HAL open science

Behavioural Contracts with Request-Response Operations

Lucia Acciai, Michele Boreale, Gianluigi Zavattaro

To cite this version:

Lucia Acciai, Michele Boreale, Gianluigi Zavattaro. Behavioural Contracts with Request-Response Operations. 12th International Conference on Coordination Models and Languages (COORDINATION) Held as part of International Federated Conference on Distributed Computing Techniques (DisCoTec), Jun 2010, Amsterdam, Netherlands. pp.16-30, 10.1007/978-3-642-13414-2_2 . hal-01054626

HAL Id: hal-01054626

https://inria.hal.science/hal-01054626

Submitted on 7 Aug 2014

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

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

Distributed under a Creative Commons Attribution 4.0 International License

Behavioural contracts with request-response operations*

Lucia Acciai[1], Michele Boreale[1], and Gianluigi Zavattaro[2]

[1] Dipartimento di Sistemi e Informatica, Università di Firenze, Italy
[2] Dipartimento di Scienze dell'Informazione, Università di Bologna, Italy

Abstract In the context of service-oriented computing, behavioural contracts are abstract descriptions of the message-passing behaviour of services. They can be used to check properties of service compositions such as, for instance, client-service compliance. Previous formal models for contracts consider unidirectional *send* and *receive* operations. In this paper, we present two models for contracts with bidirectional *request-response* operations, in the presence of unboundedly many instances of both clients and servers. The first model takes inspiration from the abstract service interface language WSCL, the second one is inspired by Abstract WS-BPEL. We prove that client-service compliance is decidable in the former while it is undecidable in the latter, thus showing an interesting expressiveness gap between the modeling of *request-response* operations in WSCL and in Abstract WS-BPEL.

1 Introduction

One interesting aspect of Service Oriented Computing (SOC) and Web Services technology is the need to describe in rigorous terms not only the format of the messages exchanged among interacting parties, but also their protocol of interaction. This specific aspect is clearly described in the Introduction of the Web Service Conversation Language (WSCL) specification [25], one of the proposals of the World Wide Web Consortium (W3C) for the description of the so-called Web Services *abstract interfaces*:

> Defining which XML documents are expected by a Web service or are sent back as a response is not enough. It is also necessary to define the order in which these documents need to be exchanged; in other words, a business level conversation needs to be specified. By specifying the conversations supported by a Web service —by defining the documents to be exchanged and the order in which they may be exchanged— the external visible behavior of a Web service, its abstract interface, is defined.

The abstract interface of services can be used in several ways. For instance, one could check the *compliance* between a client and a service, that is, a guarantee for the

* The third author is partly supported by the EU integrated projects HATS and is member of the joint INRIA/University of Bologna Research Team FOCUS.

client that the interaction with the service will in any case be completed successfully. One could also check, during the service discovery phase, the *conformance* of a concrete service to a given abstract interface by verifying whether the service implements at least the expected functionalities and does not require more.

Formal models are called for to devise rigorous forms of reasoning and verification techniques for services and abstract interfaces. To this aim, theories of *behavioural contracts* based on CCS-like process calculi [19] have been thoroughly investigated [3,6,7,8,10,11,12]. However, these models lack of expressiveness in at least one respect: they cannot be employed to describe bidirectional *request-response* interactions, in contexts where several instances of the client and of the service may be running at the same time. This situation, on the other hand, is commonly found in practice-oriented contract languages, like the abstract service interface language WSCL [25] and Abstract WS-BPEL [22].

In this paper, we begin a formal investigation of contract languages of the type described above, that is allowing bidirectional request-response interaction, taking place between instances of services and clients. We present two contract languages that, for simplicity, include only the request-response pattern[1]: the first language is inspired by WSCL while the second one by Abstract WS-BPEL. In both these models, the request-response interaction pattern is decomposed into sequences of more fundamental *send-receive-reply* steps: the client first *sends* its invocation, then the service *receives* such an invocation, and finally the service sends its *reply* message back to the client. The binding between the requesting and the responding sides (instances) of the original operation is maintained by employing naming mechanisms similar to those found in the π-calculus [20]. In both models, we do not put any restriction on the number of client or service instances that can be generated at runtime, so that the resulting systems are in general infinite-state. The difference between the two models is that in the former it is not possible to describe intermediary activities of the service between the receive and the reply steps, while this is possible in the latter. In fact, WSCL models the service behaviour in request-response interactions with a unique *ReceiveSend* primitive indicating the kind of incoming and outgoing messages. On the contrary, in Abstract WS-BPEL there exist two distinct primitives that allows one to model independently the *receive* and the *reply* steps.

We define client-service compliance on the basis of the *must testing* relation of [13]: a client and a service are compliant if any sequence of interactions between them leads the client to a successful state. Our main results show that client-service compliance is decidable in the WSCL-inspired model, while it is undecidable in the Abstract WS-BPEL model: this points to an interesting expressiveness gap between the two approches for the modeling of the request-response interaction pattern. In the former case, the decidability proof is based on a translation of contracts into Petri nets. This translation is not precise, in the sense that intermediate steps of the request-response interaction are not represented in the Petri net. However, the translation is complete, in the sense that it preserves and reflects the existence of unsuccessful computations, which is enough to reduce the original compliance problem to a decidable problem in Petri nets.

[1] As discussed in Section 4, the languages that we propose are sufficiently expressive to model also the one-way communication pattern.

This yields a practical compliance-checking procedure, obtained by adaptation of the classical Karp-Miller coverability tree construction [18].

The rest of the paper is organized as follows. In Section 2 we present the two formal models and the definition of client-service compliance. Sections 3 contains the Petri nets semantics and the proof of decidability of client-service compliance for the WSCL model. Section 4 reports on undecidability for the Abstract WS-BPEL model. We draw some conclusions and discuss further work in Section 5.

2 Behavioural contracts with request-response

We presuppose a denumerable set of contract variables *Var* ranged over by X, Y, $\cdots$, a denumerable set of names *Names* ranged over by a, b, r, s, $\cdots$. We use I, J, $\cdots$ to denote a sets of indexes.

