

Modelling session types using contracts

Giovanni Bernardi and Matthew Hennessy [†]

School of Statistics and Computer Science, The University of Dublin, Trinity College

Received July 30, 2013

Session types and contracts are two formalisms used to study client-server protocols. In this paper we study the relationship between them. The main result is the existence of a fully abstract model of session types; this model is based on a natural interpretation of these types into a subset of contracts.

Contents

1	Introduction	1
2	Session types	4
2.1	Sub-typing	6
3	Contracts	9
3.1	The contract language	9
3.2	The server pre-order	20
3.3	Must testing	25
4	Session Contracts	30
4.1	Session contracts	30
4.2	The restricted server pre-order	31
4.3	The restricted client pre-order	36
5	Modelling session types	39
5.1	Examples and applications	42
6	Conclusions	44
6.1	Summary	44
6.2	Related work	45
6.3	Future work	50

1. Introduction

Communication between processes in a distributed system often consists of a structured dialogue, following a protocol which specifies the format of the messages interchanged and, at least for binary communication, the direction of the messages. *Session types*, *ST*,

[†] This research was supported by SFI project 06 IN.1 1898.

have been introduced as an approach to the static analysis of the participants of such dialogues. They allow structured sequences of non-uniform messages to be interchanged between the participants. For example, using the notation of (Gay & Hole 2005), the type $![\mathbf{Int}]; ?[\mathbf{Real}]; \mathbf{END}$ specifies the output of a value of type \mathbf{Int} followed by the input of a value of type \mathbf{Real} , after which the dialogue is terminated. Flexibility in the permitted sequencing of messages by a process is accommodated by two choice operators; the *branching type* $\&\langle T_1, T_2 \rangle$ offers a choice to the partner in the dialogue between following either the protocol specified by the type T_1 or that specified by T_2 . On the other hand the *choice type* $\oplus\langle T_1, T_2 \rangle$ allows the process itself to follow either of the protocols specified by T_1 or T_2 .

Sub-typing, (Gay & Hole 2005), also increases the flexibility of the type system; intuitively $T_1 \preceq_{\text{ST}} T_2$ means that any participant designed with the protocol specified by T_1 in mind may also be used in a situation where the protocol specified by T_2 will be followed. Intuitively this pre-order between session types is generated by allowing more possibilities in branching types and restricting them in choice types. The reader is referred to (Caires & Pfenning 2010, Gay & Hole 2005, Honda et al. 1998) for more details on session types, including how they are associated with processes and what behaviour they guarantee.

Web services (Alonso et al. 2004, Bernardo et al. 2009) are distributed components which can be combined using standard communication protocols and machine-independent message formats to provide services to clients. To encourage reusability, descriptions of their behaviour are typically made available in searchable repositories (Oasis Standard 2011). In papers such as (Barbanera & de'Liguoro 2010, Carpineti et al. 2006, Castagna et al. 2009, Laneve & Padovani 2007) a language of *contracts* has been proposed for describing behaviours; despite a very different surface syntax, the language of contracts is very similar in style to session types. In particular there is the sequencing of messages $\alpha_1.\alpha_2$, an external choice between behaviours $\sigma_1 + \sigma_2$ reminiscent of the branching type $\&\langle T_1, T_2 \rangle$, and an internal choice between allowed behaviours $\sigma_1 \oplus \sigma_2$, reminiscent of the choice type $\oplus\langle T_1, T_2 \rangle$.

The object of this paper is to study the precise relationship between these two formalisms. In particular for *first-order* session types, which do not allow the use of communication channels in messages, we show that the theory of session types, $\langle \mathcal{ST}, \preceq_{\text{ST}} \rangle$, can be captured precisely using a natural pre-order over a subclass of contracts.

Contracts for web services serve two roles. A contract σ may describe the behaviour of a server offering some specific service. Dually a contract ρ may describe the behaviour expected of a client who wishes to avail of a particular service. Central to the theory of contracts for web services is the idea of compliance between such contracts, formalised as an asymmetric relation $\rho \dashv \sigma$; it has been defined in a variety of ways in papers such as (Castagna et al. 2009, Laneve & Padovani 2007, 2008). This leads to two natural pre-orders on contracts, defined set theoretically:

- the server pre-order: $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2$ if for every (client) contract ρ , $\rho \dashv \sigma_1$ implies $\rho \dashv \sigma_2$
- the client pre-order: $\rho_1 \sqsubseteq_{\text{CLT}} \rho_2$ if for every (server) contract σ , $\rho_1 \dashv \sigma$ implies $\rho_2 \dashv \sigma$

As we have already stated session types are more or less a syntactic variant of contracts; formally there is a straightforward translation $\mathcal{M}(T)$ of session types into contracts.

Unfortunately neither of the relations $\sqsubseteq_{\text{SRV}}, \sqsubseteq_{\text{CLT}}$ are sound with respect to sub-typing; specifically there are session types T_1, T_2 such that $T_1 \preceq_{\text{ST}} T_2$ but $\mathcal{M}(T_1)$ and $\mathcal{M}(T_2)$ are unrelated as contracts.

The problem lies in the fact that, viewed as constraints on behaviour, session types are much more constraining than contracts. We therefore isolate a subset of contracts, which we call *session contracts*, \mathcal{SC} , that are the range of the translation function \mathcal{M} . This enables us to define a sub-server relation and a sub-client relation, $\sqsubseteq_{\text{SRV}}^{\mathcal{SC}}$ and $\sqsubseteq_{\text{CLT}}^{\mathcal{SC}}$ respectively, on these contracts. Even though these relations are respectively coarser than \sqsubseteq_{SRV} and \sqsubseteq_{CLT} , it turns out that these relations are still unsound with respect to session sub-typing. But in the main result of the paper we show that by combining these pre-orders we obtain *full-abstraction*, that is a sound and complete model for session types.

Throughout the paper we follow the approach of Castagna et al. (2009), Padovani (2010), and think of “contracts” only as terms of (some dialect of) CCS without τ ’s (Nicola & Hennessy 1987); in particular, our notion of contract is completely independent of that discussed in (McNeile 2010) and (Meyer 1997).

Contributions The contributions of this paper are the following: Theorem 5.7 provides what, to the best of our knowledge, is the first fully abstract model of session types in terms of contracts. Theorem 4.14 and Theorem 4.24 are the first alternative characterisations of the server and the client pre-orders on session contracts. Corollary 3.51 is the first published proof showing the equality between a first-order compliance-based refinement and a testing based must-preorder; moreover, the corollary means that the equality holds regardless of the co-action relation used to let contracts interact.

Structure of the paper The paper is organised as follows. In the next section we give the definition of session types and the sub-typing between them; this material is taken directly from (Gay & Hole 2005), although our definition is based on first-order types; however, we allow a primitive sub-typing relation between the basic types.

In the subsequent section we study contracts. We use the language for contracts of (Padovani 2010), and we provide a formulation of the notion of compliance which differs from the one of (Padovani 2010) in that (a) it is co-inductively defined, and (b) it is parametrised over the co-action relations \bowtie . Then, disregarding the parameter \bowtie , we provide a co-inductive characterisation of the server pre-order on contracts, and we prove that it equals the must pre-order over contracts. As we reason up-to \bowtie , the result that we obtain is more general than a similar result presented in (Laneve & Padovani 2007). In Section 4 we focus on a subset of contracts called session contracts \mathcal{SC} , this time giving co-inductive characterisations to both the restricted server pre-order $\sqsubseteq_{\text{SRV}}^{\mathcal{SC}}$ and the restricted client pre-order $\sqsubseteq_{\text{CLT}}^{\mathcal{SC}}$ over them. Due to the very restricted nature of these contracts, these co-inductive characterisations are purely in terms of their syntax. In Section 5 we tackle the central question of the paper. Having defined the (obvious) translation of session types into session contracts, we explain why the two natural pre-orders $\sqsubseteq_{\text{SRV}}^{\mathcal{SC}}$ and $\sqsubseteq_{\text{CLT}}^{\mathcal{SC}}$ are unsound relative to the sub-typing on session types. Finally, we prove that when combined they provide a sound and complete model; the proof is

$S, T ::=$	Session types
END	<i>satisfaction</i>
$?[\mathbf{t}]; S$	<i>input</i>
$![\mathbf{t}]; S$	<i>output</i>
$\&\langle \mathbf{l}_1 : S_1, \dots, \mathbf{l}_n : S_n \rangle, n \geq 1$	<i>branch</i>
$\oplus\langle \mathbf{l}_1 : S_1, \dots, \mathbf{l}_n : S_n \rangle, n \geq 1$	<i>choice</i>
X	<i>type variable</i>
$\mu X. S$	<i>recursion</i>

We impose the additional proviso that in a term the \mathbf{l}_i 's are pair-wise different.

Figure 1. Session types (first-order).

greatly facilitated by their co-inductive characterisations. The paper concludes with a brief look at related work.

2. Session types

The syntax of terms for session types is given by the language $L_{\mathcal{ST}}$ in Figure 1. It presupposes a denumerable set of labels \mathbb{L} , ranged over by \mathbf{l} , and a set of basic or ground types \mathbf{BT} types ranged over by \mathbf{t} . We also use a denumerable set of variables $Vars$, ranged over by X , in order to express recursive types.

The use of variables leads to the usual notion of *free* and *bound* occurrences of variables in terms in the standard manner; we say that a term is *closed* if it contains no free variables. We also have the standard notion of *capture avoidance* substitution of terms for free variables. For the sake of clarity let us recall this definition: a substitution \mathbf{s} is a mapping from the set $Vars$ to the set of terms in $L_{\mathcal{ST}}$. Let

$$\mathbf{s} - X = \begin{cases} \mathbf{s} \setminus \{(X, \mathbf{s}(X))\} & \text{if } X \in \text{dom}(\mathbf{s}) \\ \mathbf{s} & \text{otherwise} \end{cases}$$

Then the result of applying a substitution \mathbf{s} to the term S is defined as follows:

$$S\mathbf{s} = \begin{cases} \text{END} & \text{if } S = \text{END} \\ \mathbf{s}(X) & \text{if } S = X, \text{ and } X \in \text{dom}(S) \\ X & \text{if } S = X, \text{ and } X \notin \text{dom}(S) \\ ![\mathbf{t}]; (S'\mathbf{s}) & \text{if } S = ![\mathbf{t}]; S' \\ ?[\mathbf{t}]; (S'\mathbf{s}) & \text{if } S = ?[\mathbf{t}]; S' \\ \&\langle \mathbf{l}_1 : (S_1\mathbf{s}), \dots, \mathbf{l}_n : (S_n\mathbf{s}) \rangle & \text{if } S = \&\langle \mathbf{l}_1 : S_1, \dots, \mathbf{l}_n : S_n \rangle \\ \oplus\langle \mathbf{l}_1 : (S_1\mathbf{s}), \dots, \mathbf{l}_n : (S_n\mathbf{s}) \rangle & \text{if } S = \oplus\langle \mathbf{l}_1 : S_1, \dots, \mathbf{l}_n : S_n \rangle \\ \mu X. (S'(\mathbf{s} - X)) & \text{if } S = \mu X. S' \end{cases}$$

In the final clause the application of $\mathbf{s} - X$ embodies the idea that in $\mu X. S'$ occurrences of X in the sub-term S' are bound and therefore substitutions have no effect on them.

$$\frac{}{T} \quad T \neq \mu Z. S \quad \frac{T \{ \mu Y. T / Y \}}{\mu Y. T}$$

Figure 2. Inference rules for \downarrow_{DPT} on *closed* terms.

It is easy to check that the effect of a substitution depends only on free variables; that is, $Ss_1 = Ss_2$ whenever $s_1(X) = s_2(X)$ for every free variable X occurring in S . We use the symbol $\{T/X\}$ to denote the singleton substitution $\{(X, T)\}$.

We will only use guarded recursion, which we now explain formally. Let \downarrow_{DPT} be the least fixed point of the functional on closed terms defined by the inference rules in Figure 2.

Intuitively, $T \downarrow_{\text{DPT}}$ means that the free variables in T occur after a type constructor, which differs from μ . Now we say that a term T is *guarded* if every sub-term of T that has the form $\mu X. S$ satisfies $S \downarrow_{\text{DPT}}$. Finally, we use \mathcal{ST} to denote the set of closed guarded terms, and we refer to the elements in \mathcal{ST} as *session types*.

Example 2.1. The property \downarrow_{DPT} and the property of being guarded are different. Consider the term $T = \&\langle 1: \mu X. X \rangle$; it is not a variable and the top-most constructor in it is not a recursion, therefore $T \downarrow_{\text{DPT}}$. A sub-term of T is $\mu X. X$ and clearly $\mu X. X \downarrow_{\text{DPT}}$ is false; therefore T is not guarded. \square

The advantage of only using guarded terms is that we can unfold types so as to obtain their top-most type constructor. To explain this formally first let us consider the function *depth* from terms to \mathbb{N}^∞ (the set of natural numbers augmented by ∞). This is defined as the least such function which satisfies:

$$\text{depth}(S) = \begin{cases} 1 + \text{depth}(S' \{S/X\}) & \text{if } S = \mu X. S', \\ 0 & \text{otherwise} \end{cases}$$

Note that $\text{depth}(\mu X. X) = \infty$, but one can show that when applied to terms that satisfy the predicate \downarrow_{DPT} , one always obtains a natural number.

Lemma 2.2. If $S \downarrow_{\text{DPT}}$ then $\text{depth}(S) \in \mathbb{N}$.

Proof. The proof is by rule induction on the derivation of $S \downarrow_{\text{DPT}}$. If the axiom was used then thanks to the side condition $S \neq \mu X. S'$, and so $\text{depth}(S) = 0$ because of the definition of *depth*. If the other rule was used then $S = \mu X. S'$, and the hypothesis of the rule implies $S' \{S'/X\} \downarrow_{\text{DPT}}$. Then, by definition of *depth*, $\text{depth}(S) = 1 + \text{depth}(S' \{S'/X\})$, and by the inductive hypothesis $\text{depth}(S' \{S'/X\}) \in \mathbb{N}$; hence $\text{depth}(S) \in \mathbb{N}$. \square

Proposition 2.3. The depth of any session type is finite.

Proof. Follows from the definition of \mathcal{ST} and Lemma 2.2. \square

This function *depth* will therefore provide a measure of session types over which we can perform induction.

Definition 2.4. [Unfolding (Gay & Hole 2005)]

For all $T \in \mathcal{ST}$, define $\text{UNFOLD}(T)$ as follows:

$$\text{UNFOLD}(T) = \begin{cases} \text{UNFOLD}(T' \{ \mu X.T / X \}) & \text{if } T = \mu X.T' \\ T & \text{otherwise} \end{cases}$$

□

Lemma 2.5. For every $T \in \mathcal{ST}$, $\text{UNFOLD}(T)$ is a well-defined session type.

Proof. We have to show that $\text{UNFOLD}(T)$ is closed and guarded. The proof is by induction on $\text{depth}(P)$. It relies on the fact that each step of unfolding replaces one variable with a closed and guarded term; hence the overall unfolding is closed. □

Intuitively, $\text{UNFOLD}(T)$ unfolds top-level recursive definitions until a type constructor appears, which is not μ . This will be extremely useful in manipulating session types.

We conclude this sub-section by showing some typical examples of session types, which we recall from the literature.

Example 2.6. [Math server, (Gay & Hole 2005)]

Consider the session type

$$S_1 = \mu X. \&\langle \text{plus} : ?[\text{Int}]; ?[\text{Int}]; ![\text{Int}]; X, \\ \text{eq} : ?[\text{Int}]; ?[\text{Int}]; ![\text{Bool}]; \text{END} \rangle$$

This specifies the protocol of a server which offers two services at the labels `plus` and `eq`. The first expects the input of two integers, after which an integer is returned, and then the service is once more available. The second also expects two integers, then returns a boolean, after which the session terminates.

An extension to the service is specified by the type

$$S_2 = \mu X. \&\langle \text{plus} : ?[\text{Int}]; ?[\text{Int}]; ![\text{Int}]; X, \\ \text{eq} : ?[\text{Int}]; ?[\text{Int}]; ![\text{Bool}]; \text{END}, \\ \text{neg} : ?[\text{Bool}]; ![\text{Bool}]; \text{END} \rangle$$

This provides in addition queries for negation. □

2.1. Sub-typing

There are three sources for the sub-typing relation over types. The first is some predefined pre-order over the basic types, $\mathfrak{t}_1 \preceq_{\mathfrak{g}} \mathfrak{t}_2$, which intuitively says that all data-values that have type \mathfrak{t}_1 may be safely used where data-values of type \mathfrak{t}_2 are expected.

Example 2.7. An example of sub-typing on base types is given in Figure 3, for the ensuing set of types, $\text{BT} = \{\text{Bounded}, \text{Bool}, \text{Int}, \text{Real}, \text{Num}, \text{Random}\}$. In the figure the pre-order $\preceq_{\mathfrak{g}}$ is depicted by the arrows; for instance, the arrow from type `Int` to type `Real` means that $\text{Int} \preceq_{\mathfrak{g}} \text{Real}$. □




Figure 3. A sub-type relation on a set of basic types BT ; the arrow represents the relation \preceq_g (see Example 2.7).

More generally, if $\llbracket \mathbf{t} \rrbracket$ denotes the set of values of the basic type \mathbf{t} then we can define \preceq_g by letting $\mathbf{t}_1 \preceq_g \mathbf{t}_2$ whenever $\llbracket \mathbf{t}_1 \rrbracket \subseteq \llbracket \mathbf{t}_2 \rrbracket$. The other sources are two constructs of the language: the *branch* construct allows sub-typing by extending the set of labels involved, while in the *choice* construct the set of labels may be restricted (see Example 2.8 and Example 2.9 below). Moreover, we will have the standard co-variance/contra-variance of input/output types (Pierce & Sangiorgi 1996), extended to both the *branch* and *choice* constructs.

Example 2.8. [Sub-typing on branch types]

In this example we explain how the sub-typing relates the branch types. Consider the type

$$\text{BARTENDER} = \&\langle \text{espresso} : T_1 \rangle$$

Intuitively, the BARTENDER offers only the label **espresso**, thus all the customers satisfied by BARTENDER, are satisfied by any other type that offers *at least* the label **espresso**. Let

$$\begin{aligned} \text{ITALIANBARTENDER} = \&\langle \text{espresso} : T'_1, \\ &\text{deka} : T'_2, \\ &\text{doubleespresso} : T'_3 \rangle \end{aligned}$$

Following the intuition, The ITALIANBARTENDER will satisfy all the customers satisfied by the BARTENDER; this is formalised by the sub-typing, which relates the two types as follows

$$\text{BARTENDER} \preceq_{\text{ST}} \text{ITALIANBARTENDER}$$

as long as also the continuations T_1 and T'_1 are related as well (ie. $T_1 \preceq_{\text{ST}} T'_1$).

We have shown that, intuitively, it is safe to replace a branch type with a branch type that offers more labels. \square

Example 2.9. [Sub-typing on choice types]

In this example we show how the sub-typing relate the choice types. Let ITALIANCUSTOMER describe the different coffees that a process may want to order when interacting

with a bar tender.

$$\begin{aligned} \text{ITALIANCUSTOMER} = \oplus \langle & \text{espresso} : T'_1, \\ & \text{deka} : T'_2, \\ & \text{doubleespresso} : T'_3 \rangle \end{aligned}$$

All the bar tenders that are able to satisfy this range of choices, have to offer *at least* the four labels that appear in ITALIANCUSTOMER. Now consider the type

$$\text{CUSTOMER} = \oplus \langle \text{espresso} : T'_1 \rangle$$

Since CUSTOMER chooses among fewer options than ITALIANCUSTOMER, it is safe to use a channel at type CUSTOMER in place of a channel at type ITALIANCUSTOMER. This is formalised by the sub-typing relation as follows,

$$\text{ITALIANCUSTOMER} \preceq_{\text{ST}} \text{CUSTOMER}$$

In this example we have shown that, intuitively, it is safe to replace a choice type, with a choice type that chooses among fewer labels. \square

However, because of the recursive nature of our collection of types, the formal definition of the sub-typing relation is given co-inductively.

Definition 2.10. [Type simulation]

Let $\mathcal{P}(X)$ denote the powerset of a set X and let $\mathcal{F}_{\preceq_{\text{ST}}} : \mathcal{P}(\mathcal{ST}^2) \rightarrow \mathcal{P}(\mathcal{ST}^2)$ be the function defined so that $(T, U) \in \mathcal{F}_{\preceq_{\text{ST}}}(\mathcal{R})$ whenever one of the following holds:

- (i) if $\text{UNFOLD}(T) = \text{END}$ then $\text{UNFOLD}(U) = \text{END}$
- (ii) if $\text{UNFOLD}(T) = ?[\mathbf{t}_1]; S_1$ then $\text{UNFOLD}(U) = ?[\mathbf{t}_2]; S_2$ and $(S_1, S_2) \in \mathcal{R}$ and $\mathbf{t}_1 \preceq_{\mathbf{g}} \mathbf{t}_2$
- (iii) if $\text{UNFOLD}(T) = ![\mathbf{t}_1]; S_1$ then $\text{UNFOLD}(U) = ![\mathbf{t}_2]; S_2$ and $(S_1, S_2) \in \mathcal{R}$ and $\mathbf{t}_2 \preceq_{\mathbf{g}} \mathbf{t}_1$
- (iv) if $\text{UNFOLD}(T) = \&\langle \mathbf{l}_1 : S_1, \dots, \mathbf{l}_m : S_m \rangle$ then $\text{UNFOLD}(U) = \&\langle \mathbf{l}'_1 : S'_1, \dots, \mathbf{l}'_n : S'_n \rangle$ where $m \leq n$ and $(S_i, S'_i) \in \mathcal{R}$ for all $i \in [1, \dots, m]$
- (v) if $\text{UNFOLD}(T) = \oplus \langle \mathbf{l}_1 : S_1, \dots, \mathbf{l}_m : S_m \rangle$ then $\text{UNFOLD}(U) = \oplus \langle \mathbf{l}'_1 : S'_1, \dots, \mathbf{l}'_n : S'_n \rangle$ where $n \leq m$ and $(S_i, S'_i) \in \mathcal{R}$ for all $i \in [1, \dots, n]$

If $\mathcal{R} \subseteq \mathcal{F}_{\preceq_{\text{ST}}}(\mathcal{R})$ then we say that \mathcal{R} is a co-inductive type simulation. Let \preceq_{ST} denote the greatest solution of the equation $X = \mathcal{F}_{\preceq_{\text{ST}}}(X)$; formally,

$$\preceq_{\text{ST}} = \nu \mathcal{F}_{\preceq_{\text{ST}}}$$

We call \preceq_{ST} the *type simulation*. Standard arguments ensure that the relation \preceq_{ST} exists, that it is a typing relation, and indeed the greatest type simulation. \square

Example 2.11. Let T, S denote the types $\mu X. ?[\text{Int}]; X$, $\mu X. ?[\text{Real}]; X$ respectively. We can prove that $T \preceq_{\text{ST}} S$, because the relation $\mathcal{R} = \{(T, S), (?[\text{Int}]; T, ?[\text{Real}]; S)\}$ is a type simulation, since $\text{Int} \preceq_{\mathbf{g}} \text{Real}$.

