A theory of retractable and speculative contracts

Franco Barbanera, Ivan Lanese, Ugo De'Liguoro

► To cite this version:

Franco Barbanera, Ivan Lanese, Ugo De'Liguoro. A theory of retractable and speculative contracts. Science of Computer Programming, 2018, 167, pp.25 - 50. 10.1016/j.scico.2018.06.005 . hal-01912858

HAL Id: hal-01912858

<https://hal.inria.fr/hal-01912858>

Submitted on 6 Nov 2018

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

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

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.*

^b*Dipartimento di Informatica - Scienza e Ingegneria, University of Bologna/INRIA*

^c*Dipartimento di Informatica, University of Torino.*

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.

Keywords: Behavioral contracts, backtracking, speculative execution, compliance, reversible computing

1. Introduction

Binary behavioral contracts [1, 2, 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-

*Corresponding author

Email addresses: barba@di.unica.it (Franco Barbanera), ivan.lanese@gmail.com (Ivan Lanese), ugo.deliguoro@unito.it (Ugo de'Liguoro)

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 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 **rsC** of *retractable/speculative contracts* is defined as the set of the *closed* expressions generated by the following grammar,

σ, ρ	$:=$	$\mathbf{1}$	SUCCESS
		$\sum_{i \in I} a_i. \sigma_i$	EXTERNAL INPUT CHOICE
		$\sum_{i \in I} \bar{a}_i. \sigma_i$	EXTERNAL OUTPUT CHOICE
		$\oplus_{i \in I} a_i. \sigma_i$	INTERNAL INPUT CHOICE
		$\oplus_{i \in I} \bar{a}_i. \sigma_i$	INTERNAL OUTPUT CHOICE
		x	VARIABLE
		$\text{rec } x. \sigma$	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 $\text{rec } x. \sigma$. Recursion in **rsC** is guarded and hence contractive in the usual sense. We take an equi-recursive view of recursion by equating $\text{rec } x. \sigma$ with $\sigma[\text{rec } x. \sigma / x]$.

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 $a. \sigma$ one first performs a and then

continues as specified by σ . In the syntax above, branches $a_i.\sigma_i$ and $\bar{a}_i.\sigma_i$ 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 $\bar{\alpha} = \bar{a}$ if $\alpha = a$, and $\bar{\alpha} = a$ 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 $H \ltimes \sigma$. 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 $H ::= \langle \rangle \mid H:\sigma$, 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:

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

We write just $\sigma_1 : \dots : \sigma_k$ for the stack $(\dots(\langle \rangle : \sigma_1) : \dots) : \sigma_k$. Moreover, we simply write σ for $\langle \rangle : \sigma$ 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).

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

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:

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

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

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

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

Figure 1: Rollback

shape of the server. If the server is, for instance, $\bar{a}.\bar{c} + \bar{b}.\bar{e}$ (see the first reduction sequence in Figure 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 Figure 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 $\rho \neq 1$ 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

$$\begin{aligned}
& \langle \rangle \ltimes \overline{\text{QoSday}.\text{priceMed.ok} + \text{priceLow.ok}} \parallel \langle \rangle \ltimes \sum_{\text{QoS}} \overline{\text{QoS.price}_{\text{QoS}.ok}} \quad (1) \\
\longrightarrow & \langle \rangle : \overline{\text{QoSnight.priceLow.ok}} \ltimes \text{priceMed.ok} + \text{priceLow.ok} \parallel \langle \rangle : \frac{\sum_{\text{QoS} \neq \text{QoSday}} \overline{\text{QoS.price}_{\text{QoS}.ok}}}{\text{priceHigh.ok}} \quad (2) \\
\longrightarrow & \langle \rangle \ltimes \overline{\text{QoSnight.priceLow.ok}} \parallel \langle \rangle \ltimes \sum_{\text{QoS} \neq \text{QoSday}} \overline{\text{QoS.price}_{\text{QoS}.ok}} \quad (3) \\
\longrightarrow & \langle \rangle : \circ \ltimes \text{priceLow.ok} \parallel \langle \rangle : \frac{\sum_{\text{QoS} \neq \text{QoSday}, \text{QoSnight}} \overline{\text{QoS.price}_{\text{QoS}.ok}}}{\text{priceLow.ok}} \quad (4) \\
\longrightarrow & \langle \rangle : \circ : \circ \ltimes \overline{\text{ok}} \parallel \langle \rangle : \frac{\sum_{\text{QoS} \neq \text{QoSday}, \text{QoSnight}} \overline{\text{QoS.price}_{\text{QoS}.ok} : \circ}}{\text{ok}} \quad (5) \\
\longrightarrow & \langle \rangle : \circ : \circ : \circ \ltimes \mathbf{1} \parallel \langle \rangle : \frac{\sum_{\text{QoS} \neq \text{QoSday}, \text{QoSnight}} \overline{\text{QoS.price}_{\text{QoS}.ok} : \circ : \circ}}{\mathbf{1}} \quad (6)
\end{aligned}$$

Figure 2: An example of retractable interaction

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:

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

Notice that the contract does not specify which alternative the client prefers: this aspect of the client behavior is abstracted away¹. A sample server is:

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

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

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

In particular, the interaction starts from a rollbackable interaction point (1) and proceeds as follows: (2) synchronization on QoSday and memorization of the

¹As 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.

discarded **QoSnight** branch; (3) rollback, since no choice has been offered for the requested **priceHigh**; (4) synchronization on **QoSnight** and memorization of $\circ$ (no further alternative branch is available); (5) synchronization on **priceLow**; (6) synchronization on **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.

