

Manuscript Number: SCICO-D-11-00251

Title: Model Checking Adaptive Multilevel Service Compositions

Article Type: Special Issue: FACS 2010

Keywords: processalgebra; non-interference; model-checking; webservices

Corresponding Author: Dr Sabina Rossi,

Corresponding Author's Institution:

First Author: Michele Bugliesi

Order of Authors: Michele Bugliesi; Andrea Marin; Sabina Rossi

Abstract: In this paper we present a logic-based technique for verifying both security and correctness properties of multilevel service compositions. We define modal μ -calculus formulae interpreted over service configurations. Our formulae characterize those compositions which satisfy a non-interference property and are compliant, i.e., are both deadlock and livelock free. Moreover, we use filters as prescriptions of behavior (coercions to prevent service misbehavior) and we devise a model checking algorithm for adaptive service compositions which automatically synthesizes an adapting filter.

Highlights

- > We present a technique for verifying security and correctness of multilevel service compositions.
- > We define modal μ -calculus formulae interpreted over service configurations.
- > We use filters as prescriptions of behavior, i.e., coercions to prevent service misbehavior.
- > We devise a model checking algorithm for adaptive service compositions.
- > We study a case-study of a federated login system.

Model Checking Adaptive Multilevel Service Compositions[☆]

M. Bugliesi^a, A. Marin^a, S. Rossi^{a,1,*}

^a*Università Ca' Foscari Venezia, via Torino 155, 30172
Venezia (Italy)*

Abstract

In this paper we present a logic-based technique for verifying both security and correctness properties of multilevel service compositions. We define modal μ -calculus formulae interpreted over service configurations. Our formulae characterize those compositions which satisfy a non-interference property and are compliant, i.e., are both deadlock and livelock free. Moreover, we use filters as prescriptions of behavior (coercions to prevent service misbehavior) and we devise a model checking algorithm for adaptive service compositions which automatically synthesizes an adapting filter.

Keywords: process algebra, non-interference, model-checking, web services

1. Introduction

Service Oriented Architectures (SOA) provide a software architectural style to connect loosely specified and coupled services that communicate with each other. Simple Object Access Protocol (SOAP)-based Web services are becoming the most common implementation of SOA. They are designed to support interoperable service-to-service interactions over a network. This interoperability is gained through a set of XML-based open standards, such as the *Web Service Definition Language* (WSDL). These standards provide a common approach for defining, publishing, and using web services.

[☆]Work partially supported by the MIUR Project IPODS “Interacting Processes in Open-ended Distributed Systems”.

*Corresponding author.

Email addresses: bugliesi@dais.unive.it (M. Bugliesi), marin@dais.unive.it (A. Marin), srossi@dais.unive.it (S. Rossi)

¹Full Contact details: Dipartimento di Scienze Ambientali, Informatica e Statistica (DAIS), via Torino 155 - 30172 Venezia (Italy), Tel. +39-041-2348411 Fax. +39-041-2348419

1
2
3
4
5
6
7
8
9
10 As is the case of many other applications, the information processed in web ser-
11 vices might be commercially sensitive and it is important to protect this information
12 against security threats such as disclosure to unauthorized parties. Indeed, many
13 of the features that make web services attractive, including greater accessibility
14 of data, dynamic connections, and relative autonomy (lack of human intervention)
15 are at odds with traditional security models and controls. Difficult issues and un-
16 solved problems exist, such as protecting confidentiality and integrity of data that
17 is transmitted through web services protocols in service-to-service transactions.
18

19 Traditionally, security policies are used to specify the security requirements of
20 a system and to control the access to confidential data and resources (see, e.g., [1,
21 2]). However, if this approach can be successfully applied to enforce the security
22 requirements of a centralized system, it is less suited to formulate the security needs
23 of a service composition. Indeed, there is no central authority in the web that is able
24 to fix the security labels of all services and data. Rather than manipulating stored
25 data, web services compute requested information from dynamic data available on
26 the net that need to be dynamically classified according to their stored information.
27

28 In this paper we specify service compositions in terms of behavioral contracts
29 which provide abstract descriptions of system behaviors by means of terms of some
30 process algebra. Formal theories of contracts have first been introduced in [3], and
31 then further developed along independent lines of research in [4, 5, 6], and in [7, 8].
32

33 Based upon such formal descriptions, we propose a notion of *non-interference*
34 [9] for multilevel service compositions. This is motivated by the fact that SOAs are
35 increasingly relying on complex distributed systems that share information with
36 multiple levels of security. In these systems information with mixed security levels
37 is processed and targeted to particular clients. For example, in an e-business sys-
38 tem, some data will be privileged (e.g., credit card numbers and medical records)
39 and some data will be public (e.g., stock market quotes). Such systems need to be
40 equipped with appropriate security facilities to guarantee the security requirements
41 (e.g., confidentiality or integrity) of the participants.
42

43 In this paper, we present an information flow security model [10, 11, 12, 13] for
44 service compositions to control the flow of confidential data in web services. Secu-
45 rity policies are used to specify the security requirements of service components.
46 In order to capture the dynamic nature, heterogeneity and lack of knowledge which
47 are intrinsic features of modern web services, we allow policies to be dynamically
48 specified by the service participants. In our model, for example, customers may
49 formulate their security requirements by dynamically assign types (that are secu-
50 rity annotations) to individual service components.
51

52 We also consider the property of *compliance* which is widely used in the con-
53 text of SOA as a formal device to identify well-formed service compositions, those
54 whose interactions are free of synchronization errors. Our notion of compliance
55
56
57
58

1
2
3
4
5
6
7
8
9
10 (see [7, 8]) is a strong condition that ensures the absence of deadlocks *and* live-
11 locks, and the application setting is that of choreographies: compliant choreogra-
12 phies are those whose computations never get stuck or trapped into infinite loops
13 without chances to exit.

14 We develop a method for verifying both security and correctness properties of
15 multilevel service compositions based on the use of model-checking techniques [14].

16 Model checking is a promising formal method for automatically verifying prop-
17 erties of finite-state concurrent systems [14]. A model checker normally performs
18 an exhaustive search of the state space of a system to determine if a particular prop-
19 erty, expressed in terms of a logical formula, holds. Given sufficient resources, the
20 procedure will always terminate with a yes/no answer.

21 We define modal μ -calculus [15] formulae, interpreted over service configura-
22 tions, characterizing those compositions which satisfy a non-interference property
23 and are compliant, i.e., are both deadlock and livelock free. Indeed, due to its
24 expressiveness and its conciseness the μ -calculus is well-suited to model both the
25 security and the correctness properties we are interested in. A model checker (like,
26 e.g., NCSU Concurrency Workbench) can then be used to simultaneously check
27 non-interference and compliance.

28 We also develop an algorithm for verifying adaptive service compositions. This
29 is based on the use of filters, introduced in [4], as prescriptions of behaviour (co-
30 ercions to prevent service misbehaviour). Security and correctness properties for
31 adaptive multilevel service compositions are ensured by the automatic synthesis of
32 an adapting filter.

33 Finally, we apply these results to a case-study in which a customer receives a
34 security level according to its identity that may be certified by a federated login
35 system. This consists of a set of providers that share an open standard to handle
36 user accounts, such as *OpenId*. Google, facebook, Yahoo!, Wordpress are examples
37 of providers that adhere to OpenId standard. When a web application allows for
38 a federated login, a customer can use the account owned at any of the federated
39 providers to obtain the required service. We show how the calculus presented here
40 allows one to formally describe such a system, and we apply the compliance and
41 non-interference analysis to the model. Particular attention is devoted to interpret
42 the results of such analysis in terms of good developing practices.

43 1.1. Plan of the paper

44 The paper is organized as follows: Section 2 introduces the calculus for mul-
45 tilevel service compositions. Section 3 formalizes the notion of non-interference.
46 Section 4 presents characteristic modal μ -calculus formulae for non-interference.
47 Section 5 introduces the notion of compliance and defines a modal μ -calculus for-
48 mula characterizing it. Section 6 presents an algorithm for adaptive service com-
49

<i>Type environments</i>	$\Delta ::= \emptyset \mid \Delta, u : \varsigma$	$u \in \mathcal{P} \cup \mathcal{V}, \varsigma \in \Sigma$
<i>Actions</i>	$\varphi ::= \bar{a} @ u \mid a @ u$	$a \in \mathcal{A}, u \in \mathcal{P} \cup \mathcal{V}$
<i>Contracts</i>	$\sigma ::= \mathbf{1} \mid \mathbf{x} \mid \varphi . \sigma \mid \sigma + \sigma$ $\mid \sigma [\Delta \oplus \Delta] \sigma \mid \text{rec}(\mathbf{x}) \sigma$	
<i>Compositions</i>	$C ::= p[\sigma] \mid C \parallel C$	

Table 1: Syntax

positions. A case study is presented in Section 7. Finally, Section 8 concludes the paper.

2. The Calculus

We represent service contracts as terms of a value-passing CCS-like [16] process calculus that includes recursion and operators for external and internal choice. In the algebra, parallel composition arises in *contract compositions* that we define after [8, 17, 18] as the parallel (and concurrent) composition of a set of principals executing contracts. We presuppose a denumerable set of action names \mathcal{A} , ranged over by a, b, c , a denumerable set of principal identities \mathcal{P} , ranged over by p, q, r , and a denumerable set of variables \mathcal{V} , ranged over by x, y . The actions represent the basic units of observable behavior of the underlying services, while the principal names specify the peers providing the services.

In order to specify multilevel service compositions, we assign security levels to principal identities and express both contracts and compositions as typed terms of our calculus. Formally, we assume a complete lattice $\langle \Sigma, \preceq \rangle$ of security annotations, ranged over by ς, ϱ , where \top and \perp represent the top and the bottom elements of the lattice. We denote by \sqcup and \sqcap the join and meet operators over Σ , respectively. Type environments are used to assign security levels to principals. A type environment Δ is a finite mapping from principals and variables to security annotations. Closed type environments, ranged over by Γ , assign security levels to principals only, i.e., they do not contain variables in their domain. We define $\Gamma_1 \sqcup \Gamma_2$ (resp., $\Gamma_1 \sqcap \Gamma_2$) the type environment Γ such that $\Gamma(p) = \Gamma_i(p)$ if $p \notin \text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2)$ and $p \in \text{dom}(\Gamma_i)$, and $\Gamma(p) = \Gamma_1(p) \sqcup \Gamma_2(p)$ (resp., $\Gamma_1(p) \sqcap \Gamma_2(p)$) otherwise.

2.1. Syntax

The syntax of our calculus is presented in Table 1. Contracts have the following form. Term $\mathbf{1}$ denotes a contract that has reached a successful state. The contract

$\bar{a}@p.\sigma$ describes a service that sends a message on a to principal p and then behaves as σ ; syntactically, the principal identity p may be a variable, but it must be a name when the prefix is ready to fire. Dually, the input prefix $a@u.\sigma$ waits for an input on a from a particular/any principal and then continues as σ . If u is a variable x , then the input form is a binder for x with scope σ : upon synchronization with a principal p , x gets uniformly substituted by p in σ . The contract $\sigma + \sigma'$ denotes an external choice, guided by the environment. The contract $\sigma[\Delta \oplus \Delta']\sigma'$ represents the internal choice between σ in the type environment Δ and σ' in the type environment Δ' made irrespective of the structure of the interacting components. Syntactically, Δ and Δ' may contain variables in their domain, but they must be closed when the internal choice is ready to fire. The internal choice operator we adopt in this paper allows us to model the fact that a principal may dynamically change (upgrade) the security level of his interactions with other service components through the specific type environment associated with each choice. This will be explained in the definition of the semantics of our calculus and Example 2.1. Finally, $\text{rec}(\mathbf{x})\sigma$ makes it possible to express iteration in the contract language. As usual, we assume a standard contractivity condition for recursion, requiring that recursive variables be guarded by a prefix.

Given a principal $p \in \mathcal{P}$, we say that a contract σ is *p-compatible* if for all $\bar{a}@q$ and $a@q$ occurring in σ , q is different from p .

A composition $p_1[\sigma_1] \parallel \cdots \parallel p_n[\sigma_n]$ of principals must be *well-formed* [17] to constitute a legal composition, namely: (i) the principal identities p_i 's must all be pairwise different, and (ii) each contract σ_i , executed by principal p_i , is p_i -compatible. If $C = p_1[\sigma_1] \parallel \cdots \parallel p_n[\sigma_n]$ is a legal composition, we say that C is a $\{p_1, \dots, p_n\}$ -composition (dually, that $\{p_1, \dots, p_n\}$ are the underlying principals for C). Throughout, we assume that contracts are closed (they have no free variables) and that compositions are well-formed. Also, we often omit trailing $\mathbf{1}$'s. We say that $\Gamma \triangleright C$ is a *configuration* if Γ is a type environment and C is a well-formed $\{p_1, \dots, p_n\}$ -service composition such that $\{p_1, \dots, p_n\} \subseteq \text{dom}(\Gamma)$.

