

The pairing of contracts and session types

Cosimo Laneve¹ and Luca Padovani²

¹ Department of Computer Science, University of Bologna

² Information Science and Technology Institute, University of Urbino

Dedicated to Ugo Montanari in occasion of his 65th birthday

Abstract. We pair session types and contracts using two encodings. The encoding of session types accommodates width and depth subtyping, two properties that partially hold in contracts. The encoding of contracts accommodates complex synchronization patterns, since session types own a simple control protocol. The encodings allow one to use the two formalisms interchangeably, within the context of dyadic interactions.

1 Introduction

Service Oriented technologies and Web Services have been recently proposed as a new way of distributing and organizing complex applications across the Internet. The success of these technologies has fostered the development of formal methods for statically analyzing and verifying the behavior of concurrent and distributed systems. Two such methods – session types [13, 10, 9] and contracts [3, 14, 4] – aim at describing the communication protocol implemented by services. These methods provide a foundation for statically checking that a process implements a given communication protocol and formally characterize *compliance* (when a client interacts successfully with a service) and *safe replacement* (when it is possible to replace a process with another one).

Session types and contracts find their origins in two different domains: the former ones derive from the domain of type theory and type systems, whereas contracts are more related to the study of behavioral equivalences, such as bisimulation and testing equivalence [6, 12]. Both languages are equipped with similar constructors. For example, the session type

$$S \stackrel{\text{def}}{=} \mu x. \& \langle \text{Login} : \oplus \langle \overline{\text{Wrong}} : x; \text{Ok} : \& \langle \text{VoteA} : \text{end}; \text{VoteB} : \text{end} \rangle; \text{Cheat} : \text{end} \rangle \rangle$$

and the contract

$$\mathbf{I}[\sigma] \stackrel{\text{def}}{=} \mathbf{I}[\text{rec } x. \text{Login}. (\overline{\text{Wrong}}.x \oplus \overline{\text{Ok}}. (\text{VoteA} + \text{VoteB}) \oplus \overline{\text{Cheat}})]$$

where $\mathbf{I} = \{\text{Login}, \overline{\text{Wrong}}, \overline{\text{Ok}}, \overline{\text{Cheat}}, \text{VoteA}, \text{VoteB}\}$, represent a simple service for an online ballot between two candidates A and B. Before a client is allowed to vote, he must provide a valid login token that the system uses for ensuring that preferences are expressed at most once, for otherwise the voter is identified as a cheater.

In contracts, the actions such as Login and $\overline{\text{Cheat}}$ represent atomic communications between the voter and the service and we have two binary operators for encoding

alternatives: a $+$ indicates an *external choice* (the voter chooses the candidate) whereas a \oplus indicates an *internal choice* (the service decides whether the login is valid or not). In session types we have *branches* $\&\langle \dots \rangle$ and *choices* $\oplus\langle \dots \rangle$ that play the same roles played by $+$ and \oplus . However, branches indicate the receipt of a label expressing the decision of the voter, whereas choices indicate that the service emits a label expressing the alternative it has chosen. Hence, the direction of the exchanged messages is encoded in the type of alternative, rather than in the labels themselves. Also, while actions in contracts may encode the exchange of data between the voter and the service, labels in session types are only meant to implement the control part of the protocol. In fact, typical presentations of session types admit additional constructs for representing the exchange of data, which we omit in this paper as they are irrelevant for the results that follow.

The relationship between session types and contracts also regards the semantics, although, in this case, their different origin is evident because the relations are mostly the opposite. For example, the session type

$$T \stackrel{\text{def}}{=} \&\langle \text{Login} : \oplus\langle \text{Ok} : \&\langle \text{VoteA} : \text{end}; \text{VoteB} : \text{end} \rangle \rangle \rangle$$

and the contract

$$I[\tau] \stackrel{\text{def}}{=} I[\text{Login}.\overline{\text{Ok}}.(\text{VoteA} + \text{VoteB})]$$

both describe a ballot service that does not care about cheaters and always accepts login tokens regardless their validity. It turns out that T is a *subtype* S (notation $T \leq S$) and $I[\sigma]$ is a *subcontract* of $I[\tau]$ (notation $I[\sigma] \preceq I[\tau]$). That is, the subtype relation $T \leq S$ embodies the notion of safe substitutability: every term having type S may be replaced by a term having type T without affecting the context in a sensible way. The subcontract relation $I[\sigma] \preceq I[\tau]$ embodies the notion of successful interaction (called *compliance*): the set of clients succeeding in interacting with a service of contract $I[\sigma]$ also succeed with a service of contract $I[\tau]$.

There are also some differences between the theory of session types and the one of contracts. In one direction, these differences mainly regard the so-called *width* and *depth subtyping*. Session types, much alike object-oriented type disciplines, enjoy the property $\&\langle \text{VoteA} : \text{end}; \text{VoteB} : \text{end} \rangle \leq \&\langle \text{VoteA} : \text{end} \rangle$ (the ballot service can be extended with more candidates without invalidating former voters), namely width extensions of capabilities is always possible. In contracts this is not the case. To discuss the point, consider the contracts $I[\sigma_1] = \{\text{VoteA}, \text{VoteB}\}[\text{VoteA}]$ and $I[\sigma_2] = \{\text{VoteA}, \text{VoteB}\}[\text{VoteA} + \text{VoteB}]$ and the client $K[\rho] = \{\overline{\text{VoteA}}, \overline{\text{VoteB}}, e\}[\overline{\text{VoteA}}.e + \overline{\text{VoteB}}.\overline{\text{VoteB}}.e]$. Such client tries to vote for candidate A once or for B twice. It easy to verify that $K[\rho]$ successfully interacts with $I[\sigma_1]$ whilst the interaction may fail with $I[\sigma_2]$, therefore $I[\sigma_1] \not\preceq I[\sigma_2]$. In contracts width subtyping is admitted provided the additional capabilities are not present in the interface of the smaller contract. For example $\{\text{VoteA}\}[\text{VoteA}] \preceq \{\text{VoteA}, \text{VoteB}\}[\text{VoteA} + \text{VoteB}]$ (the client $K[\rho]$ is not a valid client for $\{\text{VoteA}\}[\text{VoteA}]$ because it has a larger interface). Similar arguments may be given for depth subtyping, the property guaranteeing the safety of the replacement of a service with another one providing a *longer* communication protocol. In the other direction, the differences follow by the fact that session types embody the property that sessions are supposed to be completed *symmetrically* by both parties, whilst

contracts are biased towards clients, which are free to interrupt the interaction any time they please. Another difference is that contracts describe a more abstract synchronization pattern than session types do. For example, in the theory of contracts, a client such as

$$\{\overline{\text{Login}}, \text{Ok}, \overline{\text{VoteA}}, \overline{\text{VoteB}}, e\}[\overline{\text{Login}}.\text{Ok}.(\overline{\text{VoteA}}.e + \overline{\text{VoteB}}.e)]$$

may successfully interact with a service as $I[\tau]$, while, in session types, a branch cannot be matched by another branch. That is, session types describe a communication in which no handshaking between the interacting parties ever occurs. There is always exactly one party having control, and this party has to explicitly notify the other one about the (internal) choices it has made.

In this contribution we undertake a thorough comparison between session types and contracts for assessing a precise relationship between the two formalisms. We define two encodings, one from session types to contracts, and the other from contracts to session types. These encodings allows one to use the two formalisms interchangeably, without losing any relevant information. Therefore it is possible to argue about session types by means of the subcontract relation and, conversely, about contracts by using the deductive system of the subtyping relation. However, because of the differences between the two theories, the two encodings are not one the converse of the other. Let us discuss this issue with few examples. We encode the session type

$$T' \stackrel{\text{def}}{=} \&\langle \text{Login} : \oplus \langle \text{Ok} : \text{end}; \text{Cheat} : \text{end} \rangle \rangle$$

into the contract

$$I'[\tau'] \stackrel{\text{def}}{=} [\text{Login}, \overline{\text{Ok}}, \overline{\text{Cheat}}](\text{Login}.(\overline{\text{Ok}} \oplus \overline{\text{Cheat}}) + \overline{\text{Ok}}.\Omega + \overline{\text{Cheat}}.\Omega)$$

where the terms $\overline{\text{Ok}}.\Omega$ and $\overline{\text{Cheat}}.\Omega$ have been added in order to enforce width subtyping in the contract. If a client of $I'[\tau']$ attempts actions that are not explicitly allowed by T' then a catastrophic state – Ω – is reached, meaning that, in practice, such actions are *not* guaranteed.

To illustrate the encoding of contracts into session types we discuss the encoding of $\{\text{VoteA}, \text{VoteB}\}[\text{VoteA} \oplus \text{VoteB}]$, which eventually generates the session type

$$\begin{aligned} &\oplus \langle \{\text{VoteA}\} : \&\langle \emptyset : \text{end}; \{\text{VoteA}\} : \text{end} \rangle; \\ &\quad \{\text{VoteB}\} : \&\langle \emptyset : \text{end}; \{\text{VoteB}\} : \text{end} \rangle; \\ &\quad \{\text{VoteA}, \text{VoteB}\} : \&\langle \emptyset : \text{end}; \{\text{VoteA}\} : \text{end}; \{\text{VoteB}\} : \text{end} \rangle \end{aligned}$$

where we use sets of actions as labels. This type manifests a blow up of the input contract that is needed for compiling the complex synchronization patterns of contracts. Indeed, in the theory of contracts, $\{\text{VoteA}, \text{VoteB}\}[\text{VoteA} \oplus \text{VoteB}]$ is equivalent to $\{\text{VoteA}, \text{VoteB}\}[\text{VoteA} \oplus \text{VoteB} \oplus (\text{VoteA} + \text{VoteB})]$. That is, an internal choice between two alternatives means that one *or possibly both* are available. The encoding of contracts has to model explicitly which alternative is taken by sending a notification to the partner.

Related work. The research on contracts was inspired by “CCS without τ ’s” [7] and by Hennessy’s model of acceptance trees [11, 12]. Contracts are an alternative representation of acceptance trees. The relation \preceq was first introduced in [3], albeit it suffered from the lack of a clean semantic characterization and from the fact that it was not transitive. The version of \preceq used in this paper is the same as the one introduced in [14]. In fact, \preceq resembles the must preorder (and it reduces to the must preorder when the interfaces are large enough), but it arises from a notion of compliance that significantly differs from the notion of “passing a test” in the testing framework [6] and that more realistically describes well-behaved clients of Web services. The version of \preceq we work with is actually a stricter version of a more powerful subcontract relation that has been investigated in [4].

Session types have been originally proposed in [13] and subsequently extended for dealing with functional languages [15], asynchrony, object-orientation [9]. In this paper we take [10] as the main reference for session types because it focuses on the subtyping relation. It is worth to notice that we restrict our analysis on the control aspects of session types, whilst other features such as first-class sessions and name passing, which are described in [10], have not been investigated in the framework of contracts yet.

Structure of the paper. We present the formal syntax and semantics of contracts in Section 2 and of session types in Section 3. Sections 4 and 5 present the encoding from session types to contracts and from contracts to session types, respectively. Section 6 concludes by summarizing the main similarities and differences between contracts and session types.

