
A survey on web services composition

Schahram Dustdar* and Wolfgang Schreiner

Distributed Systems Group,
Vienna University of Technology, TU Wien

E-mail: dustdar@infosys.tuwien.ac.at

E-mail: schreiner@infosys.tuwien.ac.at

*Corresponding author

Abstract: Due to the web services' heterogeneous nature, which stems from the definition of several XML-based standards to overcome platform and language dependence, web services have become an emerging and promising technology to design and build complex inter-enterprise business applications out of single web-based software components. To establish the existence of a global component market, in order to enforce extensive software reuse, service composition experienced increasing interest in doing a lot of research effort. This paper discusses the urgent need for service composition, the required technologies to perform service composition. It also presents several different composition strategies, based on some currently existing composition platforms and frameworks, re-presenting first implementations of state-of-the-art technologies, and gives an outlook to essential future research work.

Keywords: web services; coordination; transaction; context; conversation modelling; service execution; quality of service; service composition.

Reference to this paper should be made as follows: Dustdar, S. and Schreiner, W. (2005) 'A survey on web services composition', *Int. J. Web and Grid Services*, Vol. 1, No. 1, pp.1–30.

Biographical notes: Schahram Dustdar is a Full Professor at the Distributed Systems Group, Information Systems Institute, Vienna University of Technology (TU Wien). In 1999 he co-founded Caramba Labs Software AG (CarambaLabs.com) in Vienna, a venture capital co-funded software company focused on software for collaborative processes in teams. Caramba Labs was nominated for several (international and national) awards. He has published some 70 scientific papers as conference-, journal-, and book contributions. He has written three academic books, one professional book, and co-edited six books/proceedings. More information can be found at: <http://www.infosys.tuwien.ac.at/Staff/sd>

Wolfgang Schreiner is a PhD student at the Information Systems Institute, Vienna University of Technology, TU Wien. His research interests include web services workflows and web services evolution.


1 Introduction

A web service is a software system identified by a URL, whose public interfaces and bindings are defined and described using XML. Its definition can be discovered by other software systems. These systems may then interact with the web service in a manner

prescribed by its definition, using XML-based messages conveyed by internet protocols. This definition has been published by the world wide web consortium W3C, in the Web Services Architecture document (Booth et al., 2004).

The web service model consists of three entities, the service provider, the service registry and the service consumer. Other models, such as a peer-to-peer structure, exist as will be discussed later in this paper. Figure 1 shows a graphical representation of the traditional web service model:

Figure 1 The web service model



The service provider creates or simply offers the web service. The service provider needs to describe the web service in a standard format, which in turn is XML and publish it in a central Service Registry. The service registry contains additional information about the service provider, such as address and contact of the providing company, and technical details about the service. The Service Consumer retrieves the information from the registry and uses the service description obtained to bind to and invoke the web service. The appropriate methods are depicted in Figure 1 by the keywords 'publish', 'bind' and 'find'. In order to achieve communication among applications running on different platforms and written in different programming languages, standards are needed for each of these operations.

Web services architecture is loosely coupled, service oriented. The Web Service Description Language WSDL uses the XML format to describe the methods provided by a web service, including input and output parameters, data types and the transport protocol, which is typically HTTP, to be used. The Universal Description Discovery and Integration standard UDDI suggests means to publish details about a service provider, the services that are stored and the opportunity for service consumers to find service providers and web service details. Besides UDDI, other standards have been developed as well. Dustdar and Treiber (2004) deals with web service registries in greater detail. The Simple Object Access Protocol SOAP is used for XML formatted information exchange among the entities involved in the web service model.

This paper is structured as follows: In Section 2 the need for service composition and what advantages we expect from doing research in this area is discussed. Section 3 gives a brief overview of our sample scenario that has been developed at our institute and which will be used to describe the different composition strategies that have been identified, based on the example of managing a film crew. Section 4 contains the main focus of this survey paper. Different approaches in service composition by discussing

composition platforms and frameworks that have been developed in the course of comprehensive research in the area of service composition. In Section 5 common characteristics and features of the frameworks, by which we can compare them and provide a table as a visual mean to structure our summary and conclusion are identified.

