

Service oriented architecture

Abstract

Service-oriented architectures (SOAs) provide an architectural paradigm for software development. Systems can be organized in terms of services units of software that provide functionality 'as is' to users. Functionality descriptions and other properties and quality attributes such as security or performance and usage-oriented information such as invocation protocols and locations are advertised by providers and can be looked up by potential users.

The Web Services Framework (WSF) is such an SOA [2]. The WSF provides an SOA infrastructure consisting of a description language (WSDL), an invocation protocol (SOAP), and a repository for descriptions (UDDI) based on standard Internet and Web technologies such as XML. While the first generation of the WSF has focused on the use of services 'as is', the next needs to address service composition to enable larger software systems to be assembled based on services as the basic unit [3,4]. Composition of services to processes is here the paradigm of composition. Two forms - orchestration and choreography- have recently been discussed in the WSF community as techniques for service composition and collaboration [5]. These two reflect the perspective of business

processes modelled and executed (orchestration) and of systems as interacting processes (choreography). (So far, the WSF is focused more on invocation than development. UDDI supports potential users in locating suitable services; how these services are integrated into existing software systems and how these services can be composed to larger systems is, however, not sufficiently addressed. The state-of-the-art comprises languages for orchestration and choreography, such as WS-BPEL4 or WS-CDL [5]. The basis of these languages is workflow and message exchanges, and aspects of interaction processes and patterns. Principles of component-based software development CBSD [6] are not yet integrated. We will therefore focus here on using the WSF platform as an infrastructure for service-based software systems development. The overall aim is to support service-based software development on and for the Web. Formal methods are proven to be successful for the development of safety-critical, dependable software systems. Formal models allow a higher degree of understanding of principles and mechanisms of the context, but also particular properties of the application. A formal model of orchestration and choreography and a description notation is therefore our central objectives. An important requirement arises if in particular the Web-based development of service-based systems is to be realized. The Semantic Web paradigm needs to be embraced in order to support the SOA principle of distributed development involving different organizations. The semantic Web provides a shared knowledge representation framework and platform, based on ontology at the core. Ontology can capture properties of services; they can also support composition description and reasoning. Our objective here is:

- We clarify the notions of orchestration and choreography. To this end, we will provide a formal model for service process composition in Section 3.

- We provide an ontological framework for service process composition that supports the CBSD objective of reuse in Section 4. This framework will be based on the semantically definition of service process composition.

Development of Composite Service Processes

Web-based Service Development

A service is made up of a coherent set of operations provided at a certain location. The service provider makes an abstract service interface description available that can be used by potential service users to locate and invoke the service. Services are often used 'as is' in single request response interactions, but more and more the composition of services to processes is important. Reuse is a central software engineering principle. Existing services can be reused to form business or workflow processes. The principle of architectural composition here is process

assembly. The discovery and invocation infrastructure - a registry or marketplace where potential users can search for services and an invocation protocol- with the services and their clients form a service-oriented architecture. Languages for description and composition and protocols are central elements of this architecture. Fig. 1 illustrates this infrastructure for the WSF. Software development for service-oriented architectures is a two-step process. Discovery is based on abstract computation descriptions (and other software properties),(formalized based on ontology. Assembly is about composition of services to processes. Ontologies to represent knowledge about services are essential for the Web as a development platform. Usage complements a basic service lifecycle. It is about communication and process interactions between services.

Service Process Composition - Orchestration and Choreography

The WSF provides a platform to invoke services on a 'usage as is'-basis. Real value, however, will be added if services can be connected [5]. Supporting and implementing business processes within the WSF through composed services is the requirement. Orchestration and choreography are two forms of service composition and collaboration that are currently discussed.

- Orchestration refers to a composed business process that may use both internal and external Web services to fulfill its task. The business process is controlled by one of the agents in the system. The process is described at the message level, i.e. in terms of message exchanges and execution order.

- Choreography addresses the interactions that implement the collaboration between services. Multiple agents are considered where each agent describes its own part in the interaction. Orchestration and choreography address different perspectives. Orchestration is focused on the internal behavior of a business process. Choreography is focused on the external perspective, looking at process interaction. These perspectives are the essential aspects of an SOA. We will capture these in a process model (orchestration) and an interaction model (choreography).