Definition 1 (WSCL Contracts). *The syntax of WSCL* contracts *is defined by the following grammar*

$$G ::= \mathit{invoke}(a, \sum_{i\in I} b_i.C_i) \mid \mathit{recreply}(a, \sum_{i\in I} b_i.C_i) \mid \surd \qquad C ::= \sum_{i\in I} G_i \mid C|C \mid X \mid \mathit{rec}X.C$$

where recX._ is a binder for the contract variable X. We assume guarded recursion*, that is, given a contract recX.C all the free occurrences of X in C are inside a guarded contract G.*

A client contract *is a contract C containing at least one occurrence of the guarded* success *contract* $\surd$*, while a* service contract *is a contract not containing* $\surd$*.*

G is used to denote guarded contracts, ready to perform either an invoke or a receive on a request-response operation a: the selection of the continuation C_i depends on the actual reply message b_i. A set of guarded contracts G_i can be combined into a choice $\sum_{i\in I} G_i$; if the index set I is empty, we denote this term by $\mathbf{0}$. Contracts can be composed in parallel. Note that infinite-state contract systems can be defined using recursion (see example later in the section). In the following, we use $\mathit{Names}(C)$ to denote the set of names occurring in C. Before presenting the semantics of WSCL contracts, we introduce BPEL contracts as well.

Definition 2 (BPEL Contracts). *BPEL contracts are defined like WSCL contracts in Definition 1, with the only difference that guarded contracts are as follows*

$$G \ ::= \ \mathit{invoke}(a, \sum_{i\in I} b_i.C_i) \quad | \quad \mathit{receive}(a).C \quad | \quad \mathit{reply}(a,b).C \quad | \quad \surd.$$

We now define the operational semantics of both models. We start by observing that the WSCL contract $\mathsf{recreply}(a, \sum_{i\in I} b_i.C_i)$ is the same as the BPEL contract $\mathsf{receive}(a).\sum_{i\in I}(\mathsf{reply}(a,b_i).C_i)$ that receives an invocation on the operation a and then replies with one of the messages b_i. We shall rely on a run-time syntax of contracts, which is obtained from the original one by extending the clause for guarded contract, thus $G \ ::= \ \cdots \mid \overline{a}\langle r\rangle \mid r\langle b\rangle.C$. Both terms $\overline{a}\langle r\rangle$ and $r\langle b\rangle.C$ are used to represent an

emitted and pending invocation of a request-response operation a: the name r represents a (fresh) channel r that will be used by the invoked operation to send the reply message back to the invoker. From now onwards we will call (WSCL) contract any term that could be obtained from this run-time syntax. In the following, we let $Labels \triangleq \{\tau, \surd\} \cup \{a\langle b\rangle, \overline{a}\langle b\rangle, (a) \mid a, b \in Names\}$.

Definition 3 (Operational semantics). *The operational semantics of a contract is given by the minimal labeled transition system, with labels taken from the set Labels, satisfying the following axiom and rules ($a, b, r \in Names$)*

$$\frac{G_l \xrightarrow{\alpha} G'_l \quad l \in I}{\sum_{i\in I} G_i \xrightarrow{\alpha} G'_l} \qquad \frac{r \notin Names(\sum_{i\in I} b_i.C_i)}{\mathsf{invoke}(a, \sum_{i\in I} b_i.C_i) \xrightarrow{(r)} \sum_{i\in I}(r\langle b_i\rangle.C_i) \mid \overline{a}\langle r\rangle}$$

$$\mathsf{receive}(a).C \xrightarrow{a\langle r\rangle} C\{^r/_a\} \qquad \mathsf{reply}(r,b).C \xrightarrow{\tau} C \mid \overline{r}\langle b\rangle \qquad \surd \xrightarrow{\surd} \mathbf{0}$$

$$\overline{r}\langle b\rangle \xrightarrow{\overline{r}\langle b\rangle} \mathbf{0} \qquad r\langle b\rangle.C \xrightarrow{r\langle b\rangle} C \qquad \frac{P \xrightarrow{\overline{a}\langle b\rangle} P' \quad Q \xrightarrow{a\langle b\rangle} Q'}{P|Q \xrightarrow{\tau} P'|Q'}$$

$$\frac{P \xrightarrow{\alpha} P' \quad \alpha \neq (r)}{P|Q \xrightarrow{\alpha} P'|Q} \qquad \frac{P \xrightarrow{(r)} P' \quad r \notin Names(Q)}{P|Q \xrightarrow{\alpha} P'|Q} \qquad \frac{C\{^{recX.C}/_X\} \xrightarrow{\alpha} C'}{recX.C \xrightarrow{\alpha} C'}$$

plus the symmetric version of the three rules for parallel composition. With $C\{^r/_a\}$ we denote the term obtained from C by replacing with r every occurrence of a not inside a $\mathsf{receive}(a).D$, while $C\{^{recX.C}/_X\}$ denotes the usual substitution of free contract variables with the corresponding definition.

In the following, we use $C \xrightarrow{\alpha}$ to say that there is some C' such that $C \xrightarrow{\alpha} C'$. Moreover, we use $C \longrightarrow C'$ to denote reductions, i.e. transitions that C can perform also when it is in isolation. Namely, $C \longrightarrow C'$ if $C \xrightarrow{\tau} C'$ or $C \xrightarrow{(r)} C'$ for some r.

We now formalize the notion of client-service compliance resorting to *must-testing* [13]. Intuitively, a client C is *compliant* with a service contract S if all the computations of the system $C|S$ lead to the client's success. Other notions of compliance have been put forward in the literature [6,7,8]; we have chosen this one because of its technical and conceptual simplicity (see e.g. [10]).

Definition 4 (Client-Service compliance). *Given a contract D, a* computation *is a sequence of reduction steps $D_1 \longrightarrow D_2 \longrightarrow \cdots \longrightarrow D_n \longrightarrow \cdots$. It is a* maximal computation *if it is infinite or it ends in a state D_n such that D_n has no outgoing reductions.*

A client contract C is compliant *with a service contract S if for every maximal computation $C|S \longrightarrow D_1 \longrightarrow \cdots \longrightarrow D_l \longrightarrow \cdots$ there exists k such that $D_k \xrightarrow{\surd}$.*

Example 1 (An impatient client and a latecomer service). This example shows that even very simple WSCL scenarios could result in infinite-state systems. Consider a client C that asks the box office service S for some tickets and then waits for them by listening

on *offerTicket*. Our client is impatient: at any time, it can decide to stop waiting and issue a new request. This behaviour can be described in WSCL as follows

$$C \triangleq recX.(\mathsf{invoke}(requireTicket, ok.X) + \mathsf{recreply}(offerTicket, ok.\surd))$$

Consider the box office service S, defined below, that is always ready to receive a *requireTicket* invocation and immediately responds by notifying (performing a call-back) on *offerTicket*.

$$S \triangleq recX.\mathsf{recreply}(requireTicket, ok.(\mathsf{invoke}(offerTicket, ok)|X))$$

