

Automatic Composition of *E*-services That Export Their Behavior^{*}

Daniela Berardi, Diego Calvanese, Giuseppe De Giacomo,
Maurizio Lenzerini, and Massimo Mecella

Dipartimento di Informatica e Sistemistica
Università di Roma “La Sapienza”
Via Salaria 113, 00198 Roma, Italy
lastname@dis.uniroma1.it

Abstract. The main focus of this paper is on automatic *e*-Service composition. We start by developing a framework in which the exported behavior of an *e*-Service is described in terms of its possible executions (execution trees). Then we specialize the framework to the case in which such exported behavior (i.e., the execution tree of the *e*-Service) is represented by a finite state machine. In this specific setting, we analyze the complexity of synthesizing a composition, and develop sound and complete algorithms to check the existence of a composition and to return one such a composition if one exists. To the best of our knowledge, our work is the first attempt to provide an algorithm for the automatic synthesis of *e*-Service composition, that is both proved to be correct, and has an associated computational complexity characterization.

1 Introduction

Service Oriented Computing (SOC [16]) aims at building agile networks of collaborating business applications, distributed within and across organizational boundaries.¹ *e*-Services, which are the basic building blocks of SOC, represent a new model in the utilization of the network, in which self-contained, modular applications can be described, published, located and dynamically invoked, in a programming language independent way.

The commonly accepted and *minimal* framework for *e*-Services, referred to as Service Oriented Architecture (SOA [17]), consists of the following basic roles: (i) the *service provider*, which is the subject (e.g., an organization) providing services; (ii) the *service directory*, which is the subject providing a repository/registry of service descriptions, where providers publish their services and requestors find services; and, (iii) the *service requestor*, also referred to as client, which is the subject looking for and invoking the service in order to fulfill

^{*} This work has been partially supported by MIUR through the “Fondo Strategico 2000” Project *VISPO* and the “FIRB 2001” Project *MAIS*. The work of Massimo Mecella has been also partially supported by the European Commission under Contract No. IST-2001-35217, Project *EU-PUBLI.com*.

¹ cf., Service Oriented Computing Net: <http://www.eusoc.net/>

some goals. A requestor discovers a suitable service in the directory, and then it connects to the specific service provider and uses the service.

Research on *e*-Services spans over many interesting issues regarding, in particular, composability, synchronization, coordination, and verification [21]. In this paper, we are particularly interested in automatic *e*-Service composition. *e*-Service *composition* addresses the situation when a client request cannot be satisfied by an available *e*-Service, but a *composite e*-Service, obtained by combining “parts of” available *component e*-Services, might be used. Each composite *e*-Service can be regarded as a kind of client wrt its components, since it (indirectly) looks for and invokes them. *e*-Service composition leads to enhancements of the SOA, by adding new elements and roles, such as brokers and integration systems, which are able to satisfy client needs by combining available *e*-Services.

Composition involves two different issues. The first, sometimes called *composition synthesis*, or simply *composition*, is concerned with synthesizing a new composite *e*-Service, thus producing a specification of how to coordinate the component *e*-Services to obtain the composite *e*-Service. Such a specification can be obtained either *automatically*, i.e., using a tool that implements a composition algorithm, or *manually* by a human. The second issue, often referred to as *orchestration*, is concerned with coordinating the various component *e*-Services according to some given specification, and also monitoring control and data flow among the involved *e*-Services, in order to guarantee the correct execution of the composite *e*-Service, synthesized in the previous phase.

Our main focus in this paper is on automatic composition synthesis. In order to address this issue in an effective and well-founded way, our first contribution is a general formal framework for representing *e*-Services. Note that several works published in the literature address service oriented computing from different points of views (see [11] for a survey), but an agreed comprehension of what an *e*-Service is, in an abstract and general fashion, is still lacking. Our framework, although simplified in several aspects, provides not only a clear definition of *e*-Services, but also a formal setting for a precise characterization of automatic composition of *e*-Services.

The second contribution of the paper is an effective technique for automatic *e*-Service composition. In particular, we specialize the general framework to the case where *e*-Services are specified by means of finite state machines, and we present an algorithm that, given a specification of a target *e*-Service, i.e., specified by a client, and a set of available *e*-Services, synthesizes a composite *e*-Service that uses only the available *e*-Services and fully captures the target one. We also study the computational complexity of our algorithm, and we show that it runs in exponential time with respect to the size of the input state machines.

Although several papers have been already published that discuss either a formal model of *e*-Services (even more expressive than ours, see e.g., [7]), or propose algorithms for computing composition (e.g., [15]), to the best of our knowledge, the work presented in this paper is the first one tackling simultaneously the following issues: *(i)* presenting a formal model where the problem of *e*-Service composition is precisely characterized, *(ii)* providing techniques for computing *e*-Service composition in the case of *e*-Services represented by finite

state machines, and (iii) providing a computational complexity characterization of the algorithm for automatic composition.

The rest of this paper is organized as follows. In Section 2 and 3 we define our general formal framework, and in Section 4 we define the problem of composition synthesis in such a framework. In Section 5 we specialize the general framework to the case where *e*-Services are specified by means of finite state machines, and in Section 6 we present an EXPTIME algorithm for automatic *e*-Service composition in the specialized framework. Finally, in Section 7 we consider related research work and in Section 8 we draw conclusions by discussing future work.