Referring to Example 2.6, one can also show that $S_1 \preceq_{\text{ST}} S_2$ by providing an appropriate type simulation. \square

The requirement that session types be guarded is crucial for the sub-typing relation to be well-defined. We explain this fact in the next example.

Example 2.12. [Sub-typing and guardedness]

Consider again the term $T = \&\langle 1: \mu X. X \rangle$ of Example 2.1. Suppose we wanted to check whether $T \preceq_{\text{ST}} \&\langle 1: S \rangle$ for some session type S . The definition of \preceq_{ST} requires us to check whether

$$\text{UNFOLD}(\mu X. X) \preceq_{\text{ST}} \text{UNFOLD}(S)$$

This check, though, cannot be done because $\text{UNFOLD}(\mu X. X)$ is *not* defined at all (and the unfolding of S may not be defined either). \square

Proposition 2.13. The relation \preceq_{ST} is a pre-order on \mathcal{ST} .

Proof. See (Gay & Hole 2005). \square

In (Gay & Hole 2005) the set of types \mathcal{ST} are used to give a typing system for the pi calculus, and appropriate Type Safety and Type Preservation theorems are proved. Here instead our aim is to give a model to the set of types $\langle \mathcal{ST}, \preceq_{\text{ST}} \rangle$ using contracts.

3. Contracts

We first define our language for contracts and give some examples. In the following sub-section we define a natural server based pre-order on contracts, for which we give a behavioural co-inductive characterisation. In the final sub-section we investigate a closely related pre-order based on must testing.

3.1. The contract language

This subsection is roughly divided in three parts. In the first one we define the language L_C , and we are concerned with *syntactical* properties of its terms σ 's; similarly to what we have done for session types, we introduce the unfolding of closed terms of L_C and a predicate \downarrow_{DPT} to guarantee that each contract can be unfolded (Lemma 3.1). Afterwards, we give an operational semantics (Figure 5), and we discuss important *semantic* properties that we want contracts to enjoy; this will lead to the introduction of a predicate \downarrow , and two lemmas (Lemma 3.10 and 3.11) which ensure that (a) contracts do not diverge, and (b) silent moves lead to a finite number of derivatives. In the last part of the sub-section we describe how client-server interactions are modelled by contracts (Figure 7).

A language for contracts L_C is given in Figure 4. As with session types it uses a denumerable set of recursion variables *Vars*, here lower case, but also presupposes a set *Act* of actions, ranged over by α , which contracts can perform; as we will see the special action \checkmark , which we assume is *not* in *Act*, will be used to indicate the fulfilment of a contract. Intuitively the contract $\alpha.\sigma$ performs the action α and then behaves like σ ; the sum $\sigma' + \sigma''$ is ready to behave either as σ' or as σ'' and the choice depends on the external environment. For this reason the operation $+$ is called *external* sum. The *internal* sum $\sigma' \oplus \sigma''$ represents a contract that can behave as σ' or as σ'' , and the choice is taken by the contract independently from the environment. Such a decision can

$\sigma ::=$	Contracts
NIL	<i>termination</i>
1	<i>success</i>
$\alpha.\sigma$	<i>action</i>
$\sigma + \sigma$	<i>external choice</i>
$\sigma \oplus \sigma$	<i>internal choice</i>
x	<i>contract variable</i>
$\mu x.\sigma$	<i>recursion</i>

Figure 4. Contract grammar.

be due for instance to an IF statement in the process implementing the contracts. The symbol NIL denotes an empty contract, which intuitively can never be fulfilled, while **1** denotes the contract that is always satisfied.

Recursive definitions are handled in much the same way as session types and so we do not spell out the details; we assume a definition of capture-avoiding substitution \mathbf{s} . Now we define the predicate \downarrow_{DPT} , the function *depth*, and the function UNFOLD as in the previous section.

The function *depth* is the least one that satisfies

$$\text{depth}(\sigma) = \begin{cases} 1 + \text{depth}(\sigma' \{ \sigma / x \}) & \text{if } \sigma = \mu x.\sigma', \\ 0 & \text{otherwise} \end{cases}$$

and the UNFOLD is the least function that satisfies:

$$\text{UNFOLD}(\sigma) = \begin{cases} \text{UNFOLD}(\sigma' \{ \sigma / x \}) & \text{if } \sigma = \mu x.\sigma', \\ \sigma & \text{otherwise} \end{cases}$$

These definitions are not arbitrary. As it happens, they let us prove Lemma 5.2, which, to our aim, is paramount (see Section 5).

Let \downarrow_{DPT} be the least fixed point of the rule functional defined on closed terms of L_C by the rules in Figure 2.

Similarly to what done in the previous section one can prove the following lemma.

Lemma 3.1. Let $\sigma \downarrow_{\text{DPT}}$. Then

- (i) the depth of σ is finite: $\text{depth}(\sigma) \in \mathbb{N}$
- (ii) the term $\text{UNFOLD}(\sigma)$ is defined and if σ is closed then $\text{UNFOLD}(\sigma)$ is closed.

Proof. The proof of part (ii) relies on (i) and is similar to the proof of Lemma 2.2. \square

An operational semantics for the closed terms of the language L_C is given in Figure 5. The judgements are of the form

$$\sigma \xrightarrow{\mu} \sigma', \quad \mu \in \text{Act}_{\tau\checkmark}$$

where we use $\text{Act}_{\tau\checkmark}$ as a shorthand for the set $\text{Act} \cup \{\tau, \checkmark\}$, and Act_{\checkmark} as shorthand for $\text{Act} \cup \{\checkmark\}$. The judgement $\sigma \xrightarrow{\alpha} \sigma'$, where $\alpha \in \text{Act}$ has the obvious meaning; $\sigma \xrightarrow{\tau} \sigma'$,

$$\begin{array}{c}
\frac{}{\mathbf{1} \xrightarrow{\checkmark} \text{NIL}} \text{ [A-OK]} \\
\\
\frac{}{\alpha.\sigma \xrightarrow{\alpha} \sigma} \text{ [A-PRE]} \qquad \frac{}{\mu x.\sigma \xrightarrow{\tau} \sigma \{ \mu x.\sigma / x \}} \text{ [A-UNF]} \\
\\
\frac{}{\sigma \oplus \rho \xrightarrow{\tau} \sigma} \text{ [A-IN-L]} \qquad \frac{}{\sigma \oplus \rho \xrightarrow{\tau} \rho} \text{ [A-IN-R]} \\
\\
\frac{\sigma \xrightarrow{\lambda} \sigma'}{\sigma + \rho \xrightarrow{\lambda} \sigma'} \text{ [R-EXT-L]} \qquad \frac{\rho \xrightarrow{\lambda} \rho'}{\sigma + \rho \xrightarrow{\lambda} \rho'} \text{ [R-EXT-R]} \\
\\
\frac{\sigma \xrightarrow{\tau} \sigma'}{\sigma + \rho \xrightarrow{\tau} \sigma' + \rho} \text{ [R-INT-L]} \qquad \frac{\rho \xrightarrow{\tau} \rho'}{\sigma + \rho \xrightarrow{\tau} \sigma + \rho'} \text{ [R-INT-R]}
\end{array}$$

Figure 5. Inference rules for the semantics of closed terms of $L_{\mathcal{C}}$, where $\lambda \in Act_{\checkmark}$

means that the contract σ is resolved to the contract σ' by some internal computation, while $\sigma \xrightarrow{\checkmark} \sigma'$ represents the reporting of the successful completion of a computation.

Let $\xrightarrow{\tau}^*$ denote the reflexive transitive closure of $\xrightarrow{\tau}$.

We are now ready to show two properties of UNFOLD. We will use them in the rest of the paper.

Lemma 3.2. Let σ be a contract.

- (i) If $\sigma \not\xrightarrow{\tau}$ then $\text{UNFOLD}(\sigma) = \sigma$
- (ii) $\sigma \xrightarrow{\tau}^* \text{UNFOLD}(\sigma)$

Proof. Part (i) is proved by structural induction on σ . We prove part (ii); the argument is by induction on $\text{depth}(\sigma)$. If $\text{depth}(\sigma) = 0$ then from the definition of depth it follows that $\sigma \neq \mu x.\sigma'$; by definition of UNFOLD then $\text{UNFOLD}(\sigma) = \sigma$. The reflexivity of $\xrightarrow{\tau}^*$ implies $\sigma \xrightarrow{\tau}^* \text{UNFOLD}(\sigma)$.

If $\text{depth}(\sigma) > 1$ then, due to the definition of depth , $\sigma = \mu x.\sigma'$. The definition of UNFOLD implies that $\text{UNFOLD}(\sigma) = \text{UNFOLD}(\sigma' \{ \sigma / x \})$, while the definition of depth implies $\text{depth}(\sigma) = 1 + \text{depth}(\sigma' \{ \sigma / x \})$, and therefore $\text{depth}(\sigma' \{ \sigma / x \})$ is smaller than $\text{depth}(\sigma)$. We are now allowed to use the inductive hypothesis on $\sigma' \{ \sigma / x \}$:

$$\sigma' \{ \sigma / x \} \xrightarrow{\tau}^* \text{UNFOLD}(\sigma' \{ \sigma / x \})$$

We use rule [A-UNF] (see Figure 5) to infer $\sigma \xrightarrow{\tau} \sigma' \{ \sigma / x \}$, and then the transitivity of the relation $\xrightarrow{\tau}^*$ to obtain

$$\sigma \xrightarrow{\tau}^* \text{UNFOLD}(\sigma' \{ \sigma / x \})$$

We already know that $\text{UNFOLD}(\sigma) = \text{UNFOLD}(\sigma' \{ \sigma / x \})$, and, by applying this equality

to the reduction sequence above, we get

$$\sigma \xrightarrow{\tau}^* \text{UNFOLD}(\sigma)$$

This concludes the proof. \square

Let

$$\begin{aligned} S(\sigma) &= \{ \sigma' \mid \sigma \xrightarrow{\mu} \sigma' \text{ for some } \mu \in \text{Act}_{\tau, \checkmark} \} \\ F(\sigma) &= \{ \sigma' \mid \sigma \xrightarrow{\tau}^* \sigma' \} \end{aligned}$$

One might think that $F(\sigma)$ is finite if and only if σ does not diverge. This is not the case.

Example 3.3. [Divergence and finite derivatives]

Consider the terms $\mu x. x$ and $\mu x. (\text{NIL} \oplus x)$. Both terms *diverge*, in the sense that they perform an infinite sequence of τ 's:

$$\mu x. x \xrightarrow{\tau} \mu x. x, \quad \mu x. (\text{NIL} \oplus x) \xrightarrow{\tau} \text{NIL} \oplus \mu x. (\text{NIL} \oplus x) \xrightarrow{\tau} \mu x. (\text{NIL} \oplus x)$$

On the other hand we have

$$F(\mu x. x) = S(\mu x. x) = \{ \mu x. x \}$$

and

$$\begin{aligned} F(\mu x. (\text{NIL} \oplus x)) &= \{ \mu x. (\text{NIL} \oplus x), \text{NIL} \oplus \mu x. (\text{NIL} \oplus x) \} \\ S(\mu x. (\text{NIL} \oplus x)) &= \{ \text{NIL} \oplus \mu x. (\text{NIL} \oplus x) \} \end{aligned}$$

\square

The set $F(\sigma)$ is not finite for every σ .

Example 3.4. [Infinite derivatives]

We show two terms σ, σ' such that $F(\sigma)$ and $F(\sigma')$ are infinite. Let $\sigma = \mu x. (\alpha x + x)$ and $\sigma' = \mu x. (\alpha. \text{NIL} + (x \oplus x))$. According to the rules in Figure 5 one can infer:

$$\sigma \xrightarrow{\tau} (\alpha. \sigma + \sigma) \xrightarrow{\tau} (\alpha. \sigma + (\alpha. \sigma + \sigma)) \xrightarrow{\tau} \dots$$

and

$$\sigma' \xrightarrow{\tau} \alpha. \text{NIL} + (\sigma' \oplus \sigma') \xrightarrow{\tau} \alpha. \text{NIL} + \sigma' \xrightarrow{\tau} \alpha. \text{NIL} + (\alpha. \text{NIL} + (\sigma' \oplus \sigma')) \xrightarrow{\tau} \dots$$

The root of the problem is that the inference rules [R-INT-L] and [R-INT-R] do not resolve the external sum. \square

On the other hand the finiteness of $S(\sigma)$ is easy to prove.

Lemma 3.5 (Finite branches).

The set $S(\sigma)$ is finite for every σ .

Proof. The proof is by structural induction. If $\sigma = \text{NIL}$ then plainly $S(\sigma) = \emptyset$. Otherwise we proceed as follows,

- if $\sigma = \alpha. \sigma'$ then the only applicable rule (see Figure 5) is [A-PRE], and so $S(\sigma) = \{\sigma'\}$;
- if $\sigma = \mu x. \sigma'$ then only rule [A-UNF] can be applied, so $S(\sigma) = \{\sigma' \{ \sigma / x \}\}$;

$$\bar{1} \quad \overline{\text{NIL}} \quad \overline{\alpha.\sigma} \quad \frac{\rho \ \sigma}{(\sigma \oplus \rho)} \quad \frac{\rho \ \sigma}{(\rho + \sigma)} \quad \frac{\sigma \{ \mu x. \sigma / x \}}{\mu x. \sigma}$$

Figure 6. Inference rules for \downarrow over closed terms of $L_{\mathcal{C}}$.

- if $\sigma = \sigma' + \sigma''$ then σ' and σ'' are both smaller than σ . The inductive hypothesis tells us that $S(\sigma')$ and $S(\sigma'')$ are finite. Rules [R-EXT-L], [R-EXT-R], [R-INT-L] and [R-INT-R] in Figure 5 ensure that $S(\sigma) = S(\sigma') \cup S(\sigma'')$. This implies that the cardinality of $S(\sigma)$ is finite;
- if $\sigma = \sigma' \oplus \sigma''$ the argument is alike the previous one.

□

Example 3.6. [Divergence and \downarrow_{DPT}]

Consider the term

$$\sigma = \mu x. (x \oplus x)$$

Using the rules in Figure 2 one can prove that $\sigma \downarrow_{\text{DPT}}$; consequently $\text{depth}(\sigma)$ is finite and $\text{UNFOLD}(\sigma)$ is well-defined; in particular $\text{depth}(\sigma) = 1$ and $\text{UNFOLD}(\sigma) = \sigma \oplus \sigma$. Note now that the term σ engages in an infinite sequence of internal moves

$$\sigma \xrightarrow{\tau} \sigma \oplus \sigma \xrightarrow{\tau} \sigma \oplus \sigma \oplus \sigma \xrightarrow{\tau} \sigma \oplus \sigma \oplus \sigma \oplus \sigma \xrightarrow{\tau} \dots$$

In other words the term σ diverges. Similarly, one can reason that terms as $\mu x. (\text{NIL} \oplus x)$ and $\mu x. (\alpha \oplus x)$ suffer the same issue. □

Throughout the paper we want to deal only with terms that do *not* diverge and with finite $F(\sigma)$. To isolate the σ 's that converge we use the predicate \downarrow . Formally, we define it as the least fixed point of the functional given by the inference rules in Figure 6. The predicate \downarrow is essentially a strengthened version of \downarrow_{DPT} .

Proposition 3.7. For every $\sigma \in L_{\mathcal{C}}$, $\sigma \downarrow$ implies $\sigma \downarrow_{\text{DPT}}$.

Proof. Straightforward from the definitions of the predicates. □

The predicate \downarrow is preserved by silent moves.

Lemma 3.8. Let $\sigma \downarrow$. If $\sigma \xrightarrow{\tau} \sigma'$ then $\sigma' \downarrow$.

Proof. The proof proceeds by rule induction on the derivation of $\sigma \xrightarrow{\tau} \sigma'$.

The only interesting case is when the silent move $\sigma \xrightarrow{\tau} \sigma'$ is inferred by using rule [R-INT-L] or rule [R-INT-R] (see Figure 5). Suppose rule [R-INT-L] was used. Then $\sigma = \sigma_1 + \sigma_2$, $\sigma' = \sigma'_1 + \sigma_2$, and the derivation is

$$\frac{\sigma_1 \xrightarrow{\tau} \sigma'_1}{\sigma_1 + \sigma_2 \xrightarrow{\tau} \sigma'_1 + \sigma_2}$$

We have to prove that $\sigma'_1 + \sigma_2 \downarrow$; the definition of \downarrow ensures that, to this aim, it is enough to show that (a) $\sigma'_1 \downarrow$ and that (b) $\sigma_2 \downarrow$ (see Figure 6). Point (b) follows from the

hypothesis: since $\sigma \downarrow$, the equality $\sigma = \sigma_1 + \sigma_2$ implies that $\sigma_1 \downarrow$ and that $\sigma_2 \downarrow$. The last fact is exactly point (b).

Now, by using the fact that $\sigma_1 \downarrow$, we prove point (a). The derivation shown above let us use rule induction; the lemma holds for σ_1 : if $\sigma_1 \downarrow$ and $\sigma_1 \xrightarrow{\tau} \hat{\sigma}_1$ then $\hat{\sigma}'_1 \downarrow$. We know that $\sigma_1 \downarrow$ and that $\sigma_1 \xrightarrow{\tau} \sigma'_1$, thus it must be $\sigma'_1 \downarrow$.

If rule [R-INT-R] was used the argument is similar. \square

We have also the converse.

Lemma 3.9. Let σ be a closed term of L_C . Then $\sigma \downarrow$ if and only if $\sigma \xrightarrow{\tau} \sigma'$ implies $\sigma' \downarrow$.

Proof. The *only if* side of the lemma is Lemma 3.8, so we are required to prove only that if $\sigma \xrightarrow{\tau} \sigma'$ implies $\sigma' \downarrow$, then $\sigma \downarrow$.

Let σ be a closed term of L_C such that $\sigma \xrightarrow{\tau} \sigma'$ implies $\sigma' \downarrow$. We have to show that $\sigma \downarrow$.

The proof is by structural induction on the form of σ ; the only interesting case is when the term σ is an external sum. In that case, $\sigma = \sigma_1 + \sigma_2$, so to prove that $\sigma_1 \downarrow$ it is enough to show that $\sigma_1 \downarrow$ and that $\sigma_2 \downarrow$.

We prove $\sigma_1 \downarrow$. The term σ_1 is a sub-term of σ , hence structural induction guarantees that the lemma holds for σ_1 : if $\sigma_1 \xrightarrow{\tau} \sigma'_1$ implies $\sigma'_1 \downarrow$, then $\sigma_1 \downarrow$. Assume that $\sigma_1 \xrightarrow{\tau} \sigma'_1$; thanks to the structure of σ we can derive

$$\frac{\sigma_1 \xrightarrow{\tau} \sigma'_1}{\sigma_1 + \sigma_2 \xrightarrow{\tau} \sigma'_1 + \sigma_2} \text{ [R-INT-L]}$$

Now the hypothesis of the lemma implies that $\sigma'_1 + \sigma_2 \downarrow$, so from the the definition of \downarrow it follows that $\sigma'_1 \downarrow$ and that $\sigma_2 \downarrow$.

We have shown that $\sigma_1 \xrightarrow{\tau} \sigma'_1$ implies $\sigma'_1 \downarrow$, so from the inductive hypothesis it follows that $\sigma_1 \downarrow$. Moreover, we have also shown that $\sigma_2 \downarrow$; we have proven that $\sigma \downarrow$.

The argument for rule [R-INT-L] is analogous, and left to the reader. \square

The predicate \downarrow let us give an inductive characterisation of the convergent terms.

Lemma 3.10 (Convergence).

Let σ be a closed term of the set L_C ; $\sigma \downarrow$ if and only if there exists a natural number k such that $\sigma \xrightarrow{\tau}^n \sigma'$ implies $n \leq k$.

Proof. The *only if* side is by rule induction on why $\sigma \downarrow$. We prove the *if* side, which states that if there exists a $k \in \mathbb{N}$ such that $\sigma \xrightarrow{\tau}^n \sigma'$ implies $n \leq k$, then $\sigma \downarrow$.

The argument is an induction on k . If $k = 0$ then σ cannot perform τ , and so it is either $\mathbf{1}$, NIL or a prefix $\alpha.\sigma'$. In all these cases we can easily infer $\sigma \downarrow$.

If $k > 0$ then σ performs a τ , so suppose $\sigma \xrightarrow{\tau} \sigma'$; it follows that $\sigma' \xrightarrow{\tau}^m \sigma''$ implies $m \leq k - 1$. This means that there exists a k' such that $\sigma' \xrightarrow{\tau}^m \sigma''$ implies $m \leq k'$. Since $k - 1 < k$, we can apply the inductive hypothesis to σ' ; from this application it follows that $\sigma' \downarrow$.

As yet, we have proven that $\sigma \xrightarrow{\tau} \sigma'$ implies $\sigma' \downarrow$; this allows us to apply Lemma 3.9, which ensures that $\sigma \downarrow$. \square

It is easy to see that a converging term σ has finite $F(\sigma)$.

Lemma 3.11. For every $\sigma \in L_C$ if $\sigma \downarrow$ then $F(\sigma)$ is finite.

Proof. We can prove it by rule induction on the derivation of $\sigma \downarrow$. \square

We are ready to define the set of contracts.

Definition 3.12. [Contracts]

Let \mathcal{C} denote the set of all terms σ of L_C which are closed and such that if $\sigma \xrightarrow{s}^* \sigma'$ for some $s \in Act^*$, then $\sigma \downarrow$. We refer to these terms as *contracts*. \square

Proposition 3.13. Let σ be a contract, then $UNFOLD(\sigma)$ is a well-defined contract.

Proof. We have to show that $UNFOLD(\sigma)$ is closed, and that it satisfies \downarrow . The first fact is part (ii) of Lemma 3.1. The second fact follows from part (ii) of Lemma 3.2 and Lemma 3.8. \square

Example 3.14. [e-vote, (Barbanera & de'Liguoro 2010, Laneve & Padovani 2008)]

$$\text{Ballot} = \mu x. ?\text{Login}.(!\text{Wrong}.x \oplus !\text{Ok}.(? \text{VoteA}.x + ? \text{VoteB}.x))$$

$$\text{Voter} = \mu x. !\text{Login}.(? \text{Wrong}.x + ? \text{Ok}.(! \text{VoteA}.1 \oplus ! \text{VoteB}.1))$$

A process offering contract *Ballot* implements a service for e-voting. Such a service lets a client log in. If the log in fails the services starts anew, while if the log in succeeds the two actions are offered to the environment, namely *VoteA* and *VoteB*.