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

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. $||$, define the relation $\longrightarrow$ over pairs of contracts with threads. In the LTS below, $?T$ denotes either the thread T or nothing. Hence, $\beta, ?T$ and $C | ?T$ are respectively β and C if $?T$ is nothing, and β, T and $C | T$ otherwise. Also, the duality operator extends from actions to sequences: $\overline{\alpha\beta} = \overline{\alpha}\overline{\beta}$.

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

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:

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

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

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

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

Figure 3: An example of speculative interaction

A sample computation proceeds as described in Figure 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 Figure 3 is due to rule (*comm*), since

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

and

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

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

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

whereas, on the server side,

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

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 **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 $\longrightarrow$, and $\not\rightarrow$ to specify that no $\longrightarrow$ reduction exists.

Definition 8 (Retractable Compliance Relation $\dashv^R$).

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

Figure 4: System $\triangleright$

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

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 $\dashv^{\mathbf{S}}$).

The relation $\dashv^{\mathbf{S}}$ on contracts is defined by:

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

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 $\Gamma \triangleright \rho \multimap \sigma$, where the environment Γ is a finite set of expressions of the form $\delta \multimap \gamma$, with $\rho, \sigma, \delta, \gamma \in \mathbf{rsC}$. Axioms and rules are defined in Figure 4.

Intuitively, the symbol $\multimap$ is used as syntactical counterpart for both the relations $\dashv^{\mathbf{R}}$ and $\dashv^{\mathbf{S}}$ 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

Prove($\Gamma \triangleright \rho \multimap \sigma$) =

if $\rho = 1$ then $\frac{}{\Gamma \triangleright 1 \multimap \sigma} (\text{Ax})$

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

else if $\rho = \sum_{i \in I} \alpha_i \cdot \rho_i$ and $\sigma = \sum_{j \in J} \bar{\alpha}_j \cdot \sigma_j$

and exists $k \in I \cap J$ s.t. $\mathcal{D} = \text{Prove}(\Gamma, \rho \multimap \sigma \triangleright \rho_k \multimap \sigma_k) \neq \text{fail}$

then $\frac{\mathcal{D}}{\Gamma \triangleright \rho \multimap \sigma} (+ \cdot +)$ else fail

else if $\rho = \bigoplus_{i \in I} \bar{\alpha}_i \cdot \rho_i$ and $\sigma = \sum_{j \in I \cup J} \alpha_j \cdot \sigma_j$

and for all $k \in I$ $\mathcal{D}_k = \text{Prove}(\Gamma, \rho \multimap \sigma \triangleright \rho_k \multimap \sigma_k) \neq \text{fail}$

then $\frac{(\forall k \in I) \mathcal{D}_k}{\Gamma \triangleright \rho \multimap \sigma} (\oplus \cdot +)$

else if $\rho = \sum_{j \in I \cup J} \bar{\alpha}_j \cdot \rho_j$ and $\sigma = \bigoplus_{i \in I} \alpha_i \cdot \sigma_i$

and for all $k \in I$ $\mathcal{D}_k = \text{Prove}(\Gamma, \rho \multimap \sigma \triangleright \rho_k \multimap \sigma_k) \neq \text{fail}$

then $\frac{(\forall k \in I) \mathcal{D}_k}{\Gamma \triangleright \rho \multimap \sigma} (+ \cdot \oplus)$ else fail

else fail

Figure 5: The procedure **Prove**.

case is taken into account by the rules we already have. Notice that rule $(+ \cdot +)$ 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 1. We write $\triangleright \rho \multimap \sigma$ instead of $\Gamma \triangleright \rho \multimap \sigma$ when Γ is empty.

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

The procedure **Prove** in Figure 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).

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

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

where $\gamma_1 = \overline{\text{inst.}(\overline{\text{DPLL.sol}} + \overline{\text{walksat.sol}})} \multimap \overline{\text{inst.} \sum_i \overline{\text{alg}_i.\text{sol}}}$

$\gamma_2 = \overline{\text{DPLL.sol}} + \overline{\text{walksat.sol}} \multimap \sum_i \overline{\text{alg}_i.\text{sol}}$

$\gamma_3 = \text{sol} \multimap \overline{\text{sol}}$

and where, for some i , $\text{alg}_i = \text{walksat}$.

Figure 6: A sample derivation in $\triangleright$

Theorem 3 (Speculative Soundness and Completeness).

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

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

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

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

Corollary 5 (Compliance Decidability). *The relation $\dashv\!\!\vdash$ 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 $\rho \neq 1$ 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 $\langle \rangle \ltimes 1 \parallel \langle \rangle \ltimes \sigma$) would end into $\langle \rangle \ltimes \circ \parallel \langle \rangle \ltimes \circ$. 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 $\langle \rangle \ltimes \circ \parallel \langle \rangle \ltimes \circ$. 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].

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, 2, 3]) can be seen as the subset of elements in **rsC** not containing external output choices and internal input choices, with the following operational semantics.

Definition 11 (Semantics of Session Contracts).

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

As done for **rsC** (see Section 2), we 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). 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).

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

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

Figure 7: System $\triangleright_{\text{SC}}$

Definition 13 (Compliance Relation $\dashv\!\!\!\vdash_{\text{SC}}$ for Session Contracts).

