp. 402–417.
- [7] M. Bartoletti, T. Cimoli, M. Murgia, A. Podda, L. Pompianu, Compliance and subtyping in timed session types, in: FORTE, in: LNCS, vol. 9039, Springer, 2015, pp. 161–177.
- [8] P. Prabhu, G. Ramalingam, K. Vaswani, Safe programmable speculative parallelism, in: PLDI, ACM, 2010, pp. 50–61.
- [9] A. Avizienis, J. Laprie, B. Randell, C.E. Landwehr, Basic concepts and taxonomy of dependable and secure computing, IEEE Trans. Dependable Secure Comput. 1 (1) (2004) 11–33.
- [10] C.D. Carothers, K.S. Perumalla, R. Fujimoto, Efficient optimistic parallel simulations using reverse computation, ACM Trans. Model. Comput. Simul. 9 (3) (1999) 224–253.
- [11] C.G. Quiñones, C. Madriles, J. Sánchez, P. Marcuello, A. González, D.M. Tullsen, Mitosis compiler: an infrastructure for speculative threading based on pre-computation slices, in: PLDI, ACM, 2005, pp. 269–279.
- [12] M. Dalla Preda, M. Gabbriellini, I. Lanese, J. Mauro, G. Zavattaro, Graceful interruption of request-response service interactions, in: ICSOC, in: LNCS, vol. 7084, Springer, 2011, pp. 590–600.
- [13] F. Barbanera, U. de'Liguoro, Sub-behaviour relations for session-based client/server systems, Math. Struct. Comput. Sci. 25 (6) (2015) 1339–1381.
- [14] G.T. Bernardi, M. Hennessy, Modelling session types using contracts, Math. Struct. Comput. Sci. 26 (3) (2016) 510–560.
- [15] C.A.R. Hoare, Communicating Sequential Processes, Prentice-Hall, 1985.
- [16] I.C.C. Phillips, I. Ulidowski, Reversing algebraic process calculi, J. Log. Algebraic Methods Program. 73 (1–2) (2007) 70–96.
- [17] F. Barbanera, I. Lanese, U. de'Liguoro, Retractable and speculative contracts, in: COORDINATION, in: LNCS, vol. 10319, Springer, 2017, pp. 119–137.
- [18] F. Barbanera, M. Dezani-Ciancaglini, I. Lanese, U. de'Liguoro, Retractable contracts, in: PLACES 2015, in: EPTCS, vol. 203, Open Publishing Association, 2016, pp. 61–72.
- [19] J.C.M. Baeten, J.A. Bergstra, J.W. Klop, W.P. Weijland, Term-rewriting systems with rule priorities, Theor. Comput. Sci. 67 (2–3) (1989) 283–301.
- [20] A. Laville, Comparison of priority rules in pattern matching and term rewriting, J. Symb. Comput. 11 (1991) 321–347.

- [21] M.R. Mousavi, I.C.C. Phillips, M.A. Reniers, I. Ulidowski, The meaning of ordered SOS, in: FSTTCS, in: LNCS, vol. 4337, Springer, 2006, pp. 333–344.
- [22] M.R. Mousavi, I. Phillips, M.A. Reniers, I. Ulidowski, Semantics and expressiveness of ordered SOS, *Inf. Comput.* 207 (2) (2009) 85–119.
- [23] V. Danos, J. Krivine, Reversible communicating systems, in: CONCUR, in: LNCS, vol. 3170, Springer, 2004, pp. 292–307.
- [24] I. Lanese, C.A. Mezzina, J. Stefani, Reversibility in the higher-order π -calculus, *Theor. Comput. Sci.* 625 (2016) 25–84.
- [25] I. Lanese, C.A. Mezzina, F. Tiezzi, Causal-consistent reversibility, *Bulletin of the EATCS* 114.
- [26] F. Barbanera, U. de'Liguoro, Session types for orchestrated interactions, in: ICE@DisCoTec, in: EPTCS, vol. 261, 2017, pp. 17–36.
