

Composition Constraints for Semantic Web Services^{*}

Zhengang Cheng, Munindar P. Singh, and Mladen A. Vouk[†]

Department of Computer Science

North Carolina State University

Raleigh, NC 27695-7535, USA

{zcheng, mpsingh, vouk}@eos.ncsu.edu

Abstract

Current service description and composition approaches consider simplistic method invocation. They do not accommodate ongoing interactions between service providers and consumers, nor do they support descriptions of legal protocols of interactions among them. We propose richer representations which enable us to capture more of the semantics of services than current approaches. Further, we develop algorithms by which potential problems in service compositions can be detected when services are configured leading to superior execution of composed services.

^{*}This work was partially supported by the Department of Energy through grant DE-FC02-01ER25484, and by the National Science Foundation through grants IIS-9624425 and ITR-0081742.

[†]We are indebted to the anonymous reviewers for helpful comments.

1 Introduction

Web services enable application development and integration over the Web by supporting program-to-program interactions. Relevant standards include Web Services Description Language (WSDL), Universal Description, Discovery and Integration (UDDI), and Simple Object Access Protocol (SOAP) [Christensen et al., 2001; UDDI; Box et al., 2000]. These are intended, respectively, for describing, discovering, and invoking services. While current approaches represent much progress, they carry the baggage of traditional distributed objects approaches. Services are integrated through method invocation without regard to any higher-level constraints.

Current Web service techniques are limited to services where each operation is independent. A common example is stock-quote lookup, where each stock query is unrelated to every other query. But when we move from simplistic information lookup to interactive information search or to e-commerce, especially for complex business-to-business settings, it becomes obvious that current techniques are inadequate. In particular, we face two main challenges.

- What are useful ways of structuring the composition of services? Method invocation is appropriate for closed systems, but services are inherently autonomous and often to be used in long-lived interactions. For example, a long-lived interaction occurs in e-commerce when you try to change an order because of some unexpected conditions or try to get a refund for a faulty product. Even short-lived settings involve protocols, e.g., checking if the service requester is authenticated and properly authorized before accepting its order.
- What are appropriate semantic constraints on how services may participate in different compositions? As observed above, it may be required to compose services so as to carry out certain kinds of interactions. However, if a given service does not support such interactions, then the composition can fail at run time in unexpected ways. For example, if an e-commerce service does not allow orders to be changed after a certain time has elapsed, it may be unsafe to use this service in certain settings, even if it is superior in other respects.

Current approaches, based on traditional closed systems, fall short of handling the above situations. WSDL allows us to capture the various methods but does not support constraints among those methods. Either too many methods will always be enabled or too few. However, WSDL's functionality is required to specify the methods supported by a service. Recently, Web Services Flow Language (WSFL) was proposed to describe compositions of services in the form of a workflow [Leymann, 2001]. WSFL specifications tell us how different services ought to be invoked, e.g., in terms of ordering and parallelism among them, but they don't tell us whether a particular service that is bound in the workflow will in fact deliver the right interactions. XLANG [Thatte, 2001] fits in Microsoft's BizTalk Server architecture. It describes the behavior of a single Web service. XLANG as it stands today provides a notation similar in spirit to workflow languages.

Our contribution is in capturing deeper constraints on what services are willing to offer, capturing richer requirements for service composition, and comparing the two to decide if a particular service is appropriate for the intended composition. Roughly, we think of WSDL as providing an underlying layer for our work. We enhance WSFL to define a service composition language that provides the necessary inputs to our reasoning approach. XLANG specifications are envisioned to drive automated protocol engines that will ensure that the specified message flows are obtained. Further, the intended direction is to define message exchanges among Web services. In this expanded form, XLANG will relate to our approach by providing elements of a service composition language.

Although the work reported here is still in an early stage of development, we believe it addresses important real-world concerns in the future expansion of the semantic Web. In particular, for the semantic Web technologies to penetrate real-life enterprise applications, the semantic Web will need as strong a representation of actions and processes as of data.

Section 2 introduces our representations and Section 3 applies them for verifying service compositions.

2 Representing Composition Constraints

Because of their autonomy and heterogeneity, services are naturally associated with *agents*. Agents are long-lived, persistent computations that can perceive, reason, act, and communicate [Huhns and Singh, 1998a]. Agents act with varying levels of autonomy depending on environmental constraints and their previous commitments. Each agent provides a service and interacts through the exchange of messages, which denote business documents. We propose the Agent Service Description Language (ASDL) to describe the external behavior of agent services. An ASDL specification describes the messages understood by a service along with an interaction protocol it follows. Like WSDL, ASDL can be published and accessed through a registry. We propose the Agent Service Composition Language (ASCL) to describe the composition of new services from existing services. The imported services represent an agent role with its behavior described in its ASDL description. An ASCL specification describes the interaction with other agent services.