2 Contracts

The syntax of contracts uses an infinite set of *names* \mathcal{N} ranged over by a, b, c, \dots , and a disjoint set of *co-names* $\overline{\mathcal{N}}$ ranged over by $\bar{a}, \bar{b}, \bar{c}, \dots$. Names and co-names are generically called *actions*. We let $\bar{\bar{a}} = a$ and use α, β, \dots to range over actions; we let I, J, K, \dots and R, S, \dots to range over (finite) sets of actions and we extend the operation $\bar{\cdot}$ to sets of actions so that $\bar{R} = \{\bar{\alpha} \mid \alpha \in R\}$. An infinite set of variables is also used, which is ranged over by x, y, z, \dots .

Contracts are pairs $I[\sigma]$ where I is a finite subset of $\mathcal{N} \cup \overline{\mathcal{N}}$ representing the *static interface* of the contract (all the actions occurring in σ must also occur in I), whereas σ , called *behavior*, is defined by the grammar:

$$\sigma ::= \mathbf{0} \mid \alpha.\sigma \mid \sigma \oplus \sigma \mid \sigma + \sigma \mid x \mid \text{rec } x.\sigma$$

Informally, $\mathbf{0}$ describes the inactive behavior; $\alpha.\sigma$ describes the behavior that performs an action α and then behaves like σ ; $\sigma \oplus \tau$ describes the behavior that autonomously decides whether to behave as σ or as τ ; $\sigma + \tau$ describes the behavior that lets the environment choose whether it should behave as σ or as τ ; finally, $\text{rec } x.\sigma$ describes a recursive behavior that is equivalent to $\sigma\{\text{rec } x.\sigma/x\}$. In the following we write Ω for the behavior $\text{rec } x.x$.

Behaviors retain a *transition relation* that is inductively defined by the rules

$$\begin{array}{c} \alpha.\sigma \xrightarrow{\alpha} \sigma \quad \sigma \oplus \tau \longrightarrow \sigma \quad \frac{\sigma \xrightarrow{\alpha} \sigma'}{\sigma + \tau \xrightarrow{\alpha} \sigma'} \quad \frac{\sigma \longrightarrow \sigma'}{\sigma + \tau \longrightarrow \sigma' + \tau} \\[10pt] \text{rec } x.\sigma \longrightarrow \sigma\{\text{rec } x.\sigma/x\} \end{array}$$

We write \Longrightarrow for the reflexive and transitive closure of \longrightarrow ; $\sigma \xRightarrow{\alpha} \sigma'$ for $\sigma \Longrightarrow \xrightarrow{\alpha} \sigma'$; $\sigma \xRightarrow{\alpha}$ if there exists σ' such that $\sigma \xRightarrow{\alpha} \sigma'$. We write $\sigma \uparrow$ if σ has an infinite internal computation $\sigma = \sigma_0 \longrightarrow \sigma_1 \longrightarrow \sigma_2 \longrightarrow \dots$ and $\sigma \downarrow$ if not $\sigma \uparrow$. We write $\sigma \downarrow \alpha_1 \dots \alpha_n$ if $\sigma \downarrow$ and, if $\sigma \xRightarrow{\alpha_1} \sigma'$ implies $\sigma' \downarrow \alpha_2 \dots \alpha_n$; we write $\sigma \uparrow \varphi$ otherwise. For example $\Omega \uparrow$, $\text{rec } x.a + x \uparrow$, and $\text{rec } x.(a.x + b.x) \downarrow \varphi$ for every $\varphi \in \{a, b\}^*$. We let $\text{init}(\sigma)$ be $\{\alpha \mid \sigma \xRightarrow{\alpha}\}$.

A basic use of contracts is to verify whether a client protocol is *compliant with* a service protocol. This compliance is possible if, independently of the internal choices of both client and service, the client successfully completes every interaction with the service. We now formalize the notions of “interaction” and of “successful completion”:

- The interaction of a client and a service is defined by the relation \longrightarrow over pairs of behaviors as follows:

$$\frac{\rho \longrightarrow \rho'}{\rho \parallel \sigma \longrightarrow \rho' \parallel \sigma} \quad \frac{\sigma \longrightarrow \sigma'}{\rho \parallel \sigma \longrightarrow \rho \parallel \sigma'} \quad \frac{\rho \xrightarrow{\alpha} \rho' \quad \sigma \xrightarrow{\bar{\alpha}} \sigma'}{\rho \parallel \sigma \longrightarrow \rho' \parallel \sigma'}$$

where we assume that ρ is a client contract and σ is a service contract. As usual we write \Longrightarrow for the reflexive and transitive closure of \longrightarrow .

- The successful completion of the client is modeled using a special name e . The client has successfully completed the interaction with the service if no further synchronization with the service is possible and the client can emit an e action. We assume that behaviors never manifest co-names \bar{e} .

Definition 1 (Compliance). Let $e \notin S$. The (client) contract $K[\rho]$ is compliant with the (service) contract $I[\sigma]$, written $K[\rho] \dashv I[\sigma]$, if $K \setminus \{e\} \subseteq I$ and $\rho \parallel \sigma \Longrightarrow \rho' \parallel \sigma'$ implies

1. if $\rho' \parallel \sigma' \dashv$, then $\{e\} \subseteq \text{init}(\rho')$;
2. if $\sigma' \uparrow$, then $\{e\} = \text{init}(\rho')$.

According to the notion of behavioral compliance, if a client $K[\rho]$ is compliant with a service $I[\sigma]$ then it should never attempt to perform actions that are not allowed by the interface of the service it is interacting with. If the client-service conversation terminates, then the client is in a successful state (it will emit e). For example, $a.e + b.e \dashv \bar{a} \oplus \bar{b}$ and $a.e \oplus b.e \dashv \bar{a} + \bar{b}$ but $a.e \oplus b.e \not\dashv \bar{a} \oplus \bar{b}$ because of the computation $a.e \oplus b.e \parallel \bar{a} \oplus \bar{b} \Longrightarrow a.e \parallel \bar{b} \dashv$ where the client waits for an interaction on a in vain. Similarly, the client must reach a successful state if the conversation does not terminate but the divergence is due to the service. In this case, however, the client cannot rely on any signal from the service, not even an end-of-connection one, so it is required to do nothing but terminate.

Following De Nicola and Hennessy’s approach to process semantics [6], this test induces a preorder on services on the basis of the set of clients that comply with a given service.

Definition 2 (Subcontract). A contract $I[\sigma]$ is a subcontract of $J[\tau]$, written $I[\sigma] \preceq J[\tau]$, if and only if, for every $K[\rho]$, we have $K[\rho] \dashv I[\sigma]$ implies $K[\rho] \dashv J[\tau]$. We let $I[\sigma] \simeq J[\tau]$ if both $I[\sigma] \preceq J[\tau]$ and $J[\tau] \preceq I[\sigma]$.

That is, if a client is compliant with a service $I[\sigma]$ and $I[\sigma] \preceq J[\tau]$, then the same client is also compliant with $J[\tau]$. Hence, the service $J[\tau]$ can be safely used where $I[\sigma]$ is expected. As usual it is easier to figure out inequalities: $\{a, b\}[a] \not\preceq \{a, b\}[a.b]$ because $\{\bar{a}, \bar{b}, e\}[\bar{a}.(e + \bar{b})] \dashv \{a, b\}[a]$ but $\{\bar{a}, \bar{b}, e\}[\bar{a}.(e + \bar{b})] \not\vdash \{a, b\}[a.b]$; $\{a, b\}[a] \not\preceq \{a, b\}[a + b]$ because $\{\bar{b}, e\}[e + \bar{b}] \dashv \{a, b\}[a]$ but $\{\bar{b}, e\}[e + \bar{b}] \not\vdash \{a, b\}[a + b]$.

Since the set of clients compliant with a given service is usually infinite, Definition 2 gives little insight on the properties of \preceq . This calls for a direct, coinductive characterization of \preceq , which also happens to be easier to work with in the proofs of the results that follow.

Definition 3 (Coinductive subcontract). Let $\sigma \Downarrow R$ if and only if $\sigma \Longrightarrow \sigma'$ and $R = \text{init}(\sigma')$. The relation \mathcal{R} is a coinductive subcontract if $I[\sigma] \mathcal{R} J[\tau]$ implies $I \subseteq J$ and whenever $\sigma \Downarrow$ then

1. $\tau \Downarrow$, and
2. $\tau \Downarrow R$ implies $\sigma \Downarrow R'$ and $R' \subseteq R$, and
3. $\alpha \in I$ and $\tau \xrightarrow{\alpha} \tau'$ implies that there exist $\sigma_1, \dots, \sigma_n$ such that $\sigma \xrightarrow{\alpha} \sigma_i$ for every $1 \leq i \leq n$ and $I[\bigoplus_{1 \leq i \leq n} \sigma_i] \mathcal{R} J[\tau']$.

By this definition, a contract $I[\sigma]$ such that $\sigma \uparrow$ is the smallest one with interface I . When $\sigma \Downarrow$, condition 1 constrains the larger contract $J[\tau]$ to converge as well, since clients might rely on the convergence of σ to complete successfully. Condition 2 states that $J[\tau]$ must exhibit a more deterministic behavior: the smaller the number of ready sets is, the more deterministic the contract is. Furthermore, $J[\tau]$ should expose *at least* the same capabilities as the smaller one ($R' \subseteq R$). Condition 3 is perhaps the most subtle one, as it deals with all the possible derivatives of the smaller contract. The point is that $\{a, b, c\}[a.b + a.c] \simeq \{a, b, c\}[a.(b \oplus c)]$ since, after interacting on a , a client of the service on the left side of \simeq is not aware of which state the service is in (it can be either b or c). Hence, we have to consider all of the possible derivatives after a , thus reducing to verifying $\{a, b, c\}[(b + c) \oplus b \oplus c] \mathcal{R} \{a, b, c\}[b \oplus c]$ which trivially holds.

The set-theoretic and the coinductive versions of the subcontract relation do coincide, as the following proposition states (a proof can be found in the full version of [14]).

Proposition 1. \preceq is the largest coinductive subcontract.

The theory of \preceq has been thoroughly studied in [14]; in particular it has been put in correspondence with a well-known equivalence – the *must testing* [12]. A useful property relating \preceq and \longrightarrow is in order.

Proposition 2. If $\sigma \Longrightarrow \sigma'$, then $I[\sigma] \preceq I[\sigma']$.

For every $I[\sigma]$ there exists σ' in normal form such that $I[\sigma] \simeq I[\sigma']$. This result is slightly more general than those in [5, 6, 12], where normal forms are used to demonstrate the completeness of an axiomatization and are defined for recursion-free terms only. The normal form uses particular families of sets of actions – the *acceptance set*.

Definition 4 (acceptance set [12]). The acceptance set of a behavior σ , denoted by $\mathcal{A}(\sigma)$, is defined as $\mathcal{A}(\sigma) \stackrel{\text{def}}{=} \{R \mid R \subseteq \text{init}(\sigma) \text{ and } \exists S \subseteq R : \sigma \Downarrow S\}$. We let $\mathcal{A}, \mathcal{A}', \dots$ range over acceptance sets.


