AperTO - Archivio Istituzionale Open Access dell'Università di Torino

An Algebraic Theory for Web Service Contracts

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/137478> since 2019-10-28T12:20:04Z

Publisher:

SPRINGER-VERLAG BERLIN

Published version:

DOI:10.1007/978-3-642-38613-8_21

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

An Algebraic Theory for Web Service Contracts

Cosimo Laneve¹ and Luca Padovani²

¹ Università di Bologna – INRIA Focus Team, Italy

² Università di Torino – Dipartimento di Informatica, Italy

Abstract. We study a natural notion of compliance between clients and services in terms of their BPEL (abstract) descriptions. The induced preorder shows interesting connections with the *must* preorder and has normal form representatives that are parallel-free finite-state activities, called *contracts*. The preorder also admits the notion of least service contract that is compliant with a client contract, called *principal dual contract*. Our framework serves as a foundation of Web service technologies for connecting abstract and concrete service definitions and for service discovery.

1 Introduction

Service-oriented technologies and Web services have been proposed as a new way of distributing and organizing complex applications across the Internet. These technologies are nowadays extensively used for delivering cloud computing platforms.

A large effort in the development of Web services has been devoted to their specification, their publication, and their use. In this context, the Business Process Execution Language for Web Services (BPEL for short) has emerged as the de-facto standard for implementing Web service composition and, for this reason, it is supported by the toolkits of the main software vendors (Oracle Process Manager, IBM WebSphere, and Microsoft BizTalk).

As regards publication, service descriptions should retain abstract (behavioral) definitions, which are separate from the binding to a concrete protocol or endpoint. The current standard is defined by the Web Service Description Language (WSDL) [10], which specifies the format of the exchanged messages – the *schema* –, the locations where the interactions are going to occur – the *interface* –, the transfer mechanism to be used (i.e. SOAP-RPC, or others), and basic service abstractions (one-way/asynchronous and request-response/synchronous patterns of conversations). Since these abstractions are too simple for expressing arbitrary, possibly cyclic protocols of exchanged messages between communicating parties, WSDL is not adequate to verify the behavioral compliance between parties. It is also worth to notice that the attempts, such as UDDI (Universal Description, Discovery and Integration) registries [5], provide limited support because registry items only include pointers to the locations of the service abstractions, without constraining the way these abstractions are defined or related to the actual implementations (*cf.* the `<tModel>` element). In this respect, UDDI

registries are almost useless for discovering services; an operation that is performed manually by service users and consumers.

The publication of abstract service descriptions, called *contracts* in the following, and the related ability of service searching assume the existence of a formal notion of contract equivalence and, more generally, of a formal theory for reasoning about Web services by means of their contracts. We identify three main goals of a theory of Web service contracts: (1) it should provide a formal language for describing Web services at a reasonable level of abstraction and for admitting static correctness verification of client/service protocol implementations; (2) it should provide a semantic notion of contract equivalence embodying the principle of safe Web service replacement. Indeed, the lack of a formal characterization of contracts only permits excessively demanding notions of equivalence such as nominal or structural equality; (3) it should provide tools for effectively and efficiently searching Web services in Web service repositories according to their contract.

The aim of this contribution is to provide a suitable theory of contracts for Web services by developing a semantic notion of contract equivalence. In fact, we relax the equivalence into a *subcontract preorder*, so that Web services exposing “larger” contracts can be *safely* returned as results of queries for Web services with “smaller” contracts. We will precisely define what “smaller” and “larger” mean, and we will define which safety property we wish to preserve when substituting a service exposing a contract with a service exposing a larger contract. Our investigation abstracts away from the syntactical details of schemas as well as from those aspects that are oriented to the actual implementations, such as the definition of transmission protocols; all these aspects may be easily integrated on top of the formalism. We do not commit to a particular interpretation of the actions occurring in contracts either: they can represent different typed channels on which interaction occurs or different types of messages.

To equip contracts with a subcontract preorder, we commit to a testing approach. We define client satisfaction as the ability of the client to successfully complete *every* interaction with the service; here “successfully” means that the client never gets stuck (this notion is purposefully asymmetric as client’s satisfaction is our main concern). The preorder arises by comparing the sets of clients satisfied by services.

The properties enjoyed by the subcontract preorder are particularly relevant in the context of Web services. Specifically, it is possible to determine, given a client exposing a certain behavior, the smallest (according to subcontract preorder) service contract that satisfies the client – the *principal dual contract*. This contract, acting like a *principal type* in type systems, guarantees that a query to a Web service registry is answered with the largest possible set of compatible services in the registry’s databases.

Related works. Our contracts are normal forms of τ -less CCS processes, a calculus developed by De Nicola and Hennessy in a number of contributions [13, 15, 18]. The use of formal models to describe communication protocols is not new (see for instance the exchange patterns in SSDL [20], which are based on CSP and the π -

calculus), nor is it the use of CCS processes as behavioral types (see [19] and [9]). The subcontract relation $\lesssim$ has been introduced in [17]. In [6] the authors have studied a refined version of $\lesssim$ that is more suited for orchestrations. The works that are more closely related to ours are by Castagna *et al.* [8] and the ones on *session types*, especially [14] by Gay and Hole. The authors of [8] make the assumption that client and service can be mediated by a *filter*, which prevents potentially dangerous interactions by dynamically changing the interface of the service as it is seen by the client. The present work can be seen as a special case of [8] in which the filter is static and consequently is unnecessary; at the same time, in the present work we also consider divergence, which is not addressed in [8]. With respect to [14] (and systems based on session types) our contract language is much simpler and it can express more general forms of interaction. While the language defined in [14] supports first-class sessions and name passing, it is purposefully tailored so that the transitivity problems mentioned above are directly avoided at the language level. This restricts the subcontract relation in such a way that internal and external choices can never be related (hence, $\{a, b\} : a \oplus b \preceq \{a, b\} : a + b$ does *not* hold).