2 The need for web services composition

The basic web services infrastructure presented in the previous sub-chapter suffices to implement simple interactions between a client and a web service. If the implementation of a web service's business logic involves the invocation of other web services, it is necessary to combine the functionality of several web services. In this case, we speak of a composite service (Alonso et al., 2004). The process of developing a composite service in turn is called service composition. Service composition can be either performed by composing elementary or composite services. Composite services in turn are recursively defined as an aggregation (Khalaf and Leymann, 2003) of elementary and composite services.

When composing web services, the business logic of the client is implemented by several services. This is analogous to workflow management, where the application logic is realised by composing autonomous applications.

This allows the definition of increasingly complex applications by progressively aggregating components at higher levels of abstraction. A client invoking a composite service can itself be exposed as a web service.

Since it is a widely used approach to use conventional programming languages to link components to a composite web service, and thereby bridge heterogeneous middleware platforms, it becomes necessary to develop a Service Composition Middleware to support composition in terms of abstractions and infrastructure as well (Alonso et al., 2004). Programming languages focus on APIs rather than on the actual business logic. Different approaches and the need for workflow modelling have finally led to the development of the business process execution language for web services BPEL4WS (Doulkeridis et al., 2003). A composition model and language to specify the services involved in the composition, a development environment with a graphical user interface to drag and drop web service components and a run-time environment to execute the business logic can be identified as the three main elements of a web services composition middleware. A service composition middleware requires the web services to be precisely described in their functionality, interfaces and protocols they support. Conventional middleware lacks exactly those features. Workflow Management Systems (WfMS) are highly flexible and generic but on the other hand require the components to be aware of the WfMS API. Therefore, components are system and vendor-specific and attended to additional development effort.

Alonso et al. (2004) describe six different dimensions of service composition models. A component model can make different assumptions of what a component is and what it is not. The advantage of a model making very basic assumptions, for example components only have to exchange messages via XML, is a more general model, while it has to deal with much more heterogeneity of the components.

An orchestration model defines abstractions and languages to define the order in which and the conditions under which web services are invoked. Orchestration models use process-modelling languages, such as UML activity diagrams, Petri-nets, state-charts,

rule-based orchestration, activity hierarchies, and π -calculus. Data and data access models define how data is specified and exchanged between components. The service selection model deals with static and dynamic binding that is how a web service is selected as a component, statically at design-time or dynamically during run-time. Transactions define which transactional semantics can be associated to the composition and how this is done. Finally, we must also consider a model for exception handling to handle exceptional states during the execution of the composite service without the service being aborted.

The Business Process Execution Language for Web Services (BPEL4WS) standard proposed by IBM (Virdell, 2003), Microsoft and BEA (Andrews et al., 2003) is in many cases related to former standards, such as WSFL or XLANG. BPEL descriptions are XML documents, which describe the roles involved in the message exchange, supported port types and orchestration, and correlation information as aspects of a process. BPEL4WS is a service composition model (Wohed et al., 2003), which supports both, composition and coordination protocols. It also consists of an activity-based component model, an orchestration model that allows the definition of structured activities, XML schema data types, a service selection model and a mechanism for exception, event and compensation handling.

3 Managing a film crew: a case study

According to Dustdar and Treiber (2004) the case study of managing a film crew has been selected, in order to explain the different concepts and frameworks introduced in various papers, with a consistent reference scenario that runs through the whole paper.

Basically, a film crew consists of one or more film directors, external experts and crew-members. Each of these teams provides particular services and work together in a loosely coupled way, providing their expertise on demand, while they depend on the services provided by other film teams.

A film director must coordinate the film teams, is responsible for the budget, and must enable communication among different teams. Since different phases in the film production place different requirements on the arrangement of the teams, each phase requires flexible composition and configuration of the film teams. Directors hire film crews and external experts who provide own services and capabilities necessary for shooting a film. External experts offer their expertise on several topics, such as physics, health, or law to make the film and special effects more realistic. Crew-members may act as stuntmen for example.