The contract *Voter* is a recursive client for the protocol described by the contract *Ballot*. \square

Example 3.15. [e-commerce, (Bernardi et al. 2008)]

$$\begin{aligned} \text{Customer} = & !\text{Request}.(!\text{PayDebit}.\rho' \oplus \\ & !\text{PayCredit}.\rho' \oplus \\ & !\text{PayCash}.1) \end{aligned}$$

$$\rho' = !\text{Long}.? \text{Bool}.1$$

$$\begin{aligned} \text{Bank} = & \mu x. ?\text{Request}.(? \text{PayCredit}.? \text{Long}.! \text{Bool}.x + \\ & ? \text{PayDebit}.? \text{Long}.! \text{Bool}.x + \\ & ? \text{PayCash}.x) \end{aligned}$$

The contracts above describe the conversation that should take place between a client (which offers the contract *Customer*) and a bank (which has contract *Bank*) involved in an on-line payment. The conversation unfolds as follows: the *Customer* sends a request to the bank and afterwards it chooses the payment method; the choice is taken by an internal sum and this means that the decision of the *Customer* is independent from the environment (i.e., the *Bank* contract). If the *Customer* decides to pay by cash then no other action has to be taken; while if the payment is done by debit or credit card the *Customer* has to send the card number, this is represented by the output *!Long*. After the card number has been received the *Bank* answers with a boolean. Intuitively, this

represents the fact that the bank can approve or reject the payment. The Customer protocol finishes after such boolean has been received, while the Bank starts anew. \square

3.1.1. Client-Server interactions and the compliance relation Contracts are expressive enough to encode XML based languages such as WS-BPEL activities and WSCL diagrams (Castagna et al. 2009); moreover in (Carpinetti et al. 2006) it is shown how to assign contracts to a subset of CCS processes. Intuitively, if a process, such as a server, is assigned a contract σ then it guarantees to support the behaviour described in σ . The interaction between servers and clients can be described at the level of their contracts, by defining a binary operation $\rho \parallel \sigma$ between their contracts and describing the evolution of the contracts as they interact. This interacting semantics is given in Figure 7, where the judgements are of the form $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$. It presupposes a binary relation \bowtie on Act , where $\alpha \bowtie \beta$ means that the action α can synchronise with the action β . This relation can be instantiated in various ways depending on the particular set of actions Act .

Example 3.16. Suppose we take Act to be $\{a?, a! \mid a \in A\}$ where A is a set of communicating channels. Then define

$$\alpha \bowtie_i \beta \quad \text{whenever } \alpha = a?, \beta = a! \text{ or } \alpha = a!, \beta = a? \text{ for some } a \in A$$

The relation \bowtie_i represents synchronisation on channels.

We will use a more elaborate set of actions when interpreting session types as contracts. Recall from Section 2 the set of basic types BT , the set of labels \mathbb{L} used in session types, and Example 2.7. We can define Act to be the set

$$\{?b, !b \mid b \in BT\} \cup \{!?, !! \mid ! \in \mathbb{L}\}$$

with \bowtie_c determined by

$$\alpha \bowtie_c \beta \quad \text{whenever } \begin{cases} \alpha = ?b, \beta = !b' & b' \preceq_g b \\ \alpha = !b, \beta = ?b' & b \preceq_g b' \\ \alpha = ?!, \beta = !! \\ \alpha = !!, \beta = ?! \end{cases}$$

Using the basic sub-typing relation depicted in Figure 3 the following examples should be clear:

- (i) $?Num \bowtie_c !Int$: a contract that can read a datum of type Num can read a datum of type Int because $Int \preceq_g Num$.
- (ii) $?Int \not\bowtie_c !Num$: conversely a contract ready to read a datum of type Int cannot read a datum of type Num because $Num \not\preceq_g Int$.
- (iii) $?Random \bowtie_c !Bool$: as in point (i), $Bool \preceq_g Random$ hence an interaction between the actions $?Random$ and $!Bool$ is safe.

\square

Having described how interactions between clients and servers affect their contracts, let us describe, by means of a relation, when a client (guaranteeing a) contract ρ can

$$\frac{\rho \xrightarrow{\tau} \rho'}{\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma} \text{ [P-SIL-L]} \quad \frac{\sigma \xrightarrow{\tau} \sigma'}{\rho \parallel \sigma \xrightarrow{\tau} \rho \parallel \sigma'} \text{ [P-SIL-R]}$$

$$\frac{\rho \xrightarrow{\alpha} \rho' \quad \sigma \xrightarrow{\beta} \sigma'}{\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'} \alpha \bowtie \beta \text{ [P-SYNCH]}$$

Figure 7. Inference rules for contract interaction.

safely interact with a server (guaranteeing a) contract σ . Indeed, we shall formalise the meaning of “safely”.

The central notion is that of *compliance* between contracts. This is defined co-inductively and uses the predicate on contracts $\rho \xrightarrow{\checkmark}$ which intuitively means that the contract ρ has already been satisfied. Our definition is a variation on that of *compliance* in (Laneve & Padovani 2007, 2008, Padovani 2010).

Definition 3.17. [Compliance relation]

Let $\mathcal{F}_{\dashv} : \mathcal{P}(\mathcal{C}^2) \rightarrow \mathcal{P}(\mathcal{C}^2)$ be the function defined so that $(\rho, \sigma) \in \mathcal{F}_{\dashv}(\mathcal{R})$ whenever both the following hold:

- (i) if $\rho \parallel \sigma \not\xrightarrow{\tau}$ then $\rho \xrightarrow{\checkmark}$
- (ii) if $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$ then $(\rho', \sigma') \in \mathcal{R}$

If $\mathcal{R} \subseteq \mathcal{F}_{\dashv}(\mathcal{R})$ then we say that \mathcal{R} is a co-inductive *compliance* relation. Let \dashv denote the greatest solution of the equation $X = \mathcal{F}_{\dashv}(X)$; formally,

$$\dashv = \nu \mathcal{F}_{\dashv}$$

We call \dashv the *compliance* relation. If $\rho \dashv \sigma$ we say that the contract ρ *complies with* the contract σ . \square

Notice that there is an asymmetry in the relation $\rho \dashv \sigma$; the intention is that any client running contract ρ when interacting with a server running contract σ will be satisfied, in the sense that either the interaction between client and server will go on indefinitely, or, if the interaction gets stuck, the client will end on its own in a state in which it is satisfied, $\rho \xrightarrow{\checkmark}$.

Example 3.18. [Compliance and divergent terms]

In order that the relation \dashv captures the intuition described above, it is crucial that \mathcal{C} contains no divergent terms. Had we admitted them, then for every ρ the relation

$$\{(\rho', \mu x.x) \mid \rho \xrightarrow{\tau}^* \rho'\}$$

would have been a perfectly fine co-inductive compliance. Note, though, that the client contract ρ is by no means satisfied by the server. \square

Example 3.19. [Compliance and \checkmark]

According to our definition of compliance, the client need not ever perform \checkmark . For

example, suppose $\alpha \bowtie \beta$ and consider the ensuing set

$$\{(\mu x. \alpha.x, \mu y. \beta.y)\}$$

This set is a co-inductive compliance relation, and the client contract, $\mu x. \alpha.x$, does not perform \checkmark at all. \square

Example 3.20. The fact that **NIL** cannot be satisfied is formally expressed by the fact that

$$\mathbf{NIL} \not\bowtie \sigma$$

for every contract σ . On the other hand **1** is always satisfied because we can prove that for every contract σ we have the following,

$$\mathbf{1} \dashv \sigma$$

Suppose σ is a contract which cannot interact with the action α ; by this we mean that $\sigma \xrightarrow{\tau}^* \xrightarrow{\beta}$ implies $\beta \not\bowtie \alpha$. Then

$$\mathbf{1} + \alpha.\rho \dashv \sigma$$

for every ρ , because ρ is guarded by an action that can never take place.

Referring to Example 3.14, it is routine work to check that the following relation is a co-inductive compliance.

$$\begin{aligned} \mathcal{R} = \{ & (\text{Voter}, \text{Ballot}), \\ & (?Wrong.Voter + ?Ok.(!VoteA.1 \oplus !VoteB.1)), \\ & !Wrong.Ballot \oplus !Ok.(?VoteA.Ballot + ?VoteB.Ballot)), \\ & (?Ok.(!VoteA.1 \oplus !VoteB.1), !Ok.(?VoteA.Ballot + ?VoteB.Ballot)), \\ & (!VoteA.1 \oplus !VoteB.1, ?VoteA.Ballot + ?VoteB.Ballot), \\ & (!VoteA.1, ?VoteA.Ballot + ?VoteB.Ballot), \\ & (!VoteB.1, ?VoteA.Ballot + ?VoteB.Ballot), \\ & (\mathbf{1}, \text{Ballot}) \} \end{aligned}$$

\square

The previous example shows that on the client-side, the contracts **1** and **NIL** have opposite meanings, as the former is always satisfied, while the latter is never satisfied; thus a client whose contract is **1** is not equivalent to a client whose contract is **NIL**. On the server-side the situation is different; a server with contract **1** is equivalent to a server with contract **NIL**. We prove this fact.

Proposition 3.21. For every contract ρ , $\rho \dashv \mathbf{1}$ if and only if $\rho \dashv \mathbf{NIL}$.

Proof. Suppose $\rho \dashv \mathbf{1}$; this means that there exists a co-inductive compliance \mathcal{R} that contains $(\rho, \mathbf{1})$. Since **1** offers no interaction, the contract ρ enjoys the two properties which follow,

- (i) if $\rho \xrightarrow{\tau}$ then $\rho \xrightarrow{\checkmark}$
- (ii) if $\rho \xrightarrow{\tau} \rho'$ then $(\rho', \mathbf{1}) \in \mathcal{R}$

Knowing (i), it is straightforward to show that

$$\mathcal{R}' = \{ (\rho', \text{NIL}) \mid \rho \xrightarrow{\tau}^* \rho', \rho \dashv \mathbf{1} \}$$

is a co-inductive compliance. A symmetrical argument can be used to show that also

$$\mathcal{R}' = \{ (\rho', \mathbf{1}) \mid \rho \xrightarrow{\tau}^* \rho', \rho \dashv \text{NIL} \}$$

is a co-inductive compliance. \square

The following properties of the compliance relation will be useful later in the paper.

Lemma 3.22. Let ρ, σ_1 , and σ_2 be contracts. The following hold:

- (i) if $\rho \dashv \sigma_1, \rho \dashv \sigma_2$ then $\rho \dashv \sigma_1 \oplus \sigma_2$
- (ii) if $\rho_1 \dashv \sigma, \rho_2 \dashv \sigma$ then $\rho_1 \oplus \rho_2 \dashv \sigma$

Proof. As an example we outline the proof of (i). Let \mathcal{R} be the relation defined by

$$\mathcal{R} = \{ (\rho, \sigma) \mid \rho \dashv \sigma \text{ or } \sigma = \sigma_1 \oplus \sigma_2 \text{ where } \rho \dashv \sigma_1 \text{ and } \rho \dashv \sigma_2 \}$$

It is straightforward to show that \mathcal{R} is a compliance relation, from which the result follows. \square

Proposition 3.23. For all contracts ρ, σ , we have the following

- (a) if $\rho \dashv \sigma$ then $\rho \dashv \text{UNFOLD}(\sigma)$
- (b) if $\rho \dashv \sigma$ then $\text{UNFOLD}(\rho) \dashv \sigma$

Proof. Both follow in a straightforward manner from part (ii) of Lemma 3.2 and part (ii) of Definition 3.17. \square

The converse is also true:

Proposition 3.24. For all contracts ρ, σ , we have the following

- (a) if $\rho \dashv \text{UNFOLD}(\sigma)$ then $\rho \dashv \sigma$
- (b) if $\text{UNFOLD}(\rho) \dashv \sigma$ then $\rho \dashv \sigma$

Proof. Let us look at the proof of (a). Let

$$\mathcal{R} = \{ (\rho, \sigma) \mid \rho \dashv \sigma \text{ or } \rho \dashv \text{UNFOLD}(\sigma) \}$$

The result will follow if we can prove that \mathcal{R} is a co-inductive compliance relation, as given in Definition 3.17.

- (a) Suppose $\rho \parallel \sigma \xrightarrow{\tau}$. If $\rho \dashv \sigma$ then by definition $\rho \xrightarrow{\check{\tau}}$. Otherwise

$$\rho \dashv \text{UNFOLD}(\sigma)$$

Note that $\sigma \xrightarrow{\tau}$ and therefore by part (i) of Lemma 3.2 it follows that $\text{UNFOLD}(\sigma) = \sigma$, which means, since now $\rho \dashv \sigma, \rho \xrightarrow{\check{\tau}}$.

- (b) Suppose $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$. We have to show $(\rho', \sigma') \in \mathcal{R}$, which is obvious if $\rho \dashv \sigma$. On the other hand if $\rho \dashv \text{UNFOLD}(\sigma)$ there are three cases, depending on the inference of the action $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$. If the action is due to a silent move of ρ , the result

follows from part (ii) of Definition 3.17. In the other cases the result will follow by an application of part (i) and part (ii) of Lemma 3.2; and of part (ii) of Definition 3.17. \square

3.2. The server pre-order

In this subsection we show how to compare servers in terms of their ability to satisfy clients; once more this is done in terms of their respective contracts.

Definition 3.25. [Server pre-order]

We write $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2$ whenever for every contract ρ , $\rho \dashv \sigma_1$ implies $\rho \dashv \sigma_2$. \square

This provides us with a natural subsumption-like pre-order between server-side contracts. For if $\sigma \sqsubseteq_{\text{SRV}} \sigma'$ then we are assured every client satisfied by a server running the contract σ is by definition also satisfied by a server running the contract σ' .

One consequence of Proposition 3.23 and Proposition 3.24 is that

$$\llbracket \text{UNFOLD}(\sigma) \rrbracket_{\text{SRV}} = \llbracket \sigma \rrbracket_{\text{SRV}}$$

and therefore when reasoning about the server pre-order we can work up to unfolding.

Notation A In the rest of the paper we will at times use the symbols \sum and \oplus to write contracts, for instance

$$\begin{array}{ll} \sum_{i \in [1; m]} \sigma_i & \text{in place of } \sigma_1 + \sigma_2 + \dots + \sigma_m \\ \bigoplus_{i \in [1; n]} \sigma_i & \text{in place of } \sigma_1 \oplus \sigma_2 \oplus \dots \oplus \sigma_n \end{array}$$

This is justified to the fact that

$$\begin{array}{ll} \sigma \oplus \rho =_{\text{SRV}} \rho \oplus \sigma & \sigma \oplus (\sigma' \oplus \sigma'') =_{\text{SRV}} (\sigma \oplus \sigma') \oplus \sigma'' \\ \sigma + \rho =_{\text{SRV}} \rho + \sigma & \sigma + (\sigma' + \sigma'') =_{\text{SRV}} (\sigma + \sigma') + \sigma'' \end{array}$$

where $=_{\text{SRV}}$ is the equivalence relation given in the obvious way by \sqsubseteq_{SRV} .

3.2.1. Co-inductive characterisation Here we give a co-inductive characterisation of the server pre-order \sqsubseteq_{SRV} ; this is based on a number of semantic properties of contracts, which we outline in the following lemmas.

Lemma 3.26. If $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2$ and $\sigma_2 \xrightarrow{\tau} \sigma'_2$ then $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma'_2$.

Proof. Suppose $\rho \dashv \sigma_1$, where $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2$ and $\sigma_2 \xrightarrow{\tau} \sigma'_2$; we have to show $\rho \dashv \sigma'_2$.

We know \dashv is a co-inductive compliance relation, and also that $\rho \parallel \sigma_2 \xrightarrow{\tau} \rho \parallel \sigma'_2$. So by part (ii) of Definition 3.17 the required $\rho \dashv \sigma'_2$ follows. \square

The next property involves the *acceptance sets* of contracts. For any $R \subseteq \text{Act}$ let us write $\sigma \downarrow R$ whenever $\sigma \xrightarrow{\tau}$ and $R = \{ \alpha \in \text{Act} \mid \sigma \xrightarrow{\alpha} \}$. These sets R, R', \dots are called *initials* (Eshuis & Fokkinga 2002) or *ready sets* (Laneve & Padovani 2007, 2008). If $\sigma \downarrow R$ we say that σ *converges to* R ; indeed, in our presentation only *stuck* states have ready sets.




Figure 8. An example of acceptance sets (see Example 3.28).

Definition 3.27. [Acceptance set]

For every contract σ let

$$\text{ACC}(\sigma) = \{ \mathbf{R} \mid \sigma \xrightarrow{\tau}^* \sigma', \sigma' \downarrow \mathbf{R} \}$$

We say that $\text{ACC}(\sigma)$ is the *acceptance sets* of σ . \square

One sees easily that $\text{NIL} \downarrow \emptyset$ and $\mathbf{1} \downarrow \emptyset$ (as $\checkmark \notin \text{Acts}$), so

$$\text{ACC}(\text{NIL}) = \text{ACC}(\mathbf{1}) = \{ \emptyset \}$$

Example 3.28. In Figure 8 we depict the LTS's and the acceptance sets of the ensuing contracts

$$\sigma = \alpha.\sigma_1 \oplus (\beta.\sigma_2 \oplus \gamma.\sigma_3) \quad \sigma' = (\alpha.\sigma'_1 \oplus \beta.\sigma'_2) \oplus \gamma.\sigma'_3$$

\square

Proposition 3.29. Let $\sigma \in \mathcal{C}$. The ensuing statements are true,

- (a) if $\sigma \downarrow \mathbf{R}$ then \mathbf{R} is finite
- (b) the set $\text{ACC}(\sigma)$ is finite
- (c) the set $\text{ACC}(\sigma)$ is non-empty

Proof. Part (a) follows from the fact that external sums contain a finite amount of summands. Part (b) follows from Lemma 3.11. Part (c) follows from the fact that if $\sigma \downarrow$ then σ has stuck derivatives. \square

Example 3.30. [Divergent terms and acceptance sets]

Here we show the acceptance set of a divergent term. Let $\sigma = \mu x. (\alpha.x + x)$. For every ready set \mathbf{R} we have $\sigma \not\downarrow \mathbf{R}$ because $\sigma \xrightarrow{\tau}$, and the same is true for the derivative $\alpha.\sigma + \sigma$, because it also performs a τ :

$$\sigma \xrightarrow{\tau} \alpha.\sigma + \sigma \xrightarrow{\tau} \dots$$

We thus conclude that $\text{ACC}(\sigma) = \emptyset$. Indeed, in the proof of part (c) of Proposition 3.29 the crucial hypothesis is $\sigma \downarrow$. \square

Individual acceptance sets are compared by their ability to offer interactions. We will write $R \sqsubseteq S$ whenever for every $\alpha_R \in R$ and every β such that $\beta \bowtie \alpha_R$ there exists some action $\alpha_S \in S$ such that $\beta \bowtie \alpha_S$ also. The precise meaning of this pre-order actually depends on the instantiation of the interaction relation \bowtie .

Example 3.31. Suppose that Act be the set $\{a?, a! \mid a \in A\}$ and \bowtie_i is defined as in Example 3.16. One can check that $R \sqsubseteq S$ if and only if $R \subseteq S$.

Suppose that

$$Act = \{?b, !b \mid b \in BT\} \cup \{!?, !! \mid ! \in \mathbb{L}\}$$

and \bowtie_c as in Example 3.16. It turns out that $R \sqsubseteq S$ whenever the following conditions are true,

- $!1 \in R$ implies $!1 \in S$ for every $!1 \in \mathbb{L}$
- $?1 \in R$ implies $?1 \in S$ for every $!1 \in \mathbb{L}$
- $?b_R \in R$ implies $?b_S \in S$ for some type b_S such that $b_S \preceq_g b_R$
- $!b_R \in R$ implies $!b_S \in S$ for some type b_S such that $b_R \preceq_g b_S$

□

Lemma 3.32. Suppose $\sigma_1 \sqsubseteq_{SRV} \sigma_2$ and $\sigma_2 \downarrow R$, then there is some $R' \in ACC(\sigma_1)$ such that $R' \sqsubseteq R$.

Proof. This proof proceeds by contradiction. To establish the contradiction we construct a contract ρ such that

- (a) $\rho \dashv \sigma_1$
- (b) $\rho \not\vdash \sigma_2$

Suppose there is no $R' \in ACC(\sigma_1)$ such that $R' \sqsubseteq R$; thanks to part (c) of Proposition 3.29 this fact cannot be true because $ACC(\sigma_1)$ is empty. Again by Proposition 3.29 we know $ACC(\sigma_1)$ to be finite, so let R_1, \dots, R_n be all the elements in $ACC(\sigma_1)$. From the hypothesis there are $\alpha_i \in R_i$ and $\beta_i \bowtie \alpha_i$ such that $\beta_i \not\bowtie \alpha$ whenever $\alpha \in R'$. Let the contract ρ be defined as $\beta_1.\mathbf{1} + \dots + \beta_n.\mathbf{1}$.

First notice that (b) above is true: since $\sigma_2 \downarrow R$, $\rho \parallel \sigma_2 \xrightarrow{\tau} \rho \parallel \sigma_2'$ such that $\rho \parallel \sigma_2' \not\vdash$ and $\rho \not\vdash$. This means that (ρ, σ_2) cannot be contained in any co-inductive compliance relation.

To establish (a) above it is sufficient to prove that

$$\mathcal{R} = \{(\rho, \sigma_1') \mid \sigma_1 \xrightarrow{\tau} \sigma_1'\}$$

is a co-inductive compliance relation, which is relatively straightforward. □

Yet another property of \sqsubseteq_{SRV} will be necessary to give its co-inductive characterisation; to prove this property, one more notion is in order.

Definition 3.33. [After]

For any action $\alpha \in Act$ and contract σ , let $(\sigma \text{ AFTER } \alpha)$ be the set

$$\{\sigma' \mid \sigma \xrightarrow{\tau} \sigma' \xrightarrow{\beta} \sigma' \xrightarrow{\tau} \sigma', \text{ where } \alpha \bowtie \beta\}$$

□

The proof of the following property is immediate.

Proposition 3.34. Let σ be a contract, then for every $\alpha \in Act$ the set $(\sigma \text{ AFTER } \alpha)$ is finite.

Proof. This follows because for every contract $S(\sigma)$ and $F(\sigma)$ are finite; see Lemma 3.5 and Lemma 3.11. \square

Note that in general the set $(\sigma \text{ AFTER } \alpha)$ may be empty. When it is non-empty we use $\bigoplus(\sigma \text{ AFTER } \alpha)$ to denote the internal sum of all its elements.

Lemma 3.35. Suppose $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2$ and $\sigma_2 \xrightarrow{\beta} \sigma'_2$. Then whenever $\alpha \bowtie \beta$,

- (a) the set $(\sigma_1 \text{ AFTER } \alpha)$ is non-empty
- (b) the contract $\bigoplus(\sigma_1 \text{ AFTER } \alpha)$ is smaller than σ'_2 . Formally,

$$\bigoplus(\sigma_1 \text{ AFTER } \alpha) \sqsubseteq_{\text{SRV}} \sigma'_2$$

Proof. To prove of part (a) consider the contract $\rho = \mathbf{1} + \alpha.\text{NIL}$. Since

$$\rho \parallel \sigma_2 \xrightarrow{\tau} \text{NIL} \parallel \sigma'_2$$

it follows that (ρ, σ'_2) cannot be in any co-inductive compliance relation, hence $\rho \not\preceq \sigma_2$. Therefore, from $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2$, we have that $\rho \not\preceq \sigma_1$.