2.2. Semantics

We define the dynamics of typed service compositions in terms of labelled transition systems (and a success predicate), with rules reported in Table 2. In the table, and in the whole paper, λ ranges over visible contract typed actions $\bar{a}@p$, $a@p$ and silent actions τ ; δ ranges over service composition actions $a_{p \rightarrow q}$, $\bar{a}_{p \rightarrow q}$ and τ .

The first block of rules defines the successful states of a contract and a composition, which are those that expose the successful term $\mathbf{1}$ at top level, or immediately under an external choice (up-to recursive unfoldings). Notice that a composition is successful only when all its components are successful.

Contract and composition satisfaction: $\sigma \checkmark$

$$\mathbf{1} \checkmark \quad \frac{\sigma_i \checkmark}{\sigma_1 + \sigma_2 \checkmark} \quad \frac{\sigma\{\mathbf{x} := \text{rec}(\mathbf{x})\sigma\} \checkmark}{\text{rec}(\mathbf{x})\sigma \checkmark} \quad \frac{\sigma \checkmark}{p[\sigma] \checkmark} \quad \frac{C_1 \checkmark \quad C_2 \checkmark}{C_1 \parallel C_2 \checkmark}$$

Typed contract transitions: $\Gamma \triangleright \sigma \xrightarrow{\lambda} \Gamma' \triangleright \sigma'$

$$\Gamma \triangleright a@p.\sigma \xrightarrow{a@p} \Gamma \triangleright \sigma \quad \Gamma \triangleright a@x.\sigma \xrightarrow{a@p} \Gamma \triangleright \sigma\{x := p\}$$

$$\Gamma \triangleright \bar{a}@p.\sigma \xrightarrow{\bar{a}@p} \Gamma \triangleright \sigma \quad \Gamma \triangleright \sigma_1[\Gamma_1 \oplus \Gamma_2]\sigma_2 \xrightarrow{\tau} \Gamma \sqcup \Gamma_i \triangleright \sigma_i \quad (i = 1, 2)$$

$$\frac{\Gamma \triangleright \sigma_i \xrightarrow{\lambda} \Gamma' \triangleright \sigma}{\Gamma \triangleright \sigma_1 + \sigma_2 \xrightarrow{\lambda} \Gamma' \triangleright \sigma} \quad (i = 1, 2) \quad \frac{\Gamma \triangleright \sigma\{\mathbf{x} := \text{rec}(\mathbf{x})\sigma\} \xrightarrow{\lambda} \Gamma' \triangleright \sigma'}{\Gamma \triangleright \text{rec}(\mathbf{x})\sigma \xrightarrow{\lambda} \Gamma' \triangleright \sigma'}$$

Typed composition transitions: $\Gamma \triangleright C \xrightarrow{\delta} \Gamma' \triangleright C'$

$$\frac{\Gamma \triangleright \sigma \xrightarrow{a@p} \Gamma \triangleright \sigma'}{\Gamma \triangleright q[\sigma] \xrightarrow{a_{p \rightarrow q}} \Gamma \triangleright q[\sigma']} \quad p \in \text{dom}(\Gamma), p \neq q \quad \frac{\Gamma \triangleright \sigma \xrightarrow{\bar{a}@p} \Gamma \triangleright \sigma'}{\Gamma \triangleright q[\sigma] \xrightarrow{\bar{a}_{q \rightarrow p}} \Gamma \triangleright q[\sigma']} \quad p \neq q$$

$$\frac{\Gamma \triangleright C_1 \xrightarrow{a_{p \rightarrow q}} \Gamma \triangleright C'_1 \quad \Gamma \triangleright C_2 \xrightarrow{\bar{a}_{p \rightarrow q}} \Gamma \triangleright C'_2}{\Gamma \triangleright C_1 \parallel C_2 \xrightarrow{\tau} \Gamma \triangleright C'_1 \parallel C'_2}$$

$$\frac{\Gamma \triangleright \sigma \xrightarrow{\tau} \Gamma' \triangleright \sigma'}{\Gamma \triangleright p[\sigma] \xrightarrow{\tau} \Gamma' \triangleright p[\sigma']} \quad \Gamma(q) \prec \Gamma'(q) \preceq \Gamma(p), \forall q \in \text{dom}(\Gamma') : \Gamma(q) \neq \Gamma'(q)$$

$$\frac{\Gamma \triangleright C_1 \xrightarrow{\delta} \Gamma' \triangleright C'_1}{\Gamma \triangleright C_1 \parallel C_2 \xrightarrow{\delta} \Gamma' \triangleright C'_1 \parallel C_2}$$

Table 2: Typed contract and composition transitions

1
2
3
4
5
6
7
8
9
10 The second block of rules defines the typed transitions for contracts, and are
11 mostly self-explanatory. The rule for the internal choice ensures that a service com-
12 ponent cannot downgrade the security level of other principals. Intuitively, this cor-
13 responds to the fact that downgrading the security level of a principal would mean
14 downgrading the level of confidentiality of the information it manages. Notice that
15 the type environments (Γ_1 and Γ_2) that are arguments of the internal choice are
16 required to be closed when the choice is ready to fire.

17
18 Each typed contract transition yields a corresponding transition for the princi-
19 pal hosting the contract. The typed transitions for configurations are relative to the
20 underlying set $dom(\Gamma)$ of principals. τ transitions for configurations arise from
21 the execution of a contract internal choice. In this case, the principal performing
22 the choice may modify the security level of its interactions with other components
23 by assigning different security levels to the principals with which it is going to
24 interact. In the rule we assume that a principal p cannot upgrade the level of its
25 interactions with other components above its own level. This is the meaning of
26 the condition associated with the rule for τ -typed composition transitions. This is
27 the only constraint on the principal security levels we assume. Such a constraint
28 ensures that information does not explicitly flow from high to low, but it does not
29 deal with implicit flows. Instead, we will characterize non-interference in terms of
30 the actions that typed service compositions may perform.

31
32 We will use the following shorthands. We write \Longrightarrow to denote the reflexive
33 and transitive closure of $\xrightarrow{\tau}$, and $\xRightarrow{\delta}$ for $\Longrightarrow \xrightarrow{\delta} \Longrightarrow$. We extend sequence of
34 actions $w = \delta_1 \dots \delta_n$, we write \xRightarrow{w} to note $\xRightarrow{\delta_1} \dots \xRightarrow{\delta_n}$. A computation for a
35 configuration $\Gamma \triangleright C$, is a sequence $\Gamma \triangleright C = \Gamma_0 \triangleright C_0 \xrightarrow{\tau} \Gamma_1 \triangleright C_1 \xrightarrow{\tau} \dots$ of
36 internal actions.

37
38 **Lemma 2.1** (Preserving well-formedness). *Let Γ be a type environment and C be*
39 *a service composition such that $\Gamma \triangleright C$ is a configuration. If $\Gamma \triangleright C$ is well-formed*
40 *and $\Gamma \triangleright C \xrightarrow{\tau} \Gamma' \triangleright C'$, then $\Gamma' \triangleright C'$ is well-formed.*

41
42 *Proof.* The only subtlety is that an output action within principal p , may have a
43 variable, say x as a target. In any closed contract, x must be bound at an enclos-
44 ing input prefix. Now, given that $\Gamma \triangleright C$ is well-formed, that input may never
45 synchronize with an output from p itself: hence x will never get bound to p . \square

46
47 **Example 2.1.** *Table 3 shows an example of a service contract composition. Let*
48 *Σ contain two security annotations, L (public) and H (confidential), with $L \preceq$*
49 *H . Let Γ be the type environment $C : H, T : L, A_1 : L, A_2 : L$. The typed*
50 *composition $\Gamma \triangleright S$ is well-formed and consists of four services: $C[\sigma_C]$, $T[\sigma_T]$*
51 *and $A_i[\sigma_A]$ representing a customer, a travel agency, and two airline companies,*

$$S = C[\sigma_C] \parallel T[\sigma_T] \parallel A_1[\sigma_A] \parallel A_2[\sigma_A]$$

$$\sigma_C = \overline{\text{Req}}@T.\overline{\text{Lst}}@T.(\overline{\text{Close}}@T.\mathbf{1}[_{\emptyset \oplus T:\mathbb{H}}](\overline{\text{Buy1}}@T.\overline{\text{Pay}}@T.\overline{\text{Get}}@A_1.\mathbf{1}[_{A_1:\mathbb{H} \oplus A_2:\mathbb{H}}]\overline{\text{Buy2}}@T.\overline{\text{Pay}}@T.\overline{\text{Get}}@A_2.\mathbf{1}))$$

$$\sigma_T = \text{Req}@x.\overline{\text{Inq}}@A_1.\overline{\text{Inq}}@A_2.\overline{\text{Price}}@A_1.\overline{\text{Price}}@A_2.\overline{\text{Lst}}@x.(\overline{\text{Close}}@x.\mathbf{1} + \text{Buy1}@x.\overline{\text{Ord}}@A_1.\overline{\text{Pay}}@x.\overline{\text{Conf}}@A_1.\mathbf{1} + \text{Buy2}@x.\overline{\text{Ord}}@A_2.\overline{\text{Pay}}@x.\overline{\text{Conf}}@A_2.\mathbf{1})$$

$$\sigma_A = \text{Inq}@x.\overline{\text{Price}}@x.(\overline{\text{Ord}}@x.\overline{\text{Conf}}@x.\overline{\text{Get}}@y.\mathbf{1} + \mathbf{1})$$

Table 3: Example of a travel agency

respectively. The elementary actions represent business activities that result in messages being sent or received. For example, the action $\overline{\text{Req}}@T$ undertaken by the customer results in a message being sent to the travel agency. In the example, the customer sends a request to the travel agency which then inquires the airlines to get the prices for the selected route. Each airline responds and the travel agency sends to the customer the list of the best prices. The customer decides whether to close the communication with the travel agency or to buy from one of the airlines. In the latter case the customer decides to assign a high security level (\mathbb{H}) to both the travel agency and the chosen airline company in order to safeguard the confidentiality of the purchasing data. The travel agency orders the ticket from the selected airline and takes a deposit (or a full payment) from the customer. As soon as the airline receives the confirmation of the payment, the ticket is issued to the customer. \square

3. Non-Interference

The concept of *non-interference* [9] has been introduced to formalize the absence of information flow in multilevel systems. In the context of service compositions it demands that public interactions between service components are unchanged as secret communications are varied or, more generally, that the low level behaviour of the service composition is independent of the behaviour of its high components. In this way clients are assured that the data transmitted over the air to a web server remains confidential; in other words, sensitive data cannot be intercepted and understood by eavesdroppers.

The notion of non-interference we are going to introduce is relative to the internal behaviour of service compositions, i.e., we are interested in observing the

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Typed internal composition transitions: $\Gamma \triangleright C \xrightarrow{\alpha} \Gamma' \triangleright C'$

$$\frac{\Gamma \triangleright \sigma \xrightarrow{\tau} \Gamma' \triangleright \sigma'}{\Gamma \triangleright p[\sigma] \xrightarrow{\tau} \Gamma' \triangleright p[\sigma']} \quad \frac{\Gamma \triangleright C_1 \xrightarrow{\alpha} \Gamma' \triangleright C'_1}{\Gamma \triangleright C_1 \parallel C_2 \xrightarrow{\alpha} \Gamma' \triangleright C'_1 \parallel C_2}$$

$$\frac{\Gamma \triangleright C_1 \xrightarrow{a_{p \rightarrow q}} \Gamma \triangleright C'_1 \quad \Gamma \triangleright C_2 \xrightarrow{\bar{a}_{p \rightarrow q}} \Gamma \triangleright C'_2}{\Gamma \triangleright C_1 \parallel C_2 \xrightarrow{\{a\}_{p \rightarrow q}} \Gamma \triangleright C'_1 \parallel C'_2}$$

Table 4: Typed internal composition transitions

synchronizations between service components. We thus refine the semantics of compositions in order to help (i) to distinguish a local contract move from a synchronization, and (ii) to identify the principals involved in every synchronization. This is captured by the rules collected in Table 4, where we use the relation $\xrightarrow{\alpha}$ to represent typed synchronizations between service components. The τ label now indicates an internal action to a service component, while synchronizations between different peers in a composition are represented through a label of the form $\{a\}_{p \rightarrow q}$ meaning that principals p and q synchronize on action a . We let α range over the labels $\{a\}_{p \rightarrow q}$ and τ . We denote by $\xrightarrow{\tau}$ a possible empty sequence of $\xrightarrow{\tau}$ and define $\xrightarrow{\{a\}_{p \rightarrow q}} \stackrel{\text{def}}{=} \xrightarrow{\tau} \xrightarrow{\{a\}_{p \rightarrow q}} \xrightarrow{\tau}$.

The following lemma relates the two semantics for service compositions, one expressed in terms of \longrightarrow and the other one expressed in terms of $\xrightarrow{\alpha}$.

Lemma 3.1. *Let Γ be a typed environment and C be a service composition such that $\Gamma \triangleright C$ is a configuration.*

- $\Gamma \triangleright C \xrightarrow{\tau} \Gamma' \triangleright C'$ if and only if $C = C_1 \parallel p[\sigma] \parallel C_2$, $C' = C_1 \parallel p[\sigma'] \parallel C_2$ and $\sigma \xrightarrow{\tau} \sigma'$;
- $\Gamma \triangleright C \xrightarrow{\{a\}_{p \rightarrow q}} \Gamma \triangleright C'$ if and only if $C = C_1 \parallel p[\sigma] \parallel C_2 \parallel q[\rho] \parallel C_3$, $C' = C_1 \parallel p[\sigma'] \parallel C_2 \parallel q[\rho'] \parallel C_3$, $\sigma \xrightarrow{\bar{a}@q} \sigma'$ and $\rho \xrightarrow{a@p} \rho'$.