- [27] G. Boudol, I. Castellani, M. Hennessy, A. Kiehn, Observing localities, *Theor. Comput. Sci.* 114 (1) (1993) 31–61.
- [28] L. Zhang, S. Malik, The quest for efficient boolean satisfiability solvers, in: CAV, in: LNCS, vol. 2404, Springer, 2002, pp. 17–36.
- [29] L. Xu, F. Hutter, H.H. Hoos, K. Leyton-Brown, Satzilla: portfolio-based algorithm selection for SAT, *J. Artif. Intell. Res.* 32 (2008) 565–606.
- [30] L. Padovani, Contract-based discovery of web services modulo simple orchestrators, *Theor. Comput. Sci.* 411 (2010) 3328–3347.
- [31] V. Gapeyev, M.Y. Levin, B.C. Pierce, Recursive subtyping revealed, *J. Funct. Program.* 12 (6) (2002) 511–548.
- [32] F. Barbanera, U. de'Liguoro, Two notions of sub-behaviour for session-based client/server systems, in: PPDP, ACM Press, 2010, pp. 155–164.
- [33] F. Barbanera, M. Dezani-Ciancaglini, U. de'Liguoro, Compliance for reversible client/server interactions, in: BEAT, in: EPTCS, vol. 162, 2014, pp. 35–42.
- [34] F. Barbanera, M. Dezani-Ciancaglini, U. de'Liguoro, Reversible client/server interactions, *Form. Asp. Comput.* 28 (4) (2016) 697–722.
- [35] F. Tiezzi, N. Yoshida, Towards reversible sessions, in: PLACES, in: EPTCS, vol. 155, 2014, pp. 17–24.
- [36] F. Tiezzi, N. Yoshida, Reversible session-based pi-calculus, *J. Log. Algebraic Methods Program.* 84 (5) (2015) 684–707.
- [37] C.A. Mezzina, J.A. Pérez, Reversibility in session-based concurrency: a fresh look, *J. Log. Algebraic Methods Program.* 90 (2017) 2–30.
- [38] E. de Vries, V. Koutavas, M. Hennessy, Communicating transactions (extended abstract), in: CONCUR, in: LNCS, vol. 6269, Springer, 2010, pp. 569–583.
- [39] F. Barbanera, U. de'Liguoro, Retractability, games and orchestrators for session contracts, *Log. Methods Comput. Sci.* 13 (3) (2017).
- [40] K. Honda, N. Yoshida, M. Carbone, Multiparty asynchronous session types, in: POPL, ACM Press, 2008, pp. 273–284.
- [41] C.A. Mezzina, J.A. Pérez, Causally consistent reversible choreographies: a monitors-as-memories approach, in: PPDP, ACM, 2017, pp. 127–138.
- [42] M. Dezani-Ciancaglini, P. Giannini, Reversible multiparty sessions with checkpoints, in: EXPRESS/SOS, in: EPTCS, vol. 222, 2016, pp. 60–74.
- [43] I. Castellani, M. Dezani-Ciancaglini, P. Giannini, Concurrent reversible sessions, in: CONCUR, in: LIPIcs, vol. 85, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017, pp. 30:1–30:17.
- [44] G.T. Bernardi, M. Hennessy, Using higher-order contracts to model session types, *Log. Methods Comput. Sci.* 12 (2) (2016).
- [45] I. Lanese, M. Lienhardt, C.A. Mezzina, A. Schmitt, J. Stefani, Concurrent flexible reversibility, in: ESOP, in: LNCS, vol. 7792, Springer, 2013, pp. 370–390.
- [46] A. García, C. Laneve, Deadlock detection of Java bytecode, CoRR abs/1709.04152, arXiv:1709.04152, <http://arxiv.org/abs/1709.04152>.
- [47] K. Helsgaun, Cback: a simple tool for backtrack programming in C, *Softw. Pract. Exp.* 25 (8) (1995) 905–934.