The relation $\dashv\!\!\!\vdash_{\text{SC}} \subset \text{SC} \times \text{SC}$ is defined by:

$$\rho \dashv\!\!\!\vdash_{\text{SC}} \sigma \text{ if, for each } \rho', \sigma' \text{ such that} \\
\rho \parallel \sigma \xrightarrow{*}_{\text{SC}} \rho' \parallel \sigma' \not\rightarrow_{\text{SC}} \text{ we have } \rho' = \mathbf{1}$$

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

Theorem 6. *Let $\rho, \sigma \in \text{SC}$. $\rho \dashv\!\!\!\vdash_{\text{SC}} \sigma$ iff $\triangleright_{\text{SC}} \rho \dot{\smile} \sigma$*

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 $\triangleright_{\text{H}} \rho \dot{\smile} \sigma$ iff $\rho \dashv\!\!\!\vdash_{\text{H}} \sigma$. The set SC can be looked at as the first-order restriction of HSC. It is easy to show that, for $\rho, \sigma \in \text{SC}$, $\triangleright_{\text{H}} \rho \dot{\smile} \sigma$ iff $\triangleright_{\text{SC}} \rho \dot{\smile} \sigma$.

It is also not difficult to show that, for $\rho, \sigma \in \text{SC}$, $\rho \dashv\!\!\!\vdash_{\text{SC}} \sigma$ iff $\rho \dashv\!\!\!\vdash_{\text{H}} \sigma$.

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 **rC**, which, in turn, are a strict subset of the contracts **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 **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 $\rightarrow_f$ and the speculative semantics $\rightarrow$ of pairs of contracts in SC are annotated versions of their semantics as session contracts.

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

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

$$ii) \rho \parallel \sigma \xrightarrow{*}_{\text{SC}} \rho' \parallel \sigma' \text{ iff there exist } n, \alpha_1, \dots, \alpha_n, \mathbf{C}_\rho \text{ and } \mathbf{C}_\sigma \text{ such that} \\
\rho \parallel \sigma \xrightarrow{*} \alpha_1 @ \dots @ \alpha_n @ \rho' \mid \mathbf{C}_\rho \parallel \bar{\alpha}_1 @ \dots @ \bar{\alpha}_n @ \sigma' \mid \mathbf{C}_\sigma$$

where $\longrightarrow_{\text{SC}}$ 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 $\dashv\!\parallel$ to elements in SC. Similarly, compliance for retractable session contracts is the restriction of $\dashv\!\parallel$ to elements in **rC**.

To prove such results, let $\dashv\!\parallel_{\text{rC}}$ 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 $\rho, \sigma \in \text{SC}$: $\triangleright_{\text{SC}} \rho \dashv\!\sigma$ iff $\triangleright \rho \dashv\!\sigma$
- ii) Let $\rho, \sigma \in \text{rC}$: $\triangleright_{\text{rC}} \rho \dashv\!\sigma$ iff $\triangleright \rho \dashv\!\sigma$

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 $\rho, \sigma \in \text{SC}$: $\rho \dashv\!\parallel_{\text{SC}} \sigma$ iff $\rho \dashv\!\parallel \sigma$
- ii) Let $\rho, \sigma \in \text{rC}$: $\rho \dashv\!\parallel_{\text{rC}} \sigma$ iff $\rho \dashv\!\parallel \sigma$

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 $\bar{\sigma}$ of an element σ of **rsC** is obtained, as for session contracts, by interchanging any name a with $\bar{a}$ 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 $\sigma \in \mathbf{rsCo}$. The syntactic dual $\bar{\sigma}$ of σ is defined by the following clauses:

$$\begin{array}{l} \bar{1} = 1 \quad \bar{x} = x \quad \overline{\mathbf{rec} \, x. \sigma} = \mathbf{rec} \, x. \bar{\sigma} \\ \overline{\sum_{i \in I} \alpha_i. \sigma_i} = \bigoplus_{i \in I} \bar{\alpha}_i. \bar{\sigma}_i \quad \overline{\bigoplus_{i \in I} \alpha_i. \sigma_i} = \sum_{i \in I} \bar{\alpha}_i. \bar{\sigma}_i \end{array}$$

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

From now on, in order to avoid too cumbersome definitions, any time an inductive definition on elements of $\mathbf{rsC}$ is provided, it will be tacitly assumed to be the restriction to $\mathbf{rsC}$ of the corresponding inductive definition on $\mathbf{rsCo}$. A first relevant property of duality is that any contract is compliant with its syntactic dual.

Proposition 10. *For any $\sigma \in \mathbf{rsC}$, $\sigma \dashv \bar{\sigma}$.*

Proof. Since $\bar{\sigma}$ is obtained from σ by exchanging each α with $\bar{\alpha}$ and $+$ with $\oplus$, it is easy to get a derivation of $\triangleright \bar{\sigma} \lhd \sigma$. 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 $\rho, \sigma, \sigma' \in \mathbf{rsC}$, $\rho \dashv \sigma$ and $\bar{\sigma} \dashv \sigma'$ imply $\rho \dashv \sigma'$*

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

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) $\sigma \preceq_s \sigma' \iff \bar{\sigma} \dashv \sigma'$
- ii) $\sigma \preceq_c \sigma' \iff \sigma' \dashv \bar{\sigma}$
- iii) $\sigma \preceq_s \sigma' \iff \bar{\sigma'} \preceq_c \bar{\sigma}$
- 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 $\rho \dashv \bar{\rho}$, and more precisely it is a least element in the set of the servers of ρ , that is,

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

Proof. Suppose that $\rho \dashv \sigma$ and take any contract δ such that $\delta \dashv \bar{\rho}$. Since $\bar{\bar{\rho}} = \rho$, by Proposition 11 we know that $\delta \dashv \sigma$; 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 Figure 8.

