

# Specification and Verification of Data-driven Web Services

Alin Deutsch\*

Liyang Sui\*

Victor Vianu\*

## Abstract

We study data-driven Web services provided by Web sites interacting with users or applications. The Web site can access an underlying database, as well as state information updated as the interaction progresses, and receives user input. The structure and contents of Web pages, as well as the actions to be taken, are determined dynamically by querying the underlying database as well as the state and inputs. The properties to be verified concern the sequences of events (inputs, states, and actions) resulting from the interaction, and are expressed in linear or branching-time temporal logics. The results establish under what conditions automatic verification of such properties is possible and provide the complexity of verification. This brings into play a mix of techniques from logic and automatic verification.

## 1 Introduction

Web services, viewed broadly as interactive Web applications providing access to information as well as transactions, are typically powered by databases. They have a strong dynamic component and are governed by protocols of interaction with users or programs, ranging from the low-level input-output signatures used in WSDL [27], to high-level workflow specifications (e.g., see [8, 6, 10, 26, 28, 18]). One central issue is to develop static analysis techniques to increase confidence in the robustness and correctness of complex Web services. This paper presents new verification techniques for Web services, and investigates the trade-off between the expressiveness of the Web service specification language and the feasibility of verification tasks.

In the scenario we consider, a Web service is provided by an interactive Web site that posts data, takes input from the user, and responds to the input by posting more data and/or taking some actions. The Web site can access an underlying database, as

well as state information updated as the interaction progresses. The structure of the Web page the user sees at any given point is described by a Web page schema. The contents of a Web page is determined dynamically by querying the underlying database as well as the state. The actions taken by the Web site, and transitions from one Web page to another, are determined by the input, state, and database.

The properties we wish to verify about Web services involve the sequences of inputs, actions, and states that may result from interactions with a user. This covers a wide variety of useful properties. For example, in a Web service supporting an e-commerce application, it may be desirable to verify that no product is delivered before payment of the right amount is received. Or, we may wish to verify that the specification of Web page transitions is unambiguous, (the next Web page is uniquely defined at each point in a run), that each Web page is reachable from the home page, etc. To express such properties, we rely on temporal logic. Specifically, we consider two kinds of properties. The first requires that all runs must satisfy some condition on the sequence of inputs, actions, and states. To describe such properties we use a variant of linear-time temporal logic. Other properties involve several runs simultaneously. For instance, we may wish to check that for every run leading to some Web page, there exists a way to return to the home page. To capture such properties, we use variants of the branching-time logics CTL and CTL\*.

Our results identify classes of Web services for which temporal properties in the above temporal logics can be checked, and establish their complexity. As justification for the choice of these classes, we show that even slight relaxations of our restrictions lead to undecidability of verification. Thus, our decidability results are quite tight.

**Related work** Our notion of Web service is a fairly broad one. It encompasses a large class of data-intensive Web applications equipped (implicitly or explicitly) with workflows that regulate the interaction between different partners who can be users, WSDL-style Web services, Web sites, programs and databases. We address the verification of proper-

---

\*Address: Computer Science and Engineering, UC San Diego, La Jolla, CA 92093-0114. Email: {deutsch, lsui, vianu}@cs.ucsd.edu. Supported in part by the NSF under grant number ITR-0225676 (SEEK).

ties pertaining to the runs of these workflows. Our model is related to WebML [8], a high-level conceptual model for data-driven Web applications, extended in [7] with workflow concepts to support process modeling. It is also related to other high-level workflow models geared towards Web applications (e.g. [6, 10, 28]), and more broadly to general workflows (see [26, 16, 17, 11, 4, 25]), whose focus is however quite different from ours. Non-interactive variants of Web page schemas have been proposed in prior projects such as Strudel [13], Araneus [22] and Weave [14], which target the automatic generation of Web sites from an underlying database.

More broadly, our research is related to the general area of automatic verification, and particularly reactive systems [21, 20]. Directly relevant to us is Spielmann’s work on Abstract State Machine (ASM) transducers [23, 24]. Similarly to the earlier relational transducers [3] these model database-driven reactive systems that respond to input events by taking some action, and maintain state information in designated relations. Our model of Web services is considerably more complex than ASM transducers (a detailed comparison is provided in the appendix). The techniques developed in [23, 24] remain nonetheless relevant, and we build upon them to obtain our decidability results on the verification of linear-time properties of Web services (by reducing the verification problem to finite satisfiability of E+TC formulas<sup>1</sup>). For the results on branching-time properties we use a mix of techniques from finite-model theory and temporal logic (see [12]), as well as automata-theoretic model-checking techniques developed by Kupferman, Vardi, and Wolper [19].

The paper is organized as follows. Section 2 introduces our model and specification language for Web sites. Section 3 considers the verification of linear-time temporal properties, and Section 4 focuses on branching-time properties. The main body of the paper is self contained. Due to space limitations, it emphasizes informal presentation of the results and extensive examples intended to facilitate understanding of the formal framework. An appendix is provided for the reader’s convenience, containing additional background material and proofs of most results. In addition, a demo Web site implementing our running example is available at <http://www.cs.ucsd.edu/~lsui/project/index.html>.

<sup>1</sup>E+TC is existential first-order logic augmented with a transitive closure operator, see Appendix.

## 2 Web Service Specifications

In this section we provide our model and specification language for data-driven Web services. Our model of a Web service captures the interaction of an external user<sup>2</sup> with the Web site, referred to as a “run”. Informally, a Web service specification has the following components:

- A database that remains fixed throughout each run;
- A set of state relations that change throughout the run in response to user inputs;
- A set of Web page schemas, of which one is designated as the “home page”, and another as an “error page”;
- Each Web page schema defines how the set of current input choices is generated as a query on the database and states. In addition, it specifies the state transition in response to the user’s input, the actions to be taken, as well as the next Web page schema.

Intuitively, a run proceeds as follows. First, the user accesses the home page, and the state relations are initialized to empty. When accessed, each Web page generates a choice of inputs for the user, by a query on the database and states. All input options are generated by the system except for a fixed set that represents specific user information (e.g. name, password, credit card number, etc). These are represented as constants in the input schema, whose interpretations are provided by the user throughout the run as requested. The user chooses at most one tuple among the options provided for each input. In response to this choice, a state transition occurs, actions are taken, and the next Web page schema is determined, all according to the rules of the specification. As customary in verification, we assume that all runs are infinite (finite runs can be easily represented as infinite runs by fake loops).

We now formalize the above notion of Web service. We assume fixed an infinite set of elements  $\mathbf{dom}_\infty$ . A *relational schema* is a finite set of relation symbols with associated arities, together with a finite set of constant symbols. Relation symbols with arity zero are also called propositions. A relational instance over a relational schema consists of a finite subset  $Dom$  of  $\mathbf{dom}_\infty$ , and a mapping associating to each relation symbol of positive arity a finite relation of the same arity, to each propositional symbol a

<sup>2</sup>We use the term “user” generically to stand for any partner interacting with the Web site, be it an actual user, program, another Web service, etc.

truth value, and to each constant symbol an element of  $Dom$ . We use several kinds of relational schemas, with different roles in the specification of the Web service.

We assume familiarity with first-order logic (FO) over relational vocabularies. We adopt here an active domain semantics for FO formulas, as commonly done in database theory (e.g., see [2]).

**Definition 2.1** A Web service  $\mathcal{W}$  is a tuple  $(\mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon)$ , where:

- $\mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}$  are relational schemas called database, state, input, and action schemas, respectively. The sets of relation symbols of the schemas are disjoint (but they may share constant symbols). We refer to constants in  $\mathbf{I}$  as input constants, and denote them by  $\mathbf{const}(\mathbf{I})$ .
- $\mathbf{W}$  is a finite set of Web page schemas.
- $W_0 \in \mathbf{W}$  is the home page schema, and  $W_\epsilon \notin \mathbf{W}$  is the error page schema.

We also denote by  $\mathbf{Prev}_{\mathbf{I}}$  the relational vocabulary  $\{prev_I \mid I \in \mathbf{I} - \mathbf{const}(\mathbf{I})\}$ , where  $prev_I$  has the same arity as  $I$  (intuitively,  $prev_I$  refers to the input  $I$  at the previous step in the run).

A Web page schema  $W$  is a tuple  $(\mathbf{I}_W, \mathbf{A}_W, \mathbf{T}_W, \mathcal{R}_W)$  where  $\mathbf{I}_W \subseteq \mathbf{I}, \mathbf{A}_W \subseteq \mathbf{A}, \mathbf{T}_W \subseteq \mathbf{W}$ . Then  $\mathcal{R}_W$  is a set of rules containing the following:

- For each input relation  $I \in \mathbf{I}_W$  of arity  $k > 0$ , an input rule  $Options_I(\bar{x}) \leftarrow \varphi_{I,W}(\bar{x})$  where  $Options_I$  is a relation of arity  $k$ ,  $\bar{x}$  is a  $k$ -tuple of distinct variables, and  $\varphi_{I,W}(\bar{x})$  is an FO formula over schema  $\mathbf{D} \cup \mathbf{S} \cup \mathbf{Prev}_{\mathbf{I}} \cup \mathbf{const}(\mathbf{I})$ , with free variables  $\bar{x}$ .
- For each state relation  $S \in \mathbf{S}$ , one, both, or none of the following state rules:
  - an insertion rule  $S(\bar{x}) \leftarrow \varphi_{S,W}^+(\bar{x})$ ,
  - a deletion rule  $\neg S(\bar{x}) \leftarrow \varphi_{S,W}^-(\bar{x})$ ,

where the arity of  $S$  is  $k$ ,  $\bar{x}$  is a  $k$ -tuple of distinct variables, and  $\varphi_{S,W}^\epsilon(\bar{x})$  are FO formulas over schema  $\mathbf{D} \cup \mathbf{S} \cup \mathbf{Prev}_{\mathbf{I}} \cup \mathbf{const}(\mathbf{I}) \cup \mathbf{I}_W$ , with free variables  $\bar{x}$ .

- For each action relation  $A \in \mathbf{A}_W$ , an action rule  $A(\bar{x}) \leftarrow \varphi(\bar{x})$  where the arity of  $A$  is  $k$ ,  $\bar{x}$  is a  $k$ -tuple of distinct variables, and  $\varphi(\bar{x})$  is an FO formula over schema  $\mathbf{D} \cup \mathbf{S} \cup \mathbf{Prev}_{\mathbf{I}} \cup \mathbf{const}(\mathbf{I}) \cup \mathbf{I}_W$ , with free variables  $\bar{x}$ .

- for each  $V \in \mathbf{T}_W$ , a target rule  $V \leftarrow \varphi_{V,W}$  where  $\varphi_{V,W}$  is an FO sentence over schema  $\mathbf{D} \cup \mathbf{S} \cup \mathbf{Prev}_{\mathbf{I}} \cup \mathbf{const}(\mathbf{I}) \cup \mathbf{I}_W$ .

Finally,  $W_\epsilon = (\emptyset, \emptyset, \{W_\epsilon\}, \mathcal{R}_{W_\epsilon})$  where  $\mathcal{R}_{W_\epsilon}$  consists of the rule  $W_\epsilon \leftarrow true$ .

Intuitively, the action rules of a Web page specify the actions to be taken in response to the input. The state rules specify the tuples to be inserted or deleted from state relations (with conflicts given no-op semantics, as seen below). If no rule is specified in a Web page schema for a given state relation, the state remains unchanged. The input rules specify a set of options to be presented to users, from which they can pick at most one tuple to input (this feature corresponds to menus in user interfaces). At every point in time,  $I$  contains the current input tuple, and  $prev_J$  contains the input to  $J$  in the previous step of the run (if any). The choice of this semantics for  $prev_J$  relations is somewhat arbitrary, and other choices are possible without affecting the results. For example, another possibility is to have  $prev_J$  hold the *most recent* input to  $J$  occurring anywhere in the run, rather than in the previous step. Also note that  $prev_J$  relations are really state relations with very specific functionality, and are redundant in the general model. However, they are very useful when defining tractable restrictions of the model.

**Notation** For better readability of our examples, we use the following notation: relation  $R$  is displayed as  $R$  if it is a state relation, as  $\mathbf{R}$  if it is an input relation, as  $\underline{R}$  if it is a database relation, and as  $\overline{R}$  if it is an action relation. In Example 2.2 below,  $error \in \mathbf{S}$ ,  $user \in \mathbf{D}$  and  $name, password, button \in \mathbf{I}$ .

**Example 2.2** We use as a running example throughout the paper an e-commerce Web site selling computers online. New customers can register a name and password, while returning customers can login, search for computers fulfilling certain criteria, add the results to a shopping cart, and finally buy the items in the shopping cart. A demo Web site implementing this example, together with its full specification, is provided at <http://www.cs.ucsd.edu/~lsui/project/index.html>.

The demo site implements the Web service  $(\mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, \mathbf{HP}, W_\epsilon)$ . See Figure 2 in the appendix for an overview of all Web pages of the demo, depicted in WebML style. We list here only the pages in  $\mathbf{W}$  that are mentioned in the running example:

HP	the home page
RP	the new user registration page
CP	the customer page
AP	the administrator page
LSP	a laptop search page
PIP	displays the products returned by the search
CC	allows the user to view the cart contents and order items in it
MP	an error message page

The following describes the home page HP which contains two text input boxes for the customer's user name and password respectively, and three buttons, allowing customers to register, login, respectively clear the input.