2 General Framework

Generally speaking, an *e*-Service is a software artifact (delivered over the Internet) that interacts with its clients in order to perform a specified task. A client can be either a human user, or another *e*-Service. When executed, an *e*-Service performs its task by directly executing certain actions, and interacting with other *e*-Services to delegate to them the execution of other actions. In order to address SOC from an abstract and conceptual point of view, several facets may be identified [5], each one reflecting a particular aspect of an *e*-Service during its life time. Here, we focus on two of them, namely, (i) the *e*-Service *schema*, specifying functional requirements², i.e., what an *e*-Service does; (ii) an *e*-Service instance, that is an occurrence of an *e*-Service effectively running and interacting with a client. In general, several running instances corresponding to the same *e*-Service schema may exist, each one executing independently from the others.

In order to execute an *e*-Service, the client needs to *activate* an instance from a deployed *e*-Service. In our abstract model, the client can then interact with the *e*-Service instance by repeatedly *choosing* an action and waiting for either the fulfillment of the specific task, or the return of some information. On the basis of the information returned, the client chooses the next action to invoke. In turn, the activated *e*-Service instance executes (the computation associated to) the invoked action; after that, it is ready to execute new actions. Under certain circumstances, i.e., when the client has reached his goal, he may explicitly *end* (i.e., terminate) the *e*-Service instance. However, in principle, a given *e*-Service instance may need to interact with a client for an unbounded, or even infinite, number of steps, thus providing the client with a continuous service. In this case, no operation for ending the *e*-Service instance is ever executed.

When a client invokes an *e*-Service instance *e*, it may happen that *e* does not execute all of its actions on its own, but instead it *delegates* some or all of them to other *e*-Service instances. All this is transparent to the client. To precisely capture the situations when the execution of certain actions can be delegated to other *e*-Service instances, we introduce the notion of *community* of *e*-Services, which is formally characterized by:

² An *e*-Service schema may also specify non-functional requirements, such as those concerning quality or performance. However, non-functional requirements go beyond the scope of this paper.

- a finite common set of actions Σ , called the *action alphabet*, or simply the *alphabet* of the community,
- a set of *e*-Services specified in terms of the common set of actions.

Hence, to join a community, an *e*-Service needs to export its service(s) in terms of the alphabet of the community. The added value of a community is the fact that an *e*-Service of the community may delegate the execution of some or all of its actions to other instances of *e*-Services in the community. We call such an *e*-Service *composite*. If this is not the case, an *e*-Service is called *simple*. Simple *e*-Services realize offered actions directly in the software artifacts implementing them, whereas composite *e*-Services, when receiving requests from clients, can invoke other *e*-Service instances in order to fulfill the client's needs.

Notably, the community can be used to generate (virtual) *e*-Services whose execution completely delegates actions to other members of the community. In other words, the community can be used to realize a target *e*-Service requested by the client, not simply by selecting a member of the community to which delegate the target *e*-Service actions, but more generally by suitably “composing” parts of *e*-Service instances in the community in order to obtain a virtual *e*-Service which “is coherent” with the target one. This function of composing existing *e*-Services on the basis of a target *e*-Service is known as *e*-Service composition, and is the main subject of the research reported in this paper.

3 *E*-service Schema

From the external point of view, i.e., that of a client, an *e*-Service E , belonging to a community C , is seen as a black box that exhibits a certain *exported behavior* represented as sequences of atomic *actions* of C with constraints on their invocation order. From the internal point of view, i.e., that of an application deploying E and activating and running an instance of it, it is also of interest how the actions that are part of the behavior of E are effectively executed. Specifically, it is relevant to specify whether each action is executed by E itself or whether its execution is delegated to another *e*-Service belonging to the community C with which E interacts, transparently to the client of E . To capture these two points of view we introduce the notion of *e*-Service schema, as constituted by two different parts, called *external schema* and *internal schema*, respectively.

Also *e*-Service instances can be characterized by an external and an internal view: further details can be found in [5].

3.1 External Schema

The aim of the external schema is to specify the exported behavior of the *e*-Service. For now we are not concerned with any particular specification formalism, rather we only assume that, whatever formalism is used, the external schema specifies the behavior in terms of a tree of actions, called *external execution tree*. The external execution tree abstractly represents all possible executions of all possible instances of an *e*-Service. Therefore, any instance of an *e*-Service executes a path of such a tree. In this sense, each node x of an external execution

tree represents the history of the sequence of actions of all e -Service instances³, that have executed the path to x . For every action a belonging to the alphabet Σ of the community, and that can be executed at the point represented by x , there is a (single) successor node $x.a$. The node $x.a$ represents the fact that, after performing the sequence of actions leading to x , the client chooses to execute action a , among those possible, thus getting to $x.a$. Therefore, each node represents a choice point at which the client makes a decision on the next action the e -Service should perform. We call the pair $(x, x.a)$ *edge* of the tree and we say that such an edge is *labeled* with action a . The root ε of the tree represents the fact that the e -Service has not yet executed any action. Some nodes of the execution tree are *final*: when a node is final, and only then, the client can stop the execution of the e -Service. In other words, the execution of an e -Service can correctly terminate only at these points⁴.

Notably, an execution tree does not represent the information returned to the client by the e -Service instance execution, since the purpose of such information is to let the client choose the next action, and the rationale behind this choice depends entirely on the client.

Given the external schema E^{ext} of an e -Service E , we denote with $T(E^{ext})$ the external execution tree *specified* by E^{ext} .

3.2 Internal Schema

The internal schema specifies, besides the external behavior of the e -Service, the information on which e -Service instances in the community execute each given action. As before, for now, we abstract from the specific formalism chosen for giving such a specification, instead we concentrate on the notion of *internal execution tree*. An internal execution tree is analogous to an external execution tree, except that each edge is labeled by (a, I) , where a is the executed action and I is a nonempty set denoting the e -Service instances executing a . Every element of I is a pair (E', e') , where E' is an e -Service and e' is the identifier of an instance of E' . The identifier e' uniquely identifies the instance of E' within the internal execution tree. In general, in the internal execution tree of an e -Service E , some actions may be executed also by the running instance of E itself. In this case we use the special instance identifier **this**. Note that, since I is in general not a singleton, the execution of each action can be delegated to more than one other e -Service instance.

