

Web Service Composition via the Customization of Golog Programs with User Preferences

Shirin Sohrabi and Nataliya Prokoshyna and Sheila A. McIlraith

Department of Computer Science, University of Toronto, Toronto, Canada.

{shirin,nataliya,sheila}@cs.toronto.edu

Abstract

We claim that user preferences are a key component of effective Web service composition, and one that has largely been ignored. In this paper we propose a means of specifying and intergrating user preferences into Web service composition. To this end, we propose a means of performing automated Web service composition by exploiting a flexible template of the composition in the form of a generic procedure. This template is augmented by a rich specification of user preferences that guide the instantiation of the template. We exploit the agent programming language Golog to represent our templates as Golog generic procedures and we exploit a first-order preference language to represent rich qualitative temporally-extended user preferences. From these we generate Web service compositions that realize a given generic procedure, satisfying the user's hard constraints and optimizing for the user's preferences. We prove our approach is sound and optimal. Our system, Golog-Pref, is implemented and interacting with services on the Web. The language and techniques proposed in this paper can be integrated into a variety of approaches to Web or Grid service composition.

Preamble

The work that follows first appeared in the proceedings of the 5th International Semantic Web Conference (ISWC 2006) [26] and was recently reprinted in [27]. These archival publications are the appropriate references for this work. The text that follows has been shortened to address length restrictions and modified slightly to address referee comments.

We wanted to share this work with participants in the ICAPS 2009 Workshop on Generalized Planning because it presents a particular approach to specifying and instantiating generalized plans, as well as describing a well-motivated application for generalized plans – Web service composition. The approach advocated in this paper is to represent a solution to a family of planning problems as a template or workflow scaffold, which is a form of generalized plan. In this paper, we use Golog to specify the template, in more recent work on Web service composition, we have used Hierarchical Task Networks (HTNs) [25].

A key observation of this work is that the task of creating a generalized plan is something that can often be performed manually because the template is simple, but that the instantiation of that template into a high-quality plan can necessitate a number of decision points because of nondeterminism in the template, and that this presents a significant conceptual and computational challenge. We hold that this challenge is best addressed by specifying desirable properties of the plan in terms of preferences and constraints, and

generating plan instances that optimize these desirable properties. The need to provide further advice to generalized plans was first observed by Myers and colleagues, precipitating a body of work on advisable plans that augmented HTNs with hard and soft advice on how to decompose the HTN [17, 18]. The need for optimization in the instantiation of generalized plans is particularly clearly illustrated with Web service composition where the large number of available services and the large volume of Web data translates into execution-time choice points over service selection and action groundings.

Both HTNs and Golog have a rich history in terms of their ability to specify generalized plans. Generalized plans have also been specified as automata (e.g., [4]) and as hierarchical constraint networks (e.g., [10]), and both of these techniques have been applied to Web service composition, but without any preference specification. Clearly all of these techniques, including those described below, are applicable to a diversity of applications beyond Web service composition, including business process modeling and execution, and controller synthesis.

1 Introduction

Web services provide a standardized means for diverse, distributed software applications to be published on the Web and to interoperate seamlessly. Simple Web accessible programs are described using machine-processable descriptions and can be loosely composed together to achieve complex behaviour. The weather service at www.weather.com and the flight-booking services at www.aircanada.ca, are examples of Web applications that can be described and composed as Web services. They might be coupled as part of a travel-booking service, for example.

Automated Web service composition is one of many interesting challenges facing the Semantic Web. Given computer-interpretable descriptions of: the task to be performed, the properties and capabilities of available Web services, and possibly some information about the client or user's specific constraints, *automated Web service composition* requires a computer program to automatically select, integrate and invoke multiple Web services in order to achieve the specified task in accordance with any user-specific constraints. Compositions of Web or Grid services are necessary for realizing both routine and complex tasks on the Web (resp. Grid) without the need for time-consuming manual

composition and integration of information. Compositions are also a useful way of enforcing business rules and policies in both Web and Grid computing.

Fully automated Web service composition has been characterized as akin to both an artificial intelligence (AI) planning task and to a restricted software synthesis task (e.g., [14]). A composition can be achieved using classical AI planning techniques by conceiving services as primitive or complex actions and the task description specified as a (final state) goal (e.g., [13, 28]). This approach has its drawbacks when dealing with data. In general, the search space for a composition (aka plan) is huge because of the large number of available services (actions), which grow far larger with grounding for data.

A reasonable middle ground which we originally proposed in [15, 14] is to specify a flexible template of the composition in the form of a *generic procedure* and to customize such a procedure with *user constraints*. We argued that many of the tasks performed on the Web or on intranets are repeated routinely, and the basic steps to achieving these tasks are well understood, at least at an abstract level – travel planning is one such example. Nevertheless, the realization of such tasks varies as it is tailored to individual users. As such, our proposal was to specify such tasks using a workflow template or generic procedure and to customize the procedure with user constraints at run time. Such an approach is generally of the same complexity as planning but the search space is greatly reduced, and as such significantly more efficient than planning without such generic advice.

In [14] we proposed to use an augmented version of the agent programming language Golog [20] to specify our generic procedures or workflows with sufficient nondeterminism to allow for customization. (E.g., “*book inter-city transportation, local transportation and accommodations in any order*”). User constraints (e.g., “*I want to fly with Air Canada.*”) were limited to hard constraints (as opposed to “soft”), were specified in first-order logic (FOL), and were applied to the generic procedure at run-time to generate a user-specific composition of services. A similar approach was adopted using hierarchical task networks (HTNs) to represent generic procedures or templates, and realized using the HTN planner, SHOP2 (e.g., [22]) without user customization of the HTN template.

