

Policy-based reasoning for smart web service interaction^{*}

Marco Alberti¹, Federico Chesani², Marco Gavanelli¹, Evelina Lamma¹,
Paola Mello², Marco Montali², and Paolo Torroni²

¹ ENDIF, Università di Ferrara

{marco.alberti|marco.gavanelli|evelina.lamma}@unife.it

² DEIS, Università di Bologna

{fchesani|pmello|mmontali|ptorroni}@deis.unibo.it

Abstract. We present a vision of smart, goal-oriented web services that reason about other services' policies and evaluate the possibility of future interactions. To achieve our vision, we propose a proof theoretic approach. We assume web services whose interface behaviour is specified in terms of reactive rules. Such rules can be made public, in order for other web services to answer the following question: "is it possible to inter-operate with a given web service and achieve a given goal?" In this article we focus on the underlying reasoning process, and we propose a declarative and operational abductive logic programming-based framework, called WAV^e.

1 Introduction

Service Oriented Computing (SOC) is rapidly emerging as a new programming paradigm, propelled by the wide availability of network infrastructures, such as the Internet, and by the success of its predecessor, Object Oriented programming paradigm. Web service-based technologies are an implementation of SOC, aimed at overcoming the intrinsic difficulties of integrating different platforms, operating systems, languages, etc., into new applications. It is then in the spirit of SOC to take off-the-shelf solutions, like web services, and compose them into new applications. Service composition is very attractive for its support to rapid prototyping and possibility to create complex applications from simple elements. It is the philosophy followed, e.g., by BPEL [1]: composing new applications through existing web services.

On the upside, the recent popularity of these new technologies developed into a growing presence of web services, made available through the Internet, and we can foresee a steady increase of such services also for the near future. On the downside, the lifetime of software developed with the classical methodologies of composition of existing services at design-time gets shorter and shorter. It quickly

^{*} We thank the anonymous referees for their valuable feedback and pointers. This work has been partially supported by the MIUR PRIN 2005 project *Specification and verification of agent interaction protocols*.

becomes a suboptimal choice, blind to the exploitation of new opportunities. In highly competitive markets, this can be a severe drawback.

If we adopt the SOC programming paradigm, how to exploit the potential of a growing base of web services becomes one of our strategic issue. In a domain in which being more competitive means knowing more and using all available information at best, how shall we cope with the proliferation of new services? How shall we decide to use a web service rather than another one? when new ones becomes available, shall we go for them? are there new opportunities that were not there before? It is a necessary, never-ending, heavy and thus potentially very costly decision process, but it could also be very rewarding, if we had the proper tools.

A partial answer to these questions is given by service discovery. As new services become available, they are published, for instance by registration on some yellow-pages server; existing services can then become aware of the new ones and exploit them. This solves part of the problem: as through discovery we only know that there are some some services, which possibly follow some standards, but understanding whether interacting with them will be profitable or detrimental, is far from being a trivial question. For one, it is not possible to think to try and invoke all newly discovered services and analyze the results. Beside being highly error-prone, such a method would require expensive rollbacks that are often unaffordable at run-time. Thus, alternative approaches have to be developed. This is what we intend to address in this article.

The focus of this article is the following problem: how to dynamically understand if two web services can inter-operate, without them having a-priori knowledge of each other's capabilities, but by reasoning about policies exchanged at run-time.

We present a vision of smart, goal-oriented web services that reason about other services' specifications, with the aim to separate out those that can lead to a fruitful interaction, without resorting to trial and error. We envisage a two-phase discovery activity on the side of web services. First, web services collect information about other web services, and try and understand by reasoning which ones can lead to a fruitful interaction. This activity is carried out off-line, beforehand. Then they use the available information to interact with each other. It is the same philosophy of search engines: before, collect information through web spiders, then use it when requested by the user.

In this article we focus on the reasoning involved in the off-line phase, assuming that a new web service has been found, and we must decide about the possibility to interact with it. We assume that each web service publishes, alongside with its WSDL, its *interface behaviour specifications*. By reasoning on the information available about other web services' interface behaviour, each web service can verify which goals can be reached by interaction.

To achieve our vision, we propose a proof theoretic approach, based on computational logic – in fact, on abductive logic programming. In particular, we formalise policies for web services in a declarative language which is a modification of the SCIFF language originally defined in the context of the UE IST-2001-




Fig. 1. The architecture of WAV^e

32530 project, to specify and verify social-level agent interaction. In this new language, policies can be defined by way of *social integrity constraints (ICs)*: a sort of reactive rules used to generate and reason about expectations about possible evolutions of a given interaction setting. Based on the *SCIFF* framework we propose a new declarative semantics and a new proof-procedure that combines forward, reactive reasoning with backward, goal-oriented reasoning, and is tailored to the discovery activity's off-line phase's verification problem. We have called this new framework WAV^e (Web-service Abductive Verification).

We start by showing the abstract architecture of WAV^e . In Sect. 3 we introduce a running on-line shopping scenario. In Sect. 4, we briefly introduce the language used in the framework, and in Sect. 5 we show how the scenario can be modeled in WAV^e in terms of *ICs*. Sect. 6 presents the declarative and operational semantics of WAV^e , and Sect. 7 proposes the application of WAV^e to the verification problem in the reference scenario. A brief discussion follows.

2 The Architecture of WAV^e

Fig. 1 depicts our general reference architecture. Arrows indicate the flow of policies between web services. The layered architecture of a web service, e.g. ws , has WAV^e at the top of the stack, performing reasoning based on its own knowledge and on the policies obtained from other web services, e.g. ws' . The functionalities of the various elements of the knowledge will be explained in Sect. 4. For the moment, we say that policies are identified with the IC_{ws} component. The architecture is symmetric. We represented with thick borders the modules involved in the operations carried out by ws , and its output. In order for ws' to pass

$\mathcal{IC}_{ws'}$ on to ws (and vice versa), a Rule Interchange Format (RIF) is adopted. One possibility for such a RIF could be RuleML [2]. Finally, as a result of the reasoning activity, ws produces an answer \mathcal{C} to the question: “is it possible to inter-operate with ws' and achieve goal \mathcal{G}_{ws} ?”