But because of the construction of ρ this can only be the case if $(\sigma_1 \text{ AFTER } \alpha)$ is non-empty. More specifically, if it was empty we could construct a simple co-inductive compliance relation containing the pair (ρ, σ_1) .

Now consider part (b). Suppose $\rho \dashv \bigoplus(\sigma \text{ AFTER } \alpha)$; we have to show $\rho \dashv \sigma'_2$. To do so consider the contract $\rho' = \mathbf{1} + \alpha.\rho$. Suppose we could establish

$$\rho' \dashv \sigma_1 \tag{1}$$

Because $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2$ this would mean that $\rho' \dashv \sigma_2$, from which the required $\rho \dashv \sigma'_2$ follows, by part (ii) of Definition 3.17.

It remains to prove (1) above. Let

$$\mathcal{R} = \{(\rho, \sigma') \mid \rho \dashv \sigma', \sigma' \in \mathcal{C}\} \cup \{(\rho', \sigma'_1) \mid \sigma_1 \xrightarrow{\tau} \sigma'_1\}$$

Then, because $\rho \dashv \bigoplus(\sigma \text{ AFTER } \alpha)$, it is easy to establish that \mathcal{R} is a co-inductive compliance relation. \square

We have now assembled all the required properties for our co-inductive characterisation of the server pre-order.

Definition 3.36. [Semantic sub-server relation]

Let $\mathcal{F}_{\preceq_{\text{SRV}}} : \mathcal{P}(\mathcal{C}^2) \rightarrow \mathcal{P}(\mathcal{C}^2)$ be the function defined so that $(\sigma_1, \sigma_2) \in \mathcal{F}_{\preceq_{\text{SRV}}}(\mathcal{R})$ if and only if the following conditions hold:

- (i) if $\sigma_2 \xrightarrow{\tau} \sigma'_2$ then $(\sigma_1, \sigma'_2) \in \mathcal{R}$
- (ii) for every $R \in \text{ACC}(\sigma_2)$, $\sigma_2 \downarrow R$ implies $R' \sqsubseteq R$ for some $R' \in \text{ACC}(\sigma_1)$
- (iii) if $\sigma_2 \xrightarrow{\beta} \sigma'_2$ and $\alpha \bowtie \beta$ then $(\sigma_1 \text{ AFTER } \alpha) \neq \emptyset$ and $\bigoplus(\sigma_1 \text{ AFTER } \alpha) \mathcal{R} \sigma'_2$

If $\mathcal{R} \subseteq \mathcal{F}_{\preceq_{\text{SRV}}}(\mathcal{R})$ then we say that \mathcal{R} is a co-inductive *semantic sub-server* relation. Let \preceq_{SRV} denote the greatest solution of the equation $X = \mathcal{F}_{\preceq_{\text{SRV}}}(X)$; formally,

$$\preceq_{\text{SRV}} = \nu \mathcal{F}_{\preceq_{\text{SRV}}}$$

We call \preceq_{SRV} the *semantic sub-server* relation. \square

Proposition 3.37. If $\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2$ then $\sigma_1 \preceq_{\text{SRV}} \sigma_2$.

Proof. It is sufficient to prove that the relation \sqsubseteq_{SRV} is a semantic sub-server relation; this is straightforward in view of the last three lemmas. \square

We also have the converse.

Theorem 3.38. [Co-inductive characterisation]

The server pre-order is the greatest semantic sub-server relation.

Proof. We are required to prove that for all contracts σ_1 and σ_2 , the following is true,

$$\sigma_1 \sqsubseteq_{\text{SRV}} \sigma_2 \text{ if and only if } \sigma_1 \preceq_{\text{SRV}} \sigma_2$$

Because of the previous proposition it is sufficient to prove that $\sigma_1 \preceq_{\text{SRV}} \sigma_2$ and $\rho \dashv \sigma_1$ implies $\rho \dashv \sigma_2$. This will follow if we can show that the relation

$$\mathcal{R} = \{ (\rho, \sigma) \mid \rho \dashv \sigma_1 \text{ for some } \sigma_1 \text{ such that } \sigma_1 \preceq_{\text{SRV}} \sigma \}$$

is a co-inductive compliance relation.

Suppose $\rho \mathcal{R} \sigma$. By Definition 3.17 we are required to show two things; namely that

- (a) if $\rho \parallel \sigma \xrightarrow{\tau}$ then $\rho \xrightarrow{\check{\tau}}$
- (b) if $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$ then $(\rho', \sigma') \in \mathcal{R}$

By the definition of \mathcal{R} we know that there is some contract σ_1 such that $\sigma_1 \preceq_{\text{SRV}} \sigma$ and $\rho \dashv \sigma_1$.

We prove the point (a). If $\rho \parallel \sigma \xrightarrow{\tau}$ then $\rho \xrightarrow{\tau}, \sigma \xrightarrow{\tau}$; in addition, the two contracts cannot interact, that is $\rho \xrightarrow{\alpha}$ and $\sigma \xrightarrow{\beta}$ implies $\alpha \not\bowtie \beta$. Since ρ and σ are stable both $\text{ACC}(\rho)$ and $\text{ACC}(\sigma)$ contain exactly one set each, say R and S respectively. Then rephrasing the above remark we know

$$\alpha \in R, \beta \in S \text{ implies } \alpha \not\bowtie \beta \tag{2}$$

Since $\sigma_1 \preceq_{\text{SRV}} \sigma$, by part (ii) of Definition 3.36 $\sigma_1 \xrightarrow{\tau}^* \sigma'_1$ for some σ'_1 such that $\sigma'_1 \downarrow S'$ and $S' \sqsubseteq S$. One can use (2) above to show that this means

$$\alpha \in R, \beta \in S' \text{ implies } \alpha \not\bowtie \beta$$

Also, since $\rho \dashv \sigma_1$ and $\sigma_1 \xrightarrow{\tau}^* \sigma'_1$, part (ii) of Definition 3.17 implies $\rho \dashv \sigma'_1$. But $\rho \parallel \sigma'_1 \xrightarrow{\tau}$ and therefore we have the required $\rho \xrightarrow{\check{\tau}}$.

To prove point (b) above, we have to show that if $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$ there exists a $\hat{\sigma}$ such that

$$\rho' \dashv \hat{\sigma}, \quad \hat{\sigma} \preceq_{\text{SRV}} \sigma'$$

We proceed by case analysis on the rule used to infer $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$. There are three possibilities: first suppose the inference rule [P-SIL-L] from Figure 7 is used; the premises of the rule imply that $\rho \xrightarrow{\tau} \rho'$, and so $\sigma' = \sigma$. In this case the required $\hat{\sigma}$ is σ_1 ; the definition of \mathcal{R} gives $\sigma_1 \preceq_{\text{SRV}} \sigma$ and part (ii) of Definition 3.17 gives $\rho' \dashv \sigma_1$.

The case when the rule [P-SIL-R] is used is similar; choosing the required $\hat{\sigma}$ to be σ_1 again is justified by point (i) of Definition 3.36.

Finally suppose [P-SYNCH] is employed. Now we know that

$$\rho \xrightarrow{\delta} \rho', \quad \sigma \xrightarrow{\beta} \sigma', \quad \delta \bowtie \beta$$

In this case we show that the required $\hat{\sigma}$ is $\bigoplus(\sigma_1 \text{ AFTER } \delta)$. Part (iii) of Definition 3.36 implies that $\bigoplus(\sigma_1 \text{ AFTER } \delta) \preceq_{\text{SRV}} \sigma'$ and thus it suffices to show that $\rho' \dashv \bigoplus(\sigma_1 \text{ AFTER } \delta)$.

The set $\sigma_1 \text{ AFTER } \delta$ is finite and therefore by Lemma 3.22 it is sufficient to prove $\rho' \dashv \sigma''$ for every $\sigma'' \in (\sigma_1 \text{ AFTER } \delta)$.

For such a σ'' we can derive the transition

$$\rho \parallel \sigma_1 \xrightarrow{\tau} \rho' \parallel \sigma'',$$

where one of the reductions is due to the interaction through δ . From part (ii) of Definition 3.17 it follows that $\rho \dashv \sigma''$. \square

3.3. Must testing

The compliance relation between contracts, Definition 3.17, has much in common with the idea of testing from (Nicola & Hennessy 1984). Here we explain the relationship. We recall the definition of MUST testing, and explain how it differs from the compliance relation. Despite this difference we then go on to show that the testing pre-order it induces on contracts coincides with the server pre-order.

For every contract ρ and σ a sequence of reductions

$$\rho \parallel \sigma \xrightarrow{\tau} \rho_1 \parallel \sigma_1 \xrightarrow{\tau} \rho_2 \parallel \sigma_2 \longrightarrow \dots$$

is called a *computation* of $\rho \parallel \sigma$ and each derivative $\rho_i \parallel \sigma_i$ is a *state* of the computation. Intuitively, viewing ρ as a test we say that the state $\rho_i \parallel \sigma_i$ is *successful* if $\rho_i \parallel \sigma_i \xrightarrow{\checkmark}$; and a computation is successful if it contains a successful state.

Example 3.39. For every σ all the computations of

$$\mathbf{1} + ?\mathbf{1}_1.\text{NIL} \parallel \sigma$$

are successful, because in the first state we have $\mathbf{1} + ?\mathbf{1}_1.\text{NIL} \xrightarrow{\checkmark}$. On the other hand, suppose that a contract ρ does not perform \checkmark and neither do its derivatives. Then no computation of $\rho \parallel \sigma$ is successful. \square

A computation is *maximal* if either

- (i) it is infinite, or
- (ii) it is finite and the last state is stuck, that is has the form $\rho_k \parallel \sigma_k$ where $\rho_k \parallel \sigma_k \not\xrightarrow{\tau}$

Definition 3.40. [Must testing]

For all contracts ρ, σ we write $\sigma \text{ MUST } \rho$ if all the maximal computations of $\rho \parallel \sigma$ are successful. \square

The notion of a client contract complying with a server contract differs in two ways from that of a server contract passing a client contract viewed as a test.

Example 3.41. One difference between MUST and \dashv is what happens after a contract has passed a test, that is the test has reached a successfully state; the subsequent computation is disregarded by MUST , whereas the compliance relation has to hold for *all* the states in a computation.

As an example consider $\sigma = ?\text{Real.NIL}$ and $\rho = \mathbf{1} + !\text{Int.NIL}$. Clearly $\rho \dashv \sigma$, and therefore $\sigma \text{ MUST } \rho$, because each maximal computation of $\rho \parallel \sigma$ begins in a successful state. However $\rho \not\vdash \sigma$ because $\rho \parallel \sigma \xrightarrow{\tau} \text{NIL} \parallel \text{NIL}$. \square

Example 3.42. The second difference is that the compliance relation does not require the testing contract to ever report success, provided that the communication between the contracts can continue indefinitely. As an example consider the following contracts

$$\sigma = \mu x. !\text{Bool}.x, \quad \rho = (\mu y. ?\text{Rnd}.y) + !\text{Int}.1$$

Plainly, one sees that $\rho \dashv \sigma$, and, therefore, $\rho \parallel \sigma$ is not a successful state. The only computation of $\rho \parallel \sigma$ is the infinite loop $\rho \parallel \sigma \xrightarrow{\tau} \rho \parallel \sigma$, and therefore $\rho \dashv \sigma$ holds; on the other hand $\sigma \text{ MUST } \rho$ is false. Example 3.19 contains an even simpler instance of the difference between the relation \dashv and the relation MUST . \square

The MUST relation can be used to define a well known pre-order:

Definition 3.43. [Must pre-order (Nicola & Hennessy 1984)]

Let σ_1, σ_2 be contracts; we write $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2$ if and only if for every ρ , $\sigma_1 \text{ MUST } \rho$ implies $\sigma_2 \text{ MUST } \rho$. \square

Notation B As we discussed in paragraph Notation A of Section 3.2, we use \oplus and \sum also when reasoning on the must pre-order. Formally, this is justified by the ensuing equalities,

$$\begin{aligned} \sigma \oplus \rho &=_{\text{MUST}} \rho \oplus \sigma & \sigma \oplus (\sigma' \oplus \sigma'') &=_{\text{MUST}} (\sigma \oplus \sigma') \oplus \sigma'' \\ \sigma + \rho &=_{\text{MUST}} \rho + \sigma & \sigma + (\sigma' + \sigma'') &=_{\text{MUST}} (\sigma + \sigma') + \sigma'' \end{aligned}$$

where $=_{\text{MUST}}$ is the equivalence relation given in the obvious manner by $\sqsubseteq_{\text{MUST}}$.

Notwithstanding the differences between must testing and compliance exposed in Example 3.41 and Example 3.42, it turns out that the server pre-order \sqsubseteq_{SRV} and the MUST pre-order $\sqsubseteq_{\text{MUST}}$ coincide (Corollary 3.51).

First, in a series of lemmas, we show that $\sqsubseteq_{\text{MUST}}$ satisfies the three defining properties of the semantic sub-server relation (Definition 3.36).

Lemma 3.44. Let $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2$. If $\sigma_2 \xrightarrow{\tau} \sigma'_2$ then $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma'_2$.

Proof. Take a contract ρ such that $\sigma_1 \text{ MUST } \rho$ and a maximal computation C performed by $\sigma'_2 \parallel \rho$. It is easy to see that a maximal computation from $\sigma_2 \parallel \rho$ can be obtained by prefixing C with the move $\sigma_2 \parallel \rho \xrightarrow{\tau} \sigma'_2 \parallel \rho$.

Since $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2$ it follows that this extended computation must be successful. However this implies that C itself is successful since ρ does not change during the initial extending move. \square

Lemma 3.45. Let $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2$ and $\sigma_2 \downarrow R$. There exists a $R' \in \text{ACC}(\sigma_1)$ such that $R' \sqsubseteq R$.

Proof. The proof is similar to that of Lemma 3.32, and proceeds by contradiction. For some $n \geq 1$, let

$$\text{ACC}(\sigma_1) = \{ R_1, \dots, R_n \}$$

and suppose that $R_i \not\sqsubseteq R$. This means that for every R_i there is $\alpha_i \in R_i$ and a β_i such that $\beta_i \bowtie \alpha_i$ and $\beta_i \not\bowtie \alpha$ whenever $\alpha \in R$.

Let ρ be the contract $\beta_1.\mathbf{1} + \dots + \beta_n.\mathbf{1}$. The contradiction is established by showing that

- (a) $\sigma_1 \text{ MUST } \rho$ while
- (b) $\sigma_2 \text{ MUST } \rho$ is false

Both of which we leave to the reader. Intuitively (b) follows because $\sigma_2 \downarrow R$, while (a) is a consequence of the fact that if $\sigma_1 \xrightarrow{\tau} \sigma' \xrightarrow{\tau}$ then $\sigma' \downarrow R_i$ for some $1 \leq i \leq n$. \square

Lemma 3.46. Let $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2$ and $\sigma_2 \xrightarrow{\beta} \sigma'_2$. Whenever $\alpha \bowtie \beta$

- (a) the set $(\sigma_1 \text{ AFTER } \alpha)$ is not empty
- (b) the contract $\bigoplus(\sigma_1 \text{ AFTER } \alpha)$ is smaller than σ'_2 . Formally,

$$\bigoplus(\sigma_1 \text{ AFTER } \alpha) \sqsubseteq_{\text{MUST}} \sigma'_2$$

Proof. The proof of part (a) is analogous to that of part (a) in Lemma 3.35, but the contract to be used in this case is $\rho = (\mathbf{1} \oplus \mathbf{1}) + \alpha.\text{NIL}$.

We prove point part (b) by contradiction. Suppose there is a contract ρ' such that we have $\bigoplus(\sigma_1 \text{ AFTER } \alpha) \text{ MUST } \rho'$, while $\sigma'_2 \text{ MUST } \rho'$ is false. Consider the contract $\rho = \alpha.\rho' \oplus \mathbf{1}$. Clearly $\sigma_2 \text{ MUST } \rho$ is false while $\sigma_1 \text{ MUST } \rho$ is true. This contradicts the hypothesis that $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2$. \square

Proposition 3.47. If $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2$ then $\sigma_1 \preceq_{\text{SRV}} \sigma_2$.

Proof. The previous three lemmas show that $\sqsubseteq_{\text{MUST}}$ is a semantic sub-server relation, from which the result follows. \square

To establish the converse of this result we need to develop some additional notation. The first is a generalisation of the relation $\sigma \text{ AFTER } \alpha$ to $\sigma \text{ AFTER } u$ where u is a non-empty sequence of actions from Act^* . This is defined by induction on the length of u , with the inductive case being

$$(\sigma \text{ AFTER } w\alpha) = \bigcup_{\sigma' \in (\sigma \text{ AFTER } w)} (\sigma' \text{ AFTER } \alpha)$$

Example 3.48. Let $\sigma = !\mathbf{t}_1.(?\mathbf{t}_2.\sigma_1 + ?\mathbf{t}_3.\sigma_2) + !\mathbf{t}_1.\text{NIL}$ and $?\mathbf{t}_3 \bowtie !\mathbf{t}_2$; we have the following equalities,

$$\begin{aligned} (\sigma \text{ AFTER } ?\mathbf{t}_1 !\mathbf{t}_2) &= \bigcup_{\sigma' \in (\sigma \text{ AFTER } ?\mathbf{t}_1)} (\sigma' \text{ AFTER } !\mathbf{t}_2) \\ &= \bigcup_{\sigma' \in \{\text{NIL}, ?\mathbf{t}_2.\sigma_1 + ?\mathbf{t}_3.\sigma_2\}} (\sigma' \text{ AFTER } !\mathbf{t}_2) \\ &= \{\sigma_1, \sigma_2\} \end{aligned}$$

□

Next we generalise the interaction relation $\alpha \bowtie \beta$ to non-empty sequences, $u \bowtie w$ in the obvious manner; note that this implies that u and w have the same length. Finally we need the notion of contracts performing sequence of actions. For $u \in \text{Act}^*$ let $\sigma \xRightarrow{u} \sigma'$ be the least relation which satisfies

- (a) $\sigma \xRightarrow{\varepsilon} \sigma$ for every contract σ
- (b) $\sigma \xRightarrow{u} \sigma_1$, $\sigma_1 \xrightarrow{a} \sigma'$, where $a \in \text{Act}$, implies $\sigma \xRightarrow{ua} \sigma'$
- (c) $\sigma \xRightarrow{u} \sigma_1$, $\sigma_1 \xrightarrow{\tau} \sigma'$ implies $\sigma \xRightarrow{u} \sigma'$

We have the obvious generalisation of condition (iii) in Definition 3.36:

Lemma 3.49. Suppose $\sigma_1 \mathcal{R} \sigma_2$ for some semantic sub-server relation \mathcal{R} , and $\sigma_2 \xRightarrow{u} \sigma'_2$ for some non-empty $u \in \text{Act}^*$. Then $v \bowtie u$ implies that

- (a) the set $(\sigma_1 \text{ AFTER } v)$ is not empty
- (b) the contract $\bigoplus(\sigma_1 \text{ AFTER } v)$ is related by \mathcal{R} to σ'_2 . Formally,

$$\bigoplus(\sigma_1 \text{ AFTER } v) \mathcal{R} \sigma'_2$$

Proof. By induction on the non-empty size of u ; the base case follows from part (iii) of Definition 3.36. □

Theorem 3.50. [Co-inductive characterisation]

The must pre-order is the greatest semantic sub-server relation.

Proof. We have to prove that for all contracts σ_1, σ_2

$$\sigma_1 \preceq_{\text{SRV}} \sigma_2 \text{ if and only if } \sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2.$$

Because of Proposition 3.47 it is sufficient to prove $\sigma_1 \preceq_{\text{SRV}} \sigma_2$ implies $\sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2$. So, suppose $\sigma_1 \preceq_{\text{SRV}} \sigma_2$ and $\sigma_1 \text{ MUST } \rho$; we must prove $\sigma_2 \text{ MUST } \rho$.

Consider a maximal computation of $\sigma_2 \parallel \rho$

$$\sigma_2 \parallel \rho \xrightarrow{\tau} \sigma_2^1 \parallel \rho_1 \xrightarrow{\tau} \dots \quad (3)$$

We first examine the case when this is finite, with terminal state $\sigma_2^k \parallel \rho_k$. Intuitively this finite computation can be unzipped to give the contributions from the individual components σ_2 and ρ :

$$\sigma_2 \xRightarrow{u} \sigma_2^k \quad \rho \xRightarrow{v} \rho_k \quad \text{where } v \bowtie u$$

We are required to show that one of the derivatives of ρ in $\rho \xRightarrow{v} \rho_k$ is successful. To this aim we will exhibit a suitable computation of $\sigma_1 \parallel \rho$; in particular we will show that there exists a σ'_1 such that

- (a) the composition $\sigma'_1 \parallel \rho_k$ is stuck
- (b) the computation $\sigma_1 \parallel \rho \xrightarrow{\tau} \sigma'_1 \parallel \rho_k$ exists
- (c) the derivatives of ρ in the computation of point (a) are contained in the computation $\sigma_2 \parallel \rho \xrightarrow{\tau} \sigma_2^k \parallel \rho_k$

These three points are enough to prove that in $\rho \xrightarrow{v}$ there exists a successful derivative: thanks to (a), the computation in (b) is a maximal computation of $\sigma_1 \parallel \rho$; the assumption that σ_1 MUST ρ implies that the computation in (b) contains a successful $\hat{\rho}$, and point (c) ensures that $\hat{\rho}$ is contained in $\rho \xrightarrow{v}$.

We prove one by one the points above.

- (a) Here we show that, for a suitable σ'_1 , the composition $\sigma'_1 \parallel \rho_k$ is stuck. By assumption the state $\sigma_2^k \parallel \rho_k$ is terminal; this implies that (1) both σ_2^k and ρ_k are stuck. A consequence is that their acceptance sets are singleton; say $\text{ACC}(\sigma_2^k) = \{\mathbf{R}\}$ and $\text{ACC}(\rho_k) = \{\mathbf{S}\}$; and (2) the contracts σ_2^k and ρ_k cannot interact. Formally

$$\alpha \in \mathbf{R} \text{ implies } \beta \not\bowtie \alpha \text{ for every } \beta \in \mathbf{S}$$

Consider now the contract $\bigoplus(\sigma_1 \text{ AFTER } v)$; part (a) of Lemma 3.49 implies that the set $(\sigma_1 \text{ AFTER } v)$ is not empty and part (b) of the same lemma implies that $\bigoplus(\sigma_1 \text{ AFTER } v) \preceq_{\text{SRV}} \sigma_2^k$.

Part (ii) of Definition 3.36 and (2) above imply that there exists a set

$$\mathbf{R}' \in \text{ACC}(\bigoplus(\sigma_1 \text{ AFTER } v))$$

such that

$$\alpha \in \mathbf{R}' \text{ implies } \beta \not\bowtie \alpha \text{ for every } \beta \in \mathbf{S}' \quad (4)$$

From Definition 3.27 it follows that there exists a contract σ'_1 such that $\sigma'_1 \downarrow \mathbf{R}'$ and

$$\bigoplus(\sigma_1 \text{ AFTER } v) \xrightarrow{\tau} \sigma'_1$$

The latter fact means that $\sigma'_1 \not\bowtie$ and (4) above means that σ'_1 and ρ_k cannot interact; Since (1) above proves that ρ_k is stuck we have shown that $\sigma'_1 \parallel \rho_k$ is stuck.