As regards schemas, which are currently part of BPEL contracts, it is worth mentioning that they have been the subject of formal investigation by several research projects [4, 16]. This work aims at pursuing a similar objective, but moving from the description of data to the description of behaviors.

Structure of the paper. In Section 2 we introduce BPEL abstract activities and their semantics. In Section 3 we define contracts and detail their relationship with BPEL abstract activities. In Section 4 we address the issue of service discovery in repositories. We conclude in Section 5. Due to space limitations, proofs have been omitted; they can be found in the full paper.

2 BPEL and the abstract language

We introduce the basic notions of BPEL by means of an example. The XML document in Figure 1 describes the behavior of an order service that interacts with four other partners, one of them being the customer (**purchasing**), the others being providers of price (**invoicing** service), shipment (**shipping** service), and manufacturing scheduling (**scheduling** service). The business process is made of *activities*, which can be either *atomic* or *composite*. In this example atomic activities are represented by *invocation* of operations in other partners (lines 7–9, 15–18, 22–25), *acceptance* of messages from other partners, either as incoming requests (line 3) or as responses to previous invocations (lines 10–12 and 19), and *sending* of responses to clients (line 28). Atomic activities are composed together into so-called structured activities, such as *sequential composition* (see the **sequence** fragments) and *parallel composition* (see the **flow** fragment at lines 4–27). In a **sequence** fragment, all the child activities are executed in the order in which they appear, and each activity begins the execution only after the previous one has completed. In a **flow** fragment, all the child activities are executed in

```

1 <process>
2   <sequence>
3     <receive partnerLink="purchasing" operation="sendPurchaseOrder"/>
4     <flow>
5       <links> <link name="ship-to-invoice"/> <link name="ship-to-scheduling"/> </links>
6       <sequence>
7         <invoke partnerLink="shipping" operation="requestShipping">
8           <sources> <source linkName="ship-to-invoice"/> </sources>
9         </invoke>
10        <receive partnerLink="shipping" operation="sendSchedule">
11          <sources> <source linkName="ship-to-scheduling"/> </sources>
12        </receive>
13      </sequence>
14      <sequence>
15        <invoke partnerLink="invoicing" operation="initiatePriceCalculation"/>
16        <invoke partnerLink="invoicing" operation="sendShippingPrice">
17          <targets> <target linkName="ship-to-invoice"/> </targets>
18        </invoke>
19        <receive partnerLink="invoicing" operation="sendInvoice"/>
20      </sequence>
21      <sequence>
22        <invoke partnerLink="scheduling" operation="requestProductionScheduling"/>
23        <invoke partnerLink="scheduling" operation="sendShippingSchedule">
24          <targets> <target linkName="ship-to-scheduling"/> </targets>
25        </invoke>
26      </sequence>
27    </flow>
28    <reply partnerLink="purchasing" operation="sendPurchaseOrder"/>
29  </sequence>
30 </process>

```

Fig. 1. BPEL business process for an e-commerce service.

parallel, and the whole **flow** activity completes as soon as *all* the child activities have completed. It is possible to constrain the execution of parallel activities by means of *links*. In the example, there is a link **ship-to-invoice** declared at line 5 and used in lines 8 and 17, meaning that the invocation at lines 16–18 cannot take place before the one at lines 7–9. Similarly, the link **ship-to-scheduling** means that the invocation at lines 23–25 cannot take place before the receive operation at lines 10–12 has completed. Intuitively, the presence of links limits the possible interleaving of the activities in a **flow** fragment.

Note that BPEL includes other conventional constructs not shown in the example, such as conditional and iterative execution of activities.

To pursue our formal investigation, we will now present an abstract language of processes whose operators correspond to those found in BPEL. Since we will focus on the interactions of BPEL activities with the external environment, rather than on the actual implementation of business processes, our process language overlooks details regarding internal, unobservable evaluations. For example, the BPEL activity

```

<if>
  <condition> bool-expr </condition>
  activity-True
<else> activity-False </else>
</if>

```

will be abstracted into the process $\text{activity-True} \oplus \text{activity-False}$, meaning that one of the two activities will be performed and the choice will be a conse-

quence of some unspecified internal decision. A similar observation pertains to the `<while>` activity (see Remark 2.2).

2.1 Syntax of BPEL abstract activities

Let $\mathbf{N}$ be a set of *names*, ranged over by $a, b, c, \dots$, and $\bar{\mathbf{N}}$ be a disjoint set of *co-names*, ranged over by $\bar{a}, \bar{b}, \bar{c}, \dots$; the term *action* refers to names and co-names without distinction; actions are ranged over by $\alpha, \beta, \dots$. Let $\bar{a} = a$. We use $\varphi, \psi, \dots$ to range over $(\mathbf{N} \cup \bar{\mathbf{N}})^*$ and $R, S, \dots$ to range over finite sets of actions. Let $\bar{R} \stackrel{\text{def}}{=} \{\bar{\alpha} \mid \alpha \in R\}$. The syntax of BPEL *abstract activities* is defined by the following grammar:

$P, Q, P_i ::= 0$	(empty)
$ a$	(receive)
$ \bar{a}$	(invoke)
$ \sum_{i \in I} \alpha_i ; P_i$	(pick)
$ P \mid_A Q$	(flow & link)
$ P ; Q$	(sequence)
$ \bigoplus_{i \in I} P_i$	(if)
$ P^*$	(while)