Fig. 1 does not show control elements, but only information flows. We assume that suitable interaction protocols are defined to control the flow of information (e.g. policies) between the web services. In particular, in a more comprehensive setting, ws and ws' could negotiate the exchange of policies in an incremental way, or could use the result \mathcal{C} of this reasoning activity to perform the second, on-line phase of service interaction we mentioned in the introduction. All this is outside of this picture, and of this article’s scope.

3 The *alice* & *eShop* Scenario

This scenario is inspired to the one described by the Working Group on Rule Interchange Format [3]. A similar scenario is also in [4]. We consider two entities, which we call *alice* and *eShop*.³ *eShop* is a web service which sells devices. *alice* is another web service which instead needs to obtain a device, and which is considering buying it from *eShop*. *alice* and *eShop* describe their behaviour concerning sales/payment/... of items through policies, specified as rules, which they publish using some RIF.

Before *alice* buys any item from *eShop*, *alice* checks whether her policies and *eShop*’s policies are compatible, i.e., if they allow a successful transaction regarding the sales. During this process, it turns out that *eShop* accepts credit card payments, besides other payment methods, and that *alice* can only pay by credit card; in this case, in order to proceed with the payment, she requires evidence of the shop’s membership to some trusted “Better Business Bureau” (*BBB*) association. We assume that the shop is able and ready to provide such a piece of evidence. We can thus define *eShop*’s and *alice*’s policies as follows:

- (shop1) if a customer wishes to buy an item, then (s)he should pay it either by credit card, or by cash, or by cheque;
- (shop2) if a customer wishes to buy an item, and (s)he has paid it either by credit card, or by cash, or by cheque, then *eShop* will deliver the item;
- (shop3) if a customer wishes to receive a certificate about *eShop*’s membership to the *BBB*, then the shop will send it;
- (alice1) if a shop requires that *alice* pays by credit card, *alice* expects that the shop provides evidence of its membership to the *BBB*;
- (alice2) if a shop requires that *alice* pays by credit card, and the shop has provided evidence of its membership to the *BBB*, then *alice* will pay by credit card;

In this example, we can identify two kinds of policy rules. *shop1* and *alice1* express requirements, i.e., what is needed in order to proceed with accomplishing

³ In this simplified scenario, we identify *alice* and *eShop* with their representative software counterparts which will carry out transactions on their behalf.

some request. *shop2*, *shop3* and *alice2* represent the effect of requests, i.e., they tell what has to be expected if some conditions hold and some request is received.

Using this scenario, we want to demonstrate the possibility of reaching an agreement through rules exchange. Besides, we want to show how policies support backward and forward reasoning, in the following way. Backward, pro-active reasoning starts from goals to produce (expectations about) actions or events that should be generated in order to achieve the goals. Forward, reactive reasoning starts from events and is used to generate (expectations about) actions that represent reactions to such events.

In this scenario, the goal of *alice* interacting with *eShop* is to obtain an item from *eShop*. Actions are all the messages exchanged between the two web services.

The steps that we envisage are as follows:

1. *alice* wants to obtain a device. She knows that she can have it if *eShop* delivers it to her. Thus, she sends *eShop* a request, by which she wants to know *eShop*'s policies regarding the delivery of that device;
2. *eShop* considers *alice*'s request, and composes a set of rules related to *alice*'s request (its policies), possibly deriving/filtering them from a larger set. In this example, the set contains *shop1*, *shop2*, and *shop3*. Once such a set is put together, *eShop* communicates it to *alice*;
3. *alice* reasons on (1) her goal, (2) her own policies (*alice1* and *alice2*), and (3) *eShop*'s policies. Two are the possible outcomes:
 - either *alice* infers that she and *eShop* can have a successful transaction that satisfies each other's policies and that achieves her goal,
 - or *alice* infers that there is no such a possibility.
4. possibly, at a later point, *alice* and *eShop* may engage in a transaction which (hopefully) makes *alice* achieve her goal.

Points (1) through (3) represent the off-line phase of service discovery/interaction, whereas point (4) represent the actual transaction occurring between *alice* and *eShop*. The reasoning involved in (3) is the subject of this article.

4 The WAV^e Framework

In WAV^e, the observable behaviour of the web services is represented by *events*. Since we focus on (explicit) interaction between web services, events always represent exchanged messages.

WAV^e considers two types of events: those that one can control and those that one cannot. Typically, from the standpoint of a web service *ws*, an event such as a message generated by *ws* himself will fall into the first category, a message that *ws* is expecting from another fellow web service *ws'* will fall instead into the second one. We use two different functors to keep these two categories of messages distinct from each other. Atoms denoted by functor **H** will stand for events that a web service expects to be producing itself; atoms denoted by functor **E** will stand for events that a web service is expecting, and over which

it does not have any control. Since WAV^e is about reasoning on possible future courses of events, both kinds of events represent *hypotheses* that a web service can make on possibly happening events. The notation is: $\mathbf{H}(ws, ws', M, T)$, for messages (M) that a web service ws is expecting to send to ws' at time T , and $\mathbf{E}(ws', ws, M, T)$ for messages (M) expected by ws from ws' for time T .

Web service specifications in WAV^e are relations among expected events, expressed by an Abductive Logic Program (ALP). In general, an ALP [5] is a triplet $\langle P, A, IC \rangle$, where P is a logic program, A is a set of predicates named *abducibles*, and IC is a set of integrity constraints. Roughly speaking, the role of P is to define predicates, the role of A is to fill-in the parts of P which are unknown, and the role of IC is to constrain the ways elements of A are hypothesised, or “abduced”. Reasoning in abductive logic programming is usually goal-directed (being G a goal), and it accounts to finding a set of abduced hypotheses Δ built from predicates in A such that $P \cup \Delta \models G$ and $P \cup \Delta \models IC$. In the past, a number of proof-procedures have been proposed to compute Δ (see Kakas and Mancarella [6], Fung and Kowalski [7], Denecker and De Schreye [8], etc.).