2.1 Agent Service Description Language

ASDL enables us to define the behavioral characteristics of a Web Service. A Web service that implements behavior to preserve its autonomy is considered a Web agent service. In essence, ASDL is an behavioral extension to WSDL. It describes the constraints of service invocation to capture the external visible relationship between the operations.

The external behavior of the agent is demonstrated by its interaction with other agents. Such interactions occur through message exchanges. The internal implementation and reasoning logic is governed by the autonomy of the agent. For simplicity, we describe agent behavior via a finite state machine that models the allowed operation sequences.

The behavior of a service provider describes the allowed invocation sequence of operations. Invocations of this service must satisfy this sequence. We describe the agent behavior through a set of states and transitions between the states.

The state of an agent is used to maintain semantic constraints on actions. For example, a seller may require a buyer to log in before ordering some thing. Thus we can describe legal sequences of operation invocation. An operation state has a name for referencing and the operation it represents. For example, we can represent a state for the “order” operation defined in WSDL as follows.

```
<state name="order" operation="order">
```

ASDL contains states for all operations defined in WSDL with the name and operation attributes same as the operation name. We allow empty states which only have a name, but with no operation defined. They can be used to denote semantic states such as “ordered.”

An interaction can proceed from one operation to another only if allowed by the modeled transitions. Each transition has a source operation and a destination operation. It may associate a condition to specify when the transition can occur. For example, the following specifies that the customer must successfully “Login” before “Ordering” any product. (The gating condition “expr” can involve terms from the messages exchanged in the protocol.)

<transition source="Login" target="Order" condition="expr" / >

If there are no other transitions into "order," the customer must log in before ordering.

2.2 An E-Commerce Example

Figure 1 describes the behavior of an agent who requires login first, then enables querying and ordering products, leading to checkout and payment, and finally shipping.

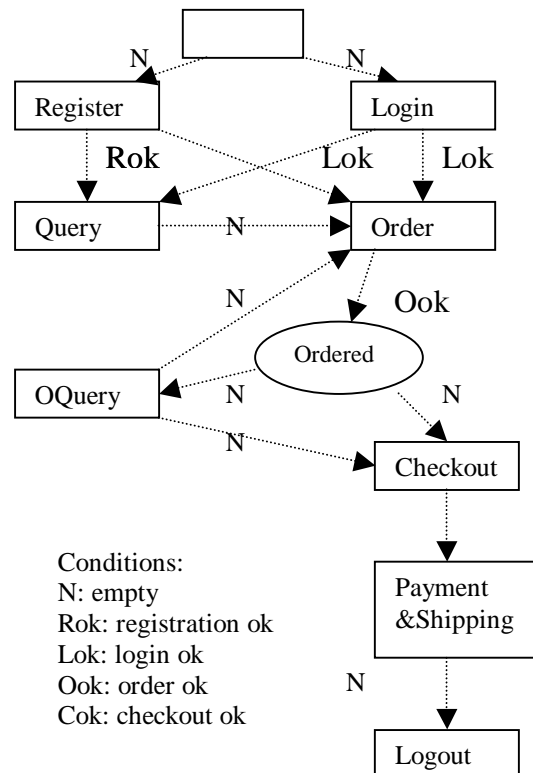


Figure 1: Behavior of A Seller Agent

We describe the seller's behavior with the behavioral extension of WSDL.

```

<behavior name="seller">
  <states>
    <state name="start" operation="" / >
    <state name="OQuery" operation="Query" / >
    <state name="Ordered" operation="" / >
  </states>
  <transitions>
    <transition source="Start" target="Register" condition="NULL" / >
    <transition source="Start" target="Login" condition="NULL" / >
    <transition source="Register" target="Query" condition="Rok" / >
    <transition source="Register" target="Order" condition="Rok" / >
    <transition source="Login" target="Query" condition="Lok" / >
  </transitions>
</behavior>

```

```

<transition source="Login" target="Order" condition="Lok" / >
<transition source="Query" target="Order" condition="NULL" / >
<transition source="Order" target="Ordered" condition="Rok" / >
<transition source="Ordered" target="OQuery" condition="NULL" / >
<transition source="Ordered" target="Checkout" condition="NULL" / >
<transition source="OQuery" target="Checkout" condition="NULL" / >
<transition source="OQuery" target="Order" condition="NULL" / >
<transition source="Checkout" target="PaymentShip" condition="Cok" / >
<transition source="PaymentShip" target="logout" condition="PoK" / >
< /transitions>
< /behavior>

```

The following is a description of the transition conditions.