4 Composition issues

Web service composition (Benatallah et al., 2001, 2002) is a very complex and challenging task, as will be seen when discussing different composition approaches in greater detail. Before performing web service composition, some basics to enable service composition have to be performed. Six different issues that have a large impact on service composition have been identified. The following sub-sections do not claim completeness of composition issues. The most important issues may be explained in a few words. The concepts, some standardisation efforts and related research work as well

have been introduced. Sun Microsystems, for example, has proposed standards for coordination (Terai et al., 2003), transaction, and context and has put them together to one comprising standard called WS-CAF (Web Services Composite Application Framework) (Bunting et al, 2003b).

Coordination, transaction, context and conversation modelling is discussed. When composing services composite service execution must be considered. Centralised and distributed service execution has been differentiated, and frameworks that fall into these categories have also been discussed. In addition to the ‘traditional’ SOA (Service Oriented Architecture) mentioned in the introduction-based infrastructure, there are also alternatives to that. Here, peer-to-peer infrastructure is dealt with and a P2P based framework called Peer-Serv is described, and suggestions to extend the classical web service model in order to precisely describe web services and facilitate quality driven service discovery (Benatallah et al., 2004a) are presented

4.1 Coordination

The basic web service infrastructure presented in Section 1 satisfies simple interactions and method invocations among interacting entities. When it comes to composing web services and building complex software systems, it is likely that interaction requires coordination of sequences of operations, to ensure correctness and consistency. New protocols and abstractions are needed and this is exactly the case of coordination. In order to provide modelling abstractions and simplify web service development, different standardisation efforts, such as WS-Coordination (Cabrera et al., 2002) by IBM or WS-CF by Sun (Bunting et al, 2003a) have been taken.

4.2 Transaction

To add guarantee to the interactions, it is of considerable importance to add a transaction (Papazoglou, 2003) protocol to the coordination framework in order to provide short-duration transactions, called atomic transactions, and long running business activities as well. In case of long running activities, it is not always possible to ensure ACID (atomicity, consistency, integrity, and durability) properties of web service transactions. Very important to name are the WS-Transaction (Cabrera et al., 2001; Freund and Storey, 2002a, 2002b) standard, proposed by IBM and WS-TXM (Bunting et al, 2003c) suggested by Sun. WS-Transaction builds upon the WS-Coordination framework mentioned above, since it defines protocols for centralised and peer-to-peer transactions, both of which require the existence of coordination.

4.3 Context

The term context is a very vague one, since many different definitions can be found in literature. In terms of web services, context maybe inferred as information utilised by the web service to adjust execution and output to provide the client with a customised and personalised behaviour (Keidl and Kemper, 2004a; Álvarez et al., 2003). Context is extensible with new types of information at any time, without any changes to the underlying infrastructure. Context may contain information such as a consumer’s name, address, and current location, the type of client device, including hard- and software that the consumer is using, or all kinds of preferences regarding the communication.

The WS-Context (Bunting et al., 2003d) standard, which is part of the WS-CAF standard proposed by Sun, specifies context, context sharing, and context management.

4.4 *Conversation modelling*

Benatallah et al. (2004b) introduce a framework that builds on currently existing standards to support developers in defining service models and richer web service abstractions. The authors identify several abstractions, which can be classified as completion or activation abstractions. Completion abstractions contain compensating operations in case an effect of a so-called forward operation has to be cancelled and resource-locking operations that lock resources for a client. Activation abstractions describe implicit and explicit transitions between states. An important type of implicit transitions is timed transitions, which occur automatically after some time. To model transitions, multiple properties are used: activation properties describe a transition's triggering features, transaction properties specify a transition's effect on the client state, and locking properties reserve certain resources for the requester for a given time. The conversation model facilitates service discovery and dynamic binding, service composition model validation, service composition skeleton generation, analysis of compositions and conversations and conversation model generation.