Definition 1 (Web service interface behaviour specification). *Given a web service ws , its web service interface behaviour specification \mathcal{P}_{ws} is an ALP, represented by the triplet*

$$\mathcal{P}_{ws} \equiv \langle \mathcal{KB}_{ws}, \mathcal{E}_{ws}, \mathcal{IC}_{ws} \rangle$$

where:

- \mathcal{KB}_{ws} is ws 's Knowledge Base,
- \mathcal{E}_{ws} is ws 's set of abducible predicates, and
- \mathcal{IC}_{ws} is ws 's set of Integrity Constraints.

\mathcal{KB}_{ws} is a set of clauses which declaratively specifies pieces of knowledge of the web service. Note that the body of \mathcal{KB}_{ws} 's clauses may contain \mathbf{E} expectations about the behaviour of the web services, as defined above. \mathcal{KB}_{ws} 's syntax is summarised in Eq. (1).

$$\begin{aligned} \mathcal{KB}_{ws} &::= [\textit{Clause}]^* \\ \textit{Clause} &::= \textit{Atom} \leftarrow \textit{Cond} \\ \textit{Cond} &::= \textit{ExtLiteral} [\wedge \textit{ExtLiteral}]^* \\ \textit{ExtLiteral} &::= \textit{Atom} \mid \textit{true} \mid \textit{Expect} \mid \textit{Constr} \\ \textit{Expect} &::= \mathbf{E}(\textit{Atom}, \textit{Atom}, \textit{Atom}, \textit{Atom}) \end{aligned} \tag{1}$$

\mathcal{E}_{ws} includes \mathbf{E} expectations, \mathbf{H} events, and predicates not defined in \mathcal{KB}_{ws} .

$$\begin{aligned} \mathcal{IC}_{ws} &::= [\textit{IC}]^* \\ \textit{IC} &::= \textit{Body} \rightarrow \textit{Head} \\ \textit{Body} &::= (\textit{Event} \mid \textit{Expect}) [\wedge \textit{BodyLit}]^* \\ \textit{BodyLit} &::= \textit{Event} \mid \textit{Expect} \mid \textit{Atom} \mid \textit{Constr} \\ \textit{Head} &::= \textit{Disjunct} [\vee \textit{Disjunct}]^* \mid \textit{false} \\ \textit{Disjunct} &::= (\textit{Expect} \mid \textit{Event} \mid \textit{Constr}) [\wedge (\textit{Expect} \mid \textit{Event} \mid \textit{Constr})]^* \\ \textit{Expect} &::= \mathbf{E}(\textit{Atom}, \textit{Atom}, \textit{Atom}, \textit{Atom}) \\ \textit{Event} &::= \mathbf{H}(\textit{Atom}, \textit{Atom}, \textit{Atom}, \textit{Atom}) \end{aligned} \tag{2}$$

Integrity Constraints (ICs) are forward rules, of the form $Body \rightarrow Head$ (Eq. (2)). The *Body* of *ICs* is a conjunction of literals and expected events; the *Head* instead is a disjunction of conjunctions of expectations, events and literals, or *false*. The syntax of \mathcal{IC}_{w_s} is a modification of that defined in [9]. In particular, unlike *SCIFF*, WAV^e treats **H** events as abducible predicates, and as such it allows them to occur in the *Head* of integrity constraints; however, this initial version of WAV^e does not yet accommodate negative expectations nor negation (\neg). We intend to consider these two features in future extensions of WAV^e .

Intuitively, the operational behaviour of integrity constraints is similar to forward rules: whenever the body becomes true, the head is also made true.

5 Modeling in WAV^e

In this section, we demonstrate web service policy modelling in WAV^e by showing the specification of *alice* and *eShop*. The first three rules represent *eShop*'s policies.

$$\begin{aligned}
& \mathbf{E}(eShop, alice, deliver(Item), T_s) \\
\rightarrow & \mathbf{E}(alice, eShop, pay(Item, cc), T_{cc}) \wedge T_{cc} < T_s \\
& \vee \mathbf{E}(alice, eShop, pay(Item, cash), T_{ca}) \wedge T_{ca} < T_s \\
& \vee \mathbf{E}(alice, eShop, pay(Item, cheque), T_{ch}) \wedge T_{ch} < T_s
\end{aligned} \tag{shop1}$$

IC shop1 says that, if *alice* expects *eShop* to deliver an *Item*, then *eShop* expects *alice* to pay by credit card, cash, or cheque, and that payment must be made before delivery.⁴ In that case, the abducibles in the head are expectations, because they represent actions that should be performed by *alice*: from *eShop*'s viewpoint, they can only be expected.

$$\begin{aligned}
& \mathbf{E}(eShop, alice, deliver(Item), T_s) \\
& \wedge \mathbf{H}(alice, eShop, pay(Item, How), T_p) \wedge T_p < T_s \\
& \wedge How::[cc, cash, cheque] \\
\rightarrow & \mathbf{H}(eShop, alice, deliver(Item), T_s).
\end{aligned} \tag{shop2}$$

IC shop2 says that, if *alice* expects *eShop* to deliver the *Item*, and *alice* has paid for it, then *eShop* will actually deliver it to *alice*. In that case, the abducible in the head is an event, because it represents an action that *eShop* should perform, and therefore it assumes that it will indeed happen (since it is its own responsibility).

$$\begin{aligned}
& \mathbf{E}(eShop, alice, give_guarantee, T_g) \\
\rightarrow & \mathbf{H}(eShop, alice, give_guarantee, T_g).
\end{aligned} \tag{shop3}$$

IC shop3 says that if *alice* expects to receive a guarantee, then *eShop* will send it. The following two rules represent *alice*'s policies.

$$\begin{aligned}
& \mathbf{E}(alice, eShop, pay(Item, cc), T_p) \\
\rightarrow & \mathbf{E}(eShop, alice, give_guarantee, T_g) \wedge T_g < T_p.
\end{aligned} \tag{alice1}$$

⁴ The alternative in the head could alternatively be expressed via a variable with domain: $\mathbf{E}(alice, eShop, pay(Item, How), T) \wedge How::[cc, cash, cheque] \wedge T < T_s$, where “::” represents a domain constraint.