It is easy to see that, in case $\mathsf{invoke}(offerTicket, \cdots)$ on the service side and $\mathsf{recreply}(offerTicket, \cdots)$ on the client side never synchronize, $C|S$ generates an infinite-state system where each state is characterized by an arbitrary number of $\mathsf{invoke}(offerTicket, \cdots)$. This infinite computation, moreover, does not traverse states in which the client can perform its $\surd$ action, thus C is not compliant with S according to Definition 4.[2]

3 Decidability of client-service compliance for WSCL contracts

We translate WSCL contract systems into place/transitions Petri nets [23], an infinite-state model in which several reachability problems are decidable (see, e.g., [15] for a review of decidable problems for finite Petri nets). The translation into Petri nets does not faithfully reproduce the operational semantics of contracts. In particular, in finite Petri nets it is not possible to represent the unbounded number of names dynamically created in contract systems to bind the reply messages to the corresponding invocations. The Petri net semantics that we present models bi-directional request-response interactions as a unique event, thus merging together the four distinct events in the operational semantics of contracts: the emission and the reception of the invocation, and the emission and the reception of the reply. We will prove that this alternative modeling preserves client-service compliance because in WSCL the invoker and the invoked contracts do not interact with other contracts during the request-response.

Another difference is that the Petri net semantics can be easily modified so that when the client contract enters in a successful state, i.e. a state with an outgoing transition $\surd$, the corresponding Petri net enters a particular successful state and blocks its execution. This way, a client contract is compliant with a service contract if and only if in the corresponding Petri net all computations are finite and finish in a *successful* state. As we will show, this last property is verifiable for finite Petri nets using the so-called *coverability* tree [18].

3.1 A Petri net semantics for WSCL contracts

We first recall the definition of Petri nets. For any set S, we let $\mathcal{M}_{fin}(S)$ be the set of the finite multisets (*markings*) over S.

[2] Other definitions of compliance, see e.g. [6], resort to should-testing [24] instead of must-testing: according to these alternative definitions C and S turn out to be compliant due to the fairness assumption characterizing the should-testing approach.

Definition 5 (Petri net). *A Petri net is a pair $N = (S,T)$, where S is the set of* places *and $T \subseteq \mathcal{M}_{fin}(S) \times \mathcal{M}_{fin}(S)$ is the set of* transitions. *A transition (c,p) is written $c \Rightarrow p$. A transition $c \Rightarrow p$ is* enabled *at a marking m if $c \subseteq m$. The* execution of the transition *produces the marking $m' = (m \setminus c) \oplus p$ (where $\setminus$ and $\oplus$ are the multiset difference and union operators). This is written as $m[\rangle m'$. A* dead marking *is a marking in which no transition is enabled. A* marked *Petri net is a triple $N(m_0) = (S,T,m_0)$, where (S,T) is a Petri net and m_0 is the* initial marking. *A computation in $N(m_0)$ leading to the marking m is a sequence $m_0[\rangle m_1[\rangle m_2 \cdots m_n[\rangle m$.*

Note that in $c \Rightarrow p$, the marking c represents the tokens to be "consumed", while the marking p represents the tokens to be "produced". The Petri net semantics that we present for WSCL contracts decomposes contract terms into multisets of terms, that represents sequential contracts at different stages of invocation. We introduce the decomposition function in Definition 7. Instrumental to this definition is the set $\mathrm{Pl}(C)$, for C a WSCL contract, defined below.

Definition 6 ($\mathrm{Pl}(C)$). *For any contract C, let $C^{(k)}$ denote the term obtained by performing k unfolding of recursive definitions in C. Let k be the minimal s.t. in $C^{(k)}$ every $recX.D$ is guarded by one of the following prefixes: $\mathsf{invoke}(\cdot,\cdot), \mathsf{receive}(\cdot), \mathsf{reply}(\cdot,\cdot), a\langle b\rangle$. $\mathrm{Pl}(C)$ is defined as follows:*

$$\mathrm{Pl}(C) \triangleq \{\textstyle\sum_{i\in I} G_i,\ a\uparrow\sum_{i\in I} b_i.C_i,\ c\downarrow\sum_{i\in I} b_i.C_i \ :\ \sum_{i\in I} G_i, \sum_{i\in I} b_i.C_i \text{ occur in } C^{(k)}, a,c \in Names(C^{(k)})\}.$$

The function $dec(\cdot)$ transforms a WSCL contract C, as given in Definition 1, into a multiset $m \in \mathrm{Pl}(C)$.

Definition 7 (Decomposition). *The* decomposition *$dec(C)$ of a WSCL contract C, as given in Definition 1, is $dec_C(C)$. The auxiliary function $dec_C(D)$ is defined below by lexicographic induction on the pair (n_1,n_2), where n_1 is the number of unguarded (i.e. not under an $\mathsf{invoke}(\cdot,\cdot), \mathsf{receive}(\cdot), \mathsf{reply}(\cdot,\cdot), a\langle b\rangle$) sub-terms of the form $recX.D'$ in D and n_2 is the syntactic size of D.*

$$dec_C(\sum_{i\in I} r\langle b_i\rangle.D_i) = a\uparrow\sum_{i\in I} b_i.D_i \quad \text{if } \overline{a}\langle r\rangle \text{ occurs in } C$$

$$dec_C(\sum_{i\in I} r\langle b_i\rangle.D_i) = c\downarrow\sum_{i\in I} b_i.D_i \quad \text{if } \overline{r}\langle c\rangle \text{ occurs in } C \text{ and } c \neq b_i \text{ for every } i \in I$$

$$dec_C(recX.D) = dec_C(D\{^{recX.D}/X\}) \qquad dec_C(\overline{a}\langle b\rangle) = \emptyset$$

$$dec_C(D_1|D_2) = dec_C(D_1) \oplus dec_C(D_2) \qquad dec_C(D) = D, \text{ otherwise}$$

There are three kinds of transitions in the Petri net we are going to define: transitions representing the emission of an invocation, transitions representing (atomically) the reception of the invocation and the emission and reception of the reply, and transitions representing (atomically) the reception of the invocation and the emission of a reply that will never be received by the invoker because it is outside the set of admitted replies. These three cases are taken into account in the definition below.