In this paper, we extend our Golog framework for Web service composition, customizing Golog generic procedures not only with hard constraints but with *soft* user constraints (henceforth referred to as *preferences*). These preferences are defeasible and may not be mutually achievable. We argue that user preferences are a critical and missing component of most existing approaches to Web service composition. User preferences are key for at least two reasons. First, the user’s task (specified as a goal and/or generic procedure with user constraints) is often under constrained. As such, it induces a family of solutions. User preferences enable a user to specify properties of solutions that make them more or less desirable. The composition system can use these to generate preferred solutions.

A second reason why user preferences are critical to Web service composition is with respect to *how* the composition

is performed. A key component of Web service composition is the selection of specific services used to realize the composition. In AI planning, primitive actions (the analogue of services) are selected for composition based on their preconditions and effects, and there is often only one primitive action that realizes a particular effect. Like actions, services are selected for composition based on functional properties such as inputs, output, preconditions and effects, but they are also selected based on domain-specific nonfunctional properties such as, in the case of airline ticket booking, whether they book flights with a carrier the user prefers, what credit cards they accept, how trusted they are, etc. By integrating user preferences into Web service composition, preferences over services (the *how*) can be specified and considered along side preferences over the solutions (the *what*).

In this paper we recast the problem of Web service composition as the task of finding a composition of services that achieves the task description (specified as a generic procedure in Golog), that achieves the user’s hard constraints, and that is *optimal* with respect to the user’s preferences. To specify user preferences, we exploit a rich qualitative preference language, based on the LPP language proposed by Bienvenu et al. to specify users’ preferences in a variant of linear temporal logic (LTL) [1, 2]. We prove the soundness of our approach and the optimality of our compositions with respect to the user’s preferences. Our system can be used to select the optimal solution from among families of solutions that achieve the user’s stated objective. Our system is implemented in Prolog and integrated with a selection of scraped Web services that are appropriate to our test domain of travel planning.

The work presented here is predicated on the assumption that Web services have been described in a computer-interpretable form. This is the starting point for most work on semantic Web services [15] and a great deal of effort has gone into the development of ontologies for precisely this purpose. In this paper, Web service descriptions are presented in FOL, *not* in one of the typical Semantic Web languages such as OWL [8] nor more specifically in terms of a semantic Web service ontology such as OWL-S [12] or WSMO [3]. Nevertheless, it is of direct significance to semantic Web services. As noted in (e.g., [12]) process models, necessary for Web service composition, cannot be expressed in OWL while preserving all and only the intended interpretations of the process model. OWL (and thus OWL-S) is not sufficiently expressive. Further OWL reasoners are not designed for the type of inference necessary for Web service composition. For both these reasons, Web service composition systems generally translate the relevant aspects of service ontologies such as OWL-S into internal representations such as the Planning Domain Definition Language (PDDL) that are more amenable to AI planning (e.g., [22, 9]). Golog served as one of the inspirations for what is now OWL-S [15] and all the OWL-S constructs have translations into Golog [16]. Further, the semantics of the OWL-S process model has been specified in situation calculus [19]. Thus, our Golog generic procedures can be expressed in OWL-S and likewise, OWL-S ontologies can be translated into our formalism. We do not have a cur-

rent implementation of this translation, but it is conceptually straightforward.

2 Situation Calculus and Golog

We use the situation calculus and FOL to describe the functional and nonfunctional properties of our Web services. We use the agent programming language Golog to specify composite Web services and to specify our generic procedures. In this section, we review the essentials of situation calculus and Golog.

The situation calculus is a logical language for specifying and reasoning about dynamical systems [20]. In the situation calculus, the *state* of the world is expressed in terms of functions and relations (fluents) relativized to a particular *situation* s , e.g., $F(\vec{x}, s)$. In this paper, we distinguish between the set of fluent predicates, \mathcal{F} , and the set of non-fluent predicates, \mathcal{R} , representing properties that do not change over time. A situation s is a *history* of the primitive actions, $a \in \mathcal{A}$, performed from a distinguished initial situation S_0 . The function $do(a, s)$ maps a situation and an action into a new situation thus inducing a tree of situations rooted in S_0 . $Poss(a, s)$ is true if action a is possible in situation s .

Web services such as the Web exposed application at www.weather.com are viewed as actions in the situation calculus and are described as actions in terms of a situation calculus basic action theory, \mathcal{D} . The details of \mathcal{D} are not essential to this paper but the interested reader is directed to [20, 19, 14] for further details.

Golog [20] is a high-level logic programming language for the specification and execution of complex actions in dynamical domains. It builds on top of the situation calculus by providing Algol-inspired extralogical constructs for assembling primitive situation calculus actions into complex actions (aka *programs*) δ . These complex actions simply serve as constraints upon the situation tree. Complex action constructs include the following:

nil – the empty program
 a – primitive action
 $\phi?$ – test action
 $\pi x. \delta$ – nondeterministic choice of argument
 $\delta_1; \delta_2$ – sequences (δ_1 is followed by δ_2)
 $\delta_1 | \delta_2$ – nondeterministic choice between δ_1 and δ_2
if ϕ **then** δ_1 **else** δ_2 **endif** – conditional
while ϕ **do** δ **endW** – loop
proc $P(\vec{v})$ δ **endProc** – procedure

We also include the construct **anyorder** $[\delta_1, \dots, \delta_n]$ which is encoded as the nondeterministic choice of all possible permutations of the sequencing of $\delta_1, \dots, \delta_n$. The conditional and while-loop constructs are defined in terms of other constructs. For the purposes of Web service composition we generally treat iteration as finitely bounded by a parameter k . Such finitely bounded programs are called *tree programs*.