Definition 5 (normal form). A behavior σ is in normal form if either $\sigma = \Omega$ or $\sigma = x$ or $\sigma = \bigoplus_{R \in \mathcal{A}(\sigma)} \sum_{\alpha \in R} \alpha. \sigma_\alpha$ or $\sigma = \text{rec } x. \bigoplus_{R \in \mathcal{A}(\sigma)} \sum_{\alpha \in R} \alpha. \sigma_\alpha$ and each σ_α is in normal form.

Let $\sigma(\alpha) \stackrel{\text{def}}{=} \bigoplus_{\sigma \xrightarrow{\alpha} \sigma'} \sigma'$ and let $\text{nf}(\sigma)$ be a family of behavior names defined as follows

$$\text{nf}(\sigma) \stackrel{\text{def}}{=} \begin{cases} \Omega & \text{if } \sigma \uparrow \\ \bigoplus_{R \in \mathcal{A}(\sigma)} \sum_{\alpha \in R} \alpha. \text{nf}(\sigma(\alpha)) & \text{otherwise} \end{cases}$$

It is not obvious that $\sigma(\alpha)$ is well defined and that behavior names $\text{nf}(\sigma)$ may be folded into a finite behavior (by using recursion and variables). In order to prove these facts, we represent behaviors as syntax trees where variables are pointers to the corresponding binder [1]. For example, the figure on the right shows the syntax tree of $\sigma = \text{rec } x. a. \text{rec } y. b. (x + y)$.

As usual, every node in syntax trees corresponds to a closed term (i.e. a behavior) that is unique up to the name of bound variables. For example, in the above tree, the node b corresponds to the behavior $b.(\sigma + (\text{rec } y. b. (\sigma + y)))$. The formal definition of behavior associated with a node is omitted because standard.



Lemma 1. Let $\lfloor \sigma \rfloor \stackrel{\text{def}}{=} \{\tau \mid \exists \alpha : \alpha. \tau \text{ occurs in the syntax tree of } \sigma\}$. If $\sigma \xRightarrow{\varphi} \sigma'$, then $\lfloor \sigma' \rfloor \subseteq \lfloor \sigma \rfloor$.

Proof. We prove that the property holds by induction on the derivation of $\sigma \longrightarrow \sigma'$ or $\sigma \xrightarrow{\alpha} \sigma'$ (symmetric cases are omitted). The lemma follows by a straightforward induction on the length of the derivation $\sigma \xRightarrow{\varphi} \sigma'$.

- ($\sigma = \alpha. \sigma' \xrightarrow{\alpha} \sigma'$) We distinguish two subcases: if σ does occur in σ' , then $\lfloor \sigma' \rfloor = \lfloor \sigma \rfloor$; if σ does *not* occur in σ' , then $\lfloor \sigma' \rfloor = \lfloor \sigma \rfloor \setminus \{\sigma'\} \subset \lfloor \sigma \rfloor$.
- ($\sigma = \sigma' \oplus \tau \longrightarrow \sigma'$) We conclude immediately $\lfloor \sigma' \rfloor \subseteq \lfloor \sigma' \oplus \tau \rfloor = \lfloor \sigma \rfloor$.
- ($\sigma = \tau + \tau', \tau \xrightarrow{\alpha} \sigma'$) By induction hypothesis we have $\lfloor \sigma' \rfloor \subseteq \lfloor \tau \rfloor$ hence $\lfloor \sigma' \rfloor \subseteq \lfloor \tau \rfloor \subseteq \lfloor \tau + \tau' \rfloor = \lfloor \sigma \rfloor$.
- ($\sigma = \tau + \tau'', \tau \longrightarrow \tau', \sigma' = \tau' + \tau''$) By induction hypothesis we have $\lfloor \tau' \rfloor \subseteq \lfloor \tau \rfloor$, hence $\lfloor \sigma' \rfloor = \lfloor \tau' + \tau'' \rfloor \subseteq \lfloor \tau + \tau'' \rfloor = \lfloor \sigma \rfloor$.
- ($\sigma = \text{rec } x. \tau, \sigma \longrightarrow \tau\{\sigma/x\} = \sigma'$) We conclude immediately $\lfloor \sigma' \rfloor = \lfloor \sigma \rfloor$. \square

It is worth to notice that Lemma 1 does not hold if behaviors are extended with a parallel operator “|”. For example if $\sigma = \text{rec } x. a \mid (b.x + 0)$, then $\sigma(b) = a \mid \sigma$, which is not a subtree in the syntax tree of σ .

Lemma 2. Let $D(\sigma) \stackrel{\text{def}}{=} \{\sigma(\varphi) \mid \sigma \xRightarrow{\varphi} \cdot\}$. Then $D(\sigma)$ is finite.

Proof. Let D_0, D_1, \dots be the family of sets defined as follows:

$$D_0 \stackrel{\text{def}}{=} \{\sigma\} \quad D_{i+1} \stackrel{\text{def}}{=} D_i \cup \{\sigma' \mid \exists \sigma \in D_i : \exists \alpha : \sigma \xRightarrow{\alpha} \sigma'\}$$

We only need to show that there exists n such that $D_n = D_{n+1}$, because each $\sigma(\varphi)$ is obtained by joining some continuations σ' , for some subtrees $\alpha.\sigma'$ in σ . Let n be the cardinality of $[\sigma]$ (n is the number of *distinct* subtrees in the syntax tree of σ and it exists because σ is finite). By contradiction, assume that there exists $\sigma_{n+1} \in D_{n+1} \setminus D_n$. Then there exist $\sigma_1, \dots, \sigma_n$ and $\alpha_1, \dots, \alpha_n, \alpha_{n+1}$ such that

$$\sigma \xRightarrow{\alpha_1} \sigma_1 \xRightarrow{\alpha_2} \dots \xRightarrow{\alpha_n} \sigma_n \xRightarrow{\alpha_{n+1}} \sigma_{n+1}$$

and $\sigma_i \in D_i \setminus D_{i-1}$ for every $1 \leq i \leq n$. We have $\sigma_{n+1} \in [\sigma_n]$. By Lemma 1 we deduce $\sigma_{n+1} \in [\sigma]$. Since $[\sigma]$ has only n distinct subtrees, σ_{n+1} is reachable from σ with at most n reductions, thus $\sigma_{n+1} \in D_n$, which contradicts $\sigma_{n+1} \in D_{n+1} \setminus D_n$. \square

Theorem 1. *For every $I[\sigma]$ there exists σ' in normal form such that $I[\sigma] \simeq I[\sigma']$. The behavior σ' is the folding of $\text{nf}(\sigma)$.*

Proof. By Lemma 2 there are finitely many $\sigma(\varphi)$. These $\sigma(\varphi)$ are in one-to-one correspondence with the behavior names that may be recursively invoked by $\text{nf}(\sigma)$. Therefore the set of such behavior names is finite and $\text{nf}(\sigma)$ may be folded into a behavior with recursion and variables.

To prove the subcontract equivalence, let \mathcal{R} be the least relation containing \preceq and such that

1. if $I[\sigma] \preceq J[\tau]$ and τ' is the normal form of τ , then $I[\sigma] \mathcal{R} J[\tau']$;
2. if $I[\sigma] \preceq J[\tau]$ and σ' is the normal form of σ , then $I[\sigma'] \mathcal{R} J[\tau]$.

We prove that \mathcal{R} is a subcontract relation; the theorem follows directly.

Let $I[\sigma] \mathcal{R} J[\tau']$, where τ' is the normal form of τ . Then $I[\sigma] \preceq J[\tau]$. As regards condition 1 in the definition of coinductive subcontract, notice that from $I[\sigma] \preceq J[\tau]$ we have $\sigma \downarrow$ implies $\tau \downarrow$, hence $\tau' \downarrow$ from the definition of normal form. As regards condition 2, let $\tau' \Downarrow R$, then $R \in \mathcal{A}(\tau)$, hence $\tau \Downarrow R'$ and $R' \subseteq R$ by definition of $\mathcal{A}(\tau)$. From $I[\sigma] \preceq J[\tau]$ we obtain $\sigma \Downarrow S$ and $S \subseteq R'$, hence we conclude $S \subseteq R$. As regards condition 3, let $\tau \xRightarrow{\alpha} \tau''$. Then $J[\tau'(\alpha)] \preceq J[\tau'']$ and, by definition, $\tau'(\alpha)$ is in normal form. Then $\tau \xRightarrow{\alpha}$ and, by $I[\sigma] \preceq J[\tau]$, $\sigma \xRightarrow{\alpha}$. Therefore $I[\sigma(\alpha)] \preceq J[\tau(\alpha)]$, hence by definition of \mathcal{R} , $I[\sigma(\alpha)] \mathcal{R} J[\tau''']$, where τ''' is the normal form of $\tau(\alpha)$. By definition of normal form $\tau''' = \tau'(\alpha)$ and we conclude.

Let $I[\sigma'] \mathcal{R} J[\tau]$, where σ' is the normal form of σ . Then $I[\sigma] \mathcal{R} J[\tau]$. As regards condition 1 in the definition of coinductive subcontract, $\sigma' \downarrow$ implies $\sigma \downarrow$, hence $\tau \downarrow$. As regards condition 2, let $\tau \Downarrow R$. From $I[\sigma] \preceq J[\tau]$ we have $\sigma \Downarrow S$ and $S \subseteq R$ and we conclude by observing that $\sigma' \Downarrow S$. As regards condition 3, let $\tau \xRightarrow{\alpha} \tau'$. Then there exist $\sigma_1, \dots, \sigma_n$ such that $I[\bigoplus_{1 \leq i \leq n} \sigma_i] \preceq J[\tau']$. Since $I[\sigma(\alpha)] \preceq I[\bigoplus_{1 \leq i \leq n} \sigma_i]$ and the normal form of $\sigma(\alpha)$ is equal to $\sigma'(\alpha)$ we conclude $I[\sigma'(\alpha)] \mathcal{R} J[\tau']$ by definition of \mathcal{R} . \square

By Theorem 1, from now on we let $\text{nf}(\sigma)$ be the normal form of σ .

3 Session types

Similarly to a contract, a session type is meant to describe the type of a communication as a sequence of actions and branching points. The syntax of session types uses an infinite set of labels \mathcal{L} ranged over by ℓ, ℓ', \dots , and an infinite set of variables ranged over by x, y, z, \dots . Session types are defined by the grammar:

$$S ::= \text{end} \quad | \quad \oplus \langle \ell_i : S_i^{i \in I} \rangle \quad | \quad \& \langle \ell_i : S_i^{i \in I} \rangle \quad | \quad x \quad | \quad \mu x. S$$

The session type end describes a completed communication; the session type $\oplus \langle \ell_i : S_i^{i \in I} \rangle$ describes a *choice* where a process autonomously decides to proceed according to one of the continuations S_i 's. Before doing so, the process notifies the partner of the communication by sending a label ℓ_i ; the session type $\& \langle \ell_i : S_i^{i \in I} \rangle$ describes a *branch* where a process is ready to proceed according to any of the continuations S_i 's. Before doing so, the process waits for a label ℓ_i from the process it is interacting with. The label uniquely identifies the continuation. The session type $\mu x. S$ describes a recursive type, much like a recursive contract.