An internal execution tree *induces* an external execution tree: given an internal execution tree T_{int} we call *offered external execution tree* the external execution tree T_{ext} obtained from T_{int} by dropping the part of the labeling denoting the e -Service instances, and therefore keeping only the information on the actions. An internal execution tree T_{int} *conforms to* an external execution tree T_{ext} if T_{ext} is equal to the offered external execution tree of T_{int} .

³ In what follows, we omit the terms “schema” and “instance” when clear from the context.

⁴ Typically, in an e -Service, the root is final, to model that the computation of the e -Service may not be started at all by the client.

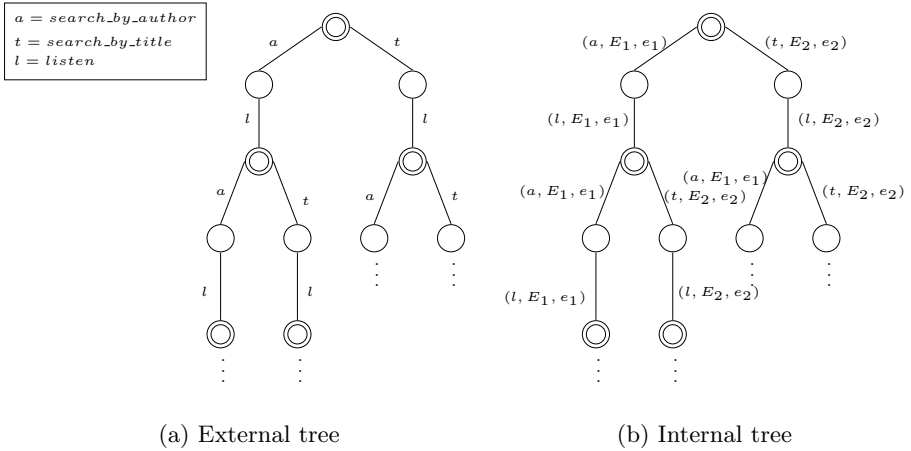


Fig. 1. E -service execution trees

Given an e -Service E , the internal schema E^{int} of E is a *specification* that uniquely represents an internal execution tree. We denote such an internal execution tree by $T(E^{int})$.

We now formally define when an e -Service of a community correctly delegates actions to other e -Services of the community. We need a preliminary definition: given the internal execution tree T_{int} of an e -Service E , and a path p in T_{int} starting from the root, we call the *projection* of p on an instance e' of an e -Service E' the path obtained from p by removing each edge whose label (a, I) is such that I does not contain e' , and collapsing start and end node of each removed edge.

We say that the internal execution tree T_{int} of an e -Service E is *coherent* with a community C if:

- for each edge labeled with (a, I) , the action a is in the alphabet of C , and for each pair (E', e') in I , E' is a member of the community C ;
- for each path p in T_{int} from the root of T_{int} to a node x , and for each pair (E', e') appearing in p , with e' different from **this**, the projection of p on e' is a path in the external execution tree T'_{ext} of E' from the root of T'_{ext} to a node y , and moreover, if x is final in T_{int} , then y is final in T'_{ext} .

Observe that, if an e -Service of a community C is simple, i.e., it does not delegate actions to other e -Service instances, then it is trivially coherent with C . Otherwise, it is composite and hence delegates actions to other e -Service instances. In the latter case, the behavior of each one of such e -Service instances must be correct according to its external schema.

Example 1. Figure 1(a) shows (a portion of) an (infinite) external execution tree representing an e -Service that allows for searching and listening to mp3 files⁵.

⁵ Final nodes are represented by two concentric circles.

In particular, the client may choose to search for a song by specifying either its author(s) or its title (action `search_by_author` and `search_by_title`, respectively). Then the client selects and listens to a song (action `listen`). Finally, the client chooses whether to perform those actions again.

Figure 1(b)⁶ shows (a portion of) an (infinite) internal execution tree, conforming to the previous external execution tree, where all the actions are delegated to e -Services of the community. In particular, the execution of `search_by_title` action and its subsequent `listen` action are delegated to instance e_2 of e -Service E_2 , and `search_by_author` action and its subsequent `listen` action to instance e_1 of e -Service E_1 . \square

4 Composition Synthesis

When a user requests a certain service from an e -Service community, there may be no e -Service in the community that can deliver it directly. However, it may still be possible to synthesize a new composite e -Service, which suitably delegates action execution to the e -Services of the community, and when suitably orchestrated, provides the user with the service he requested. Formally, given an e -Service community C and the external schema E^{ext} of a target e -Service E expressed in terms of the alphabet Σ of C , a *composition* of E wrt C is an internal schema E^{int} such that (i) $T(E^{int})$ conforms to $T(E^{ext})$, (ii) $T(E^{int})$ delegates all actions to the e -Services of C (i.e., **this** does not appear in $T(E^{int})$), and (iii) $T(E^{int})$ is coherent with C .

The problem of *composition existence* is the problem of checking whether there exists some internal schema E^{int} that is a composition of E wrt C . Observe that, since for now we are not placing any restriction of the form of E^{int} , this corresponds to checking if there exists an internal execution tree T_{int} such that (i) T_{int} conforms to $T(E^{ext})$, (ii) T_{int} delegates all actions to the e -Services of C , and (iii) T_{int} is coherent with C .

The problem of *composition synthesis* is the problem of synthesizing an internal schema E^{int} for E that is a composition of E wrt C .