IC *alice1* says that, if *eShop* expects *alice* to pay for an *Item* by credit card, then *alice* expects that *eShop* will have provided a guarantee by the time she pays.

$$\begin{aligned} & \mathbf{E}(alice, eShop, pay(Item, cc), T_p) \\ & \wedge \mathbf{H}(eShop, alice, give_guarantee, T_g) \wedge T_g < T_p \quad (\text{alice2}) \\ & \rightarrow \mathbf{H}(alice, eShop, pay(Item, cc), T_p). \end{aligned}$$

IC *alice2* says that, if *eShop* expects *alice* to pay for an *Item* by credit card, and *eShop* has provided *alice* with a guarantee, then *alice* will pay the *Item* by credit card. Finally, the following clause is part of \mathcal{KB}_{alice}

$$\begin{aligned} & have(alice, Item, T) \leftarrow \\ & \mathbf{E}(eShop, alice, deliver(Item), T_d) \wedge T_d \leq T. \quad (\text{alice3}) \end{aligned}$$

Clause *alice3* says that, in order for *alice* to have an *Item* at time *T*, then *alice* expects *eShop* to deliver the *Item* by time *T*.

6 Declarative and Operational Semantics

We have assumed that all web services have their own interface behaviour specified in the language of *ICs*. This behaviour could be thought of as an extension of WSDL, that could be used by other fellow web services to reason about the specifications, or to check if inter-operability is possible. We are currently studying an XML-like extension of RuleML [2] to represent *ICs*.

Another approach would be to obtain web services' interface behaviour through an appropriate request protocol, in which *ICs* are (interactively) exchanged so that each web service may disclose *ad hoc*, customised information on demand.

In this work, we make the simplifying assumption that all information regarding the interface behaviour is provided at once. The web service will then try and prove that a fruitful interaction is possible based on what it receives.

The web service initiating the interaction has a goal \mathcal{G} , which is a given state of affairs. A typical goal could be to access a resource, to retrieve some information, or to obtain a service from another web service. \mathcal{G} will often be an expectation (of obtaining a service, accessing a resource, or gathering information), but in general it can be any conjunction of expectations, CLP constraints, and any other literals, in the syntax of \mathcal{IC}_{ws} *Head Disjuncts* (Eq. 2).

The verification of a web service *ws* about the possibility to achieve a goal \mathcal{G} by interacting with another fellow web service *ws'* makes use of \mathcal{KB}_{ws} , \mathcal{IC}_{ws} , \mathcal{G} , and of the information obtained about *ws'*'s policies, $\mathcal{IC}_{ws'}$ (see Fig. 1). The idea is to obtain, through abductive reasoning, a set of expectations about a possible course of events that together with \mathcal{KB}_{ws} entails $\mathcal{IC}_{ws} \cup \mathcal{IC}_{ws'}$ and \mathcal{G} .

Note that we do not assume that *ws* knows $\mathcal{KB}_{ws'}$, as the \mathcal{KB} is not part of the interface. However, in general integrity constraints can involve predicates defined in the knowledge base. For example, they can contain predicates defining parameters, deadlines, coefficients, etc., or other knowledge only available to *ws'*. If the interface behaviour provided by *ws'* involves predicates defined in $\mathcal{KB}_{ws'}$, unknown to *ws*, we have two alternatives:

- either ws' provides ws with the necessary information, e.g. with (part of) its $\mathcal{KB}_{ws'}$;
- or ws will have to make assumptions about such unknown predicates.

We take the second option, and consider unknowns that are neither **H** events nor **E** expectations as literals that can be abduced, and we keep them in a set Δ . We then have the following two equations that define the set of abductive answers representing possible interaction between ws and ws' achieving \mathcal{G} :

$$\mathcal{KB}_{ws} \cup \mathbf{HAP} \cup \mathbf{EXP} \cup \Delta \models \mathcal{G} \quad (3)$$

$$\mathcal{KB}_{ws} \cup \mathbf{HAP} \cup \mathbf{EXP} \cup \Delta \models \mathcal{IC}_{ws} \cup \mathcal{IC}_{ws'}. \quad (4)$$

where **HAP** is a conjunction of **H** atoms, **EXP** is a conjunction of **E** atoms, and Δ a conjunction of abducible atoms.

We ground the notion of entailment on a model theoretic semantics defined for Abductive Disjunctive Logic Programs [10]. Different semantics have been proposed for logic programs with disjunctions. Among them, answer set semantics [11] adopts an exclusive interpretation of disjunction, whereas possible model semantics [10] adopts an inclusive one (and recovers the former by additional constraints imposing mutual exclusion among the literals in the disjunctive head of a clause).

In the possible model semantics, the disjunctive program generates several (non-disjunctive) split programs, obtained by separating the disjuncts in the head of rules. Given a disjunctive logic program P , a *split program* is defined as a (ground) logic program obtained from P by replacing every (ground) rule

$$r : L_1 \vee \dots \vee L_l \leftarrow \Gamma$$

from P with the rules in a non-empty subset of $Split_r$, where

$$Split_r = \{L_i \leftarrow \Gamma \mid i = 1, \dots, l\}.$$

By definition, P has multiple split programs in general.

Example 1. The following program can be split into three split programs

$$\begin{aligned} \mathbf{E}(p) \vee \mathbf{H}(p) &\leftarrow \mathbf{E}(p). \\ goal &\leftarrow \mathbf{E}(p). \end{aligned}$$

where the first clause is respectively substituted by $\{\mathbf{E}(p) \leftarrow \mathbf{E}(p)\}$, $\{\mathbf{H}(p) \leftarrow \mathbf{E}(p)\}$ and $\{\mathbf{E}(p) \leftarrow \mathbf{E}(p), \mathbf{H}(p) \leftarrow \mathbf{E}(p)\}$.

A *possible model* for a disjunctive logic program P is then defined as an answer set of a split program of P .

The inclusive interpretation of disjunctions adopted by the possible model semantics fits better with our case, since more than one disjunct in the head of an integrity constraint can be true at the same time, as in the following example.⁵

⁵ For the sake of simplicity, in this example and in the following one, we specify a single argument for expectations and events.