Following [10], session types are taken *contractive*, namely every occurrence of a variable is guarded by at least a choice or a branch. Unlike [10], however, we omit session types describing the communications of actual data (not merely labels) during a session. Typically, such communications are represented by session types of the form $!t.S$ (for “send a message of type t on the channel and then continue as S ”) or $?t.S$ (for “receive a message of type t from the channel and then continue as S ”). However, as far as the comparison of session types and contracts is concerned, the presence of such actions is not particularly relevant, since we are going to focus on the control part of an interaction. The addition of $!t$ and $?t$ actions to the above grammar does not pose any particular problem for all of the results that follow, but at the same time it does not contribute any significant insight in the comparison.

The semantics of session types is defined in terms of a subtyping relation \leq , which is presented by means of a deductive system or as a coinductive definition [10]. In this contribution we stick to this latter presentation that turns out to be easier to compare with the subcontract relation defined in the previous section. More precisely, $S \leq T$, whenever every process with type S can be used where a process with type T is expected (in this respect, \leq differs from \preceq as it is oriented as a real subtyping relation). Because of the structure of session types, \leq is quite straightforward to define: basically every choice in S must be defined on a subset of the labels occurring in the corresponding choice in T , whereas any branch in T must be defined on a superset of the labels occurring in the corresponding branch in S . Let

$$\text{unfold}(S) \stackrel{\text{def}}{=} \begin{cases} \text{unfold}(T\{S/x\}) & \text{if } S = \mu x. T \\ S & \text{otherwise} \end{cases}$$

That is, $\text{unfold}(\cdot)$ removes topmost recursions until it reveals a branch or a choice. It is worth to notice that $\text{unfold}(\cdot)$ is always defined because session types are contractive.

Definition 6 (Subtyping). *The relation \mathcal{R} is a coinductive subtyping if $S \mathcal{R} T$ implies*

1. *if $\text{unfold}(S) = \text{end}$, then $\text{unfold}(T) = \text{end}$,*

2. if $\text{unfold}(S) = \oplus \langle \ell_i : S_i^{i \in I} \rangle$, then $\text{unfold}(T) = \oplus \langle \ell_j : T_j^{j \in J} \rangle$ and $I \subseteq J$ and for every $i \in I$ we have $S_i \mathcal{R} T_i$,
3. if $\text{unfold}(S) = \& \langle \ell_i : S_i^{i \in I} \rangle$, then $\text{unfold}(T) = \& \langle \ell_j : T_j^{j \in J} \rangle$ and $J \subseteq I$ and for every $j \in J$ we have $S_j \mathcal{R} T_j$.

Let \leq be the largest coinductive subtyping.

For example, we have $\oplus \langle \ell : S \rangle \leq \oplus \langle \ell : S; \ell' : T \rangle$ since the smaller session type represents a process that behaves more deterministically than the one represented by the larger session type. Dually we have $\& \langle \ell : S; \ell' : T \rangle \leq \& \langle \ell : S \rangle$ since the process represented by the smaller session type offers a wider range of continuations (it is ready to accept a superset of the labels accepted by the process represented by the larger session type). Unlike contracts, however, we have $\& \langle \ell : S \rangle \not\leq \text{end}$ and $\oplus \langle \ell : S \rangle \not\leq \text{end}$. In the first case the process *waits for* a label that never arrives, in the second case it *sends* a label that the matching party is not ready to receive.

Session types are naturally equipped with a $\text{dual}(\cdot)$ function computing the dual type.

$$\begin{aligned}
\text{dual}(\text{end}) &= \text{end} \\
\text{dual}(\oplus \langle \ell_i : S_i^{i \in I} \rangle) &= \& \langle \ell_i : \text{dual}(S_i)^{i \in I} \rangle \\
\text{dual}(\& \langle \ell_i : S_i^{i \in I} \rangle) &= \oplus \langle \ell_i : \text{dual}(S_i)^{i \in I} \rangle \\
\text{dual}(x) &= x \\
\text{dual}(\mu x. S) &= \mu x. \text{dual}(S)
\end{aligned}$$

As for $\text{unfold}(\cdot)$, $\text{dual}(\cdot)$ is always defined because session types are contractive. If S is the session type representing the behavior of a process, then $\text{dual}(S)$ represents the behavior of a “canonical” process that interacts successfully with the first one. The theory of session session types does not formalize the notion of successful interaction, but we will be able to reason about it by means of one of our encodings in Section 5.

The duality between choices and branches in session types is formalized by the following proposition. Let $\text{in}(S)$ (respectively, $\text{out}(S)$) be the set of labels occurring in branches (respectively, choices) of S . Then subtyping induces a relation on labels $\text{in}(\cdot)$ and $\text{out}(\cdot)$.

Proposition 3. *If $S \leq T$ then (1) $\text{out}(S) \subseteq \text{out}(T)$, (2) $\text{in}(T) \subseteq \text{in}(S)$, and (3) $\text{dual}(T) \leq \text{dual}(S)$.*

Notice that from the previous proposition and from the fact that $\text{dual}(\cdot)$ is the inverse of itself ($\text{dual}(\text{dual}(S)) = S$), we have $S \leq T$ if and only if $\text{dual}(T) \leq \text{dual}(S)$.

4 Encoding session types into contracts

In order to encode session types into contracts we preliminarily need to set a correspondence between labels and names. For simplicity we let $\mathcal{L} = \mathcal{N}$. Therefore, in the following, we will not distinguish between labels and names and we will address them with ℓ, a, \dots .

$ \begin{aligned} \llbracket \text{end} \rrbracket_I &= \sigma_e \\ \llbracket \oplus \langle \ell_i : S_i^{i \in 1..n} \rangle \rrbracket_I &= \bigoplus_{i \in 1..n} \bar{\ell}_i \cdot \llbracket S_i \rrbracket_I \\ \llbracket \& \langle \ell_i : S_i^{i \in 1..n} \rangle \rrbracket_I &= \sum_{i \in 1..n} \ell_i \cdot \llbracket S_i \rrbracket_I + \sum_{\alpha \in I \setminus \{\ell_1, \dots, \ell_n\}} \alpha \cdot \Omega \\ \llbracket x \rrbracket_I &= x \\ \llbracket \mu x. S \rrbracket_I &= \text{rec } x. \llbracket S \rrbracket_I \end{aligned} $

Table 1. Encoding of session types into contracts.

The encoding $\llbracket S \rrbracket_I$ of a session type S is defined with respect to an interface I (Table 1). The terminal behavior σ_e is \mathbf{e} if the session type being encoded regards a client, or $\mathbf{0}$ if it regards a service. The only case that deserves discussion is branch. The contract $\llbracket \& \langle \ell_i : S_i^{i \in 1..n} \rangle \rrbracket_I$ has as many initial actions as are allowed by the interface, irrespective of the labels of the branches. This enforces width subtyping in the resulting contracts, which is otherwise false, in general. In facts, $\{a, b\}[a] \not\preceq \{a, b\}[a+b]$ because the client $\{\bar{a}, \bar{b}\}[\bar{a} \cdot \mathbf{e} + \bar{b}]$ is compliant with the first service, but not with the second one. By translating $S = \& \langle a : S' \rangle$ as simply $\{a, b\}[a. \llbracket S' \rrbracket]$ (the action b is in the interface because it occurs in S'), then a client such as $\{\bar{a}, \bar{b}\}[\bar{a} \cdot \mathbf{e} + \bar{b} \cdot \rho]$ could fail when one replaces $\llbracket S \rrbracket_{\{a, b\}}$ with $\llbracket T \rrbracket_{\{a, b\}}$ and $T \leq S$, for instance when $T = \& \langle a : S' ; b : T' \rangle$. To avoid these cases, one has to exclude $\{\bar{a}, \bar{b}\}[\bar{a} \cdot \mathbf{e} + \bar{b} \cdot \rho]$ from those clients that are compliant with $\llbracket S \rrbracket$ because it attempts to do an action (\bar{b}) that is not explicitly allowed by S . This exclusion follows by translating a branch into an external choice “+” with a summand for every action in the interface: the actions that do not appear as labels of the branch are given a continuation Ω . Therefore, the contract $\llbracket S \rrbracket_{\{a, b\}}$ does not admit, in general, a client $\{\bar{a}, \bar{b}\}[\bar{a} \cdot \mathbf{e} + \bar{b} \cdot \rho]$ but only a client $\{\bar{a}, \bar{b}\}[\bar{a} \cdot \mathbf{e} + \bar{b} \cdot \mathbf{e}]$ (this client has to terminate after an interaction on b , without requiring any further capability to the service). The contract Ω is the smallest one with a given interface: a client compliant with Ω will also comply with any other service, in particular with those resulting from the encoding of smaller session types.

Example 1. According to the encoding, we have

- $\llbracket \& \langle a : \text{end} \rangle \rrbracket_{\{a\}} = a;$
- $\llbracket \& \langle a : \text{end} \rangle \rrbracket_{\{a, b\}} = a + b \cdot \Omega;$
- $\llbracket \& \langle a : \text{end}; b : \text{end} \rangle \rrbracket_{\{a, b\}} = a + b;$
- $\llbracket \oplus \langle a : \text{end} \rangle \rrbracket_{\{a\}} = \llbracket \oplus \langle a : \text{end} \rangle \rrbracket_{\{a, b\}} = \bar{a};$
- $\llbracket \oplus \langle a : \text{end}; b : \text{end} \rangle \rrbracket_{\{a, b\}} = \bar{a} \oplus \bar{b}.$

We notice that $\& \langle a : \text{end}; b : \text{end} \rangle \leq \& \langle a : \text{end} \rangle$ and $\{a\}[a] \preceq \{a, b\}[a + b \cdot \Omega] \preceq \{a, b\}[a+b]$. Similarly, $\oplus \langle a : \text{end} \rangle \leq \oplus \langle a : \text{end}; b : \text{end} \rangle$ and $\{a, b\}[\bar{a} \oplus \bar{b}] \preceq \{a, b\}[\bar{a}]$. \blacksquare

More generally, the encoding of Table 1 makes the subtyping relation and the sub-contract relation agree.

Theorem 2. *Let $J = \text{in}(T) \cup \overline{\text{out}(T)}$ and $I = \text{in}(S) \cup \overline{\text{out}(T)}$. Then $S \leq T$ if and only if $J[\llbracket T \rrbracket_J] \preceq I[\llbracket S \rrbracket_I]$.*

Proof. (“only if” part) Let \mathcal{R} be the least relation such that

1. if $J \subseteq I$ then $J[\Omega] \mathcal{R} I[\sigma]$;
2. if $S \leq T$ and $J = \text{in}(T) \cup \overline{\text{out}(T)}$ and $I = \text{in}(S) \cup \overline{\text{out}(T)}$, then $J[\llbracket T \rrbracket_J] \mathcal{R} I[\llbracket S \rrbracket_I]$.