An *e -Service Integration System* delivers possibly composite e -Services on the basis of user requests, exploiting the available e -Services of a community C . When a client requests a new e -Service E , he presents his request in the form of an external e -Service schema E^{ext} for E , and expects the e -Service Integration System to execute an instance of E . To do so, first a *composer* module makes the composite e -Service E available for execution, by synthesizing an internal schema E^{int} of E that is a composition of E wrt the community C . Then, following the internal execution tree $T(E^{int})$ specified by E^{int} , an *orchestration engine* activates an (internal) instance of E , and orchestrates the different available e -Services, by activating and interacting with their external view, so as to fulfill the client's needs.

⁶ In the figure, each action is delegated to exactly one instance of an e -Service schema. Hence, for simplicity, we have denoted a label $(a, \{(E_i, e_i)\})$ simply by (a, E_i, e_i) , for $i = 1, 2$.

The orchestration engine is also in charge of terminating the execution of component e -Service instances, offering the correct set of actions to the client, as defined by the external execution tree, and invoking the action chosen by the client on the e -Service that offers it.

All this happens in a transparent manner for the client, who interacts only with the e -Service Integration System and is not aware that a composite e -Service is being executed instead of a simple one.

5 E -services as Finite State Machines

Till now, we have not referred to any specific form of e -Service schemas. In what follows, we consider e -Services whose schema (both internal and external) can be represented using only a *finite number of states*, i.e., using (deterministic) Finite State Machines (FSMs).

The class of e -Services that can be captured by FSMs are of particular interest. This class allows us to address an interesting set of e -Services, that are able to carry on rather complex interactions with their clients, performing useful tasks. Indeed, several papers in the e -Service literature adopt FSMs as the basic model of exported behavior of e -Services [7,6]. Also, FSMs constitute the core of statecharts, which are one of the main components of UML and are becoming a widely used formalism for specifying the dynamic behavior of entities.

In the study we report here, we make the simplifying assumption that the number of instances of an e -Service in the community that can be involved in the internal execution tree of another e -Service is bounded and fixed a priori. In fact, wlog we assume that it is equal to one. If more instances correspond to the same external schema, we simply duplicate the external schema for each instance. Since the number of e -Services in a community is finite, the overall number of instances orchestrated by the orchestrator in executing an e -Service is finite and bounded by the number of e -Services belonging to the community. Within this setting, in the next section, we show how to solve the composition problem, and how to synthesize a composition that is a FSM. Instead, how to deal with an unbounded number of instances remains open for future work.

We consider here e -Services whose external schemas can be represented with a finite number of states. Intuitively, this means that we can factorize the sequence of actions executed at a certain point into a finite number of states, which are sufficient to determine the future behavior of the e -Service. Formally, for an e -Service E , the external schema of E is a FSM $A_E^{ext} = (\Sigma, S_E, s_E^0, \delta_E, F_E)$, where:

- Σ is the alphabet of the FSM, which is the alphabet of the community;
- S_E is the set of states of the FSM, representing the finite set of states of the e -Service E ;
- s_E^0 is the initial state of the FSM, representing the initial state of the e -Service;
- $\delta_E : S_E \times \Sigma \rightarrow S_E$ is the (partial) transition function of the FSM, which is a partial function that given a state s and an action a returns the state resulting from executing a in s ;




Fig. 2. E -service specification as FSM

- $F_E \subseteq S_E$ is the set of final states of the FSM, representing the set of states that are final for the e -Service E , i.e., the states where the interactions with E can be terminated.

The FSM A_E^{ext} is an external schema in the sense that it specifies an external execution tree $T(A_E^{ext})$. Specifically, given A_E^{ext} we define $T(A_E^{ext})$ inductively on the level of nodes in the tree, by making use of an auxiliary function $\sigma(\cdot)$ that associates to each node of the tree a state in the FSM. We proceed as follows:

- ε , as usual, is the root of $T(A_E^{ext})$ and $\sigma(\varepsilon) = s_E^0$;
- if x is a node of $T(A_E^{ext})$, and $\sigma(x) = s$, for some $s \in S_E$, then for each a such that $s' = \delta_E(s, a)$ is defined, $x \cdot a$ is a node of $T(A_E^{ext})$ and $\sigma(x \cdot a) = s'$;
- x is final iff $\sigma(x) \in F_E$.

Figure 2(a) shows a FSM that is a specification for the external execution tree of Figure 1(a). Note that in general there may be several FSMs that may serve as such a specification.

Since we have assumed that each e -Service in the community can contribute to the internal execution tree of another e -Service with at most one instance, in specifying internal execution trees we do not need to distinguish between e -Services and e -Service instances. Hence, when the community C is formed by n e -Services E_1, \dots, E_n , it suffices to label the internal execution tree of an e -Service E by the action that caused the transition and a subset of $[n] = \{1, \dots, n\}$ that identifies which e -Services in the community have contributed in executing the action. The empty set \emptyset is used to (implicitly) denote **this**.

We are interested in internal schemas, for an e -Service E , that have a finite number of states, i.e., that can be represented as a Mealy FSM (MFSM) $A_E^{int} = (\Sigma, 2^{[n]}, S_E^{int}, s_E^0, \delta_E^{int}, \omega_E^{int}, F_E^{int})$, where:

- $\Sigma, S_E^{int}, s_E^0, \delta_E^{int}, F_E^{int}$, have the same meaning as for A_E^{ext} ;
- $2^{[n]}$ is the output alphabet of the MFSM, which is used to denote which e -Service instances execute each action;
- $\omega_E^{int} : S_E^{int} \times \Sigma \rightarrow 2^{[n]}$ is the output function of the MFSM, that, given a state s and an action a , returns the subset of e -Services that executes action a when the e -Service E is in the state s ; if such a set is empty then **this** is implied; we assume that the output function ω_E^{int} is defined exactly when δ_E^{int} is so.