- (b) We are required to exhibit the computation $\sigma_1 \parallel \rho \xrightarrow{\tau} \sigma'_1 \parallel \rho_k$. Since $\bigoplus(\sigma_1 \text{ AFTER } v) \xrightarrow{\tau} \sigma'_1$ there exists a $\sigma''_1 \in (\sigma_1 \text{ AFTER } v)$ such that $\sigma''_1 \xrightarrow{\tau} \sigma'_1$. From the definition of $(\sigma_1 \text{ AFTER } v)$ it follows that $\sigma_1 \xrightarrow{w} \sigma''_1$ for some $w \in \text{Act}^*$ such that $w \bowtie v$, and this implies that $\sigma_1 \xrightarrow{w} \sigma'_1$. Zipping this action sequence together with $\rho \xrightarrow{v} \rho_k$ we obtain the computation

$$\sigma_1 \parallel \rho \xrightarrow{\tau} \sigma'_1 \parallel \rho_k \not\bowtie$$

We remark that the computation above is finite and cannot be extended, hence it is maximal.

- (c) The derivatives of ρ in the computation

$$\sigma_1 \parallel \rho \xrightarrow{\tau} \sigma'_1 \parallel \rho_k$$

are contained in the computation $\sigma_2 \parallel \rho \xrightarrow{\tau} \sigma_2^k \parallel \rho_k$ because the former computation has been obtained by zipping $\rho \xrightarrow{v} \rho_k$ with a computation made by σ_1 .

Now suppose that the maximal computation (3) above is infinite. Then the result of unzipping gives infinite traces u, v such that

$$\sigma_2 \xrightarrow{u} \quad \rho \xrightarrow{v}$$

Let us denote the finite prefixes of these traces of length k by $u(k), v(k)$ respectively. By Lemma 3.49 we know that $\sigma_1 \text{ AFTER } v(k)$ is non-empty, for every $k \geq 0$. This means that the LTS generated by σ_1 is infinite.

Now consider the sub-LTS consisting of all nodes σ which can be reached from σ_1 using a weak move $\sigma_1 \xrightarrow{w(k)} \sigma'$ where $w(k)$ is some trace satisfying $w(k) \bowtie v(k)$. This sub-LTS is therefore infinite. It is also finite-branching and so by König's lemma it has an infinite path. By following this path from the root we get $\sigma_1 \xrightarrow{w}$ such that $w(k) \bowtie v(k)$, for every $k \geq 0$.

This infinite computation can now be zipped with $\rho \xrightarrow{v}$ to obtain an infinite computation from $\sigma_1 \parallel \rho$. Since $\sigma_1 \text{ MUST } \rho$ it follows that there is some $\sigma_2^k \parallel \rho_k$ in the maximal computation (3) above which is successful, and therefore $\sigma_2 \text{ MUST } \rho$. \square

Corollary 3.51. The server pre-order equals the must pre-order. Formally

$$\sqsubseteq_{\text{SRV}} = \sqsubseteq_{\text{MUST}}$$

Proof. It is a consequence of Theorem 3.38 and Theorem 3.50. \square

Thus far, we have never written explicitly the parameter \bowtie used to infer the reductions of parallel compositions of contracts. Now we make the parameter \bowtie explicit; this lets us emphasise the true import of Corollary 3.51. For every $\bowtie \subseteq \text{Act}^2$, we denote with $\xrightarrow{\tau}_{\bowtie}$ the relation inferred by using the rule in Figure 7 and the \bowtie in the subscript. By replacing $\xrightarrow{\tau}_{\bowtie}$ in Definition 3.17, and Definition 3.40, we define two relations \dashv_{\bowtie} and MUST_{\bowtie} ; in turns this lets us generalise the pre-orders we have studied,

- for every $\bowtie \subseteq \text{Act}^2$, we write $\sigma_1 \sqsubseteq_{\text{SRV}}^{\bowtie} \sigma_2$ if and only if $\rho \dashv^{\bowtie} \sigma_1$ implies that $\rho \dashv^{\bowtie} \sigma_2$;
- for every $\bowtie \subseteq \text{Act}^2$, we write $\sigma_1 \sqsubseteq_{\text{MUST}}^{\bowtie} \sigma_2$ if and only if $\sigma_1 \text{ MUST}^{\bowtie} \rho$ implies that $\sigma_2 \text{ MUST}^{\bowtie} \rho$.

Since in proving Corollary 3.51 we have used no hypothesis on \bowtie , we can rephrase it as follows,

$$\text{for every } \bowtie \subseteq \text{Act}^2. \sqsubseteq_{\text{SRV}}^{\bowtie} = \sqsubseteq_{\text{MUST}}^{\bowtie}$$

This means that the equality between the server pre-order and the must pre-order does not depend on the \bowtie at hand. That is, the equality holds regardless of the co-action relation used in rule [P-SYNCH].

4. Session Contracts

Here we specialise the contract language, to a sub-language which will be the target of our interpretation of the session types from Section 2. This is the topic of the first sub-section. We then go on to re-examine the server pre-order as it applies to this sub-language; in particular we show that it can also be characterised co-inductively, this time using purely

$\sigma, \rho ::=$	$\mathbf{1}$ $\sum_{i \in I} ?l_i.\sigma$ $\bigoplus_{i \in I} !l_i.\sigma$ $?t.\sigma$ $!t.\sigma$ x $\mu x.\sigma$	Session Contracts <i>success</i> <i>external choice, I finite, non-empty</i> <i>internal choice, I finite, non-empty</i> <i>input</i> <i>output</i> <i>contract variable</i> <i>recursion</i>
--------------------	--	---

We impose the additional proviso that in a term the l_i 's are pair-wise different.

Figure 9. Session contract grammar.

syntactical criteria. In the final section we give a similar co-inductive characterisation to a related sub-client pre-order.

4.1. Session contracts

The syntax for the language L_{SC} is given in Figure 9.

Definition 4.1. [Session contracts]

We use \mathcal{SC} to denote the set of closed terms σ of L_{SC} such that if for some $s \in Act^*$, $\sigma \xrightarrow{s} \sigma'$, then $\sigma' \downarrow$. We refer to these terms as *session contracts*. \square

Note that \mathcal{SC} is a subset of the more general language of contracts \mathcal{C} , but

- external choices are restricted to inputs on labels
- internal choices are restricted to outputs on labels

Note also that \mathbf{NIL} is not a session contract. Instead we have chosen $\mathbf{1}$ to be the base contract, for reasons which will become apparent. Moreover, we already reasoned that a server contract $\mathbf{1}$ has the same behaviour as \mathbf{NIL} (Proposition 3.21).

Session contracts, due to their restrictive syntax, enjoy some properties which we will use in the next sub-sections, and which we prove now.

Lemma 4.2. Let σ be a session contract. Then

- (i) $\sigma \xrightarrow{\checkmark}$ if and only if $\sigma = \mathbf{1}$
- (ii) $\sigma \xrightarrow{\tau} \sigma' \xrightarrow{\checkmark}$ if and only if $\text{UNFOLD}(\sigma) = \mathbf{1}$
- (iii) $\sigma \xrightarrow{\alpha}$ if and only if $\sigma \neq \mathbf{1}$

Proof. Part (i) follows from the restrictive syntax of session contract. The proof of part (ii) requires two arguments. The *if* side, $\text{UNFOLD}(\sigma) = \mathbf{1}$ implies $\sigma \xrightarrow{\tau} \sigma' \xrightarrow{\checkmark}$, is justified by part (ii) of Lemma 3.2. The *only if* side, $\sigma \xrightarrow{\tau} \sigma' \xrightarrow{\checkmark}$ implies $\text{UNFOLD}(\sigma) = \mathbf{1}$, can be proven by induction on the length of the sequence $\xrightarrow{\tau}$; the base case being part (i) of this lemma. Part (iii) can be proven by structural induction. \square

4.2. The restricted server pre-order

Definition 3.25 applies equally well to session contracts, but it is inappropriate as it compares session contracts from the point of view of satisfying clients who may use the more general contracts from Section 3. Instead let us restrict attention to clients who also only run the more restricted session contracts.

Definition 4.3. [Restricted server pre-order]

For $\sigma_1, \sigma_2 \in \mathcal{SC}$ let $\sigma_1 \sqsubseteq_{\text{SRV}}^{\mathcal{SC}} \sigma_2$ whenever $\rho \dashv \sigma_1$ implies $\rho \dashv \sigma_2$ for every ρ in \mathcal{SC} . \square

This relation is more generous than \sqsubseteq_{SRV} , in that it allows implementation refinement (Padovani 2010) to happen, as the following example shows.

Example 4.4.

$$?\text{latte.1} \sqsubseteq_{\text{SRV}}^{\mathcal{SC}} ?\text{moka.1} + ?\text{latte.1}$$

If a session contract ρ can interact with $?\text{latte.1}$ then, modulo unfolding, it has to be defined by an internal sum. Moreover this sum can only contain one summand, and therefore ρ complies also with $?\text{moka.1} + ?\text{latte.1}$.

Consider now the more general contract $\rho' = !\text{latte.1} + !\text{moka.NIL}$. One can check that

$$\rho' \dashv ?\text{latte.1}$$

whereas $\rho' \not\vdash ?\text{moka.1} + ?\text{latte.1}$. It therefore follows that

$$?\text{latte.1} \not\sqsubseteq_{\text{SRV}} ?\text{moka.1} + ?\text{latte.1}$$

\square

Example 4.5. [e-vote, ballot refinement]

We give a more concrete instance of the previous example. Recall Example 3.14 and consider the session contract

$$\begin{aligned} \text{BallotB} = & \mu x. ?\text{Login}. (!\text{Wrong.1} \oplus \\ & !\text{Ok}. (?\text{VoteA.1} + ?\text{VoteB.1} + ?\text{VoteC.1} + ?\text{VoteD.1})) \end{aligned}$$

BallotB offers to a voter more options than Ballot, and intuitively it should be possible to use a server that guarantees BallotB in place of a server that guarantees Ballot. This is not the case if the contracts are compared with \sqsubseteq_{SRV} , because $\text{Ballot} \not\sqsubseteq_{\text{SRV}} \text{BallotB}$. On the other hand, if we restrict our attention to session contracts, and thus to the pre-order $\sqsubseteq_{\text{SRV}}^{\mathcal{SC}}$, we have $\text{Ballot} \sqsubseteq_{\text{SRV}}^{\mathcal{SC}} \text{BallotB}$. \square

When comparing session contracts relative to this pre-order it will be convenient to work modulo unfolding, which is possible because of the following result:

Proposition 4.6. Let σ_1, σ_2 be session contracts; $\sigma_1 \sqsubseteq_{\text{SRV}}^{\mathcal{SC}} \sigma_2$ if and only if $\text{UNFOLD}(\sigma_2) \sqsubseteq_{\text{SRV}}^{\mathcal{SC}} \text{UNFOLD}(\sigma_1)$.

Proof. Follows from Proposition 3.23 and 3.24. \square

Proposition 4.7 (Bottom element).

The pre-order $\sqsubseteq_{\text{SRV}}^{\mathcal{SC}}$ enjoys the following properties,

- (i) it has a bottom element
- (ii) if σ_{\perp} is a bottom element of $\sqsubseteq_{\text{SRV}}^{\text{SC}}$ then $\text{UNFOLD}(\sigma_{\perp}) = \mathbf{1}$

Proof. To prove part (i) we show that $\mathbf{1}$ is a bottom element of $\sqsubseteq_{\text{SRV}}^{\text{SC}}$, that is $\mathbf{1} \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma$ for every session contract σ . Let ρ be a session contract such that $\rho \dashv \mathbf{1}$. The session contract $\mathbf{1}$ offers no interaction. Therefore, because of the restricted syntax of session contracts, ρ must also be, modulo unfolding, the simple contract $\mathbf{1}$. Now fix a session contract σ . Clearly $\mathbf{1} \dashv \sigma$, therefore from an application of Proposition 3.24 it follows that $\rho \dashv \sigma$.

To prove part (ii) let σ_{\perp} be an arbitrary bottom element of $\sqsubseteq_{\text{SRV}}^{\text{SC}}$. We are required to show that $\text{UNFOLD}(\sigma_{\perp}) = \mathbf{1}$. From the definition of bottom element follows $\sigma_{\perp} \sqsubseteq_{\text{SRV}}^{\text{SC}} \mathbf{1}$. An application of the previous proposition gives $\text{UNFOLD}(\sigma_{\perp}) \sqsubseteq_{\text{SRV}}^{\text{SC}} \mathbf{1}$. But now an analysis of the possible syntactic structure of $\text{UNFOLD}(\sigma_{\perp})$ quickly yields that it must be $\mathbf{1}$ itself. \square

Part (ii) is relevant because $\mathbf{1}$ is not the only bottom element; for example it is also true that $\mu x. \mathbf{1} \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma$ for every σ .

We now proceed, as in Section 3.2.1, to give a co-inductive characterisation of the restricted server pre-order, this time taking advantage of their restricted syntactic structure.

Definition 4.8. [Syntactic sub-server relation]

Let $\mathcal{F}_{\lesssim_{\text{SRV}}^{\text{syn}}} : \mathcal{P}(\text{SC}^2) \longrightarrow \mathcal{P}(\text{SC}^2)$ be defined by letting $(\sigma_1, \sigma_2) \in \mathcal{F}_{\lesssim_{\text{SRV}}^{\text{syn}}}(\mathcal{R})$ whenever one of the following holds:

- (i) $\text{UNFOLD}(\sigma_1) = \mathbf{1}$
- (ii) $\text{UNFOLD}(\sigma_2) = ?\mathbf{t}_2.\sigma'_2$ and $\text{UNFOLD}(\sigma_1) = ?\mathbf{t}_1.\sigma'_1$ with $\mathbf{t}_1 \lesssim_{\mathbf{g}} \mathbf{t}_2$ and $\sigma'_1 \mathcal{R} \sigma'_2$
- (iii) $\text{UNFOLD}(\sigma_2) = !\mathbf{t}_2.\sigma'_2$ and $\text{UNFOLD}(\sigma_1) = !\mathbf{t}_1.\sigma'_1$ with $\mathbf{t}_2 \lesssim_{\mathbf{g}} \mathbf{t}_1$ and $\sigma'_1 \mathcal{R} \sigma'_2$
- (iv) $\text{UNFOLD}(\sigma_2) = \sum_{j \in J} ?\mathbf{l}_j.\sigma_j^2$ and $\text{UNFOLD}(\sigma_1) = \sum_{i \in I} ?\mathbf{l}_i.\sigma_i^1$ with $I \subseteq J$ and $\sigma_i^1 \mathcal{R} \sigma_i^2$
- (v) $\text{UNFOLD}(\sigma_2) = \bigoplus_{j \in J} !\mathbf{l}_j.\sigma_j^2$ and $\text{UNFOLD}(\sigma_1) = \bigoplus_{i \in I} !\mathbf{l}_i.\sigma_i^1$ with $J \subseteq I$ and $\sigma_j^1 \mathcal{R} \sigma_j^2$

If $\mathcal{R} \subseteq \mathcal{F}_{\lesssim_{\text{SRV}}^{\text{syn}}}(\mathcal{R})$ then we say that \mathcal{R} is a co-inductive *syntactic sub-server* relation. Let $\lesssim_{\text{SRV}}^{\text{syn}}$ denote the greatest solution of the equation $X = \mathcal{F}_{\lesssim_{\text{SRV}}^{\text{syn}}}(X)$; formally,

$$\lesssim_{\text{SRV}}^{\text{syn}} = \nu \mathcal{F}_{\lesssim_{\text{SRV}}^{\text{syn}}}$$

We call $\lesssim_{\text{SRV}}^{\text{syn}}$ the *syntactic sub-server* relation. \square

We first show that the set based relation $\sqsubseteq_{\text{SRV}}^{\text{SC}}$ is contained in $\lesssim_{\text{SRV}}^{\text{syn}}$; this will follow if we can show that the former satisfies the properties defining the latter.

Lemma 4.9. Let $\sigma_1, \sigma_2 \in \text{SC}$, $\sigma_1 = \text{UNFOLD}(\sigma_1)$, $\sigma_2 = \text{UNFOLD}(\sigma_2)$ and $\sigma_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma_2$. One of the following is true,

- (a) if $\sigma_1 = !\mathbf{t}_1.\sigma'_1$ then $\sigma_2 = !\mathbf{t}_2.\sigma'_2$, $\mathbf{t}_2 \lesssim_{\mathbf{g}} \mathbf{t}_1$ and $\sigma'_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma'_2$
- (b) if $\sigma_1 = ?\mathbf{t}_1.\sigma'_1$ then $\sigma_2 = ?\mathbf{t}_2.\sigma'_2$, $\mathbf{t}_1 \lesssim_{\mathbf{g}} \mathbf{t}_2$ and $\sigma'_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma'_2$
- (c) if $\sigma_1 = \sum_{i \in I} ?\mathbf{l}_i.\sigma_i^1$ then $\sigma_2 = \sum_{j \in J} ?\mathbf{l}_j.\sigma_j^2$, $I \subseteq J$ and $\sigma_i^1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma_i^2$
- (d) if $\sigma_1 = \bigoplus_{i \in I} !\mathbf{l}_i.\sigma_i^1$ then $\sigma_2 = \bigoplus_{j \in J} !\mathbf{l}_j.\sigma_j^2$ with $J \subseteq I$ and $\sigma_j^1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma_j^2$

Proof. The proof is by case analysis on the structure of σ_1 , and the argument depends greatly on the restricted syntax of session contracts. We prove part (a) and leave the proof of the other parts to the reader.

Suppose $\sigma_1 = !\mathbf{t}_1.\sigma'_1$. Then $?\mathbf{t}_1.\mathbf{1} \dashv \sigma_1$ and because $\sigma_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma_2$ it follows that $?\mathbf{t}_1.\mathbf{1} \dashv \sigma_2$. Since $?\mathbf{t}_1.\mathbf{1}$ is stuck, σ_2 has to engage in an action $!\mathbf{t}_2$ such that $?\mathbf{t}_1 \bowtie_c !\mathbf{t}_2$; this lets us prove that $\mathbf{t}_2 \preceq_g \mathbf{t}_1$. In reason of the syntax and the hypothesis $\sigma_2 = \text{UNFOLD}(\sigma_2)$, the equality $\sigma_2 = !\mathbf{t}_2.\sigma'_2$ must hold.

We also have to prove that $\sigma'_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma'_2$. Pick a session contract ρ such that $\rho \dashv \sigma'_1$. Clearly $?\mathbf{t}_1.\rho \dashv \sigma_1$, and thus $?\mathbf{t}_1.\rho \dashv \sigma_2$. Since $?\mathbf{t}_1 \bowtie_c !\mathbf{t}_2$, we apply rule [P-SYNCH] to infer $?\mathbf{t}_1.\rho \parallel \sigma_2 \xrightarrow{\tau} \rho \parallel \sigma'_2$. From the definition of compliance it follows that $\rho \dashv \sigma'_2$. \square

Proposition 4.10. For session contracts, $\sigma_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma_2$ implies $\sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma_2$.

Proof. We prove that $\sqsubseteq_{\text{SRV}}^{\text{SC}}$ is a pre-fixed point of the function $\mathcal{F}_{\preceq_{\text{SRV}}^{\text{syn}}}$ of Definition 4.8, that is $\sigma_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma_2$ implies $(\sigma_1, \sigma_2) \in \mathcal{F}_{\preceq_{\text{SRV}}^{\text{syn}}}(\sqsubseteq_{\text{SRV}}^{\text{SC}})$.

Suppose $\sigma_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma_2$. Then by Proposition 4.6 it follows that $\text{UNFOLD}(\sigma_1) \sqsubseteq_{\text{SRV}}^{\text{SC}} \text{UNFOLD}(\sigma_2)$. Now if $\text{UNFOLD}(\sigma_1) = \mathbf{1}$ by definition $(\sigma_1, \sigma_2) \in \mathcal{F}_{\preceq_{\text{SRV}}^{\text{syn}}}(\sqsubseteq_{\text{SRV}}^{\text{SC}})$. Otherwise we can apply Lemma 4.9 to the pair $\text{UNFOLD}(\sigma_1), \text{UNFOLD}(\sigma_2)$. This provides the required information to satisfy the requirements (ii) to (v) in Definition 4.8, thereby ensuring that $(\sigma_1, \sigma_2) \in \mathcal{F}_{\preceq_{\text{SRV}}^{\text{syn}}}(\sqsubseteq_{\text{SRV}}^{\text{SC}})$. \square

Lemma 4.11. Let \mathcal{R} be a co-inductive syntactic sub-server relation and let $\sigma_1 \mathcal{R} \sigma_2$. If $\sigma_2 \xrightarrow{\tau} \sigma'_2$ then $\sigma_1 \mathcal{R} \sigma'_2$.

Proof. First note that from Definition 4.8 it follows that

$$\text{UNFOLD}(\sigma_1) \mathcal{R} \text{UNFOLD}(\sigma_2) \quad (5)$$

There are two different cases to be discussed, depending on the unfolding of σ_2 being σ_2 itself or not.

- (a) If $\text{UNFOLD}(\sigma_2) \neq \sigma_2$ then $\text{UNFOLD}(\sigma_2) = \text{UNFOLD}(\sigma'_2)$. The equality and (5) above imply $\text{UNFOLD}(\sigma_1) \mathcal{R} \text{UNFOLD}(\sigma'_2)$; the latter fact means that $\sigma_1 \mathcal{R} \sigma'_2$.
- (b) If $\text{UNFOLD}(\sigma_2) = \sigma_2$ then σ_2 must be an internal sum, say $\sigma_2 = \bigoplus_{i \in I} !\mathbf{1}_i.\sigma_i^2$, because the contract σ_2 can perform a silent move and cannot unfold. This implies that σ'_2 is the internal sum $\bigoplus_{k \in K} !\mathbf{1}_k.\sigma_k^2$ for some $K \subseteq I$. From Definition 4.8 it follows that

$$\text{UNFOLD}(\sigma_1) = \bigoplus_{j \in J} !\mathbf{1}_j.\sigma_j^2$$

with $I \subseteq J$. Since $\text{UNFOLD}(\sigma'_2) = \sigma'_2$ and $K \subseteq I \subseteq J$ one sees easily that

$$\text{UNFOLD}(\sigma_1) \mathcal{R} \text{UNFOLD}(\sigma'_2)$$

thus $\sigma_1 \mathcal{R} \sigma'_2$. \square

Lemma 4.12. Let \mathcal{R} be a co-inductive syntactic sub-server relation, $\sigma_1 \mathcal{R} \sigma_2$ and $\sigma_2 \downarrow R$. There exists a $R' \in \text{Acc}(\sigma_1)$ such that $R' \sqsubseteq R$.