In Bultan et al. (2003), the authors propose a specification for designing and analysing service composition. Peers, modelled as Mealy machines, communicate through asynchronous messaging and maintain a queue for each incoming messages.

Related work in web service conversation modelling include the WSCL (Web Services Conversation Language) specification, the WSCI (Web Service Choreography Interface) specification, and also WS-Coordination and WS-Transaction, although they do not provide means to model business conversations, but rather propose specific conversations that can be used to coordinate interacting parties and provide middleware properties (Benatallah et al., 2004a).

4.5 *Execution monitoring*

We distinguish between centralised and distributed execution of composite web services. Centralised execution is similar to the client-server paradigm. In this case, the server is the central scheduler that controls the execution of the components of the composite web service. The e-flow platform developed by HP, described in the previous section, works with a centralised-scheduler (Sun et al., 2003).

The distributed paradigm in contrast expects the participating web services to share their execution context. Each of the hosts running a web service has its own coordinator, which has to collaborate with the coordinators of the other hosts, to guarantee a correct ordered execution of the services. SELF-SERV (Sheng et al., 2002) developed by the University of New South Wales uses such a distributed execution system. A hybrid form of the distributed and centralised paradigms may be coordinator that controls not only one but a set of web services.

4.6 *Infrastructure*

Ran (2003) suggests another model for discovering web services based on QoS constraints by extending the basic web services model discussed in Section 1. 48% of the

UDDI entries are unusable. Pointers are missing, broken or contain inaccurate information. Furthermore, service discovery in UDDI is limited to functional requirements only. The approach taken in Ran (2003) does not touch the existing web service model but extends it by adding another entity called the ‘QoS certifier’. The process of publishing, finding, and binding to a web service remains the same. But it adds to the role of the QoS certifier to verify the service provider’s QoS claims before the registration of a service in the UDDI registry can take place. In the new model, a service provider also has to provide QoS information, and not just the functional aspects. The certifier then may approve the claim of the provider, or may down grade it.

When a service consumer looks for a service, it may add some QoS constraints to the functional requirements, and thus, enforce a much finer search. Since there may be many services that meet the functional requirements of the user, the QoS constraints may help to filter the proper services. If no service is found, the user may also consider trade-offs by reducing the QoS constraints.

In order to realise the proposed extension of the web service model, it is necessary to extend the UDDI data structure by an additional element called ‘qualityInformation’. The structure provides the description of different QoS aspects, such as availability, reliability or performance. Like the ‘bindingTemplate’ element it refers to the ‘tModel’ defined in the UDDI registry as references to QoS taxonomies. If we now consider a simple SOAP request to a UDDI registry, then it may now, related to the new model, contain QoS parameters to find the proper services that meet these requirements:

```
<SOAP-ENV:Envelope>
  <SOAP-ENV:Body>
    <find service>
      ...
      <qualityInformation>
        <availability>0.9</availability>
      </qualityInformation>
    </find service>
  </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

The registry may respond with service references matching the search criteria:

```
<SOAP-ENV:Envelope>
<SOAP-ENV:Body>
  <serviceList>
    <serviceInfos>
      <serviceInfo serviceKey="...">
        <qualityInformation>
          <availability>0.99</availability>
        </qualityInformation>
      </serviceInfo>
      <serviceInfo serviceKey="...">
        <qualityInformation>
          <availability>0.91</availability>
        </qualityInformation>
      </serviceInfo>
    </serviceInfos>
  </serviceList>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

In Ran (2003) the author organises different QoS aspects in runtime related QoS, such as scalability, capacity, performance (response time, latency, throughput), reliability, availability, flexibility, exception handling, and accuracy; transaction support related QoS, like regulatory, supported standard, stability, cost, and completeness; and security related QoS, for example authentication, authorisation, confidentiality, trace ability and audit ability, data encryption and non-repudiation.

Apart from conventional web service model described in Section 1, the idea of a P2P based web service infrastructure is very common in literature.