Each construct is called with the name of the corresponding BPEL construct. The activity 0 represents the completed process, it performs no visible action. The activity a represents the act of waiting for an incoming message. Here we take the point of view that a stands for a particular operation implemented by the process. The activity $\bar{a}$ represents the act of invoking the operation a provided by another partner. The activity $\sum_{i \in I} \alpha_i ; P_i$ represents the act of waiting for any of the α_i operations to be performed, i belonging to a *finite* set I . Whichever operation α_i is performed, it first disables the remaining ones and the continuation P_i is executed. If $\alpha_i = \alpha_j$ and $i \neq j$, then the choice whether executing P_i or P_j is implementation dependent. The process $P \mid_A Q$, where A is a set of names, represents the parallel composition (**flow**) of P and Q and the creation of a private set A of **link** names that will be used by P and Q to synchronize; an example will be given shortly. The n -ary version $\prod_{i \in 1..n}^A P_i$ of this construct may also be considered: we stick to the binary one for simplicity. The process $P ; Q$ represents the sequential composition of P followed by Q . Again we only provide a binary operator, where the BPEL one is n -ary. The process $\bigoplus_{i \in I} P_i$, again with I finite, represents an internal choice performed by the process, that results into one of the I exclusive continuations P_i . Finally, P^* represents the repetitive execution of process P so long as an internally verified condition is satisfied.

The pick activity $\sum_{i \in 1..n} \alpha_i ; P_i$ and the if activity $\bigoplus_{i \in 1..n} P_i$ will be also written $\alpha_1 ; P_1 + \dots + \alpha_n ; P_n$ and $P_1 \oplus \dots \oplus P_n$, respectively. In the following we treat (**empty**), (**receive**), and (**invoke**) as special cases of (**pick**), while at the same time keeping the formal semantics just as easy. In particular, we write 0 for $\sum_{\alpha \in \emptyset} \alpha ; P_\alpha$ and α as an abbreviation for $\sum_{\alpha \in \{\alpha\}} \alpha ; 0$ (tailing 0 are always omitted). Let also $\text{actions}(P)$ be the set of actions occurring in P .

Table 1. Legend for the operations of the BPEL process in Figure 1.

Name	Operation
a	<code>sendPurchaseOrder</code>
b	<code>requestShipping</code>
c	<code>sendSchedule</code>
d	<code>initiatePriceCalculation</code>
e	<code>sendShippingPrice</code>
f	<code>sendInvoice</code>
g	<code>requestProductionScheduling</code>
h	<code>sendShippingSchedule</code>
x	<code>ship-to-invoice</code>
y	<code>ship-to-scheduling</code>

Example 2.1. Table 1 gives short names to the operations used in the business process shown in Figure 1. Then the whole BPEL activity can be described by the term

$$a ; \left(\bar{b} ; ((\bar{x} \mid_{\emptyset} c ; \bar{y}) \mid_{\{x\}} \bar{d} ; x ; \bar{e} ; f) \mid_{\{y\}} \bar{g} ; y ; \bar{h} \right) ; \bar{a}$$

where we use names for specifying links. Such names are restricted so that they are not visible from outside. Indeed, they are completely internal to the process and should not be visible in the process' contract. $\diamond$

Remark 2.1. The BPEL specification defines a number of static analysis requirements beyond the mere syntactic correctness of processes whose purpose is to “detect any undefined semantics or invalid semantics within a process definition” [2]. Several of these requirements regard the use of links. For example, it is required that no link must cross the boundary of a repeatable construct (**while**). It is also required that link ends must be used exactly once (hence $0 \mid_{\{a\}} a$ is invalid because $\bar{a}$ is never used), and the dependency graph determined by links must be acyclic (hence $a.\bar{b} \mid_{\{a,b\}} b.\bar{a}$ is invalid because it contains cycles). These constraints may be implemented by restricting the arguments to the above abstract activities and then using static analysis techniques. $\blacksquare$

2.2 Operational semantics of BPEL abstract activities

The operational semantics of BPEL abstract activities is defined by means of a completion predicate and of a labelled transition system. Let $P\checkmark$, read *P has completed*, be the least predicate such that

$$0\checkmark \quad \frac{P\checkmark \quad Q\checkmark}{P \mid_A Q\checkmark} \quad \frac{P\checkmark \quad Q\checkmark}{P ; Q\checkmark}$$

Let μ range over actions and the special name ε denote internal moves. The operational semantics of processes is described by the following rules plus the

symmetric of the rules for $|$.

$$\begin{array}{c}
\text{(ACTION)} \quad \sum_{i \in I} \alpha_i ; P_i \xrightarrow{\alpha_i} P_i \quad \text{(IF)} \quad \bigoplus_{i \in I} P_i \xrightarrow{\varepsilon} P_i \\
\\
\text{(FLOW)} \quad \frac{P \xrightarrow{\mu} P' \quad \mu \notin A \cup \bar{A}}{P \mid_A Q \xrightarrow{\mu} P' \mid_A Q} \quad \text{(LINK)} \quad \frac{P \xrightarrow{\alpha} P' \quad Q \xrightarrow{\bar{\alpha}} Q' \quad \alpha \in A \cup \bar{A}}{P \mid_A Q \xrightarrow{\varepsilon} P' \mid_A Q'} \\
\\
\text{(SEQ)} \quad \frac{P \xrightarrow{\mu} P'}{P ; Q \xrightarrow{\mu} P' ; Q} \quad \text{(SEQ-END)} \quad \frac{P \checkmark \quad Q \xrightarrow{\mu} Q'}{P ; Q \xrightarrow{\mu} Q'} \quad \text{(WHILE-END)} \quad \frac{P^* \xrightarrow{\varepsilon} 0}{P^* \xrightarrow{\mu} P' ; P^*} \quad \text{(WHILE)} \quad \frac{P \xrightarrow{\mu} P'}{P^* \xrightarrow{\mu} P' ; P^*}
\end{array}$$

We briefly describe the rules. The process $\sum_{i \in I} \alpha_i ; P_i$ has as many α -labelled transitions as the number of actions in $\{\alpha_i \mid i \in I\}$. After a visible transition, only the selected continuation is allowed to execute. The process $\bigoplus_{i \in I} P_i$ may internally choose to behave as one of the P_i , with $i \in I$. The process $P \mid_A Q$ allows P and Q to internally evolve autonomously, or to emit messages, or to synchronize with each other on names in the set A . It completes when both P and Q have completed. The process $P ; Q$ reduces according to the reductions of P first, and of Q when P has completed. Finally, the process P^* may either complete in one step by reducing to 0 , or it may execute P one more time followed by P^* . The choice among the two possibilities is performed internally.