Proof. Using Lemma 3.2 we know $\text{UNFOLD}(\sigma_2) = \sigma_2$, since $\sigma_2 \not\stackrel{\tau}{\rightarrow}$. From Definition 4.8 it follows that $\text{UNFOLD}(\sigma_1) \mathcal{R} \sigma_2$. Now, according to the cases in Definition 4.8 and a case analysis on the form of σ_2 , one can show that $\text{UNFOLD}(\sigma_1) \xrightarrow{\tau}^* \sigma'_1$ for some σ'_1 which satisfies the required properties. We leave the details of the case analysis to the reader.

Finally, the proof that $\sigma_1 \xrightarrow{\tau}^* \sigma'_1$ amounts in two steps. We apply Lemma 3.2, which ensures that $\sigma_1 \xrightarrow{\tau}^* \text{UNFOLD}(\sigma_1)$. Now we know that

$$\sigma_1 \xrightarrow{\tau}^* \text{UNFOLD}(\sigma_1) \xrightarrow{\tau}^* \sigma'_1$$

so the transitivity of $\xrightarrow{\tau}^*$ gives the result. \square

Lemma 4.11 and Lemma 4.12 proves that the pre-order $\preceq_{\text{SRV}}^{\text{syn}}$ enjoys two of the properties of the pre-order \preceq_{SRV} . The third property of \preceq_{SRV} , though, is not afforded by $\preceq_{\text{SRV}}^{\text{syn}}$.

Example 4.13. Let $\sigma_1 = ?\text{latte.1}$ and $\sigma_2 = ?\text{latte.1} + ?\text{moka.1}$; in Example 4.4 we have shown that $\sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma_2$. On the one hand, we can prove the following things, $\sigma_2 \xrightarrow{?moka} \mathbf{1}$ and $!moka \bowtie_c ?moka$; on the other hand

$$(\sigma_1 \text{ AFTER } !moka) = \emptyset$$

As in Example 4.4 the crucial fact is that the pre-order $\preceq_{\text{SRV}}^{\text{syn}}$ allows implementation refinement (Padovani 2010). \square

Theorem 4.14. [Co-inductive characterisation]

For session contracts σ_1, σ_2 , $\sigma_1 \sqsubseteq_{\text{SRV}}^{\text{SC}} \sigma_2$ if and only if $\sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma_2$.

Proof. The only if part of the theorem is Proposition 4.10 while the if part, that is the set inclusion $\preceq_{\text{SRV}}^{\text{syn}} \subseteq \sqsubseteq_{\text{SRV}}^{\text{SC}}$, follows from the fact that the relation

$$\mathcal{R} = \{ (\rho, \sigma_2) \mid \sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma_2, \rho \dashv \sigma_1 \text{ for some } \sigma_1 \in \mathcal{SC} \}$$

contains (ρ, σ_2) and is a compliance. We prove the latter.

We have to show that \mathcal{R} satisfies the two properties in Definition 3.17. Let $(\rho, \sigma) \in \mathcal{R}$; by definition there exists a σ_1 such that $\rho \dashv \sigma_1$ and $\sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma$.

To prove point (i) of Definition 3.17 assume $\rho \parallel \sigma \not\stackrel{\tau}{\rightarrow}$. This implies that σ and ρ are both stuck, so $\text{ACC}(\rho) = \{\text{S}\}$ and $\text{ACC}(\sigma) = \{\text{R}\}$, and that

$$\alpha \in \text{R} \text{ implies } \beta \not\bowtie_c \alpha \text{ whenever } \beta \in \text{S}$$

An application of Lemma 4.12 and of the definition of acceptance set gives a σ'_1 such that $\sigma_1 \xrightarrow{\tau}^* \sigma'_1 \downarrow \text{R}'$ and $\text{R}' \sqsubseteq \text{R}$. The last inequality implies that

$$\alpha \in \text{R}' \text{ implies } \beta \not\bowtie_c \alpha \text{ whenever } \beta \in \text{S}$$

therefore, $\rho \parallel \sigma'_1 \not\stackrel{\tau}{\rightarrow}$. Part (ii) of Definition 3.17 and the assumption $\rho \dashv \sigma_1$ imply that the session contract ρ complies with σ'_1 so $\rho \stackrel{\tau}{\rightarrow}$.

What we have left to do now is to show that if $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$ then $(\rho', \sigma') \in \mathcal{R}$, that is there exists a $\hat{\sigma} \in \mathcal{SC}$ such that

$$\rho' \dashv \hat{\sigma}, \quad \hat{\sigma} \preceq_{\text{SRV}}^{\text{syn}} \sigma'$$

Assume $\rho \parallel \sigma \xrightarrow{\tau} \rho' \parallel \sigma'$. The argument depends on the rule used to infer this silent move (see Figure 7). If rule [P-SIL-L] was used then $\sigma' = \sigma$ and $\rho \xrightarrow{\tau} \rho'$; let $\hat{\sigma} = \sigma_1$. Then we already know that $\hat{\sigma} \preceq_{\text{SRV}}^{\text{syn}} \sigma'$, and part (ii) of Definition 3.17 implies $\rho' \dashv \hat{\sigma}$. If rule [P-SIL-R] was applied then $\rho' = \rho$ and $\sigma \xrightarrow{\tau} \sigma'$. In this case an application of Lemma 4.11 implies $\sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma'$. We know by assumption that $\rho \dashv \sigma_1$, so the $\hat{\sigma}$ we are looking for is σ_1 .

If rule [P-SYNCH] was applied then

$$\rho \xrightarrow{\alpha} \rho', \quad \sigma \xrightarrow{\beta} \sigma', \quad \alpha \bowtie_c \beta.$$

Since ρ performs an observable action part (iii) of Lemma 4.2 implies $\rho \not\xrightarrow{\checkmark}$.

Let us turn our attention to $\text{UNFOLD}(\sigma_1)$. The assumption $\rho \dashv \sigma_1$ together with Proposition 3.23 implies that (a) $\rho \dashv \text{UNFOLD}(\sigma_1)$. The assumption $\sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma$ and Definition 4.8 imply that (b) $\text{UNFOLD}(\sigma_1) \preceq_{\text{SRV}}^{\text{syn}} \sigma$.

We know that $\rho \dashv \text{UNFOLD}(\sigma_1)$ ((a) above), and that $\rho \not\xrightarrow{\checkmark}$; together with part (i), these facts force $\text{UNFOLD}(\sigma_1)$ to offer an action δ such that $\delta \bowtie_c \alpha$. Thus, for some σ'_1 ,

$$\text{UNFOLD}(\sigma_1) \xrightarrow{\delta} \sigma'_1$$

An application of rule [P-SYNCH] ensures that

$$\rho \parallel \text{UNFOLD}(\sigma_1) \xrightarrow{\tau} \rho' \parallel \sigma'_1$$

Now (a) and part (ii) of Definition 3.17 imply that $\rho' \dashv \sigma'_1$. We choose σ'_1 as candidate $\hat{\sigma}$.

To finish the proof we have to show that $\sigma'_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma'$. The argument is a case analysis on the action β . Four cases are to be discussed, but, as they are all similar, we give a detailed account only of two of them.

If $\beta = ?\tau_2$ then (b) above and case (ii) of Definition 4.8 ensure that σ' is unique, and so is σ'_1 as well. The same definition implies also that $\sigma'_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma'$.

If, for some label 1 , $\beta = ?1$ then the definition of \bowtie_c implies that $\alpha = !1$, and the assumption $\delta \bowtie_c \alpha$ implies $\delta = ?1$. We have proven that $\delta = \beta$. Now (b) above and case (iv) of Definition 4.8 imply that $\sigma'_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma'$. \square

We conclude this subsection with a summary of our knowledge on the pre-orders which compare contracts on the server side of the compliance relation.

Corollary 4.15. The following equalities and inequalities hold

$$\begin{aligned} \sqsubseteq_{\text{MUST}} &= \sqsubseteq_{\text{SRV}} \\ \sqsubseteq_{\text{SRV}}^{\text{SC}} &\not\subseteq \sqsubseteq_{\text{SRV}} \\ \sqsubseteq_{\text{SRV}} &\not\subseteq \sqsubseteq_{\text{SRV}}^{\text{SC}} \\ \sqsubseteq_{\text{SRV}}^{\text{SC}} &= \preceq_{\text{SRV}}^{\text{syn}} \end{aligned}$$

Proof. The (in)equalities are consequence respectively of Corollary 3.51, Example 4.4, the fact that $\text{SC} \subset \mathcal{C}$, and Theorem 4.14. \square

4.3. The restricted client pre-order

We introduce a new pre-order which compares the capacity of clients to be satisfied by servers. The structure of this sub-section is similar to that of the previous one on the restricted server pre-order.

Definition 4.16. [Restricted client pre-order]

For $\rho_1, \rho_2 \in \mathcal{SC}$ let $\rho_1 \sqsubseteq_{\text{CLT}}^{\text{SC}} \rho_2$ whenever $\rho_1 \dashv \sigma$ implies $\rho_2 \dashv \sigma$ for every σ in \mathcal{SC} . \square

Also on the restricted client pre-order we can reason modulo unfolding.

Proposition 4.17. For every session contract ρ_1 and ρ_2 , $\rho_1 \sqsubseteq_{\text{CLT}}^{\text{SC}} \rho_2$ if and only if $\text{UNFOLD}(\rho_1) \sqsubseteq_{\text{CLT}}^{\text{SC}} \text{UNFOLD}(\rho_2)$.

Proof. Follows from Proposition 3.23 and 3.24. \square

Example 4.18. We have shown in Example 4.4 that $?\text{latte.1} \sqsubseteq_{\text{SRV}}^{\text{SC}} ?\text{moka.1} + ?\text{latte.1}$. A similar argument, this time applied to client-side session contracts, can be used to show that

$$?\text{latte.1} \sqsubseteq_{\text{CLT}}^{\text{SC}} ?\text{moka.1} + ?\text{latte.1}$$

Similarly to what happens for server contracts, if we turn our attention to general contracts then the session contracts above are no longer related. Let us see why. The client $?\text{latte.1}$ complies with the server $!\text{latte.1} + !\text{moka.1}$, because the action $!\text{moka}$ will never be performed by the server. On the other hand

$$?\text{moka.1} + ?\text{latte.1} \not\sqsubseteq_{\text{CLT}}^{\text{SC}} !\text{latte.1} + !\text{moka.1} \xrightarrow{\tau} ?\text{moka.1} \parallel \mathbf{1} \not\rightarrow$$

and $?\text{moka.1} \not\rightarrow$; this proves that

$$?\text{moka.1} + ?\text{latte.1} \not\sqsubseteq_{\text{CLT}}^{\text{SC}} !\text{latte.1} + !\text{moka.1}$$

Had we defined in the obvious way the pre-order \sqsubseteq_{CLT} on contracts, then the argument above would have proven that

$$!\text{latte.1} \not\sqsubseteq_{\text{CLT}} ?\text{moka.1} + ?\text{latte.1} \tag{6}$$

We have therefore shown that

$$\sqsubseteq_{\text{CLT}}^{\text{SC}} \not\subseteq \sqsubseteq_{\text{CLT}}$$

\square

We have seen in Proposition 4.7 that the session contract $\mathbf{1}$ is a bottom element in the restricted server pre-order. The client pre-order enjoys the dual property.

Proposition 4.19 (Top element).

The pre-order $\sqsubseteq_{\text{CLT}}^{\text{SC}}$ enjoys the following two properties,

- (i) it has a top element
- (ii) if σ_{\top} is a top element of $\sqsubseteq_{\text{CLT}}^{\text{SC}}$ then $\text{UNFOLD}(\sigma_{\top}) = \mathbf{1}$

Proof. Since $\mathbf{1} \dashv \sigma$ for every contract σ , the session contract $\mathbf{1}$ it is a top element in the restricted client pre-order. Moreover, reasoning as in Proposition 4.7 we can show that if σ_{\top} is an arbitrary top element then $\text{UNFOLD}(\sigma_{\top}) = \mathbf{1}$. \square

Definition 4.20. [Syntactic sub-client relation]

Let $\mathcal{F}_{\preceq_{\text{CLT}}^{\text{syn}}} : \mathcal{P}(\mathcal{SC}^2) \rightarrow \mathcal{P}(\mathcal{SC}^2)$ be defined so that $(\rho_1, \rho_2) \in \mathcal{F}_{\preceq_{\text{CLT}}^{\text{syn}}}(\mathcal{R})$ whenever one of the following is true:

- (i) $\text{UNFOLD}(\rho_2) = \mathbf{1}$
- (ii) $\text{UNFOLD}(\rho_2) = ?\mathbf{t}_2.\rho'_2$ and $\text{UNFOLD}(\rho_1) = ?\mathbf{t}_1.\rho'_1$ with $\mathbf{t}_1 \preceq_{\mathbf{g}} \mathbf{t}_2$ and $\rho'_1 \mathcal{R} \rho'_2$
- (iii) $\text{UNFOLD}(\rho_2) = !\mathbf{t}_2.\rho'_2$ and $\text{UNFOLD}(\rho_1) = !\mathbf{t}_1.\rho'_1$ with $\mathbf{t}_2 \preceq_{\mathbf{g}} \mathbf{t}_1$ and $\rho'_1 \mathcal{R} \rho'_2$
- (iv) $\text{UNFOLD}(\rho_2) = \sum_{j \in J} \mathbf{1}_j.\rho_j^2$ and $\text{UNFOLD}(\rho_1) = \sum_{i \in I} \mathbf{1}_i.\rho_i^1$ with $I \subseteq J$ and $\rho_i^1 \mathcal{R} \rho_i^2$
- (v) $\text{UNFOLD}(\rho_2) = \bigoplus_{j \in J} \mathbf{1}_j.\sigma_j^2$ and $\text{UNFOLD}(\rho_1) = \bigoplus_{i \in I} \mathbf{1}_i.\sigma_i^1$ with $J \subseteq I$ and $\sigma_j^1 \mathcal{R} \sigma_j^2$

If $\mathcal{R} \subseteq \mathcal{F}_{\preceq_{\text{CLT}}^{\text{syn}}}(\mathcal{R})$ then we say that \mathcal{R} is a co-inductive syntactic sub-client relation. Let $\preceq_{\text{CLT}}^{\text{syn}}$ denote the greatest solution of the equation $X = \mathcal{F}_{\preceq_{\text{CLT}}^{\text{syn}}}(X)$; formally,

$$\preceq_{\text{CLT}}^{\text{syn}} = \nu \mathcal{F}_{\preceq_{\text{CLT}}^{\text{syn}}}$$

We call $\preceq_{\text{CLT}}^{\text{syn}}$ the sub-client relation. \square

Lemma 4.21. Let $\rho_1, \rho_2 \in \mathcal{SC}$, $\rho_1 = \text{UNFOLD}(\rho_1)$, $\rho_2 = \text{UNFOLD}(\rho_2)$ and $\rho_1 \sqsubseteq_{\text{CLT}}^{\text{SC}} \rho_2$. Then

- (a) if $\rho_2 = !\mathbf{t}_2.\rho'_2$ then $\rho_1 = !\mathbf{t}_1.\rho'_1$, $\mathbf{t}_2 \preceq_{\mathbf{g}} \mathbf{t}_1$ and $\rho'_1 \sqsubseteq_{\text{CLT}}^{\text{SC}} \rho'_2$
- (b) if $\rho_2 = ?\mathbf{t}_2.\rho'_2$ then $\rho_1 = ?\mathbf{t}_1.\rho'_1$, $\mathbf{t}_1 \preceq_{\mathbf{g}} \mathbf{t}_2$ and $\rho'_1 \sqsubseteq_{\text{CLT}}^{\text{SC}} \rho'_2$
- (c) if $\rho_2 = \sum_{j \in J} \mathbf{1}_j.\rho_j^2$ then $\rho_1 = \sum_{i \in I} \mathbf{1}_i.\rho_i^1$ with $I \subseteq J$ and $\rho_i^1 \sqsubseteq_{\text{CLT}}^{\text{SC}} \rho_i^2$
- (d) if $\rho_2 = \bigoplus_{j \in J} \mathbf{1}_j.\rho_j^2$ then $\rho_1 = \bigoplus_{i \in I} \mathbf{1}_i.\rho_i^1$ with $J \subseteq I$ and $\rho_j^1 \sqsubseteq_{\text{CLT}}^{\text{SC}} \rho_j^2$

Proof. The proof is almost the same as Lemma 4.9, the difference being that here we look at left-hand side of the compliance relation, and that we need a result of the fact that for each session contract ρ there exists a session contract σ such that $\rho \dashv \sigma$. This is proven in (Barbanera & de'Liguoro 2010, Section 2.1). \square

Proposition 4.22. For every session contract ρ_1 and ρ_2 , if $\rho_1 \sqsubseteq_{\text{CLT}}^{\text{SC}} \rho_2$ then $\rho_1 \preceq_{\text{CLT}}^{\text{syn}} \rho_2$.

Proof. The argument is similar to the one of Proposition 4.10, but here we use the function $\mathcal{F}_{\preceq_{\text{CLT}}^{\text{syn}}}$ and Lemma 4.21. \square

Lemma 4.23. Let \mathcal{R} be a co-inductive sub-client relation and let $\rho_1 \preceq_{\text{CLT}}^{\text{syn}} \rho_2$. If $\rho_2 \xrightarrow{\tau} \rho'_2$ then $\rho_1 \preceq_{\text{CLT}}^{\text{syn}} \rho'_2$.

Proof. The proof is similar to the proof of Lemma 4.11. \square

Theorem 4.24. [Co-inductive characterisation]

Let $\rho, \sigma \in \mathcal{SC}$. Then $\rho \preceq_{\text{CLT}}^{\text{syn}} \sigma$ if and only if $\rho \sqsubseteq_{\text{CLT}}^{\text{SC}} \sigma$.

Proof. In view of Proposition 4.22 we have to prove only the inclusion $\preceq_{\text{CLT}}^{\text{syn}} \subseteq \sqsubseteq_{\text{CLT}}^{\text{SC}}$. It is enough to show that

$$\mathcal{R} = \{ (\rho_2, \sigma) \mid \rho_1 \preceq_{\text{CLT}}^{\text{syn}} \rho_2, \rho_1 \dashv \sigma \text{ for some } \rho_1 \in \mathcal{SC} \}$$

is a co-inductive compliance. Let $\rho \mathcal{R} \sigma$; by definition of \mathcal{R} there exists a session contract ρ_1 such that $\rho_1 \preceq_{\text{CLT}}^{\text{syn}} \rho$ and $\rho_1 \dashv \sigma$.

We prove part (i) of Definition 3.17. Assume $\rho \parallel \sigma \not\stackrel{\tau}{\rightarrow}$; we have to show that $\rho \not\stackrel{\checkmark}{\rightarrow}$. To this aim it is sufficient to show

$$\text{UNFOLD}(\rho_1) = \mathbf{1} \quad (7)$$

We explain why this fact suffices. Assume (7). Since $\text{UNFOLD}(\rho_1) \preceq_{\text{CLT}}^{\text{syn}} \rho$ we know that

$$(\rho_1, \rho) \in \mathcal{F}_{\preceq_{\text{CLT}}^{\text{syn}}} (\preceq_{\text{CLT}}^{\text{syn}})$$

This is possible only thanks to case (i) of Definition 4.20, and therefore $\text{UNFOLD}(\rho) = \mathbf{1}$. Since $\rho \not\stackrel{\tau}{\rightarrow}$, part (i) of Lemma 3.2 implies $\rho = \text{UNFOLD}(\rho)$, and so, now, an application of part (ii) of Lemma 4.2 ensures $\rho \not\stackrel{\checkmark}{\rightarrow}$.

We prove (7). The argument revolves around the unfolding of ρ_1 . To begin with, note two facts: one, the assumption $\rho_1 \preceq_{\text{CLT}}^{\text{syn}} \rho$ and Definition 4.20 imply $\text{UNFOLD}(\rho_1) \preceq_{\text{CLT}}^{\text{syn}} \rho$; and the other, the assumption $\rho_1 \dashv \sigma$ and Proposition 3.23 imply $\text{UNFOLD}(\rho_1) \dashv \sigma$.

The fact that $\rho \parallel \sigma \not\stackrel{\tau}{\rightarrow}$ can be used to prove

$$\alpha \in \mathbf{R} \text{ implies } \alpha \not\bowtie_c \beta \text{ for every } \beta \in \mathbf{S} \quad (8)$$

From the definition of acceptance set and $\rho_1 \xrightarrow{\tau}^* \text{UNFOLD}(\rho_1)$ (part (ii) of Lemma 3.2) it follows

$$\text{ACC}(\text{UNFOLD}(\rho_1)) \subseteq \text{ACC}(\rho_1) \quad (9)$$

Now we prove that $\text{UNFOLD}(\rho_1) = \mathbf{1}$. Fix a *stuck* derivative ρ'_1 of $\text{UNFOLD}(\rho_1)$:

$$\text{UNFOLD}(\rho_1) \xrightarrow{\tau}^* \rho'_1 \not\stackrel{\tau}{\rightarrow}$$

Such a stuck state exists because of the restricted syntax of session contracts. Further, since $\rho_1 \not\stackrel{\tau}{\rightarrow}$, by definition we have $\rho_1 \downarrow \mathbf{R}$ for some \mathbf{R} . Point (9) implies that $\mathbf{R} \in \text{ACC}(\rho_1)$, and so point (8), together with $\rho_1 \not\stackrel{\tau}{\rightarrow}$ and $\sigma \not\stackrel{\tau}{\rightarrow}$, implies that $\rho'_1 \parallel \sigma \not\stackrel{\tau}{\rightarrow}$. We have already seen that $\text{UNFOLD}(\rho_1) \dashv \sigma$, so part (ii) of Definition 3.17 now imply that $\rho'_1 \not\stackrel{\checkmark}{\rightarrow}$. We can now apply part (ii) of Lemma 4.2 to obtain $\text{UNFOLD}(\rho_1) = \mathbf{1}$.

As yet we have proven that (ρ, σ) respects part (i) of Definition 3.17. The argument to show that also part (ii) of Definition 3.17 holds is similar to the one used in Theorem 4.14. The difference amounts in the use of Definition 4.20 in place of Definition 4.8. We leave the details to the reader. \square

5. Modelling session types

The interpretation of session types as contracts is expressed as a function from the language $L_{\mathcal{ST}}$ in Section 2 to the language $L_{\mathcal{SC}}$ in Section 4. The function is little more than a syntactic transformation.

Let $\mathcal{M} : L_{\mathcal{ST}} \rightarrow L_{\mathcal{SC}}$ be defined by:

$$\mathcal{M}(S) = \begin{cases} \mathbf{1} & \text{if } S = \text{END} \\ !\mathbf{t}.\mathcal{M}(S) & \text{if } S = ![\mathbf{t}]; S \\ ?\mathbf{t}.\mathcal{M}(S) & \text{if } S = ?[\mathbf{t}]; S \\ \sum_{i \in [1;n]} ?\mathbf{l}_i.\mathcal{M}(S_i) & \text{if } S = \&\langle \mathbf{l}_1 : S_1, \dots, \mathbf{l}_n : S_n \rangle \\ \bigoplus_{i \in [1;n]} !\mathbf{l}_i.\mathcal{M}(S_i) & \text{if } S = \oplus\langle \mathbf{l}_1 : S_1, \dots, \mathbf{l}_n : S_n \rangle \\ \mu x.\mathcal{M}(S') & \text{if } S = \mu X.S' \\ x & \text{if } S = X \end{cases}$$

It is easy to see that \mathcal{M} maps session types, \mathcal{ST} , to session contracts, \mathcal{SC} ; indeed it defines a bijection between these sets:

- for every $\sigma \in \mathcal{SC}$ there exists some session type T such that $\mathcal{M}(T) = \sigma$
- if $\mathcal{M}(T_1) = \mathcal{M}(T_2)$ then $T_1 = T_2$

where $T_1 = T_2$ denotes syntactic identity. Further, substitution is preserved by \mathcal{M} .