Definition 8 (Petri net semantics). *Let C be a WSCL contract system as in Definition 1. We define $Net(C)$ as the Petri net (S,T) where:*

$\{\sum_{i\in I} G_i\} \Rightarrow \{a\uparrow\sum_{j\in J} b_j.C_j\}$	if $G_k = \mathsf{invoke}(a, \sum_{j\in J} b_j.C_j)$ for some $k \in I$
$\{a\uparrow\sum_{j\in J} b_j.C_j, \sum_{i\in I} G_i\} \Rightarrow dec(C_y) \oplus dec(D_z)$	if $\begin{cases} G_k = \mathsf{recreply}(a, \sum_{l\in L} c_l.D_l) \text{ for some } k \in I \text{ and} \\ b_y = c_z \text{ for some } y \in J,\ z \in L \end{cases}$
$\{a\uparrow\sum_{j\in J} b_j.C_j, \sum_{i\in I} G_i\} \Rightarrow \{c_z\downarrow\sum_{j\in J} b_j.C_j\} \oplus dec(D_z)$	if $\begin{cases} G_k = \mathsf{recreply}(a, \sum_{l\in L} c_l.D_l) \text{ for some } k \in I \text{ and} \\ \text{there exists } z \in L \text{ s.t. } c_z \neq b_j \text{ for every } j \in J \end{cases}$

Table1. Transitions schemata for the Petri net semantics of WSCL contracts.

- $S = \mathrm{Pl}(C)$;
- $T \subseteq \mathcal{M}_{fin}(S) \times \mathcal{M}_{fin}(S)$ *includes all the transitions that are instances of the transitions schemata in Table 1.*

We define the marked net $Net^{\mathrm{m}}(C)$ *as the marked net* (S, T, m_0)*, where the initial marking is* $m_0 = dec(C)$.

Example 2. Consider the client C and the service S introduced in Example 1 and let

$$C' \triangleq \mathsf{invoke}(requireTicket, ok.C) + \mathsf{recreply}(offerTicket, ok.\surd)$$
$$S' \triangleq \mathsf{recreply}(requireTicket, ok.(\mathsf{invoke}(offerTicket, ok)|S)).$$

The marked net $Net^{\mathrm{m}}(S|C)$ is depicted on the right. A bunch of unreachable places (like $ok\downarrow ok.\surd, ok\uparrow ok.\surd, \ldots$) have been omitted for the sake of clarity.

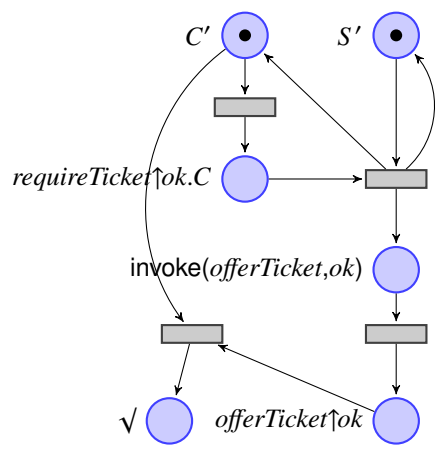

We divide the proof of the correspondence between the operational and the Petri net semantics of WSCL contracts in two parts: we first prove a *soundness* result showing that all Petri net computations reflect computations of contracts, and then a *completeness* result showing that contract computations leading to a state in which there are no uncompleted request-response interactions are reproduced in the Petri net.

In the proof of the soundness result we use the following structural congruence rule to remove empty contracts and in order to rearrange the order of contracts in parallel compositions. Let $\equiv$ be the minimal congruence for contract systems such as

$$C|\mathbf{0} \equiv C \qquad C|D \equiv D|C \qquad C|(D|E) \equiv (C|D)|E \qquad recX.C \equiv C\{^{recX.C}/_{X}\}$$

As usual, we have that the structural congruence respects the operational semantics.

Proposition 1. *Let C and D be two contract systems such that $C \equiv D$. If $C \xrightarrow{\alpha} C'$, then there exists D' such that $D \xrightarrow{\alpha} D'$ and $C' \equiv D'$.*

The following result establishes a precise relationship between the form of m and the form of C when $dec(C) = m$.

Lemma 1. *Let C be a* WSCL *contract system and suppose $dec(C) = m$. The following holds:*

1. *if $m = \{\sum_{i\in I} G_i\} \oplus m'$ then $C \equiv \sum_{i\in I} G_i \mid D$, for some D such that $dec(D) = m'$;*
2. *if $m = \{a\uparrow\sum_{j\in J} b_j.C_j\} \oplus m'$ then $C \equiv \sum_{j\in J} r\langle b_j\rangle.C_j \mid \overline{a}\langle r\rangle \mid D$, for some D and r such that $r \notin Names(D)$ and $dec(D) = m'$;*
3. *if $m = \{c\downarrow\sum_{j\in J} b_j.C_j\} \oplus m'$ then $c \neq b_j$ for each $j \in J$ and $C \equiv \sum_{j\in J} r\langle b_j\rangle.C_j \mid \overline{r}\langle c\rangle \mid D$, for some D and r such that $r \notin Names(D)$ and $dec(D) = m'$.*

Proposition 2. *Let C be a* WSCL *contract. Consider the Petri net $Net(C) = (S,T)$ and a marking m of $Net(C)$. We have that m is dead if and only if D has no outgoing reductions, for every D such that $dec(D) = m$.*

In order to prove that the Petri net semantics preserves client-service compliance, we need to introduce the notion of *success* marking. A *success* marking m contain at least one token in a place corresponding to a successful client state, formally, $m(\sum_{i\in I} G_i) > 0$ for some contract $\sum_{i\in I} G_i$ such that $G_k = \surd$, for some $k \in I$.

We are now ready to prove the *soundness* result.

Theorem 1 (Soundness). *Let C be a* WSCL *contract. Consider the Petri net $Net(C) = (S,T)$ and let m be a marking of $Net(C)$. If $m[\rangle m'$ then for each D such that $dec(D) = m$ there exists a computation $D \triangleq D_0 \longrightarrow D_1 \longrightarrow \cdots \longrightarrow D_l$, with $dec(D_l) = m'$. Moreover, if m is not a success marking then there exists no $j \in \{0,\cdots,l-1\}$ such that $D_j \xrightarrow{\surd}$.*

Proof: The proof proceeds by case analysis on the three possible kinds of transition.