P2P systems bring more availability, scalability, and extensibility (Wang et al., 2003). Peer-Serv is a framework that is composed of several providers, brokers and requestors to support web services in a P2P environment (Wang et al., 2003). Service providers and service requestors have the same role as the service providers and the service consumers described in the web service model. Service brokers are responsible for maintaining the status of the peers. A service broker maintains providers and requestors as well to form a community. Several communities may then interact with each other over their brokers. Peer-Serv uses P2P technology for service execution as well as for service publishing and service querying.

5 Composition approaches

Five categories of composition strategies have identified here. Static and dynamic composition strategies concern the time when web services are composed. They are equivalent to design-time and run-time composition. Model driven service composition, business rule driven service composition, and declarative composition have been dealt with. Another topic concerns automated (Narayanan and McIlraith, 2002) and manual service composition. A survey of ontology-based composition and semantic web services composition (Rao and Su, 2004; Agarwal et al., 2003) as counter-part to manual composition techniques such as provided by BPEL has been provided. This section closes with an overview of context based service discovery and composition.

5.1 Static vs. dynamic service composition

Static composition takes place during design-time when the architecture and the design of the software system are planned. The components to be used are chosen, linked together, and finally compiled and deployed. This may work fine as long as the web service environment – business partners, and service components does not, or only rarely, change. Microsoft Biztalk and Bea WebLogic are examples of static composition engines (Sun et al., 2003).

If other businesses provide newer services, or the old services are replaced by other ones, inconsistencies might be caused. In that case, it is unavoidable to change the software architecture and bind to other services or, in the worst case, even change the process definition and redesigns the system. In this case, static composition may be too restrictive and components should automatically adapt to unpredictable changes (Sun et al., 2003). e-flow from HP or the StarWSCoP, both described later in this section implement a dynamic composition engine.

The service environment is a highly flexible and dynamic environment. New services become available on a daily basis and the number of service providers is constantly growing. Ideally, service processes should be able to transparently adapt to environment changes, and to adapt to customer requirements with minimal user intervention.

HP has developed a system called e-flow for specifying, enacting and monitoring composite e-services (Casati and Shan, 2001). HP defines e-services as the means by which an enterprise offers its products, services, resources, and know-how via the internet. E-services may be used by persons (business to consumer), enterprises (business to business), and electronic devices and is thus, a broader term than e-commerce or e-business.

Composite services (Casati et al., 2000) are modelled as business processes enacted by a service process engine. E-flow provides features that support service process specification and management, including a simple service composition language, exception handling, ACID transactions and security management. E-flow enables process specifications that automatically configure at run-time, has simple process modification semantics and enforces consistency rules to avoid run-time errors, as well as authorisation rules.

E-flow runs on top of E-services Platforms (ESPs), such as HP e-speak or Sun Jini, which basically correspond to the web service model introduced in Section 1. They allow the development, deployment and secure delivery of e-services to businesses and customers. Service providers may register service descriptions of services they offer and may monitor and manage service execution. Services may be discovered by service consumers by specifying search criteria and then be invoked. Service invocation is restricted to authorised users.

In e-flow, a composite service is described as a process schema that composes basic or composite services. Composite services are modelled by a graph defining the flow of service invocations similar to UML activity diagrams, which is also translated into a XML structure. E-flow additionally allows defining low-level graphs to specify the order in which certain methods of one specific service are invoked. It also includes “transactional regions” to ensure the atomic execution of parts of the process.

The e-flow engine notifies completions of method nodes, whole service nodes and then determines the next service node to be executed. It is connected to the ESP via the e-flow adapter to get notified whenever the service environment changes, that is, when service definitions change or new providers and services are registered within the ESP. In order to manage and even take advantage of the frequent changes in the service environment, the services need to be adaptive – to adjust to changes in the environment with minimal user intervention. E-flow implements several features, such as dynamic service discovery, dynamic conversation selection, multi-service nodes, and generic nodes.