if ϕ **then** δ_1 **else** δ_2 **endif** $\stackrel{\text{def}}{=} [\phi?; \delta_1] \mid [\neg\phi?; \delta_2]$
while $_1(\phi)$ δ **endWhile** $\stackrel{\text{def}}{=} \text{if } \phi \text{ then } \delta \text{ endif}^1$
while $_k(\phi)$ δ **endWhile** $\stackrel{\text{def}}{=} \text{if } \phi \text{ then } [\delta; \text{while}_{k-1}(\phi)\delta \text{endWhile}] \text{ endif}$

These constructs can be used to write programs in the language of the domain theory, or more specifically, they can be used to specify both composite Web services and also generic procedures for Web service composition. E.g.²,

```
bookAirTicket( $\vec{x}$ ) ; if far then bookRentalCar( $\vec{y}$ )
else bookTaxi( $\vec{y}$ ) endif
bookRentalCar( $\vec{x}$ ) ; bookHotel( $\vec{y}$ ).
```

In order to understand how we modify Golog to incorporate user preferences, the reader must understand the basics of Golog semantics. There are two popular semantics for Golog programs: the original evaluation semantics [20] and a related single-step transition semantics that was proposed for on-line execution of concurrent Golog programs [5]. The transition semantics is axiomatized through two predicates $Trans(\delta, s, \delta', s')$ and $Final(\delta, s)$. Given an action theory \mathcal{D} , a program δ and a situation s , $Trans$ defines the set of possible successor configurations (δ', s') according to the action theory. $Final$ defines whether a given program has legally terminated in a given situation. $Trans$ and $Final$ are defined for every complex action. A few examples follow. (See [5] for details):

$$\begin{aligned} Trans(nil, s, \delta', s') &\equiv False \\ Trans(a, s, \delta', s') &\equiv Poss(a[s], s) \wedge \delta' = nil \\ &\quad \wedge s' = do(a[s], s) \\ Trans(\phi?, s, \delta', s') &\equiv \phi[s] \wedge \delta' = nil \wedge s' = s \\ Trans([\delta_1; \delta_2], s, \delta', s') &\equiv Final(\delta_1, s) \wedge Trans(\delta_2, s, \delta', s') \\ &\quad \vee \exists \delta'' . \delta' = (\delta''; \delta_2) \wedge Trans(\delta_1, s, \delta'', s') \\ Trans([\delta_1 \mid \delta_2], s, \delta', s') &\equiv Trans(\delta_1, s, \delta', s') \\ &\quad \vee Trans(\delta_2, s, \delta', s') \\ Trans(\pi(x)\delta, s, \delta', s') &\equiv \exists x. Trans(\delta_x^v, s, \delta', s') \\ Final(nil, s) &\equiv TRUE \\ Final(a, s) &\equiv FALSE \\ Final([\delta_1; \delta_2], s) &\equiv Final(\delta_1, s) \wedge Final(\delta_2, s) \end{aligned}$$

Thus, given the program $bookCar(\vec{x}); bookHotel(\vec{y})$, if the action $bookCar(\vec{x})$ is possible in situation s , then

$$Trans([bookCar(\vec{x}); bookHotel(\vec{y})], s, bookHotel(\vec{y}), do(bookCar(\vec{x}), s))$$

describes the only possible transition according to the action theory. $do(bookCar(\vec{x}), s)$ is the transition and $bookHotel(\vec{y})$ is the remaining program to be executed. Using the transitive closure of $Trans$, denoted $Trans^*$, one can define a *Do* predicate as follows. This *Do* is equivalent to the original evaluation semantics Do [5].

$$Do(\delta, s, s') \stackrel{\text{def}}{=} \exists \delta'. Trans^*(\delta, s, \delta', s') \wedge Final(\delta', s'). \quad (1)$$

Given a domain theory, \mathcal{D} and Golog program δ , program execution must find a sequence of actions \vec{a} (where \vec{a} is a vector of actions) such that: $\mathcal{D} \models Do(\delta, S_0, do(\vec{a}, S_0))$. $Do(\delta, S_0, do(\vec{a}, S_0))$ denotes that the Golog program δ , starting execution in S_0 will legally terminate in situation $do(\vec{a}, S_0)$, where $do(\vec{a}, S_0)$ abbreviates

¹**if-then-endif** is the obvious variant of **if-then-else-endif**.

²Following convention we will generally refer to fluents in situation-suppressed form, e.g., $at(Toronto)$ rather than $at(Toronto, s)$. Reintroduction of the situation term is denoted by $[s]$. Variables are universally quantified unless otherwise noted.

$do(a_n, do(a_{n-1}, \dots, do(a_1, S_0)))$). Thus, given a generic procedure, described as a Golog program δ , and an initial situation S_0 , we would like to infer a terminating situation $do(\vec{a}, S_0)$ such that the vector \vec{a} denotes a sequence of Web services that can be performed to realize the generic procedure.

3 Specifying User Preferences

In this section, we describe the syntax of the first-order language we use for specifying user preferences. This description follows the LPP language we proposed in [1, 2] for preference-based planning. The semantics of the language is described in the situation calculus. We provide an informal description here, directing the reader to [1, 2] for further details. Our language is richly expressive, enabling the expression of static as well as temporal preferences, and action-centric as well as state-centric preferences. Unlike many preference languages, it induces a total order over the compositions that differ with respect to the preferences, which avoids the high degree of incomparability experienced by many other non-quantitative preference languages, and simplifies computation of preferred compositions. Our language is qualitative, rather than ordinal or quantitative. Unlike many ordinal preference languages, our language provides a facility to stipulate the relative strength of preferences. We claim that its qualitative nature facilitates elicitation.