A Bank Account Example

An online banking example shall illustrate our service process framework. Login; !(Balance+Lodgement+ Transfer); Logout is a process expression describing an interaction process of an online banking user starting with a login, then repeatedly executing balance enquiries, lodgements or money transfer, before login out. In Fig. 2, four services are described in a pseudo-codenotation. Each of these services implements a process internally (orchestration).(The interactions resulting from the service invocations (import, overlined,e.g. Balance) and service provision (export, normal, e.g. Login) are the result of service choreography. For instance, Account Process is a client of Bank Account and Login Server; Bank Account is a client of Account Registry.

"Services and Processes - a Formal Operational Model

Description and composition are central design activities. In this section, we develop a formal model and an abstract language that form an operational framework for both activities. We formalize orchestration and choreography and develop a semantical framework that defines and supports composition activities. This operational semantically framework serves to capture requirements and form an underlying layer for the ontological framework.

Orchestration and Choreography Description Orchestration.

We can derive the following core requirements for an orchestration notation from languages such as WS-BPEL.

-basic elements: message-based actions in two form - invocations for external services and receive/reply actions if the service is available to others,

•process language: sequence, choice, iteration, and concurrency are the service/process combinators,

•abstraction and export interface: a process can be provided as aWeb service,

•state and data: variables and parameters to actions are needed. The focus of orchestration is illustrated in Fig. 3. The business process itself and the Web services that implement the process are separated. This keeps the process logic apart from its implementation. The process is executed byan orchestration engine which invokes the respective services.

We capture the foundations of orchestration in form of a process model focusing on service composition. A process description is about control flow and the determination of the execution order. We will start with abstractactions to concentrate on control flow first - data aspects and also interactions will be added later. A process description can serve different purposes:

-to define a business process in terms of actions and control flow,

-to describe the external, observable interaction pattern that a service can engage in a composed system (if the process is made available as a service.(Our process model is based on the principles and notation of then-calculus. Service process expressions, or processes are inductively formed based on a basic process names, named process expressions, and the combinatory sequence

, ;parallel composition I , non-deterministic choice+, and iteration

. !A named process expression $P(s_1, \dots, s_k)$ is defined by a service process expression on based services s_1, \dots, s_k and the combinators. The process definition is recursive. Based on basic processes (which are Web services,(composite services can be defined, i.e. expressions such as $P = s_1; s_2; Q$ can be used. We also use the notation $P-s_1--7;s_2 Q$ to emphasise the transitional character of processes. Example 3.1 Login; !(Balance+Lodgement+Transfer); Logout is a business process for an online bank account. This orchestration example ignores the import/export classification of process elements necessary for choreography. We now adds data by refining the notion of actions. For service s and data item x , $s(x)$ is the receive action,

$s \xrightarrow{\quad\quad\quad} x$

is the reply action, andlet $y =$

$s \quad\quad\quad x$

in P is the invoke action. Receive and reply actions areneeded to faciliate service provision. Invoke is

needed to use other services. The invocation provides a scope for the returned result y of the interaction.

Choreography.

Similar to our orchestration discussion, we note the main requirements for choreography description notation [5:]:

- basic activities: request/response action for local activities, invoke to call external services,
- structured activities: loop, sequence, choice, and concurrency,
- infrastructure: channels/connections between ports that represent services.

The focus of choreography is observable interaction behavior, not execution, see Fig. 3. The orchestration model is a process model with its focus on controlflow and execution order. The choreography model is an interaction model about process interaction, i.e. synchronization and data exchange. Essential in modeling process interaction is to add data flow between processes. Web services are connected through a network. The network end points that represent services are called ports - service names will act as port names. Services (and their ports) can be receivers and senders of data, i.e. read from or write to communication channels set up between the ports. Assume a service port s and a data item x . Then, $s(x)$ is the receive action and

$s \xrightarrow{\quad\quad\quad} x$

is the send action. Note, that in contrast to orchestration, we have abstracted here from the difference between provider actions (receive/reply) and client actions (invoke). The expression

$\text{Balance}(\text{bal})$ asks service Balance for the current balance of account ace and then receives the balance bal . An interaction is the activation of a remote service. Two forms shall be provided. Assume a process expression P .

- request-response: for each service s in P a write-read sequence

$s \xrightarrow{\quad\quad\quad} x$

$;$ $s(y)$ (where y is the returned result from an external service).