- NULL: This represents an empty condition.
- Rok, meaning registration OK, is a boolean expression based on the RegResult output message, and can be evaluated when the buyer receive the RegResult message.
- Lok, meaning Login OK, is a boolean expression based on the LoginResult output message.
- Ook, meaning Order OK, is a boolean expression based on the OrderResult output message. It implies that the product is successfully ordered.
- Cok, meaning Checkout OK, implies that an invoice is successfully send to the buyer.

The state *OQuery* represents that the buyer has ordered something and can checkout at any time.

2.3 Agent Service Composition Language

Agent Service Composition Language (ASCL) specifications describe the logic of how a new service is composed from existing services. For reasons of space, we don't include details of ASCL syntax here. Suffice it to state that it enables services to be bound in, and various flow primitives (sequencing, branching, and so on) used to specify desired compositions.

3 Verification of Service Composition

Using the above representations, we can determine is service protocols will be violated by the desired compositions. This can be done when services are bound—that is, at configuration rather than at run time.

We have developed some algorithms for reasoning based on the above representations. We lack the space to discuss these in detail. Briefly, these algorithms enable us to construct behavior state diagrams, find legal operations, build an execution graph. We present the most interesting algorithm, which is for checking compliance of operation invocations.

To verify compliance at the operation level, we need the following data structures and functions.

- Build a service transition graph for each imported services.
- Find legal operations. Based on the service diagram, we can find what are legal next states for a given current state.

- Remember history of invocation.

Given the above, following algorithm verifies a composition by traversing all possible execution paths.

1. Do a depth first search on the resulted operation graph.
2. For each invocation of imported services, find the last state from the history, find the legal set of operation-states for the last state and current service from the corresponding service transition graph, check whether current operation is in the legal operation-state set. If not, there is an error. Otherwise add current operation to the history.
3. Continue until the whole operation graph is traversed.

Say a buyer finds the seller's service specification from the registry and would like to use the service to purchase a product from the seller. Figure 1 describes the seller's behavior in ASDL. It is up to the designer of the buyer agent on how to utilize the seller service. The verification algorithm is used for ensure that all possible execution paths in a service composition are legal. The following are some simplified examples to show how a service composition can be verified.

3.1 Sequential Composition

Suppose the buyer would like to implement the buy activity to buy products. For brevity, unnecessary XML tags are removed.

```

    <consume role="seller" urlref="seller.xml" / >
    <activity name="buy">
    <sequence>
    <operation name="login" performedby="seller" / >
    <operation name="query" performedby="seller" / >
    <operation name="order" performedby="seller" / >
    < /sequence>
    < /activity>

```

Here the buyer invokes the login, query, order operations of seller in sequence. Since login, query, order are on the execution path, it is obvious that the above sequence is valid according to the seller's behavior.

3.2 Complex Composition

Now we consider the case where a buyer uses two seller services.

```

    <consume role="S0" urlref="seller0.xml" / >
    <consume role="S1" urlref="seller1.xml" / >
    <activity name="buy">
    <operation name="register" performedby="S1" / >
    <fork condition="expr">
    <thread name="t1">
    <sequence>
    <operation name="login" performedby="S0" / >
    <if condition="expr">
    <sequence>

```

```

                <operation name="query" performedby="S1" / >
                <operation name="order" performedby="S0" / >
            </sequence>
        </if>
        <operation name="order" performedby="S0" / >
    </sequence>
< /thread>
<thread name="t2">
    <sequence>
        <operation name="query" performedby="S0" / >
        <switch variable="name">
            <case condition="expr2">
                <sequence>
                    <operation name="order" performedby="S0" / >
                    <operation name="query" performedby="S1" / >
                < /sequence>
            </case>
            <case condition="expr3">
                <sequence>
                    <operation name="query" performedby="S0" / >
                    <operation name="checkout" performedby="S1" / >
                < /sequence>
            </case>
        < /switch>
        <operation name="order" performedby="S0" / >
    < /sequence>
< /thread>
< /fork>
< /activity>

```

The above activity contains a *fork* construct. The threads in the *fork* construct are sequentialized. Based on the algorithm, we can build the execution graph as figure 2.