Illustrative example: To help illustrate our preference language, consider the task of travel planning. A generic procedure, easily specified in Golog, might say: *In any order, book inter-city transportation, book local accommodations and book local transportation.* With this generic procedure in hand an individual user can specify their hard constraints (e.g., *Lara needs to be in Chicago July 29-Aug 5, 2009.*) together with a list of preferences described in the language to follow.

To understand the preference language, consider the composition we are trying to generate to be a situation – a sequence of actions or Web services executed from the initial situation. A user specifies his or her preferences in terms of a single, so-called *General Preference Formula*. This formula is an aggregation of preferences over constituent properties of situations (i.e., compositions). The basic building block of our preference formula is a *Trajectory Property Formula* which describes properties of (partial) situations (i.e., compositions).

Definition 1 (Trajectory Property Formula (TPF))

A trajectory property formula is a sentence drawn from the smallest set \mathcal{B} where:

1. $\mathcal{F} \subset \mathcal{B}$
2. $\mathcal{R} \subset \mathcal{B}$
3. $f \in \mathcal{F}$, then $final(f) \in \mathcal{B}$
4. If $a \in \mathcal{A}$, then $occ(a) \in \mathcal{B}$
5. If φ_1 and φ_2 are in \mathcal{B} , then so are $\neg\varphi_1$, $\varphi_1 \wedge \varphi_2$, $\varphi_1 \vee \varphi_2$, $(\exists x)\varphi_1$, $(\forall x)\varphi_1$, **next**(φ_1), **always**(φ_1), **eventually**(φ_1), and **until**(φ_1, φ_2).

final(f) states that fluent f holds in the final situation, **occ**(a) states that action a occurs in the present situation, and

next(φ_1), **always**(φ_1), **eventually**(φ_1), and **until**(φ_1, φ_2) are basic LTL constructs.

TPFs establish properties of preferred situations (i.e., compositions of services). By combining TPFs using boolean connectives we are able to express a wide variety of properties of situations. E.g.³

$$\mathbf{final}(at(Home)) \quad (P1)$$

$$(\exists \vec{c}).\mathbf{occ}'(bookAir(\vec{c}, Economy, Direct)) \wedge member(\vec{c}, StarAlliance) \quad (P2)$$

$$\mathbf{always}(\neg((\exists \vec{h}).hotelBooked(\vec{h}) \wedge hilton(\vec{h}))) \quad (P3)$$

$$(\exists \vec{h}, \vec{r}).(\mathbf{occ}'(bookHotel(\vec{h}, \vec{r})) \wedge starsGE(\vec{r}, 3) \wedge paymentOption(\vec{h}, Visa)) \quad (P4)$$

P1 states that the user is at home in the final situation. P2 states that at some point the user books a direct economy flight with a Star Alliance carrier. Recall there was no stipulation in the generic procedure regarding the mode of transportation between cities or locally. P3 states that a Hilton hotel never be booked while P4 states that at some point the user books a hotel that accept Visa credit cards and has a rating of 3 or more.

Atomic Preference Formulae (APFs) define a preference ordering over alternative properties of situations. Each alternative being ordered comprises two components: the property of the situation, specified by a TPF, and a *value* term which stipulates the relative strength of the preference.

Definition 2 (Atomic Preference Formula (APF))

Let \mathcal{V} be a totally ordered set with minimal element v_{min} and maximal element v_{max} . An atomic preference formula is a formula $\varphi_0[v_0] \gg \varphi_1[v_1] \gg \dots \gg \varphi_n[v_n]$, where each φ_i is a TPF, each $v_i \in \mathcal{V}$, $v_i < v_j$ for $i < j$, and $v_0 = v_{min}$. When $n = 0$, atomic preference formulae correspond to TPFs.

An APF expresses a preference over alternatives. Note that v_{min} is the most preferred and v_{max} is the least preferred. In what follows, we let $\mathcal{V} = [0, 1]$, but we could instead choose a strictly qualitative set like $\{best < good < indifferent < bad < worst\}$ since the operations on these values are limited to *max* and *min*. The following APFs express an ordering over Lara's preferences.

$$P2[0] \gg (\exists \vec{c}, \vec{w}).\mathbf{occ}'(bookAir(\vec{c}, Economy, \vec{w})) \wedge member(\vec{c}, StarAlliance)[0.2] \gg \mathbf{occ}'(bookAir(Delta, Economy, Direct))[0.5] \quad (P5)$$

$$(\exists \vec{t}).\mathbf{occ}'(bookCar(National, \vec{t}))[0] \gg (\exists \vec{t}).\mathbf{occ}'(bookCar(Alamo, \vec{t}))[0.2] \gg (\exists \vec{t}).\mathbf{occ}'(bookCar(Avis, \vec{t}))[0.8] \quad (P6)$$

$$(\exists \vec{c}).\mathbf{occ}'(bookCar(\vec{c}, SUV))[0] \gg (\exists \vec{c}).\mathbf{occ}'(bookCar(\vec{c}, Compact))[0.2] \quad (P7)$$

P5 states that Lara prefers direct economy flights with a Star Alliance carrier, followed by economy flights with

³To simplify the examples many parameters have been suppressed. For legibility, variables are bold faced, **eventually**(**occ**(φ)) is abbreviated to **occ'**(φ), and preference formulae are referenced by their labels.

a Star Alliance carrier, followed by direct economy flights with Delta airlines. P6 and P7 are preference over cars. Lara strongly prefers National and then Alamo over Avis, followed by all other car-rental companies. Finally she slightly prefers an SUV over a compact with any other type of car a distant third.