We prove that \mathcal{R} is a coinductive subcontract. Let $J[\llbracket T \rrbracket_J] \mathcal{R} I[\llbracket S \rrbracket_I]$. By Proposition 3, $J \subseteq I$. By definition, $\llbracket \cdot \rrbracket$ always yields behaviors σ such that $\sigma \downarrow$. Therefore condition 1 in Definition 3 is trivially satisfied. As regards conditions 2 and 3, we reason by cases on S . By definition of \leq , $S \leq T$ if and only if $\text{unfold}(S) \leq \text{unfold}(T)$. Therefore we may restrict to cases where neither S nor T are recursive types.

- ($S = \text{end}$) Then $T = \text{end}$. Condition 2 follows immediately; condition 3 is vacuous.
- ($S = \oplus \langle \ell_h : S_h \mid h \in H \rangle$) Then $T = \oplus \langle \ell_k : T_k \mid k \in K \rangle$ and $H \subseteq K$ and $S_h \leq T_h$, for every $h \in H$. As regards condition 2, if $\llbracket S \rrbracket_I \Downarrow R$ then $\overline{\ell_h} \in R$ for some $h \in H$. Since $H \subseteq K$, $\llbracket T \rrbracket_J \Downarrow \{\ell_h\}$ and $\{\overline{\ell_h}\} \subseteq R$. As regards conditions 3, by Proposition 2, it suffices to consider the transitions $\llbracket S \rrbracket_I \xrightarrow{\overline{\ell_h}} \llbracket S_h \rrbracket_I$ for every $h \in H$. Since $H \subseteq K$, $\llbracket T \rrbracket_J \xrightarrow{\overline{\ell_h}} \llbracket T_h \rrbracket_J$ and $\llbracket T_h \rrbracket_J \mathcal{R} \llbracket S_h \rrbracket_I$ follows by $S_h \leq T_h$;
- ($S = \& \langle \ell_k : S_k \mid k \in K \rangle$) Then $T = \& \langle \ell_h : T_h \mid h \in H \rangle$ with $H \subseteq K$. Condition 2 follows because $\llbracket T \rrbracket_J \Downarrow \{\ell_h \mid h \in H\}$ and $\llbracket S \rrbracket_I \Downarrow \{\ell_k \mid k \in K\}$ and $\{\ell_h \mid h \in H\} \subseteq \{\ell_k \mid k \in K\}$. As regards condition 3, by Proposition 2, it suffices to consider $\llbracket S \rrbracket_I \xrightarrow{\ell_k} \llbracket S_k \rrbracket_I$ with $k \in J$. There are two subcases: either (i) $\ell_k \notin \{\ell_h \mid h \in H\}$ or (ii) $\ell_k \in \{\ell_h \mid h \in H\}$. In subcase (i), $\llbracket T \rrbracket_J \xrightarrow{\ell_k} \Omega$ and $J[\Omega] \mathcal{R} I[\llbracket S_k \rrbracket_I]$ by definition of \mathcal{R} . In subcase (ii), $\llbracket T \rrbracket_J \xrightarrow{\ell_k} \llbracket T_k \rrbracket_J$. From $S_k \leq T_k$ we derive $J[\llbracket T_k \rrbracket_J] \mathcal{R} I[\llbracket S_k \rrbracket_I]$ by definition of \mathcal{R} .

(“if” part) We prove that if $S \not\leq T$, then $J[\llbracket T \rrbracket_J] \not\mathcal{R} I[\llbracket S \rrbracket_I]$. We reason by cases on the shape of S and T . It is sufficient to consider those cases in which none of the types begins with a recursion and $S \leq T$ holds directly, without looking at the session types in the choices and branches possibly found in S and T .

- ($S = \text{end}$) It must be $T \neq \text{end}$. Then $\llbracket S \rrbracket_I \Downarrow \{e\}$ whereas $\llbracket T \rrbracket_J \Downarrow R$ implies $R \neq \emptyset$ and $e \notin R$;
- ($S = \oplus \langle \ell_i : S_i \mid i \in I \rangle$) If $T = \text{end}$, then we can reason as for the case $S = \text{end}$ and conclude that $J[\llbracket T \rrbracket_J] \not\mathcal{R} I[\llbracket S \rrbracket_I]$. If $T = \& \langle \ell_j : T_j \mid j \in J \rangle$, then $\llbracket T \rrbracket_J$ has only one ready set $J = \text{in}(T) \cup \overline{\text{out}(T)}$, which contains at least two actions and hence cannot be smaller than $\{\ell_i\}$, for every $i \in I$. If $T = \oplus \langle \ell_j : T_j \mid j \in J \rangle$ and $I \not\subseteq J$, then there exists $i \in I$ such that $i \notin J$. So $\llbracket S \rrbracket_I \Downarrow \{\ell_i\}$ whereas $\llbracket T \rrbracket_J \Downarrow R$ implies $R \neq \emptyset$ and $\overline{\ell_i} \notin R$;
- ($S = \& \langle \ell_i : S_i \mid i \in I \rangle$) If $T = \text{end}$ we conclude immediately since no ℓ_i is equal to e . If $T = \oplus \langle \ell_j : T_j \mid j \in J \rangle$, then from $I = \text{in}(S) \cup \overline{\text{out}(T)}$ we have that $j \in J$ implies $\overline{\ell_j} \in I$. Now we have $\llbracket S \rrbracket_I \xrightarrow{\overline{\ell_j}} \Omega$, whereas $\llbracket T \rrbracket_J \xrightarrow{\overline{\ell_j}} \sigma$ implies $\sigma \downarrow$. If $T = \& \langle \ell_j : T_j \mid j \in J \rangle$ and $J \not\subseteq I$, then there exists $j \in J$ such that $j \notin I$. If $J \not\subseteq I$ there is nothing to prove. If $J \subseteq I$, then $\llbracket S \rrbracket_I \xrightarrow{\overline{\ell_j}} \Omega$, whereas $\llbracket T \rrbracket_J \xrightarrow{\overline{\ell_j}} \sigma$ implies $\sigma \downarrow$. \square

The following proposition shows that internal moves in the encoding of a session type S correspond to reducing its choices and determining a session type $S' \leq S$.

Proposition 4. *If $\llbracket S \rrbracket_1 \Longrightarrow \sigma$, then $\sigma = \llbracket S' \rrbracket_1$ and $S' \leq S$.*

Proof. It is sufficient to consider the case $\llbracket S \rrbracket_1 \longrightarrow \sigma$, then the lemma follows directly. We reason by cases on the structure of S , there are only two possibilities:

($S = \oplus \langle \ell_i : S_i \mid i \in K \rangle$) Then $\llbracket S \rrbracket_1 = \bigoplus_{i \in I} \bar{\ell}_i. \llbracket S_i \rrbracket_1 \longrightarrow \bigoplus_{i \in J} \bar{\ell}_j. \llbracket S_j \rrbracket_1 = \llbracket \oplus \langle \ell_j : S_j \mid j \in J \rangle \rrbracket_1$ with $J \subseteq I$. We conclude by observing that $\oplus \langle \ell_j : S_j \mid j \in J \rangle \leq S$.
 ($S = \mu x. S'$) Then $\llbracket S \rrbracket_1 = \text{rec } x. \llbracket S' \rrbracket_1 \longrightarrow \llbracket S' \rrbracket_1 \{ \text{rec } x. \llbracket S' \rrbracket_1 / x \} = \llbracket S' \{ S/x \} \rrbracket_1$ and this is the only possible reduction. We conclude because $S' \{ S/x \} \leq S$. \square

The encoding in Table 1 allows us to relate session types and their duals. The relation is exactly the compliance of Definition 1.

Theorem 3. *Let $I = \text{in}(S) \cup \overline{\text{out}(S)}$ and $J = \text{in}(T) \cup \overline{\text{out}(T)}$ and $S \leq \text{dual}(T)$. Then $\llbracket S \rrbracket_1 \dashv \llbracket T \rrbracket_J$.*

Proof. Let \mathcal{R} be the least relation such that if $S \leq \text{dual}(T)$ and $\text{in}(S) \cup \overline{\text{out}(S)} \subseteq I$ and $\text{in}(T) \cup \overline{\text{out}(T)} \subseteq J$, then $\llbracket S \rrbracket_1 \mathcal{R} \llbracket T \rrbracket_J$. We prove that \mathcal{R} is a compliance relation. Let $\llbracket S \rrbracket_1 \mathcal{R} \llbracket T \rrbracket_J$. By Proposition 4, if $\llbracket S \rrbracket_1 \longrightarrow \rho$, then $\rho = \llbracket S' \rrbracket_1$ and $S' \leq S \leq \text{dual}(T)$. Symmetrically, if $\llbracket T \rrbracket_J \longrightarrow \sigma$, then $\sigma = \llbracket T' \rrbracket_J$ and $T' \leq T$, hence, by Proposition 3(3), $S \leq \text{dual}(T) \leq \text{dual}(T')$. It follows that we may restrict our analysis to cases where neither the encoding of S nor the encoding of T can perform internal moves:

($S = \text{end}$) Then $T = \llbracket \text{end} \rrbracket_J$. Therefore $\llbracket S \rrbracket_1 \parallel \llbracket T \rrbracket_J \dashv \dashv$ and $\llbracket S \rrbracket_1 \xrightarrow{e}$;
 ($S = \oplus \langle \ell : S' \rangle$) Then $T = \& \langle \ell_i : T_i \mid i \in K \rangle$, and $\ell = \ell_i$ for some $i \in K$ and $S' \leq \text{dual}(T_i)$. Now $\llbracket S \rrbracket_1 \xrightarrow{\bar{\ell}}$ and $\llbracket T \rrbracket_J \xrightarrow{\ell}$, hence $\llbracket S \rrbracket_1 \parallel \llbracket T \rrbracket_J \longrightarrow \llbracket S' \rrbracket_1 \parallel \llbracket T_i \rrbracket_J$ and we conclude $\llbracket S' \rrbracket_1 \mathcal{R} \llbracket T_i \rrbracket_J$ by definition of \mathcal{R} .
 ($S = \& \langle \ell_i : S_i \mid i \in I \rangle$) Dual of the previous case. \square

5 Encoding contracts into session types

The encoding of contracts to session types is partially defined. In particular, we will restrict the encoding to behaviors that are convergent. The reason for this restriction derives from the fact that session types describe communications whose end is agreed by both parties having type end. The behavior Ω , on the other hand, represents a service that may not pay any attention to the communication it is involved in, and it has no corresponding session type.

A behavior σ is *strongly convergent* if, for every sequence φ , $\sigma \downarrow \varphi$. The following encodings are defined on contracts with *strongly convergent* behaviors. It is also convenient to restrict the encoding of client contracts to those whose behavior never leads to $\mathbf{0}$ without emitting e . For example, the behavior $a.e + b.\mathbf{0}$ describes a client that succeeds if the service proposes \bar{a} , but that fails if the service proposes \bar{b} . Such a behavior has no counterpart in session types, where failures are not explicitly represented.