1. If $m[\rangle m'$ by applying the first kind of transition then $m = \{\sum_{j\in J} G_j\} \oplus m''$, with $G_k = \mathsf{invoke}(a, \sum_{i\in I} b_i.C_i)$ for a $k \in J$. Moreover, $m' = \{a\uparrow\sum_{i\in I} b_i.C_i\} \oplus m''$.
   By Lemma 1, $D \equiv \sum_{j\in J} G_j \mid C' \longrightarrow \overline{a}\langle r\rangle \mid \sum_{i\in I} r\langle b_i\rangle.C_i \mid C' \triangleq D'$ and $dec(D') = m'$.
2. If $m[\rangle m'$ by applying the second kind of transition then $m = \{a\uparrow\sum_{j\in J} b_j.C_j, \sum_{i\in I} G_i\} \oplus m''$, with $G_k = \mathsf{recreply}(a, \sum_{l\in L} c_l.D_l)$, for some $k \in J$, and $b_y = c_z$ for some $y \in J$ and $z \in L$. By Lemma 1, if $dec(D) = m$ then $D \equiv \overline{a}\langle r\rangle \mid \sum_{j\in J} r\langle b_j\rangle.C_j \mid \sum_{i\in I} G_i \mid C'$, for any C' such that $dec(C') = m''$. Therefore, by $G_k = \mathsf{recreply}(a, \sum_{l\in L} c_l.D_l)$:

$$\begin{aligned} D &\longrightarrow \textstyle\sum_{j\in J} r\langle b_j\rangle.C_j \mid \sum_{l\in L} \mathsf{reply}(r,c_l).D_l \mid C' \\ &\longrightarrow \textstyle\sum_{j\in J} r\langle b_j\rangle.C_j \mid \overline{r}\langle c_z\rangle \mid D_z \mid C' \\ &\longrightarrow C_y \mid D_z \mid C' \quad \triangleq \; D' \end{aligned}$$

   with $dec(D') = dec(C_y) \oplus dec(D_z) \oplus m'' = m'$. Notice that each intermediate state in the reduction sequence from D to D' cannot perform a successful transition.

3. If $m[\rangle m'$ by applying the third kind of transition then $m = \{a\uparrow\sum_{j\in J} b_j.C_j, \sum_{i\in I} G_i\} \oplus m''$, with $G_k = \mathsf{recreply}(a, \sum_{l\in L} c_l.D_l)$, for some $k \in J$, and there is $z \in L$ such that $b_y \neq c_z$ for each $y \in J$. By Lemma 1, if $dec(D) = m$ then $D \equiv \overline{a}\langle r\rangle \mid \sum_{j\in J} r\langle b_j\rangle.C_j \mid \sum_{i\in I} G_i \mid C'$, for any C' such that $dec(C') = m''$. Therefore, by $G_k = \mathsf{recreply}(a, \sum_{l\in L} c_l.D_l)$, we get

$$\begin{aligned} D &\longrightarrow \textstyle\sum_{j\in J} r\langle b_j\rangle.C_j \mid \sum_{l\in L} \mathsf{reply}(r,c_l).D_l \mid C' \\ &\longrightarrow \textstyle\sum_{j\in J} r\langle b_j\rangle.C_j \mid \overline{r}\langle c_z\rangle \mid D_z \mid C' \triangleq D' \end{aligned}$$

with $dec(D') = \{c_z\downarrow\sum_{j\in J} b_j.C_j\} \oplus dec(D_z) \oplus m'' = m'$. Notice that each intermediate state in the reduction sequence from D to D' cannot perform a successful transition.

□

Definition 9 (Stable contracts). *A WSCL contract C (in the run-time syntax) is said* stable *if it contains neither unguarded* reply*(r,b) actions nor pairs of matching terms of the form $\overline{r}\langle b\rangle$ and $r\langle b\rangle$.*

Notice that any *initial* WSCL contract (according to the syntax of Definition 1) is stable.

Lemma 2. *Suppose C is stable and that $C \longrightarrow C'$. Then, there exists C'' stable such that $C' \longrightarrow C_1 \longrightarrow \cdots \longrightarrow C_l \longrightarrow C''$ ($l \geq 0$) and for each $i = 1, \cdots, l$ it holds that $C_i \not\xrightarrow{\surd}$.*

Proof: If C' is not stable then it may contain both unguarded reply actions and pairs of the form $\overline{r}\langle b\rangle$ and $r\langle b\rangle$. According to the operational semantics of contracts, all unguarded reply actions and $\overline{r}\langle b\rangle$ and $r\langle b\rangle$ can be consumed performing a sequence of reductions. Therefore a stable contract C'' can be reached from C' without traversing any state capable of $\xrightarrow{\surd}$. □

We now move to the completeness part.

Theorem 2 (Completeness). *Let C be a WSCL contract and let D be a contract reachable from C through the computation $C = C_0 \longrightarrow C_1 \longrightarrow \cdots \longrightarrow C_n = D$. If D is stable then there exists a computation $m_0[\rangle m_1[\rangle m_2 \cdots m_{l-1}[\rangle m_l$ of the marked Petri net $Net^{\mathrm{m}}(C)$ such that $dec(D) = m_l$. Moreover, if there exists no $k \in \{0, \cdots, n\}$ such that $C_k \xrightarrow{\surd}$ then for every $j \in \{0, \cdots, l\}$ we have that m_j is not a success marking.*

Proof: The proof is by induction on the length n of the derivation $C = C_0 \longrightarrow C_1 \longrightarrow \cdots \longrightarrow C_n = D$. The base case ($n = 0$) is trivial. In the inductive case there are two possible cases: C_{n-1} is stable or it is not stable. In the first case the proof is straightforward. In the second case, there are two possible scenarios to be considered: either C_n contains an unguarded action reply(r,b) term, or it contains a pair of matching terms $\overline{r}\langle b\rangle$ and $r\langle b\rangle$. We consider the first of these two cases, the second one can be treated similarly.

Let C_{n-1} be a non stable contract containing an unguarded action reply(r,b). This action cannot appear unguarded in the initial contract C: let C_j, with $j > 0$, be the first contract traversed during the computation of C in which the action reply(r,b) appears unguarded. Hence, we have that $C_{j-1} \longrightarrow C_j$ consists of the execution of a receive

action. We now consider a different computation from C to D obtained by rearranging the order of the steps in the considered computation $C = C_0 \longrightarrow C_1 \longrightarrow \cdots \longrightarrow C_n = D$. Namely, let $C = C_0 \longrightarrow C_1 \longrightarrow \cdots \longrightarrow C_{l-1} \longrightarrow C'_l \longrightarrow \cdots \longrightarrow C'_{n-2} \longrightarrow C_{n-1} \longrightarrow C_n = D$ be the computation obtained by delaying as much as possible the execution of the receive action generating the unguarded action $\mathsf{reply}(r,b)$. In the new computation, this action appears for the first time in the contract C_{n-1}. Moreover, C'_{n-2} must be a stable contract otherwise C_n is not stable. Hence, we can straightforwardly prove the thesis by applying the inductive hypothesis to the shorter computation $C = C_0 \longrightarrow C_1 \longrightarrow \cdots \longrightarrow C_{l-1} \longrightarrow C'_l \longrightarrow \cdots \longrightarrow C'_{n-2}$ leading to the stable contract C'_{n-2}. □

As a simple corollary of the last two theorems, we have that client-service compliance is preserved by the Petri net semantics.