General Preference Formulae (GPFs) extend our language to conditional, conjunctive, and disjunctive preferences, allowing the user aggregate APFs in nontrivial ways.

Definition 3 (General Preference Formula (GPF))

A formula Φ is a general preference formula if one of the following holds:

- Φ is an APF
- Φ is $\gamma : \Psi$, where γ is a TPF and Ψ is a GPF [Conditional]
- Φ is one of
 - $\Psi_0 \& \Psi_1 \& \dots \& \Psi_n$ [General Conjunction]
 - $\Psi_0 \mid \Psi_1 \mid \dots \mid \Psi_n$ [General Disjunction]

where $n \geq 1$ and each Ψ_i is a GPF.

Continuing our example:

$$(\forall \vec{h}, \vec{c}, \vec{e}, \vec{w}). \text{always}(\neg \text{hotelBooked}(\vec{h}) : \neg \text{occ}'(\text{bookAir}(\vec{c}, \vec{e}, \vec{w}))) \quad (\text{P8})$$

$$\text{far} : P5 \quad (\text{P9})$$

$$P3 \& P4 \& P6 \& P7 \& P8 \& P9 \quad (\text{P10})$$

P8 states that Lara prefers not to book her air ticket until she has a hotel booked. P9 conditions Lara's airline preferences on her destination being far away. (If it is not far, she will not fly and the preferences are irrelevant.) Finally, P10 aggregates previous preferences into one formula.

Semantics: Informally, the semantics of our preference language is achieved through assigning a weight to a situation s with respect to a GPF, Φ , written $w_s(\Phi)$. This weight is a composition of its constituents. For TPFs, a situation s is assigned the value v_{min} if the TPF is satisfied in s , v_{max} otherwise. Recall that in our example above $v_{min} = 0$ and $v_{max} = 1$, though they could equally well have been a qualitative e.g., [excellent, abysmal]. Similarly, given an APF, and a situation s , s is assigned the weight of the best TPF that it satisfies within the defined APF. Returning to our example above, for P6 if a situation (composition) booked a car from Alamo rental car, it would get a weight of 0.2. Finally GPF semantics follow the natural semantics of boolean connectives. As such General Conjunction yields the maximum of its constituent GPF weights and General Disjunction yields the minimum of its constituent GPF weights. For a full explanation of the situation calculus semantics, please see [1]. Here we also define further aggregations that can be performed. These are mostly syntactic sugar that are compelling to the user and we omit them for space.

We conclude this section with the following definition which shows us how to compare two situations (and thus two compositions) with respect to a GPF:

Definition 4 (Preferred Situations) A situation s_1 is at least as preferred as a situation s_2 with respect to a GPF Φ , written $\text{pref}(s_1, s_2, \Phi)$ if $w_{s_1}(\Phi) \leq w_{s_2}(\Phi)$.

4 Web Service Composition

In this section, we define the notion of Web Service Composition (WSC) with generic procedures and customizing user preferences, present an algorithm for computing these compositions and prove properties of our algorithm. Our definition relies on the definition of Do from (1) in Section 2.

Definition 5 (WSC w/ User Preferences (WSCP))

A Web service composition problem with user preferences is described as a 5-tuple $(\mathcal{D}, O, \delta, C, \Phi)$ where:

- \mathcal{D} is a situation calculus basic action theory describing functional properties of the Web services,
- O is a FOL theory describing the non-functional properties of the Web services⁴,
- δ is a generic procedure described in Golog,
- C is a formula expressing hard user constraints, and
- Φ is a GPF describing user preferences.

A Web Service Composition (WSC) is a sequence of Web services \vec{a} such that

$$\mathcal{D} \wedge O \models \exists s. Do(\delta, S_0, s) \wedge s = do(\vec{a}, S_0) \wedge C(s)$$

A preferred WSC (WSCP) is a sequence of Web services \vec{a} such that

$$\mathcal{D} \wedge O \models \exists s. Do(\delta, S_0, s) \wedge s = do(\vec{a}, S_0) \wedge C(s) \wedge \nexists s'. [Do(\delta, S_0, s') \wedge C(s') \wedge \text{pref}(s', s, \Phi)]$$

A WSC is a sequence of Web services, \vec{a} , whose execution starting in the initial situation enforces the generic procedure and hard constraints terminating successfully in $do(\vec{a}, s)$. A WSCP yields a most preferred terminating situation.

4.1 Computing Preferred Compositions

A Golog program places constraints on the situation tree that evolves from S_0 . As such, any implementation of Golog is effectively doing planning in a constrained search space, searching for a legal termination of the Golog program. The actions that define this terminating situation are the plan. In the case of composing web services, this plan is a web service composition.

To compute a preferred composition, WSCP, we search through this same constrained search space to find the *most preferred* terminating situation. Our approach, embodied in a system called GologPref, searches for this optimal terminating situation by modifying the PPLAN approach to planning with preferences proposed in [1]. In particular, GologPref performs best-first search through the constrained search space resulting from the Golog program, $\delta; C$. The search is guided by an admissible evaluation function that evaluates partial plans with respect to whether they satisfy the preference formula, Φ . The admissible evaluation function is the optimistic evaluation of the preference formula, with the pessimistic evaluation and the plan length used as tie breakers where necessary, in that order.