- execute-reply: for each services in P a read-write sequence $s(x)$;

$s \xrightarrow{\quad\quad\quad} \sim x$

where f is some internal service functionality. These interactions are the basic building blocks of the process life cycle. Input services names in a process expression need to be bound to a concrete service that can execute the service functionality. Finding suitable services that match each individual service requirements and managing the connections is part of the interaction model and its matching and connection support. So far, the concurrent composition of processes $A \mid B$ does not allow interactions. A transition rule (called reaction rule in then-calculus) can capture interaction and describe the data flow in these interactions - see details in below. A shared channel can be created

that forms a connection between two agents. Usually, the port names act as channel names (e.g. the π -calculus requires matching port names to establish a connection; we will loosen this constraint later on). Choreography is often about fixed connections. Process calculi, however, also cater for connections that are created dynamically. Using the π -calculus' scope extrusion, dynamic architectures can be modeled.

Composition Support

Descriptions are needed to publish services in repositories or to capture requirements for these services. We will provide a simple development and deployment model for services in form of a life cycle model, before addressing techniques needed for individual activities in that lifecycle.

Life Cycle and Activities.

Description and matching are design activities. Essential is, however, the support of the full process life cycle. Binding individual service names to existing services, i.e. composing a process instance and executing this instance are as important as description and matching. The foundations of these aspects are given in form of a choreography or interaction model that describes bindings, connections, and interactions between services. Each service is a family of ports sC , sI , sR that address the needs of the different life cycle stages. Port sC is a contract port, representing an interface that captures abstract properties. sI and sR are connector ports for interactions! sI handles service invocation and input and sR handles the service reply. We express the service life cycle in an annotated process notation

for the requestor with annotations for requesting, invoking, and result. Dual to the requestor view there is a provider view

with annotations for providing, executing and replying. In the requestor view, Req
 sC

is an annotated output action of service

s . A process can request Req a service using contract port sC . Connector port references sI and sR are subsequently sent for further interactions. If matching between a request or port type and a provider port type is successful, then the requestor and the provider process can be composed, i.e. requestor can interact with the provided service repeatedly. The requestor would invoke Inv the service at port sI and receive a result Res at port sR

Matching.

Matching is central in composition. An existing service that is reused and integrated, for instance into a business process, must match the requirements in order to allow the business process to fulfill its task.

- Import process patterns describe how a process expects to use other services.
- Export process patterns describe how provided services have to be used.

These are elements of an orchestrated business process. Orchestration elements are more relevant to matching than choreography aspects such as interaction, which is more deployment-oriented.

The specification of processes describes the ordering of observable activities. We use a simulation notion to define process matching. The requested process is the import process pattern that the client expects the provider to support. A provider process P matches (or simulates) a requested process R if there exists a binary relation S over the set of processes such that if whenever RSP and $R \rightarrow S R$ then there exists P such that $P \rightarrow S RSP$.

This definition originates from the simulation definition of the π -calculus [7]. [The provider needs to be able to simulate a request, i.e. needs to meet the request pattern of the client. Dynamic binding of concrete services to the process names is possible. This definition is about potential interaction.

Example 3.2 A provided service process $!(\text{Balance} + \text{Lodgement} + \text{Transfer})$ matches the expected support of process $!(\text{Bal} + \text{Ldg})$. If the pairs $\text{Balance}/\text{Bal}$ and $\text{Lodgement}/\text{Ldg}$ match (e.g. equal signatures), then the provider matches (or simulates) the requested process.

Connection and Interaction.

Composition consists of two activities: matching and connection. Successful matching can result in a connection between service ports. From the perspective of a business process, concrete services are connected to the abstract business process elements.

So far, we have been looking at matching of abstract process descriptions. We now focus on the computational side of compositions. The connection of matching services shall now be formalized using an operational semantics. In the composition process we can distinguish a contract phase where both process instances try to form a contract based on abstract descriptions. The connection phase establishes a connector channel for interaction between these services. We will capture contract and connector establishment in form of transition rules. This formalises the connection of provider and client in the WSF - a virtual link between URIs that are used by the SOAP protocol.

For a parallel composition

$C \mid nC(nl).P$ of a client business process element and a provider service, both processes commit themselves to a communication along a (virtual) channel between ports m and nC . A contract rule formalizes the process of matching and commitment 2 :

Services and Processes - an Ontological Framework

Supporting service development is ideally supported through ontology technology for shared representation of knowledge - here service descriptions. We illustrate how description and composition of services processes can be represented in a description logic that underlies a Web ontology language.