Remark 2.2. According to the operational semantics, P^* may execute the activity P an arbitrary number of times. This is at odds with concrete BPEL activities having P^* as abstract counterpart. For example, the BPEL activity

```

<while>
  <condition> bool-expr </condition>
  activity
</while>

```

executes **activity** as long as the **bool-expr** condition is true. Representing such BPEL activity with **activity**^{*} means *overapproximating* it. This abstraction is crucial for the decidability of our theory. $\blacksquare$

We illustrate the semantics of BPEL abstract activities through a couple of examples:

1. $(\bar{a} \oplus \bar{b} \mid_{\{a,b\}} a \oplus b) ; \bar{c} \xrightarrow{\varepsilon} (\bar{a} \mid_{\{a,b\}} a \oplus b) ; \bar{c}$ by (IF), (FLOW), and (SEQ). By the same rules, it is possible to have $(\bar{a} \mid_{\{a,b\}} a \oplus b) ; \bar{c} \xrightarrow{\varepsilon} (\bar{a} \mid_{\{a,b\}} b) ; \bar{c}$, which cannot reduce anymore ($\bar{a} \mid_{\{a,b\}} b$ is a *deadlocked* activity).
2. let $\Psi \stackrel{\text{def}}{=} 0 ; (0 \oplus 0)^*$. Then, according to rules (SEQ-END), (IF), and (WHILE), $\Psi \xrightarrow{\varepsilon} \Psi$ and $\Psi \xrightarrow{\varepsilon} 0$.

Let $\xRightarrow{\varepsilon}$ be the reflexive, transitive closure of $\xrightarrow{\varepsilon}$ and $\xRightarrow{\alpha}$ be $\xRightarrow{\varepsilon} \xrightarrow{\alpha} \xRightarrow{\varepsilon}$; let also $P \xrightarrow{\mu}$ (resp. $P \xRightarrow{\alpha}$) if there exists P' such that $P \xrightarrow{\mu} P'$ (resp. $P \xRightarrow{\alpha} P'$); we let $P \xrightarrow{\mu}$ if not $P \xrightarrow{\mu}$.

A relevant property of our BPEL abstract calculus is that the model of every activity is always finite. This result is folklore (the argument is similar to the one for CCS* [7]).

Lemma 2.1. *Let $\text{Reach}(P) = \{Q \mid \text{there are } \mu_1, \dots, \mu_n \text{ with } P \xrightarrow{\mu_1} \dots \xrightarrow{\mu_n} Q\}$. Then, for every activity P , the set $\text{Reach}(P)$ is always finite.*

We introduce a number of auxiliary definitions that will be useful in Section 3. By Lemma 2.1 these notions are trivially decidable.

Definition 2.1. *We introduce the following notation:*

- $P \uparrow$ if there is an infinite sequence of ε -transitions $P \xrightarrow{\varepsilon} \xrightarrow{\varepsilon} \dots$ starting from P . Let $P \downarrow$ if not $P \uparrow$.
- $\text{init}(P) \stackrel{\text{def}}{=} \{\alpha \mid P \xRightarrow{\alpha}\}$;
- we say that P has ready set R , notation $P \Downarrow R$, if $P \xRightarrow{\varepsilon} P'$ and $R = \text{init}(P')$;
- let $P \xRightarrow{\alpha}$. Then $P(\alpha) \stackrel{\text{def}}{=} \bigoplus_{P \xRightarrow{\varepsilon} \alpha_{P'}} P'$. We call $P(\alpha)$ the continuation of P after α .

The above definitions are almost standard, except for $P(\alpha)$ (that we already used in [17]). Intuitively, $P(\alpha)$ represents the residual behavior of P after an action α , from the point of view of the party that is interacting with P . Indeed, the party does not know which, of the possibly multiple, α -labelled branches P has taken. For example $(a ; b+a ; c+b ; d)(a) = b \oplus c$ and $(a ; b+a ; c+b ; d)(b) = d$.

2.3 The compliance preorder

We proceed defining a notion of equivalence between activities that is based on their observable behavior. To this aim, we introduce a special name **e** for denoting the successful termination of an activity (“**e**” stands for **end**). By *compliance* between a “client” activity T and a “service” activity P we mean that every interaction between T and P where P stops communicating with T is such that T has reached a successfully terminated state.

Definition 2.2 (Compliance). *Let $A_P = \{a \mid a \in \text{actions}(P) \cup \overline{\text{actions}(P)}\}$ and $\mathbf{e} \notin A_P$. The (client) activity T is compliant with the (service) activity P , written $T \dashv P$, if $P \mid_{A_P} T \xRightarrow{\varepsilon} P' \mid_{A_P} T'$ implies:*

1. if $P' \mid_{A_P} T' \not\xrightarrow{\varepsilon}$, then $\{\mathbf{e}\} \subseteq \text{init}(T')$, and
2. if $P' \uparrow$, then $\{\mathbf{e}\} = \text{init}(T')$.

We write $P \sqsubseteq Q$, called compliance preorder, if and only if $T \dashv P$ implies $T \dashv Q$ for every T . Let $\approx \stackrel{\text{def}}{=} \sqsubseteq \cap \sqsupseteq$.