Page HP

Inputs  $\mathbf{I}_{HP}$  : name, password, button( $x$ )

Input Rules:

$$\text{Options}_{\text{button}(x)} \leftarrow x = \text{"login"} \vee x = \text{"register"} \vee x = \text{"clear"}$$

State Rules:

$$\text{error}(\text{"failed login"}) \leftarrow \neg \text{user}(\text{name}, \text{password}) \wedge \text{button}(\text{"login"})$$

Target Web Pages  $\mathbf{T}_{HP}$ : HP, RP, CP, AP, MP

Target Rules:

$$\text{HP} \leftarrow \text{button}(\text{"clear"})$$

$$\text{RP} \leftarrow \text{button}(\text{"register"})$$

$$\text{CP} \leftarrow \text{user}(\text{name}, \text{password}) \wedge \text{button}(\text{"login"}) \wedge \text{name} \neq \text{"Admin"}$$

$$\text{AP} \leftarrow \text{user}(\text{name}, \text{password}) \wedge \text{button}(\text{"login"}) \wedge \text{name} = \text{"Admin"}$$

$$\text{MP} \leftarrow \neg \text{user}(\text{name}, \text{password}) \wedge \text{button}(\text{"login"})$$

End Page HP

Notice how the three buttons are modeled by a single input relation `button`, whose argument specifies the clicked button. The corresponding input rule restricts it to a login, clear or register button only. As will be seen shortly (Definition 2.3), each input relation may contain at most one tuple at any given time, corresponding to the user's pick from the set of tuples defined by the associated input rule. This guarantees that no two buttons may be clicked simultaneously. The user name and password are modeled as input constants, as their value is not supposed to change during the session. If the login button is clicked, the first state rule looks up the name/password combination in the database table `user`. If the lookup fails, the second state rule records the login failure in the state relation `error`, and the last target rule fires a transition to the message page MP. Notice how the "Admin" user enjoys special treatment: upon login, she is directed to the admin page AP, whereas all

other users go to the customer page CP. Assume that the CP page allows users to follow either a link to a desktop search page, or a laptop search page LSP. We illustrate only the laptop search functionality of the search page LSP (see the online demo for the full version, which also allows users to search for desktops).

Page LSP

Inputs  $\mathbf{I}_{LSP}$  : laptopsearch( $ram, hdisk, display$ ), button( $x$ )

Input Rules:

$$\text{Options}_{\text{button}(x)} \leftarrow x = \text{"search"} \vee x = \text{"view cart"} \vee x = \text{"logout"}$$

$$\text{Options}_{\text{laptopsearch}(r, h, d)} \leftarrow \text{criteria}(\text{"laptop"}, \text{"ram"}, r) \wedge \text{criteria}(\text{"laptop"}, \text{"hdd"}, h) \wedge \text{criteria}(\text{"laptop"}, \text{"display"}, d)$$

State Rules:

$$\text{userchoice}(r, h, d) \leftarrow \text{laptopsearch}(r, h, d) \wedge \text{button}(\text{"search"})$$

Target Web Pages  $\mathbf{T}_{LSP}$ : HP, PIP, CC

Target Rules:

$$\text{HP} \leftarrow \text{button}(\text{"logout"})$$

$$\text{PIP} \leftarrow \exists r \exists h \exists d \text{laptopsearch}(r, h, d) \wedge \text{button}(\text{"search"})$$

$$\text{CC} \leftarrow \text{button}(\text{"view cart"})$$

End Page SP

Notice how the second input rule looks up in the database the valid parameter values for the search criteria pertinent to laptops. This enables users to pick from a menu of legal values instead of providing arbitrary ones.  $\square$

We next define the notion of "run" of a Web service. Essentially, a run specifies the fixed database and consecutive Web pages, states, inputs, and actions. Thus, a run over database instance  $D$  is an infinite sequence  $\{\langle V_i, S_i, I_i, P_i, A_i \rangle\}_{i \geq 0}$ , where  $V_i \in \mathbf{W}$ ,  $S_i$  is an instance of  $\mathbf{S}$ ,  $I_i$  is an instance of  $\mathbf{I}_{V_i}$ ,  $P_i$  is an instance of  $\mathbf{prev}_{\mathbf{I}}$ , and  $A_i$  is an instance of  $\mathbf{A}_{V_i}$ . In particular, the input constants play a special role. Their interpretation is not fixed a priori, but is instead provided by the user as the run progresses. We will need to make sure this occurs in a sound fashion. For example, a formula may not use an input constant before its value has been provided. We will also prevent the Web service from asking the user repeatedly for the value of the same input constant. To formalize this, we will use the following notation. For each  $i \geq 0$ ,  $\kappa_i$  denotes the set of input constants occurring in some  $\mathbf{I}_{V_j}$  in the run,  $j \leq i$ , and  $\sigma_i$  denotes the mapping associating to each  $c \in \kappa_i$  the unique  $I_j(c)$  where  $j \leq i$  and  $c \in \mathbf{I}_{V_j}$ .

**Definition 2.3** Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  be a Web service and  $D$  a database instance over schema  $\mathbf{D}$ . A run of  $\mathcal{W}$  for database  $D$  is an infinite

sequence  $\{(V_i, S_i, I_i, P_i, A_i)\}_{i \geq 0}$  where  $V_i \in \mathbf{W}$ ,  $S_i$  is an instance of  $\mathbf{S}$ ,  $I_i$  is an instance of  $\mathbf{I}_{V_i}$ ,  $P_i$  is an instance of  $\mathbf{prev}_I$ ,  $A_i$  is an instance of  $\mathbf{A}_{V_i}$  and:

- $V_0 = W_0$ , and  $S_0, A_0, P_0$  are empty;
- for each  $i \geq 0$ ,  $V_{i+1} = W_\epsilon$  if one of the following holds:
  - (i) some formula used in a rule of  $V_i$  involves a constant  $c \in \mathbf{I}$  that is not in  $\kappa_i$ ;
  - (ii)  $\mathbf{I}_{V_i} \cap \kappa_{i-1} \neq \emptyset$ ;
  - (iii) there are distinct  $W, W' \in \mathbf{T}_{V_i}$  for which  $\varphi_{W, V_i}$  and  $\varphi_{W', V_i}$  are both true when evaluated on  $D, S_i, I_i$  and  $P_i$ , and interpretation  $\sigma_i$  for the input constants occurring in the formulas;

Otherwise,  $V_{i+1}$  is the unique  $W \in \mathbf{T}_{V_i}$  for which  $\varphi_{W, V_i}$  is true when evaluated on  $D, S_i, I_i, P_i$  and  $\sigma_i$  if such  $W$  exists; if not,  $V_{i+1} = V_i$ .

- for each  $i \geq 0$ , and for each relation  $R$  in  $\mathbf{I}_{V_i}$  of arity  $k > 0$ ,  $I_i(R) \subseteq \{v\}$  for some  $v \in \text{Options}_R$ , where  $\text{Options}_R$  is the result of evaluating  $\varphi_{R, V_i}$  on  $D, S_i, P_i$  and  $\sigma_i$ ;
- for each  $i \geq 0$ , and for each proposition  $R$  in  $\mathbf{I}_{V_i}$ ,  $I_i(R)$  is a truth value;
- for each  $i \geq 0$ , and for each constant  $c$  in  $\mathbf{I}_{V_i}$ ,  $I_i(c)$  is an element in  $\mathbf{dom}_\infty$ ;
- for each  $i \geq 0$ , and for each relation  $\text{prev}_I$  in  $\mathbf{prev}_I$ ,  $P_i(\text{prev}_I) = I_{i-1}(I)$  if  $I \in \mathbf{I}_{V_{i-1}}$  and  $P_i(\text{prev}_I)$  is empty otherwise.
- for each  $i \geq 0$ , and relation  $S$  in  $\mathbf{S}$ ,  $S_{i+1}(S)$  is the result of evaluating

$$\begin{aligned} & (\varphi_{S, V_i}^+(\bar{x}) \wedge \neg \varphi_{S, V_i}^-(\bar{x})) \vee \\ & (S(\bar{x}) \wedge \varphi_{S, V_i}^-(\bar{x}) \wedge \varphi_{S, V_i}^+(\bar{x})) \vee \\ & (S(\bar{x}) \wedge \neg \varphi_{S, V_i}^-(\bar{x}) \wedge \neg \varphi_{S, V_i}^+(\bar{x})) \end{aligned}$$

on  $D, S_i, I_i, P_i$  and  $\sigma_i$ , where  $\varphi_{S, V_i}^\epsilon(\bar{x})$  is taken to be false if it is not provided in the Web page schema ( $\epsilon \in \{+, -\}$ ). In particular,  $S$  remains unchanged if no insertion or deletion rule is specified for it.

- for each  $i \geq 0$ , and relation  $A$  in  $\mathbf{A}_{V_{i+1}}$ ,  $A_{i+1}(A)$  is the result of evaluating  $\varphi_{A, V_i}$  on  $D, S_i, I_i, P_i$  and  $\sigma_i$ .

Note that the state and actions specified at step  $i + 1$  in the run are those triggered at step  $i$ . This

choice is convenient for technical reasons. As discussed above, input constants are provided an interpretation as a result of user input, and need not be values already existing in the database. Once an interpretation is provided for a constant, it can be used in the formulas determining the run. For example, such constants might include **name**, **password**, **credit-card**, etc. The error Web page serves an important function, since it signals behavior that we consider anomalous. Specifically, the error Web page is reached in the following situations: (i) the value of an input constant is required by some formula before it was provided by the user; (ii) the user is asked to provide a value for the same input constant more than once; and, (iii) the specification of the next Web page is ambiguous, since it produces more than one Web page. Once the error page is reached, the run loops forever in that page. We call a run *error free* if the error Web page is not reached, and we call a Web service error-free if it generates only error-free runs. Clearly, it would be desirable to verify that a given Web service is error-free. As we will see, this can be expressed in the temporal logics we consider and can be checked for restricted classes of Web services.

### 3 Verifying linear-time temporal properties

In this section we consider the verification of properties that must be satisfied by *all* runs of a Web service. Such properties are expressed using a variant of linear-time temporal logic, adapted from [12, 1, 24]. We begin by defining this logic, that we denote LTL-FO (first-order linear temporal logic).

**Definition 3.1** [12, 1, 24] *The language LTL-FO (first-order linear-time temporal logic) is obtained by closing FO under negation, disjunction, and the following formula formation rule: If  $\varphi$  and  $\psi$  are formulas, then  $\mathbf{X}\varphi$  and  $\varphi\mathbf{U}\psi$  are formulas. Free and bound variables are defined in the obvious way. The universal closure of an LTL-FO formula  $\varphi(\bar{x})$  with free variables  $\bar{x}$  is the formula  $\forall \bar{x}\varphi(\bar{x})$ . An LTL-FO sentence is the universal closure of an LTL-FO formula.*

Note that quantifiers cannot be applied to formulas containing temporal operators, except by taking the universal closure of the entire formula, yielding an LTL-FO sentence.

Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  be a Web service. To express properties of runs of  $\mathcal{W}$ , we use LTL-FO sentences over schema  $\mathbf{DUSUIUPrev}_I\mathbf{UAUW}$ ,

where each  $W \in \mathbf{W}$  is used as a propositional variable. The semantics of LTL-FO formulas is standard, and we describe it informally. Let  $\forall \bar{x}\varphi(\bar{x})$  be an LTL-FO sentence over the above schema. The Web service  $\mathcal{W}$  satisfies  $\forall \bar{x}\varphi(\bar{x})$  iff every run of  $\mathcal{W}$  satisfies it. Let  $\rho = \{\alpha_i\}_{i \geq 0}$  be a run of  $\mathcal{W}$  for database  $D$ , and let  $\rho_j$  denote  $\{\alpha_i\}_{i \geq j}$ , for  $j \geq 0$ . Note that  $\rho = \rho_0$ . Let  $\text{Dom}(\rho)$  be the active domain of  $\rho$ , i.e. the set of all elements occurring in relations or as constants in  $\rho$ . The run  $\rho$  satisfies  $\forall \bar{x}\varphi(\bar{x})$  iff for each valuation  $\nu$  of  $\bar{x}$  in  $\text{Dom}(\rho)$ ,  $\rho_0$  satisfies  $\varphi(\nu(\bar{x}))$ . The latter is defined by structural induction on the formula. An FO sentence  $\psi$  is satisfied by  $\alpha_i = \langle V_i, S_i, I_i, P_i, A_i \rangle$  if the following hold:

- the set of input constants occurring in  $\psi$  is included in  $\kappa_i$ ;
- the structure  $\alpha'_i$  satisfies  $\psi$ , where  $\alpha'_i$  is the structure obtained by augmenting  $\alpha_i$  with interpretation  $\sigma_i$  for the input constants. Furthermore,  $\alpha_i$  assigns *true* to  $V_i$  and *false* to all other propositional symbols in  $\mathbf{W}$ .

The semantics of Boolean operators is the obvious one. The meaning of the temporal operators  $\mathbf{X}$ ,  $\mathbf{U}$  is the following (where  $\models$  denotes satisfaction and  $j \geq 0$ ):

- $\rho_j \models \mathbf{X}\varphi$  iff  $\rho_{j+1} \models \varphi$ ,
- $\rho_j \models \varphi \mathbf{U} \psi$  iff  $\exists k \geq j$  such that  $\rho_k \models \psi$  and  $\rho_l \models \varphi$  for  $j \leq l < k$ .

Observe that the above temporal operators can simulate all commonly used operators, including  $\mathbf{B}$  (before),  $\mathbf{G}$  (always) and  $\mathbf{F}$  (eventually). Indeed,  $\varphi \mathbf{B} \psi$  is equivalent to  $\neg(\neg\varphi \mathbf{U} \neg\psi)$ ,  $\mathbf{G}\varphi \equiv \text{false} \mathbf{B} \varphi$ , and  $\mathbf{F}\varphi \equiv \text{true} \mathbf{U} \varphi$ . We use the above operators as shorthand in LTL-FO formulas whenever convenient.

LTL-FO sentences can express many interesting properties of a Web service. A useful class of properties pertains to the navigation between pages.

**Example 3.2** The following property states that whenever page  $P$  is reached in a run, page  $Q$  will be eventually reached as well:

$$\mathbf{G}(\neg P) \vee \mathbf{F}(P \wedge \mathbf{F} Q) \quad (1)$$

□

Another important class of properties describes the flow of the interaction between user and service.