Proof. The proof follows by induction on the derivations. □

In order to define our notion of non-interference, we need to be able to distinguish the component interactions at a given security clearance. As transitions are typed, we can assign a security level to them as follows: the level of a synchronization depends on the level of the principals performing it. More precisely, by abuse

of notation, we denote by $\Gamma(\{a\}_{p \rightarrow q})$ the level of the synchronization $\{a\}_{p \rightarrow q}$ in the type environment Γ and define it as:

$$\Gamma(\{a\}_{p \rightarrow q}) = \Gamma(p) \sqcap \Gamma(q).$$

Thus a ς -level synchronization between two components is performed by principals whose security clearance is higher or equal to ς .

3.1. Behavioural Observations

We define behavioural observations in terms of equivalences that are parametric with respect to the security level $\varsigma \in \Sigma$ of the behaviour we want to observe. Such equivalences are relations over configurations that equate service compositions exhibiting the same ς -level component interactions.

Our behavioural equivalences are defined as a variant of the notion of *weak bisimulation* [16], an observation equivalence which allows one to observe the non-deterministic structure of the LTSs and focuses only on the observable actions.

In the following, we write $\Gamma_1 =_{\varsigma} \Gamma_2$ whenever $\{p \in \text{dom}(\Gamma_1) \mid \Gamma_1(p) \preceq \varsigma\} = \{p \in \text{dom}(\Gamma_2) \mid \Gamma_2(p) \preceq \varsigma\}$.

Definition 3.1 (Weak bisimulation on ς -low actions). Let $\varsigma \in \Sigma$. A *weak bisimulation on ς -low actions* is the largest symmetric relation \approx_{ς} over configurations such that whenever $\Gamma_1 \triangleright C_1 \approx_{\varsigma} \Gamma_2 \triangleright C_2$ with $\Gamma_1 =_{\varsigma} \Gamma_2$

- if $\Gamma_1 \triangleright C_1 \xrightarrow{\alpha} \Gamma'_1 \triangleright C'_1$ with $\alpha = \tau$ or $\Gamma_1(\alpha) \preceq \varsigma$, then there exist Γ'_2 and C'_2 such that $\Gamma_2 \triangleright C_2 \xrightarrow{\alpha} \Gamma'_2 \triangleright C'_2$ with $\Gamma'_1 \triangleright C'_1 \approx_{\varsigma} \Gamma'_2 \triangleright C'_2$ and $\Gamma'_1 =_{\varsigma} \Gamma'_2$;
- if $\Gamma_1 \triangleright C_1 \xrightarrow{\alpha} \Gamma'_1 \triangleright C'_1$ with $\alpha \neq \tau$ and $\Gamma_1(\alpha) \not\preceq \varsigma$, then there exist Γ'_2 and C'_2 such that
 - either $\Gamma_2 \triangleright C_2 \xrightarrow{\alpha} \Gamma'_2 \triangleright C'_2$
 - or $\Gamma_2 \triangleright C_2 \xrightarrow{\tau} \Gamma'_2 \triangleright C'_2$
 with $\Gamma'_1 \triangleright C'_1 \approx_{\varsigma} \Gamma'_2 \triangleright C'_2$ and $\Gamma'_1 =_{\varsigma} \Gamma'_2$.

We write $\Gamma \models C_1 \approx_{\varsigma} C_2$ when $\Gamma \triangleright C_1 \approx_{\varsigma} \Gamma \triangleright C_2$. □

$$M = C[\sigma_C] \parallel F_1[\sigma_F] \parallel F_2[\sigma_F] \parallel S[\sigma_S]$$

$$\sigma_C = \overline{\text{Inq}}@F_1.\overline{\text{Inq}}@F_2.\text{Plan}@F_1.\text{Plan}@F_2.$$

$$(\overline{\text{Agree}}@F_1.\overline{\text{Close}}@F_2.\mathbf{1}_{[F_1:\text{H} \oplus F_2:\text{H}]}\overline{\text{Agree}}@F_2.\overline{\text{Close}}@F_1.\mathbf{1}))$$

$$\sigma_F = \text{Inq}@x.\overline{\text{LookUp}}@S.\overline{\text{Quote}}@x.\overline{\text{Plan}}@C.(\text{Agree}@x.\mathbf{1} + \text{Close}@x.\mathbf{1})$$

$$\sigma_S = \text{LookUp}@x.\overline{\text{Quote}}@x.\mathbf{1}$$

$$M' = C[\sigma'_C] \parallel F_1[\sigma_F] \parallel F_2[\sigma_F] \parallel S[\sigma_S]$$

$$\sigma'_C = \overline{\text{Inq}}@F_1.\overline{\text{Inq}}@F_2.\text{Plan}@F_1.\text{Plan}@F_2.$$

$$(\overline{\text{Close}}@F_2.\overline{\text{Agree}}@F_1.\mathbf{1}_{[F_1:\text{H} \oplus F_2:\text{H}]}\overline{\text{Close}}@F_1.\overline{\text{Agree}}@F_2.\mathbf{1}))$$

Table 5: Example of a financial consulting platform

3.2. Non-interference

The notion of non-interference is inspired by [12] and is expressed in terms of a restriction operator $\cdot|_\varsigma$ which allows one to represent a service composition prevented from performing internal synchronizations of a level higher than ς . The semantics of $C|_\varsigma$ is described by the following rule:

$$\frac{\Gamma \triangleright C \xrightarrow{\alpha} \Gamma' \triangleright C'}{\Gamma \triangleright C|_\varsigma \xrightarrow{\alpha} \Gamma' \triangleright C'|_\varsigma} \quad \Gamma(\alpha) \preceq \varsigma$$

Definition 3.2 (Non-interference). Let $\varsigma \in \Sigma$, Γ be a type environment and C be a service composition such that $\Gamma \triangleright C$ be a configuration. We say that the service composition C satisfies the non-interference property with respect to the level ς in the type environment Γ , denoted $C \in \mathcal{NI}_{\Gamma, \varsigma}$, if

$$\Gamma \models C \approx_\varsigma C|_\varsigma. \quad \square$$

Example 3.1. Consider again the service composition S in the type environment Γ of Example 2.1. The property $S \in \mathcal{NI}_{\Gamma, \perp}$ holds. \square

Example 3.2. Consider the service composition depicted in Table 5: it consists of a client C , two financial consulting services F_1 and F_2 and a stock quote service provider S . The client inquires the financial services to get investment advices. The financial services consult the stock quote service provider in order to look up information on the financial quotes. Then, the financial services send their investment recommendations to the client which may decide whether or not adhere to the investment plan proposed by one of the financial services and close the connection with the other one.

Let $\Sigma = \{L, H\}$ with $L \preceq H$ and Γ be the type environment $C : H, F_1 : L, F_2 : L, S : L$. In this case we have that $M \notin \mathcal{NT}_{\Gamma, L}$. Indeed, there is a direct causality between the high level actions $\{\text{Agree}\}_{C \rightarrow F_i}$ and the low level action $\{\text{Close}\}_{C \rightarrow F_j}$ with $i \neq j$, performed after the clients makes the choice. As a consequence, if the client decides to accept the proposal of F_1 then F_2 knows that the customer has agreed to proceed with investment recommendation of F_1 by just observing that the action $\{\text{Close}\}_{C \rightarrow F_2}$ has been performed. The service composition can be made secure by letting $\{\text{Close}\}_{C \rightarrow F_j}$ be executed independently of $\{\text{Agree}\}_{C \rightarrow F_i}$ as in the composition M' which is obtained from M by replacing the contract σ_C with σ'_C . \square

4. Modal Formulae for Non-Interference

In this section we present a method for verifying whether $\Gamma \models C \approx_\varsigma C|_\varsigma$ which consists in defining a modal μ -calculus formula $\phi^{\approx_\varsigma}(\Gamma \triangleright C)$ such that $\Gamma' \triangleright C'$ satisfies $\phi^{\approx_\varsigma}(\Gamma \triangleright C)$ iff $\Gamma \triangleright C \approx_\varsigma \Gamma' \triangleright C'$.

The modal μ -calculus [15] is a small, yet expressive process logic. We consider modal μ -calculus formulae constructed according to the following grammar:

$$\phi ::= \mathbf{true} \mid \mathbf{false} \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \langle \alpha \rangle \phi \mid [\alpha] \phi \mid X \mid \mu X. \phi \mid \nu X. \phi$$

where X ranges over an infinite set of variables and α over the labels $\{a\}_{p \rightarrow q}$ and τ . The *fixpoint operators* μX and νX bind the variable X and we adopt the usual notion of closed formula. For a finite set M of formulae, we write $\bigwedge M$ and $\bigvee M$ for the conjunction and disjunction of the formulae in M , where $\bigwedge \emptyset = \mathbf{true}$ and $\bigvee \emptyset = \mathbf{false}$.

Modal μ -calculus formulae are interpreted over configurations modelled by LTS's. Let $\Gamma \triangleright C$ be a configuration and $LTS(\Gamma \triangleright C) = (S_{\Gamma \triangleright C}, Act, \xrightarrow{\quad})$ where $S_{\Gamma \triangleright C}$ is the set of states reachable from $\Gamma \triangleright C$ through $\xrightarrow{\alpha}$ and $\alpha \in Act$ denotes the action $\{a\}_{p \rightarrow q}$ or τ . The subset of configurations in $S_{\Gamma \triangleright C}$ that satisfy a formula ϕ , noted by $M_{\Gamma \triangleright C}(\phi)(\rho)$, is inductively defined in Table 6. ρ is an *environment*, i.e., it is a partial mapping $\rho : Var \dashv \twoheadrightarrow 2^{S_{\Gamma \triangleright C}}$ that interprets at least

$M_{\Gamma \triangleright C}(\mathbf{true})(\rho)$	$= S_{\Gamma \triangleright C}$
$M_{\Gamma \triangleright C}(\mathbf{false})(\rho)$	$= \emptyset$
$M_{\Gamma \triangleright C}(\phi_1 \wedge \phi_2)(\rho)$	$= M_{\Gamma \triangleright C}(\phi_1)(\rho) \cap M_{\Gamma \triangleright C}(\phi_2)(\rho)$
$M_{\Gamma \triangleright C}(\phi_1 \vee \phi_2)(\rho)$	$= M_{\Gamma \triangleright C}(\phi_1)(\rho) \cup M_{\Gamma \triangleright C}(\phi_2)(\rho)$
$M_{\Gamma \triangleright C}(\langle \alpha \rangle \phi)(\rho)$	$= \{\Gamma' \triangleright C' \mid \exists \Gamma'' \triangleright C'' : \Gamma' \triangleright C' \xrightarrow{\alpha} \Gamma'' \triangleright C''$ $\quad \wedge \Gamma'' \triangleright C'' \in M_{\Gamma \triangleright C}(\phi)(\rho)\}$
$M_{\Gamma \triangleright C}([\alpha] \phi)(\rho)$	$= \{\Gamma' \triangleright C' \mid \forall \Gamma'' \triangleright C'' : \Gamma' \triangleright C' \xrightarrow{\alpha} \Gamma'' \triangleright C''$ $\quad \Rightarrow \Gamma'' \triangleright C'' \in M_{\Gamma \triangleright C}(\phi)(\rho)\}$
$M_{\Gamma \triangleright C}(X)(\rho)$	$= \rho(X)$
$M_{\Gamma \triangleright C}(\mu X. \phi)(\rho)$	$= \bigcap \{x \subseteq S_{\Gamma \triangleright C} \mid M_{\Gamma \triangleright C}(\phi)(\rho[X \mapsto x]) \subseteq x\}$
$M_{\Gamma \triangleright C}(\nu X. \phi)(\rho)$	$= \bigcup \{x \subseteq S_{\Gamma \triangleright C} \mid M_{\Gamma \triangleright C}(\phi)(\rho[X \mapsto x]) \supseteq x\}$

Table 6: Semantics of modal mu-calculus

the free variables of ϕ by subsets of $S_{\Gamma \triangleright C}$. For a set $x \subseteq S_{\Gamma \triangleright C}$ and a variable X , we write $\rho[X \mapsto x]$ for the environment that maps X to x and $Y \neq X$ to $\rho(Y)$ if ρ is defined on Y .