E-flow specifies a service selection rule, which analyses user requirements to create several input parameters, in order to discover appropriate services. If more services are found, the best match is taken. Dynamic conversation selection allows the user to select the conversation within a service node at run-time, which may be useful when services are discovered also at run-time, and service interfaces are unknown. Then e-flow may select an appropriate conversation from the conversation repository.

The semantics of multi-service nodes is useful, when it is necessary to invoke multiple, parallel instances of the same type of service.

If a composite service has been defined, this may be fine for several users, but some may need additional functionality. To cope with these needs, e-flow adds the notion of dynamic service creation by including generic service nodes. A user may select additional services from a service pool, which are then put into a service list and passed to the generic service node in order to configure it. Hence, multi-service nodes allow the activation of several instances of the same service node, while generic nodes allow the dynamic selection of different service nodes.

Adaptive processes should reduce the need for human intervention in the process definition. Nevertheless, there might still be cases in which service process modifications are unavoidable to handle exceptions, improve the process or incorporate new laws or policies. E-flow distinguishes between ad hoc changes and bulk changes. Ad hoc changes are usually applied to a single process instance. Either the process schema, or the state of the process instance may be modified. Ad hoc changes are usually made in exceptional changes and are not expected to occur a second time. Bulk changes are made when some or even all of the running instances are affected. This may be the case when a service no longer exists, or a service definition is changed. Then running instances will have to be notified of the change in the environment, in order to cope with the new situation.

In Sun et al. (2003) the Web Services Composition Platform, StarWSCoP (Star Web Services Composition Platform), is introduced. It focuses on dynamic web services composition. The composition consists of four steps:

- service providers publish their services at a web service registry
- the Service Composition Engine decomposes user requirements into an abstract service and sends a SOAP request to the registry to find the proper services
- the web service registry returns a set of concrete services
- the service composition engine sends a SOAP request to the concrete services and binds to them.

StarWSCoP includes several modules: an intelligent system to decompose user requirements into the description of an abstract service, a service registry to provide a web service repository; a service discovery engine to find proper services that meet user requirements in the service registry; a composition engine that schedules the composite services to be executed in order, a wrapper to achieve interoperability of heterogeneous services of different providers that have been developed separately; a service execution information library to store trace information of composite web service execution; a QoS estimation to estimate real-time QoS metrics of the composite service; and an event monitor to monitor events and notify the composition engine.

To support QoS based dynamic web services composition, WSDL is extended with QoS attributes, such as time, cost, or reliability:

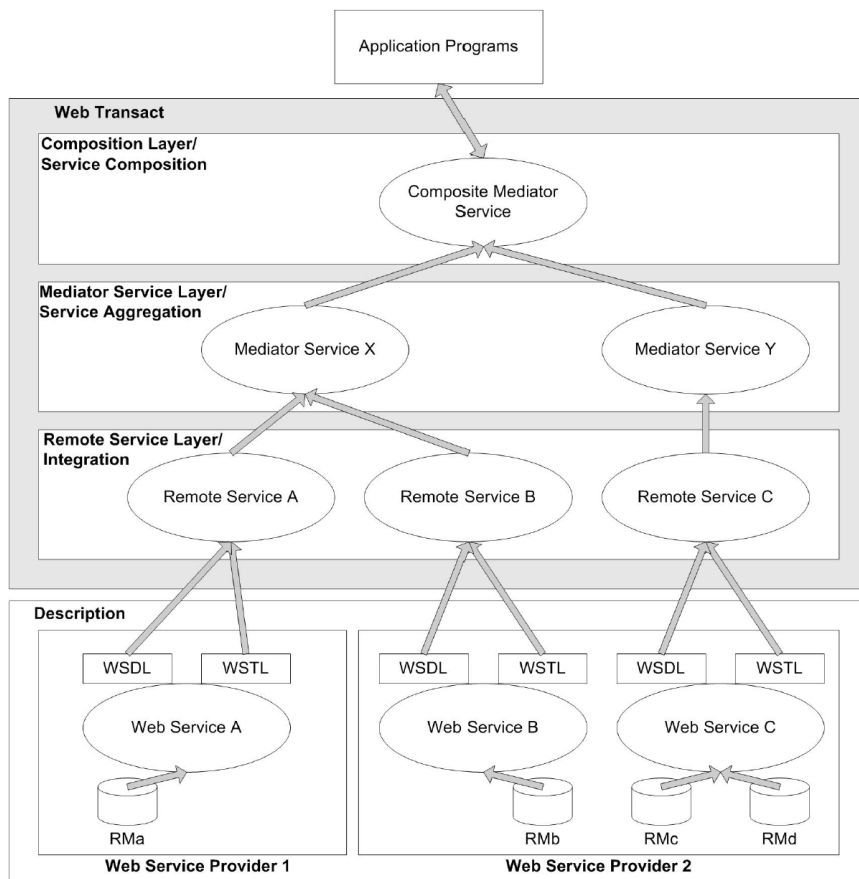
```
<portType name="GetTotalFilmCost">
  <operation name="getTotalFilmCost"
    cost="20"
    time="50"
    reliability="90%">
    ...
  </operation>
</portType>
```