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

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

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

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

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

Theorem 15 (Soundness and Completeness of $\blacktriangleright$). $\blacktriangleright \sigma \ll \sigma' \iff \sigma \preceq_s \sigma'$

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

Figure 8: The formal system $\blacktriangleright$

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

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

System $\blacktriangleright$ 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 $\rho \Vdash \sigma$ 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 Figure 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 $\mathbf{SC}$. Moreover, the restriction of $\preceq_s$ to $\mathbf{rC}$ provides a suitable notion of subcontract for $\mathbf{rC}$ (which has never been studied before).

Figure 9: Subcontract preorder: a sample

5. Complexity Issues

The decision algorithm **Prove** for compliance of Figure 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.

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

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

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 **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 **Prove**($\Gamma \triangleright \rho_k \dashv \sigma_k$), with $k \neq 0$ and $\rho_k \dashv \sigma_k \notin \Gamma$, produces two immediate calls: **Prove**($\Gamma_1 \triangleright \rho_k \dashv \sigma_{k-1}$) and **Prove**($\Gamma_1 \triangleright \rho_{k-1} \dashv \sigma_k$). These latter two calls, produce, besides others, a call of **Prove**($\Gamma_2 \triangleright \rho_{k-1} \dashv \sigma_{k-1}$) and one of **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**”).

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

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 **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.

$\frac{(\text{Ax}_\infty)}{\triangleright \mathbf{1} \multimap \sigma}$	$\frac{(\cdot + \cdot +_\infty)}{\frac{\rho \multimap \sigma}{\alpha.\rho + \rho' \multimap \bar{\alpha}.\sigma + \sigma'}}$
$\frac{(\oplus \cdot +_\infty)}{(\forall h \in I) \quad \rho_h \multimap \sigma_h}$	$\frac{(\cdot + \cdot \oplus_\infty)}{(\forall h \in I) \quad \rho_h \multimap \sigma_h}$
$\frac{\oplus_{i \in I} \bar{\alpha}_i.\rho_i \multimap \sum_{j \in I \cup J} \alpha_j.\sigma_j}{\sum_{j \in I \cup J} \bar{\alpha}_j.\rho_j \multimap \oplus_{i \in I} \alpha_i.\sigma_i}$	$\frac{\sum_{j \in I \cup J} \bar{\alpha}_j.\rho_j \multimap \oplus_{i \in I} \alpha_i.\sigma_i}{\sum_{j \in I \cup J} \bar{\alpha}_j.\rho_j \multimap \oplus_{i \in I} \alpha_i.\sigma_i}$

Figure 10: The non-well founded system $\triangleright_\infty$

A quadratic decision algorithm.

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

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

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

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

In Figure 11 we present a decision algorithm **Decide**_{||}, based on the procedures **P** and **P**⁺. Whereas the procedure **Prove** returns, if any, a derivation in $\triangleright$, **Decide**_{||} just checks for the existence of a derivation in $\triangleright_\infty$ (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 $(\cdot + \cdot \oplus_\infty)$ are universal and nodes corresponding to $(\cdot + \cdot +_\infty)$ are existential; **P**(A, F, L, b) attempts to prove all statements in its goal list L, while **P**⁺(A, F, L, b) succeeds if at least one goal in L is satisfiable.

The procedure **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 $\triangleright_\infty$ using a *depth-first* technique. **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 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 $(\cdot + \cdot +_\infty)$ 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.

$$\text{Decide}_{\perp\parallel}(\rho \smile \sigma) = \text{let } (A, F, b) = \mathbf{P}(\emptyset, \emptyset, [\rho \smile \sigma], \mathbf{ok}) \\ \text{in } b = \mathbf{ok}$$

where

$$\mathbf{P}(A, F, [\], b) = (A, F, b)$$
$$\mathbf{P}(A, F, (\rho \smile \sigma) : \mathbf{x}s, \mathbf{b}) =$$

```

-1-   if  $\rho = 1$  then  $P(A, F, xs, b)$ 
-2-   else if  $\rho \sqcap \sigma \in A$  then  $P(A, F, xs, b)$ 
-3-   else if  $\rho \sqcap \sigma \in F$  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  $(A_0, F_0, b_0) = P^+(A \cup \{\rho \sqcap \sigma\}, F, [\rho_{i_1} \sqcap \sigma_{i_1} \dots \rho_{i_n} \sqcap \sigma_{i_n}], b)$ 
-6-           in if  $b_0 = fail$  then  $(A_0, F_0, fail)$ 
-7-           else  $P(A_0, F_0, xs, b_0)$ 
-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  $(A_0, F_0, b_0) = P(A \cup \{\rho \sqcap \sigma\}, F, [\rho_{i_1} \sqcap \sigma_{i_1} \dots \rho_{i_n} \sqcap \sigma_{i_n}], b)$ 
-10-          in if  $b_0 = fail$  then  $(A_0, F_0, fail)$ 
-11-          else  $P(A_0, F_0, xs, b_0)$ 
-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  $(A_0, F_0, b_0) = P(A \cup \{\rho \sqcap \sigma\}, F, [\rho_{i_1} \sqcap \sigma_{i_1} \dots \rho_{i_n} \sqcap \sigma_{i_n}], b)$ 
-14-          in if  $b_0 = fail$  then  $(A_0, F_0, fail)$ 
-15-          else  $P(A_0, F_0, xs, b_0)$ 
-16-   else if  $\rho = rec\ x. \rho'$  then  $P(A, F, (\{rec\ x. \rho' / x\} \rho' \sqcap \sigma), xs, b)$ 
-17-   else if  $\sigma = rec\ x. \sigma'$  then  $P(A, F, (\rho \sqcap \{rec\ x. \sigma' / x\} \sigma'), xs, b)$ 
-18-   else  $(A, F \cup \{\rho \sqcap \sigma\}, fail)$ 

```

and where

$$\mathbf{P}^+(\mathbf{A}, \mathbf{F}, [\rho \smile \sigma], \mathbf{b}) = \mathbf{P}(\mathbf{A}, \mathbf{F}, [\rho \smile \sigma], \mathbf{b})$$
$$\mathbf{P}^+(A, F, (\rho \smile \sigma) : x_s, b) =$$

```

-19-   let  (A0, F0, b0) = P(A, F, [ρ ↦ σ], b) in
-20-       if b0 = fail then P+(A ∪ A0, F ∪ F0, xs, ok)
-21-       else (A0, F0, b0)

```

Figure 11: The quadratic decision procedure for compliance.