Intuitively, **true** and **false** hold for all, resp. no, states and \wedge and \vee are interpreted by conjunction and disjunction, $\langle \alpha \rangle \phi$ holds for a configuration $\Gamma' \triangleright C' \in S_{\Gamma \triangleright C}$ if there exists $\Gamma'' \triangleright C''$ reachable from $\Gamma' \triangleright C'$ with action α and satisfying ϕ , and $[\alpha] \phi$ holds for $\Gamma' \triangleright C'$ if all configurations $\Gamma'' \triangleright C''$ reachable from $\Gamma' \triangleright C'$ with action α satisfy ϕ . The interpretation of a variable X is as prescribed by the environment. The formula $\mu X. \phi$, called *least fixpoint formula*, is interpreted by the smallest subset x of $S_{\Gamma \triangleright C}$ that recurs when ϕ is interpreted with the substitution of x for X . Similarly, $\nu X. \phi$, called *greatest fixpoint formula*, is interpreted by the largest such set. Existence of such sets follows from the well-known Knaster-Tarski fixpoint theorem. As the meaning of a closed formula ϕ does not depend on the environment, we sometimes write $M_{\Gamma \triangleright C}(\phi)$ for $M_{\Gamma \triangleright C}(\phi)(\rho)$ where ρ is an arbitrary environment.

The set of configurations *satisfying* a closed formula ϕ is defined as $Conf(\phi) = \{\Gamma \triangleright C \mid \Gamma \triangleright C \in M_{\Gamma \triangleright C}(\phi)\}$. We also refer to (closed) *equation systems*:

$$Eqn : X_1 = \phi_1 \quad \dots \quad X_n = \phi_n$$

where X_1, \dots, X_n are distinct variables and ϕ_1, \dots, ϕ_n are formulae having at most X_1, \dots, X_n as free variables.

An environment $\rho : \{X_1, \dots, X_n\} \rightarrow 2^{S_{\Gamma \triangleright C}}$ is a *solution* of an equation system Eqn , if $\rho(X_i) = M_{\Gamma \triangleright C}(\phi_i)(\rho)$. The fact that solutions always exist, is again a

consequence of the Knaster-Tarski fixpoint theorem. In fact the set of environments that are candidates for solutions, $Env_{\Gamma \triangleright C} = \{\rho \mid \rho : \{X_1, \dots, X_n\} \rightarrow 2^{S_{\Gamma \triangleright C}}\}$, together with the lifting \sqsubseteq of the inclusion order on $2^{S_{\Gamma \triangleright C}}$, defined by

$$\rho \sqsubseteq \rho' \text{ iff } \rho(X_i) \subseteq \rho'(X_i) \text{ for } i \in [1..n]$$

forms a complete lattice. Now, we can define the *equation functional* $Func_{\Gamma \triangleright C}^{Eqn} : Env_{\Gamma \triangleright C} \rightarrow Env_{\Gamma \triangleright C}$ by $Func_{\Gamma \triangleright C}^{Eqn}(\rho)(X_i) = M_{\Gamma \triangleright C}(\phi_i)(\rho)$ for $i \in [1..n]$, the fixpoints of which are just the solutions of Eqn . $Func_{\Gamma \triangleright C}^{Eqn}$ is monotonic and there is the largest solution $\nu Func_{\Gamma \triangleright C}^{Eqn}$ of Eqn (with respect to \sqsubseteq), which we denote by $M_{\Gamma \triangleright C}(Eqn)$. This definition interprets equation systems on the configurations reachable by a given initial configuration $\Gamma \triangleright C$. We lift this to configurations by agreeing that a configuration satisfies an equation system Eqn , if its initial state is in the largest solution of the first equation. Thus the set of configurations satisfying the equation system Eqn is $Conf(Eqn) = \{\Gamma \triangleright C \mid \Gamma \triangleright C \in M_{\Gamma \triangleright C}(Eqn)(X_1)\}$.

The relation \approx_ς can be characterized as the greatest fixpoint $\nu Func_{\approx_\varsigma}$ of the monotonic functional $Func_{\approx_\varsigma}$ on the complete lattice of relations \mathcal{R} over configurations ordered by set inclusion, where $(\Gamma_1 \triangleright C_1, \Gamma_2 \triangleright C_2) \in Func_{\approx_\varsigma}(\mathcal{R})$ if and only if points (1) and (2) of Definition 3.1 hold. Thus a relation \mathcal{R} is a weak bisimulation on ς -low actions if and only if $\mathcal{R} \subseteq Func_{\approx_\varsigma}(\mathcal{R})$, i.e., \mathcal{R} is a *post-fixpoint* of $Func_{\approx_\varsigma}$. By the Knaster-Tarski fixpoint theorem, $\nu Func_{\approx_\varsigma}$ is the union of all the post-fixpoints of $Func_{\approx_\varsigma}$, i.e., it is the largest weak bisimulation on ς -low actions. If we restrict to the complete lattice of relations $\mathcal{R} \subseteq S_{\Gamma_1 \triangleright C_1} \times S_{\Gamma_2 \triangleright C_2}$ we obtain a monotonic functional $Func_{\approx_\varsigma}^{(\Gamma_1 \triangleright C_1, \Gamma_2 \triangleright C_2)}$ whose greatest fixpoint is exactly $\nu Func_{\approx_\varsigma} \cap (S_{\Gamma_1 \triangleright C_1} \times S_{\Gamma_2 \triangleright C_2})$, and this is enough to determine if $\Gamma_1 \triangleright C_1 \approx_\varsigma \Gamma_2 \triangleright C_2$.

Let $\Gamma \triangleright C$ be finite-state, $\Gamma_1 \triangleright C_1, \dots, \Gamma_n \triangleright C_n$ its $|S_{\Gamma \triangleright C}| = n$ states, and $\Gamma_1 \triangleright C_1 = \Gamma \triangleright C$ its initial state. To derive a formula characterizing $\Gamma \triangleright C$ up to \approx_ς we construct a *characteristic equation system* [19]:

$$Eqn_{\approx_\varsigma} : X_{\Gamma_1 \triangleright C_1} = \phi_{\Gamma_1 \triangleright C_1}^{\approx_\varsigma}, \dots, X_{\Gamma_n \triangleright C_n} = \phi_{\Gamma_n \triangleright C_n}^{\approx_\varsigma}$$

consisting of one equation for each service configuration $\Gamma_1 \triangleright C_1, \dots, \Gamma_n \triangleright C_n \in S_{\Gamma \triangleright C}$. We define the formulae $\phi_{\Gamma_i \triangleright C_i}^{\approx_\varsigma}$ such that the largest solution $M_{\Gamma \triangleright C}(Eqn_{\approx_\varsigma})$ of Eqn_{\approx_ς} associates the variables $X_{\Gamma_i \triangleright C_i}$ just with the states $\Gamma'_i \triangleright C'_i$ of $S_{\Gamma \triangleright C}$ which are weakly bisimilar on ς -low actions to $\Gamma_i \triangleright C_i$, i.e., such that it holds $M_{\Gamma \triangleright C}(Eqn_{\approx_\varsigma})(X_{\Gamma_i \triangleright C_i}) = \{\Gamma'_i \triangleright C'_i \in S_{\Gamma \triangleright C} \mid \Gamma_i \triangleright C_i \approx_\varsigma \Gamma'_i \triangleright C'_i\}$. Theorem 4.1 shows the exact form of such formulae. First we define:

$$\langle\langle \alpha \rangle\rangle_{\Gamma, \varsigma} \phi \stackrel{\text{def}}{=} \begin{cases} \langle\langle \alpha \rangle\rangle \phi & \text{if } \Gamma(\alpha) \preceq \varsigma \text{ or } \alpha = \tau \\ \langle\langle \alpha \rangle\rangle \phi \vee \langle\langle \tau \rangle\rangle \phi & \text{if } \Gamma(\alpha) \not\preceq \varsigma \text{ and } \alpha \neq \tau \end{cases}$$

where $\langle\langle\tau\rangle\rangle\phi \stackrel{\text{def}}{=} \mu X. \phi \vee \langle\tau\rangle X$ and $\langle\langle\alpha\rangle\rangle\phi \stackrel{\text{def}}{=} \langle\langle\tau\rangle\rangle\langle\alpha\rangle\langle\langle\tau\rangle\rangle\phi$. Let $\xrightarrow{\alpha}_{\Gamma, \varsigma}$ note either $\xrightarrow{\alpha}_{\rightarrow}$ or $\xrightarrow{\tau}_{\rightarrow}$. Then $\langle\langle\alpha\rangle\rangle_{\Gamma, \varsigma}$, $\langle\langle\tau\rangle\rangle$ and $\langle\langle\alpha\rangle\rangle$ correspond to $\xrightarrow{\alpha}_{\rightarrow}$, $\xrightarrow{\tau}_{\rightarrow}$ and $\xrightarrow{\alpha}_{\rightarrow}$, since

- $M_{\Gamma \triangleright C}(\langle\langle\alpha\rangle\rangle_{\Gamma, \varsigma}\phi)(\rho) = \{\Gamma' \triangleright C' \mid \exists \Gamma'' \triangleright C'' : \Gamma' \triangleright C' \xrightarrow{\alpha}_{\rightarrow} \Gamma'' \triangleright C'' \wedge \Gamma'' \triangleright C'' \in M_{\Gamma \triangleright C}(\phi)(\rho)\}$
- $M_{\Gamma \triangleright C}(\langle\langle\tau\rangle\rangle\phi)(\rho) = \{\Gamma' \triangleright C' \mid \exists \Gamma'' \triangleright C'' : \Gamma' \triangleright C' \xrightarrow{\tau}_{\rightarrow} \Gamma'' \triangleright C'' \wedge \Gamma'' \triangleright C'' \in M_{\Gamma \triangleright C}(\phi)(\rho)\}$
- $M_{\Gamma \triangleright C}(\langle\langle\alpha\rangle\rangle\phi)(\rho) = \{\Gamma' \triangleright C' \mid \exists \Gamma'' \triangleright C'' : \Gamma' \triangleright C' \xrightarrow{\alpha}_{\rightarrow} \Gamma'' \triangleright C'' \wedge \Gamma'' \triangleright C'' \in M_{\Gamma \triangleright C}(\phi)(\rho)\}$.

Theorem 4.1. Let $\phi_{\Gamma_i \triangleright C_i}^{\approx_\varsigma}$ be the formula

$$\begin{aligned} & \wedge \{ \wedge \{ \langle\langle\alpha\rangle\rangle_{\Gamma, \varsigma} X_{\Gamma'_i \triangleright C'_i} \mid \Gamma_i \triangleright C_i \xrightarrow{\alpha}_{\rightarrow} \Gamma'_i \triangleright C'_i \} \} \\ & \wedge \{ \wedge \{ [\alpha] \vee \{ X_{\Gamma'_i \triangleright C'_i} \mid \Gamma_i \triangleright C_i \xrightarrow{\alpha}_{\rightarrow} \Gamma'_i \triangleright C'_i \} \} \}. \end{aligned}$$

Then $M_{\Gamma \triangleright C}(\text{Eqn}_{\approx_\varsigma})(X_{\Gamma_i \triangleright C_i}) = \{\Gamma'_i \triangleright C'_i \in S_{\Gamma \triangleright C} \mid \Gamma_i \triangleright C_i \approx_\varsigma \Gamma'_i \triangleright C'_i\}$.

Proof. The proof is standard [19] and is based on the observation that $\text{Env}_{\Gamma \triangleright C}$, the set of candidates for solutions of $\text{Eqn}_{\approx_\varsigma}$ is order-isomorphic to $2^{S_{\Gamma \triangleright C} \times S_{\Gamma \triangleright C}}$, the set of relations that are candidates to be weak bisimulation on ς -low actions between $S_{\Gamma \triangleright C} \times S_{\Gamma \triangleright C}$. Actually, the mapping $\alpha : \text{Env}_{\Gamma \triangleright C} \rightarrow 2^{S_{\Gamma \triangleright C} \times S_{\Gamma \triangleright C}}$ defined by:

$$\alpha(\rho) = \{(\Gamma_i \triangleright C_i, \Gamma'_i \triangleright C'_i) \in S_{\Gamma \triangleright C} \times S_{\Gamma \triangleright C} \mid \Gamma'_i \triangleright C'_i \in \rho(X_{\Gamma_i \triangleright C_i})\}$$

is an order isomorphism between $\text{Env}_{\Gamma \triangleright C}$ and $2^{S_{\Gamma \triangleright C} \times S_{\Gamma \triangleright C}}$, the inverse of which is the mapping $\beta : 2^{S_{\Gamma \triangleright C} \times S_{\Gamma \triangleright C}} \rightarrow \text{Env}_{\Gamma \triangleright C}$ defined by $\beta(\mathcal{R})(X_{\Gamma_i \triangleright C_i}) = \{\Gamma'_i \triangleright C'_i \in S_{\Gamma \triangleright C} \mid (\Gamma_i \triangleright C_i, \Gamma'_i \triangleright C'_i) \in \mathcal{R}\}$.