The MFSM A_E^{int} is an internal schema in the sense that it specifies an internal execution tree $T(A_E^{int})$. Given A_E^{int} we, again, define the internal execution tree $T(A_E^{int})$ by induction on the level of the nodes, by making use of an auxiliary function $\sigma^{int}(\cdot)$ that associates each node of the tree with a state in the MFSM, as follows:

- ε is, as usual, the root of $T(A_E^{int})$ and $\sigma^{int}(\varepsilon) = s_E^0$;
- if x is a node of $T(A_E^{int})$, and $\sigma^{int}(x) = s$, for some $s \in S_E^{int}$, then for each a such that $s' = \delta_E^{int}(s, a)$ is defined, $x \cdot a$ is a node of $T(A_E^{int})$ and $\sigma^{int}(x \cdot a) = s'$;
- if x is a node of $T(A_E^{int})$, and $\sigma^{int}(x) = s$, for some $s \in S_E^{int}$, then for each a such that $\omega_E^{int}(s, a)$ is defined (i.e., $\delta_E^{int}(s, a)$ is defined), the edge $(x, x \cdot a)$ of the tree is labeled by $\omega_E^{int}(s, a)$;
- x is final iff $\sigma^{int}(x) \in F_E^{int}$.

As an example, Figure 2(b) shows a MFSM that is a specification of an internal execution tree that conforms to the external execution tree specified by the FSM of Figure 2(a). Indeed the MFSM in the figure compactly represents the e -Service whose internal execution tree is shown in Figure 1(b). In general, an external schema specified as FSM and its corresponding internal schema specified as MFSM may have different structures, as the example shows.

Given an e -Service E whose external schema is an FSM and whose internal schema is an MFSM, checking whether E is well formed, i.e., whether the internal execution tree conforms to the external execution tree, can be done using standard finite state machine techniques. Similarly for coherency of E with a community of e -Services whose external schemas are FSMs. In this paper, we do not go into the details of these problems, and instead we concentrate on composition.

6 Automatic E -service Composition

We address the problem of actually checking the existence of a composite e -Service in the FSM-based framework introduced above. We show that if a composition exists then there is one where the internal schema is constituted by a MFSM, and we show how to actually synthesize such a MFSM. The basic tool we use to show such results is reducing the problem of composition existence into satisfiability of a suitable formula of Deterministic Propositional Dynamic Logic (DPDL), a well-known logic of programs developed to verify properties of program schemas [12].

We start from a set \mathcal{P} of atomic propositions and a set \mathcal{A} of *deterministic* atomic actions and we define DPDL formulas as follows:

$$\phi \longrightarrow P \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \langle r \rangle \phi \mid [r] \phi$$

where P is an atomic proposition in \mathcal{P} and r is a regular expression over the set of actions in \mathcal{A} . That is, DPDL formulas are composed from atomic propositions by applying arbitrary propositional connectives, and modal operators $\langle r \rangle \phi$ and $[r] \phi$.

The meaning of the latter two is, respectively, that there exists an execution of r reaching a state where ϕ holds, and that all terminating executions of r reach a state where ϕ holds. Let u be an abbreviation for $(\bigcup_{a \in \mathcal{A}} a)^*$. Then $[u]$ represents the *master modality*, which can be used to state universal assertions.

A DPDL interpretation is a (deterministic) Kripke structure of the form $\mathcal{I} = (\Delta^{\mathcal{I}}, \{a^{\mathcal{I}}\}_{a \in \mathcal{A}}, \{P^{\mathcal{I}}\}_{P \in \mathcal{P}})$, where $a^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ is a partial function from elements of $\Delta^{\mathcal{I}}$ to elements of $\Delta^{\mathcal{I}}$, and $P^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ are all the elements of $\Delta^{\mathcal{I}}$ where P is true. Such interpretation is then extended to formulas and complex actions in a standard way, see [12] for details.

DPDL enjoys two properties that are of particular interest for our aims. The first is the *tree model property*, which says that every model of a formula can be unwound to a (possibly infinite) tree-shaped model (considering domain elements as nodes and partial functions interpreting actions as edges). The second is the *small model property*, which says that every satisfiable formula admits a finite model whose size (in particular the number of domain elements) is at most exponential in the size of the formula itself.

Given the target e -Service E_0 whose external schema is a FSM A_0 and a community of e -Services formed by n component e -Services E_1, \dots, E_n whose external schemas are FSMs A_1, \dots, A_n respectively, we build a DPDL formula Φ as follows. As set of atomic propositions \mathcal{P} in Φ we have (i) one proposition s_j for each state s_j of A_j , $j = 0, \dots, n$, that is true if A_j is in state s_j ; (ii) propositions F_j , $j = 0, \dots, n$, denoting whether A_j is in a final state; and (iii) propositions $moved_j$, $j = 1, \dots, n$, denoting whether (component) automaton A_j performed a transition. As set of atomic actions \mathcal{A} in Φ we have the actions in Σ (i.e., $\mathcal{A} = \Sigma$). The formula Φ is built as a conjunction of the following formulas.