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 $\rho \curvearrowright \sigma$, 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:

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

Corollary 21 (Complexity of **Decide**_⊥).

*The complexity of the algorithm **Decide**_⊥ 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 **Decide**_⊥ 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 **rC** of the system in Figure 4. Obviously, when applied to elements of **rC**, the clauses -8- and -12- of **Decide**_⊥ 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.

$$\begin{array}{c} \alpha.\sigma \oplus \sigma' \xrightarrow{\tau}_U \alpha.\sigma \qquad \alpha.\sigma \xrightarrow{\alpha}_U \sigma \qquad \alpha.\sigma + \sigma' \xrightarrow{\alpha}_U \sigma \\ \hline \frac{\rho \xrightarrow{\tau}_U \rho'}{\rho \parallel \sigma \longrightarrow_U \rho' \parallel \sigma} \qquad \frac{\sigma \xrightarrow{\tau}_U \sigma'}{\rho \parallel \sigma \longrightarrow_U \rho \parallel \sigma'} \qquad \frac{\rho \xrightarrow{\alpha}_U \rho' \quad \sigma \xrightarrow{\bar{\alpha}}_U \sigma'}{\rho \parallel \sigma \longrightarrow_U \rho' \parallel \sigma'} \end{array}$$

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 $\llbracket \cdot \rrbracket : \text{rsCo} \cup (\text{rsCo} \times \text{rsCo}) \longrightarrow \text{TC}$ is defined inductively as follows:

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

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.

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

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 $\sigma, \rho, \sigma', \rho' \in \text{rsC}$.

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

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.

$$\frac{\text{std}(X)}{\alpha_\tau.X \xrightarrow{\alpha_\tau} \underline{\alpha_\tau}.X} \quad \frac{X \xrightarrow{\alpha_\tau} X'}{\underline{\alpha_\tau}.X \xrightarrow{\alpha_\tau} \underline{\alpha_\tau}.X'} \quad \frac{X \xrightarrow{\alpha_\tau} X' \quad \text{std}(Y)}{X + Y \xrightarrow{\alpha_\tau} X' + Y}$$

$$\frac{X \xrightarrow{\tau} X'}{X \parallel Y \xrightarrow{\tau} X' \parallel Y} \quad \frac{X \xrightarrow{\alpha} X' \quad Y \xrightarrow{\bar{\alpha}} Y'}{X \parallel Y \xrightarrow{\tau} X' \parallel Y'}$$

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.

Figure 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.

$\llbracket a.c + b.d \rrbracket = a.c + b.d$		$\llbracket \bar{a}.\bar{c} + \bar{b}.(e \oplus f) \rrbracket = \bar{a}.\bar{c} + \bar{b}.(e + \tau.f)$	
$\langle \rangle \ltimes a.c + b.d \parallel \langle \rangle \ltimes \bar{a}.\bar{c} + \bar{b}.(e \oplus f)$	$a.c + b.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + \tau.f)$	$\longrightarrow_f a.c \ltimes d \parallel \bar{a}.\bar{c} \ltimes e \oplus f$	$\longrightarrow a.c + \underline{b}.d \parallel \bar{a}.\bar{c} + \underline{\bar{b}}.(\tau.e + \tau.f)$
$\longrightarrow_f a.c \ltimes d \parallel \bar{a}.\bar{c} \ltimes e$	$\longrightarrow a.c + \underline{b}.d \parallel \bar{a}.\bar{c} + \underline{\bar{b}}.(\tau.e + \tau.f)$	$\rightsquigarrow a.c + \underline{b}.d \parallel \bar{a}.\bar{c} + \underline{\bar{b}}.(\tau.e + \tau.f)$	$\rightsquigarrow a.c + b.d \parallel \bar{a}.\bar{c} + \bar{b}.(e + \tau.f)$
$\longrightarrow_b \langle \rangle \ltimes a.c \parallel \langle \rangle \ltimes \bar{a}.\bar{c}$	$\longrightarrow \underline{a}.c + b.d \parallel \underline{\bar{a}}.\bar{c} + \bar{b}.(e + \tau.f)$	$\longrightarrow_f \circ \ltimes c \parallel \circ \ltimes \bar{c}$	$\longrightarrow \underline{a}.c + b.d \parallel \underline{\bar{a}}.\bar{c} + \bar{b}.(e + \tau.f)$
$\longrightarrow_f \circ \ltimes \circ \ltimes \mathbf{1} \parallel \circ \ltimes \circ \ltimes \mathbf{1}$	$\longrightarrow \underline{a}.\underline{c} + b.d \parallel \underline{\bar{a}}.\underline{\bar{c}} + \bar{b}.(e + \tau.f)$		

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

Definition 24 (Barbs).