According to the notion of compliance, if the client-service conversation terminates, then the client is in a successful state (it will emit an **e**-name). For example, $a ; \mathbf{e} + b ; \mathbf{e} \dashv \bar{a} \oplus \bar{b}$ and $a ; \mathbf{e} \oplus b ; \mathbf{e} \dashv \bar{a} + \bar{b}$ but $a ; \mathbf{e} \oplus b ; \mathbf{e} \not\vdash \bar{a} \oplus \bar{b}$ because of the computation $a ; \mathbf{e} \oplus b ; \mathbf{e} \mid_{\{a,b\}} \bar{a} \oplus \bar{b} \Longrightarrow a ; \mathbf{e} \mid_{\{a,b\}} \bar{b} \not\rightarrow$ 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, every reachable state of the client must be such that the only possible action is e . The practical justification of such a notion of compliance derives from the fact that connection-oriented communication protocols (like those used for interaction with Web services) typically provide for an explicit end-of-connection signal. Consider for example the client behavior $e + \bar{a}; e$. Intuitively this client tries to send a request on the name a , but it can also succeed if the service rejects the request. So $e + \bar{a}; e \dashv \mathbf{0}$ because the client can detect the fact that the service is not ready to interact on a . The same client interacting with a diverging service would have no way to distinguish a service that is taking a long time to accept the request from a service that is perpetually performing internal computations, hence $e + \bar{a}; e \not\sqsubseteq \Psi$. As a matter of facts, the above notion of compliance makes Ψ the “smallest service” – the one a client can make the least number of assumptions on (this property will be fundamental in the definition of principal dual contract in Section 4). That is $\Psi \sqsubseteq P$, for every P . As another example, we notice that $a; b + a; c \sqsubseteq a; (b \oplus c)$ since, after interacting on a , a client of the smaller service is not aware of which state the service is in (it can be either b or c). Had we picked only one a -derivative of the smaller contract behavior, we would have failed to relate it with the a -derivative of the larger contract, since both $b \not\sqsubseteq b \oplus c$ and $c \not\sqsubseteq b \oplus c$.

As by Definition 2.2, it is difficult to understand the general properties of the compliance preorder because of the universal quantification over all (client) activities T . For this reason, it is convenient to provide an alternative characterization of $\sqsubseteq$ which turns out to be the following:

Definition 2.3. A coinductive compliance is a relation $\mathcal{R}$ such that $P \mathcal{R} Q$ and $P \downarrow$ implies

1. $Q \downarrow$, and
2. $Q \downarrow \mathcal{R}$ implies $P \downarrow \mathcal{S}$ for some $\mathcal{S} \subseteq \mathcal{R}$, and
3. $Q \xRightarrow{\alpha}$ implies $P \xRightarrow{\alpha}$ and $P(\alpha) \mathcal{R} Q(\alpha)$.

Let $\preceq$ be the largest coinductive compliance relation.

The pre-order $\preceq$ corresponds to the *must-testing preorder* [15] and is also an alternative definition of $\sqsubseteq$:

Theorem 2.1. $P \sqsubseteq Q$ if and only if $P \preceq Q$.

3 Contracts

In this section we discuss how to associate a behavioral description, called *contract*, to a BPEL abstract activity. The ultimate goal is being able to reason about properties of BPEL activities by means of the respective contracts.

Contracts use a set of contract names, ranged over $C, C', C_1, \dots$. A contract is a tuple

$$(C_1 = \sigma_1, \dots, C_n = \sigma_n, \sigma)$$

where $C_j = \sigma_j$ are contract name definitions, σ is the main term, and we assume that there is no chain of definitions of the form $C_{n_1} = C_{n_2}, C_{n_2} = C_{n_3}, \dots, C_{n_k} = C_{n_1}$. The syntax of σ_j and σ is given by

$$\sigma ::= C \mid \alpha; \sigma \mid \sigma + \sigma \mid \sigma \oplus \sigma$$

where $C \in \{C_1, \dots, C_n\}$. The contract $\alpha; \sigma$ represents sequential composition in the restricted form of prefixing. The operations $+$ and $\oplus$ correspond to **pick** and **if** of BPEL activities, respectively. These operations are assumed to be associative and commutative; therefore we will write $\sigma_1 + \dots + \sigma_n$ and $\sigma_1 \oplus \dots \oplus \sigma_n$ without confusion and will sometimes shorten these contracts as $\sum_{i \in 1..n} \sigma_i$ and $\bigoplus_{i \in 1..n} \sigma_i$, respectively. The contract name C is used to model recursive behaviors such as $C = a; C$. In what follows we will leave contract name definitions implicit and identify a contract $(C_1 = \sigma_1, \dots, C_n = \sigma_n, \sigma)$ with its main body σ . We will write $\text{cnames}(\sigma)$ for the set $\{C_1, \dots, C_n\}$ and use the following abbreviations:

- $0 \stackrel{\text{def}}{=} C_0$, where $C_0 = C_0 + C_0$ represents a terminated activity;
- $\Omega \stackrel{\text{def}}{=} C_\Omega$, where $C_\Omega = C_\Omega \oplus C_\Omega$ represents divergence, that is a non-terminating activity.

Note that, even if apparently simpler, the contract language *is not* a sublanguage of BPEL abstract activities. For example, Ω cannot be written as a term in the syntax of Section 2.1. Nevertheless, in the following, we demonstrate that contracts provide alternative descriptions (with respect to the preorder $\sqsubseteq$) to BPEL abstract activities.