Corollary 1 (Compliance preservation). *Let C and S be respectively a* WSCL *client and service contract, as in Definition 1. We have that C is compliant with S if and only if in the marked Petri net $Net^{\mathrm{m}}(C|S)$ all the maximal computations traverse at least one success marking.*

Proof: ($\Rightarrow$). Trivial by Theorem 1.

($\Leftarrow$). Suppose that in $Net^{\mathrm{m}}(C|S)$ all the maximal computations traverse at least one success marking and suppose by contradiction that C is not compliant with S. This means that there is a maximal computation from $C|S$ that does not traverse a state D such that $D \xrightarrow{\surd}$. This computation can either end in a state D' with no outgoing reductions or can be infinite.

In the first case we get a contradiction by Theorem 2. Indeed there would be a maximal computation from $Net^{\mathrm{m}}(C|S)$ traversing only non-success markings.

Consider the second case. From the infinite sequence of reductions, we can build an infinite set of maximal computations of arbitrary length, starting from C and ending in a stable state (Lemma 2) without traversing a succes state. By Theorem 2, for each of these maximal computations there exists a corresponding maximal computation in the net $Net^{\mathrm{m}}(S|C)$ that does not traverse a success marking. We can arrange these computations so as to form a tree where m' is a child of m iff $m[\rangle m'$: this is an infinite, but finite-branching, tree. By König's lemma, in $Net^{\mathrm{m}}(S|C)$ there exists then an infinite computation that does not traverse a success marking and we get a contradiction. □

3.2 Verifying client-service compliance using the Petri net semantics

In the light of Corollary 1, checking whether C is compliant with S reduces to verifying if all the maximal computations in $Net^{\mathrm{m}}(C|S)$ traverse at least one success marking. In order to verify this property, we proceed as follows:

- we first modify the net semantics in such a way that the net computations block if they reach a success markings;
- we define a (terminating) algorithm for checking whether in the modified Petri net all the maximal computations are finite and end in a success marking.

The modified Petri net semantics simply adds one place that initially contains one token. All transitions consume such a token, and reproduce it only if they do not introduce tokens in success places, i.e., places $\sum_{i \in I} G_i$ such that $G_k = \surd$ for some $k \in I$

Definition 10 (Modified Petri net semantics). *Let C be a* WSCL *contract and $Net(C) = (S,T)$ the corresponding Petri net as defined in Definition 8. We define $ModNet(C)$ as the Petri net (S',T') where:*

- *$S' = S \cup \{run\}$, where run is an additional place;*
- *for each transition $c \Rightarrow p \in T$, then T' contains a transition that consumes the multiset $c \uplus \{run\}$ and produces either p, if p contains a place $\sum_{i \in I} G_i$ such that $G_k = \surd$ for some $k \in I$, or $p \uplus \{run\}$, otherwise.*

The marked modified net *$ModNet^{\mathrm{m}}(C)$ is defined as the net $ModNet(C)$ with initial marking m_0 where*

$$m_0 = \begin{cases} dec(C) \uplus \{run\} & \text{if } dec(C) \text{ is not a success marking} \\ dec(C) & \text{otherwise.} \end{cases}$$

We now state an important relationship between $Net^{\mathrm{m}}(C)$ and $ModNet^{\mathrm{m}}(C)$. It can be proved by relying on the definition of modified net.

Proposition 3. *Let C be a* WSCL *contract, $Net^{\mathrm{m}}(C)$ (resp. $ModNet^{\mathrm{m}}(C)$) the corresponding Petri net (resp. modified Petri net). We have that all the maximal computations of $Net^{\mathrm{m}}(C)$ traverse at least one success marking if and only if in $ModNet^{\mathrm{m}}(C)$ all the maximal computations are finite and end in a success marking.*

We now present the algorithm for checking whether in a Petri net all the maximal computations are finite and end in a success marking. In the algorithm and in the proof, we utilize the following preorder over multisets on Places(C): $m \preceq m'$ iff for each p, $m(p) \leq m'(p)$. It can be shown that this preorder is a *well-quasi-order*, that is, in any infinite sequence of multisets there is a pair of multisets m and m' such that $m \preceq m'$ (see e.g. [16]).

Theorem 3. *Let C be a* WSCL *contract as in Definition 1 and let $ModNet^{\mathrm{m}}(C) = (S,T,m_0)$ be the corresponding modified Petri net. The algorithm described in Table 2 always terminates. Moreover, it returns TRUE iff all the maximal computations in $ModNet^{\mathrm{m}}(C)$ are finite and end in a success marking.*

4 Undecidability of client-service compliance for BPEL contracts

We now move to the proof that client-service compliance is undecidable for BPEL contracts. The proof is by reduction from the termination problem in Random Access Machines (RAMs) [21], a well known Turing powerful formalism based on registers containing nonnegative natural numbers. The registers are used by a program, that is a set of indexed instructions I_i which are of two possible kinds:

- $i : Inc(r_j)$ that increments the register r_j and then moves to the execution of the instruction with index $i+1$ and
- $i : DecJump(r_j, s)$ that attempts to decrement the register r_j; if the register does not hold 0 then the register is actually decremented and the next instruction is the one with index $i+1$, otherwise the next instruction is the one with index s.

1. If the initial marking m_0 is not a success marking then label it as the root and tag it "new".
2. While "new" markings exist do the following:
   (a) Select a "new" marking m.
   (b) If no transitions are enabled at m, return FALSE.
   (c) While there exist enabled transitions at m, do the following for each of them:
      i. Obtain a marking m' that results from firing the transition.
      ii. If on the path from the root to m there exists a marking m'' such that $m'(p) \geq m''(p)$ for each place p then return FALSE.
      iii. If m' is not a success marking introduce m' as a node, draw an arc from m to m', and tag m' "new".
   (d) Remove the tag "new" from the marking m.
3. Return TRUE.

Table2. An algorithm for checking the coverability of success markings.

Without loss of generality we assume that given a program $I_1, \cdots, I_n$, it starts by executing I_1 with all the registers empty (i.e. all registers contain 0) and terminates trying to perform the first undefined instruction I_{n+1}.