**Example 3.3** Assume that the Web service in Example 2.2 allows the user to pick a product and records the pick in a state relation

$\text{pick}(\text{product\_id}, \text{price})$ . There is also a payment page  $\text{PP}$ , with input relation  $\text{pay}(\text{amount})$  and “authorize payment” button. Clicking this button authorizes the payment of  $\text{amount}$  for the product with identifier recorded in state  $\text{pick}$ , on behalf of the user whose name was provided by the constant  $\text{name}$  (recall page  $\text{HP}$  from Example 2.2). Also assume the existence of an order confirmation page  $\text{OCP}$ , containing the actions  $\text{conf}(\text{user\_id}, \text{price})$  and  $\text{ship}(\text{user\_id}, \text{product\_id})$ . The following property involving state, action, input and database relations requires that any shipped product be previously paid for:

$$\forall \text{pid}, \text{price} [\xi(\text{pid}, \text{price}) \mathbf{B} \neg(\text{conf}(\text{name}, \text{price}) \wedge \text{ship}(\text{name}, \text{pid}))] \quad (2)$$

where  $\xi(\text{pid}, \text{price})$  is the formula

$$\begin{aligned} & \text{PP} \wedge \text{pay}(\text{price}) \wedge \text{button}(\text{“authorize payment”}) \\ & \quad \wedge \text{pick}(\text{pid}, \text{price}) \\ & \quad \wedge \exists \text{pname} \text{ catalog}(\text{pid}, \text{price}, \text{pname}) \quad (3) \end{aligned}$$

□

We consider the following verification scenario. Given a Web service, we would first like to verify, as a minimum soundness check, that it is error free. If it is, we may wish to verify some additional temporal properties expressed by LTL-FO sentences.

It is easily seen that it is undecidable if a Web service is error free, or if it satisfies a LTL-FO formula, using Trakhtenbrot’s theorem. To obtain decidability, we must restrict both the Web services and the LTL-FO sentences. We use a restriction proposed in [23, 24] for ASM transducers, limiting the use of quantification in state, action, and target rule formulas to “input-bounded” quantification, and limiting formulas of input rules to be existential. The restriction is formulated in our framework as follows. Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  be a Web service. The set of *input-bounded* FO formulas over the schema  $\mathbf{D} \cup \mathbf{S} \cup \mathbf{I} \cup \mathbf{A} \cup \mathbf{W} \cup \text{Prev}_\mathbf{I}$  is obtained by replacing in the definition of FO the quantification formation rule by the following:

- if  $\varphi$  is a formula,  $\alpha$  is a current or previous input atom using a relational symbol from  $\mathbf{I} \cup \text{Prev}_\mathbf{I}$ ,  $\bar{x} \subseteq \text{free}(\alpha)$ , and  $\bar{x} \cap \text{free}(\beta) = \emptyset$  for every state or action atom  $\beta$  in  $\varphi$ , then  $\exists \bar{x}(\alpha \wedge \varphi)$  and  $\forall \bar{x}(\alpha \rightarrow \varphi)$  are formulas.

A Web service is input-bounded iff all formulas in state, action, and target rules are input bounded, and all input rules use  $\exists^*$ FO formulas in which all

state atoms are ground. An LTL-FO sentence over the schema of  $\mathcal{W}$  is input-bounded iff all of its FO subformulas are input-bounded.

**Example 3.4** All rules on pages HP,SP in Example 2.2 are input-bounded. Property (1) in Example 3.2 is trivially input-bounded, as it contains no quantifiers. Property (2) in Example 3.3, however, is not input-bounded because  $pname$  appears in no input atom. We turn this into an input-bounded property by modeling the catalog database relation with two relations prod\_prices( $pid, price$ ) and prod\_names( $pid, pname$ ). We can now rewrite Property (2) to the input-bounded sentence

$$\forall pid, price [\zeta'(pid, price) \mathbf{B} \neg(\overline{\text{conf}}(\text{name}, price) \wedge \overline{\text{ship}}(\text{name}, pid))] \quad (4)$$

where  $\zeta'(pid, price)$  is short for

$$\text{PP} \wedge \text{pay}(price) \wedge \text{button}(\text{"authorize payment"}) \wedge \text{pick}(pid, price) \wedge \text{prod\_prices}(pid, price) \quad (5)$$

□

We can now state the main result of this section:

**Theorem 3.5** *The following are decidable:*

- (i) *given a Web service  $\mathcal{W}$  with input-bounded rules, whether it is error free;*
- (ii) *given an error-free Web service  $\mathcal{W}$  with input-bounded rules and an input-bounded LTL-FO sentence  $\varphi$  over the schema of  $\mathcal{W}$ , whether  $\mathcal{W}$  satisfies  $\varphi$ .*

Furthermore, both problems are PSPACE-complete for schemas with fixed bound on the arity, and in EX-PSPACE for schemas with no fixed bound on the arity.

The proof of Theorem 3.5 is outlined in the appendix. The lower bound is an easy reduction from Quantified Boolean Formula [15]. The upper bound is much more involved, and is based on a reduction to the satisfiability problem for the logic E+TC (existential FO augmented with a transitive closure operator). This requires modifying and extending an analogous reduction used in [23, 24] for ASM transducers (see appendix).

**Remark 3.6 (Sessions and Database Updates)** *Recall that our model prohibits database updates within a run. However, in practice it is very useful to update the database at various points in the interaction with users. In our running example, it*

*makes sense to do so when a user logs out. In fact, the Web site implementing the full example specifies database update rules triggered at logout. This implicitly assumes that the runs to be verified consist of interactions of a single user between login and logout. Indeed, these are natural boundaries of sessions to be verified, and can be specified implicitly within the temporal formula to be verified. However, other definitions of sessions are possible (see also the discussion in the Conclusions).*

**Boundaries of decidability** One may wonder whether our restrictions can be relaxed without affecting decidability of verification. Unfortunately, even small relaxations of these restrictions can lead to undecidability. Specifically, we consider the following: (i) relaxing the requirement that state atoms be ground in formulas defining input options, by allowing state atoms with variables, (ii) relaxing the input-bounded restriction by allowing a very limited form of non input-bounded quantification in the form of state projections, (iii) allowing  $prev_I$  relations to record *all* previous inputs to  $I$  rather than just the most recent one, and (iv) extending LTL-FO formulas with path quantification.

We begin with extension (i) and show undecidability even for a fixed LTL-FO formula and input options defined by quantifier-free FO formulas using just database and state relations.

**Theorem 3.7** *There exists a fixed input-bounded LTL-FO formula  $\varphi$  for which it is undecidable, given an input-bounded Web service  $\mathcal{W}$  with input options defined by quantifier-free FO formulas over database and state relations, whether  $\mathcal{W} \models \varphi$ .*

**Proof:** The proof is by reduction of the question of whether a Turing Machine halts on input  $\epsilon$ , see Appendix. □

We next consider extension (ii): we relax input-boundedness of rules by allowing projections of state relations. We call a Web service *input-bounded with state projections* if all its formulas are input-bounded, excepting state rules that allow insertions of the form:

$$S(\bar{x}) \leftarrow \exists \bar{y} S'(\bar{x}, \bar{y})$$

where  $S$  and  $S'$  are state relations. We can show the following.

**Theorem 3.8** *It is undecidable, given a simple, input-bounded Web service  $\mathcal{W}$  with state projections and input-bounded LTL-FO sentence  $\varphi$ , whether  $\mathcal{W} \models \varphi$ .*

**Proof:** The proof is by reduction of the implication problem for functional and inclusion dependencies, known to be undecidable [9], see Appendix.  $\square$

We now deal with extension (iii). We say that a Web service has *lossless input* if the  $prev_I$  relations record *all* previous inputs to  $I$  in the current run.

**Theorem 3.9** *It is undecidable, given an input-bounded Web service  $\mathcal{W}$  with lossless input and an input-bounded LTL-FO formula  $\varphi$ , whether  $\mathcal{W} \models \varphi$ .*

The proof uses a straightforward reduction of the undecidability of finite validity of FO formulas and is omitted.

The undecidability of extension (iv) is shown in the next section, after the notation on branching-time logics is introduced.

## 4 Branching-time properties

In this section we consider the verification of temporal properties of Web services involving quantification over runs. This allows expressing properties involving multiple runs of a Web service rather than just individual ones, such as “at any point in a run there is a way to return to the home page”. To specify such properties, we use variants of the logics CTL and CTL\* [12] extending the logic LTL-FO used in the previous section. These variants, denoted CTL-FO and CTL\*-FO, are defined in the appendix. They extend LTL-FO with path quantifiers **E** and **A**, subject to certain syntactic restrictions. Informally,  $\mathbf{E}\varphi$  stands for “there exists a continuation of the current run that satisfies  $\varphi$ ” and  $\mathbf{A}\varphi$  means “every continuation of the current run satisfies  $\varphi$ ”.

**Example 4.1** The following CTL\*-FO sentence expresses the fact that in every run, whenever a product  $pid$  is bought by a user, it will eventually ship, but until that happens, the user can still cancel the order for  $pid$ .

$$\forall pid \forall price \mathbf{AG}(\xi'(pid, price) \rightarrow A((\mathbf{EF}\overline{\text{cancel}}(\text{name}, pid))\mathbf{U}(\overline{\text{ship}}(\text{name}, pid))))$$

where  $\xi'$  is the formula defined in Example 3.4 (5).  $\square$

As noted earlier, the decidability results of the previous section do not extend to CTL<sup>(\*)</sup>-FO sentences, even if they are restricted to be input bounded (by requiring every FO subformula to be input bounded). Indeed, the following can be shown (see Appendix).

**Theorem 4.2** *It is undecidable, given a simple, input-bounded Web service  $\mathcal{W}$  and input-bounded CTL-FO sentence  $\varphi$ , whether  $\mathcal{W} \models \varphi$ .*

The proof further shows that a single alternation of path quantifiers is sufficient to yield undecidability, since one alternation is enough to express validity of FO sentences in the prefix class  $\exists^*\forall^*\text{FO}$ , known to be undecidable [5].

We next consider several restrictions leading to decidability of the verification problem for CTL<sup>(\*)</sup>-FO sentences.

### Propositional input-bounded Web services

The first restriction further limits input-bounded Web services by requiring all states and actions to be propositional. Furthermore, no rules can use  $\mathbf{Prev}_I$  atoms. We call such Web services *propositional*. In a propositional Web service, inputs can still be parameterized in the Web service specification. The CTL\* formulas we consider are propositional and use input, action, state, and Web page symbols, viewed as propositions. Satisfaction of such a CTL\* formula by a Web service is defined as for CTL\*-FO, where truth of propositional symbols in a given configuration  $\langle V, S, I, A \rangle$  is defined as follows: a Web page symbol is true iff it equals  $V$ , a state symbol  $s$  is true iff  $s \in S$ , an input symbol  $I$  is true iff  $I \in \mathbf{I}_V$ , and an action symbol  $a$  is true iff  $a \in A$ .

**Example 4.3** CTL<sup>(\*)</sup>-FO is particularly useful for specifying navigational properties of Web services. Note that these services do not necessarily have to be propositional; we could abstract their predicates to propositional symbols, thus concentrating only on reachability properties.<sup>3</sup> For our running example, abstracting all non-input atoms to propositions, we could ask whether from any page it is possible to navigate to the home page HP using the following CTL sentence:

$$\mathbf{AGEF}(\text{HP})$$

The following CTL property states that, after login, the user can reach a page where he can authorize payment for a product:<sup>4</sup>

$$\mathbf{AG}((\text{HP} \wedge \text{button}(\text{“login”})) \rightarrow \mathbf{EF}(\text{button}(\text{“authorize payment”})))$$

where  $\text{button}(\text{“login”})$ ,  $\text{button}(\text{“authorize payment”})$  denote the corresponding propositions. In the specification of the abstracted service, we can still allow in the home page HP a state rule that checks successful login:

$$\overline{\text{logged\_in}} \leftarrow \overline{\text{users}}(\text{name}, \text{password}) \wedge \text{button}(\text{“login”}).$$

<sup>3</sup>This is in the spirit of program verification, where program variables are first abstracted to booleans, in order to check CTL\* properties such as liveness.

<sup>4</sup>The most important property in electronic commerce ☺

□

For a given Web service  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$ , we denote by  $\Sigma_{\mathcal{W}}$  the propositional vocabulary consisting in all symbols in  $\mathbf{S} \cup \mathbf{I} \cup \mathbf{A} \cup \mathbf{W}$ . By abuse of notation, we use the same symbol for a relation  $R$  in the vocabulary of  $\mathcal{W}$  and for the corresponding propositional symbol in  $\Sigma_{\mathcal{W}}$ . We first show the following (the proof is provided in the appendix):

**Theorem 4.4** *Given a propositional, input-bounded Web service  $\mathcal{W}$  and a CTL\* formula  $\varphi$  over  $\Sigma_{\mathcal{W}}$ , it is decidable whether  $\mathcal{W} \models \varphi$ . The complexity of the decision procedure is CO-NEXPTIME if  $\varphi$  is in CTL, and EXPSPACE if  $\varphi$  is in CTL\*.*

The complexity of the decision problem of Theorem 4.4 can be decreased under additional assumptions. The following result (see Appendix) focuses on verification of navigational properties of Web sites, expressed by CTL\* formulas over alphabet  $\mathbf{W}$ .

**Corollary 4.5** *Let  $\mathbf{S}$  be a fixed set of state propositions and  $\mathbf{D}$  a fixed database schema. Given a propositional, input-bounded, error-free Web service  $\mathcal{W}$  with states  $\mathbf{S}$  and database schema  $\mathbf{D}$ , and a CTL\* formula  $\varphi$  over  $\mathbf{W}$ , it is decidable in PSPACE whether  $\mathcal{W} \models \varphi$ .*

Another special case of interest involves Web services that are entirely propositional. Thus, the database plays no role in the specification: inputs, states, and actions are all propositional, and the rules do not use the database. Let us call such a Web service *fully propositional*. We can show the following (see Appendix):

**Theorem 4.6** *Given a fully propositional Web service  $\mathcal{W}$  and a CTL\* formula  $\varphi$  over  $\Sigma_{\mathcal{W}}$ , it is decidable in PSPACE whether  $\mathcal{W} \models \varphi$ .*

One may wonder if the restrictions of Theorem 4.4 can be relaxed without compromising the decidability of verification. In particular, it would be of interest if one could lift some of the restrictions on the propositional nature of states and actions. Unfortunately, we have shown that allowing parameterized actions leads to undecidability of verification, even for CTL formulas whose only use of action predicates is to check emptiness. The proof is by reduction of the implication problem for functional and inclusion dependencies. We omit the details.