Ontologies for Web Services and Processes

The formal model (see Section 3) goes beyond what we need for the ontological framework to support the development of service-based software systems. Ontologies are needed to support composition through matching of patterns and processes, i.e. port orientation and other interaction and choreography aspects are not relevant since they address the deployment infrastructure. The ontological framework therefore abstracts the underlying formal operational model, which defines the development and deployment infrastructure. We will develop the ontological framework in terms of a description logic [8]. [Description logic as the underlying logic of the Semantic Web is particularly interesting for the software engineering context due to a correspondence between description logic and dynamic logic (a modal logic of programs) [9]. [This correspondence is based on a similarity between quantified constructors) expressing quantified relations between concepts) and modal constructors(expressing safety and liveness properties of programs).

A Basic Process Ontology

Ontologies are formal frameworks that provide knowledge description and reasoning techniques. The starting point in defining ontology is to decide what the basic ontology elements (concepts and roles) represent. Here, the ontology shall formalize process-based, i.e. state-transition based software systems.

- Concepts are classes of objects with the same properties. Individuals are named objects. Concepts represent software system properties in this context. Systems are dynamic. Descriptions of properties are inherently based on underlying notions of state and state change.

- Roles in general are relations between concepts. Here, they shall represent two different kinds of relations. Transitional roles represent service operations in form of accessibility relations on states, i.e. they represent services resulting in state changes. Descriptive roles represent properties of a state such as invariant descriptions like service name and description or pre and post conditions (if they are part of the description format.)

- Constructors allow more complex concepts to be constructed in form of concept descriptions. Classical constructors include conjunction and negation \sim . Hybrid constructors are based on a concept and a role. The constructor $\forall R.C$ - called value restriction - is interpreted based on either an accessibility relation R to a new state C for transitional roles, or on a property R satisfying a constraint C for descriptive roles. The dual $\exists R.C$ is called existential quantification.

In Fig. 4, the service process ontology is shown. A state is an abstract concept that is described in terms of elements of auxiliary domains through descriptive roles such as invariant and mutable state properties (formal conditions, textual descriptions, etc.). The two essential state concepts are pre and post, which denote abstract pre- and post-states for service process transitions (not to be confused with pre- and post conditions). For example,

$\forall \text{outSign.int}$ specifies a post-state by associating an output signature int . Throughout this paper, we use a description logic notation, e.g. for a given concept pre , we could constrain input signatures, $\forall \text{inSign}(\text{int}, \text{int})$. This notation is equivalent to a triple expression $(\text{pre}, \text{inSign}, (\text{int}, \text{int}))$, which would be used in RDF (on which OWL is based).

Orchestration and Choreography

We introduced the representation of basic services in a description logic-based ontology. An ontology that captures service processes and their composition, however, requires an extension of classical description logics [8]. So far, roles -that represent service operations - are atomic. We define the combinators

'|', '!', ';and+' as role constructors for sequential composition, transitive closure (iteration), intersection (parallel composition without interaction),(and union (non-deterministic choice) of service processes, respectively. We also use ||| for sequential composition to emphasize the functional character of roles. Role constructors allow us to integrate process description and composition into

an ontology framework. The description logic expression $\text{Balance} + \text{Lodgement} + \text{Transfer}$.post describes a process. Axioms and inference rules allow us to capture activity-related properties in the logic, e.g. in order to reason about matching. For example, $\forall R. \exists S. C \phi :: \}$

$\forall R; S.C$ is an axiom that describes the conversion between logical operators and role expression combinators. We need to integrate data and process parameters into the logic. We

introduce data in form of names. Names stand for individual data elements.

- We denote a name n by a role nN , interpreted by an identity relation

$\{ \{ nI, nI \} \}$ for the interpretation nI of n .

- An operation R is a parameterised role $R I \sim D \times S \times S$ for domain D of a name and states S .

- A parameterized role R applied to a name nN , represented here as an identity relation, i.e. $R ||| nN$, forms a transitional role, i.e. $R ||| nN \sim S \times S$.

1°. 1° Com position Support Matching.