The proof follows by showing that $\text{Func}_{\approx_\varsigma}$ and $\text{Func}_{\Gamma \triangleright C}^{\text{Eqn}_{\approx_\varsigma}}$ are equal up to the isomorphism induced by (α, β) , i.e., such that

$$\text{Func}_{\Gamma \triangleright C}^{\text{Eqn}_{\approx_\varsigma}} = \beta \circ \text{Func}_{\approx_\varsigma} \circ \alpha.$$

Then their largest fixpoints are also related by the isomorphism, which yields $M_{\Gamma \triangleright C}(\text{Eqn}_{\approx_\varsigma})(X_{\Gamma_i \triangleright C_i}) = \{\Gamma'_i \triangleright C'_i \in S_{\Gamma \triangleright C} \mid \Gamma_i \triangleright C_i \approx_\varsigma \Gamma'_i \triangleright C'_i\}$. \square

Characteristic formulae, i.e., *single* formulae characterizing configurations can be constructed by applying simple semantics-preserving transformation rules on equation systems. These rules are similar to the ones used by A. Mader in [20]

as a mean of solving Boolean equation systems (with alternation) by Gauss elimination. Hence, since for any equation system Eqn there is a formula ϕ such that $Conf(Eqn) = Conf(\phi)$, we obtain that:

Theorem 4.2. *For any finite-state configuration $\Gamma \triangleright C$ there is a modal μ -calculus formula $\phi^{\approx_\varsigma}(\Gamma \triangleright C)$ such that $Conf(\phi^{\approx_\varsigma}(\Gamma \triangleright C)) = \{\Gamma' \triangleright C' \in S_{\Gamma \triangleright C} \mid \Gamma' \triangleright C' \approx_\varsigma (\Gamma' \triangleright C')|_\varsigma\}$, that is the set of all the states reachable from $\Gamma \triangleright C$ and satisfying $\mathcal{NT}_{\Gamma, \varsigma}$ \square*

5. A Modal Formula for Compliance

Compliance is a basic property that characterizes the correct behavior of concurrent distributed systems. It is used widely in the context of SOA as a formal device to identify well behaving service compositions, those whose interactions are free of synchronization errors.

In this paper we refer to the notion of compliance for contract service compositions studied in [21]. Intuitively, a composition of services is compliant if it is deadlock and livelock free, i.e., it does not get stuck nor does it get trapped into infinite loops with no exit states. This notion is independent of the security levels of the principals involved in the component synchronizations. Therefore, we omit trailing type environments in the definitions below, and write, e.g., $C \Longrightarrow C'$ to denote a transition of the form $\Gamma \triangleright C \Longrightarrow \Gamma' \triangleright C'$ for some type environments Γ and Γ' .

Definition 5.1 (Compliance). Let C be a contract service composition. We say that C is *compliant*, noted $C \downarrow$, if for every C' such that $C \Longrightarrow C'$ there exists C'' such that $C' \Longrightarrow C''$ and $C'' \checkmark$. \square

In other words, the notion of compliance ensures that at each intermediate step of the computation in a service composition, each component has a way to reach a successful state (either autonomously, or via synchronizations). This is enough to avoid both deadlocks and livelocks.

Example 5.1. *Consider the service composition S of Example 2.1. It holds that $\Gamma \triangleright S$ is both compliant, i.e., $S \downarrow$, and non interfering, i.e., $S \in \mathcal{NT}_{\Gamma, \bar{L}}$. \square*

The notion of compliance can be equivalently expressed in terms of $\xrightarrow{\alpha}$ where α denotes a synchronization $\{a\}_{p \rightarrow q}$ or τ . More precisely, let $\gamma = \alpha_1, \dots, \alpha_n$. We denote by $\xrightarrow{\gamma}$ the sequence of transitions $\xrightarrow{\alpha_1} \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n}$. Again we write $C \xrightarrow{\gamma} C'$ to denote a derivation $\Gamma \triangleright C \xrightarrow{\gamma} \Gamma' \triangleright C'$ for some type environments Γ and Γ' .

Proposition 5.1. *Let C be a contract service composition. It holds that C is compliant, $C \downarrow$, if and only if every C' such that $C \xrightarrow{\gamma'} C'$ for some $\gamma' \in Act^*$ there exist C'' and $\gamma'' \in Act^*$ such that $C' \xrightarrow{\gamma''} C''$ and $C'' \checkmark$.*

Proof. The proof follows by induction on the derivations and Lemma 3.1. \square

The modal μ -calculus formula that characterizes compliance is defined as follows:

$$\phi^c \stackrel{\text{def}}{=} \mu X. \left(\bigwedge_{\alpha \in Act} ([\alpha]X) \wedge \phi \right)$$

where

$$\phi \stackrel{\text{def}}{=} \mu X. \left((\checkmark) \vee \bigvee_{\alpha \in Act} (\langle \alpha \rangle X) \right) \wedge \neg \mu X. \left(\bigvee_{\alpha \in Act} (\langle \alpha \rangle X) \right)$$

The sub-formula $\neg \mu X. (\bigvee_{\alpha \in Act} (\langle \alpha \rangle X))$ will ensure that any configuration satisfying ϕ^c doesn't get trapped into infinite loops without chances to reach a successful state. The next theorem characterizes the set of service configurations satisfying ϕ^c . A complete proof is given in [22].

Theorem 5.1. *Consider the modal μ -calculus formula ϕ^c defined above. It holds that $Conf(\phi^c) = \{\Gamma \triangleright C \mid C \downarrow \text{ and } \Gamma \text{ is a type environment}\}$.*

Proof. The fact that $Conf(\phi^c) \supseteq \{\Gamma \triangleright C \mid C \downarrow \text{ and } \Gamma \text{ is a type environment}\}$ easily follows by contradiction. The other direction follows from the fact that, by definition of ϕ^c , for each configuration $\Gamma \triangleright C \in Conf(\phi^c)$ it holds that $\Gamma \triangleright C \in Conf(\phi')$ where ϕ' is the sub-formula $\mu X. ((\checkmark) \vee \bigvee_{\alpha \in Act} (\langle \alpha \rangle X))$. \square

Corollary 5.1. *A composition C is compliant if and only if $\Gamma \triangleright C \in Conf(\phi^c)$ for some type environment Γ .* \square

As a consequence of Theorems 4.2 and 5.1 we have:

Corollary 5.2. *Let $\varsigma \in \Sigma$, $\Gamma \triangleright C$ be a configuration and*

$$\Phi_{\Gamma \triangleright C}^{\varsigma} \stackrel{\text{def}}{=} \phi^{\approx_{\varsigma}}(\Gamma \triangleright C) \wedge \phi^c.$$

It holds that $\Gamma \triangleright C \in Conf(\Phi_{\Gamma \triangleright C}^{\varsigma})$ if and only if both $C \in \mathcal{NT}_{\Gamma, \varsigma}$ and $C \downarrow$. \square

Using this method we can for instance exploit the model checker NCSU Concurrency Workbench (see [23]) to check whether both $C \in \mathcal{NT}_{\Gamma, \varsigma}$ and $C \downarrow$.

Transitions for filters

$$\begin{array}{c}
\delta.f \xrightarrow{\delta} f \quad \frac{f\{X := \text{rec}(\mathbf{X})f\} \xrightarrow{\delta} f'}{\text{rec}(\mathbf{X})f \xrightarrow{\delta} f'} \quad \frac{f \xrightarrow{\delta} f_{\delta} \quad g \xrightarrow{\delta} g_{\delta}}{f \otimes g \xrightarrow{\delta} f_{\delta} \otimes g_{\delta}} \\
\frac{f \xrightarrow{\delta} f_{\delta} \quad g \xrightarrow{\delta} g_{\delta}}{f \times g \xrightarrow{\delta} f_{\delta} \times g_{\delta}} \quad \frac{f \xrightarrow{\delta} f_{\delta} \quad g \not\xrightarrow{\delta}}{f \times g \xrightarrow{\delta} f_{\delta}} \quad \frac{f \not\xrightarrow{\delta} \quad g \xrightarrow{\delta} g_{\delta}}{f \times g \xrightarrow{\delta} g_{\delta}}
\end{array}$$

Transitions for filtered peers

$$\begin{array}{c}
\frac{\Gamma \triangleright p[\sigma] \xrightarrow{\delta} \Gamma \triangleright p[\sigma'] \quad f \xrightarrow{\delta} f'}{\Gamma \triangleright f(p[\sigma]) \xrightarrow{\delta} \Gamma \triangleright f'(p[\sigma'])} \\
\frac{\Gamma \triangleright p[\sigma] \xrightarrow{\tau} \Gamma' \triangleright p[\sigma']}{\Gamma \triangleright f(p[\sigma]) \xrightarrow{\tau} \Gamma' \triangleright f(p[\sigma'])} \quad \frac{\Gamma \triangleright p[\sigma] \checkmark}{\Gamma \triangleright f(p[\sigma]) \checkmark}
\end{array}$$

Table 7: Dynamics of Filtered contract service compositions

6. An Adaptive Algorithm

The model checking technique is based on the idea that the state transition graph of a finite-state system defines a Kripke structure, and efficient algorithms can be given for checking if the state graph defines a model of a given specification expressed in an appropriate temporal logic. In the explicit state approach the Kripke structure is represented extensionally, using conventional data structures such as adjacency matrices and linked lists so that each state and transition is enumerated explicitly. Moreover, in the global calculation approach, given a structure M and formula ϕ , the model checking algorithms calculate $\phi^M = \{s : M, s \models \phi\}$ that is the set of all states in M satisfying ϕ . We show how such algorithms can be exploited to develop an adaptive model checking technique for service compositions which adapts, when it is possible, the composition under investigation in such a way that it satisfies both non-interference and compliance. We use the filters, introduced in [4] and revised in [21], as prescriptions of behaviour.

Filters. A filter is the specification of the legal flow of actions for an individual contract. The syntax is as follows, while the semantics is defined in Table 7.

$$f \in \mathcal{F} := \mathbf{0} \mid \delta.f \mid f \times f \mid f \otimes f \mid X \mid \text{rec}(\mathbf{X})f$$

Definition 6.1 (Filter pre-order). The filter pre-order $f \leq g$ is the largest relation such that if $f \xrightarrow{\delta} f_\delta$ then $g \xrightarrow{\delta} g_\delta$ and $f_\delta \leq g_\delta$. \square

We note $(\mathcal{F}, \sqsubseteq)$ the partial order induced by \leq : by abuse of notation, we identify a filter f with its equivalence class $[f]_\sim$, where \sim is the symmetric closure of \leq . The union and intersection of filters represent the glb and lub operators for $(\mathcal{F}, \sqsubseteq)$. Furthermore, if we assume a finite alphabet A of actions, the set of filters \mathcal{F}_A insisting on A forms a complete lattice with $\mathbf{0}$ as bottom and the identity filter $I_A \stackrel{\text{def}}{=} \text{rc}(\mathbf{X}) \prod_{\delta \in A} \delta.X$ as top element.

The application $\Gamma \triangleright f(p[\sigma])$ blocks any action from $\Gamma \triangleright p[\sigma]$ that is not explicitly enabled by f . Filters may be composed to help shape a service composition. Given a set π of principals, a composite π -filter F is a finite map from the principals in π to filters: $\{p \rightarrow f[p] \mid p \in \pi\}$. A π -filter may be applied to a π -composition:

$$\Gamma \triangleright F(p_1[\sigma_1] \parallel \cdots \parallel p_n[\sigma_n]) ::= \Gamma \triangleright F[p_1](p_1[\sigma_1]) \parallel \cdots \parallel \Gamma \triangleright F[p_n](p_n[\sigma_n])$$

When we write $\Gamma \triangleright F(C)$ we tacitly assume that the underlying set of principals for both F and C is π . The operators of union and intersection, as well as the ordering on filters extends directly to composite filters, as expected. Namely, for F and G π -filters and for $\bullet \in \{\times, \otimes\}$:

$$\begin{aligned} F \leq_\pi G & \quad \text{iff} \quad F[p] \leq G[p] \quad \text{for all } p \in \pi \\ (F \bullet_\pi G)[p] & \stackrel{\text{def}}{=} F[p] \bullet G[p] \quad \text{for all } p \in \pi \end{aligned}$$

We generalize the syntax of service compositions by allowing the term $\Gamma \triangleright F(C)$ to account for the application of filters on the components of C . The dynamics of filtered service compositions derives directly by combining the transitions in Tables 2 and 7.

Relevance. Below we present an algorithm that given a configuration $\Gamma \triangleright C$ infers a composite filter F that fixes $\Gamma \triangleright C$, whenever such F exists. The algorithm is so structured as to guarantee two important properties on the inferred filter. On the one hand, the filter is as permissive as possible, in that it is the greatest (with respect to the pre-order \leq) among the filters that fix $\Gamma \triangleright C$. On the other side, the inferred filter is *relevant*, i.e., minimal in size: for any computation state reached by the service configuration via a series of τ transitions (local moves or synchronizations), the filter only enables actions that may be attempted at that state (either directly, or via a local choice), by one of the components of the service configuration.

Definition 6.2 (Relevance). Let π be a set of principals and \mathcal{C} be a non-empty set of π -configurations. A filter f is *p-relevant* in \mathcal{C} , written $f \propto_p \mathcal{C}$, if whenever

$f \xrightarrow{\delta} \widehat{f}$ one has $\delta \in \{a_{\rightarrow p}, \bar{a}_{p \rightarrow -}\}$ and there exists $\Gamma \triangleright C \in \mathcal{C}$ such that $\Gamma \triangleright C \xrightarrow{\alpha} \widehat{f}$ with $\alpha \in \{\{a\}_{\rightarrow p}, \{a\}_{p \rightarrow -}\}$ and $\widehat{f} \propto_p \{\Gamma' \triangleright C' \mid \Gamma \triangleright C \xrightarrow{\alpha} \Gamma' \triangleright C'\}$. A composite π -filter F is *relevant* for \mathcal{C} , written $F \propto \mathcal{C}$, if $F(p) \propto_p \mathcal{C}$ for all $p \in \pi$. A composite π -filter is relevant for a π -configuration $\Gamma \triangleright C$ if $F \propto \{\Gamma \triangleright C\}$. \square

The Algorithm. We describe an algorithm that synthesizes the \sqsubseteq -greatest relevant filter that fixes $\Gamma \triangleright C$, if it exists, when $\Gamma \triangleright C$ does not satisfy $\Phi_{\Gamma \triangleright C}^s$.