The preference formula Φ and the constraints C are evaluated over intermediate situations (partial compositions) by exploiting *progression* as described in [1]. Informally, progression takes a situation and a temporal logic formula (TLF), evaluates the TLF with respect to the state of the

⁴the content of \mathcal{D} and O would typically come from an OWL-S, SWSO, or other semantic Web service ontology.

situation, and generates a new formula representing those aspects of the TLF that remain to be satisfied in subsequent situations.

```

GologPref(init, pgm, pref)
frontier ← INITFRONTIER(init, pgm, pref)
while frontier ≠ ∅
  current ← REMOVEFIRST(frontier)
  % establishes current values for progPgm, partialPlan,
  % state, progPref
  if progPgm=nil and optW=pessW
    return partialPlan, optW
  end if
  neighbours ← EXPAND(progPgm, partialPlan,
    state, progPref)
  frontier ← SORTNMERGEBYVAL(neighbours, frontier)
end while
return [], ∞

EXPAND(progPgm, partialPlan, state, progPref) returns a list of
new nodes to add to the frontier. If partialPlan=nil then EXPAND
returns []. Otherwise, EXPAND uses Golog's Trans to
determine all the executable actions that are legal transitions of
progPgm in state and to compute the remaining program for each.
It returns a list which contains, for each of these executable actions
a a node (optW, pessW, newProgPgm, newPartialPlan, new-
State, newProgPref) and for each a leading to a terminating state,
a second node (realW, realW, nil, newPartialPlan, newState, new-
ProgPref).

```

Figure 1: A sketch of the GologPref algorithm.

Fig 1 provides a sketch of the basic GologPref algorithm following from PPLAN. The full GologPref algorithm takes as input a 5-tuple $(\mathcal{D}, O, \delta, C, \Phi)$. For ease of explication, our algorithm sketch in Fig 1 explicitly identifies the initial situation of \mathcal{D} , *init*, the Golog program, δ ; C which we refer to as *pgm* and Φ , which we refer to as *pref*. GologPref returns a sequence of Web services, i.e. a plan, and the weight of that plan. The *frontier* is a list of nodes of the form $[optW, pessW, pgm, partialPlan, state, pref]$, sorted by optimistic weight, pessimistic weight, and then by length. The frontier is initialized to the input program and the empty partial plan, its *optW*, *pessW*, and *pref* corresponding to the progression and evaluation of the input preference formula in the initial state.

On each iteration of the **while** loop, GologPref removes the first node from the frontier and places it in *current*. If the Golog program of *current* is *nil* then the situation associated with this node is a terminating situation. If it is also the case that *optW=pessW*, then GologPref returns *current*'s partial plan and weight. Otherwise, it calls the function **EXPAND** with *current*'s node as input.

EXPAND returns a new list of nodes to add to the frontier. If *progPgm* is *nil* then no new nodes are added to the frontier. Otherwise, **EXPAND** generates a new set of nodes of the form $[optW, pessW, prog, partialPlan, state, pref]$, one for each action that is a legal Golog transition of *pgm* in *state*. For actions leading to terminating states, **EXPAND** also generates a second node of the same form but with *optW* and *pessW* replaced by the actual weight achieved by the plan.

The new nodes generated by **EXPAND** are then sorted by *optW, pessW*, then length and merged with the remainder of the frontier. If we reach the empty frontier, we exit the **while** loop and return the empty plan.

We now prove the correctness of our algorithm.

Theorem 1 (Soundness and Optimality)

Let $\mathcal{P}=(\mathcal{D}, O, \delta, C, \Phi)$ be a Web service composition problem, where δ is a tree program. Let \vec{a} be the plan returned by Golog-Pref from input \mathcal{P} . Then \vec{a} is a WSCP of $(\mathcal{D}, O, \delta, C, \Phi)$.

Proof sketch: We prove that the algorithm terminates appealing to the fact that δ is a tree program. Then we prove that \vec{a} is a WSC by cases over *Trans* and *Final*. Finally we prove that \vec{a} is also optimal, by exploiting the correctness of progression of preference formulae proven in [1], the admissibility of our evaluation function, and the bounded size of the search space generated by the Golog program δ ; C .

4.2 Integrated Optimal Web Service Selection

Most Web service composition systems use AI planning techniques and as such generally ignore the important problem of Web service selection or discovery, assuming it will be done by a separate matchmaker. The work presented here is significant because it enables the selection of services for composition based, not only on their inputs, outputs, pre-conditions and effects but also based on other nonfunctional properties. As such, users are able to specify properties of services that they desire along side other properties of their preferred solution, and services are selected that optimize for the users preferences in the context of the overall composition.

To see how selection of services can be encoded in our system, we reintroduce the service parameter \vec{u} which was suppressed from the example preferences in Section 3. Revisiting P2, we see how the selection of a service \vec{u} is easily realized within our preference framework with preference P2'.

$$\begin{aligned}
(\exists \vec{c}, \vec{u}). \mathbf{occ}'(\mathit{bookAir}(\vec{c}, \mathit{Economy}, \mathit{Direct}, \vec{u})) \\
\wedge \mathit{member}(\vec{c}, \mathit{StarAlliance}) \\
\wedge \mathit{serviceType}(\vec{u}, \mathit{AirTicketVendor}) \\
\wedge \mathit{sellsTickets}(\vec{u}, \vec{c}) \quad (\text{P2}')
\end{aligned}$$

P2' causes GologPref to prefer booking air tickets with an air ticket vendor that sells the tickets of a carrier that is a member of Star Alliance.

5 Implementation and Application

We have implemented the generation of Web Service compositions using generic procedures and customizing user preferences as described in previous sections. Our implementation, GologPref, builds on an implementation of PPLAN[1] and an implementation of IndiGolog [20] both in SWI Prolog⁵.