In general, if a client is unable to handle a particular action, like b in the example, it should simply omit that action from its behavior. We say that a (client) contract $I[\rho]$ is *canonical* if, whenever $\rho \xrightarrow{\varphi} \rho'$ is maximal, then $\varphi = \varphi' \mathbf{e}$ and $\mathbf{e} \notin \text{names}(\varphi')$. For example $\{a\}[a.\mathbf{e}]$, $\{a\}[\text{rec } x.a.x]$, and $\emptyset[\Omega]$ are canonical; $\{a, b\}[a.\mathbf{e} + b.\mathbf{0}]$ and $\{a\}[\text{rec } x.a + x]$ are not canonical.

The encoding of contracts into session types is reported in Table 2. We write $S \bowtie \mathcal{A}$ (\mathcal{A} is an acceptance set) if for every $R \in \mathcal{A}$, either $\mathbf{e} \in R$ or $S \cap R \neq \emptyset$. We use an injective map $\widehat{\cdot}$ from sets of actions to labels that we leave unspecified.

<p><i>Encoding of client contracts/behaviors:</i></p> $\mathcal{C}[\llbracket K[\rho] \rrbracket] = \oplus \langle \widehat{K \setminus \{\mathbf{e}\}} : \mathcal{C}[\llbracket \text{nf}(\rho) \rrbracket]_{K \setminus \{\mathbf{e}\}} \rangle$ $\mathcal{C}[\llbracket \bigoplus_{R \in \mathcal{A}} \sum_{\alpha \in R} \alpha.\rho_\alpha \rrbracket]_K = \& \langle \widehat{S} : \oplus \langle \widehat{\emptyset} : \text{end}; \widehat{\alpha} \rangle : \mathcal{C}[\llbracket \rho_\alpha \rrbracket]_K \rangle_{R \in \mathcal{A}, \alpha \in S \cap R, S \subseteq K, S \bowtie \mathcal{A}}$ <p style="text-align: center;">if $\mathbf{e} \in R$</p> $\mathcal{C}[\llbracket \text{rec } x.\rho \rrbracket]_K = \mu x. \mathcal{C}[\llbracket \rho \rrbracket]_K$ $\mathcal{C}[\llbracket x \rrbracket]_K = x$	
<p><i>Encoding of service contracts/behaviors:</i></p> $\mathcal{S}[\llbracket I[\sigma] \rrbracket] = \& \langle \widehat{K} : \mathcal{S}[\llbracket \text{nf}(\sigma) \rrbracket]_{K^{\subseteq I}} \rangle$ $\mathcal{S}[\llbracket \bigoplus_{R \in \mathcal{A}} \sum_{\alpha \in R} \alpha.\sigma_\alpha \rrbracket]_K = \oplus \langle \widehat{R \cap K} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\alpha} \rangle : \mathcal{S}[\llbracket \sigma_\alpha \rrbracket]_K \rangle_{\alpha \in R \cap K, R \in \mathcal{A}}$ $\mathcal{S}[\llbracket \text{rec } x.\rho \rrbracket]_K = \mu x. \mathcal{S}[\llbracket \rho \rrbracket]_K$ $\mathcal{S}[\llbracket x \rrbracket]_K = x$	

Table 2. Encoding of contracts into session types.

The encoding distinguishes between client and service contracts. One reason is that, unlike sessions, which are supposed to be completed *symmetrically* by both parties, the theory of contracts is biased towards clients, which are free to interrupt the interaction any time they please. The other reason is that contracts describe a more abstract synchronization pattern than session types do, and the encoding of contracts into session types has to render this synchronization pattern, which amounts to a little handshaking protocol.

The session types corresponding to the client contract $K[\rho]$ and the service contract $I[\sigma]$ are denoted by $\mathcal{C}[\llbracket K[\rho] \rrbracket]$ and $\mathcal{S}[\llbracket I[\sigma] \rrbracket]$, respectively. The labels in the types represent finite sets of actions in the source contracts. Let us discuss the encoding of a service contract. The first step in the generation of the session type is the offering of an interface to the client. Any client that asks no more capabilities than those offered by the service can connect. Hence, the service offers as many interfaces as the number of the subsets of its own interface. For every offered interface, the continuation session type is a specialized encoding of the service's contract, restricted to the offered interface: $\mathcal{S}[\llbracket \sigma \rrbracket]_K$. For the proper encoding of the behavior, we resort to the normal form (Definition 5). A behavior of the form $\bigoplus_{R \in \mathcal{A}} \sum_{\alpha \in R} \alpha.\sigma_\alpha$ is one in which the service can be

in as many states as the cardinality of \mathcal{A} . In each state $R \in \mathcal{A}$, the service is ready to perform any of the actions in R . So, the service begins by communicating to the client the state it is in (restricted to the interface of the connected client). Then, the service accepts a singleton action, among those that are available, indicating the choice of the client, or the special action $\widehat{\emptyset}$ denoting the fact that the client has decided to terminate at this stage. We notice that the term

$$\oplus \langle \widehat{R \cap K} : \& \langle \widehat{\emptyset} : \text{end}; \{\widehat{\alpha}\} : \mathcal{S}[\sigma_\alpha]_K \rangle_{\alpha \in R \cap K} \rangle_{R \in \mathcal{A}}$$

is well formed: every label of a branch or a choice has exactly one continuation because the behavior $\bigoplus_{R \in \mathcal{A}} \sum_{\alpha \in R} \alpha \cdot \sigma_\alpha$ is in normal form.

In some sense the encoding of a service contract is “kind” as it tries to accommodate the largest number of clients (not only those connecting with exactly the same interface as the service, but also with smaller ones). Conversely, the encoding of a client contract is selfish in that it only encodes the client behavior, whose only purpose is to successfully achieve its task. The first action of the client is the selection of the service interface that matches with its own (containing the co-actions of the client’s interface). Then, at each interaction, the client must be ready to accept a set \bar{S} of actions from the service, representing the state the service is in. However, not all such states are suitable for the client. In particular, let \mathcal{A} be the acceptance set of the client; the client will only accept those states \bar{S} such that, for every $R \in \mathcal{A}$, either $e \in R$ (the client is ready to terminate) or $S \cap R \neq \emptyset$ (the client and the service can synchronize). Said otherwise, the fewer sets \bar{S} the client can accept from the service, the more demanding it is. Once a set has been accepted, the client may choose the special label $\widehat{\emptyset}$, signaling its intention to terminate the interaction, or it may choose a label $\widehat{\alpha}$, signaling the intention of synchronizing on α . Because of the way the sets \bar{S} are accepted by the client, it is always possible to pursue at least one of such possibilities.

There are discretions in the two encodings $\mathcal{C}[\cdot]$ and $\mathcal{S}[\cdot]$: the alternation between \oplus and $\&$ may be reversed without changing the following correctness results.

Example 2. Consider the service contracts

$$I[\sigma] \stackrel{\text{def}}{=} \{a, b\}[a \oplus b] \quad \text{and} \quad I[\tau] \stackrel{\text{def}}{=} \{a, b\}[a + b]$$

and notice that $I[\sigma] \preceq I[\tau]$. Their respective encodings are

$$\begin{aligned} \mathcal{S}[I[\sigma]] &= \& \langle \widehat{\emptyset} : \oplus \langle \widehat{\emptyset} : \& \langle \widehat{\emptyset} : \text{end} \rangle \rangle; \\ &\quad \widehat{\{a\}} : \oplus \langle \widehat{\emptyset} : \& \langle \widehat{\emptyset} : \text{end} \rangle; \widehat{\{a\}} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\{a\}} : \text{end} \rangle \rangle; \\ &\quad \widehat{\{b\}} : \oplus \langle \widehat{\emptyset} : \& \langle \widehat{\emptyset} : \text{end} \rangle; \widehat{\{b\}} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\{b\}} : \text{end} \rangle \rangle; \\ &\quad \widehat{\{a, b\}} : \oplus \langle \widehat{\{a\}} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\{a\}} : \text{end} \rangle; \\ &\quad \quad \widehat{\{b\}} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\{b\}} : \text{end} \rangle; \\ &\quad \quad \widehat{\{a, b\}} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\{a\}} : \text{end}; \widehat{\{b\}} : \text{end} \rangle \rangle \rangle \\ \mathcal{S}[I[\tau]] &= \& \langle \widehat{\emptyset} : \oplus \langle \widehat{\emptyset} : \& \langle \widehat{\emptyset} : \text{end} \rangle \rangle; \\ &\quad \widehat{\{a\}} : \oplus \langle \widehat{\{a\}} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\{a\}} : \text{end} \rangle \rangle; \\ &\quad \widehat{\{b\}} : \oplus \langle \widehat{\{b\}} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\{b\}} : \text{end} \rangle \rangle; \\ &\quad \widehat{\{a, b\}} : \oplus \langle \widehat{\{a, b\}} : \& \langle \widehat{\emptyset} : \text{end}; \widehat{\{a\}} : \text{end}; \widehat{\{b\}} : \text{end} \rangle \rangle \rangle \end{aligned}$$

and now we have $\mathcal{S}[\llbracket \mathbf{I}[\tau] \rrbracket] \leq \mathcal{S}[\llbracket \mathbf{I}[\sigma] \rrbracket]$. Notice that, in the encoding of $\mathbf{I}[\sigma]$, we have $\mathcal{A}(\sigma) = \{\{a\}, \{b\}, \{a, b\}\}$. \blacksquare

More generally, the encodings of related contracts are related session types, except that the direction of the preorder is reversed.

Theorem 4. *Let σ and τ strongly convergent. Then $\mathbf{I}[\sigma] \preceq \mathbf{J}[\tau]$ if and only if $\mathcal{S}[\llbracket \mathbf{J}[\tau] \rrbracket] \leq \mathcal{S}[\llbracket \mathbf{I}[\sigma] \rrbracket]$.*

Proof. (“only if” part) Let \mathcal{R} be the least relation such that

- if $\mathbf{I}[\sigma] \preceq \mathbf{J}[\tau]$ and σ and τ are strongly convergent, then $\mathcal{S}[\llbracket \mathbf{J}[\tau] \rrbracket] \mathcal{R} \mathcal{S}[\llbracket \mathbf{I}[\sigma] \rrbracket]$; additionally, if $\mathbf{K} \subseteq \mathbf{I}$ then $\mathcal{S}[\llbracket \mathbf{nf}(\tau) \rrbracket_{\mathbf{K}}] \mathcal{R} \mathcal{S}[\llbracket \mathbf{nf}(\sigma) \rrbracket_{\mathbf{K}}]$.

We prove that \mathcal{R} is a coinductive subtyping. Let $S \mathcal{R} T$. We have two possibilities, according to the definition of \mathcal{R} .