In order to simplify the notation, in this section we introduce a notation corresponding to standard input and output prefixes of CCS [19][3]. Namely, we model simple synchronization as a request-response interaction in which there is only one possible reply message. Assuming that this unique reply message is *ok* (with $ok \in Names$ not necessarily fresh), we introduce the following notation:

$$\overline{a}.P \;=\; \mathsf{invoke}(a,ok.P) \qquad\qquad a.P \;=\; \mathsf{receive}(a).\mathsf{reply}(a,ok).P$$

In order to reduce RAM termination to client-service compliance, we define a client contract that simulates the execution of a RAM program, and a service contract that represent the registers, such that the client contract reaches the success $\surd$ if and only if the RAM program terminates.

Given a RAM program $I_1, \cdots, I_n$, we consider the client contract C as follows

$$\prod_{i \in \{1,\cdots,n\}} [\![I_i]\!] \mid inst_{n+1}.\surd$$

$$[\![I_i]\!] = \begin{cases} recX.(inst_i.\overline{inc_j}.ack.(\overline{inst_{i+1}} \mid X)) & \text{if } I_i = (i : Inc(r_j)) \\ recX.(inst_i.\overline{dec_j}.(ack.(\overline{inst_{i+1}} \mid X) + zero.(\overline{inst_s} \mid X))) & \text{if } I_i = (i : DecJump(r_j, s)) \end{cases}$$

An increment instruction $Inc(r_j)$ is modeled by a recursive contract that invokes the operation inc_j, waits for an acknowledgement on *ack*, and then invokes the service corresponding to the subsequent instruction. On the contrary, a decrement instruction $DecJump(r_j, s)$ invokes the operation dec_j and then waits on two possible operations: *ack* or *zero*. In the first case the service corresponding to the subsequent instruction

[3] The input and output prefixes correspond also to the representation of the one-way interaction pattern in contract languages such as those in [10,6,11].

with index $i+1$ is invoked, while in the second case the service corresponding to the target of the jump is invoked instead.

We now move to the modeling of the registers. Each register r_j is represented by a contract representing the initially empty register in parallel with a service responsible to model every unit subsequently added to the register

$$\begin{aligned}[\![r_j]\!] = {} & recX.(dec_j.\overline{zero}.X + inc_j.\overline{unit_j}.\mathsf{invoke}(u_j,ok.\overline{ack}.X)) \mid \\ & recX.unit_j.(X \mid \mathsf{receive}(u_j).\overline{ack}.recY.(dec_j.\mathsf{reply}(u_j,ok) + \\ & \qquad inc_j.\overline{unit_j}.\mathsf{invoke}(u_j,ok.\overline{ack}.Y))).\end{aligned}$$

The idea of the encoding is to model numbers with chains of nested request-response interactions. When a register is incremented, a new instance of a contract is spawn invoking the operation $unit_j$, and a request-response interaction is opened between the previous instance and the new one. In this way, the previous instance blocks waiting for the reply. When an active instance receives a request for decrement, it terminates by closing the request-response interaction with its previous instance, which is then reactivated. The contract that is initially active represents the empty register because it replies to decrement requests by performing an invocation on the *zero* operation.

We extend structural congruence $\equiv$, introduced in Section 3, to $\equiv_{ren}$ to admit the injective renaming of the operation name

$$C \equiv_{ren} D \text{ if there exists an injective renaming } \sigma \text{ such that } C\sigma \equiv D$$

Clearly, injective renaming preserves the operational semantics.

Proposition 4. *Let C and D be two contract systems such that $C \equiv_{ren} D$. If $C \xrightarrow{\alpha} C'$, then there exists D' and a label α' obtained by renaming the operation names in α such that $D \xrightarrow{\alpha'} D'$ and $C' \equiv_{ren} D'$.*

Now, we introduce $\{\!\{r_j, c\}\!\}$ that we use to denote the modeling of the register r_j when it holds the value c. Namely, $\{\!\{r_j, 0\}\!\} = [\![r_j]\!]$, while if $c > 0$ then

$$\{\!\{r_j, c\}\!\} = \begin{cases} b_0\langle ok\rangle.\overline{ack}.recX.(dec_j.\overline{zero}.X + inc_j.\overline{unit_j}.\mathsf{invoke}(u_j,ok.\overline{ack}.X)) \mid \\ b_1\langle ok\rangle.\overline{ack}.recY.(dec_j.\overline{b_0}\langle ok\rangle + inc_j.\overline{unit_j}.\mathsf{invoke}(u_j,ok.\overline{ack}.Y)) \mid \\ \cdots \mid \\ recY.(dec_j.\overline{b_{c-1}}\langle ok\rangle + inc_j.\overline{unit_j}.\mathsf{invoke}(u_j,ok.\overline{ack}.Y)) \mid \\ recX.unit_j.(X \mid \mathsf{receive}(u_j).\overline{ack}.recY.(dec_j.\mathsf{reply}(u_j,ok) + \\ \qquad inc_j.\overline{unit_j}.\mathsf{invoke}(u_j,ok.\overline{ack}.Y))) \end{cases}$$

In the following theorem, stating the correctness of our encoding, we use the following notation: $(i, c_1, \cdots, c_m)$ to denote the state of a RAM in which the next instruction to be executed is I_i and the registers $r_1, \cdots, r_m$ respectively contain the values $c_1, \cdots, c_m$, and $(i, c_1, \cdots, c_m) \rightarrow_R (i', c'_1, \cdots, c'_m)$ to denote the change of the state of the RAM R due to the execution of the instruction I_i.

Theorem 4. *Consider a RAM R with instructions $I_1, \cdots, I_n$ and registers $r_1, \cdots, r_m$. Consider also a state $(i, c_1, \cdots, c_m)$ of the RAM R and a corresponding contract C such that $C \equiv_{ren} \overline{inst_i} | [\![I_1]\!] | \cdots | [\![I_n]\!] | inst_{n+1}.\surd | \{\!\{r_1, c_1\}\!\} | \cdots | \{\!\{r_m, c_m\}\!\}$. We have that*

– *either the RAM computation has terminated, thus* $i = n+1$
– *or* $(i, c_1, \cdots, c_m) \rightarrow_R (i', c'_1, \cdots, c'_m)$ *and there exists* $l > 0$ *such that* $C \longrightarrow C_1 \longrightarrow \cdots \longrightarrow C_l$ *and*
  - $C_l \equiv_{ren} \overline{inst_{i'}} | [\![I_1]\!] | \cdots | [\![I_n]\!] | inst_{n+1}.\surd | \{\!\{r_1, c'_1\}\!\} | \cdots | \{\!\{r_m, c'_m\}\!\}$
  - *for each* k $(1 \leq k < l)$: $C_k \overset{\surd}{\not\rightarrow}$

As a corollary we get that client-service compliance is undecidable.