- The formulas representing $A_0 = (\Sigma, S_0, s_0^0, \delta_0, F_0)$:
 - $[u](s \rightarrow \neg s')$ for all pairs of states $s \in S_0$ and $s' \in S_0$, with $s \neq s'$; these say that propositions representing different states are disjoint (cannot be true simultaneously).
 - $[u](s \rightarrow \langle a \rangle \text{true} \wedge [a]s')$ for each a such that $s' = \delta_0(s, a)$; these encode the transitions of A_0 .
 - $[u](s \rightarrow [a]\text{false})$ for each a such that $\delta(s, a)$ is not defined; these say when a transition is not defined.
 - $[u](F_0 \leftrightarrow \bigvee_{s \in F_0} s)$; this highlights final states of A_0 .
- For each component FSM $A_i = (\Sigma, S_i, s_i^0, \delta_i, F_i)$, the following formulas:
 - $[u](s \rightarrow \neg s')$ for all pairs of states $s \in S_i$ and $s' \in S_i$, with $s \neq s'$; these again say that propositions representing different states are disjoint.
 - $[u](s \rightarrow [a](moved_i \wedge s' \vee \neg moved_i \wedge s))$ for each a such that $s' = \delta_i(s, a)$; these encode the transitions of A_i , conditioned to the fact that the component A_i is actually required to make a transition a in the composition.
 - $[u](s \rightarrow [a]\neg moved_i)$ for each a such that $\delta_i(s, a)$ is not defined; these say that when a transition is not defined, A_i cannot be asked to execute in the composition.
 - $[u](F_i \leftrightarrow \bigvee_{s \in F_i} s)$; this highlights final states of A_i .

– Finally, the following formulas:

- $s_0^0 \wedge \bigwedge_{i=1, \dots, n} s_i^0$; this says that initially all e -Services are in their initial state; note that this formula is not prefixed by $[u]$.
- $[u](\langle a \rangle \mathbf{true} \rightarrow [a] \bigvee_{i=1, \dots, n} \mathit{moved}_i)$, for each $a \in \Sigma$; these say that at each step at least one of the component FSM has moved.
- $[u](F_0 \rightarrow \bigwedge_{i=1, \dots, n} F_i)$; this says that when the target e -Service is in a final state also all component e -Services must be in a final state.

Theorem 1. *The DPDL formula Φ , constructed as above, is satisfiable if and only if there exists a composition of E_0 wrt E_1, \dots, E_n .*

Proof (sketch). “ \Leftarrow ” Suppose that there exists some internal schema (without restriction on its form) E_0^{int} which is a composition of E_0 wrt E_1, \dots, E_n . Let $T_{int} = T(E_0^{int})$ be the internal execution tree defined by E_0^{int} .

Then for the target e -Service E_0 and each component e -Service E_i , $i = 1, \dots, n$, we can define mappings σ and σ_i from nodes in T_{int} to states of A_0 and A_i , respectively, by induction on the level of the nodes in T_{int} as follows.

- base case: $\sigma(\varepsilon) = s_0^0$ and $\sigma_i(\varepsilon) = s_i^0$.
- inductive case: let $\sigma(x) = s$ and $\sigma_i(x) = s_i$, and let the node $x \cdot a$ be in T_{int} with the edge $(x, x \cdot a)$ labeled by (a, I) , where $I \subseteq [n]$ and $I \neq \emptyset$ (notice that **this** may not occur since T_{int} is specified by a composition). Then we define

$$\sigma(x \cdot a) = s' = \delta_0(s, a)$$

and

$$\sigma_i(x \cdot a) = \begin{cases} s_i' = \delta_i(s_i, a) & \text{if } i \in I \\ s_i & \text{if } i \notin I \end{cases}$$

Once we have σ and σ_i in place we can define a model $\mathcal{I} = (\Delta^{\mathcal{I}}, \{a^{\mathcal{I}}\}_{a \in \Sigma}, \{P^{\mathcal{I}}\}_{P \in \mathcal{P}})$ of Φ as follows:

- $\Delta^{\mathcal{I}} = \{x \mid x \in T_{int}\}$;
- $a^{\mathcal{I}} = \{(x, x \cdot a) \mid x, x \cdot a \in T_{int}\}$, for each $a \in \Sigma$;
- $s^{\mathcal{I}} = \{x \in T_{int} \mid \sigma(x) = s\}$, for all propositions s corresponding to states of A_0 ;
- $s_i^{\mathcal{I}} = \{x \in T_{int} \mid \sigma_i(x) = s_i\}$, for all propositions s_i corresponding to states of A_i ;
- $\mathit{moved}_i^{\mathcal{I}} = \{x \cdot a \mid (x, x \cdot a) \text{ is labeled by } I \text{ with } i \in I\}$, for $i = 1, \dots, n$;
- $F_0^{\mathcal{I}} = \{x \in T_{int} \mid \sigma(x) = s \text{ with } s \in F_0\}$;
- $F_i^{\mathcal{I}} = \{x \in T_{int} \mid \sigma_i(x) = s_i \text{ with } s_i \in F_i\}$, for $i = 1, \dots, n$.

It is easy to check that, being T_{int} specified by a composition E_{int} , the above model indeed satisfies Φ .

“ \Rightarrow ” Let Φ be satisfiable and $\mathcal{I} = (\Delta^{\mathcal{I}}, \{a^{\mathcal{I}}\}_{a \in \Sigma}, \{P^{\mathcal{I}}\}_{P \in \mathcal{P}})$ be a tree-like model. From \mathcal{I} we can build an internal execution tree T_{int} for E_0 as follows.

- the nodes of the tree are the elements of $\Delta^{\mathcal{I}}$; actually, since \mathcal{I} is tree-like we can denote the elements in $\Delta^{\mathcal{I}}$ as nodes of a tree, using the same notation that we used for internal/external execution tree;
- nodes x such that $x \in F_0^{\mathcal{I}}$ are the final nodes;
- if $(x, x \cdot a) \in a^{\mathcal{I}}$ and for all $i \in I$, $x \cdot a \in moved_i^{\mathcal{I}}$ and for all $j \notin I$, $x \cdot a \notin moved_j^{\mathcal{I}}$, then $(x, x \cdot a)$ is labeled by (a, I) .