Given a retractable client-server pair $H_1 \ltimes \rho \parallel H_2 \ltimes \sigma$ and a name or coname α , we say that $H_1 \ltimes \rho \parallel H_2 \ltimes \sigma$ has a barb at α , written $H_1 \ltimes \rho \parallel H_2 \ltimes \sigma \downarrow_\alpha$, when defined by the following structural induction on $H_1 \ltimes \rho \parallel H_2 \ltimes \sigma$:

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

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 :

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

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 $(H_1 \ltimes \rho \parallel H_2 \ltimes \sigma, X) \in R$:

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

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

Theorem 23.

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

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.

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

It is easy to check, instead, that $\langle \rangle \ltimes a.c + b.d \parallel \langle \rangle \ltimes \bar{a}.\bar{c} + \bar{b}.(e \oplus f)$ 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 then 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, 2, 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, 36, 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.

Acknowledgments

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 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.

References

- [1] S. Carpineti, G. Castagna, C. Laneve, L. Padovani, A formal account of contracts for Web Services, in: WS-FM, no. 4184 in LNCS, Springer, 2006, pp. 148–162.
- [2] C. Laneve, L. Padovani, The must preorder revisited: An algebraic theory for web services contracts, in: CONCUR, Vol. 4703 of LNCS, Springer, 2007, pp. 212–225.
- [3] G. Castagna, N. Gesbert, L. Padovani, A theory of contracts for web services, ACM Trans. on Prog. Lang. and Sys. 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, Vol. 1381 of LNCS, Springer, 1998, pp. 22–138.
- [5] H. Hüttel, I. Lanese, V. T. Vasconcelos, L. Caires, M. Carbone, P.-M. Deniélou, 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, Vol. 5201 of LNCS, 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, Vol. 9039 of LNCS, 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. Dep. Sec. 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. Gabbrielli, I. Lanese, J. Mauro, G. Zavattaro, Graceful interruption of request-response service interactions, in: ICSOC, Vol. 7084 of LNCS, Springer, 2011, pp. 590–600.
- [13] F. Barbanera, U. de’Liguoro, Sub-behaviour relations for session-based client/server systems, MSCS 25 (6) (2015) 1339–1381.
- [14] G. T. Bernardi, M. Hennessy, Modelling session types using contracts, Mathematical Structures in Computer Science 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. of Logic and Alg. Progr. 73 (1-2) (2007) 70–96.
- [17] F. Barbanera, I. Lanese, U. de’Liguoro, Retractable and speculative contracts, in: COORDINATION, Vol. 10319 of LNCS, Springer, 2017, pp. 119–137.
- [18] F. Barbanera, M. Dezani-Ciancaglini, I. Lanese, U. de’Liguoro, Retractable contracts, in: PLACES 2015, Vol. 203 of EPTCS, 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. Symbolic Computation 11 (1991) 321–347.
- [21] M. R. Mousavi, I. C. C. Phillips, M. A. Reniers, I. Ulidowski, The meaning of ordered SOS, in: FSTTCS, Vol. 4337 of LNCS, 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*, Vol. 3170 of LNCS, 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*, Vol. 261 of EPTCS, 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*, Vol. 2404 of LNCS, 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, *Theoretical Computer Science* 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*, Vol. 162 of EPTCS, 2014, pp. 35–42.
- [34] F. Barbanera, M. Dezani-Ciancaglini, U. de’Liguoro, Reversible client/server interactions, *Formal Asp. Comput.* 28 (4) (2016) 697–722.
- [35] F. Tiezzi, N. Yoshida, Towards reversible sessions, in: *PLACES*, Vol. 155 of EPTCS, 2014, pp. 17–24.
- [36] F. Tiezzi, N. Yoshida, Reversible session-based pi-calculus, *J. Log. Algebr. Meth. Program.* 84 (5) (2015) 684–707.
- [37] C. A. Mezzina, J. A. Pérez, Reversibility in session-based concurrency: A fresh look, *J. Log. Algebr. Meth. Program.* 90 (2017) 2–30.
- [38] E. de Vries, V. Koutavas, M. Hennessy, Communicating transactions - (extended abstract), in: *CONCUR*, Vol. 6269 of LNCS, Springer, 2010, pp. 569–583.

- [39] F. Barbanera, U. de'Liguoro, Retractability, games and orchestrators for session contracts, Logical Methods in Computer Science Volume 13, Issue 3.
- [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, Vol. 222 of EPTCS, 2016, pp. 60–74.
- [43] I. Castellani, M. Dezani-Ciancaglini, P. Giannini, Concurrent reversible sessions, in: CONCUR, Vol. 85 of LIPIcs, 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, Logical Methods in Computer Science 12 (2).
- [45] I. Lanese, M. Lienhardt, C. A. Mezzina, A. Schmitt, J. Stefani, Concurrent flexible reversibility, in: ESOP, Vol. 7792 of LNCS, Springer, 2013, pp. 370–390.
- [46] A. Garcia, C. Laneve, Deadlock detection of java bytecode, CoRR abs/1709.04152. arXiv:1709.04152.
URL <http://arxiv.org/abs/1709.04152>
- [47] K. Helsgaun, Cback: A simple tool for backtrack programming in C, Softw., Pract. Exper. 25 (8) (1995) 905–934.

Appendix A. Proofs

Appendix A.1. Proofs of Section 3

Proof of Theorem 1

It is straightforward to check the following

Fact 24. *i) $\text{Prove}(\Gamma \triangleright \rho \multimap \sigma) \neq \text{fail}$ iff $\Gamma \triangleright \rho \multimap \sigma$.*

ii) $\text{Prove}(\Gamma \triangleright \rho \multimap \sigma) = \mathcal{D} \neq \text{fail}$ implies $\mathcal{D}$ has conclusion $\Gamma \triangleright \rho \multimap \sigma$.