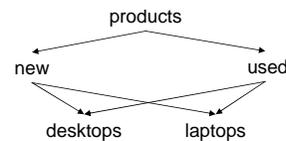


Figure 1: Fragment of  $R_I$  for Example 4.8

**Web services with input-driven search** The restrictions considered so far require states of a Web service to be propositional, and do not allow the use of  $\mathbf{Prev}_I$  atoms. Although adequate for some verification tasks, this is a serious limitation in many situations, since no values can be passed on from one Web page to another. We next alleviate some of this limitation by considering Web services that allow limited use of  $\mathbf{Prev}_I$  atoms. This can model commonly arising applications involving a user-driven search, going through consecutive stages of refinement. More formally:

**Definition 4.7** *A Web service with input-driven search is an input-bounded Web service  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  where:*

- $\mathbf{I}$  consists of a single unary relation  $I$
- $\mathbf{S}$  consists of propositional states including not-start
- $\mathbf{A}$  is propositional
- $\mathbf{D}$  includes a constant  $i_0$  and a designated binary relation  $R_I$
- the state rule for not-start is  $\text{not-start} \leftarrow \neg \text{not-start}$
- the input option rule for  $I$  is in all Web pages of the form

$$\begin{aligned} \text{Options}_I(y) &\leftarrow (\neg \text{not-start} \wedge y = i_0) \\ &\vee (\text{not-start} \wedge \exists x(\text{prev}_I(x) \wedge R_I(x, y)) \wedge \varphi(y)) \end{aligned}$$

where  $\varphi(y)$  is a quantifier-free formula over  $\mathbf{D} \cup \mathbf{S}$  with free variable  $y$ .

Note that *not-start* is false at the start of the computation and true thereafter. To initialize the search, the first input option is the constant  $i_0$ . Subsequently (when *not-start* is true), if  $x$  was the previously chosen input, the allowed next inputs are the  $y$ 's for which  $R_I(x, y) \wedge \varphi(y)$  holds, where  $R_I$  is the special input search relation and  $\varphi$  places some additional condition on  $y$  involving the database and the propositional states.

**Example 4.8** Consider a variation of a computer-selling Web site which doesn't just partition its products into desktops and laptops, but rather uses the more complex classification depicted in Figure 1. The user can search the hierarchy of categories, and will only see a certain category if it is currently in stock, as reflected by the database. The propositional state new is set on the page which offers the choice between new and used products. The page schemas for new and old computers are reused, so when generating the options, the Web site must consult state new to distinguish among new and old products. We can abstract this Web site as a Web service with input-driven search, in which the binary database relation  $R_I$  is a graph which contains as a subgraph the one in Figure 1, and in which the unary database relations such as newDesktop, usedDesktop, usedLaptop contain the in-stock products. Here is the input rule corresponding to the desktop search page:

$$\begin{aligned} Options_I(y) \leftarrow & (-not-start \wedge y = i_0) \vee \\ & not-start \wedge \exists x(prev_I(x) \wedge R_I(x, y)) \wedge \\ & (new \wedge \underline{newDesktop}(y) \vee \neg new \\ & \wedge \underline{usedDesktop}(y)) \end{aligned} \quad (6)$$

□

We can show the following.

**Theorem 4.9** *Given a Web service with input-driven search  $\mathcal{W}$  and a CTL\* formula  $\varphi$ , it is decidable whether  $\mathcal{W} \models \varphi$  in EXPTIME if  $\varphi$  is in CTL, and 2-EXPTIME if  $\varphi$  is in CTL\*.*

**Proof:** We reduce the problem of checking whether  $\mathcal{W} \models \varphi$  to the satisfiability problem for CTL<sup>(\*)</sup> formulas, known to be EXPTIME-complete for CTL and 2-EXPTIME complete for CTL\* [12], see Appendix. □

## 5 Conclusions

We have identified a practically appealing and fairly tight class of Web services and linear-time temporal formulas for which verification is decidable. The complexity of verification is PSPACE-complete (for fixed database arity). This is quite reasonable as static analysis goes<sup>5</sup>. For branching-time properties, we identified decidable restrictions for which the complexity of verification ranges from PSPACE to 2-EXPTIME. To obtain these results, we used a mix of techniques from logic and automatic verification.

<sup>5</sup>Recall that even testing inclusion of two conjunctive queries is NP-complete!

Other interesting aspects of Web service verification could not be addressed in this paper and are left for future work. We mention a few of them.

**Specifying and verifying sessions** As discussed in Section 3, in practical Web service applications it is not always realistic to assume that verification applies to *all* possible runs of the service. This may be due to various reasons: there may be a need to verify properties of complex services in a modular fashion, the restrictions needed for decidability may only hold for certain portions of runs, etc. Let us call portions of runs to be verified *sessions*. Some sessions can be specified implicitly within the temporal formula to be verified, while others may require explicit definition by other means. It is of interest to understand what types of sessions can be verified by our approach. For instance, in our running example, the default assumption is that sessions consist of single-user runs beginning at login and ending at logout. However, other types of sessions can be fit to our restrictions, including multi-user sessions (as long as no database updates occur within the session and only a bounded number of new users register).

**Interacting Web services** An important aspect of Web services is the interaction of multiple services and their composition into more complex services (as in e-service composition, see [18]). On the face of it, our model concerns the behavior of a single Web service interacting with its environment. However, it can also capture to some extent the interaction and composition of multiple Web services. For example, external calls to a service viewed as a black box can be modeled simply by an extra database relation with a limited access pattern. In terms of verification, certain properties of the sequence of messages exchanged by Web services (called *conversations* in the framework of WSDL [27]) can be specified using our temporal formulas. We plan to further explore to what extent interacting Web services can be modeled in our framework and their properties verified by our techniques.

### Algorithms and heuristics for verification

While our positive results provide complexity-theoretic upper bounds on Web service verification, significant work is still needed in order to obtain practical algorithms and heuristics. We plan to explore this issue, including the use of theorem-proving techniques in conjunction with the logic and automata-based approaches suggested in the paper.

## Acknowledgement

We wish to thank Caroline Cruz and Dayou Zhou for their help in implementing the demo Web site for our running example.

## References

- [1] S. Abiteboul, L. Herr, and J. V. den Bussche. Temporal versus first-order logic to query temporal databases. In *Proc. ACM PODS*, pages 49–57, 1996.
- [2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [3] S. Abiteboul, V. Vianu, B. Fordham, and Y. Yesha. Relational transducers for electronic commerce. *JCSS*, 61(2):236–269, 2000. Extended abstract in PODS 98.
- [4] A. J. Bonner and M. Kifer. An overview of transaction logic. *Theor. Comput. Sci.*, 133(2):205–265, 1994.
- [5] E. Borger, E. Gradel, and Y. Gurevich. *The Classical Decision Problem*. Springer, 1997.
- [6] BPML.org. Business process modeling language. <http://www.bpml.org>.
- [7] M. Brambilla, S. Ceri, S. Comai, P. Fraternali, and I. Manolescu. Specification and design of workflow-driven hypertexts. *Journal of Web Engineering*, 1(1), 2002.
- [8] S. Ceri, P. Fraternali, A. Bongio, M. Brambilla, S. Comai, and M. Matera. *Designing data-intensive Web applications*. Morgan-Kaufmann, 2002.
- [9] A. K. Chandra and M. Vardi. The implication problem for functional and inclusion dependencies is undecidable. *SIAM J. Comp.*, 14(3):671–677, 1985.
- [10] DAML-S Coalition (A. Ankolekar et al). DAML-S: Web service description for the semantic Web. In *The Semantic Web - ISWC*, pages 348–363, 2002.
- [11] H. Davulcu, M. Kifer, C. R. Ramakrishnan, and I. V. Ramakrishnan. Logic based modeling and analysis of workflows. In *Proc. ACM PODS*, pages 25–33, 1998.
- [12] E. A. Emerson. Temporal and modal logic. In J. V. Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*, pages 995–1072. North-Holland Pub. Co./MIT Press, 1990.
- [13] M. F. Fernández, D. Florescu, A. Y. Levy, and D. Suciu. Declarative specification of web sites with Strudel. *VLDB Journal*, 9(1):38–55, 2000.
- [14] D. Florescu, K. Yagoub, P. Valduriez, and V. Isarny. WEAVE: A data-intensive web site management system (software demonstration). In *Proc. of the Conf. on Extending Database Technology (EDBT)*, 2000.
- [15] M. R. Garey and D. S. Johnson. *Computers and Intractability*. Freeman, 1979.
- [16] D. Georgakopoulos, M. F. Hornick, and A. P. Sheth. An overview of workflow management: From process modeling to workflow automation infrastructure. *Distributed and Parallel Databases*, 3(2):119–153, 1995.
- [17] D. Harel. On the formal semantics of statecharts. In *Proc. LICS*, pages 54–64, 1987.
- [18] R. Hull, M. Benedikt, V. Christophides, and J. Su. E-Services: a look behind the curtain. In *Proc. ACM PODS*, pages 1–14, 2003.
- [19] O. Kupferman, M. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *J. of ACM*, 47(2):312–360, 2000.
- [20] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer Verlag, 1991.
- [21] Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer Verlag, 1995.
- [22] G. Mecca, P. Merialdo, and P. Atzeni. Araneus in the era of XML. *IEEE Data Engineering Bulletin*, 22(3):19–26, 1999.
- [23] M. Spielmann. Abstract State Machines: Verification problems and complexity. Ph.D. thesis, RWTH Aachen, 2000.
- [24] M. Spielmann. Verification of relational transducers for electronic commerce. *JCSS*, 66(1):40–65, 2003. Extended abstract in PODS 2000.

- [25] D. Wodtke and G. Weikum. A formal foundation for distributed workflow execution based on state charts. In *Proc. ICDT*, pages 231–246, 1997.
- [26] Workflow management coalition, 2001. <http://www.wfmc.org>.
- [27] Web Services Description Language(WSDL 1.1), 2001. <http://www.w3.org/TR/2001/NOTE-wsdl-20010315>.
- [28] Web Services Flow Language(WSFL 1.0), 2001. <http://www-3.ibm.com/software/solutions/webservices/pdf/WSFL.pdf>.

## A Appendix

In this appendix we briefly review background material on ASM transducers and temporal logics, and provide proofs for most of our results. Figure 2 represents the Web pages of our demo, depicted in WebML style.

### A.1 Background on ASM transducers and logic

**ASM transducers** Like Web services, and similarly to the earlier relational transducers [3], the ASM transducers studied by Spielmann in [23, 24] model database-driven reactive systems that respond to input events by taking some action, and maintain state information in designated relations. In terms of our framework, the ASM relational transducer can be viewed as a simplified Web service consisting of a single Web page. Like Web services, ASM transducers use database and state relations (called *memory* relations), as well as action and input relations. At each step, the transducer receives from the environment inputs consisting of arbitrary relations over the input vocabulary, whose elements come from the underlying database. The transducer reacts to the inputs with a state transition and by producing output relations. The control of the transducer is defined by rules similar to ours. The temporal properties to be verified are expressed by LTL-FO formulas.

In order to achieve decidability of verification, Spielmann considers several possible restrictions. The one of interest to us is *input-bounded ASM with bounded input flow*, denoted  $ASM^I$ . This requires the following: (i) each input relation received in any single step has cardinality bounded by a constant, and (ii) the rules used in the specification, as well as the LTL-FO formula to be verified, are input bounded.

The definition of input-bounded rule and formula are the same as ours, except that ASM rules use no **Prev<sub>I</sub>** atoms.

In summary, the main differences between our input-bounded Web services and  $ASM^I$  transducers are:

- Web services use multiple Web pages and specify transitions among them,
- Web service inputs are restricted by input options defined by certain  $\exists^*FO$  formulas,
- The input vocabulary of Web services may contain input constants whose values are progressively supplied by users and need not come from the database, and
- input-bounded Web services allow the use of **Prev<sub>I</sub>** atoms, treated the same as input atoms.

Spielmann’s proof of decidability of input-bounded LTL-FO properties of  $ASM^I$  transducers makes use of several logics for which finite satisfiability is decidable:

- $FO^W$ , the *witness-bounded* fragment of FO;
- $FO^W + \text{posTC}$ , the extension of  $FO^W$  with the positive occurrences of the transitive closure operator, and
- E+TC, the existential fragment of FO+TC.

The main idea of the proof is to reduce the problem of checking the existence of a run of the transducer violating the desired property to that of checking finite satisfiability of a formula in one of the above logics. Specifically, this is done by an ingenious polynomial reduction to the finite satisfiability problem for  $FO^W + \text{posTC}$ . Next, it is shown in [23] that finite satisfiability of  $FO^W + \text{posTC}$  is polynomially reducible to the finite satisfiability of E+TC, and the latter is in PSPACE for fixed database arity and in EXPSpace for arbitrary arity.

Before providing more details on the reductions, we briefly review the above logics, using the terminology of [24]. We start with some notation. A finite set of constant symbols and variables is also called a *witness set*. For a witness set  $W$  and a variable  $x$  not in  $W$ , let  $(x \in W)$  abbreviate the formula  $(\bigvee_{v \in W} x = v)$ . Intuitively,  $(x \in W)$  holds iff the interpretation of  $x$  matches the interpretation of some symbols in  $W$ .

**Definition A.1** *The witness-bounded fragment of FO, denoted by  $FO^W$ , is obtained from FO by replacing the formula-formation rule for first-order quantification with the following rules for witness-bounded quantification:*



**(WBQ)** If  $W$  is a witness set,  $x$  is a variable not in  $W$ , and  $\varphi$  is a formula, then  $(\exists x \in W)\varphi$  and  $(\forall x \in W)\varphi$  are formulas.

The free and bound variables of  $FO^W$  formulas are defined as usual. In particular,  $x$  occurs bound in  $(\exists x \in W)\varphi$  and  $(\forall x \in W)\varphi$ , whereas all variables in the witness set  $W$  are free. Thus,  $FO^W$  can be viewed as a fragment of FO where formulas of the form  $(\exists x \in W)\varphi$  and  $(\forall x \in W)\varphi$  are mere for  $\exists x(x \in W \wedge \varphi)$  and  $\forall x(x \in W \rightarrow \varphi)$ , respectively.

Let FO+TC denote FO augmented with the transitive closure operator TC, and E+TC denote the existential fragment of FO+TC.