It is possible to show that: (i) T_{int} conforms to $T(A_0)$, (ii) T_{int} delegates all actions to the e -Services of E_1, \dots, E_n , and (iii) T_{int} is coherent with E_1, \dots, E_n . Since we are not placing any restriction on the kind of specification allowed for internal schemas, it follows that there exists an internal schema E_{int} that is a composition of E_0 wrt E_1, \dots, E_n . \square

Observe that the size of Φ is polynomially related to the size of A_0, A_1, \dots, A_n . Hence, from the EXPTIME-completeness of satisfiability in DPDL and from Theorem 1 we get the following complexity result.

Theorem 2. *Checking the existence of an e -Service composition can be done in EXPTIME.*

Observe that, because of the small model property, from Φ one can always obtain a model which is at most exponential in the size of Φ . From such a model one can extract an internal schema for E_0 that is a composition of E_0 wrt E_1, \dots, E_n , and has the form of a MFSM. Specifically, given a finite model $\mathcal{I} = (\Delta^{\mathcal{I}}, \{a^{\mathcal{I}}\}_{a \in \Sigma}, \{P^{\mathcal{I}}\}_{P \in \mathcal{P}})$, we define such an MFSM $A_c = (\Sigma, 2^{[n]}, S_c, s_c^0, \delta_c, \omega_c, F_c)$ as follows:

- $S_c = \Delta^{\mathcal{I}}$;
- $s_c^0 = d_0$ where $d_0 \in (s_0^0 \wedge \bigwedge_{i=1, \dots, n} s_i^0)^{\mathcal{I}}$;
- $s' = \delta_c(s, a)$ iff $(s, s') \in a^{\mathcal{I}}$;
- $I = \omega_c(s, a)$ iff $(s, s') \in a^{\mathcal{I}}$ and for all $i \in I$, $s' \in moved_i^{\mathcal{I}}$ and for all $j \notin I$, $s' \notin moved_j^{\mathcal{I}}$;
- $F_c = F_0^{\mathcal{I}}$.

As a consequence of this, we get the following result.

Theorem 3. *If there exists a composition of E_0 wrt E_1, \dots, E_n , then there exists one which is a MFSM of at most exponential size in the size of the external schemas A_0, A_1, \dots, A_n of E_0, E_1, \dots, E_n respectively.*

Proof (sketch). By Theorem 1, if A_0 can be obtained by composing A_1, \dots, A_n , then the DPDL formula Φ constructed as above is satisfiable. In turn, if Φ is satisfiable, for the small-model property of DPDL there exists a model \mathcal{I} of size at most exponential in Φ , and hence in A_0 and A_1, \dots, A_n . From \mathcal{I} we can construct a MFSM A_c as above. It is possible to show that the internal execution tree $T(A_c)$ defined by A_c satisfies all the conditions required for A_c to be a composition, namely: (i) $T(A_c)$ conforms to $T(A_0)$, (ii) $T(A_c)$ delegates all actions to the e -Services of E_1, \dots, E_n , and (iii) $T(A_c)$ is coherent with E_1, \dots, E_n . \square

In [4] a detailed example is provided, that explains the composition synthesis algorithm step by step.

From a practical point of view, because of the correspondence between Propositional Dynamic Logics (which DPDL belongs to) and Description Logics (DLs [8]), one can use current highly optimized DL-based systems [3]⁷ to check the existence of *e*-Service compositions. Indeed, these systems are based on tableaux techniques that construct a model when checking for satisfiability, and from such a model one can construct a MFSM that is the composition.

7 Related Work

Up to now, research on *e*-Services has mainly concentrated on the issues of (*i*) service description and modeling, and (*ii*) service composition, including synthesis and orchestration.

Current research in description and modeling of *e*-Services is mainly founded on the work on workflows, which model business processes as sequences of (possibly partially) automated activities, in terms of data and control flow among them (e.g., [18,19]). In [14] *e*-Services are represented as statecharts, and in [7], an *e*-Service is modeled as a Mealy machine, with input and output messages, and a queue is used to buffer messages that were received but not yet processed.

In our paper, we model *e*-Services as finite state machines, even if we do not consider communication delays and therefore any concept of message queuing is not taken into account. Indeed, from the survey of [11], it stems that the most practical approaches for modeling and describing *e*-Services are the ones based on specific forms of state machines. Additionally, our model of *e*-Service is oriented towards representing the interactions between a client and an *e*-Service. Therefore, our focus is on action sequences, rather than on message sequences as in [7], or on actions with input/output parameters as in [15].

Orchestration requires that the composite *e*-Service is specified in a precise way, considering both the specification of how various component *e*-Services are linked and the internal process flow of the component one. In [11], different technologies, standards and approaches for specification of composite *e*-Services are considered, including BPEL4WS, BPML, AZTEC, etc. In [11] three different kinds of composition are identified: (*i*) peer-to-peer, in which the individual *e*-Services are equals, (*ii*) the mediated approach, based on a hub-and-spoke topology, in which one service is given the role of process mediator, and (*iii*) the brokered approach, where process control is centralized but data can pass between component *e*-Services. With respect to such a classification, the approach proposed in this paper belongs to the mediated one. Analogously, research works reported in [9,19,13] can be classified into the mediated approach to composition. Conversely in [10] the enactment of a composite *e*-Service is carried out in a decentralized way, through peer-to-peer interactions. In [7], a peer-to-peer

⁷ In fact, current DL-based systems cannot handle Kleene star. However, since in Φ , $*$ is only used to mimic universal assertions, and such systems have the ability of handling universal assertions, they can indeed check satisfiability of Φ .

approach is considered, and the interplay between a composite *e*-Service and component ones is analyzed, also in presence of unexpected behavior.