From figure 2, we can find the following possible execution paths:

- Path 1: *S1:register, S0:login, S0:checkout, S0:query, S0:order, S0:query, S0:order*. Path 1 is not valid. The operation *S0:checkout*, is illegal since seller *S0* requires the buyer to *order* something before *checkout*.
- Path 2: *S1:register, S0:login, S1:query, S0:order, S0:checkout, S0:query, S0:order, S0:query, S0:order*. Path 2 is legal for both service provider *S0* and *S1*.
- Path 3: *S1:register, S0:login, S0:checkout, S0: query, S0: query, S1: checkout, S0:order*. Path 3 is illegal for the same reason as path 1.
- Path 4: *S1:register,S0:login, S1:query, S0:order, S0:checkout, S0: query, S0: query, S1: checkout, S0:order*. Path 4 is illegal, since the buyer has not ordered anything from seller *S1*, before trying to *checkout*.

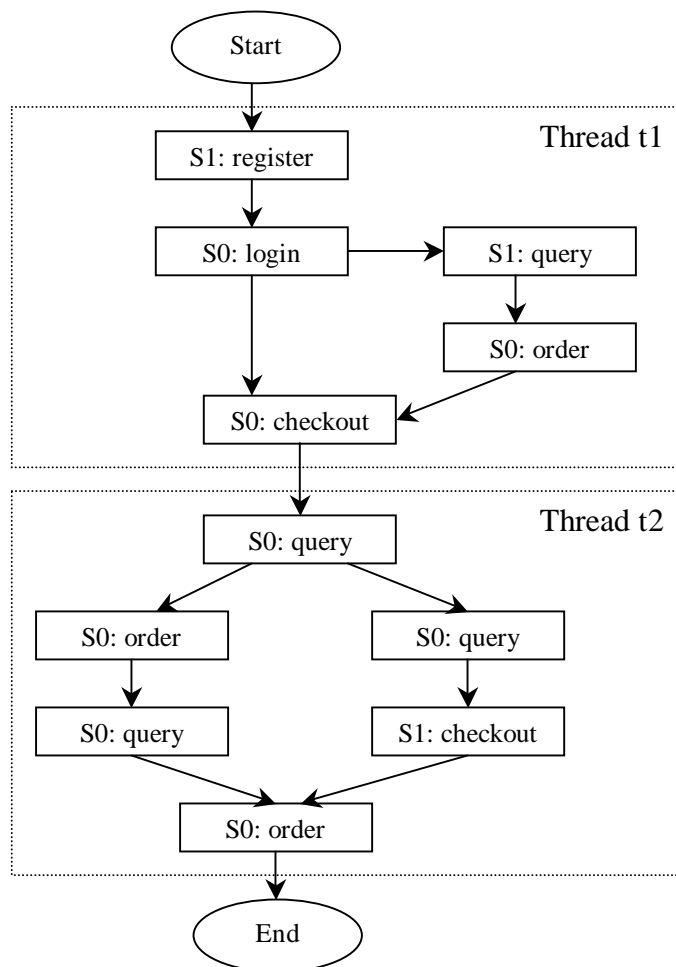


Figure 2: A Composed Service

4 Discussion

Although still in its infancy, the approach developed here seeks to facilitate the design and implementation of complex Web services as compositions of other Web services. Addressing this challenge is potentially of immense practical value. We believe it will be crucial to the expansion of semantic Web services into mainstream business process applications.

Our work raises Web services to the level of an autonomous agent, which can make independent decisions. ASDL exposes some behavior to the outside world. Our verification algorithms check the validity of a composed service, thereby detecting potential problems during the design phase of a composed service. We have developed a prototype tool to automatically check the validity of services.

4.1 Literature

Besides Web services, the work described here touches upon extensive bodies of research on the semantic Web, workflow modeling, protocols, and agent-based techniques. We lack the space to review these in detail

(some were cited in the above presentation), but mention representative work here.