GologPref interfaces with Web services through the implementation of domain-specific scrapers developed using

⁵See [20] for a description of the translation of \mathcal{D} to Prolog.

AgentBuilder 3.2, and AgentRunner 3.2, Web agent design applications developed by Fetch Technologies ©. Among the sites we have scraped are Mapquest, and several air, car and hotel services. The information gathered is collected in XML and then processed by GologPref.

We tested GologPref in the domain of travel planning. Our tests serve predominantly as a proof of the concept and to illustrate the utility of GologPref.

Our generic procedure which is represented in Golog was very simple, allowing flexibility in how it could be instantiated. What follows is an example of the Prolog encoding of a GologPref generic procedure.

```
anyorder[bookAcc,bookCityToCityTrans,bookLocalTrans]

proc(bookAcc(Location, Day, Num),
  [ stayWithFriends(Location) |
    bookHotel(Location, Day, Num) ]).

proc(bookLocalTrans(Location, SDay, RDay),
  [ getRide(Location, SDay, RDay) |
    walk(Location) |
    bookCar(Location, SDay, RDay) ]).

proc(bookCityToCityTrans(Orig, Des, SDay, RDay),
  [ getRide(Orig, Des, SDay, RDay) |
    bookAir(Orig, Des, SDay, RDay) |
    bookCar(Orig, Des, SDay, RDay) ]).
```

We tested our GologPref generic procedure with 3 different user profiles: Jack the impoverished university student, Lara the picky frequent flyer, and Conrad the corporate executive who likes timely luxury travel. Each user lived in Toronto and wanted to be in Chicago for specific days. A set of rich user preferences were defined for each user along the lines of those illustrated in Section 3. These preferences often required access to different Web information, such as driving distances.

Not surprisingly, in all cases, GologPref found the optimal WSC for the user. Compositions varied greatly ranging from Jack who arranged accommodations with friends; checked out the distance to his local destinations and then arranged his local transportation (walking since his local destination was close to where he was staying); then once his accommodations were confirmed, booking an economy air ticket Toronto-Chicago with one stop on US Airways with Expedia. Lara on the other hand, booked a hotel (not Hilton), booked an intermediate-sized car with National, and a direct economy air ticket with Star Alliance partner Air Canada via the Air Canada Web site. The optimality and the diversity of the compositions, all from the same generic procedure, illustrate the flexibility afforded by the WSCP approach.

Figure 2 shows the number of nodes expanded relative to the search space size for 6 test scenarios. The full search space represents all possible combinations of city-to-city transportation, accommodations and local transportation available to the users which could have been considered. These results illustrate the effectiveness of the heuristic used to find optimal compositions.

CASE #	NODES		TIME (SEC)	FULL SEARCH SPACE
	EXPANDED	CONSIDERED		
1	104	1700	14.38	28,512
2	102	1647	13.71	28,512
3	27	371	2.06	28,512
4	27	368	2.09	28,512
5	99	1692	14.92	28,512
6	108	1761	14.97	28,512

Figure 2: Test results for 6 scenarios run under 64bit Ubuntu Linux with 2.66 GHz CPU.

6 Summary and Related Work

In this paper we argued that the integration of user preferences into Web service composition was a key missing component of Web service composition. Building on our previous framework for Web service composition via generic procedures [14] and our work on preference-based planning [1], we proposed a system for Web service composition with user preferences. Key contributions of this paper include: characterization of the task of Web service composition with generic procedures and user preferences, provision of a previously developed language for specifying user preferences, provision of the GologPref algorithm that integrates preference-based reasoning into Golog, a proof of the soundness and optimality of GologPref with respect to the user’s preferences, and a working implementation of our GologPref algorithm. A notable side effect of our framework is the seamless integration of Web service selection with the composition process.

We tested GologPref on 6 diverse scenarios applied to the same generic procedure. Results illustrated the diversity of compositions that could be generated from the same generic procedure. The number of nodes expanded by the heuristic search was several orders of magnitude smaller than the grounded search space, illustrating the effectiveness of the heuristic and the Golog program in guiding search.

A number of researchers have advocated using AI planning techniques to address the task of Web service composition including using regression-based planners [13], planners based on model checking (e.g., [28]), highly optimized hierarchical task network (HTN) planners such as SHOP2 (e.g., [21]), and a combination of classical and HTN planning called XPLAN [9]. Like Golog, HTNs afford the user the ability to define a generic procedure or *template* of how to perform a task.

Sirin et al. incorporated simple service preferences into the SHOP2 HTN planner to achieve dynamic service binding [22]. Their preference language is significantly less expressive than the one presented here and is restricted to the task of service selection rather than solution optimization. Nevertheless, it is a promising start. Also related is the work by Fritz and the third author in which they *precompiled* a subset of the preference language presented here into Golog programs that were then integrated with a decision-theoretic Golog (DTGolog) program [6]. The main objective of this work was to provide a means of integrating qualitative and quantitative preferences for agent programming. While both

used a form of Golog, the form and processing of preferences was quite different.

Since the original publication of this work, preference-based planning has been the subject of much interest, spurred on in great part by three tracks on planning with preferences at the 2006 International Planning Competition (IPC-5). A number of preference-based planners were developed, including one by a subset of the authors, all based on the competition's PDDL3 language [7]. The most notable new work that is directly related to this paper is that of [11]. In this paper, the authors propose a prototype HTN preference-based planner, **SCUP**, tailored to the task of Web service composition and that uses PDDL3 as its preference specification language.