The *DAML-S Coalition* [2] is defining a specific ontology and a related language for *e*-Services, with the aim of composing them in automatic way. In [20] the issue of service composition is addressed, in order to create composite services by re-using, specializing and extending existing ones; in [15] composition of *e*-Services is addressed by using GOLOG and providing a semantics of the composition based on Petri Nets. In [1] a way of composing *e*-Services is presented, based on planning under uncertainty and constraint satisfaction techniques, and a request language, to be used for specifying client goals, is proposed. *e*-Service composition is indeed a form of program synthesis as is planning. The main conceptual difference is that, while in planning we are typically interested in synthesizing a *new* sequence of actions (or more generally a program, i.e., an execution tree) that achieves the client goal, in *e*-Service composition, we try to obtain (the execution tree of) the target *e*-Service by *reusing* in a suitable way fragments of the executions of the component *e*-Services.

8 Conclusions

The main contribution of this paper wrt research on service oriented computing is in tackling *simultaneously* the following issues: (*i*) presenting a formal model where the problem of *e*-Service composition is precisely characterized, (*ii*) providing techniques for computing *e*-Service composition in the case of *e*-Services represented by finite state machines, and (*iii*) providing a computational complexity characterization of the algorithm for automatic composition.

In the future we plan to extend our work both in practical and theoretical directions. On one side, we are developing a DL-based prototype system that implements the composition technique presented in the paper. Such system will enable us to test how the complexity of composition in our framework impacts real world applications. On the theoretical side, we will address open issues such as the characterization of a lower bound for the complexity of the composition problem. Additionally, we plan to extend our setting, by taking into account the possibility that the target *e*-Service is underspecified, as well as the presence of communication delays and of an unbounded number of active instances.

References

1. M. Aiello, M.P. Papazoglou, J. Yang, M. Carman, M. Pistore, L. Serafini, and P. Traverso. A Request Language for Web-Services Based on Planning and Constraint Satisfaction. In *Proc. of VLDB-TES 2002*.
2. A. Ankolekar, M. Burstein, J. Hobbs, O. Lassila, D. Martin, D. McDermott, S. McIlraith, S. Narayanan, M. Paolucci, T. Payne, and K. Sycara. DAML-S: Web Service Description for the Semantic Web. In *Proc. of ISWC 2002*.
3. F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider. *The Description Logic Handbook: Theory, Implementation and Applications*. CUP, 2003.

4. D. Berardi, D. Calvanese, G. De Giacomo, M. Lenzerini, and M. Mecella. Automatic Composition of *e*-Services. Technical Report DIS 22-03 (<http://www.dis.uniroma1.it/~berardi/publications/techRep/TR-22-03.pdf>).
5. D. Berardi, D. Calvanese, G. De Giacomo, M. Lenzerini, and M. Mecella. A Foundational Vision of *e*-Services. In *Proc. of WES 2003*.
6. D. Berardi, F. De Rosa, L. De Santis, and M. Mecella. Finite State Automata as Conceptual Model for *e*-Services. In *Proc. of IDPT 2003*, to appear.
7. T. Bultan, X. Fu, R. Hull, and J. Su. Conversation Specification: A New Approach to Design and Analysis of *E*-Service Composition. In *Proc. of WWW 2003*.
8. D. Calvanese, G. De Giacomo, M. Lenzerini, and D. Nardi. Reasoning in Expressive Description Logics. *Handbook of Automated Reasoning*, ESP, 2001.
9. F. Casati and M.C. Shan. Dynamic and Adaptive Composition of *e*-Services. *Information Systems*, 6(3), 2001.
10. M.C. Fauvet, M. Dumas, B. Benatallah, and H.Y. Paik. Peer-to-Peer Traced Execution of Composite Services. In *Proc. of VLDB-TES 2001*.
11. R. Hull, M. Benedikt, V. Christophides, and J. Su. *E*-Services: A Look Behind the Curtain. In *Proc. of PODS 2003*.
12. D. Kozen and J. Tiuryn. Logics of programs. *Handbook of Theoretical Computer Science — Formal Models and Semantics*, ESP, 1990.
13. M. Mecella and B. Pernici. Building Flexible and Cooperative Applications Based on *e*-Services. Technical Report DIS 21-2002 (http://www.dis.uniroma1.it/~mecella/publications/mp_techreport_212002.pdf).
14. M. Mecella, B. Pernici, and P. Craca. Compatibility of *e*-Services in a Cooperative Multi-Platform Environment. In *Proc. of VLDB-TES 2001*.
15. S. Narayanan and S. McIlraith. Simulation, Verification and Automated Composition of Web Services. In *Proc. of WWW 2002*.
16. M. Papazoglou. Agent-Oriented Technology in Support of *e*-Business. *Communications of the ACM*, 44(4):71-77, 2001.
17. T. Pilioura and A. Tsalgatidou. *e*-Services: Current Technologies and Open Issues. In *Proc. of VLDB-TES 2001*.
18. H. Schuster, D. Georgakopoulos, A. Cichocki, and D. Baker. Modeling and Composing Service-based and Reference Process-based Multi-enterprise Processes. In *Proc. of CAiSE 2000*.
19. G. Shegalov, M. Gillmann, and G. Weikum. XML-enabled Workflow Management for *e*-Services across Heterogeneous Platforms. *VLDB Journal*, 10(1), 2001.
20. J. Yang and M.P. Papazoglou. Web Components: A Substrate for Web Service Reuse and Composition. In *Proc. of CAiSE 2002*.
21. J. Yang, W.J. van den Heuvel, and M.P. Papazoglou. Tackling the Challenges of Service Composition in *e*-Marketplaces. In *Proc. of RIDE-2EC 2002*.