**Definition A.2** [24] *The witness-bounded fragment of transitive-closure logic, denoted by  $(FO^W + TC)$ , is obtained from  $(FO+TC)$  by replacing the formula-formation rule for the first-order quantification with the rule **(WBQ)**. An occurrence of a TC operator in a  $(FO^W + TC)$  formula is called **positive** if the occurrence is in the scope of an even number of negation. By  $(FO^W + posTC)$  we denote the set of those  $(FO^W + TC)$  formulas in which every occurrence of a TC operator is positive.*

Let  $\mathbf{T}$  be an  $ASM^I$  transducer and  $\forall \bar{x}\varphi(\bar{x})$  the universal closure of an input-bounded LTL-FO formula. Clearly, every run of  $\mathbf{T}$  satisfies  $\forall \bar{x}\varphi(\bar{x})$  iff it is not the case that

(†) there exists a run of  $\mathbf{T}$  satisfying  $\psi \equiv \exists \bar{x}\neg\varphi(\bar{x})$

Thus, it is enough to solve (†). We outline the main steps of the reduction of (†) to the finite satisfiability of a  $FO^W + posTC$  sentence, provided in [23, 24].

Although (†) has the flavor of a satisfiability test, there is an immediate difficulty: transducer runs are infinite, whereas finite satisfiability involves finite structures. This is dealt with by the following:

**Periodic Run Lemma:**  $\mathbf{T}$  has a run satisfying  $\psi$  iff it has a periodic run satisfying  $\psi$ .

Since a periodic run can be represented by a finite prefix, such runs are representable by finite structures.

The next step provides a definition by an FO+TC formula of the finite structures representing periodic runs of  $\mathbf{T}$  that satisfy  $\psi$ . Intuitively, the formula describes the connection between consecutive configurations of the transducer by FO formulas and uses the transitive closure operator to describe the entire run and verify satisfaction of  $\psi$ . However, a difficulty arises: the formula uses universal quantification. This problem is alleviated by the following:

**Local Run Lemma:** *this states, intuitively, that an approximate description of the runs of  $\mathbf{T}$  is sufficient when checking satisfiability of  $\psi$ . Specifically, the description is exact on the inputs of the runs, but provides only the correct description of the restrictions of memory and action relations to a designated set  $C$  of constants. The set  $C$  consists of the database constants as well as constants standing for witnesses to the existentially quantified variables  $\bar{x}$  in  $\psi = \exists \bar{x}\neg\varphi$ . Such a “run” is called a local run of  $\mathbf{T}$ . The lemma shows that there exists a run of  $\mathbf{T}$  satisfies  $\psi$  iff there exists a local run of  $\mathbf{T}$  satisfying  $\psi$ . Thus, it is sufficient to consider only local runs of  $\mathbf{T}$  when checking satisfiability of  $\psi$ .*

The Local Run Lemma allows replacing the quantifiers of the FO+TC formula by witness-bounded quantifiers, with  $C$  serving as a witness set. With some work, this yields an equivalent formula in  $FO^W + posTC$  constructed in polynomial time from  $\mathbf{T}$  and  $\psi$ .

## A.2 Branching-time Temporal Logics

We next provide the definitions of the syntax and semantics of the branching-time logics CTL-FO and CTL\*-FO, adapted from [1, 24]. These are extensions of the well-known languages CTL and CTL\* (see [12]). We also review the notion of satisfaction of a CTL<sup>(\*)</sup> formula by a Kripke structure.

**Definition A.3** *Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\varepsilon \rangle$  be a Web service. The set of CTL-FO\* formulas over  $\mathcal{W}$  is the set of state formulas defined inductively together with the set of path formulas as follows:*

1. each FO formula over the vocabulary of  $\mathcal{W}$  is a state formula;
2. if  $\varphi$  and  $\psi$  are state formulas then so are  $\varphi \wedge \psi$ ,  $\varphi \vee \psi$ , and  $\neg\varphi$ ;
3. if  $\varphi$  is a path formula, then  $E\varphi$  and  $A\varphi$  are state formulas;
4. each state formula is also a path formula;
5. if  $\varphi$  and  $\psi$  are path formulas then so are  $\varphi \wedge \psi$ ,  $\varphi \vee \psi$ , and  $\neg\varphi$ ;
6. if  $\varphi$  and  $\psi$  are path formulas then so are  $\mathbf{X}\varphi$  and  $\varphi\mathbf{U}\psi$ .

The set of CTL-FO formulas over  $\mathcal{W}$  is defined by replacing (4)-(6) above by the rule:

- if  $\varphi$  and  $\psi$  are state formulas then  $\mathbf{X}\varphi$ , and  $\varphi\mathbf{U}\psi$  are path formulas.

The set of CTL<sup>(\*)</sup>-FO sentences consists of the universal closures of CTL<sup>(\*)</sup>-FO formulas.

Note that, as in the case of LTL-FO, first-order quantifiers cannot be applied to formulas using temporal operators or path quantifiers. The formula is closed at the very the end by universally quantifying all remaining free variables, yielding an CTL<sup>(\*)</sup>-FO sentence. The semantics of the temporal operators is the natural extension of LTL-FO with path quantifiers. Informally,  $\mathbf{E}\varphi$  stands for "there exists a continuation of the current run that satisfies  $\varphi$ " and  $\mathbf{A}\varphi$  means "every continuation of the current run satisfies  $\varphi$ ". More formally, satisfaction of a CTL<sup>(\*)</sup>-FO sentence by a Web service  $\mathcal{W}$  is defined using the tree corresponding to the runs of  $\mathcal{W}$  on a given database  $D$ . The nodes of the tree consist of all prefixes of runs of  $\mathcal{W}$  on  $D$  (the empty prefix is denoted *root* and is the root of the tree). A prefix  $\pi$  is a child of a prefix  $\pi'$  iff  $\pi$  extends  $\pi'$  with a single configuration. We denote the resulting infinite tree by  $\mathcal{T}_{\mathcal{W},D}$ . Note that  $\mathcal{T}_{\mathcal{W},D}$  has only infinite branches (so no leafs) and each such infinite branch corresponds to a run of  $\mathcal{W}$  on database  $D$ . Satisfaction of an CTL<sup>(\*)</sup>-FO sentence by  $\mathcal{T}_{\mathcal{W},D}$  is the natural extension of the classical notion of satisfaction of CTL<sup>(\*)</sup> formulas by infinite trees labeled with propositional variables (e.g., see [12]), and is provided below. The main difference is that propositional variables are not explicitly provided; instead, the relevant FO formulas have to be evaluated on the current configuration (last of the prefix defining the node). We say that a Web service  $\mathcal{W}$  satisfies  $\varphi$ , denoted  $\mathcal{W} \models \varphi$ , iff  $\mathcal{T}_{\mathcal{W},D} \models \varphi$  for every database  $D$ .

We review the classical notion of satisfaction of a CTL<sup>(\*)</sup> formula by a Kripke structure (see[12]). The languages CTL<sup>(\*)</sup> are the restrictions of CTL<sup>(\*)</sup>-FO where all FO formulas are propositions.

**Definition A.4** Let  $AP = \{p_1, p_2, \dots, p_n\}$  be a finite set of atomic propositional symbols. A Kripke structure over  $AP$  is a 4-tuple  $K = (S, S^0, R, L)$  where:

- $S$  is a finite set of states.
- $S^0 \in S$  is an initial state.
- $R$  is a total binary relation on  $S$  ( $R \subseteq S \times S$ ), called the transition relation.
- $L: S \rightarrow 2^{AP}$  assigns to each state the set of atomic propositions which are true in that state.

A path  $\rho$  in Kripke structure  $K$  is an infinite sequence of states  $(s_0, s_1, \dots)$  such that  $(s_i, s_{i+1}) \in R$  for every  $i \geq 0$ . Let  $\rho_i$  denote the suffix path  $(s_i, s_{i+1}, s_{i+2}, \dots)$ . The notation  $K, s \models p$  indicates

that a CTL<sup>\*</sup> state formula  $p$  holds at state  $s$  of the Kripke Structure  $K$ . Similarly,  $K, \rho \models \psi$  indicates that a CTL<sup>\*</sup> path formula  $\psi$  holds at a path of  $\rho$  of the Kripke Structure  $K$ . We write  $s \models p$  or  $\rho \models \psi$  when it is obvious from the context which structure is concerned.

The notion of satisfaction of a CTL<sup>\*</sup> formula by a Kripke structure is defined as follows:

1.  $s \models p$  iff  $p \in L(s)$ .
2.  $s \models \neg p$  iff  $p \notin L(s)$ .
3.  $s \models \varphi_1 \wedge (\forall)\varphi_2$  iff  $s \models \varphi_1$  and(or)  $s \models \varphi_2$ .  
 $s \models \mathbf{E}\psi$  iff there exists an infinite path  $\rho' = (s, s_1, s_2, \dots)$  in  $K$ , starting from  $s$ , such that  $\rho' \models \psi$ .  
 $s \models \mathbf{A}\psi$  iff for every infinite path  $\rho' = (s, s_1, s_2, \dots)$  in  $K$  starting from  $s$ ,  $\rho' \models \psi$ .
4.  $\rho_j \models \psi$  iff  $s' \models \psi$  where  $s'$  is the first state in  $\rho_j$ .
5.  $\rho_j \models \psi_1 \wedge (\forall)\psi_2$  iff  $\rho_j \models \psi_1$  and(or)  $\rho_j \models \psi_2$ .  
 $\rho_j \models \neg\psi$  iff  $\rho_j \not\models \psi$ .
6.  $\rho_j \models \psi_1 \mathbf{U}\psi_2$  iff there exists  $i \geq j$  such that  $\rho_i \models \psi_2$  and  $\rho_k \models \psi_1$  for all  $j \leq k < i$ .  
 $\rho_j \models \mathbf{X}\psi$  iff  $\rho_{j+1} \models \psi$ .

A formula of CTL is also interpreted using the CTL<sup>\*</sup> semantics. The complexity of checking whether a CTL<sup>(\*)</sup> formula is satisfied by a Kripke structure (model checking) is in PTIME for CTL and PSPACE-complete for CTL<sup>\*</sup>. The satisfiability problem for CTL<sup>(\*)</sup> formulas is EXPTIME-complete for CTL and 2-EXPTIME complete for CTL<sup>\*</sup>. See [12] for a concise survey on temporal logics, and further references.

## A.3 Some proofs

### A.3.1 Proof of Theorem 3.5

We outline the main steps needed to prove Theorem 3.5. To begin, we note that (i) can be reduced to (ii).

**Lemma A.5** For each Web service  $\mathcal{W}$  with input-bounded rules there exists an error-free Web service  $\mathcal{W}'$  with input bounded rules, of size quadratic in  $\mathcal{W}$ , such that  $\mathcal{W}$  is error free iff  $\mathcal{W}' \models \varphi$ , for some fixed input-bounded LTL-FO sentence  $\varphi$ .

**Proof:** Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  be a Web service with input-bounded rules. Intuitively, we wish to construct a Web service  $\mathcal{W}'$  with a new Web page schema  $W'_\epsilon$  that is reached according to the rules of the service (and without generating an error), exactly when the error page  $W_\epsilon$  would be reached in

the original Web service. Then it is enough to verify that  $W'_\epsilon$  is never reached in any run of  $\mathcal{W}'$ . To this end, let  $\mathcal{W}' = \langle \mathbf{D}, \mathbf{S}', \mathbf{I}, \mathbf{A}, \mathbf{W}', W_0, W_\epsilon \rangle$  where  $\mathbf{S}' = \mathbf{S} \cup \{c \mid c \text{ is an input constant}\}$ , and  $\mathbf{W}'$  contains a new Web page schema  $W'_\epsilon$  defined identically to  $W_\epsilon$ , and for each Web page schema  $W = \langle \mathbf{I}_W, \mathbf{A}_W, \mathbf{T}_W, \mathcal{R}_W \rangle$  of  $\mathbf{W}$  different from  $W_0$  and  $W_\epsilon$ , a Web page schema  $W' = \langle \mathbf{I}_W, \mathbf{A}_W, \mathbf{T}'_W, \mathcal{R}'_W \rangle$ , where  $\mathbf{T}'_W = \mathbf{T}_W \cup \{W'_\epsilon\}$  and  $\mathcal{R}'_W$  consists of the following rules:

- all state, input, and action rules of  $\mathcal{R}_W$ ;
- a state rule  $c \leftarrow true$ , for each input constant  $c \in \mathbf{I}_W$ .
- each target rule  $V \leftarrow \varphi_{V,W}$ , where  $V \in \mathbf{T}_W$ , is replaced by  $V \leftarrow \varphi_{V,W} \wedge \neg\psi$ , where  $\psi$  is the disjunction of all formulas  $\varphi_{V',W}$  of target rules  $V' \leftarrow \varphi_{V',W}$  in  $\mathcal{R}_W$ , for  $V' \neq V$ ;
- $W'_\epsilon \leftarrow \psi \vee \xi$  where  $\psi$  is the disjunction of all formulas  $\varphi_{V,W} \wedge \varphi_{V',W}$  where  $V \neq V'$  and  $V \leftarrow \varphi_{V,W}$ ,  $V' \leftarrow \varphi_{V',W}$  are target rules in  $\mathcal{R}_W$ , and  $\xi$  is the disjunction of all formulas  $\varphi_{V,W} \wedge \neg c$  where  $V \leftarrow \varphi_{V,W}$  is a target rule in  $\mathcal{R}_W$  and  $c$  is an input constant occurring in some input rule in  $V$  but not occurring in  $\mathbf{I}_W$ , or occurring in some other formula of  $V$ , but not in  $\mathbf{I}_W$  or  $\mathbf{I}_V$ .

Finally,  $W'_0$  is a special case. It is defined as above if input rules of  $W_0$  contain no input constants, and the other formulas contain only input constants in  $\mathbf{I}_{W_0}$ ; otherwise, it is defined as  $\langle \emptyset, \emptyset, \{W'_\epsilon\}, \{W'_\epsilon \leftarrow true\} \rangle$ .

Note that  $W'_\epsilon$  is reached either if the original target rules are ambiguous (stated by  $\psi$ ) or if the next Web page requires some input constant not yet provided (formula  $\xi$ ). It is easily verified that  $\mathcal{W}'$  is error free and  $\mathcal{W}'$  is input bounded if  $\mathcal{W}$  is input bounded. Also,  $\mathcal{W}$  is error free iff the page  $W'_\epsilon$  is never reached in any run of  $\mathcal{W}'$ , i.e.  $\mathcal{W}'$  satisfies the input-bounded LTL-FO sentence  $\mathbf{G} \neg W'_\epsilon$ .  $\square$

It follows from Lemma A.5 that to prove Theorem 3.5 it is enough to establish PSPACE-hardness for (i), and inclusion in PSPACE for (ii). The following establishes PSPACE-hardness.