We also have two follow-up pieces of work [24] and [23] in which we specify flexible templates in the form of an HTN rather than a Golog generic procedure. In [24] we proposed a qualitative language very similar to the preference language discussed in this paper but tailored to HTN planning. In [23] we extended PDDL3 with HTN-specific preference constructs. The proposed planners employ state of the art heuristic guided search and algorithms that exploit HTN-specific preferences and control. In contrast to the work presented here, optimality is not guaranteed without exhaustive search. In future work, we would like to improve the GologPref algorithm with the addition of more informative inadmissible heuristics coupled with branch and bound search. We would also like to exploit a recent extension to the LPP preference language to include preferences over the occurrence of Golog complex actions [2].

References

- [1] M. Bienvenu, C. Fritz, and S. McIlraith. Planning with qualitative temporal preferences. In *KR*, pages 134–144, 2006.
- [2] M. Bienvenu, C. Fritz, and S. McIlraith. Specifying and generating preferred plans. 2009. Submitted for publication.
- [3] J. De Bruijn, H. Lausen, A. Polleres, and D. Fensel. The web service modeling language WSM^L: An overview. Technical report, DERI, 2006.
- [4] D. Calvanese, G. De Giacomo, M. Lenzerini, M. Mecella, and F. Patrizi. Automatic service composition and synthesis: the Roman Model. *IEEE Data Eng. Bull.*, 31(3):18–22, 2008.
- [5] G. De Giacomo, Y. Lespérance, and H. Levesque. ConGolog, a concurrent programming language based on the situation calculus. *AIJ*, 121(1–2):109–169, 2000.
- [6] C. Fritz and S. McIlraith. Decision-theoretic golog with qualitative preferences. In *KR*, pages 153–163, Lake District, UK, June 2006.
- [7] A. Gerevini and D. Long. Plan constraints and preferences for PDDL3. Technical Report 2005-08-07, University of Brescia, 2005.
- [8] I. Horrocks, P. Patel-Schneider, and F. van Harmelen. From *SHIQ* and RDF to OWL: The making of a web ontology language. *Journal of Web Semantics*, 1(1):7–26, 2003.
- [9] M. Klusch, A. Gerber, and M. Schmidt. Semantic web service composition planning with OWLS-Xplan. In *AAAI-05 Fall Symposium*, 2005.
- [10] C. A. Knoblock, S. Minton, J. Luis Ambite, M. Muslea, J. Oh, and M. Frank. Mixed-initiative, multi-source information assistants. In *WWW*, pages 697–707, 2001.
- [11] N. Lin, U. Kuter, and E. Sirin. Web service composition with user preferences. In *ESWC-08*, pages 629–643, 2008.
- [12] D. Martin, M. Burstein, D. McDermott, S. McIlraith, M. Paolucci, K. Sycara, D. McGuinness, E. Sirin, and N. Srinivasan. Bringing semantics to web services with OWL-S. *World Wide Web Journal*, 10(3):243–277, 2007.
- [13] D. V. McDermott. Estimated-regression planning for interactions with web services. In *AIPS*, pages 204–211, 2002.
- [14] S. McIlraith and T. Son. Adapting golog for composition of semantic web services. In *KR*, pages 482–493, Toulouse, France, April 22-25 2002.
- [15] S. McIlraith, T. Son, and H. Zeng. Semantic Web services. In *IEEE Intelligent Systems (Special Issue on the Semantic Web)*, volume 16, March/April 2001.
- [16] Sheila McIlraith and Ronald Fadel. Planning with complex actions. In *NMR-02*, pages 356–364, 2002.
- [17] K. Myers. Strategic advice for hierarchical planners. In *Principles of Knowledge Representation and Reasoning: Proc. of the 5th International Conference (KR96)*, pages 112–123. Morgan Kaufmann, 1996.
- [18] K. Myers. Planning with conflicting advice. In *AIPS*, pages 355–362, 2000.
- [19] S. Narayanan and S. McIlraith. Simulation, verification and automated composition of web services. In *Proc. of the 11th International World Wide Web Conference (WWW-2002)*, May 2002.
- [20] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, Cambridge, MA, 2001.
- [21] E. Sirin, B. Parsia, and J. Hendler. Template-based composition of semantic web services. In *AAAI-05 Fall Symposium on Agents and the Semantic Web*, 2005.
- [22] E. Sirin, B. Parsia, D. Wu, J. Hendler, and D. Nau. HTN planning for web service composition using SHOP2. *Journal of Web Semantics*, 1(4):377–396, 2005.
- [23] S. Sohrabi, J. Baier, and S. McIlraith. HTN planning with preferences. In *IJCAI*, 2009.
- [24] S. Sohrabi and S. A. McIlraith. On planning with preferences in HTN. In *12th International Workshop on Non-Monotonic Reasoning (NMR-08)*, pages 241–248, Sydney, Australia, 2008.
- [25] S. Sohrabi and S. A. McIlraith. Optimizing web service composition while enforcing regulations. In *ISWC*, 2009. To appear.
- [26] S. Sohrabi, N. Prokoshyna, and S. A. McIlraith. Web service composition via generic procedures and customizing user preferences. In *ISWC*, pages 597–611, 2006.
- [27] S. Sohrabi, N. Prokoshyna, and S. A. McIlraith. Web service composition via the customization of Golog programs with user preferences. In *Conceptual Modeling: Foundations and Applications: Essays in Honor of John Mylopoulos*, pages 319–334. Springer-Verlag, 2009.
- [28] P. Traverso and M. Pistore. Automatic composition of semantic web services into executable processes. In *ISWC*, 2004.