- Semantic Web. DARPA Agent Markup Language (DAML) [Hendler and McGuinness, 2001] enables the creation of ontologies in the description of specific Web sites. DAML-S [Ankolekar et al., 2001] is a Web service ontology from the semantic Web community [Berners-Lee et al., 2001]. DAML-S provides a core set of markup language constructs for describing the properties and capabilities of Web services. Some emerging approaches add structure to service descriptions through the use of ontologies, e.g., [Trastour et al., 2001].
- Workflow and process modeling. Klein & Bernstein develop a richer approach for describing and indexing services based on process models [2001]. Verharen develops a contract specification language, CoLa, to specify transactions and contracts [1997]. Verharen's approach captures obligations involving actions, but does not allow the obligations to be manipulated dynamically. This is a possible line of extension for the present work.
- Protocols. Barbuceanu and Fox [1995] develop a language, COOL, for describing coordination among agents. Their approach is based on modeling conversations through FSMs, where the states denote the possible states a conversation can be in, and the transitions represent the flow of the conversation through message exchange. They try to handle exceptions through error recovery rules. HP's Conversation Definition Language [Govindarajan et al.] has similar goals to ASDL. CDL provides an XML schema for defining valid sequences of documents exchanged between Web services. Like ASDL, it uses conversations to model the externally visible interaction model of the Web service.
- Agent-based exception handling. Klein & Dellarocas exploit a knowledge base of generic exception detection, diagnosis, and resolution expertise [1999]. Specialized agents are dedicated to exception handling. Our approach is complementary, since it applies at design time and does not require extensive intelligence.

4.2 Directions

On the practical side, we are working on our prototype to enhance its representational capabilities for services. On the theoretical side, it will be helpful to explicitly incorporate extended transactions in our models to capture richer constraints on service behavior.

References

- Anupriya Ankolekar, Mark Burstein, Jerry R. Hobbs, Ora Lassila, David L. Martin, Sheila A. McIlraith, Srini Narayanan, Massimo Paolucci, Terry Payne, Katia Sycara, and Honglei Zeng. DAML-S: Semantic markup for web services. In *Proceedings of the International Semantic Web Working Symposium (SWWS)*, July 2001.
- Mihai Barbuceanu and Mark S. Fox. COOL: A language for describing coordination in multi agent systems. In *Proceedings of the International Conference on Multiagent Systems*, pages 17–24, 1995.
- Tim Berners-Lee, James Hendler, and Ora Lassila. The semantic web. *Scientific American*, 284(5):34–43, 2001.

- Don Box, David Ehnebuske, Gopal Kakivaya, Andrew Layman, Noah Mendelsohn, Henrik Frystyk Nielsen, Satish Thatte, and Dave Winer. Simple object access protocol (SOAP) 1.1, 2000. www.w3.org/TR/SOAP.
- Erik Christensen, Francisco Curbera, Greg Meredith, and Sanjiva Weerawarana. Web services description language (WSDL) 1.1, 2001. www.w3.org/TR/wsdl.
- Kannan Govindarajan, Alan Karp, Harumi Kuno, Dorothea Beringer, and Arindam Banerji. Conversation definitions: Defining interfaces of web services. <http://www.w3.org/2001/03/WSWS-pop/paper20>.
- James Hendler and Deborah L. McGuinness. DARPA agent markup language. *IEEE Intelligent Systems*, 15 (6):72–73, 2001.
- Michael N. Huhns and Munindar P. Singh. Agents and multiagent systems: Themes, approaches, and challenges. In *Huhns and Singh [1998b]*, chapter 1, pages 1–23. 1998a.
- Michael N. Huhns and Munindar P. Singh, editors. *Readings in Agents*. Morgan Kaufmann, San Francisco, 1998b.
- Mark Klein and Abraham Bernstein. Searching for services on the semantic web using process ontologies. In *Proceedings of the International Semantic Web Working Symposium (SWWS)*, July 2001.
- Mark Klein and Chrysanthos Dellarocas. Exception handling in agent systems. In *Proceedings of the 3rd International Conference on Autonomous Agents*, pages 62–68, Seattle, 1999.
- Frank Leymann. Web services flow language. TR WSFL 1.0, IBM Software Group, May 2001.
- Satish Thatte. XLANG, Web services for business process design, 2001. www.gotdotnet.com/team/xml-wsspecs/xlang-c/default.htm.
- David Trastour, Claudio Bartolini, and Javier Gonzalez-Castillo. A semantic web approach to service description for matchmaking of services. In *Proceedings of the International Semantic Web Working Symposium (SWWS)*, July 2001.
- UDDI. UDDI technical white paper, 2000. www.uddi.org/pubs/Iru-UDDI-Technical-White-Paper.pdf.
- Egon M. Verharen. *A Language-Action Perspective on the Design of Cooperative Information Agents*. Catholic University, Tilburg, Holland, 1997.