Subsumption is the central inference technique in description logic. The subsumption $C1$ of concepts is the subset-relationship of the corresponding object classes. Equally, we define subsumption for roles $R1 R2$. We defineservice process matching in the expected way. A process $P(nI, \dots, nk)$ matches a process $R(mI, \dots, ml)$, if $P(nI, \dots, nk)$ simulates $R(mI, \dots, ml)$. Subsumption on roles, however, is input/output-oriented, whereas the simulation needs to consider internal states of the composite role execution. For each request in a process, there needs to be a corresponding provided service. Although not the same, matching is a sufficient condition for subsumption. If the process expression $P(nI, \dots, nk)$ simulates the process $R(mI, \dots, ml)$, then R

P. Matching can be ontologically supported by constructive axioms [10]. [Connection and Interaction.

We have formulated the operational semantics of interaction in form of process calculus-style contract and connector rules. In terms of the ontology, services were so far described as transitional

roles and we considered system states that describe service (and process) properties such as pre and post-states to define transitional process behavior. We formalize composition and interaction in the ontology framework through inference rules. In order to address interaction, we need to look at a special kind of a parallel composition transition. This transition is based on the synchronization of concurrent services through data exchange. We can characterize properties of interactions between two services, here a reformulation of the contract rule without annotations and matching constraints,

This rule for parallel composition complements other constructor-specific axioms and rules that we can derive from dynamic logic and process calculi such as the axiom $Vp: q.C \parallel Vp: r.C$ for the sequence. These axioms and inference rules form an application-specific extension of description logic that allow us to infer more properties about service processes and their interactions.

Related Work

Composition of services is an active area of research [3,4,11]. In particular the need to address semantics in the context of composition has been recognized. In [3], an ontological framework for service composition is presented based on OWL-S (a rich services ontology, formerly known as DAML-S) as the underlying service ontology [13]. Their application area is the Grid services context and knowledge-based advice systems. OWL-S is different from our ontological

framework in its process model. OWL-S represents services as concepts in the ontology, not as transitions. Therefore, the bridge to dynamic logics cannot be exploited in the way we proposed. Another OWL-S based approach is taken in [4]. Here, the logical side is strengthened. OWL-S descriptions are converted in linear logic and architecture based on a logic-based planner and a semantic reasoner are proposed. The ultimate aim, as in our description-logic based approach, is the exploitation of logic reasoning for service composition. Our approach differs from the discussed OWL-S-based approaches and other service ontologies such as WSMO [12] in that the ontological model captures services and processes in a more intrinsic way. OWL-S and WSMO address a wider range of properties, which suggests an integration of these approaches with our composition framework. We have aimed at reflecting the current discussion on orchestration and choreography in our technical and ontological models here.

Semantic Web services are a subject that our approach needs to be related to. OWL-S [13] and WSMO [12] are examples of ontological frameworks that support matching of semantically described services. Both focus on the semantically description of services including abstract descriptions, quality of-service aspects, and functional abstractions such as pre- and post conditions. We can use pre- and post conditions as abstractions for ports, enabling the design-by-contract approach [14]. Dynamic logic is a suitable logical framework that subsumes pre- and post condition specification [15]. The connection between description logic and dynamic logic allows us to integrate these contracts easily into our framework. Similar to signatures, we can associate

(descriptive) pre- and post condition roles to pre- and post states, respectively. Two service operations described by pre- and post conditions and represented by contract ports nC and mC match, if the requestor's precondition is weakened and the post condition strengthened [15,16].

Again, we would need to integrate reasoning about services matching with subsumption. Subsumption is interpreted as a subset relationship on sets of states that satisfy pre- or post-state descriptions. We present a matching inference rule for transitional roles. We define the matching rule

Conclusions

Organizing software systems as service-oriented architectures is a new architectural paradigm. The Web Services Framework is currently the most important platform supporting this paradigm. In order to make the paradigm more successful as a software development approach, the focus on deployment, invocation, and reply has to shift towards more re-use and composition. Service reuse is one of our objectives; process assembly of reusable services is the principle of architectural composition in this context. What is needed is a component-style composition framework for services - in particular an ontological framework to make it work as a development approach for distributed and shared platforms such as the Web. We have developed a formal framework based on ontologies (and underlying logics) and process-based service composition, applicable to the Semantic Web and Web Services platforms. We have looked at core aspects of such a framework here - a complete and

formal treatment was beyond the scope of this paper. One of the lessons we have learned is that a comprehensive formal framework for service-oriented architectures is needed to address predictable assembly, reuse, maintenance, and change and evolution management. We consider our contribution part of a methodology for service-based software development and deployment.