The operational semantics of contracts is defined by the rules below:

$$\begin{array}{c} \alpha; \sigma \xrightarrow{\alpha} \sigma \quad \sigma \oplus \rho \xrightarrow{\varepsilon} \sigma \\[10pt] \frac{\sigma \xrightarrow{\varepsilon} \sigma'}{\sigma + \rho \xrightarrow{\varepsilon} \sigma' + \rho} \quad \frac{\sigma \xrightarrow{\alpha} \sigma'}{\sigma + \rho \xrightarrow{\alpha} \sigma'} \quad \frac{C = \sigma \quad \sigma \xrightarrow{\mu} \sigma'}{C \xrightarrow{\mu} \sigma'} \end{array}$$

plus the symmetric of rules $+$ and $\oplus$. Note that $+$ evaluates the branches as long as they can perform invisible actions. This rule is absent in BPEL abstract activities because, there, the branches are always guarded by an action.

Example 3.1. The Web service conversation language WSCL [3] describes *conversations* between two parties by means of an activity diagram (Figure 2). The diagram is made of *interactions* connected with each other by *transitions*. An interaction is a basic one-way or two-way communication between the client and the server. Two-way communications are just a shorthand for two sequential one-way interactions. Each interaction has a *name* and a list of *document types* that can be exchanged during its execution. A transition connects a *source* interaction

Fig. 2. Contract of a simple e-commerce service as a WSCL diagram.

with a *destination* interaction. A transition may be *labeled* by a document type if it is active only when a message of that specific document type was exchanged during the previous interaction.

The diagram in Figure 2 describes the conversation of a service requiring clients to login before they can issue a query. After the query, the service returns a catalog. From this point on, the client can decide whether to purchase an item from the catalog or to logout and leave. In case of purchase, the service may either report that the purchase is successful, or that the item is out-of-stock, or that the client's payment is refused. By interpreting names as message types, this e-commerce service can be described by the tuple:

$$\begin{aligned}
 (& C_1 = \text{Login}; (\overline{\text{InvalidLogin}}; C_1 \oplus \overline{\text{ValidLogin}}; C_2) , \\
 & C_2 = \text{Query}; \overline{\text{Catalog}}; (C_2 + C_3 + C_4) , \\
 & C_3 = \text{Purchase}; (\overline{\text{Accepted}} \\
 & \quad \oplus \overline{\text{InvalidPayment}}; (C_3 + C_4) \\
 & \quad \oplus \overline{\text{OutOfStock}}; (C_2 + C_4)) , \\
 & C_4 = \text{Logout} , \\
 & C_1)
 \end{aligned}$$

There is a strict correspondence between unlabeled (respectively, labeled) transitions in Figure 2 and external (respectively, internal) choices in the contract. Recursion is used for expressing iteration (the cycles in the figure) so that the client is given another chance whenever an action fails for some reason. $\diamond$

We can relate BPEL abstract activities and contracts by means of the corresponding transition systems. To this aim, let X and Y range over BPEL abstract activities *and* contracts. Then, X and Y interact according to the rules

$$\begin{array}{c}
 \frac{X \xrightarrow{\varepsilon} X'}{X \parallel Y \xrightarrow{\varepsilon} X' \parallel Y} \quad \frac{Y \xrightarrow{\varepsilon} Y'}{X \parallel Y \xrightarrow{\varepsilon} X \parallel Y'} \quad \frac{X \xrightarrow{\alpha} X' \quad Y \xrightarrow{\bar{\alpha}} Y'}{X \parallel Y \xrightarrow{\varepsilon} X' \parallel Y'}
 \end{array}$$

It is possible to extend the definition of compliance to contracts and, by Definition 2.2, obtain a relation that allows us to compare activities and contracts without distinction, and similarly for $\preceq$. To be precise, the relation $X \sqsubseteq Y$ is smaller (in principle) than the relation $\sqsubseteq$ given in Definition 2.2 because, as we have said, the contract language is not a sublanguage of that of activities and, therefore, the set of tests that can be used for comparing X and Y is larger. Nonetheless, the process language used in [12] includes both BPEL abstract activities and contracts and since $\sqsubseteq$ is equivalent to must-testing, then we may

safely use the same symbol $\sqsubseteq$ for both languages. This is a key point in our argument, which will allow us to define, for every activity P , a contract σ_P such that $P \approx \sigma_P$. In particular, let C_P be the set of contract name definitions defined as follows

$$C_P = \begin{cases} \Omega & \text{if } P \uparrow \\ \bigoplus_{P \downarrow \mathbb{R}} \sum_{\alpha \in \mathbb{R}} \alpha ; C_{P(\alpha)} & \text{otherwise} \end{cases}$$

A relevant property of C_P is an immediate consequence of Lemma 2.1.

Lemma 3.1. *For every P , the set $\text{cnames}(C_P)$ is finite.*

The construction of the contract C_P with respect to a BPEL abstract activity is both correct and complete with respect to compliance:

Theorem 3.1. $P \approx C_P$.

4 Service discovery and dual contracts

We now turn our attention to the problem of querying a database of Web service contracts. To this aim, the relation $\sqsubseteq$ (and the must-testing) turns out to be too strong (see below). Following [17], we switch to more informative service contracts than what described in Section 3. In particular, we consider pairs $I : \sigma$, where I is the interface, i.e. the set of actions performed by the service, and σ is as in Section 3 (it is intended that the names occurring in σ are included into I). It is reasonable to think that a similar extension applies to client contracts: clients, which are defined by BPEL activities as well, are abstracted by terms in the language of Section 2 and, in turn, their behavior is defined by a term in the contract language, plus the interface.