By Fact 24, we only need to show that the procedure **Prove** always terminates. Note that, in all recursive calls $\text{Prove}(\Gamma, \rho \multimap \sigma \triangleright \rho_k \multimap \sigma_k)$ inside $\text{Prove}(\Gamma \triangleright \rho \multimap \sigma)$, 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 $\langle \rangle \times \rho \parallel \langle \rangle \times \sigma \xrightarrow{*} H_1 \times \rho' \parallel H_2 \times \sigma' \not\rightarrow$, then $H_1 = H_2 = \langle \rangle$.*

Proof. Clearly $H_1 \times \rho' \parallel H_2 \times \sigma' \not\rightarrow$ 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

$$\langle \rangle \times \rho_{k_i} \parallel \langle \rangle \times \sigma_{k_i} \xrightarrow{*} H_i \times \rho'_i \parallel H'_i \times \sigma'_i \not\rightarrow$$

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:

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

and $\langle \rangle \ltimes \sum_{i \in I'} \alpha_i \cdot \rho_i \parallel \langle \rangle \ltimes \sum_{j \in J'} \bar{\alpha}_j \cdot \sigma_j$ 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

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

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

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

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

which imply

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

In both the cases we conclude that $\rho \not\ll^R \sigma$.

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 $\dashv^R$ can be coinductively defined, or equivalently that $\dashv^R = \bigcap_n \dashv_n^R$ where $\dashv_0^R$ is the trivial relation $\mathbf{rsC} \times \mathbf{rsC}$, and for all $n > 0$, $\mathbf{1} \dashv_n^R \sigma$ and we have $\rho \dashv_n^R \sigma$ 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 $\exists k \in I \cap J$. $\rho_k \dashv_{n-1}^R \sigma_k$;
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 $\forall k \in I$. $\rho_k \dashv_{n-1}^R \sigma_k$;
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 $\forall k \in J$. $\rho_k \dashv_{n-1}^R \sigma_k$.

We write:

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

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

Observing that $\dashv_{n+1}^R \subseteq \dashv_n^R$, 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 $\Gamma \models_n^R \rho \smile \sigma$ for all n , then $\Gamma \models^R \rho \smile \sigma$.*

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

Theorem 2 (Retractable Soundness and Completeness)

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

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

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

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

$$\frac{(+ \cdot +) \quad \Gamma, \alpha.\rho + \rho' \smile \bar{\alpha}.\sigma + \sigma' \triangleright \rho \smile \sigma}{\Gamma \triangleright \alpha.\rho + \rho' \smile \bar{\alpha}.\sigma + \sigma'}$$

then we have to show that $\models_n^R \Gamma$ implies $\alpha.\rho + \rho' \dashv_n^R \bar{\alpha}.\sigma + \sigma'$. By induction over n we know that $\Gamma \models_{n-1}^R \alpha.\rho + \rho' \smile \bar{\alpha}.\sigma + \sigma'$; from this and the fact that $\models_n^R \Gamma$ implies $\models_{n-1}^R \Gamma$, we obtain that $\alpha.\rho + \rho' \dashv_{n-1}^R \bar{\alpha}.\sigma + \sigma'$, and hence that $\models_{n-1}^R \Gamma, (\alpha.\rho + \rho' \smile \bar{\alpha}.\sigma + \sigma')$. By induction over $\mathcal{D}$ it follows that $\rho \dashv_{n-1}^R \sigma$, which implies $\alpha.\rho + \rho' \dashv_n^R \bar{\alpha}.\sigma + \sigma'$ by definition of $\dashv_n^R$, 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**($\triangleright \rho \smile \sigma$) always terminates. By Lemma 27 and Fact 24, $\rho \dashv^R \sigma$ implies that **Prove**($\triangleright \rho \smile \sigma$) $\neq$ **fail**, and hence $\triangleright \rho \smile \sigma$. $\square$

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. $\rho = \sum_{i \in I} \alpha_i.\rho_i$, $\sigma = \sum_{j \in J} \bar{\alpha}_j.\sigma_j$ and $\exists k \in I \cap J$. $\rho_k \dashv^S \sigma_k$;
3. $\rho = \bigoplus_{i \in I} \bar{\alpha}_i.\rho_i$, $\sigma = \sum_{j \in J} \alpha_j.\sigma_j$, $I \subseteq J$ and $\forall k \in I$. $\rho_k \dashv^S \sigma_k$;
4. $\rho = \sum_{i \in I} \bar{\alpha}_i.\rho_i$, $\sigma = \bigoplus_{j \in J} \alpha_j.\sigma_j$, $I \supseteq J$ 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 $\rho = \sum_{i \in I} \alpha_i.\rho_i$, $\sigma = \sum_{j \in J} \bar{\alpha}_j.\sigma_j$, $I \cap J = \{k_1, \dots, k_n\}$ and $\rho_{k_i} \not\vdash^S \sigma_{k_i}$

for $1 \leq i \leq n$. Then we get $\rho_{k_i} \parallel \sigma_{k_i} \xrightarrow{*} \mathbf{C}_i^\rho \parallel \mathbf{C}_i^\sigma \not\rightarrow$

for $1 \leq i \leq n$, where $\mathbf{C}_i^\rho \neq \mathbf{C} \mid \alpha_1 @ \dots @ \alpha_n @ \mathbf{1}$.

This implies

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

where $\prod_{i \in I} \mathbf{C}_i$ denotes the parallel composition of $\mathbf{C}_i$ for each $i \in I$. Terms $\sum_{i \in I \setminus J} \alpha_i.\rho_i$ and $\sum_{j \in J \setminus I} \bar{\alpha}_j.\sigma_j$ 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 $\rho_{k_i} \parallel \sigma_{k_i} \xrightarrow{*} \mathbf{C}_i^\rho \parallel \mathbf{C}_i^\sigma \not\rightarrow$, 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 $\rho \not\parallel^S \sigma$. 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

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

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

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

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 $\rho \not\parallel^S \sigma$.