Lemma 5.1. Let $S, T \in \mathcal{ST}$. Then $\mathcal{M}(S \{ T/X \}) = (\mathcal{M}(S)) \{ \mathcal{M}(T)/\mathcal{M}(X) \}$.

Proof. The proof is by structural induction on S . □

The interpretation also commutes with the two functions $depth(-)$ and $UNFOLD(-)$:

Lemma 5.2. For every $T \in \mathcal{ST}$ and $\sigma \in \mathcal{SC}$ the ensuing properties are true,

- (i) $dpt(T) = dpt(\mathcal{M}(T))$
- (ii) $UNFOLD(\mathcal{M}(T)) = \mathcal{M}(UNFOLD(T))$
- (iii) $UNFOLD(\mathcal{M}^{-1}(\sigma)) = T$ if and only if $UNFOLD(\sigma) = \mathcal{M}(T)$

Proof. The proofs of the first two points are by induction on $dpt(T)$, the proof of (ii) using (i) and the previous lemma. The third point follows immediately from (ii). □

As we have shown, the difficulty is to find a natural pre-order on session contracts which accurately reflects the sub-typing relation on session contracts. There are two obvious candidates, the restricted server pre-order and the restricted client pre-order on session contracts. The difficulty lies in the interpretation of END .

Example 5.3. Recall that $\mathcal{M}(\text{END}) = \mathbf{1}$. In the restricted server pre-order the session contract $\mathbf{1}$ is a least element, being smaller or equal to every other session contract. On the other hand, for session types $\text{END} \preceq_{\text{ST}} T$ if and only if $UNFOLD(T) = \text{END}$. Consequently the relation $\sqsubseteq_{\text{SRV}}^{\mathcal{SC}}$ is an unsound model for sub-typing between session types. For example:

$$\mathbf{1} \sqsubseteq_{\text{SRV}}^{\mathcal{SC}} !\mathbf{t}.\mathbf{1}, \quad \text{END} \not\preceq_{\text{ST}} ![\mathbf{t}]; \text{END}$$

The restricted client pre-order presents the dual issue as it relates every session contract to $\mathbf{1}$; it is one of the top element. Once again a model based on $\sqsubseteq_{\text{CLT}}^{\mathcal{SC}}$ would be unsound:

$$!\mathbf{t}.\mathbf{1} \sqsubseteq_{\text{CLT}}^{\mathcal{SC}} \mathbf{1}, \quad ![\mathbf{t}]; \text{END} \not\preceq_{\text{ST}} \text{END}$$

□

The main result of the paper is that the bijection \mathcal{M} gives a fully abstract interpretation of sub-typing between session types in terms of session contracts, provided we combine these two set-based pre-orders.

Definition 5.4. [Session contract pre-order]

For $\sigma_1, \sigma_2 \in \mathcal{SC}$ let $\sigma_1 \sqsubseteq^{\mathcal{SC}} \sigma_2$ whenever $\sigma_1 \sqsubseteq_{\text{SRV}}^{\mathcal{SC}} \sigma_2$ and $\sigma_1 \sqsubseteq_{\text{CLT}}^{\mathcal{SC}} \sigma_2$. \square

Example 5.5. It is instructive to see the behaviour of $\mathbf{1}$, the image of END under \mathcal{M} , relative to this combined pre-order. First suppose $\sigma \sqsubseteq^{\mathcal{SC}} \mathbf{1}$ for some session contract σ . This implies $\sigma \sqsubseteq_{\text{SRV}}^{\mathcal{SC}} \mathbf{1}$ and therefore, as we have shown in Proposition 4.7, σ must be a bottom element relative to $\sqsubseteq_{\text{SRV}}^{\mathcal{SC}}$ and $\text{UNFOLD}(\sigma)$ must be $\mathbf{1}$. A similar argument, using the pre-order $\sqsubseteq_{\text{CLT}}^{\mathcal{SC}}$ ensures that if $\mathbf{1} \sqsubseteq^{\mathcal{SC}} \sigma$ then $\text{UNFOLD}(\sigma)$ must also be $\mathbf{1}$.

In other words modulo unfolding the only session contract related to $\mathbf{1}$ via $\sqsubseteq^{\mathcal{SC}}$ is $\mathbf{1}$ itself. \square

Proposition 5.6 (Completeness).

For session contracts, $\sigma_1 \sqsubseteq^{\mathcal{SC}} \sigma_2$ implies $\mathcal{M}^{-1}(\sigma_1) \preceq_{\text{ST}} \mathcal{M}^{-1}(\sigma_2)$.

Proof. Let \mathcal{R} be the relation over session types defined by

$$\mathcal{R} = \{(\mathcal{M}^{-1}(\sigma_1), \mathcal{M}^{-1}(\sigma_2)) \mid \sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma_2, \sigma_1 \preceq_{\text{CLT}}^{\text{syn}} \sigma_2\}$$

If we prove that \mathcal{R} is a type simulation, that is it satisfies the properties given in Definition 2.10, then the result will follow because of Theorems 4.14 and 4.24.

The proof proceeds by a case analysis on the structure of $\text{UNFOLD}(\sigma_1)$; we give the details of two cases.

— Suppose $\text{UNFOLD}(\mathcal{M}^{-1}(\sigma_1)) = \text{END}$. According to Definition 2.10 we have to show that

$$\text{UNFOLD}(\mathcal{M}^{-1}(\sigma_2)) = \text{END}.$$

Because of part (iii) of Lemma 5.2 we know that $\text{UNFOLD}(\sigma_1) = \mathbf{1}$; moreover in Example 5.5 above we have already reasoned that $\text{UNFOLD}(\sigma_2)$ must be $\mathbf{1}$.

— Suppose $\text{UNFOLD}(\mathcal{M}^{-1}(\sigma_1)) = ![\mathbf{t}_1]; S_1$. We are required to prove that

$$\text{UNFOLD}(\mathcal{M}^{-1}(\sigma_2)) = ![\mathbf{t}_2]; S_2, \tag{10}$$

for some \mathbf{t}_2 and S_2 such that $\mathbf{t}_2 \preceq_{\mathbf{g}} \mathbf{t}_1$ and $(\mathcal{M}(S_1), \mathcal{M}(S_2)) \in \preceq_{\text{SRV}}^{\text{syn}} \cap \preceq_{\text{CLT}}^{\text{syn}}$.

Again by Lemma 5.2 (iii) we know that $\text{UNFOLD}(\sigma_1) = ![\mathbf{t}_1].\mathcal{M}(S_1)$. As $\sigma_1 \preceq_{\text{SRV}}^{\text{syn}} \sigma_2$, Definition 4.8 implies that $\text{UNFOLD}(\sigma_2) = ![\mathbf{t}_2].\sigma'_2$ for some \mathbf{t}_2 such that $\mathbf{t}_2 \preceq_{\mathbf{g}} \mathbf{t}_1$ and some σ'_2 such that $\mathcal{M}(S_1) \preceq_{\text{SRV}}^{\text{syn}} \sigma'_2$. Now letting S_2 denote $\mathcal{M}^{-1}(\sigma'_2)$, another application of Lemma 5.2 (iii) ensures that (10) above is satisfied. By the definition of S_2 we also have the requirement $\mathcal{M}(S_1) \preceq_{\text{SRV}}^{\text{syn}} \mathcal{M}(S_2)$.

It remains to show $\mathcal{M}(S_1) \preceq_{\text{CLT}}^{\text{syn}} \mathcal{M}(S_2)$. But this follows from $\sigma_1 \preceq_{\text{CLT}}^{\text{syn}} \sigma_2$, by part (iii) of Definition 4.20.

The proof for the remaining cases is similar to the argument already shown, and left to the reader. \square

Theorem 5.7. [Full abstraction]

For all session types, $T_1 \preceq_{\text{ST}} T_2$ if and only if $\mathcal{M}(T_1) \sqsubseteq^{\mathcal{SC}} \mathcal{M}(T_2)$.

Proof. Thanks to the completeness theorem, Theorem 5.6, it is sufficient to prove that $T_1 \preceq_{\text{ST}} T_2$ implies $\mathcal{M}(T_1) \sqsubseteq_{\text{SRV}}^{\text{SC}} \mathcal{M}(T_2)$ and $\mathcal{M}(T_1) \sqsubseteq_{\text{CLT}}^{\text{SC}} \mathcal{M}(T_2)$. As an example we outline the proof of the former. Because of Theorem 4.14 it is sufficient to show that the relation \mathcal{R} given by

$$\mathcal{R} = \{ (\sigma_1, \sigma_2) \mid \mathcal{M}^{-1}(\sigma_1) \preceq_{\text{ST}} \mathcal{M}^{-1}(\sigma_2) \}$$

is a syntactic sub-server relation, that is $\mathcal{R} \subseteq \mathcal{F}_{\preceq_{\text{SRV}}^{\text{syn}}}(\mathcal{R})$, where $\mathcal{F}_{\preceq_{\text{SRV}}^{\text{syn}}}$ is given in Definition 4.8.

Suppose $\sigma_1 \mathcal{R} \sigma_2$. The proof is a case analysis.

- If $\text{UNFOLD}(\sigma_1) = \mathbf{1}$ we have nothing to prove because condition (i) of Definition 4.8 does not require anything.
- If $\text{UNFOLD}(\sigma_1) = ?\mathbf{t}_1.\sigma'_1$ we have to show that

$$\text{UNFOLD}(\sigma_2) = ?\mathbf{t}_2.\sigma'_2$$

with $\mathbf{t}_1 \preceq_{\mathbf{g}} \mathbf{t}_2$ and $\sigma'_1 \mathcal{R} \sigma'_2$. An application of part (iii) of Lemma 5.2 shows that $\text{UNFOLD}(\mathcal{M}^{-1}(\sigma_1)) = ?[\mathbf{t}_1]; \mathcal{M}^{-1}(\sigma'_1)$. The fact that $\mathcal{M}^{-1}(\sigma_1) \preceq_{\text{ST}} \mathcal{M}^{-1}(\sigma_2)$ let us use Definition 2.10 to deduce that $\mathcal{M}^{-1}(\sigma_2) = ?[\mathbf{t}_2]; \mathcal{M}^{-1}(\sigma'_2)$ for some \mathbf{t}_2 such that $\mathbf{t}_1 \preceq_{\mathbf{g}} \mathbf{t}_2$, and some $\mathcal{M}^{-1}(\sigma'_2)$ such that $\mathcal{M}^{-1}(\sigma'_1) \preceq_{\text{ST}} \mathcal{M}^{-1}(\sigma'_2)$. From the last inequality and the definition of \mathcal{R} it follows that $\sigma'_1 \mathcal{R} \sigma'_2$. Since we have proven the conditions on the input actions $\mathbf{t}_1, \mathbf{t}_2$ and on the continuations σ'_1, σ'_2 we have left only to show that the structure of $\text{UNFOLD}(\sigma_2)$ is the required one; this follows from another application of part (iii) of Lemma 5.2.

The other cases are analogous and left to the reader. □

Corollary 5.8. The relation \sqsubseteq^{SC} is decidable.

Proof. To begin with, note that \mathcal{M} is defined by structural induction, so it is decidable. The corollary then follows from Corollary 2 of (Gay & Hole 2005), which ensures that the relation \preceq_{ST} is decidable, and our Theorem 5.7, whereby we can prove the isomorphism $\preceq_{\text{ST}} \cong \sqsubseteq^{\text{SC}}$. □

5.1. Examples and applications

In this subsection we give a series of examples in order to discuss the results we obtained. The first two example are of theoretical nature, whereas the last one shows an application.

Example 5.9. [Type simulations and the weak simulation relation]

At this stage, a natural question arises, which concerns the relationship between type simulations and weak simulations (Milner 1999). Assume the standard definition of the weak simulation (Milner 1999); we use the symbol \lesssim to denote the greatest weak simulation relation.

We begin by showing that, even though two session types are in a co-inductive types simulation, their images through \mathcal{M} need not be in a weak simulation. Consider the relation

$$\mathcal{R} = \{ (\oplus\langle \mathbf{1}_1 : \text{END}, \mathbf{1}_2 : \text{END} \rangle, (\oplus\langle \mathbf{1}_1 : \text{END} \rangle)), (\text{END}, \text{END}) \}$$

The standard co-inductive proof technique let one prove that the relation \mathcal{R} is a type simulation. On the other hand, the definition of \mathcal{M} implies that

$$\begin{aligned}\mathcal{M}(\oplus\langle \mathbf{1}_1 : \text{END}, \mathbf{1}_2 : \text{END} \rangle) &= !\mathbf{1}_1.\mathbf{1} \oplus !\mathbf{1}_2.\mathbf{1} \\ \mathcal{M}(\oplus\langle \mathbf{1}_1 : \text{END} \rangle) &= !\mathbf{1}_1.\mathbf{1}\end{aligned}$$

Then $\mathcal{M}(\oplus\langle \mathbf{1}_1 : \text{END}, \mathbf{1}_2 : \text{END} \rangle) \not\lesssim \mathcal{M}(\oplus\langle \mathbf{1}_1 : \text{END} \rangle)$ because $!\mathbf{1}_1.\mathbf{1} \oplus !\mathbf{1}_2.\mathbf{1} \xrightarrow{\tau} \xrightarrow{\mathbf{1}_2}$, while $!\mathbf{1}_1.\mathbf{1} \not\xrightarrow{\mathbf{1}_2}^*$. We have proven that $S_1 \preceq_{\text{ST}} S_2$ does not imply $\mathcal{M}(S_1) \lesssim \mathcal{M}(S_2)$.

Looking at the foregoing argument, one might be tempted to reason that if $S_1 \preceq_{\text{ST}} S_2$ then $\mathcal{M}(S_2) \lesssim \mathcal{M}(S_1)$. We prove that this is not the case. We can prove that

$$?\mathbf{1}_1.\mathbf{1} \sqsubseteq^{\text{SC}} ?\mathbf{1}_2.\mathbf{1} + ?\mathbf{1}_1.\mathbf{1}$$

An application of \mathcal{M}^{-1} gives us:

$$\begin{aligned}\mathcal{M}^{-1}(?\mathbf{1}_1.\mathbf{1}) &= \&\langle \mathbf{1}_1 : \text{END} \rangle \\ \mathcal{M}^{-1}(?\mathbf{1}_2.\mathbf{1} + ?\mathbf{1}_1.\mathbf{1}) &= \&\langle \mathbf{1}_1 : \text{END}, \mathbf{1}_2 : \text{END} \rangle\end{aligned}$$

A look at the definition of \preceq_{ST} , Definition 2.10, lets one prove that for every type simulation \mathcal{R}

$$(\&\langle \mathbf{1}_1 : \text{END}, \mathbf{1}_2 : \text{END} \rangle, \&\langle \mathbf{1}_1 : \text{END} \rangle) \notin \mathcal{R}$$

and, therefore,

$$\&\langle \mathbf{1}_1 : \text{END}, \mathbf{1}_2 : \text{END} \rangle \not\preceq_{\text{ST}} \&\langle \mathbf{1}_1 : \text{END} \rangle$$

□

Example 5.10. [e-vote, revisited]

In this example we use Theorem 5.7 in conjunction with Theorem 2 of (Gay & Hole 2005), in order to show how the set based pre-order \sqsubseteq^{SC} can be used to guarantee that a process P_a can be safely replaced by a suitable process P_b .

Consider two contracts *BallotA* and *BallotB* such that *BallotA* \sqsubseteq^{SC} *BallotB*. Let *BallotA* = $\mathcal{M}^{-1}(\text{BallotA})$ and *BallotB* = $\mathcal{M}^{-1}(\text{BallotB})$. From Theorem 5.7 it follows that

$$\text{BallotA} \preceq_{\text{ST}} \text{BallotB} \tag{11}$$

Let \perp_c denote the *coinductive duality relation* defined as in Definition 9 of (Gay & Hole 2005). Suppose now that $\text{BLTSRVA}(x^+)$, $\text{BLTSRVB}(x^+)$ and $\text{VOTER}(x^-)$ are pi calculus processes (as in (Gay & Hole 2005)) such that

$$\begin{aligned}\{x^+ : \text{BallotA}\} &\vdash \text{BLTSRVA}(x^+), \\ \{x^+ : \text{BallotB}\} &\vdash \text{BLTSRVB}(x^+), \\ \{x^- : \text{Voter}\} &\vdash \text{VOTER}(x^-)\end{aligned}$$

for some session type *Voter* such that *Voter* \perp_c *BallotA*. By means of the typing rules

of (Gay & Hole 2005), it is possible to derive

$$\frac{\frac{\vdots}{\{x^+ : \text{Ballot}A\} \vdash \text{BLTSRVA}(x^+)} \quad \frac{\vdots}{\{x^- : \text{Voter}\} \vdash \text{VOTER}(x^-)}}{\frac{\{x^+ : \text{Ballot}A\}, x^- : \text{Voter} \vdash \text{BLTSRVA}(x^+) \mid \text{VOTER}(x^-)}{\vdash (\nu x : \text{Ballot}A) \text{BLTSRVA}(x^+) \mid \text{VOTER}(x^-)}} \quad \begin{array}{l} \text{[T-PAR]} \\ \text{[T-NEWS]} \end{array}$$

Then (11) above and Theorem 2 of (Gay & Hole 2005) can be used to guarantee that if process $\text{BLTSRVB}(x^+)$ is used in place of process $\text{BLTSRVA}(x^+)$, then no communication error will happen along the channel x .

One can use non-recursive versions of the contracts seen in Examples 3.14 and 4.5 to obtain contracts that satisfy the assumptions above:

$$\begin{aligned} \text{Ballot}A &= ?\text{Login}.(!\text{Wrong}.1 \oplus !\text{Ok}.(? \text{Vote}A.1 + ? \text{Vote}B.1)) \\ \text{Ballot}B &= ?\text{Login}.(!\text{Wrong}.1 \oplus \\ &\quad !\text{Ok}.(? \text{Vote}A.1 + ? \text{Vote}B.1 + ? \text{Vote}C.1 + ? \text{Vote}D.1)) \\ \text{Voter} &= \mathcal{M}^{-1}(!\text{Login}.(? \text{Wrong}.1 + ? \text{Ok}.(! \text{Vote}A.1 \oplus ! \text{Vote}B.1))) \end{aligned}$$

□

Example 5.11. [Protocol conformance]

As already remarked, the language for contracts is a sublanguage of CCS without τ 's (Nicola & Hennessy 1987), and consequently contracts are suitable for specifying communication protocols.

Assume a protocol Pr to be specified by a contract σ , and let Q be a process (in the sense of (Gay & Hole 2005)), which is well-typed under the environment Γ . Assume also that $\Gamma(x) = S$ for some channel x .

We want to answer the following question:

(Q) “Does the session type S conform to the protocol specification σ ?”

Clearly, as long as the notion of conformance is not mathematically defined, it is not possible to give an answer (at least not a meaningful one).

In light of Theorem 5.7, we propose the following definition of conformance. Assume the standard definition of weak bisimilarity equivalence (Milner 1999); we denote this relation \approx . We say that a session type S conforms to a protocol specification σ if and only if $\mathcal{M}(S) \approx \sigma$.

To answer the question (Q) now one has only to prove that $\mathcal{M}(S) \approx \sigma$ or to show a counter example to this statement.

For example, if we had given a specification of the protocol POP3 (Rose 1988) with a contract σ , then we would have been able to check whether the session type POP3 of (Gay et al. 2003) conforms to σ .

In order for the notion of conformance we have given to be of any practical consequence, one last thing has to be ascertained. We have to prove that weak bisimilarity equivalence, when restricted to session contracts, is *decidable*. We leave this as an open problem worth further investigation. □

Behavioural	=	Co-inductive
\sqsubseteq_{SRV}	=	\preceq_{SRV}
$\sqsubseteq_{\text{SRV}}^{\text{SC}}$	=	$\preceq_{\text{SRV}}^{\text{syn}}$
$\sqsubseteq_{\text{CLT}}^{\text{SC}}$	=	$\preceq_{\text{CLT}}^{\text{syn}}$
\sqsubseteq_{SC}	\cong	\preceq_{ST}

Figure 10. Our behavioural pre-orders for contracts, the relations characterising them, and the isomorphism with sub-typing.

6. Conclusions

6.1. Summary

In this paper we have used contracts (Padovani 2010) to give a fully abstract model for first-order session types ordered by their sub-typing relation (Gay & Hole 2005).

In view of the interpretation \mathcal{M} , we have shown that the “natural” server refinement on contract (Definition 3.25) is not a *complete* model for the sub-typing, as it does not allow the refinement $a \sqsubseteq a + b$. Example 4.18 shows that also the “natural” client pre-order on contracts has the same issue. This has lead us to the identification of a subset of contracts which we call session contracts. We have then shown that the “natural” server and client refinements on session contracts are *unsound* models of the sub-typing (Example 5.3). This result explains why to obtain a fully-abstract model, it is *necessary* to introduce yet another pre-order, that is the intersection of the server and the client pre-orders on session contracts (Definition 5.4).

We believe that our work

- provides the first fully abstract model of session types in terms of contracts;
- shows the first alternative characterisation of the server and the client pre-orders on session contracts;
- contains the first published proof that shows the equality between the must pre-order and a refinement for servers based on the compliance relation; moreover, it shows that the equality holds regardless of the \bowtie relation used to let contracts interact.

The table in Figure 10 sketches our knowledge on the refinements that we have investigated. By means of examples we have shown that the relations in different rows are different.

If we bear in mind the the debate on the “right” notion of compliance for contracts, Section 3.3 and Corollary 3.51 are worthy of comment. The result states that as long as we are concerned with refinements for the servers, it does not matter whether we pick as definition of compliance Definition 3.17 or Definition 3.40, because the resulting refinements are the same. This means also that the must pre-order on contracts does not provide a complete model for the sub-typing.

6.2. Related work

Roughly speaking, the refinements for contracts that have been proposed thus far in the literature can be related either to the must testing (Nicola & Hennessy 1984), or to the fair testing (Rensink & Vogler 2007). According to this criterion, we divide this section in two parts; first we compare the refinements we have investigated with the pre-orders used in the papers that have influenced us most, namely (Barbanera & de'Liguoro 2010, Laneve & Padovani 2007, 2008, Padovani 2010). The theories presented in these papers are related the must pre-order; thus we will refer to them as *must-theories*. Afterwards, we compare our work with two theories inspired by the fair testing, which are given in (Bravetti & Zavattaro 2009, Padovani 2011). We refer to those theories as *fair-theories*. As we will see, the must-theories bear some similarities with our results; whereas the fair-theories turns out to provide refinements that are different from the ones we have studied.