Definition 4.1 (Subcontract relation). *Let $I : \sigma \lesssim J : \tau$ if $I \subseteq J$ and, for every ρ such that $\text{actions}(\rho) \setminus \{e\} \subseteq I$ and $\rho \dashv \sigma$ implies $\rho \dashv \tau$. Let $\approx$ be $\lesssim \cap \gtrsim$.*

Let us comment on the differences between $I : \sigma \lesssim J : \tau$ and $\sigma \sqsubseteq \tau$. We notice that $I : \sigma \lesssim J : \tau$ only if $I \subseteq J$. This apparently natural prerequisite has substantial consequences on the properties of $\lesssim$ because it ultimately enables width and depth extensions, which are not possible in the $\sqsubseteq$ preorder. For instance, we have $\{a\} : a \lesssim \{a, b\} : a + b$ whilst $a \not\sqsubseteq a + b$ (width extension). Similarly we have $\{a\} : a \lesssim \{a, b\} : a ; b$ whilst $a \not\sqsubseteq a ; b$ (depth extension). These extensions are desirable when searching for services, since every service offering more methods than required is a reasonable result of a query. The precise relationship between $\lesssim$ and $\sqsubseteq$ is expressed by the following statement.

Proposition 4.1. $I : \sigma \approx J : \tau$ if and only if $\sigma \approx \tau$ and $I = J$.

The basic problem for querying Web service repositories is that, given a client $\kappa : \rho$, one wishes to find all the service contracts $I : \sigma$ such that $\text{actions}(\rho) \setminus \{e\} \subseteq I$ and $\rho \dashv \sigma$. We attack this problem in two steps: first of all, we compute *one*

particular service contract $\bar{K} \setminus \{\bar{e}\} : D_\rho^K$ such that $\rho \dashv D_\rho^K$; second, we take all the services in the registry whose contract is larger than this one. In order to maximize the number of service contracts returned as answer to the query, the dual of a (client) contract $K : \rho$ should be a contract $\bar{K} \setminus \{\bar{e}\} : D_\rho^K$ such that it is the smallest service contract that satisfies the client contract $K : \rho$. We call such contract the *principal dual contract* of $K : \rho$.

In defining the principal dual contract, it is convenient to restrict the definition to those client's behaviors ρ that never lead to 0 without emitting e . For example, the behavior $a ; e + b$ describes a client that succeeds if the service proposes $\bar{a}$, but that fails if the service proposes $\bar{b}$. As far as querying is concerned, such behavior is completely equivalent to $a ; e$. As another example, the degenerate client behavior 0 is such that no service will ever satisfy it. In general, if a client is unable to handle a particular action, like b in the first example, it should simply omit that action from its behavior. We say that a (client) contract $K : \rho$ is *canonical* if, whenever $\rho \xrightarrow{\varphi} \rho'$ is maximal, then $\varphi = \varphi' e$ and $e \notin \text{actions}(\varphi')$. For example $\{a, e\} : a ; e$, $\{a\} : C$, where $C = a ; C$, and $\emptyset : \Omega$ are canonical; $\{a, b, e\} : a ; e + b$ and $\{a\} : C'$, where $C' = a \oplus C'$, are not canonical.

Observe that Lemma 2.1 also applies to contracts. Therefore it is possible to extend the notions in Definition 2.1, by replacing activities with contracts.

Definition 4.2 (Dual contract). Let $K : \rho$ be a canonical contract. The dual of $K : \rho$ is $\bar{K} \setminus \{\bar{e}\} : D_\rho^K$ where D_ρ^K is the contract name defined as follows:

$$D_\rho^K \stackrel{\text{def}}{=} \begin{cases} \Omega & \text{if } \text{init}(\rho) = \{e\} \\ \sum_{\substack{\rho \Downarrow_R \\ R \setminus \{e\} \neq \emptyset}} \left(\underbrace{0 \oplus \bigoplus_{\alpha \in R \setminus \{e\}} \bar{\alpha}}_{\text{if } e \in R} ; D_{\rho(\alpha)}^K \right) + \text{OTH}_{K \setminus \text{init}(\rho)} & \text{otherwise} \end{cases}$$

where $\text{OTH}_S \stackrel{\text{def}}{=} 0 \oplus \underbrace{\bigoplus_{\alpha \in S} \bar{\alpha}}_{\text{if } S \neq \emptyset} ; \Omega$

Few comments about D_ρ^K , when $\text{init}(\rho) \neq \{e\}$, follow. In this case, the behavior ρ may autonomously transit to different states, each one offering a particular ready set. Thus the dual behavior leaves the choice to the client: this is the reason for the external choice in the second line. Once the state has been chosen, the client offers to the service a spectrum of possible actions: this is the reason for the internal choice underneath the sum $\sum$.

The contract $\text{OTH}_{K \setminus \text{init}(\rho)}$ covers all the cases of actions that are allowed by the interface and that are not offered by the client. The point is that the dual operator must compute the principal (read, the smallest) service contract that satisfies the client, and the smallest convergent behavior with respect to a nonempty (finite) interface S is $0 \oplus \bigoplus_{\alpha \in S} \bar{\alpha} ; \Omega$. The 0 summand accounts for the possibility that none of the actions in $K \setminus \text{init}(\rho)$ is present. The external choice “+” distributes the proper dual contract over the internal choice of all the actions in $K \setminus \text{init}(\rho)$. For example, $D_{a ; e}^{\{a, \bar{a}, e\}} = \bar{a} ; \Omega + (0 \oplus a ; \Omega)$. The dual of a divergent

(canonical) client $\{A\} : C$, where $C = a ; e \oplus C$, is also well defined: $D_{C''}^{\{a,e\}} = \bar{a} ; \Omega$. We finally observe that the definition also accounts for duals of nonterminating clients, such as $\{A\} : C'$, where $C' = a ; C'$. In this case, $D_{C'}^{\{a\}} = \bar{a} ; D_{C'}^{\{a\}}$.

Similarly to the definition of contract names C_P , it is possible to prove that D_ρ^K is well defined.

Lemma 4.1. *For every $\kappa : \rho$, the set $\text{cnames}(D_\rho^K)$ is finite.*

The property that $\bar{\kappa} \setminus \{\bar{e}\} : D_\rho^K$ is the least dual contract of $\kappa : \rho$ follows.