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 $\triangleright \rho \dashv \sigma$. $\square$

Appendix 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 $\rho \in \mathbf{SC}$. 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) $\bar{a} \cdot \rho_1 \oplus \rho_2 \xrightarrow{\tau} \bar{a} \cdot \rho_1$ if and only if $H \times \bar{a} \cdot \rho_1 \oplus \rho_2 \xrightarrow{\tau} H \times \bar{a} \cdot \rho_1$;
- ii) $\alpha \cdot \rho' \xrightarrow{\alpha} \rho'$ if and only if $H \times \alpha \cdot \rho' \xrightarrow{\alpha} H : \circ \times \rho'$;
- iii) $a \cdot \rho_1 + \rho_2 \xrightarrow{a} \rho_1$ if and only if $H \times a \cdot \rho_1 + \rho_2 \xrightarrow{a} H : \rho_2 \times \rho_1$

iv) $\rho = \mathbf{1}$ if and only if $(H \ltimes \rho \not\rightarrow^Q \text{ and } H \ltimes \rho \not\rightarrow^T)$.

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

Fact 31. Let $\rho, \sigma \in \mathbf{SC}$. $H_1 \ltimes \rho \parallel H_2 \ltimes \sigma \rightarrow_f H'_1 \ltimes \rho' \parallel H'_2 \ltimes \sigma'$ implies $\rho', \sigma' \in \mathbf{SC}$

We can now proceed with the proof of Propositions 7.

Proof of Proposition 7

(i) The inclusion $\mathbf{SC} \subseteq \mathbf{rsC}$ holds by definition. Hence, given $\rho, \sigma \in \mathbf{SC}$ we have $\rho, \sigma \in \mathbf{rsC}$ and $H_1 \ltimes \rho, H_2 \ltimes \sigma \in \mathbf{rsCH}$.

($\Rightarrow$) By induction on the length of the reduction sequence $\xrightarrow{*}_{\mathbf{SC}}$, 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 $H_1 \ltimes \alpha.\rho_1 + \rho_2 \parallel H_2 \ltimes \bar{\alpha}.\sigma_1 + \sigma_2$ 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 $\rho \parallel \sigma \xrightarrow{*}_{\mathbf{SC}} \rho' \parallel \sigma'$, hence the thesis follows since ρ' and σ' can match the synchronization. $\square$

Appendix A.2. Proofs of Section 4

Proof of Proposition 11

Like $\dashv^R$, $\dashv$ can be coinductively defined, since, by Corollary 4, $\dashv^R = \dashv^S = \dashv$. So $\dashv = \bigcap_n \dashv_n$ where $\dashv_n = \dashv_n^R$.

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 $\dashv_n$.

$\rho = 1$ Immediate.

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

We have then that $\bar{\sigma} = \bigoplus_{j \in J} \alpha_j \cdot \bar{\sigma}_j$. So, by $\bar{\sigma} \dashv_n \sigma'$ and by definition of $\dashv_n$ we have that $\sigma' = \sum_{h \in H} \bar{\alpha}_h \cdot \sigma'_h$, $J \subseteq H$ 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 $\rho \dashv_n \sigma'$.

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

We have then that $\bar{\sigma} = \bigoplus_{j \in J} \bar{\alpha}_j \cdot \bar{\sigma}_j$. So, by $\bar{\sigma} \dashv_n \sigma'$ and by definition of $\dashv_n$ we have that $\sigma' = \sum_{h \in H} \alpha_h \cdot \sigma'_h$, $J \subseteq H$ 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 $\rho \dashv_n \sigma'$.

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

We have then that $\bar{\sigma} = \sum_{j \in J} \bar{\alpha}_j \cdot \bar{\sigma}_j$. 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 $\bar{\sigma}' = \sigma'$.

(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 $\sigma = \sigma'$. $\square$

Appendix A.3. Proofs of Section 5

Proof sketch of Lemma 19

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

Otherwise, we observe that if an occurrence $\Gamma', \rho' \multimap \sigma' \triangleright \rho' \multimap \sigma'$ of rule (HYP) occurs in a given derivation of $\triangleright \rho \multimap \sigma$, then the expression $\rho' \multimap \sigma'$ 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 $\triangleright_{\infty} \rho \multimap \sigma$ 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 $\triangleright_{\infty} \rho \multimap \sigma$.

($\Leftarrow$) Take an infinite derivation of $\triangleright_{\infty} \rho \multimap \sigma$ (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 $\triangleright \rho \multimap \sigma$ by replacing all the expressions of the form $\rho' \multimap \sigma'$ 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 $\triangleright_{\infty}$ 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 $\mathcal{O}(n^2)$. $\square$

Appendix 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 $(\langle \rangle \ltimes \rho \parallel \langle \rangle \ltimes \sigma, \llbracket \rho \parallel \sigma \rrbracket)$. 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 barbs 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 $H \ltimes \rho \parallel H' \ltimes \rho'$. The possible contracts with history $H \ltimes \rho$ (resp. $H' \ltimes \rho'$) 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 $\sigma' = \sum_{i \in I} \pi_i.\sigma_i$, 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 $\sum_{i \in J} \pi_i.\sigma_i$ 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 $\llbracket \hat{\sigma} \rrbracket = \sum_{i \in J} \pi_i.\sigma_i$. 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. $\square$