Example 2. Let us consider the program:

$$\begin{aligned} \mathbf{E}(p) \vee \mathbf{H}(p) &\leftarrow \text{true}. \\ \mathbf{H}(p) &\leftarrow \mathbf{E}(p). \\ \text{goal} &\leftarrow \mathbf{E}(p). \end{aligned}$$

We would like to have an explanation for *goal*, where $\mathbf{E}(p)$ is assumed, and $\mathbf{H}(p)$ is also true because of clause $\mathbf{H}(p) \leftarrow \mathbf{E}(p)$. This is, instead, avoided by following an answer set approach, which adopts an exclusive interpretation of disjunctions.

Furthermore, in [10] possible model semantics was also applied to provide a model theoretic semantics for Abductive Extended Disjunctive Logic Programs (AEDP), which is our case. Semantics is given to AEDP in terms of *possible belief sets*. Given an AEDP $\Pi = \langle P, \mathcal{A} \rangle$ (where P is a disjunctive logic program and \mathcal{A} is the set of abducible literals), a possible belief set S of Π is a possible model of the disjunctive program $P \cup E$, where P is extended with a set E of abducible literals ($E \subseteq \mathcal{A}$).

Definition 2 (Answer to a goal \mathcal{G}). *An answer E to a (ground) goal \mathcal{G} is a set E of abducible literals constituting the abductive portion of a possible belief set S (i.e., $E = S \cap \mathcal{A}$) that entails \mathcal{G} .*

We rely upon possible belief set semantics, but we adopt a new notion for minimality with respect to abducible literals. In [10], a possible belief set S is \mathcal{A} -minimal if there is no possible belief set T such that $T \cap \mathcal{A} \subset S \cap \mathcal{A}$.

Example 3. Consider, again, the program of Example 1. The possible belief sets are the belief sets of each of the split programs: the first gives $\{\text{goal}, \mathbf{E}(p)\}$, and the others give $\{\text{goal}, \mathbf{E}(p), \mathbf{H}(p)\}$. Only the first explanation is \mathcal{A} -minimal under set inclusion, according to [10], but we cannot rely upon such definition for minimality since we would discard explanations which are, instead, correct.

We restate, then, the notion of \mathcal{A} -minimality as follows.

Definition 3 (\mathcal{A} -minimal possible belief set). *A possible belief set S is \mathcal{A} -minimal iff there is no possible belief set T for the same split program such that $T \cap \mathcal{A} \subset S \cap \mathcal{A}$.*

More intuitively, the notion of minimality with respect to hypotheses that we introduce is checked by considering the *same* split program, and by checking whether with a smaller set of abducibles the same consequences can be made true, but in the same split program. For the case depicted in Example 3, then, both the two possible belief sets are \mathcal{A} -minimal, according to Definition 3.

Finally, we provide a model-theoretic notion of explanation to an observation, in terms of answer to a goal, as follows.

Definition 4 (\mathcal{A} -minimal answer to a goal). *E is an \mathcal{A} -minimal answer to a goal \mathcal{G} iff $E = S \cap \mathcal{A}$ for some possible \mathcal{A} -minimal belief set S that entails \mathcal{G} .*

We can now proceed with defining what kind of interaction is possible/fruitful, given two web services and a goal.

Definition 5 (Possible interaction about \mathcal{G}). *A possible interaction about a goal \mathcal{G} between two web services ws and ws' is an \mathcal{A} -minimal set $\mathbf{HAP} \cup \mathbf{EXP} \cup \Delta$ such that Eq. 3 and 4 hold.*

Among all possible interactions about \mathcal{G} , some of them are fruitful, and some are not. An interaction only based on expectations which will not be matched by corresponding events is not a fruitful one: for example, the goal of ws might not have a corresponding event, thus the goal is not actually reached, but only *expected*. Or, one of the web services could be waiting for a message from the other fellow, which will never arrive, thus undermining the inter-operability.

We select, among the possible interactions, those whose history satisfies all the expectations of both the web services. After the abductive phase, we have a verification phase in which there are no abducibles, and in which the previously abducted predicates \mathbf{H} and \mathbf{E} are now considered as defined by atoms in \mathbf{HAP} and \mathbf{EXP} , and they have to match. If among the possible interactions there exists one satisfying

$$\mathbf{HAP} \cup \mathbf{EXP} \models \mathbf{E}(X, Y, Action, T) \leftrightarrow \mathbf{H}(X, Y, Action, T) \quad (5)$$

then ws has found a sequence of actions that obtains the goal \mathcal{G} .

Definition 6 (Possible interaction achieving \mathcal{G}). *Given two web services, ws and ws' , and a goal \mathcal{G} , a possible interaction achieving \mathcal{G} is a possible interaction about \mathcal{G} satisfying Eq. 5.*

Intuitively, the “ \rightarrow ” implication in Eq. 5 is there to avoid situations in which a web service waits forever for an event that the other web service will never produce. The “ \leftarrow ” implication is there to avoid that one web service sends unexpected messages, which in the best case may not be understood (and in the worst scenarios it may lead to faulty, unpredictable behaviour of the parties involved).

6.1 Operational Semantics

The operational semantics is a modification of the *SCIFF* proof-procedure [12]. *SCIFF* is a transition system, whose state is given by the following tuple:

$$T \equiv \langle R, CS, PSIC, \Delta A, \mathbf{PEND}, \mathbf{HAP}, \mathbf{FULF}, \mathbf{VIOL} \rangle$$

The set of expectations \mathbf{EXP} is partitioned into the fulfilled (\mathbf{FULF}), violating (\mathbf{VIOL}), and pending (\mathbf{PEND}) expectations. The other elements are: the resolvent (R), the abducted literals that are not expectations (ΔA), the constraint store (CS), a set of implications, inherited from the IFF [7], called *partially solved integrity constraints* ($PSIC$), and the history of happened events (\mathbf{HAP}).