1. ($S = \mathcal{S}[\llbracket \mathbf{J}[\tau] \rrbracket]$, $T = \mathcal{S}[\llbracket \mathbf{I}[\sigma] \rrbracket]$, and $\mathbf{I}[\sigma] \preceq \mathbf{J}[\tau]$) By definition of $\mathcal{S}[\llbracket \cdot \rrbracket]$ we have $S = \&\langle \widehat{\mathbf{K}} : \mathcal{S}[\llbracket \mathbf{nf}(\tau) \rrbracket_{\mathbf{K}}] \rangle^{\mathbf{K} \subseteq \mathbf{J}}$ and $T = \&\langle \widehat{\mathbf{K}} : \mathcal{S}[\llbracket \mathbf{nf}(\sigma) \rrbracket_{\mathbf{K}}] \rangle^{\mathbf{K} \subseteq \mathbf{I}}$. From $\mathbf{I} \subseteq \mathbf{J}$ we have $\{\widehat{\mathbf{K}} \mid \mathbf{K} \subseteq \mathbf{I}\} \subseteq \{\widehat{\mathbf{K}} \mid \mathbf{K} \subseteq \mathbf{J}\}$, hence each label in the topmost branch of T also occurs as a label in the topmost branch of S . Now take $\mathbf{K} \subseteq \mathbf{I}$. We have to show that $\mathcal{S}[\llbracket \mathbf{nf}(\tau) \rrbracket_{\mathbf{K}}] \mathcal{R} \mathcal{S}[\llbracket \mathbf{nf}(\sigma) \rrbracket_{\mathbf{K}}]$, but this follows immediately from the definition of \mathcal{R} and from $\mathbf{I}[\sigma] \preceq \mathbf{J}[\tau]$.
2. ($S = \mathcal{S}[\llbracket \mathbf{nf}(\tau) \rrbracket_{\mathbf{K}}]$, $T = \mathcal{S}[\llbracket \mathbf{nf}(\sigma) \rrbracket_{\mathbf{K}}]$, $\mathbf{K} \subseteq \mathbf{I}$, and $\mathbf{I}[\sigma] \preceq \mathbf{J}[\tau]$) We have several subcases depending on the shape of the normal form of σ and τ . Assume that $\mathbf{nf}(\sigma) = \bigoplus_{\mathbf{R} \in \mathcal{A}(\sigma)} \sum_{\alpha \in \mathbf{R}} \alpha \cdot \sigma_{\alpha}$ and $\mathbf{nf}(\tau) = \bigoplus_{\mathbf{S} \in \mathcal{A}(\tau)} \sum_{\beta \in \mathbf{S}} \beta \cdot \tau_{\beta}$. Then $S = \bigoplus \langle \widehat{\mathbf{S} \cap \mathbf{K}} : \&\langle \widehat{\emptyset} : \mathbf{end}; \{\beta\} : \mathcal{S}[\llbracket \tau_{\beta} \rrbracket_{\mathbf{K}}] \rangle^{\beta \in \mathbf{S} \cap \mathbf{K}} \rangle_{\mathbf{S} \in \mathcal{A}(\tau)}$ and $T = \bigoplus \langle \widehat{\mathbf{R} \cap \mathbf{K}} : \&\langle \widehat{\emptyset} : \mathbf{end}; \{\alpha\} : \mathcal{S}[\llbracket \sigma_{\alpha} \rrbracket_{\mathbf{K}}] \rangle^{\alpha \in \mathbf{R} \cap \mathbf{K}} \rangle_{\mathbf{R} \in \mathcal{A}(\sigma)}$. From $\mathbf{I}[\sigma] \preceq \mathbf{J}[\tau]$ we have that $\{\mathbf{S} \cap \mathbf{K} \mid \mathbf{S} \in \mathcal{A}(\tau)\} \subseteq \{\mathbf{R} \cap \mathbf{K} \mid \mathbf{R} \in \mathcal{A}(\sigma)\}$, hence each label in the topmost choice of S also occurs as a label in the topmost choice of T . Take $\mathbf{S} \in \mathcal{A}(\tau)$. The label $\widehat{\mathbf{S} \cap \mathbf{K}}$ occurs in both S and T , hence the corresponding branches have exactly the same set of labels $\{\widehat{\emptyset}\} \cup \{\{\alpha\} \mid \alpha \in \mathbf{S} \cap \mathbf{K}\}$. Let $S' = \mathcal{S}[\llbracket \tau_{\alpha} \rrbracket_{\mathbf{K}}]$ and $T' = \mathcal{S}[\llbracket \sigma_{\alpha} \rrbracket_{\mathbf{K}}]$. From $\mathbf{I}[\mathbf{nf}(\sigma)] \preceq \mathbf{J}[\mathbf{nf}(\tau)]$ and $\alpha \in \mathbf{I}$ we have that $\mathbf{nf}(\tau) \xrightarrow{\alpha} \text{implies } \mathbf{nf}(\sigma) \xrightarrow{\alpha}$ and $\mathbf{I}[\sigma_{\alpha}] \preceq \mathbf{J}[\tau_{\alpha}]$, so we conclude $S' \mathcal{R} T'$ by definition of \mathcal{R} . Assume that $\mathbf{nf}(\sigma) = \mathbf{rec } x. \sigma'$, where σ' is in normal form. Then $\mathbf{I}[\sigma] \simeq \mathbf{I}[\sigma' \{ \mathbf{nf}(\sigma) / x \}]$ and $\sigma' \{ \mathbf{nf}(\sigma) / x \}$ is itself in normal form. Since σ is strongly convergent, we eventually reach some σ'' in normal form such that σ'' does not begin with a recursion and $\mathbf{I}[\sigma] \simeq \mathbf{I}[\sigma'']$. Similarly for τ . Hence, we easily reduce to the previous subcase.

(“if” part) We prove that $\mathbf{I}[\sigma] \not\preceq \mathbf{J}[\tau]$ implies $\mathcal{S}[\llbracket \mathbf{J}[\tau] \rrbracket] \not\leq \mathcal{S}[\llbracket \mathbf{I}[\sigma] \rrbracket]$. It suffices to consider all the possibilities by which $\mathbf{I}[\sigma] \not\preceq \mathbf{J}[\tau]$ directly:

1. if $\mathbf{I} \not\subseteq \mathbf{J}$, then there exists $\alpha \in \mathbf{I}$ such that $\alpha \notin \mathbf{J}$. Then $\{\alpha\}$ is a label occurring in the topmost branch of $\mathcal{S}[\llbracket \mathbf{I}[\sigma] \rrbracket]$ but not occurring in the topmost branch of $\mathcal{S}[\llbracket \mathbf{J}[\tau] \rrbracket]$, hence $\mathcal{S}[\llbracket \mathbf{J}[\tau] \rrbracket] \not\leq \mathcal{S}[\llbracket \mathbf{I}[\sigma] \rrbracket]$;
2. let $\mathbf{R}_1, \dots, \mathbf{R}_n$ be the ready sets of σ and assume that there exists \mathbf{R} such that $\tau \Downarrow \mathbf{R}$ and for every $1 \leq i \leq n$ we have $\mathbf{R}_i \not\subseteq \mathbf{R}$, that is for every $1 \leq i \leq n$ there exists

- $\alpha_i \in R_i \cap I$ and $\alpha_i \notin R$. From $I \subseteq J$ we know that both $\mathcal{S}[I[\sigma]]$ and $\mathcal{S}[J[\tau]]$ have the label \widehat{I} in their corresponding topmost branches with continuations $\mathcal{S}[\mathbf{nf}(\sigma)]_I$ and $\mathcal{S}[\mathbf{nf}(\tau)]_I$ respectively. Then, by the fact that $\widehat{\cdot}$ is injective, it follows that $\widehat{R \cap I} \notin \{\widehat{R_i \cap I} \mid R_i \in \mathcal{A}(\sigma)\}$, hence $\widehat{R \cap I}$ is a label occurring in the topmost choice of $\mathcal{S}[\mathbf{nf}(\tau)]_I$ but not occurring in the topmost choice of $\mathcal{S}[\mathbf{nf}(\sigma)]_I$, so we conclude $\mathcal{S}[J[\tau]] \not\leq \mathcal{S}[I[\sigma]]$;
3. assume that there exists $\alpha \in I$ such that $\tau \xrightarrow{\alpha}$ and $\sigma \not\xrightarrow{\alpha}$. Then $\tau \Downarrow S$ where $\alpha \in S$ whereas $\sigma \Downarrow R$ implies $\alpha \notin R$. Hence we can reason as for the previous case and conclude $\mathcal{S}[J[\tau]] \not\leq \mathcal{S}[I[\sigma]]$. \square

The strict correspondence between \leq and \preceq allows one to use them interchangeably, as discussed in the example below.

Example 3. Consider the client contracts

$$\mathbf{K}[\rho] \stackrel{\text{def}}{=} \{\bar{a}, \bar{b}, \mathbf{e}\}[\bar{a}.\mathbf{e} + \bar{b}.\mathbf{e}] \quad \text{and} \quad \mathbf{K}[\rho'] \stackrel{\text{def}}{=} \{\bar{a}, \bar{b}, \mathbf{e}\}[\bar{a}.\mathbf{e} \oplus \bar{b}.\mathbf{e}]$$

and notice that $\mathbf{K}[\rho] \dashv I[\sigma]$ whereas $\mathbf{K}[\rho'] \not\vdash I[\sigma]$, where $I[\sigma]$ is the service contract defined in Example 2. The respective encodings of these two client contracts are

$$\begin{aligned} \mathcal{C}[\mathbf{K}[\rho]] &= \oplus \langle \widehat{\{a, b\}} : \& \langle \widehat{\{a\}} : \oplus \langle \widehat{\{a\}} : \text{end} \rangle \\ &\quad \widehat{\{b\}} : \oplus \langle \widehat{\{b\}} : \text{end} \rangle \\ &\quad \widehat{\{a, b\}} : \oplus \langle \widehat{\{a\}} : \text{end}; \widehat{\{b\}} : \text{end} \rangle \rangle \rangle \\ \mathcal{C}[\mathbf{K}[\rho']] &= \oplus \langle \widehat{\{a, b\}} : \& \langle \widehat{\{a, b\}} : \oplus \langle \widehat{\{a\}} : \text{end}; \widehat{\{b\}} : \text{end} \rangle \rangle \rangle \end{aligned}$$

and now we notice that $\mathcal{S}[I[\sigma]] \leq \text{dual}(\mathcal{C}[\mathbf{K}[\rho]])$, namely it is safe to use $\mathcal{S}[I[\sigma]]$ to interact successfully with $\mathcal{C}[\mathbf{K}[\rho]]$. On the other hand $\mathcal{S}[I[\sigma]] \not\leq \text{dual}(\mathcal{C}[\mathbf{K}[\rho']])$ because $\mathcal{S}[I[\sigma]]$ may be in a state where only a or only b are available, whilst the client $\mathbf{K}[\rho']$ autonomously decides which of the two actions to execute. Notice however that $\mathcal{S}[I[\tau]] \leq \text{dual}(\mathcal{C}[\mathbf{K}[\rho']])$. \blacksquare

More generally, if $\mathbf{K}[\rho] \dashv I[\sigma]$ holds then it is safe to use the service $\mathcal{S}[I[\sigma]]$ to interact with the client $\mathcal{C}[\mathbf{K}[\rho]]$ and in fact $\text{dual}(\mathcal{C}[\mathbf{K}[\rho]])$ is the *principal service type* that interacts successfully with $\mathcal{C}[\mathbf{K}[\rho]]$.