From now on we reason under the assumption that in our definition of compliance the synchronisation relation \bowtie_i is used.

6.2.1. Must-theories The comparison with (Laneve & Padovani 2007, 2008) is complicated by the fact that in these papers compliance judgements take the form $I_1[\rho] \dashv I_2[\sigma]$ where I_1, I_2 are finite sets of actions representing in some sense the interfaces of the processes guaranteeing the contracts; moreover, for a contract $I[\sigma]$ to be valid its interface I has to contain all the action names (including \checkmark) that appear in the behaviour σ . Let us refer to the pairs $I[\sigma]$ as *constrained contracts*.

In (Laneve & Padovani 2007) a theorem is proven, which resembles our Corollary 3.51:

Theorem 2 Let $I = \text{names}(\tau)$. $I[\sigma] \preceq^{lp07} I[\tau]$ if and only if $\sigma \sqsubseteq_{\text{MUST}} \tau$.

This theorem is weaker than Corollary 3.51. First of all, as constrained contracts are pairs (ie. a set of actions and a behaviour), the server refinement \preceq^{lp07} (Laneve & Padovani 2007, see Definition 2) *cannot* be directly compared with the must pre-order; formally

$$\preceq^{lp07} \not\subseteq \sqsubseteq_{\text{MUST}}, \quad \sqsubseteq_{\text{MUST}} \not\subseteq \preceq^{lp07} \quad (12)$$

One may argue that this has no significance, as it is still easy to prove that

$$\text{if } \sigma_1 \sqsubseteq_{\text{MUST}} \sigma_2 \text{ then } \text{names}(\sigma_2)[\sigma_1] \preceq^{lp07} \text{names}(\sigma_2)[\sigma_2]$$

The converse implication, though, relies heavily on the interfaces of the constrained contracts at hand, and in general is not true:

$$I[\sigma] \preceq^{07} I'[\sigma'] \text{ does not imply that } \sigma \sqsubseteq_{\text{MUST}} \sigma' \quad (13)$$

For instance, we can prove the following

$$\begin{array}{l} \emptyset[\text{NIL}] \preceq^{lp07} \{!a\}![a.\text{NIL}] \\ \text{NIL} \not\subseteq_{\text{MUST}} !a.\text{NIL} \end{array}$$

It follows that *in the sense shown above* the pre-order \preceq^{lp07} is coarser than the must pre-order.

Since the relation \preceq^{lp07} cannot be compared directly with $\sqsubseteq_{\text{MUST}}$ (12), and is somehow coarser than $\sqsubseteq_{\text{MUST}}$ (13), to the best of our knowledge, our Corollary 3.51 is the first

published proof of the equality between a first-order compliance-based refinement and a testing based must-preorder.

Also there is a difference between our compliance relation and the compliance of (Laneve & Padovani 2007). The compliance used in (Laneve & Padovani 2007), which we denote \dashv^{lp07} is related to our compliance in the sense that

Proposition. If $I[\rho] \dashv^{lp07} J[\sigma]$ for some I, J then $\rho \dashv \sigma$

The converse of the proposition is not true; we explain why. We have $?a. \checkmark . \text{NIL} \dashv !a. \checkmark . \text{NIL}$, while $\checkmark \in \text{names}(!a. \checkmark . \text{NIL})$ thus $?a. \checkmark . \text{NIL} \not\vdash^{lp07} !a. \checkmark . \text{NIL}$, because Definition 1 of (Laneve & Padovani 2007) requires that \checkmark be not among the names of the contract on the right hand side of the compliance.

Now we compare our work with (Laneve & Padovani 2008); we denote \dashv^{lp08} the compliance relation given by (Laneve & Padovani 2008, Definition 1), and \preceq^{lp08} the pre-order given there in Definition 2.

Our compliance relation and \dashv^{lp08} are related in a way analogous to what we have seen for \dashv^{07} .

Proposition. If $I[\rho] \dashv^{lp08} J[\sigma]$ for some I, J then $\rho \dashv \sigma$

The converse is not true, $?a. \checkmark . \text{NIL} + ?b. \checkmark . \text{NIL} \dashv !a. \checkmark . \text{NIL}$, and from $\{a, b\} \not\subseteq \{a\}$ it follows

$$\{a, b, \checkmark\}[?a. \checkmark . \text{NIL} + ?b. \checkmark . \text{NIL}] \not\vdash^{lp08} \{a, \checkmark\}![a. \checkmark . \text{NIL}].$$

We can prove that the relation \preceq^{lp08} extends \preceq^{07} to the terms that contain \checkmark ; for instance we can prove that

$$\begin{aligned} \{\checkmark\}[\checkmark . \text{NIL}] &\preceq^{lp08} \emptyset[\text{NIL}] \\ \{\checkmark\}[\checkmark . \text{NIL}] &\preceq^{lp08} \{\ell, \checkmark\}[\ell. \checkmark . \text{NIL}] \end{aligned}$$

(Laneve & Padovani 2008) presents the first comparison between session types and contracts; in particular, it is shown that the subcontract relation \preceq^{lp08} together with two interpretations similar to \mathcal{M} provide two sound models for the subtyping on a subset of our session types. These interpretations are denoted $\llbracket - \rrbracket_{\mathbf{1}}$ and $\llbracket - \rrbracket_{\text{NIL}}$. Their proposed full abstraction result, (Laneve & Padovani 2008, see Theorem 2), though, appears not to be true. According to that definition and the interpretation $\llbracket - \rrbracket_{\text{NIL}}$

$$\emptyset[\text{NIL}] \preceq^{lp08} \{\ell\}[\ell. \text{NIL}]$$

Their Theorem 2 therefore implies $\&\langle \ell : \text{END} \rangle \preceq_{\text{ST}} \text{END}$, which is not true. On the other hand if $\llbracket - \rrbracket_{\mathbf{1}}$ is used then there are two issues. According to Theorem 2 the pair $(\text{END}, \&\langle \ell : \text{END} \rangle)$ is interpreted as $(\emptyset[\checkmark . \text{NIL}], \{\ell\}[\ell. \checkmark . \text{NIL}])$. Then

- (a) neither $\emptyset[\checkmark . \text{NIL}]$ nor $\{\ell\}[\ell. \checkmark . \text{NIL}]$ are constrained contracts, because their interfaces do not contain all the action names which appear in the respective behaviours; moreover
- (b) even if the interpretation was correct, Theorem 2 would be false because

$$\{\checkmark\}[\checkmark . \text{NIL}] \preceq^{lp08} \{\ell, \checkmark\}[\ell. \checkmark . \text{NIL}]$$

while, as stated above, $\&\langle \ell : \text{END} \rangle \preceq_{\text{ST}} \text{END}$ is not true.

Our study of the restricted pre-orders on session contracts (Definition 4.3 and Definition 4.16) is clearly inspired by (Barbanera & de'Liguoro 2010). In (Barbanera & de'Liguoro 2010) the language for session types is the same one as we used, whereas the subset of contracts in which session types are embedded is the set of *session behaviours*. The set of session behaviours is bigger than the set of session contracts because of the lack of distinction between labels and base types. It is possible to write session behaviours as

$$?Int.1+?1_1.1$$

which are image of no session type according to the given interpretation $\llbracket - \rrbracket$. Note, though, that $\llbracket - \rrbracket = \mathcal{M}$, so the range of $\llbracket - \rrbracket$ is the set of session contracts, and our Theorem 5.7 proves that $\llbracket - \rrbracket$ and their pre-order \preceq : (Barbanera & de'Liguoro 2010, Definition 3.4) provides a complete model for the sub-typing. The completeness of \preceq : was only conjectured in (Barbanera & de'Liguoro 2010).

Their approach is complementary to ours, in that they provide a co-inductive characterisation of the pre-order \preceq :, which turns out to equal the intersection of their sub-server and sub-client pre-orders. In contrast, we have studied the (restricted) server and the client pre-orders independently, providing their co-inductive characterisations; we have then (1) explained why it is necessary to use the intersection of the two pre-orders to obtain a fully-abstract model; and (2) proven that the intersection of these pre-orders is a sound *and complete* model of the sub-typing.

The comparison with (Padovani 2010) is straightforward. Our definition of compliance and the definition in that paper (Definition 2.1) are different; nevertheless, we can prove that the resulting relations coincide (as long as we instantiate \bowtie to the relation \bowtie_i). As a consequence, also our server pre-order $\sqsubseteq_{\text{SRV}}^{\bowtie_i}$ coincides with the strong subcontract relation (Padovani 2010, see Definition 2.2).

6.2.2. Fair-theories In (Bravetti & Zavattaro 2009, Padovani 2011) the notion of correctness requires *all* the components of a composition to be successful (ie. be able of performing \checkmark) *at the same time* in order for the whole composition to be successful. This requirement implies that the compositions which contain terms as

$$NIL, \quad \mu x. a.x$$

cannot be correct, because the contracts above do not perform \checkmark at all. This phenomenon renders the viability of contracts (Padovani 2011, Definition 3.1) a non trivial matter; on the contrary, in our theory every contract is viable wrt. MUST and \neg .

The pre-orders on session contracts and session types that we have studied allow the refinement

$$a \oplus b \sqsubseteq a \tag{14}$$

For instance we can prove the following facts

$$\begin{array}{l} \mu x. (!\text{espresso}.x \oplus !\text{moka}.1) \quad \sqsubseteq^{\text{SC}} \quad \mu x. !\text{espresso}.x \\ \mu X. \oplus \langle \text{livelock}: X, \text{stop}: \text{END} \rangle \quad \preceq_{\text{ST}} \quad \mu X. \oplus \langle \text{livelock}: X \rangle \end{array}$$

In (Padovani 2011) it is pointed out that the refinements shown above are not sound

with respect to the fair-testing (Rensink & Vogler 2007); this will let us prove that the refinements proposed in (Bravetti & Zavattaro 2009, Padovani 2011) are not contained in our relations \sqsubseteq_{SRV} and \sqsubseteq^{SC} .

Now we give the detailed comparisons with the mentioned papers. The language used in (Padovani 2011) is similar to our session contracts, the differences being that actions are decorated with a role tag $\mathbf{p}, \mathbf{q}, \dots$; and there is a special session type FAIL. Also, sessions are *multiparty*, that is they are general compositions of session types (tagged with a role), for instance

$$\mathbf{p}_1 : T_1 \parallel \mathbf{p}_2 : T_1 \parallel \dots \parallel \mathbf{p}_k : T_k$$

The notion of correct session type composition is given in (Padovani 2011, Definition 2.1), and it is used to define a set-theoretical sub-typing relation on session types (Padovani 2011, Definition 2.2), which is denoted \leq . We can prove

$$\mu x. (\mathbf{p!livelock}.x \oplus \mathbf{p!stop}.1) \not\leq \mu x. \mathbf{p!livelock}.x$$

because the term $\mu x. \mathbf{p!livelock}.x$ cannot reach a successful state at all. This means that (14) is not sound for \leq .

Also the following facts are true:

$$\begin{array}{l} \mu x. \mathbf{p!livelock}.x \leq \mu x. \mathbf{p!stop}.x \\ \mu x. \mathbf{p!livelock}.x \not\sqsubseteq^{SC} \mu x. \mathbf{p!stop}.x \end{array}$$

The first fact is true because no composition containing the session type $\mu x. \mathbf{p!livelock}.x$ can be correct, as this term does not perform \checkmark at all; whereas the second fact follows from the definition of \mathcal{M} and Definition 2.10. It thus follows that

$$\leq \not\sqsubseteq \sqsubseteq^{SC}, \quad \sqsubseteq^{SC} \not\sqsubseteq \leq$$

Similar to (Padovani 2011), in (Bravetti & Zavattaro 2009) general compositions of contracts are allowed

$$[C_1] \parallel [C_2] \parallel \dots \parallel [C_n]$$

and for a composition to be successful *all* its components have to be successful (ie. be able of performing \checkmark) *at the same time*.

As for the contract language, the main difference with our framework is that the theory of (Bravetti & Zavattaro 2009) involves output persistent contracts (Bravetti & Zavattaro 2009, see Definition 4); for instance the term $!a.1 + b.1$ is a contract in our theory, that is ruled out in (Bravetti & Zavattaro 2009), as it is not output persistent.

Definition 12 of that paper introduces the subcontract relations \preceq_O on output persistent contracts, where the parameter O is a the set of *output* actions that the compositions used as tests can show. The comparison between our server pre-orders and the pre-orders \preceq_O is complicated by two aspects,

- if $C' \preceq_O C$ then it is safe to use C' in place of C ; in view of this, we will compare our pre-orders with *the inverse* of the pre-orders \preceq_O ;
- a priori, it is not clear how to choose the parameter O . To solve this complication we treat \preceq_O as a function of O , and briefly discuss its monotonicity.

The function \preceq_O is not monotonically increasing, as $\emptyset \subseteq \{\bar{a}\}$, while

$$\begin{array}{l} a.\text{NIL} \preceq_{\emptyset} a.\mathbf{1} \\ a.\text{NIL} \not\preceq_{\{\bar{a}\}} a.\mathbf{1} \end{array}$$

On the other hand \preceq_O is monotonically *decreasing* in O .

Proposition. If $O \subseteq O'$ then $\preceq_{O'} \subseteq \preceq_O$.

This proposition gives us two criteria to reason on all the pre-orders \preceq_O :

- for every O the pairs in $\preceq_{\mathcal{N}}$ are in \preceq_O , where $\mathcal{N} = \{\bar{a} \mid a \in \text{Act}\}$;
- for every O the pairs *not* in \preceq_{\emptyset} are *not* in \preceq_O .

As for the restriction on output persistent contracts, in the oncoming discussion we will use only contracts that enjoy that property; thus our arguments are sound.

We have the following facts

$$\begin{array}{l} \mathbf{1} + a.\mathbf{1} \sqsubseteq_{\text{SRV}} a.\mathbf{1} \\ \mathbf{1} + a.\mathbf{1} \not\preceq_{\emptyset}^{-1} a.\mathbf{1} \end{array}$$

where the test used to prove the second fact is $\mathbf{1}$. Also the ensuing facts are true,

$$\begin{array}{l} a.\text{NIL} \preceq_{\mathcal{N}}^{-1} \text{NIL} \\ a.\text{NIL} \not\sqsubseteq_{\text{SRV}} \text{NIL} \end{array}$$

The first fact holds because no system containing $a.\text{NIL}$ can be correct. The test that we use to prove the second fact is $\bar{a}.\mathbf{1}$. Thus far we have proven

$$\sqsubseteq_{\text{SRV}} \not\subseteq \preceq_{\emptyset}, \quad \preceq_{\mathcal{N}} \not\subseteq \sqsubseteq_{\text{SRV}}$$

Similarly to what done for the relation \leq of (Padovani 2011), we can prove also that

$$\mu x. (\tau.!\text{livelock}.x + \tau.!\text{stop}.\mathbf{1}) \not\preceq_{\emptyset}^{-1} \mu x.!\text{livelock}.x$$

Thus our relation \sqsubseteq^{SC} is coarser than \preceq_{\emptyset} . As $\preceq_{\mathcal{N}}$ relates also terms more general than session contracts, we have the following facts

$$\sqsubseteq^{\text{SC}} \not\subseteq \preceq_{\emptyset}, \quad \preceq_{\mathcal{N}} \not\subseteq \sqsubseteq^{\text{SC}}$$

6.3. Future work

Models for session types and sub-typing Session types haven been originally used to type dialects of the pi-calculus (Honda et al. 1998). Recently session types have been proposed also to type other formalisms, for instance orchestration charts (Fantechi & Najm 2008). In that paper session types have behaviours described by typecharts, which are essentially LTS's (see Section 4.1 there), and the sub-typing \preceq , is defined in the following simulation-like manner:

$T_1 \preceq T_2$ if and only if the following conditions are true,

- if $T_2 \xrightarrow{?m} T'_2$ implies that there exists a T'_1 such that $T_1 \xrightarrow{?m} T'_1$ and $T'_1 \preceq T'_2$
- if $T_1 \xrightarrow{!m} T'_1$ implies that there exists a T'_2 such that $T_2 \xrightarrow{!m} T'_2$ and $T'_1 \preceq T'_2$

In view of Example 5.9 it should be clear that the relation \sqsubseteq^{SC} does not model \preceq . We leave as an open problem the use of \sqsubseteq^{SC} to model the inverse of \preceq ; the complication being that the definition above does not account for the τ actions.

Higher-order language The restriction to first-order session types is a severe limitation on our results, and we intend to extend them to the full language of session types in (Gay & Hole 2005). To this end it will be necessary to use a higher-order version of the language that we have used in the paper; that fact that we have used the parameter \bowtie to express the notion of co-action will be of some help here.

Refinements for clients By and large, the results in the literature on contracts for web-services pertain to refinements for either the server side of binary connections, or peers of multi-party connections. To the best of our knowledge, the first paper that avails of a refinement for clients is (Barbanera & de’Liguoro 2010). Assuming the obvious definition of the refinements for the clients due to the must testing and to the compliance relation, denoted respectively \sqsubseteq_{MUST}^{CLT} and \sqsubseteq_{CLT} , we can show that in our framework these pre-orders differ from the respective refinement for servers,

$$\begin{array}{ccccc} \sqsubseteq_{MUST}^{CLT} & \not\sqsubseteq & \sqsubseteq_{MUST} & \not\sqsubseteq & \sqsubseteq_{MUST}^{CLT} \\ \sqsubseteq_{CLT} & \not\sqsubseteq & \sqsubseteq_{SRV} & \not\sqsubseteq & \sqsubseteq_{CLT} \end{array}$$

These facts call for an investigation of the client refinements \sqsubseteq_{MUST}^{CLT} and \sqsubseteq_{CLT} , which we leave for future work.

Decidability refinements In Corollary 5.8 have briefly discussed the decidability of the session pre-order. We leave as open problem the investigation of the decidability of the other refinements we have discussed.

Acknowledgements The authors wish to acknowledge the anonymous reviewers for the criticisms and the suggestions on the previous drafts of this paper.

References

- Alonso, G., Casati, F., Kuno, H. A. & Machiraju, V. (2004), *Web Services - Concepts, Architectures and Applications*, Data-Centric Systems and Applications, Springer.
- Barbanera, F. & de’Liguoro, U. (2010), Two notions of sub-behaviour for session-based client/server systems, *in* T. Kutsia, W. Schreiner & M. Fernández, eds, ‘PPDP’, ACM, pp. 155–164.
- Bernardi, G., Bugliesi, M., Macedonio, D. & Rossi, S. (2008), A theory of adaptable contract-based service composition, *in* V. Negru, T. Jebelean, D. Petcu & D. Zaharie, eds, ‘SYNASC’, IEEE Computer Society, pp. 327–334.
- Bernardo, M., Padovani, L. & Zavattaro, G., eds (2009), *Formal Methods for Web Services, 9th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2009, Bertinoro, Italy, June 1-6, 2009, Advanced Lectures*, Vol. 5569 of *Lecture Notes in Computer Science*, Springer.

- Bravetti, M. & Zavattaro, G. (2009), Contract-based discovery and composition of web services, *in* Bernardo et al. (2009), pp. 261–295.
- Caires, L. & Pfenning, F. (2010), Session types as intuitionistic linear propositions, *in* P. Gastin & F. Laroussinie, eds, ‘CONCUR’, Vol. 6269 of *Lecture Notes in Computer Science*, Springer, pp. 222–236.
- Carpinetti, S., Castagna, G., Laneve, C. & Padovani, L. (2006), A formal account of contracts for web services, *in* M. Bravetti, M. Núñez & G. Zavattaro, eds, ‘WS-FM’, Vol. 4184 of *Lecture Notes in Computer Science*, Springer, pp. 148–162.
- Castagna, G., Gesbert, N. & Padovani, L. (2009), ‘A theory of contracts for web services’, *ACM Trans. Program. Lang. Syst.* **31**(5), 1–61. Supersedes the article in POPL ’08.
- Eshuis, R. & Fokkinga, M. M. (2002), ‘Comparing refinements for failure and bisimulation semantics’, *Fundam. Inform.* **52**(4), 297–321.
- Fantechi, A. & Najm, E. (2008), Session types for orchestration charts, *in* D. Lea & G. Zavattaro, eds, ‘COORDINATION’, Vol. 5052 of *Lecture Notes in Computer Science*, Springer, pp. 117–134.
- Gay, S. J. & Hole, M. (2005), ‘Subtyping for session types in the pi calculus’, *Acta Inf.* **42**(2-3), 191–225.
- Gay, S., Vasconcelos, V. & Ravara, A. (2003), Session types for inter-process communication, Technical Report TR-2003-133, Department of Computing Science, University of Glasgow.
- Honda, K., Vasconcelos, V. T. & Kubo, M. (1998), Language primitives and type discipline for structured communication-based programming, *in* C. Hankin, ed., ‘ESOP’, Vol. 1381 of *Lecture Notes in Computer Science*, Springer, pp. 122–138.
- Laneve, C. & Padovani, L. (2007), The must preorder revisited, *in* ‘Proceedings of the 18th international conference on Concurrency Theory’, Springer-Verlag, Berlin, Heidelberg, pp. 212–225.
- Laneve, C. & Padovani, L. (2008), The pairing of contracts and session types, *in* P. Degano, R. D. Nicola & J. Meseguer, eds, ‘Concurrency, Graphs and Models’, Vol. 5065 of *Lecture Notes in Computer Science*, Springer, pp. 681–700.
- McNeile, A. (2010), A framework for the semantics of behavioral contracts, *in* ‘Proceedings of the Second International Workshop on Behaviour Modelling: Foundation and Applications’, BM-FA ’10, ACM, New York, NY, USA, pp. 3:1–3:5.
- Meyer, B. (1997), *Object-Oriented Software Construction, 2nd Edition*, Prentice-Hall.
- Milner, R. (1999), *Communicating and mobile systems - the Pi-calculus*, Cambridge University Press.
- Nicola, R. D. & Hennessy, M. (1984), ‘Testing equivalences for processes’, *Theoretical Computer Science* **34**, 83–133.
- Nicola, R. D. & Hennessy, M. (1987), CCS without τ ’s, *in* ‘TAPSOFT, Vol.1’, Vol. 249 of *Lecture Notes in Computer Science*, Springer, pp. 138–152.
- Oasis Standard (2011), ‘Universal Description, Discovery, and Integration’.
URL: <http://uddi.xml.org/>
- Padovani, L. (2010), ‘Contract-based discovery of web services modulo simple orchestrators’, *Theor. Comput. Sci.* **411**(37), 3328–3347.

- Padovani, L. (2011), Fair Subtyping for Multi-Party Session Types, *in* ‘Proceedings of the 13th Conference on Coordination Models and Languages’, Vol. LNCS 6721, Springer, pp. 127–141.
- Pierce, B. C. & Sangiorgi, D. (1996), ‘Typing and subtyping for mobile processes’, *Mathematical Structures in Computer Science* **6**(5), 409–453.
- Rensink, A. & Vogler, W. (2007), ‘Fair testing’, *Information and Computation* **205**(2), 125–198.
- Rose, M. (1988), ‘Post Office Protocol: Version 3’, RFC 1081. Obsoleted by RFC 1225.