A classical application of *SCIFF* is on-line checking of compliance of agent interaction to protocols. In fact, *SCIFF* was initially developed to specify and verify agent interaction protocols on-the-fly, under the assumption of open agent environments adopted by other noteworthy agent research work [13]. *SCIFF* processes events drawing from **HAP** and generates (abduces) expectations; then it checks that all expectations are fulfilled by at least one happened event. The declarative semantics of *SCIFF* contains in fact a requirement $\mathbf{E}(X) \rightarrow \mathbf{H}(X)$ – differently from WAV^e , which has a double implication (Eq. 5). In *SCIFF*, as soon as new **H** events are processed, a transition *fulfilment* labels the relevant matching expectations as *fulfilled* and moves them to the set **FULF**. At the end of the derivation, if some expectation remains in the set **PEND**, a failure node is generated, and other alternative branches will be explored in backtracking, if there exist any.

WAV^e extends *SCIFF* and abduces **H** events as well as expectations. The events history is not taken as input, but all possible interactions are hypothesised. Moreover, in WAV^e events not matched by an expectation (which are perfectly acceptable in the multi-agent scenario addressed by *SCIFF*) cannot be part of a *possible interaction achieving* the goal.

The two phases in the declarative semantics (generation of possible interactions and their test for conformance) are condensed into one single derivation process, thanks to a new transition adopted in WAV^e . The *expected* transition, symmetrical to *fulfilment*, labels each **H** events with an *expected* flag as soon as an expectation matching it is abduced. At the end of the derivation, **H** with *expected* status = false will cause failure.

Otherwise, if the WAV^e derivation in a program \mathcal{P} for a goal \mathcal{G} succeeds with set of expectation $\mathbf{EXP} \cup \mathbf{HAP} \cup \Delta$, we write $\mathcal{P} \vdash_{\mathbf{EXP} \cup \mathbf{HAP} \cup \Delta} \mathcal{G}$.

6.2 Soundness and completeness results

WAV^e is a conservative modification of the *SCIFF* proof-procedure, which is sound and complete under reasonable assumptions [14]. In the following, we give the soundness and completeness statements, and briefly explain why the soundness and completeness results proven for *SCIFF* also hold with the new declarative and operational semantics of WAV^e .

Theorem 1 (Soundness). *If $\mathcal{P} \vdash_{\mathbf{EXP} \cup \mathbf{HAP} \cup \Delta} \mathcal{G}$ then $\mathbf{EXP} \cup \mathbf{HAP} \cup \Delta$ is a possible interaction achieving \mathcal{G} .*

Theorem 2 (Completeness). *If there exists a possible interaction $\mathbf{EXP} \cup \mathbf{HAP} \cup \Delta$ achieving a goal \mathcal{G} , then $\exists \mathbf{EXP}' \cup \mathbf{HAP}' \cup \Delta' \subseteq \mathbf{EXP} \cup \mathbf{HAP} \cup \Delta$ such that $\mathcal{P} \vdash_{\mathbf{EXP}' \cup \mathbf{HAP}' \cup \Delta'} \mathcal{G}$.*

Note that WAV^e introduces two main additions to *SCIFF*. The first one is the “ \rightarrow ” implication of Eq. 5, which makes the declarative relation between events and expectations symmetric. The operational semantics introduces a new transition, completely symmetric to that devoted to check fulfilment; the extension of the proofs for this are therefore straightforward.

The second ‘important’ addition is the notion of \mathcal{A} -minimality introduced in the declarative semantics. By choosing \mathcal{A} -minimal answers, therefore restricting the set of models considered, we do not affect completeness, which still holds by virtue of the completeness of **SCIFF**. Concerning soundness, instead, we have to prove that, given a goal \mathcal{G} , for each abductive WAV^e computation:

$$\mathcal{KB}_{ws} \cup \mathcal{IC}_{ws} \cup \mathcal{IC}_{ws'} \vdash_{\mathbf{HAP} \cup \mathbf{EXP} \cup \Delta} \mathcal{G} \quad (6)$$

the computed set $\mathbf{HAP} \cup \mathbf{EXP} \cup \Delta$ corresponds to a possible interaction achieving \mathcal{G} . Again, the proof can exploit the soundness result of **SCIFF** [15], apart from the \mathcal{A} -minimality requirement introduced here. This further condition requires to prove that the computed set $E = \mathbf{HAP} \cup \mathbf{EXP} \cup \Delta$ for goal \mathcal{G} corresponds to an (\mathcal{A} -minimal) possible interaction for \mathcal{G} . First, it can easily be proven that a WAV^e computation corresponds to a (WAV^e) computation into a *single* split program obtained from the original one. Furthermore, it can be proven that a WAV^e computation corresponds to a (WAV^e) computation into a *renamed* split program obtained from the former by considering only the applied clauses and the fired social integrity constraints, and by duplicating and properly renaming them as many times as each of them is applied or fires. Let us denote $P^{split} = \mathcal{KB}_{ws}^{split} \cup \mathcal{IC}_{ws}^{split} \cup \mathcal{IC}_{ws'}^{split}$ such a *renamed* split program.

Example 4. Let us consider the following program:

$$\begin{aligned} \mathbf{E}(X) \vee \mathbf{H}(X) &\leftarrow \mathbf{E}(X). \\ \text{goal} &\leftarrow \mathbf{E}(p). \\ \text{goal} &\leftarrow \mathbf{E}(q). \end{aligned}$$

This program has two different successful derivations for *goal*, originated, respectively, by the following two renamed split programs:⁶

$$\begin{array}{c} \frac{P_1^{split}}{\mathbf{H}(p) \leftarrow \mathbf{E}(p). \\ \text{goal} \leftarrow \mathbf{E}(p).} \qquad \frac{P_2^{split}}{\mathbf{H}(q) \leftarrow \mathbf{E}(q). \\ \text{goal} \leftarrow \mathbf{E}(q).} \end{array}$$

Thanks to the soundness of **SCIFF** we have that E is a possible interaction for \mathcal{G} , given the considered *renamed* split program. We still have to prove that this set is \mathcal{A} -minimal. This part can be proven by reductio ad absurdum. Suppose there exists a smaller set $E' \subset E$ and that E' is a possible interaction for \mathcal{G} in the *renamed* split program. Due to **SCIFF** completeness, then there also exists a (WAV^e) computation which computes a set $E'' \subseteq E'$ for \mathcal{G} . But this is not possible, by the way the *renamed* split program has been built. Contradiction!