Corollary 2. *Consider a RAM* R *with instructions* $I_1, \cdots, I_n$ *and registers* $r_1, \cdots, r_m$. *Consider the client contract* $C = \overline{inst_1} | [\![I_1]\!] | \cdots | [\![I_n]\!] | inst_{n+1}.\surd$ *and the service contract* $S = \{\!\{r_1, 0\}\!\} | \cdots | \{\!\{r_m, 0\}\!\}$. *We have that* C *is compliant with* S *if and only if* R *terminates.*

5 Conclusion

We have presented two models of contracts with bidirectional request-response interaction, studied a notion of compliance based on must testing and established an expressiveness gap between the two models showing that compliance is decidable in the first one while it is undecidable in the second one.

As for future work, we plan to investigate the (un)decidability of other definitions of compliance present in the literature. In fact, the must-testing approach —the one that we consider in this paper— has been adopted in early works about service compliance (see e.g. [10]). More recent papers consider more sophisticated notions. For instance, the should-testing approach [24] adopted, e.g., in [9] admits also infinite computations if in every reached state there is always at least one path leading to a success state. Other approaches require the successful completion of all the services in the system (see e.g. [8,12]) in order to deal with multiparty service compositions in which there is no distinction between a client and a service.

Moreover, it would be interesting to apply the techniques presented in this paper to more sophisticated orchestration languages, like the recently proposed calculi based on the notion of *session* [5,4]. For instance, in [2], a type system is presented ensuring a client progress property – basically, absence of deadlock – in a calculus where interaction between (instances of) the client and the service is tightly controlled via session channels. It would be interesting to check to what extent the decidability techniques presented here apply to this notion of progress. Also connections with *behavioural types* [17,1] deserve attention. In the setting of process calculi, these types are meant to provide behavioural abstractions that are in general more tractable than the original process. In the present paper, the translation function of WSCL contracts into Petri nets can be seen too as a form of behavioural abstraction. In the case of tightly controlled interactions (sessions) [2], BPP processes, a proper subset of Petri nets featuring no synchronization [14], have been seen to be sufficient as abstractions. For general pi-processes, full CCS with restriction is in general needed. One would like to undertake a systematic study of how communication capabilities in the original language (unconstrained interaction vs. sessions vs. request-response vs....) trades off with tractability of the behavioural abstractions (CCS vs. BPP vs. Petri nets vs. ...).

References

1. Acciai, L., Boreale, M.: Spatial and behavioural Types in the pi-calculus. In Proc. of CONCUR'08, LNCS 5201:372-386 (2008). Full version to appear in *Inf. and Comp.*
2. Acciai, L., Boreale, M.: A Type System for Client Progress in a Service-Oriented Calculus. In Concurrency, Graphs and Models, Essays Dedicated to Ugo Montanari on the Occasion of His 65th Birthday. LNCS 5065:642–658 (2008)
3. Boreale, M., Bravetti, M.: Advanced mechanisms for service composition, query and discovery. LNCS (2010). To appear.
4. Boreale, M., Bruni, R., De Nicola, R., Loreti, M.: Sessions and Pipelines for Structured Service Programming. In Proc. of FMOODS'08, LNCS 5051:19–38 (2008)
5. Boreale, M., Bruni, R., Caires, L., De Nicola, R., Lanese, I., Loreti, M., Martins, F., Montanari, U., Ravara, A., Sangiorgi, D., Vasconcelos, V.T., Zavattaro, G.: SCC: A Service Centered Calculus. In Proc. of WS-FM'06, LNCS 4184:38–57 (2006)
6. Bravetti, M., Zavattaro, G.: Contract based Multi-party Service Composition, In Proc. of FSEN'07, LNCS 4767207–222 (2007)
7. Bravetti, M., Zavattaro, G.: A Theory for Strong Service Compliance, In Proc. of Coordination'07, LNCS 4467:96–112 (2007)
8. Bravetti, M., Zavattaro, G.: Towards a Unifying Theory for Choreography Conformance and Contract Compliance, In Proc. of SC'07, LNCS 4829:34–50 (2007)
9. Bravetti, M. and Zavattaro, G.: Contract-Based Discovery and Composition of Web Services. In Proc. of SFM'09. LNCS 5569: 261–295 (2009)
10. Carpineti, S., Castagna, G., Laneve, C., Padovani, L.: A Formal Account of Contracts for Web Services, In Proc. of WS-FM'06, LNCS 4184:148–162 (2006)
11. Castagna, G., Gesbert, N., Padovani, L.: A Theory of Contracts for Web Services, In Proc. of POPL'08, ACM Press 261–272 (2008)
12. Castagna, G. and Padovani, L.: Contracts for Mobile Processes, In Proc. of Concur'09, LNCS 5710:211–228 (2009)
13. De Nicola, R., Hennessy, M.: Testing equivalences for processes. Theor. Comput. Sci, 34:83–133 (1984)
14. Esparza, J.: Petri Nets, Commutative Context-Free Grammars, and Basic Parallel Processes. Fundam. Inform. 31(1):13–25 (1997)
15. Esparza, J., Nielsen, M.: Decidability Issues for Petri Nets - a survey. Bulletin of the EATCS 52:244–262 (1994)
16. Finkel, A., Schnoebelen, Ph.: Well-Structured Transition Systems Everywhere! Theor. Comput. Sci., 256(1-2): 63–92 (2001)
17. Igarashi, A., Kobayashi, N.: A generic type system for the Pi-calculus. Theor. Comput. Sci. 311(1-3), 121–163 (2004)
18. Karp , R.M., Miller, R.E.: Parallel Program Schemata., Journal of Computer and System Sciences 3:147–195 (1969)
19. Milner, R.: Communication and concurrency. Prentice-Hall (1989)
20. Milner, R., Parrow, J., Walker, D.: A calculus of mobile processes. Information and Computation, volume 100, pages 1–40 (1992)
21. Minsky, M.L.: Computation: finite and infinite machines. Prentice-Hall, Englewood Cliffs (1967)
22. OASIS: Web Services Business Process Execution Language (WSBPEL). Standard available at: `www.oasis-open.org/committees/wsbpel` (2007)
23. Petri, C.A.: *Kommunikation mit Automaten*. Ph. D. Thesis. University of Bonn (1962)
24. Rensink A., Vogler W.: Fair testing. Inf. Comput. 205(2), 125-198 (2007)
25. W3C: Web Services Conversation Language (WSCL). Standard proposal available at: `http://www.w3.org/TR/wscl10` (2002)