Theorem 4.1. *Let $\kappa : \rho$ be a canonical contract. Then:*

1. $\rho \dashv D_\rho^K$;
2. if $\bar{\kappa} \setminus \{\bar{e}\} \subseteq s$ and $\rho \dashv \sigma$, then $\bar{\kappa} \setminus \{\bar{e}\} : D_\rho^K \lesssim s : \sigma$.

A final remark is about the computational complexity of the discovery algorithm. Determining $\lesssim$ is EXPTIME-complete in the size of the contracts [1], which has to be multiplied by the number of $\lesssim$ -checks (to find a compliant service in the repository) to obtain the overall cost.

5 Conclusions

In this contribution we have studied a formal theory of Web service abstract (behavioral) definitions as normal forms of a natural semantics for BPEL activities. Our abstract definitions may be effectively used in any query-based system for service discovery because they support a notion of principal dual contract. This operation is currently done in an ad hoc fashion using search engines or similar technologies.

Several future research directions stem from this work. On the technical side, a limit of our technique is that BPEL activities are “static”, *i.e.* they cannot create other services on the fly. This constraint implies the finiteness of models and, for this reason, it is possible to effectively associate an abstract description to activities. However, this impacts on scalability, in particular when services adapt to peaks of requests by creating additional services. It is well-known that such an additional feature makes models to be infinite states and requires an approximate inferential process to extract abstract descriptions from activities. Said otherwise, extending our technique to full CCS or π -calculus amounts to defining abstract finite models such that Theorem 3.1 does not hold anymore. For this reason, under- and over-estimations for services and clients, respectively, must be provided.

Another interesting technical issue concerns the extension of our study to other semantics for BPEL activities, such as the preorder in [6], or even to weak bisimulation (which has a polynomial computational cost). Perhaps one may use axiomatizations of these equivalences for determining the class of contracts. However it is not clear whether they admit a principal dual contract or not.

It is also interesting to prototyping our theory and experimenting it on some existing repository, such as <http://www.service-repository.com/>. To this aim we might use tools that have been already developed for the must testing, such as the concurrency workbench [11].

References

1. L. Aceto, A. Ingolfssdottir, and J. Srba. The algorithmics of bisimilarity. In D. Sangiorgi and J. Rutten, editors, *Advanced Topics in Bisimulation and Coinduction*, volume 52 of *Cambridge Tracts in Theoretical Computer Science*, chapter 3, pages 100–172. Cambridge University Press, 2011.
2. A. Alves et al. *Web Services Business Process Execution Language Version 2.0*, Jan. 2007. <http://docs.oasis-open.org/wsbpel/2.0/CS01/wsbpel-v2.0-CS01.html>.
3. A. Banerji, C. Bartolini, D. Beringer, V. Chopella, et al. *Web Services Conversation Language (WSCL) 1.0*, Mar. 2002. <http://www.w3.org/TR/2002/NOTE-wscl10-20020314>.
4. V. Benzaken, G. Castagna, and A. Frisch. CDuce: an XML-centric general-purpose language. *SIGPLAN Notices*, 38(9):51–63, 2003.
5. D. Beringer, H. Kuno, and M. Lemon. *Using WSCL in a UDDI Registry 1.0*, 2001. UDDI Working Draft Best Practices Document, <http://xml.coverpages.org/HP-UDDI-wscl-5-16-01.pdf>.
6. M. Bravetti and G. Zavattaro. A theory of contracts for strong service compliance. *Mathematical Structures in Computer Science*, 19:601–638, 5 2009.
7. N. Busi, M. Gabbriellini, and G. Zavattaro. On the expressive power of recursion, replication and iteration in process calculi. *Mathematical Structures in Computer Science*, 19(6):1191–1222, 2009.
8. G. Castagna, N. Gesbert, and L. Padovani. A Theory of Contracts for Web Services. *ACM Transactions on Programming Languages and Systems*, 31(5), 2009.
9. S. Chaki, S. K. Rajamani, and J. Rehof. Types as models: model checking message-passing programs. *SIGPLAN Not.*, 37(1):45–57, 2002.
10. E. Christensen, F. Curbera, G. Meredith, and S. Weerawarana. *Web Services Description Language (WSDL) 1.1*, 2001. <http://www.w3.org/TR/2001/NOTE-wsdl-20010315>.
11. R. Cleaveland, J. Parrow, and B. Steffen. The concurrency workbench: a semantics-based tool for the verification of concurrent systems. *ACM Trans. Program. Lang. Syst.*, 15(1):36–72, 1993.
12. R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theor. Comput. Sci.*, 34:83–133, 1984.
13. R. De Nicola and M. Hennessy. CCS without τ 's. In *Proceedings of TAPSOFT'87/CAAP'87*, LNCS 249, pages 138–152. Springer, 1987.
14. S. Gay and M. Hole. Subtyping for session types in the π -calculus. *Acta Informatica*, 42(2-3):191–225, 2005.
15. M. Hennessy. *Algebraic Theory of Processes*. Foundation of Computing. MIT Press, 1988.
16. H. Hosoya and B. C. Pierce. XDuce: A statically typed XML processing language. *ACM Trans. Internet Techn.*, 3(2):117–148, 2003.
17. C. Laneve and L. Padovani. The *must* preorder revisited – an algebraic theory for web services contracts. In *CONCUR'07*, LNCS 4703, pages 212–225. Springer, 2007.
18. R. Milner. *A Calculus of Communicating Systems*. Springer, 1982.
19. H. R. Nielson and F. Nielson. Higher-order concurrent programs with finite communication topology (extended abstract). In *Proceedings of POPL'94*, pages 84–97. ACM Press, 1994.
20. S. Parastatidis and J. Webber. *MEP SSDL Protocol Framework*, Apr. 2005. <http://ssdl.org>.