We will next demonstrate the operational functioning of verification in WAV^e in the *alice & eShop* scenario.

⁶ For the sake of keeping a lightweight notation, we do not show renamed variables.

7 Verification in WAV^e

In the following, the sets \mathbf{EXP}_a^N and \mathbf{HAP}_a^N represent the evolution of *alice*'s expectations and events as WAV^e's derivation progresses; N is an incremental index. Let g be the following goal of *alice*'s:

$$g \leftarrow \text{have}(\text{alice}, \text{device}, 50). \quad (\text{goal})$$

Then, by unfolding of clause *alice3*,

$$\mathbf{EXP}_a^0 = \{ \mathbf{E}(\text{eShop}, \text{alice}, \text{deliver}(\text{device}), T_s) \wedge T_s < 50 \} \quad (\text{by } (\text{alice3}))$$

To this expectation, *eShop* will react by expecting a payment:

$$\begin{aligned} \mathbf{EXP}_a^1 = \{ & \mathbf{E}(\text{eShop}, \text{alice}, \text{deliver}(\text{device}), T_s) \wedge T_s < 50 \\ & \wedge (\mathbf{E}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cc}), T_{cc}) \wedge T_{cc} < T_s \\ & \vee \mathbf{E}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cash}), T_{ca}) \wedge T_{ca} < T_s \\ & \vee \mathbf{E}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cheque}), T_{ch}) \wedge T_{ch} < T_s) \} \quad (\text{by } (\text{shop1})) \end{aligned}$$

Since the expectation containing the payment by *cc* is the only one which generates an expectation matching a rule of *alice* ((*alice1*)), the first expectation among the three payment alternatives is selected (the other branches eventually fail by Eq. 5, because no matching **H** is abduced). This choice triggers (*alice1*):

$$\begin{aligned} \mathbf{EXP}_a^2 = \{ & \mathbf{E}(\text{eShop}, \text{alice}, \text{deliver}(\text{device}), T_s) \wedge T_s < 50 \\ & \wedge \mathbf{E}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cc}), T_{cc}) \wedge T_{cc} < T_s \\ & \wedge \mathbf{E}(\text{eShop}, \text{alice}, \text{give_guarantee}, T_g) \wedge T_g < T_{cc} \} \quad (\text{by } (\text{alice1})) \end{aligned}$$

Then (*shop3*) fires, and abduces the happening of *give_guarantee* event. We then have:

$$\begin{aligned} \mathbf{EXP}_a^3 = \{ & \mathbf{E}(\text{eShop}, \text{alice}, \text{deliver}(\text{device}), T_s) \wedge T_s < 50 \\ & \wedge \mathbf{E}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cc}), T_{cc}) \wedge T_{cc} < T_s \\ & \wedge \mathbf{E}(\text{eShop}, \text{alice}, \text{give_guarantee}, T_g) \wedge T_g < T_{cc} \} \quad (\text{by } (\text{alice1})) \\ \mathbf{HAP}_a^3 = \{ & \mathbf{H}(\text{eShop}, \text{alice}, \text{give_guarantee}, T_g) \wedge T_g < T_{cc} \} \quad (\text{by } (\text{shop3})) \end{aligned}$$

Given the guarantee, *alice* will pay by credit card (rule (*alice2*) fires):

$$\begin{aligned} \mathbf{EXP}_a^4 = \{ & \mathbf{E}(\text{eShop}, \text{alice}, \text{deliver}(\text{device}), T_s) \wedge T_s < 50 \\ & \wedge \mathbf{E}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cc}), T_{cc}) \wedge T_{cc} < T_s \\ & \wedge \mathbf{E}(\text{eShop}, \text{alice}, \text{give_guarantee}, T_g) \wedge T_g < T_{cc} \} \\ \mathbf{HAP}_a^4 = \{ & \mathbf{H}(\text{eShop}, \text{alice}, \text{give_guarantee}, T_g) \wedge T_g < T_{cc} \\ & \wedge \mathbf{H}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cc}), T_{cc}) \wedge T_{cc} < T_s \} \quad (\text{by } (\text{alice2})) \end{aligned}$$

Having received the payment, *eShop*'s policy would be to deliver the device:

$$\begin{aligned} \mathbf{EXP}_a^5 = \{ & \mathbf{E}(\text{eShop}, \text{alice}, \text{deliver}(\text{device}), T_s) \wedge T_s < 50 \\ & \wedge \mathbf{E}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cc}), T_{cc}) \wedge T_{cc} < T_s \\ & \wedge \mathbf{E}(\text{eShop}, \text{alice}, \text{give_guarantee}, T_g) \wedge T_g < T_{cc} \} \\ \mathbf{HAP}_a^5 = \{ & \mathbf{H}(\text{eShop}, \text{alice}, \text{give_guarantee}, T_g) \wedge T_g < T_{cc} \\ & \wedge \mathbf{H}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cc}), T_{cc}) \wedge T_{cc} < T_s \\ & \wedge \mathbf{H}(\text{eShop}, \text{alice}, \text{deliver}(\text{device}), T_s) \wedge T_s < 50 \} \quad (\text{by } (\text{shop2})) \end{aligned}$$

Summarising, *alice* devised the following set of events, which should let her achieve her goal (have the desired device) while respecting both of *alice*'s and *eShop*'s policies.

$$\begin{aligned} C_a = \{ & \mathbf{H}(\text{eShop}, \text{alice}, \text{give_guarantee}, T_g) \wedge T_g < T_p \\ & \wedge \mathbf{H}(\text{alice}, \text{eShop}, \text{pay}(\text{device}, \text{cc}), T_p) \wedge T_p < T_s \\ & \wedge \mathbf{H}(\text{eShop}, \text{alice}, \text{deliver}(\text{device}), T_s) \wedge T_s < 50 \} \end{aligned}$$

8 Discussion

WAV^e is a framework intended for describing declaratively the behavioural interface of web services, and for testing operationally the possibility of fruitful interaction between them. WAV^e answers the question “does there exist a viable interaction, between two given web services, which achieves a given goal \mathcal{G} ?” In case of success, WAV^e produces a set of expectations about events. WAV^e is particularly suitable for highly dynamic environments, in which inter-operability is an unknown that has to be checked.