As discussed above, a global model checking algorithm applied to a configuration $\Gamma \triangleright C$ and the modal formula $\Phi_{\Gamma \triangleright C}^s$ calculates the set of states in the reduction graph (tracing the states reached by means of synchronizations or internal moves) of $\Gamma \triangleright C$ satisfying $\Phi_{\Gamma \triangleright C}^s$. This is the input of our algorithm. The reduction graph can be represented as a directed graph $G = (V, E)$ with labelled edges and vertices. The vertices in V represent the reachable states of $\Gamma \triangleright C$. With each $\mathbf{v} \in V$ we associate two fields: *state* $[\mathbf{v}]$ gives the computation state (i.e., the derivative $\Gamma' \triangleright C'$ of the initial state $\Gamma \triangleright C$) associated with \mathbf{v} ; *result* $[\mathbf{v}]$ is a tag SUCC or FAIL depending on whether the corresponding configuration satisfies $\Phi_{\Gamma \triangleright C}^s$ or not as calculated by the model checker. An edge in E is a triple $(\mathbf{u}, \mathbf{v})_\alpha$ representing the transition $state[\mathbf{u}] \xrightarrow{\alpha} state[\mathbf{v}]$. Reduction graphs may be stored in a adjacency list representation, so that the set of outgoing edges for each $\mathbf{u} \in V$ can be retrieved as *Adj* $[\mathbf{u}]$: thus $(\mathbf{u}, \mathbf{v})_\alpha \in E$ iff $(\alpha, \mathbf{v}) \in Adj[\mathbf{u}]$. We also write *Adj* $[\mathbf{u}, \alpha]$ for the set $\{\mathbf{v} \in V \mid (\mathbf{u}, \mathbf{v})_\alpha \in E\}$. Vertices with no outgoing edges are called leaves. We denote by *root* $[G]$ the vertex representing the initial state $\Gamma \triangleright C$.

The first step consists in *re-labelling* the graph G calculated by the model-checker in such a way that the result label at each vertex \mathbf{u} is set to FAIL if there exists at least one silent transition from \mathbf{u} to a FAIL vertex; it is set to SUCC if either there are no silent transitions from \mathbf{u} to a FAIL vertex and there exists a silent transition from \mathbf{u} to a SUCC vertex or there exists one non-silent and non-conflicting transition from \mathbf{u} to a SUCC vertex. The procedure iteratively examines all the vertices in the graph until it reaches a fixed point. This computation is accomplished by the `PushLabels` procedure and uses the following auxiliary definitions. Let *locs* (α) be $\{p, q\}$ in case $\alpha = \{a_{p \rightarrow q}\}$, and \emptyset in case $\alpha = \tau$. Then, let $G = (V, E)$ be a reduction graph, and $\alpha = \{a_{p \rightarrow q}\}$.

- A path $\varpi = (\mathbf{u}, \mathbf{u}_1)_{\alpha_1}, \dots, (\mathbf{u}_{n-1}, \mathbf{v})_{\alpha_n}$ from \mathbf{u} to \mathbf{v} in G is α -free if $locs(\alpha) \cap locs(\alpha_i) = \emptyset$ for all i 's.
- A vertex \mathbf{v} is a α -free descendant of \mathbf{u} in G (dually, \mathbf{u} is a α -free ancestor of \mathbf{v}) if there is a α -free path from \mathbf{u} to \mathbf{v} .
- A vertex \mathbf{u} yields a conflict on α if \mathbf{u} has two distinct α -free descendants \mathbf{v}_1 and \mathbf{v}_2 such that $(\mathbf{v}_1, \mathbf{w}_1)_\alpha$ and $(\mathbf{v}_2, \mathbf{w}_2)_\alpha \in E$ and $result[\mathbf{w}_1] \neq result[\mathbf{w}_2]$.

Procedure `PushLabels` (G)

Input: A reduction graph $G = (V, E)$

Output: The graph G updated

$done := false;$

while $\neg done$ **do**

$done := true;$

foreach $u \in V$ **do**

$succ := false; fail := false;$

if $Adj[u, \tau] \neq \emptyset$ **then**

if $\exists v \in Adj[u, \tau] : result[v] = FAIL$ **then**

$fail := true;$

else if $\exists v \in Adj[u, \tau] : result[v] = SUCC$ **then**

$succ := true;$

else if $\exists (\alpha, v) \in Adj[u] \wedge result[v] = SUCC \wedge \neg Conflict(\alpha, u)$

then

$succ := true;$

if $succ \wedge result[u] \neq SUCC$ **then**

$result[u] := SUCC; done := false;$

else if $fail \wedge result[u] \neq FAIL$ **then**

$result[u] := FAIL; done := false$

-
- A vertex v has a conflict on α in G , noted $Conflict_G(\alpha, v)$ if v has a α -free ancestor yielding a conflict on α .

Intuitively, our algorithm will prune G by banning all the ‘bad’ synchronizations, and by preserving all the ‘good’ synchronizations that lead to nodes satisfying both non-interference and compliance. Due to the presence of internal choices, the same synchronization can look good at one point, but actually be bad. The definition of conflict formally captures this notion of ambiguous synchronizations.

Proposition 6.1. *After the call to `PushLabels` (G), the following conditions hold for every node u in G :*

1. $result[u] = FAIL$ iff either there exists no $(u, v)_\alpha \in E$ such that $result[v] = SUCC$ and $\neg Conflict_G(\alpha, u)$ or there exists $(u, v)_\tau \in E$ such that $result[v] = FAIL$;
2. $result[u] = SUCC$ iff there exists no $(u, v)_\tau \in E$ such that $result[v] = FAIL$ and there exists either $(u, v)_\tau \in E$ such that $result[v] = SUCC$ or $(u, v)_\alpha \in E$ with $\alpha \neq \tau$, $\neg Conflict_G(\alpha, u)$ and $result[v] = SUCC$.

Function SuccessGraph (G)

Input: A reduction graph $G = (V, E)$
Output: $G' = (V', E')$ the success sub-graph of G
 $V' := (\text{result}[\text{root}[G]] = \text{SUCC}) ? \{\text{root}[G]\} : \emptyset; E' := \emptyset;$
 $done := \text{false};$
while $\neg done$ **do**
 $done := \text{true};$
 foreach $(\mathbf{u}, \mathbf{v})_\alpha \in E \setminus E'$ **do**
 if $\mathbf{u} \in V' \wedge \text{result}[\mathbf{v}] = \text{SUCC} \wedge \neg \text{Conflict}(\alpha, \mathbf{u})$ **then**
 $V' := V' \cup \{\mathbf{v}\}; E' := E' \cup \{(\mathbf{u}, \mathbf{v})_\alpha\};$
 $done := \text{false}$
return $G' = (V', E');$

Proof. The proof easily follows by construction. \square

We say that a path ϖ in G is *successful* if $\text{result}[\mathbf{u}] = \text{SUCC}$ for every node \mathbf{u} in ϖ , otherwise ϖ is *unsuccessful*. A node \mathbf{u} is *root-successful* if it is reachable from $\text{root}[G]$ via a successful path, otherwise it is *root-unsuccessful*. The next step of the algorithm computes the sub-graph of G that only includes the root-successful vertices. This computation is accomplished by the SuccessGraph function.

Proposition 6.2. *Let $G' = (E', V')$ be generated by SuccessGraph (G). Then $\mathbf{u} \in V'$ if and only if \mathbf{u} is root-successful in G .* \square

The final step of the algorithm synthesizes the filter out of the success graph, in case this is not empty. Let $G' = \text{SuccessGraph}(G)$, π be the underlying set of principals, and $F^{Alg}[\Phi_{\Gamma \triangleright C}^\zeta] = \text{ExtractFilter}_\pi(\text{root}[G], \emptyset, G')$. A complete proof of the next theorem is given in [24].

Theorem 6.1 (Soundness and maximality). *Let $\Gamma \triangleright C$ be a π -composition. Then $\Gamma \triangleright F^{Alg}[\Phi_{\Gamma \triangleright C}^\zeta](C)$ is such that*

- $F^{Alg}[\Phi_{\Gamma \triangleright C}^\zeta](C) \in \mathcal{NT}_{\Gamma, \zeta}$
- $F^{Alg}[\Phi_{\Gamma \triangleright C}^\zeta](C) \downarrow$.

Also, if a filter F fixes $\Gamma \triangleright C$ and is relevant for $\Gamma \triangleright C$, then $F \leq F^{Alg}[\Phi_{\Gamma \triangleright C}^\zeta]$.

Proof. The proof follows by construction and by Propositions 6.1 and 6.2. \square

Function $\text{ExtractFilter}_\pi(\mathbf{u}, U, G)$

Input: $G = (V, E)$ a success graph. $\mathbf{u} \in V, U \subseteq V$

Output: F , an π -composite filter

$F[p] := \mathbf{0}$ for all $p \in \pi$;

if $\text{state}[\mathbf{u}] \checkmark$ **then**

\lfloor **return** F ;

if $\mathbf{u} \in U$ **then**

\lfloor $\text{rec}[\mathbf{u}] := \text{true}$; **return** $(X_{\mathbf{u}}, \dots, X_{\mathbf{u}})$;

foreach $(\alpha, \mathbf{v}) \in \text{Adj}[\mathbf{u}]$ **do**

$F_{\mathbf{v}} := \text{ExtractFilter}_\pi(\mathbf{v}, U \cup \{\mathbf{u}\}, G)$;

foreach $p \in \pi$ **do**

if $\alpha = \{a_{p \rightarrow _}\}$ **then**

$F[p] := F[p] \times \bar{a}_{p \rightarrow _}.F_{\mathbf{v}}[p]$;

else if $\alpha = \{a_{_ \rightarrow p}\}$ **then**

$F[p] := F[p] \times a_{_ \rightarrow p}.F_{\mathbf{v}}[p]$;

else

$\lfloor F[p] := F[p] \times F_{\mathbf{v}}[p]$;

if $\text{rec}[\mathbf{u}] = \text{true}$ **then**

foreach $p \in \pi : X_{\mathbf{u}} \in \text{fv}(F[p])$ **do**

$\lfloor F[p] := \text{rec}(X_{\mathbf{u}}) F[p]$;

return F ;

7. A case study

In this section we address the problem of formalizing the interactions among an authentication service (AS), a web service (WS) and a user. First, we introduce the system description and discuss how our formalism copes with it. Then, we show that the non-interference framework presented in Section 3 can be used to prevent the developer to introduce architecture flaws while projecting the WS. The advantage of such a formal approach is shown to be twofold: first, it gives an unambiguous method to describe a possibly complex software architecture consisting of several cooperating services. Second, it allows one to decide if a model satisfies certain security properties in a completely algorithmic way.

7.1. System definition

Authentication services are web services that allow web applications to partially or totally avoid managing users' accounts. The advantages of using authentication services are: the possibility for the user to own just one account to access

1
2
3
4
5
6
7
8
9
10 several independent web services and the simplifications of the design of the web
11 service since the storage and managing of user accounts is left to a specialised ap-
12 plication. Most of the implementations of such a mechanism relies on a project
13 called *OpenID*. A user that wishes to obtain an OpenID, just needs to register with
14 one of the providers that support the project: Google, Yahoo! and Wordpress are
15 probably the most common. Many web applications and services that require au-
16 thentication, at the login phase, redirect the user to one of the providers handling
17 OpenID. Then, the authentication server recognises the user and sends back both
18 to the user and to the web application that the authentication has been carried out
19 successfully. This process is illustrated in Figure 1 and consists of nine steps:

- 22 1. The WA offers the end user a set of login options
- 23 2. The user selects to log-in with a third party account, choosing among OpenID-
24 compatible providers (e.g., Google, facebook and Yahoo!)
- 25 3. The WA sends a *discovery* request to the selected provider in order to obtain
26 a login web service endpoint
- 27 4. The provider answers with a eXtensible Resource Descriptor Sequence (XRDS)
28 document that contains the address(es) of the login web service endpoint
- 29 5. The WA contacts the login endpoint to send a request for authentication ser-
30 vice
- 31 6. The user is redirected to the chosen provider where the authentication step
32 can be done
- 33 7. The user provides login information and, upon success, receives a confirma-
34 tion and is again redirected to the web application
- 35 8. The AS sends to the WA a token which confirms user's identity
- 36 9. Finally, the WA is able to exchange confidential information to the user

37
38
39
40
41 Note that, with respect to the previous examples, in this case the confidentiality is
42 a requirement of the web application instead of the user.