**Lemma A.6** *Checking whether an input-bounded Web service is error free is PSPACE-hard.*

**Proof:** The proof is by reduction from Quantified Boolean Formula (QBF), known to be PSPACE-complete [15]. Let  $\varphi$  be a quantified Boolean formula (we can assume  $\varphi$  uses just  $\vee, \neg, \exists$ ). Consider the Web service  $\mathcal{W}_\varphi = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  where:

- $\mathbf{D} = \{R : 1, 0, 1\}$ ,  $\mathbf{S} = \emptyset$ ,  $\mathbf{I} = \{I_0 : 1, I_1 : 1\}$ ,  $\mathbf{A} = \emptyset$ ,  $\mathbf{W} = \{W_0, W_1, W_2\}$ ;

- $W_0 = \langle \{I_0, I_1\}, \emptyset, \{W_1, W_2\}, \mathcal{R}_{W_0} \rangle$  where  $\mathcal{R}_{W_0}$  consists of the input rules  $Options_{I_i}(x) \leftarrow R(x)$ , for  $i \in \{0, 1\}$  and the target rules

$$W_i \leftarrow I_0(0) \wedge I_1(1) \wedge 0 \neq 1 \wedge \varphi', i \in \{1, 2\},$$

where  $\varphi'$  is defined from  $\varphi$  as follows:

- each propositional variable  $x_i$  is replaced by  $(x_i = 1)$ ;
- disjunction and negation remain unchanged;
- $\exists x\psi$  becomes  $\exists x((I_0(x) \vee I_1(x)) \wedge \psi)$ .

- $\{W_1, W_2\}$  are arbitrary.

Clearly,  $\mathcal{W}_\varphi$  is input-bounded and of size polynomial in  $\varphi$ , and it is error free iff there is no run for which  $I_0 = \{0\}$ ,  $I_1 = \{1\}$ , and  $\varphi'$  is true. Obviously, there exists a run for which  $I_0 = \{0\}$  and  $I_1 = \{1\}$ . But then  $\varphi'$  has the same value as  $\varphi$ . Therefore,  $\mathcal{W}$  is error-free iff  $\varphi$  is false.  $\square$

The proof of containment in PSPACE (for fixed schema arity) requires adapting the proof of Spielmann's analogous result for  $ASM^I$  transducers, outlined in Appendix A.1. This is done in three steps:

1. We first extend the standard ASM model to allow for input option rules and the use of *prev<sub>I</sub>* atoms in all rules. We denote the resulting model, further restricted to be input-bounded with bounded input flow, by  $ASM^{IR}$ . We then show decidability of verification of input-bounded LTL-FO properties of  $ASM^{IR}$  transducers, via a direct reduction to finite satisfiability of E+TC (the reduction to  $FO^W + TC$  in [24] is no longer possible here).
2. Next, we define a special type of Web service, called “simple”, that corresponds directly to  $ASM^{IR}$ .
3. Finally, we reduce the general verification problem to verification of simple Web services.

The following establishes (1).

**Theorem A.7** *It is decidable in PSPACE, given an  $ASM^{IR}$  transducer  $T$  and input-bounded LTL-FO sentence  $\varphi$  over the vocabulary of  $T$ , whether  $T \models \varphi$ .*

**Proof: (Sketch)** We adapt Spielmann’s proof that input-bounded LTL-FO properties of  $ASM^I$  transducers can be checked in PSPACE [23]. The structure of the proof is similar. Recall the stages of Spielmann’s proof, outlined in Appendix A.1. The Periodic Run Lemma continues to hold. So does the Local Run Lemma, but the proof needs to be adapted to the extended model. Specifically, the definition of local run and the proof have to be modified to take into account the restriction on the inputs defined by the input option rules of the  $ASM^{IR}$  model. Next, we need to pursue the reduction to satisfiability of E+TC formulas. Due to the restrictions on the inputs and the presence of **Prev**<sub>I</sub> atoms, the FO formula defining the connection between successive configurations of a local run has to be modified. In particular, the formula can no longer be expressed solely by an  $FO^W$  formula since the existentially quantified component reflecting the input restrictions cannot be replaced by witness-bounded quantification. Therefore, the intermediate reduction to  $FO^W + \text{posTC}$  is no longer possible. Instead, we can obtain a direct reduction to satisfiability of an E+TC formula.  $\square$

For the second stage of the proof, we show that  $ASM^{IR}$  transducers correspond to a special kind of Web service defined next.

**Definition A.8** A Web service  $\langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  is simple if  $\mathbf{W} = \{W_0\}$  and  $\mathbf{I}$  has no constants.

Note that a simple Web service is necessarily error free, so  $W_\epsilon$  plays no role.

The following is immediate from the definition of  $ASM^{IR}$  transducers.

**Lemma A.9** For each simple, input-bounded Web service  $\mathcal{W}$  and input-bounded LTL-FO sentence  $\varphi_{\mathcal{W}}$  over the vocabulary of  $\mathcal{W}$ , there exist an  $ASM^{IR}$  transducer  $T$  and an input-bounded LTL-FO sentence  $\varphi_T$  over the vocabulary of  $T$  such that the sizes of  $T$  and  $\varphi_T$  are linear in  $\mathcal{W}$  and  $\varphi_{\mathcal{W}}$ , and  $\mathcal{W} \models \varphi_{\mathcal{W}}$  iff  $T \models \varphi_T$ .

Finally, for the third stage of the proof, we reduce the verification of error-free Web services to that of simple Web services.

**Lemma A.10** Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  be an error-free, input-bounded Web service and  $\varphi$  an input-bounded LTL-FO sentence. There exists a simple input-bounded Web service  $\mathcal{W}'$ , of size linear in  $\mathcal{W}$ , and an input-bounded LTL-FO sentence  $\varphi'$ , of size linear in  $\varphi$  and  $\mathcal{W}$ , such that  $\mathcal{W} \models \varphi$  iff  $\mathcal{W}' \models \varphi'$ .

**Proof:** In brief, the reduction has to overcome two obstacles: (i) simulating multiple Web schemas of  $\mathcal{W}$  with a single Web page schema in  $\mathcal{W}'$ , and (ii) eliminating the constants from the input schema of  $\mathcal{W}$ . It is easy to deal with (i): we just simulate the behavior of different Web pages and transitions using new propositional state variables corresponding to the Web pages. Overcoming (ii) makes essential use of the assumption that  $\mathcal{W}$  is error free. Indeed, this guarantees that the value of each input constant is only provided once, and that no formula makes use of such constants before they are provided. This allows to assume that the input constants are provided prior to the run, as part of the database.

More precisely, let  $\mathcal{W}' = \langle \mathbf{D}', \mathbf{S}', \mathbf{I}', \mathbf{A}', \mathbf{W}', W_0, W_\epsilon \rangle$  where:

- $\mathbf{D}' = \mathbf{D} \cup \text{const}(\mathbf{I})$ , where  $\text{const}(\mathbf{I})$  denotes the set of input constant symbols in  $\mathbf{I}$
- $\mathbf{S}' = \mathbf{S} \cup \mathbf{W}$ , where each  $W \in \mathbf{W}$  is taken to be a propositional symbol;
- $\mathbf{I}' = \mathbf{I} - \text{const}(\mathbf{I})$ ;
- $\mathbf{A}' = \mathbf{A}$ ;
- $\mathbf{W}' = \{W_0\}$ .

The set of rules  $\mathcal{R}_{W_0}$  of  $W_0$  are defined as follows. The input rules are the following. For each relational input  $I$  of  $\mathbf{I}$  we add to  $\mathcal{R}_{W_0}$  the input rule  $Options_I(\bar{x}) \leftarrow \xi$ , where  $\xi$  is the disjunction of all formulas  $\varphi_{I,W}(\bar{x}) \wedge W$  for which  $Options_I(\bar{x}) \leftarrow \varphi_{I,W}(\bar{x})$  is an input rule of the page  $W$  in  $\mathcal{W}$ . We define the state rules next. For each state rule  $(\neg)S(\bar{x}) \leftarrow \varphi_{S,W}^{\epsilon}(\bar{x})$  of  $\mathbf{W}$ , we add a state rule  $(\neg)S(\bar{x}) \leftarrow \varphi_{S,W}^{\epsilon}(\bar{x}) \wedge W$  to  $\mathcal{R}_{W_0}$ . In addition, for each target rule  $V \leftarrow \varphi_{V,W}$  of  $\mathbf{W}$  we add to  $\mathcal{R}_{W_0}$  the state rules  $V \leftarrow \varphi_{V,W} \wedge W$  and  $\neg W \leftarrow \varphi_{V,W} \wedge W$ . The action rules of  $\mathcal{R}_{W_0}$  consist of all rules  $A(\bar{x}) \leftarrow \varphi(\bar{x}) \wedge W$  for which  $A(\bar{x}) \leftarrow \varphi(\bar{x})$  is an action rule of Web page schema  $W$  in  $\mathbf{W}$ . Finally, the only target rule of  $\mathcal{R}_{W_0}$  is  $W_0 \leftarrow \text{true}$ .

The LTL-FO sentence  $\varphi'$  is defined as follows. Let  $\varphi = \forall \bar{x} \psi(\bar{x})$ . Then  $\varphi'$  is the formula

$$\forall \bar{x} [\mathbf{G}(\xi) \rightarrow \psi],$$

where  $\xi$  is the conjunction of the input-bounded FO formulas  $\forall \bar{x} \forall \bar{y} ((I(\bar{x}) \wedge I(\bar{y})) \rightarrow \bar{x} = \bar{y})$  for each relational input  $I \in \mathbf{I}$ .  $\square$

### A.3.2 Other proofs

**Proof of Theorem 3.7** The proof is by reduction of the question of whether a Turing Machine (TM) M

halts on input  $\epsilon$ . Let  $M$  be a deterministic TM with a left-bounded, right-infinite tape. We construct a Web service with a single page (excepting the error page). The idea is to represent configurations of  $M$  using a 4-ary state relation  $T$ . The first two coordinates of  $T$  represent a successor relation on a subset of the active domain of the database. A tuple  $T(x, y, u, v)$  says that the content of the  $x$ -th cell is  $u$ , the next cell is  $y$ , and  $v$  is a state  $p$  iff  $M$  is in state  $p$  and the head is on cell  $x$ . Otherwise,  $v$  is some special symbol  $\#$ . The moves of  $M$  are simulated by modifying  $T$  accordingly.  $M$  halts on input  $\epsilon$  iff there exists a run of  $\mathcal{W}$  on some database such that some halting state  $h$  is reached. Thus,  $M$  does not accept  $\epsilon$  iff for every run,  $T(x, y, u, h)$  does not hold for any  $x, y, u$ , that is,  $\mathcal{W} \models \forall x \forall y \forall u \mathbf{G}(\neg T(x, y, u, h))$ .

We now outline the construction of  $\mathcal{W}$  in more detail. The database schema of  $\mathcal{W}$  consists of a unary relation  $D$  and a constant  $min$ . The state relations are the following:

- $T$ , a 4-ary relation;
- $Cell$ ,  $Max$ , and  $Head$ , unary relations;
- propositional states used to control the computation: *initialized*, *simul*

The input relations are  $I$  and  $H$ , both unary.

The first phase of the simulation constructs the initial configuration of  $M$  on input  $\epsilon$ , and the tape that the current run will make available for the computation. This phase makes use of the unary input relation  $I$ . Intuitively, the role of  $I$  is to pick a new value from the active domain, that has not yet been used to identify a cell, and use it to identify a new cell of the tape. The state relation  $Cell$  keeps track of the values previously chosen, to prevent them from being chosen again. The state relation  $Max$  keeps track of the most recently inserted value.

The rules implementing the initialization are the following (the symbol  $b$  denotes the blank symbol of  $M$  and  $q_0$  is the start state):

$$\begin{array}{ll}
Options_I(y) & \leftarrow D(y), y \neq min, \\
& \quad \neg Cell(y), \neg simul \\
T(min, y, b, q_0) & \leftarrow I(y), \neg initialized \\
Cell(min) & \leftarrow \neg initialized \\
Head(min) & \leftarrow \neg initialized \\
initialized & \leftarrow \neg initialized \\
T(x, y, b, \#) & \leftarrow I(y), Max(x) \\
Cell(y) & \leftarrow I(y) \\
\neg Max(x) & \leftarrow Max(x) \\
Max(y) & \leftarrow I(y) \\
simul & \leftarrow \forall y \neg I(y)
\end{array}$$

The state *simul* signals the transition to the simulation phase. Notice that this happens either if the input options for  $I$  become empty (because we have used the entire active domain) or because the input is empty at any point. In the simulation phase,  $T$  is updated to reflect the consecutive moves of  $M$ . The simulation is aborted if  $T$  runs out of tape. We illustrate the simulation with an example move. Suppose  $M$  is in state  $p$ , the head is at cell  $x$ , the content of the cell is 0, and the move of  $M$  in this configuration consists of overwriting 0 with 1, changing states from  $p$  to  $q$ , and moving right. The rules simulating this move are the following:

$$\begin{array}{ll}
Options_H(x, y, 0, p) & \leftarrow simul, Head(x), T(x, y, 0, p) \\
\neg T(x, y, 0, p) & \leftarrow simul, H(x, y, 0, p) \\
T(x, y, 1, \#) & \leftarrow simul, H(x, y, 0, p) \\
\neg T(y, z, u, \#) & \leftarrow simul, H(x, y, 0, p), T(y, z, u, \#) \\
T(y, z, u, q) & \leftarrow simul, H(x, y, 0, p), T(y, z, u, \#) \\
\neg Head(x) & \leftarrow simul, H(x, y, 0, p) \\
Head(y) & \leftarrow simul, H(x, y, 0, p)
\end{array}$$

Such rules are included for every move of  $M$ . It is easy to see that this correctly simulates the moves of  $M$ . Note that if the input  $H$  is empty,  $T$  does not change. Finally, if the head reaches the last value provided in  $T$ , the transducer goes into an infinite loop in which, again,  $T$  stays unchanged. Thus,  $T(x, y, u, h)$  holds in some run iff the computation of  $M$  on  $\epsilon$  is halting. Equivalently,  $M$  does not halt on  $\epsilon$  iff the transducer satisfies the formula  $\varphi = \forall x \forall y \forall u \mathbf{G}(\neg T(x, y, u, h))$ .  $\square$

**Proof of Theorem 3.8** The proof is by reduction of the implication problem for functional and inclusion dependencies, known to be undecidable [9]. Let  $\Delta$  be a set of FDs and IDs over a relation  $S$ , and  $f$  an FD over the same relation. We construct a simple, input-bounded Web service  $\mathcal{W}$  with state projections and an input-bounded LTL-FO sentence  $\varphi$  such that  $\Delta \models f$  iff  $\mathcal{W} \models \varphi$ .

Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  where  $\mathbf{D} = R$ ,  $\mathbf{A} = \emptyset$ ,  $\mathbf{I} = \{I, done\}$  where  $I$  has the same arity as  $S$  and  $done$  is propositional, and  $\mathbf{S}$  consists of the following relations:

- the relation  $S$ ;
- two propositions  $stop_1, stop_2$ ;
- for each ID  $\sigma$  of the form  $[X] \subseteq [Y]$  in  $\Delta$ , a relation  $S_X$  of arity  $|X|$ , a relation  $S_Y$  of arity  $|Y|$ , a relation  $S_X^\sigma$  of arity  $|X|$ , and a proposition  $viol_\sigma$ ;

- for each FD  $\sigma$  of the form  $X \rightarrow A$  in  $\Delta \cup \{f\}$  a relation  $S_{XA}$  of arity  $|XA|$ , a relation  $S_{XA_1A_2}^\sigma$  of arity  $|XA| + 1$ , and a proposition  $viol_\sigma$ .

Next, let  $\mathbf{W} = W_0$  and let  $W_0$  be defined as follows. The input option rule for  $I$  defines the cross-product of the active domain given by the database relation  $R$ . The state rules consist of the following:

$$\begin{aligned} S(\bar{x}) &\leftarrow I(\bar{x}) \wedge \neg stop_1 \\ stop_1 &\leftarrow done \\ stop_2 &\leftarrow stop_1 \end{aligned}$$

for each ID  $\sigma$  of the form  $[X] \subseteq [Y]$  in  $\Delta$ , the rules:

$$\begin{aligned} S_X &\leftarrow \pi_X(S) \\ S_Y &\leftarrow \pi_Y(S) \\ S_X^\sigma(\bar{x}) &\leftarrow S_X(\bar{x}) \wedge \neg S_Y(\bar{x}) \wedge stop_2 \\ viol_\sigma &\leftarrow \exists \bar{x} S_X^\sigma(\bar{x}) \end{aligned}$$

for each FD  $\sigma$  of the form  $X \rightarrow A$  in  $\Delta \cup \{f\}$ , the rules:

$$\begin{aligned} S_{XA} &\leftarrow \pi_{XA}(S) \\ S_{XA_1A_2}^\sigma(\bar{x}, a_1, a_2) &\leftarrow S_{XA}(\bar{x}a_1) \wedge S_{XA}(\bar{x}a_2) \\ &\quad \wedge a_1 \neq a_2 \wedge stop_2 \\ viol_\sigma &\leftarrow \exists \bar{x} \exists a_1 \exists a_2 S_{XA_1A_2}^\sigma(\bar{x}, a_1, a_2) \end{aligned}$$

Intuitively, the state relation  $S$  is populated by repeated inputs, until  $done$  is set to true, which is remembered in the state propositions  $stop_1$  and  $stop_2$  ( $stop_2$  is needed for timing reasons, to ensure that violations are not tested too early). The rules check for violations of the dependencies in  $\Delta$ , so that  $viol_\sigma$  is set to true iff  $S$  violates  $\sigma$ .

Note that all rules are input bounded, except those consisting of projections of state relations. Next, let  $\xi$  be the input-bounded LTL-FO sentence

$$G(\neg done) \vee [F(done) \wedge (F(\bigvee_{\sigma \in \Delta} viol_\sigma) \vee G(\psi_f))]$$

where  $\psi_f$  is the formula  $\neg S_{XA_1A_2}^f(\bar{x}, a_1, a_2)$  whose universal closure states that the FD  $f = X \rightarrow A$  is satisfied. Finally, let  $\varphi$  be the universal closure of  $\xi$ . Intuitively,  $\varphi$  states that either  $done$  is never set to true, or it is set to true and at least one of the constraints of  $\Delta$  is violated, or  $f$  is satisfied. Thus,  $\mathcal{W} \models \varphi$  iff  $\Delta \models f$ .  $\square$

**Proof of Theorem 4.2** Using path quantifiers, one can easily simulate first-order quantification by considering runs that provide values for the quantified variables as inputs. This allows to use a reduction of finite validity of FO sentences to the above verification problem. We illustrate the reduction

for FO sentences of the form  $\exists x \forall y \alpha(x, y)$  where  $\alpha$  is a quantifier free formula over relational vocabulary  $\{R\}$ . Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  be a simple Web service where  $\mathbf{D} = \{R\}$ ,  $\mathbf{I} = \{X, Y\}$  ( $X, Y$  are unary relations),  $\mathbf{A} = \emptyset$ ,  $\mathbf{W} = \{W_0\}$ ,  $\mathbf{S} = \{S_X, S_Y, done_x, true_\alpha\}$  ( $S_X, S_Y$  are unary and the other states are propositional). The input option rules are:

$$\begin{aligned} Options_X(x) &\leftarrow (\psi_{dom}(x) \wedge \neg done_x) \\ &\quad \vee (done_x \wedge S_X(x)) \\ Options_Y(y) &\leftarrow done_x \wedge \psi_{dom}(y) \end{aligned}$$

where  $\psi_{dom}(x)$  defines the active domain provided by  $R$ . The state rules are the following:

$$\begin{aligned} S_X(x) &\leftarrow X(x) \\ done_x &\leftarrow \neg done_x \\ true_\alpha &\leftarrow \exists x \exists y (X(x) \wedge Y(y) \wedge \alpha(x, y)) \end{aligned}$$

Note that a path in  $\mathcal{T}_{\mathcal{W}}$  starts at  $root$ , then proceeds to the start configuration of a run on some database  $D$ . The first input provided is a value of  $x$ , which is remembered in the state relation  $S_X$ . In the next configuration,  $done_x$  is true, the same value of  $x$  as previously chosen is provided again via input  $X$ , and an arbitrary value is provided for  $y$  by the input relation  $Y$ . In the following configuration  $true_\alpha$  is true if  $\alpha(x, y)$  is satisfied for the chosen values of  $x, y$ . Let  $\varphi$  be the CTL-FO sentence  $\text{EXAX}(true_\alpha)$ . Clearly,  $\mathcal{W} \models \varphi$  iff  $\exists x \forall y \alpha(x, y)$  is valid. Note that  $\mathcal{W}$  and  $\varphi$  are input bounded (in fact  $\varphi$  is propositional, so in CTL). This proof is easily extended along the same lines to the general case.

**Proof of Theorem 4.4** The proof has two stages. First, we show that there is a bound on the size of databases that need to be considered when checking for violations of  $\varphi$  (or equivalently, satisfaction of  $\neg\varphi$ ). Second, we prove that for a given database  $D$  there exists a Kripke structure  $\mathcal{K}_{\mathcal{W}, D}$  over alphabet  $\Sigma_{\mathcal{W}}$ , of size exponential in  $\Sigma_{\mathcal{W}}$ , such that  $\mathcal{T}_{\mathcal{W}, D} \models \neg\varphi$  iff  $\mathcal{K}_{\mathcal{W}, D} \models \neg\varphi$ . This allows us to use known model-checking techniques for CTL<sup>(\*)</sup> on Kripke structures to verify whether  $\mathcal{T}_{\mathcal{W}, D} \models \neg\varphi$ .

We start with the following:

**Lemma A.11** *Let  $\mathcal{W}$  be a propositional, input bounded Web service, and  $\varphi$  a CTL\* formula over  $\Sigma_{\mathcal{W}}$ . Then  $\mathcal{W} \not\models \varphi$  iff there exists a database instance  $D$  of size exponential in  $\mathcal{W}$ , such that  $\mathcal{T}_{\mathcal{W}, D} \models \neg\varphi$ .*

**Proof:** Let  $\mathcal{W} = \langle \mathbf{D}, \mathbf{S}, \mathbf{I}, \mathbf{A}, \mathbf{W}, W_0, W_\epsilon \rangle$  be a propositional, input-bounded Web service and  $\varphi$  a CTL\* formula over  $\Sigma_{\mathcal{W}}$ . For each configuration  $\rho =$

$\langle V, S, I, A \rangle$  of  $\mathcal{W}$  we denote by  $\lambda(\rho)$  the set of propositions of  $\Sigma_{\mathcal{W}}$  true in  $\rho$  (as defined above). We also denote by  $\lambda$  the extension of this mapping to trees of configurations  $\mathcal{T}_{\mathcal{W}, D}$ , where  $\lambda(\text{root}) = \emptyset$ . Obviously, if  $\lambda(\mathcal{T}_{\mathcal{W}, D_1}) = \lambda(\mathcal{T}_{\mathcal{W}, D_2})$  then  $\mathcal{T}_{\mathcal{W}, D_1}$  and  $\mathcal{T}_{\mathcal{W}, D_2}$  satisfy the same CTL\* formulas over  $\Sigma_{\mathcal{W}}$ . Now suppose that  $\mathcal{W} \not\models \varphi$ , so there is some  $D$  such that  $\mathcal{T}_{\mathcal{W}, D} \models \neg\varphi$ . We show that there exists a database  $D_0$  of size exponential in  $\mathcal{W}$ , such that  $\lambda(\mathcal{T}_{\mathcal{W}, D}) = \lambda(\mathcal{T}_{\mathcal{W}, D_0})$ , so  $\mathcal{T}_{\mathcal{W}, D_0} \models \neg\varphi$ . This is done by showing that  $\lambda(\mathcal{T}_{\mathcal{W}, D}) = \lambda(\mathcal{T}_{\mathcal{W}, D_0})$  iff  $D_0$  satisfies a particular FO sentence  $\xi$  in the prefix class  $\exists^*\forall^*\text{FO}$  with a number of variables exponential in  $\mathcal{W}$ . Since  $\xi$  is satisfied by  $D$ , it is satisfiable. But this implies that  $\xi$  has a model  $D_0$  whose domain has a number of elements equal to the number of existential variables of  $\xi$ , so exponential in  $\mathcal{W}$  (see [5]). Thus, the size of  $D_0$  is also exponential in  $\mathcal{W}$  (for bounded database schema arity).

We next describe  $\xi$ . Note that, because states are propositional, the sets of propositions true in successors of a configuration  $\rho$  of a run of  $\mathcal{W}$  on  $D$  depend only on  $D$  and  $\lambda(\rho)$ . Thus,  $\lambda(\rho)$  uniquely determines the set  $\{\lambda(\bar{\rho}) \mid \langle \rho, \bar{\rho} \rangle \in \mathcal{T}_{\mathcal{W}, D}\}$ . Consider a pair  $\langle \Sigma, \bar{\Sigma} \rangle = \langle \lambda(\rho), \lambda(\bar{\rho}) \rangle \in \lambda(\mathcal{T}_{\mathcal{W}, D})$ . Note that, since  $\rho$  and  $\bar{\rho}$  are configurations of  $\mathcal{W}$ , each of  $\Sigma$  and  $\bar{\Sigma}$  contain a single Web page symbol and input, state, and action symbols compatible with the schema of the respective Web page. Suppose  $V$  is the Web page symbol in  $\Sigma$ , let  $I_1, \dots, I_k$  be the input predicates in  $\mathbf{I}_V$ , and let  $\varphi_1, \dots, \varphi_k$  be the  $\exists^*\text{FO}$  formulas defining the input options for  $I_1, \dots, I_k$  in  $V$ . We construct a quantifier-free FO sentence  $\varphi_{\langle \Sigma, \bar{\Sigma} \rangle}(\bar{x}_1, \dots, \bar{x}_k)$  on  $\mathbf{D}$  such that  $\lambda(\bar{\rho}) = \bar{\Sigma}$  whenever  $\bar{\rho}$  is the next configuration from  $\rho$  resulting from the choice of inputs  $\bar{x}_1, \dots, \bar{x}_k$  from the options available for  $I_1, \dots, I_k$ . For simplicity, we show the construction for the case when all user inputs are non-empty. The construction can be easily adapted to account for empty inputs.

Thus,  $\langle \Sigma, \bar{\Sigma} \rangle \in \lambda(\mathcal{T}_{\mathcal{W}, D_0})$  iff  $D_0 \models \exists \bar{x}_1 \dots \exists \bar{x}_k [\varphi_1(\bar{x}_1) \wedge \dots \wedge \varphi_k(\bar{x}_k) \wedge \varphi_{\langle \Sigma, \bar{\Sigma} \rangle}(\bar{x}_1, \dots, \bar{x}_k)]$ . To ensure that only valid pairs  $\langle \Sigma, \bar{\Sigma} \rangle$  occur in  $\lambda(\mathcal{T}_{\mathcal{W}, D_0})$ , it must also be the case that  $D_0 \models \forall \bar{x}_1 \dots \forall \bar{x}_k [(\varphi_1(\bar{x}_1) \wedge \dots \wedge \varphi_k(\bar{x}_k)) \rightarrow \bigvee_{\bar{\Sigma}} \varphi_{\langle \Sigma, \bar{\Sigma} \rangle}(\bar{x}_1, \dots, \bar{x}_k)]$ .

Then  $\xi$  is the conjunction of all such formulas for all pairs in  $\lambda(\mathcal{T}_{\mathcal{W}, D})$ , yielding a formula in the prefix class  $\exists^*\forall^*\text{FO}$ . Since there can be exponentially many such pairs,  $\xi$  is exponential in  $\mathcal{W}$ .