WAV^e uses and extends a technology initially developed for on-line compliance verification of agent interaction to protocols [9]. SCIFF and the protocol specification language based on social integrity constraints were motivated and inspired by conspicuous work done in the context of agent interaction in open societies, notably work by Singh [13] and colleagues. The extension of such a work to the context of web services, centering around the concept of policies, as proposed in this work, seems to be very promising. The idea of policies for web services and policy-based reasoning is one that many other authors also adopt. We will cite work by Finin and colleagues [16], and by Bradshaw and colleagues [17], the first one with an emphasis on representation of actions, the latter on the deontic semantic aspects of web service interaction. We acknowledge the importance of action modelling and we point that the idea of expected behaviour of web services can have a deontic reading. In fact, previous work on SCIFF has been devoted to investigating and clarifying the interesting links between deontic operators and expectation-based reasoning [18]. The distinguishing features of WAV^e, compare to most work of literature, are its logical underpinning and its sound and complete operational characterisation. It is in our agenda to carry out an extensive empiric evaluation of WAV^e based on interesting cases and scenarios such as those proposed in related work, and on the existing implementation of the SCIFF framework.⁷

Another direction of current work relates to the actual use of the answers of WAV^e by web services after they manage a successful derivation. In principle, the sequence of events produced by WAV^e could be instantiated into a concrete sequence of messages, which will guarantee the achievement of \mathcal{G} , under ideal external conditions. But this is true only if the policies disclosed by both web services are a faithful representation of their actual behaviour. This may not be the case, as for example policies may depend on sensible data, and web services may be not allowed to disclose full information to the outside. In that case nothing warrants that the course of action produced by WAV^e will be satisfactory for either web service. We might then have to resort to further steps. For example both web services could “formally” agree that a certain course of events in an acceptable option, possibly after another mutual verification phase. This is subject of ongoing work. Finally, we are currently investigating the exchange of policies between web services, for which a suitable interaction

⁷ See <http://lia.deis.unibo.it/research/sciff>.

protocol needs to be devised. We are thinking of specifying such a protocol for exchanging the policies in the same language WAV^e uses to specify policies.

References

1. Andrews, T., Curbera, F., Dholakia, H., Golland, Y., Klein, J., Leymann, F., Liu, K., Roller, D., Smith, D., Thatte, S., Trickovic, I., Weerawarana, S.: Business process execution language for web services version 1.1 (2003) <http://www.ibm.com/developerworks/library/ws-bpel/>.
2. Adi, A., Stoutenburg, S., Tabet, S., eds.: *Proc. 1st Int. Conference on Rules and Rule Markup Languages for the Semantic Web, LNCS 3791*, Springer Verlag (2005)
3. Working Group on Rule Interchange Format: Use cases and requirements. <http://www.w3.org/2005/rules/wg/ucr/draft-20060323.html> (2006)
4. Bry, F., Eckert, M.: Twelve theses on reactive rules for the web. In: Proceedings of the Workshop on Reactivity on the Web, Munich, Germany (2006)
5. Kakas, A.C., Kowalski, R.A., Toni, F.: Abductive Logic Programming. *Journal of Logic and Computation* **2** (1993) 719–770
6. Kakas, A.C., Mancarella, P.: On the relation between Truth Maintenance and Abduction. In: *Proc. 1st Pacific Rim International Conference on Artificial Intelligence*, Ohmsha Ltd. (1990) 438–443
7. Fung, T.H., Kowalski, R.A.: The IFF proof procedure for abductive logic programming. *Journal of Logic Programming* **33** (1997) 151–165
8. Denecker, M., Schreye, D.D.: SLDNFA: an abductive procedure for abductive logic programs. *Journal of Logic Programming* **34** (1998) 111–167
9. Alberti, M., Chesani, F., Gavanelli, M., Lamma, E., Mello, P., Torroni, P.: The SOCS computational logic approach for the specification and verification of agent societies. In: *Global Computing 2004, LNAI 3267*, Springer-Verlag (2005) 324–339
10. Sakama, C., Inoue, K.: Abductive logic programming and disjunctive logic programming: their relationship and transferability. *Journal of Logic Programming* **44** (2000) 75–100
11. Gelfond, M., Lifschitz, V.: Classical negation in logic programming and disjunctive databases. *New Generation Computing* **9** (1991) 365–385
12. Alberti, M., Gavanelli, M., Lamma, E., Mello, P., Torroni, P.: The SCIFF abductive proof-procedure. In: *AI*IA, LNAI 3673*, Springer-Verlag (2005) 135–147
13. Singh, M.: Agent communication language: rethinking the principles. *IEEE Computer* (1998) 40–47
14. Alberti, M., Chesani, F., Gavanelli, M., Lamma, E., Mello, P., Torroni, P.: Verifiable agent interaction in abductive logic programming: the SCIFF proof-procedure. Technical Report DEIS-LIA-06-001, DEIS, University of Bologna (Italy) (2006)
15. Gavanelli, M., Lamma, E., Mello, P.: Proof of properties of the SCIFF proof-procedure. Technical Report CS-2005-01, DI, Università di Ferrara (2005) Available at <http://www.ing.unife.it/informatica/tr/CS-2005-01.pdf>.
16. Kagal, L., Finin, T.W., Joshi, A.: A policy based approach to security for the semantic web. In *Proc. 2nd ISWC. LNCS 2870*, Springer (2003) 402–418
17. Uszok, A., Bradshaw, J.M., Jeffers, R., Tate, A., Dalton, J.: Applying KAoS services to ensure policy compliance for semantic web services workflow composition and enactment. In *Proc. 3rd ISWC. LNCS 3298*, Springer (2004) 425–440
18. Alberti, M., Gavanelli, M., Lamma, E., Mello, P., Sartor, G., Torroni, P.: Mapping deontic operators to abductive expectations. *Computational and Mathematical Organization Theory* (2006) To appear.