43 44 45 7.2. *An ideal model*

46 We consider an ideal model of the federated authentication system based on
47 OpenID, i.e., in the first step we do not consider the possibility of errors in the WA
48 development. The security level lattice is $\Sigma = \{L, h_1, h_2, H\}$ and Figure 2 shows
49 relation \preceq . Let us assume that a WA is able to provide certain types of information
50 that can be classified according to three security levels: L, h_1, h_2 . Unknown users'
51 principals are associated with the lowest security level L , while authenticated users
52 are associated with one of the higher security levels h_1 and h_2 . The web application
53 security level is the highest: H . Observe that, the security levels are assigned
54 according to the WA confidentiality requirements. We first give the definition of the
55
56
57
58

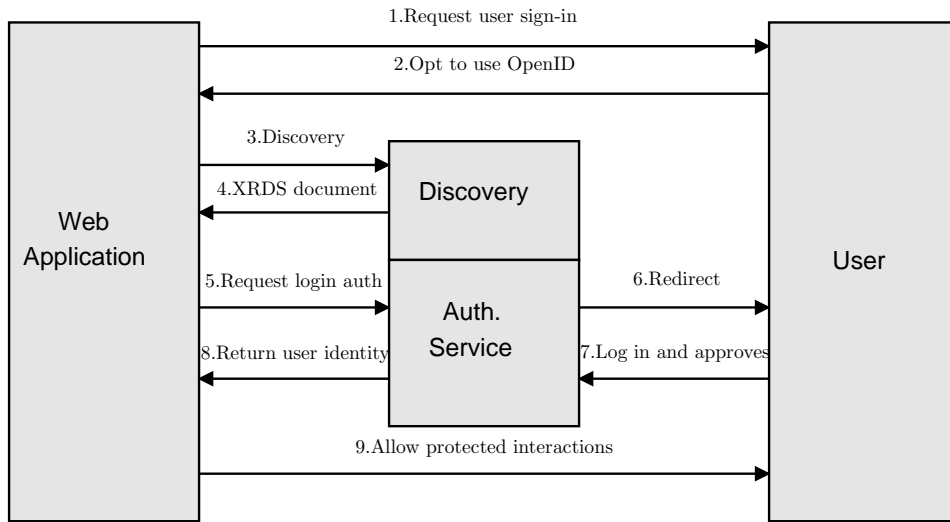


Figure 1: OpenID login authentication for web applications.




Figure 2: Lattice $\langle \Sigma, \preceq \rangle$ used in the Example of Section 7.

$$\begin{aligned}\sigma_D &= \text{DiscReq}@x.\overline{\text{DiscAS}}@x.1 \\ \sigma_{AS} &= \text{Req}_C@y.\overline{\text{Redirect}}@C.\text{GetAccount}@C. \\ &\quad (\overline{\text{Ok}}@C.\overline{\text{Token}}@y.1 \mid_{\emptyset \oplus \emptyset} \overline{\text{Fail}}@C.\overline{\text{NoToken}}@y.1)\end{aligned}$$

Table 8: Contracts executed by the Authentication Service and the Discovery Service in the model for federated login system.

contracts executed by the discovery (σ_D) and authentication (σ_{AS}) web services. Their definitions in terms of the calculus are given in Table 8. Process σ_D is rather simple, since it simply returns a WA endpoint ($\overline{\text{DiscAS}}$) upon a request DiscReq . Observe that, endpoint AS is going to be used as a principal name by the WA's contract. Since in our calculus we do not explicitly model the variable passing between synchronisations, the action synchronising with $\overline{\text{DiscAS}}@x$ is part of an external choice with all the possible values of the output (see [16]). To clarify, consider the definition of σ_{AS} in Table 8: note that the first input Req_C contains the label of principal C that is contacted by the following output action. Therefore, a complete definition of the AS process would be:

$$\sigma'_{AS} = \sigma_{ASC_1} + \sigma_{ASC_2} + \dots + \sigma_{ASC_n},$$

where σ_{ASC_i} is the definition of σ_{AS} of Table 8 rewritten for $C = C_i$, and C_1, \dots, C_n are all the possible (finite) instances of user principal.

However, as illustrated by the definition of σ_{AS} in Table 8, for the sake of a compact notation, we omit to specify all the possible external choices. The AS process, after receiving a request to authenticate user C (Req_C) redirects C to the login service and waits for the account information (Account). If the authentication is successful, then C is informed by receiving an Ok message and the authentication token is sent to the WA. Otherwise, both are informed that the authentication process failed.

Table 9 formalises the definition of contract σ_{WA} performed by principal WA . Its behaviour strictly resembles the process depicted by Figure 1. The first line of the definition σ_{WA} corresponds to steps 1 and 2 of the authentication procedure. User's principal C can decide between an internal authentication procedure or one based on OpenID . In the first case, the WS directly accepts the account information (Account) and based on them, it decides if the user is not recognised (sending of message $\overline{\text{NoAuth}}$, if it can access information with security level h_1 (and hence the process continues as σ_{Prof1}) or with security level h_2 (and hence the process continues as σ_{Prof2}). The offered services depends on the user authentication and

$$\begin{aligned} \sigma_{WA} &= \text{ConfReq}@x.\overline{\text{Choices}}@x.(\text{Internal}@x.\sigma_{AUT} + \text{OpenID}@x.\sigma_{OID}) \\ \sigma_{AUT} &= \text{Account}@x.(\overline{\text{Passed}}@x.(\sigma_{Prof_1} \lfloor x:h_1 \oplus x:h_2 \rfloor \sigma_{Prof_2}) \\ &\quad \lfloor \emptyset \oplus \emptyset \rfloor \overline{\text{NoAuth}}@x.\sigma_{Ser_3}) \\ \sigma_{OID} &= \overline{\text{DiscReq}}@D.\overline{\text{Disc}}_{AS}@D.\overline{\text{Req}}_x@AS.(\text{Token}@AS. \\ &\quad (\sigma_{Prof_1} \lfloor x:h_1 \oplus x:h_2 \rfloor \sigma_{Prof_2}) + \text{NoToken}@AS.\sigma_{Ser_3}) \\ \sigma_{Prof_1} &= \sigma_{Ser_1} + \sigma_{NoSer_2} + \sigma_{Ser_3} \\ \sigma_{Prof_2} &= \sigma_{NoSer_1} + \sigma_{Ser_2} + \sigma_{Ser_3} \\ \sigma_{Ser_i} &= \text{Op}_i@x.\overline{\text{Reply}}_i@x.1 \quad i = 1, 2, 3 \\ \sigma_{NoSer_i} &= \text{Op}_i@x.\overline{\text{Deny}}_i@x.1 \quad i = 1, 2 \end{aligned}$$

Table 9: Contract executed by the Web Application.

$$\begin{aligned} \sigma_C &= \overline{\text{ConfReq}}@WA.\overline{\text{Choices}}@WA. \\ &\quad (\overline{\text{Internal}}@WA.\sigma_{IA} \lfloor \emptyset \oplus \emptyset \rfloor \overline{\text{OpenID}}@x.\sigma_{OID}) \\ \sigma_{IA} &= \overline{\text{Account}}@WA.(\overline{\text{Passed}}@WA.\sigma_{Act} \\ &\quad + \text{NoAuth}@x.\overline{\text{Op}}_3@x.\overline{\text{Reply}}_3@x.1) \\ \sigma_{OID} &= \overline{\text{Redirect}}@z.\overline{\text{GetAccount}}@z.(\text{Ok}@z.\sigma_{Act} \\ &\quad + \text{Fail}@z.\overline{\text{Op}}_3@x.\overline{\text{Reply}}_3@x.1) \\ \sigma_{Act} &= \overline{\text{Op}}_3@WA.\overline{\text{Reply}}@WA.3.1 \lfloor \emptyset \oplus \emptyset \rfloor (\overline{\text{Op}}_2@WA.(\overline{\text{Reply}}_2@WA.1 + \\ &\quad \text{Deny}_2@WA.1) \lfloor \emptyset \oplus \emptyset \rfloor \overline{\text{Op}}_1@WA.(\overline{\text{Reply}}_1@WA.1 + \text{Deny}_1@WA.1)) \end{aligned}$$

Table 10: Contract executed by the User.

are σ_{Ser_1} for users with level h_1 , σ_{Ser_2} for users with level h_2 , whereas σ_{Ser_3} is available for both (and also non-registered users). Observe that upon a successful authentication, user's security level is upgraded by the internal choice operator. Finally, Table 10 gives the definition of the user process σ_C .

Initially, user's contract σ_C synchronises with WA and chooses one of the two

available authentication methods. In both cases, it sends the account information ($\overline{\text{Account}}$ and $\overline{\text{Getaccount}}$) and then proceeds with a service request (σ_{Act}). Note that, if the authentication method is OpenID then a `Redirect` input allows the principal to identify the authentication service endpoint. Finally, the user performs an internal choice among the three possible requests Op_i , $i = 1, 2, 3$. Observe that, in order to satisfy the compliance property, we consider the possibilities that in case operations 1 or 2 are requested, then WA may answer with a denial of access according to the user profile. It can be proved that the service composition $M = WA[\sigma_{WA}] \parallel C[\sigma_C] \parallel D[\sigma_D] \parallel AS[\sigma_{AS}]$ with the type system $\Gamma = \{WA : H, C : L, D : L, AS : L\}$ is compliant. However, if we consider the definitions of Table 9 omitting σ_{NoSer_i} , (intuitively, the WA simply does not accept requests from unauthorised users), then configuration M would be non-compliant. In fact, a user that cannot ask for service i ($i = 1, 2$) could anyway perform an internal choice for it (see the definition of σ_{Act}) and hence cause a deadlock in its execution. It is worthwhile pointing out that this corresponds to what is known to be a good software engineering practice, i.e., handling with correct procedures (maybe of denial) all the possible requests from the users. Finally, $M \in \mathcal{NT}_{\Gamma, L}$ since once the security level of the customer is upgraded to h_1 or h_2 , all the synchronisations between WA and C have the same level of security.

7.3. Compliance and non-interference analysis in an insecure WA

Web services are vulnerable to many attacks that can overcome the security measures designed by the developer. This is mainly due to the fact that most of the web service endpoints are completely exposed in a untrusted environment, i.e., the Internet. In this part we consider the possibility that a user, for malicious purposes, is able to send to the WA an ill-formed request that causes a failure in the procedure that is being carried out. In practice, this corresponds to a badly formed XML document or to a SQL injection attack. It is well-known that correct security implementations require that the WA reveals, as a consequence of a failure, the lowest amount of information about its activity so that the attacker cannot infer any information about the confidential activities that are being processed (see, e.g., [25]). We show how our non-interference framework is able to formally detect if a WA is designed according to this security principle.

The revised version of the model is depicted in Table 11 where the definition of the AS and D contracts are omitted since they are identical to those of Table 8.

In this model, we assume that the WA may receive the malicious request at two epochs: just before the user chooses the required service (before synchronisations on Op_i) and just before it receives the answer for that request. In both cases, the WA procedure aborts. Although, a WA that suffers this problem clearly contains flaws in its design (probably missing checks on input), these are very common