Theorem 5. $\mathbf{K}[\rho] \dashv I[\sigma]$ if and only if $\mathcal{S}[I[\sigma]] \leq \text{dual}(\mathcal{C}[\mathbf{K}[\rho]])$.

Proof. (“only if” part) Let \mathcal{R} be the least relation such that

- if $\mathbf{K}[\rho] \dashv I[\sigma]$ and ρ and σ are strongly convergent, then $\mathcal{S}[I[\sigma]] \mathcal{R} \text{dual}(\mathcal{C}[\mathbf{K}[\rho]])$ and $\mathcal{S}[\mathbf{nf}(\sigma)]_{\overline{\mathbf{K} \setminus \{\mathbf{e}\}}} \mathcal{R} \text{dual}(\mathcal{C}[\mathbf{nf}(\rho)]_{\mathbf{K} \setminus \{\mathbf{e}\}})$.

It is sufficient to show that \mathcal{R} is a coinductive subtyping. Let $S \mathcal{R} T$. We have two possibilities.

1. ($S = \mathcal{S}[I[\sigma]]$, $T = \text{dual}(\mathcal{C}[\mathbf{K}[\rho]])$, and $\mathbf{K}[\rho] \dashv I[\sigma]$) Then $S = \& \langle \widehat{H} : \mathcal{S}[\mathbf{nf}(\sigma)]_H \rangle^{H \subseteq I}$ and $T = \& \langle \widehat{\mathbf{K} \setminus \{\mathbf{e}\}} : \text{dual}(\mathcal{C}[\mathbf{nf}(\rho)]_{\mathbf{K} \setminus \{\mathbf{e}\}}) \rangle$. From $\mathbf{K} \setminus \{\mathbf{e}\} \subseteq I$ we have that the label in the topmost branch of T also occurs as a label in the topmost label of S . We must show that $\mathcal{S}[\mathbf{nf}(\sigma)]_H \mathcal{R} \text{dual}(\mathcal{C}[\mathbf{nf}(\rho)]_{\mathbf{K} \setminus \{\mathbf{e}\}})$, but this is obvious by definition of \mathcal{R} .

2. ($S = \mathcal{S}[\llbracket \mathbf{nf}(\sigma) \rrbracket_{\bar{K}}]$, $T = \mathbf{dual}(\mathcal{C}[\llbracket \mathbf{nf}(\rho) \rrbracket_K])$, and $(K \cup \{e\})[\rho] \dashv I[\sigma]$) Let $\mathbf{nf}(\sigma) = \bigoplus_{R \in \mathcal{A}(\sigma)} \sum_{\alpha \in R} \alpha \cdot \sigma_\alpha$ and $\mathbf{nf}(\rho) = \bigoplus_{R \in \mathcal{A}(\rho)} \sum_{\alpha \in R} \alpha \cdot \rho_\alpha$. Then

$$\begin{aligned} S &= \bigoplus \langle \widehat{R \cap \bar{K}} : \& \langle \widehat{\emptyset} : \mathbf{end}; \{\widehat{\alpha}\} : \mathcal{S}[\llbracket \mathbf{nf}(\sigma_\alpha) \rrbracket_{\bar{K}}] \rangle_{\alpha \in R \cap \bar{K}} \rangle_{R \in \mathcal{A}(\sigma)} \\ T &= \bigoplus \langle \widehat{S} : \& \langle \widehat{\emptyset} : \mathbf{end}; \{\widehat{\alpha}\} : \mathbf{dual}(\mathcal{C}[\llbracket \mathbf{nf}(\rho_\alpha) \rrbracket_K]) \rangle_{\alpha \in S \cap R} \rangle_{R \in \mathcal{A}(\rho), \alpha \in S \cap R} \rangle_{S \subseteq K, S \bowtie \mathcal{A}(\rho)} \end{aligned}$$

Let $R \in \mathcal{A}(\sigma)$ be a ready set of the service. Then $\overline{R \cap \bar{K}} = \bar{R} \cap K \subseteq K$. Furthermore, by definition of \bowtie , we have that $R \cap \bar{K} \bowtie \mathcal{A}(\rho)$ if and only if for every $R' \in \mathcal{A}(\rho)$ we have either $e \in R'$ or $R \cap \bar{K} \cap R' \neq \emptyset$. But $R' \subseteq K$, hence $R \cap \bar{K} \cap R' = \bar{R} \cap R'$. So $R \cap \bar{K} \bowtie \mathcal{A}(\rho)$ is a direct consequence of the hypothesis $(K \cup \{e\})[\rho] \dashv I[\sigma]$. Furthermore, from the same hypothesis and from $\mathbf{nf}(\rho) \xrightarrow{\alpha} \mathbf{nf}(\sigma) \xrightarrow{\bar{\alpha}}$ it follows that $(K \cup \{e\})[\rho_\alpha] \dashv I[\sigma_\alpha]$, hence we conclude $\mathcal{S}[\llbracket \mathbf{nf}(\sigma_\alpha) \rrbracket_{\bar{K}}] \mathcal{R} \mathbf{dual}(\mathcal{C}[\llbracket \mathbf{nf}(\rho_\alpha) \rrbracket_K])$ by definition of \mathcal{R} .

(“if” part) Trivial by the definitions of the encoding, the details are left to the reader. \square

By combining Theorem 5 and Proposition 3(3) we derive an interesting dual result, by which $\mathbf{dual}(\mathcal{S}[\llbracket I[\sigma] \rrbracket])$ is the *principal client type* that interacts successfully with $\mathcal{S}[\llbracket I[\sigma] \rrbracket]$.

Corollary 1. $K[\rho] \dashv I[\sigma]$ if and only if $\mathcal{C}[K[\rho]] \leq \mathbf{dual}(\mathcal{S}[\llbracket I[\sigma] \rrbracket])$.

6 Discussion

The encoding of session types into contracts (Section 4) is simple and almost homomorphic with respect to the operations. The branch is the only operation whose encoding requires some care, by adding extra $\ell.\Omega$ subterms in the generated contract for those labels ℓ not mentioned in the session type. That is, the encoding renders clearly that processes involved in a session may only perform actions that are explicitly allowed by the type of the partner process. In other words, session types describe communications where at any time exactly one of the two interacting parties has control and no hand-shaking occurs. Interestingly, by interpreting Ω as a catastrophic state of a process and by considering a slightly weaker notion of subcontract relation called *safe must* in [2], we have $\alpha.\Omega \preceq \mathbf{0}$, meaning that, in practice, action α is *not* guaranteed.

The encoding of contracts into session types (Section 5) manifests an exponential blow up of the encoded contract. This indicates that contracts are more abstract than session types and are capable of expressing more complex synchronization scenarios. Part of this added expressiveness derives from the operators $+$ and \oplus , which can be used for composing arbitrary behaviors in a liberal way. As we have anticipated in the introduction, the encoding of $\{a, b\}[a \oplus b]$ explicitly notifies the client if only a is available, or if only b is available, or if both a and b are available. This is possible only if the service has centralized control over its own resources and can decide about the availability of a and b . If, on the other hand, the service is a collection of possibly distributed processes (as in service choreographies) it may be impractical or impossible to provide such information to the client.

The lesson we learn is that there is a strict correspondence between contracts and session types that allows one to use them interchangeably *within the context of dyadic interactions*, where both parties provide centralized control on the communication protocol. Contracts go beyond session types in that they permit to characterize *arbitrary processes* in addition to *sessions*. For instance, the process $\bar{a} \mid \bar{b}$, which stands for the parallel composition of two smaller processes sending messages \bar{a} and \bar{b} , can be seen as having two unrelated sessions with type $\oplus\langle a : \text{end} \rangle$ and $\oplus\langle b : \text{end} \rangle$. On the other hand, the *whole* process can be typed according to one of the following behaviors

$$\Omega \quad \bar{a}.\Omega + \bar{b}.\Omega \quad \bar{a}.\bar{b} + \bar{b}.\Omega \quad \bar{a}.\Omega + \bar{b}.\bar{a} \quad \bar{a}.\bar{b} + \bar{b}.\bar{a}$$

representing regular, increasingly accurate approximations of the process behavior.

Recently type systems with sessions have been extended so as to guarantee stronger progress properties [8]. Such type systems must necessarily consider a broader perspective that takes into account the mutual dependencies and interactions between sessions. In this respect, we plan to investigate whether the additional expressiveness of contracts already provides the required machinery for dealing with these issues.

References

1. A. Asperti and C. Laneve. Interaction systems I: The theory of optimal reductions. *Mathematical Structures in Computer Science*, 4(4):457–504, 1994.
2. M. Boreale, R. De Nicola, and R. Pugliese. Basic observables for processes. *Information and Computation*, 149(1):77–98, 1999.
3. S. Carpineti, G. Castagna, C. Laneve, and L. Padovani. A formal account of contracts for Web Services. In *WS-FM'06, 3rd Int. Workshop on Web Services and Formal Methods*, number 4184 in LNCS, pages 148–162. Springer, 2006.
4. G. Castagna, N. Gesbert, and L. Padovani. A theory of contracts for Web Services. In *POPL'08: Proceedings of 35th ACM SIGPLAN SIGACT Symposium on Principles of Programming Languages*. ACM, 2008.
5. R. De Nicola. Two complete axiom systems for a theory of communicating sequential processes. *Information and Control*, 64(1–3):136–172, 1985.
6. R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133, 1984.
7. R. De Nicola and M. Hennessy. CCS without τ 's. In *TAPSOFT'87/CAAP'87*, number 249 in LNCS, pages 138–152. Springer, 1987.
8. M. Dezani-Ciancaglini, U. de' Liguoro, and N. Yoshida. On Progress for Structured Communications. In *TGC'07*, LNCS. Springer, 2007. To appear.
9. M. Dezani-Ciancaglini, D. Mostrous, N. Yoshida, and S. Drossopoulou. Session Types for Object-Oriented Languages. In *ECOOP'06*, volume 4067 of LNCS, pages 328–352. Springer, 2006.
10. S. Gay and M. Hole. Subtyping for session types in the π -calculus. *Acta Informatica*, 42(2–3):191–225, 2005.
11. M. Hennessy. Acceptance trees. *JACM: Journal of the ACM*, 32(4):896–928, 1985.
12. M. Hennessy. *Algebraic Theory of Processes*. Foundation of Computing. MIT Press, 1988.
13. K. Honda. Types for dyadic interaction. In *CONCUR'93*, number 715 in LNCS, pages 509–523, 1993.

14. C. Laneve and L. Padovani. The *must* preorder revisited – an algebraic theory for web services contracts. In *CONCUR'07, 18th Int. Conference on Concurrency Theory*, number 4703 in LNCS, pages 212–225. Springer, 2007. Full version available at http://www.sti.uniurb.it/padovani/Papers/lncs_4703_full.pdf.
15. V. Vasconcelos, S. Gay, and A. Ravara. Type checking a multithreaded functional language with session types. *TCS: Theoretical Computer Science*, 368, 2006.