In order to define the sentence  $\varphi_{\langle \Sigma, \bar{\Sigma} \rangle}$  we need the following notation. For each FO sentence  $\psi$  let  $\psi_{\Sigma}$  be the sentence obtained by replacing in  $\psi$  every proposition  $p \in \Sigma$  by *true* and  $p \notin \Sigma$  by *false*. Further,

for each input-bounded formula  $\psi$  let the quantifier-free version of  $\psi$ , denoted  $\psi^{qf}$ , be defined as follows. Intuitively,  $\psi^{qf}$  eliminates the quantifiers by taking advantage of the fact that each input  $I$  consists, after the user's choice, of at most a single tuple  $\bar{x}_I$  ( $\bar{x}_I$  is a sequence of  $m$  distinct variables, where  $m$  is the arity of  $I$ ). The formula  $\psi^{qf}$  reformulates  $\psi$  using these tuples. Specifically, let  $\psi'$  be obtained by replacing each input-bounded quantification  $\exists \bar{x}(\alpha \wedge \beta)$  and  $\forall \bar{x}(\alpha \rightarrow \beta)$  by  $\alpha \wedge \beta$ .

Next, let  $\psi^{qf}$  be obtained by first bringing  $\psi'$  to DNF (disjunctions of conjunctions), then applying to each disjunct  $\delta$  the following procedure yielding  $\delta'$ . Let  $eq$  ( $neq$ ) be the (in)equalities occurring in  $\delta$ . For each input relation  $I$  occurring in  $\delta$  and each  $i$ ,  $1 \leq i \leq m$ , let  $\theta(I, i)$  be the set of terms occurring in the  $i$ -th position of  $I$  in a positive occurrence  $I(\bar{x})$  in  $\delta$ . Let  $\equiv$  be the reflexive, transitive closure of the following relation on the terms of  $\delta$ :  $\{(x, y) \mid x = y \in eq\} \cup \{(x, y) \mid x, y \in \theta(I, i) \text{ for some } I \text{ and } i\}$ . If for some  $x, y$  it is the case that  $x \equiv y$  and  $x \neq y$  is in  $neq$ , then  $\delta' = \text{false}$ . Otherwise, define the following equivalence relation on the pairs  $(I, i)$  of input atoms  $I$  and positions  $i$  of  $I$ :  $(I, i) \equiv (J, j)$  iff there exist terms  $x, y$  so that  $x \equiv y$ ,  $x \in \theta(I, i)$ , and  $y \in \theta(J, j)$ . For each variable  $y$  in  $\delta$ , let  $\nu(y)$  be one arbitrarily chosen  $(x_I)_i$  for which  $y \in \theta(I, i)$ . Let  $\delta'$  be obtained as follows:

1. add to  $\delta$  the conjunction of all equalities  $(x_I)_i = (y_J)_j$  where  $(I, i) \equiv (J, j)$ ,  $c = c'$  where  $c, c'$  are constants and  $c \equiv c'$ , and  $(x_I)_i = c$  for some arbitrarily chosen  $c \in \theta(I, i)$ , if such exists.
2. for each negative occurrence  $\neg I(z_1, \dots, z_m)$  of an input atom, add the conjunct consisting of the disjunction  $\bigvee_{i=1}^m (\nu(z_i) \neq (x_I)_i)$ ;
3. delete all input atoms;
4. replace each variable  $y$  by  $\nu(y)$  in the remaining atoms.

Finally,  $\psi^{qf}$  is the disjunction of all resulting  $\delta'$ .

We can now define  $\varphi_{\langle \Sigma, \bar{\Sigma} \rangle}$ . This is constructed using the rules of  $\mathcal{W}$  for the unique Web page  $V \in \Sigma$ . Consider first a proposition  $p$  in  $\Sigma_V = \mathbf{S} \cup \mathbf{I} \cup \mathbf{A}_V \cup \mathbf{T}_V$  (recall that  $\mathbf{T}_V$  denotes the target Web pages of  $V$ ). We associate to  $p$  and  $\neg p$  formulas  $\beta_p$  and  $\beta_{\neg p}$  defined using the rules for  $(\neg)p$  that apply to the Web pages  $V$  and  $\bar{V} \in \mathbf{T}_V$  (recall that the set of rules for  $V$  is denoted  $\mathcal{R}_V$ ):

- if  $p \in \mathbf{S}$  and  $p \leftarrow \gamma$  and  $\neg p \leftarrow \delta$  are in  $\mathcal{R}_V$  then  $\beta_p$  is  $\tau_{\Sigma}^{qf}$ , where  $\tau = (\gamma \wedge \neg\delta) \vee (p \wedge \neg\delta)$ , and  $\beta_{\neg p}$  is  $\pi_{\Sigma}^{qf}$  where  $\pi = (\neg p \wedge \neg\gamma) \vee (\neg p \wedge \gamma \wedge \delta) \vee (\delta \wedge \neg\gamma)$

- if  $p \in \mathbf{I}$  and  $p \in \mathbf{I}_{\bar{V}}$  then  $\beta_p = \text{true}$  and  $\beta_{\neg p} = \text{false}$ . Otherwise,  $\beta_p = \text{false}$  and  $\beta_{\neg p} = \text{true}$ .
- if  $p \in \mathbf{A}_V$  and  $p \leftarrow \gamma$  is the rule for  $p$  in  $\mathcal{R}_V$ , then  $\beta_p = \gamma_{\Sigma}^{gf}$  where and  $\beta_{\neg p} = \neg\gamma_{\Sigma}^{gf}$ ;
- if  $p \in \mathbf{T}_V$  and  $p \leftarrow \gamma$  is the rule for  $p$  in  $\mathcal{R}_V$ , then  $\beta_p = \gamma_{\Sigma}^{gf}$  and  $\beta_{\neg p} = \neg\gamma_{\Sigma}^{gf}$ ;

Finally,  $\varphi_{\langle \Sigma, \bar{\Sigma} \rangle}$  is the quantifier-free formula  $\bigwedge_{p \in \Sigma} \beta_p \wedge \bigwedge_{p \in (\Sigma_V - \bar{\Sigma})} \beta_{\neg p}$ .  $\square$

The next stage towards the proof of Theorem 4.4 is to reduce the verification problem for a fixed database to a model checking problem of a CTL<sup>(\*)</sup> formula on a Kripke structure. We therefore show the following.

**Lemma A.12** *For each Web service  $\mathcal{W}$  over database schema  $\mathbf{D}$ , each database instance  $D$  over  $\mathbf{D}$ , and each CTL<sup>(\*)</sup> formula  $\varphi$  over  $\Sigma_{\mathcal{W}}$ , one can construct, in time polynomial in  $D$  and exponential in  $\mathcal{W}$ , a Kripke structure  $\mathcal{K}_{\mathcal{W},D}$  over  $\Sigma_{\mathcal{W}}$ , of size exponential in  $\Sigma_{\mathcal{W}}$ , such that  $\mathcal{T}_{\mathcal{W},D} \models \varphi$  iff  $\mathcal{K}_{\mathcal{W},D} \models \varphi$ .*

**Proof:** The Kripke structure  $\mathcal{K}_{\mathcal{W},D}$  has one node labeled for each set of propositions  $\Sigma \subseteq \Sigma_{\mathcal{W}}$  labeling a node in  $\lambda(\mathcal{T}_{\mathcal{W},D})$ . There is an edge  $\langle \Sigma, \bar{\Sigma} \rangle$  iff there is a node labeled  $\Sigma$  with a child labeled  $\bar{\Sigma}$  in  $\lambda(\mathcal{T}_{\mathcal{W},D})$ . Clearly,  $\mathcal{K}_{\mathcal{W},D}$  can be obtained by expanding  $\lambda(\mathcal{T}_{\mathcal{W},D})$  until no new labels are found. Each edge involves evaluating the formulas of  $\mathcal{W}$  on  $\Sigma$  and  $D$ , which is polynomial in  $\Sigma$  and  $D$  and exponential in  $\mathcal{W}$ . The maximum number of edges is exponential in  $\Sigma_{\mathcal{W}}$ .  $\square$

Lemmas A.11 and A.12 provide the proof of Theorem 4.4: to check that  $\mathcal{W} \not\models \varphi$ , first guess a database  $D$  of size exponential in  $\mathcal{W}$ , then construct from  $D$  and  $\mathcal{W}$ , in time exponential in  $\mathcal{W}$ , the Kripke structure  $\mathcal{K}_{\mathcal{W},D}$ . Finally, checking that  $\mathcal{K}_{\mathcal{W},D} \models \neg\varphi$  is in polynomial time wrt  $\mathcal{K}_{\mathcal{W},D}$  and  $\neg\varphi$  if  $\varphi$  is in CTL, and in polynomial space if  $\varphi$  is in CTL\*. Overall, checking  $\mathcal{W} \models \varphi$  is in CO-NEXPTIME if  $\varphi$  is in CTL, and in EXPTIME if  $\varphi$  is in CTL\*.

**Proof of Corollary 4.5** The decision procedure is similar to that for Theorem 4.4. Since  $\mathbf{S}$  is fixed and  $\varphi$  refers only to  $\mathbf{W}$ , it is enough to retain, in labels of  $\lambda(\mathcal{T}_{\mathcal{W},D})$  only the states and Web page names. Since  $\mathcal{W}$  is error free, there is exactly one Web page name per label. It follows that the number of pairs  $\langle \Sigma, \bar{\Sigma} \rangle$  occurring in  $\lambda(\mathcal{T}_{\mathcal{W},D})$  is quadratic in  $\mathbf{W}$ , so the formula  $\xi$  has polynomially many variables, and the size of the database  $D_0$  is polynomial in  $\mathcal{W}$ . The Kripke

structure  $\mathcal{K}_{\mathcal{W},D_0}$  can now be constructed in PSPACE wrt  $\mathcal{W}$ , and checking  $\varphi$  can be done in PSPACE wrt  $\mathcal{K}_{\mathcal{W},D_0}$  and  $\varphi$ . Altogether, checking that  $\mathcal{W} \models \varphi$  is done in PSPACE wrt  $\mathcal{W}$  and  $\varphi$ .  $\square$

**Proof of Theorem 4.6** In the case of a fully propositional Web service  $\mathcal{W}$ , the Kripke structure  $\mathcal{K}_{\mathcal{W},D}$  is independent of  $D$  (let us denote it by  $\mathcal{K}_{\mathcal{W}}$ ). However, unlike the case considered in Corollary 4.5,  $\mathcal{K}_{\mathcal{W}}$  is exponential wrt  $\mathcal{W}$  so cannot be constructed in PSPACE. We therefore need a more subtle approach, that circumvents the explicit construction of  $\mathcal{K}_{\mathcal{W}}$ . To do so, we adopt techniques developed in the context of model checking for concurrent programs (modeled by propositional transition systems). Specifically, the model checking algorithm developed by Kupferman, Vardi and Wolper in [19] can be adapted to fully propositional Web services. The algorithm uses a special kind of tree automaton, called *hesitant alternating tree automaton* (HAA) (see [19] for the definition). As shown in [19], for each CTL\* formula  $\varphi$  one can construct an HAA  $A_{\varphi}$  accepting precisely the trees (with degrees in a specified finite set) that satisfy  $\varphi$ . In particular, for a given Kripke structure  $\mathcal{K}$ , one can construct a product HAA  $\mathcal{K} \times A_{\varphi}$  that is nonempty iff  $\mathcal{K} \models \varphi$ . The nonemptiness test can be rendered efficient using the crucial observation that nonemptiness of  $\mathcal{K} \times A_{\varphi}$  can be reduced to the nonemptiness of a corresponding word HAA over a 1-letter alphabet, which is shown to be decidable in linear time, unlike the general nonemptiness problem for alternating tree automata. Finally, it is shown that  $\mathcal{K} \times A_{\varphi}$  need not be constructed explicitly. Instead, its transitions can be generated on-the-fly from  $\mathcal{K}$  and  $\varphi$ , as needed in the nonemptiness test for the 1-letter word HAA corresponding to  $\mathcal{K} \times A_{\varphi}$ . This yields a model checking algorithm of space complexity polynomial in  $\varphi$  and polylogarithmic in  $\mathcal{K}$ .

In our case,  $\mathcal{K}$  is  $\mathcal{K}_{\mathcal{W}}$ , and the input consists of  $\varphi$  and  $\mathcal{W}$  instead of  $\varphi$  and  $\mathcal{K}_{\mathcal{W}}$ . The previous approach can be adapted by pushing further the on-the-fly generation of  $\mathcal{K}_{\mathcal{W}} \times A_{\varphi}$  by also generating on-the-fly the relevant edges of  $\mathcal{K}_{\mathcal{W}}$  from  $\mathcal{W}$  when needed. This yields a polynomial space algorithm for checking whether  $\mathcal{W} \models \varphi$ , similar to the algorithm with the same complexity obtained in [19] for model checking of concurrent programs.  $\square$

**Proof of Theorem 4.9** We reduce the problem of checking whether  $\mathcal{W} \models \varphi$  to the satisfiability problem for CTL<sup>(\*)</sup> formulas. As mentioned in Appendix A.2, this is known to be EXPTIME-complete for CTL and 2-EXPTIME complete for CTL\*. We consider Kripke

structures over the alphabet  $\Sigma_{\mathcal{W}} \cup \mathbf{D}$ . Intuitively, each node of the Kripke structure represents a configuration, and its label represents the relevant information about the configuration: the set of propositions in  $\Sigma_{\mathcal{W}}$  that hold, and the type of the current input with respect to the database, i.e. the set of relations  $Q$  in  $\mathbf{D} - \{R_I\}$  for which  $y^k \in Q$ , where  $k$  is the arity of  $Q$  and  $y$  the current input. Note that the types of different inputs are independent of each other because inputs are unary, so every Kripke structure can be viewed as representing an input choice relation  $R_I$  together with type assignments for the elements of  $R_I$ . In addition, in order for a Kripke structure to represent an actual run of  $\mathcal{W}$ , the assignments of literals of  $\Sigma_{\mathcal{W}}$  to nodes has to be consistent with the rules of  $\mathcal{W}$ . However, this can be easily expressed by a CTL formula  $\rho$  computable in polynomial time from  $\mathcal{W}$ . It follows that  $\mathcal{W} \models \varphi$  iff  $\rho \wedge \neg\varphi$  is unsatisfiable. The latter is a CTL formula if  $\varphi$  is in CTL, and a CTL\* formula if  $\varphi$  is in CTL\*.  $\square$