$$\begin{aligned}
M &= AS[\sigma_{AS}] \parallel C[\sigma_C] \parallel X[\sigma_X] \parallel D[\sigma_D] \parallel WA[\sigma_{WA}] \\
\sigma_{WA} &= \text{ConfReq}@x.\overline{\text{Choices}}@x.(\text{Internal}@x.\sigma_{AUT} + \text{OpenID}@x.\sigma_{OID}) \\
\sigma_{AUT} &= \text{Account}@x.(\overline{\text{Passed}}@x.(\sigma_{Prof_1} \lfloor x:h_1 \oplus x:h_2 \rfloor \sigma_{Prof_2}) \\
&\quad \lfloor \emptyset \oplus \emptyset \rfloor \overline{\text{NoAuth}}@x.\sigma_{Prof_3}) \\
\sigma_{OID} &= \overline{\text{DiscReq}}@D.\overline{\text{Disc}}_{AS}@D.\overline{\text{Req}}_x@AS.(\text{Token}@AS. \\
&\quad (\sigma_{Prof_1} \lfloor x:h_1 \oplus x:h_2 \rfloor \sigma_{Prof_2}) + \text{NoToken}@AS.\sigma_{Prof_3}) \\
\sigma_{Prof_1} &= \sigma_{Ser_1} + \sigma_{NoSer_2} + \sigma_{Ser_3} + \text{IllMsg}@y.\overline{\text{Msg}}_0@y.\overline{\text{Close}}@x.1 \\
\sigma_{Prof_2} &= \sigma_{NoSer_1} + \sigma_{Ser_2} + \sigma_{Ser_3} + \text{IllMsg}@y.\overline{\text{Msg}}_0@y.\overline{\text{Close}}@x.1 \\
\sigma_{Prof_3} &= \sigma_{Ser_3} + \text{IllMsg}@y.\overline{\text{Msg}}_0@y.\overline{\text{Close}}@x.1 \\
\sigma_{Ser_i} &= \text{Op}_i@x.(\overline{\text{Reply}}_i@x.1 + \text{IllMsg}@y.\overline{\text{Msg}}_i@y.\overline{\text{Close}}@x.1) \quad i = 1, 2, 3 \\
\sigma_{NoSer_i} &= \text{Op}_i@x.\overline{\text{Deny}}_i@x.1 \quad i = 1, 2 \\
\sigma_C &= \overline{\text{ConfReq}}@WA.\overline{\text{Choices}}@WA. \\
&\quad (\overline{\text{Internal}}@WA.\sigma_{IA} \lfloor \emptyset \oplus \emptyset \rfloor \overline{\text{OpenID}}@x.\sigma_{OID}) \\
\sigma_{IA} &= \overline{\text{Account}}@WA.(\overline{\text{Passed}}@WA.\sigma_{Act} \\
&\quad + \text{NoAuth}@x.\overline{\text{Op}}_3@x.\overline{\text{Reply}}_3@x.1) \\
\sigma_{OID} &= \overline{\text{Redirect}}@z.\overline{\text{GetAccount}}@z.(\text{Ok}@z.\sigma_{Act} \\
&\quad + \text{Fail}@z.\overline{\text{Op}}_3@x.\overline{\text{Reply}}_3@x.1) \\
\sigma_{Act} &= (\overline{\text{Op}}_3@WA.(\overline{\text{Reply}}_3@WA.1 + \overline{\text{Close}}@WA.1) \lfloor \emptyset \oplus \emptyset \rfloor \\
&\quad (\overline{\text{Op}}_2@WA.(\overline{\text{Reply}}_2@WA.1 + \overline{\text{Deny}}_2@WA.1 + \overline{\text{Close}}@WA.1) \\
&\quad \lfloor \emptyset \oplus \emptyset \rfloor \overline{\text{Op}}_1@WA.(\overline{\text{Reply}}_1@WA.1 + \overline{\text{Deny}}_1@WA.1 + \\
&\quad \overline{\text{Close}}@WA.1))) + \overline{\text{Close}}@WA.1 \\
\sigma_X &= \overline{\text{IllMsg}}@WA.(\overline{\text{Msg}}_0@WA.1 + \overline{\text{Msg}}_1@WA.1 \\
&\quad + \overline{\text{Msg}}_2@WA.1 + \overline{\text{Msg}}_3@WA.1) + 1
\end{aligned}$$

Table 11: Insecure contracts for WA and C.

1
2
3
4
5
6
7
8
9
10 and it is dangerous to assume that a service is immune from them. Therefore, the
11 developer should take care that, when this type of attack happens, the system does
12 not violate its confidentiality requirements and remains compliant. Contract σ_X
13 executed by attacker's principal X tries to send to the WA a malicious request (e.g.,
14 SQL injection) by synchronising on `I11Msg` and if this happens, then it catches one
15 of the possible error messages returned. Observe that, before an operation request
16 $\text{Op}_i, i = 1, 2, 3$, the WA returns an error message Msg_0 , whereas after that it returns
17 an error message Msg_i that depends on the service being processed.

18
19 We now show that the contract depicted by Table 11 are compliant but does
20 not yield the non-interference property for security level L . The compliance can
21 be straightforwardly derived from the observation that upon an `I11Msg` input the
22 contract stops its activity just after sending an error message to the principal that
23 caused it, and closing the connection with the customer. Symmetrically, user's con-
24 tract always considers the possibility of an abortion of the service by synchronising
25 on `Close`. To prove that $M \notin \mathcal{NT}_{\Gamma, L}$ it suffices to show that in M when X syn-
26 chronises on Msg_1 this is surely preceded by a synchronisation on Op_1 of level h_1 ,
27 therefore the attacker can infer the profile of the user being logged and that it re-
28 quired an operation Op_1 (the same holds for Op_2). This vulnerability is well-known
29 by practitioners, and is usually caused by un-handled exceptions. Conversely, ex-
30 ceptions should be always handled and the principal that causes them should re-
31 ceived a short error message that must not depend on the server status [25]. In
32 fact, in the model of Table 11, if we assume $\text{Msg}_i = \text{Msg}$ it is possible to automat-
33 ically prove that $M \in \mathcal{NT}_{\Gamma, L}$. Although in this case the non-interference analysis
34 leads to an intuitive result, this is not always true. As an instance, consider the
35 model depicted in Table 11 with $\text{Msg}_i = \text{Msg}$ for $i = 0, 1, 2, 3$ but without the
36 possibility to send a malicious message just before Op_i synchronisations. At a first
37 glance, this modification corresponds to an improvement of the system security,
38 however $M \notin \mathcal{NT}_{\Gamma, L}$ holds. The point is that an observer can infer that a high
39 security level synchronisation occurred by observing how the error is handled. In
40 fact, if `Close` synchronisation is done at level L (and hence is observable), then
41 for sure neither Op_1 nor Op_2 synchronisations occurred, otherwise a synchronisa-
42 tion of security level $\{h_1, h_2\}$ has been done. Therefore, the system reveals more
43 information about its confidential activities with respect to the system that can be
44 attacked also just before user's choice of Op_i .
45
46
47
48
49
50
51

52 8. Conclusion

53
54 Some research efforts on model checking web services have already been pro-
55 posed [26, 27, 28, 29]. The most related paper that we are aware of is by Naka-
56 jima [30] who introduces a lattice-based security labelling into BPEL in order to
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10 detect potential insecure information leakage. The paper discusses how both the
11 safety and security aspects can be analyzed in a single framework using the model-
12 checking verification techniques. The main difference with our approach is that
13 the notion of security considered in [30] is built upon a simple lattice-based model
14 for security labels. Instead, our approach deals with more flexible security policies
15 which can be dynamically specified by the service participants. As far as correct-
16 ness is concerned, [30] considers safety properties such as deadlock freedom and
17 specific progress properties. Our model instead deals also with the property of
18 livelock freedom.

19
20 In conclusion, we have developed a formal method for the analysis of both
21 information flow security and compliance of contract service compositions. This
22 is based on the characterization of such properties in terms of modal μ -calculus
23 formulae. This allows us to use a model checker, like the NCSU Concurrency
24 Workbench, in order to simultaneously check non-interference and compliance.
25 An algorithm for adaptable service compositions is also proposed. It computes the
26 greatest relevant filter fixing them.
27
28

29 30 **References**

- 31
32 [1] A. H. Anderson, An introduction to the web services policy language (wspl),
33 in: 5th IEEE International Workshop on Policies for Distributed Systems and
34 Networks (POLICY 2004), IEEE Computer Society, 2004, pp. 189–192.
35
36 [2] K. Bhargavan, C. Fournet, A. D. Gordon, Verifying policy-based web services
37 security, *ACM Trans. Program. Lang. Syst.* 30 (6) (2008) 1–59.
38
39 [3] S. Carpineti, G. Castagna, C. Laneve, L. Padovani, A Formal Account of
40 Contracts for Web Services, in: *Proc. of the International Workshop on Web
41 Services and Formal Methods (WS-FM’06)*, Vol. 4184 of LNCS, Springer,
42 2006, p. 148162.
43
44 [4] G. Castagna, N. Gesbert, L. Padovani, A Theory of Contracts for Web Ser-
45 vices, in: *Proc. of the annual Symposium on Principles of Programming Lan-
46 guages (POPL’08)*, ACM press, 2008, pp. 261–272.
47
48 [5] G. Castagna, N. Gesbert, L. Padovani, A Theory of Contracts for Web
49 Services, *ACM Transactions on Programming Languages and Systems*
50 (TOPLAS) 31 (2009) 53–61.
51
52 [6] C. Laneve, L. Padovani, The *must* Preorder Revisited, in: *Proc. of the In-
53 ternational Conference on Concurrency Theory (CONCUR’07)*, Vol. 4703 of
54 LNCS, Springer, 2007, pp. 212–225.
55
56
57
58

- 1
2
3
4
5
6
7
8
9
10 [7] M. Bravetti, G. Zavattaro, A Foundational Theory of Contracts for Multi-
11 party Service Composition, *Fundamenta Informaticae* 89 (4) (2009) 451–478.
12
13 [8] M. Bravetti, G. Zavattaro, Towards a unifying theory for choreography con-
14 formance and contract compliance, in: *Proc. of 6th International Symposium*
15 *on Software Composition (SC'07)*, Vol. 4829 of LNCS, Springer, 2007, pp.
16 34–50.
17
18 [9] J. A. Goguen, J. Meseguer, Security Policies and Security Models, in: *Proc.*
19 *of the IEEE Symposium on Security and Privacy (SSP'82)*, IEEE Computer
20 Society, 1982, pp. 11–20.
21
22 [10] A. Zakinthinos, E. S. Lee, A General Theory of Security Properties, in: *Proc.*
23 *of the IEEE Symposium on Security and Privacy (SSP'97)*, IEEE Computer
24 Society, 1997, pp. 74–102.
25
26 [11] J. McLean, Security Models and Information Flow, in: *Proc. of the IEEE*
27 *Symposium on Security and Privacy (SSP'90)*, IEEE Computer Society,
28 1990, pp. 180–187.
29
30 [12] R. Focardi, S. Rossi, Information Flow Security in Dynamic Contexts, *Jour-*
31 *nal of Computer Security* 14 (1) (2006) 65–110.
32
33 [13] A. Sabelfeld, A. C. Myers, Language-Based Information-Flow Security,
34 *IEEE Journal on Selected Areas in Communication* 21 (1) (2003) 5–19.
35
36 [14] E. M. Clarke, O. Grumberg, D. A. Peled, *Model checking*, The MIT Press,
37 1999.
38
39 [15] D. Kozen, Results on the Propositional μ -calculus, *Theoretical Computer Sci-*
40 *ence* 27 (1983) 333–354.
41
42 [16] R. Milner, *Communication and Concurrency*, Vol. 92 of Prentice Hall Inter-
43 *national Series in Computer Science*, Prentice Hall, 1989.
44
45 [17] M. Bravetti, G. Zavattaro, Contract Compliance and Choreography Confor-
46 mance in the Presence of Message Queues, in: *Proc. of the International*
47 *Workshop on Web Services and Formal Methods (WS-FM'08)*, Vol. 5387
48 *of LNCS*, Springer, 2008, pp. 37–54.
49
50 [18] M. Tarek, C. Boutrous-Saab, S. Rampacek, Verifying correctness of web ser-
51 vices choreography, in: *ECOWS'06*, 2006, pp. 306–318.
52
53 [19] M. Müller-Olm, Derivation of Characteristic Formulae, *Electronic Notes in*
54 *Theoretical Computer Science* 18.
55
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9
10 [20] A. Mader, Modal μ -calculus, Model Checking, and Gauss Elimination, in:
11 Proc. of International Conference on Tools and Algorithms for Construction
12 and Analysis of Systems (TACAS'95), Vol. 1019 of LNCS, Springer, 1995,
13 pp. 72–88.
14
- 15 [21] G. Bernardi, M. Bugliesi, D. Macedonio, S. Rossi, A Theory of Adaptable
16 Contract-Based Service Composition, in: Proc. of International Symposium
17 on Symbolic and Numeric Algorithms for Scientific Computing, Workshop
18 on Global Computing Models and Technologies (GlobalComp'08), IEEE
19 Computer Society, 2008, pp. 327–334.
20
- 21 [22] T. Basciutti, Model-Checking Web Services, Master's thesis, Department of
22 Computer Science, University Ca' Foscari of Venice (2010).
23
- 24 [23] R. Cleaveland, S. Sims, The NCSU Concurrency Workbench, in: Proc. of In-
25 ternational Conference on Computer Aided Verification (CAV'96), Vol. 1102
26 of LNCS, Springer, 1996, pp. 394–397.
27
- 28 [24] G. Bernardi, A Theory of Adaptable Contract-Based Service Compositions,
29 Master's thesis, Department of Computer Science, University Ca' Foscari of
30 Venice (2009).
31
- 32 [25] A. Singhal, T. Winograd, K. Scarfone, Guide to Secure Web Services, Tech.
33 Rep. 800-95, National Institute of Standards and Technology (2007).
34
- 35 [26] F. Abouzaid, J. Mullins, Model-checking Web Services Orchestrations using
36 BP-calculus, Electronic Notes in Theoretical Computer Science 255 (2009)
37 3–21.
38
- 39 [27] G. Dai, X. Bai, C. Zhao, A Framework for Model Checking Web Service
40 Compositions Based on BPEL4WS, in: Proc. of the IEEE International Con-
41 ference on e-Business Engineering (ICEBE'07), IEEE Computer Society,
42 2007, pp. 165–172.
43
- 44 [28] S. Nakajima, Model-Checking Behavioral Specification of BPEL Applica-
45 tions, Electronic Notes in Theoretical Computer Science 151 (2006) 89–105.
46
- 47 [29] H. Schlingloff, A. Martens, K. Schmidt, Modeling and Model Checking Web
48 Services, Electronic Notes in Theoretical Computer Science 126 (2005) 3–26.
49
- 50 [30] S. Nakajima, Model-Checking of Safety and Security Aspects in Web Ser-
51 vice Flows, in: Proc. of the International Conference of Web Engineering
52 (ICWE'04), Vol. 3140 of LNCS, Springer, 2004, pp. 488–501.
53
54
55
56
57
58
59
60
61
62
63
64
65