An ontology-based layer is added to UDDI to achieve semantic match for web services.

A second Wrapper is added to the Composition platform to conceal differences or inconsistencies in data types, security policy or content of web services. The Wrapper adds a Communication manager to translate messages between different transport protocols like HTTP or SMTP, a Security Manager to traverse firewalls and handle authentication and authorisation, a Content Manager to convert between different document representations, a Conversation Manager for conversational differences and a QoS Monitor to monitor QoS metrics of web services.

Web services composition shares many requirements with business process management (Pires et al., 2003). Web services, as well as business processes need to coordinate the sequence of service invocations, manage data flow, and manage execution of compositions as transaction units. The difference to business processes is that web services are autonomous, heterogeneous units, which makes the implementation of web services composition harder. Each service provider has its own business rules. Pires et al. (2003) introduce a framework called 'WebTransact' which provides the necessary infrastructure to build reliable service compositions. WebTransact is composed of a multi-layered framework containing a Service Composition Layer, a Service Aggregation Layer, an Integration Layer, and a Description Layer shown in Figure 2.

Figure 2 The multi-layered architecture of WebTransact



Starting at the bottom (Description Layer) WebTransact uses WSDL for describing the service functionalities and adds a web Services Transaction Language WSTL on top of WSDL enhancing it with functionalities facilitating web services composition. In short, WSTL describes the transaction support for a web. The code example below shows a WSDL code snippet and the corresponding WSTL extension:

```
<operation name="getTotalFilmCost">
  <input message="tns:cancelSoapIn"/>
  <output message="tns:cancelSoapOut"/>
</operation>

<wstl:transactionDefinitions>
<wstl:transactionBehavior operationName="getTotalFilmCost"
  type="retriable" />
</wstl:transactionDefinitions>
```

The integration, or remote service layer defines a remote service as a logical unit of work that performs a set of operations at a particular site. Each operation has a well-defined transactional behaviour. WebTransact defines four types of transaction behaviour to accommodate the levels of transaction support: compensational if the effects of an operation after its execution can be undone by invoking another operation, virtual-compensational for operations whose underlying system supports 2PC, retriable if it is guaranteed that an operation can succeed after a finite set of repeated executions, and pivot if it is neither compensational nor retriable.

The service aggregation, or mediator service layer aggregates semantically equivalent remote services providing a homogeneous view of heterogeneous remote services. Semantically equivalent services provide different WSDL descriptions, but offer the same functionality. Thus, mediator services are simply virtual services that delegate method invocations to the corresponding concrete remote service. Operations of mediator services also have a well-defined transactional behaviour, which can be either compensational, retriable, or pivot. If semantically equivalent remote services have different transactional behaviours, the mediator service will use the least restrictive behaviour among the behaviours of the aggregated remote services.

Each mediator service exposes one single interface to make composition easier. A remote service provides mapping information and content description of a web service to the mediator service. When it comes to composition, a composite mediator service describes the interaction patterns for a set of cooperating mediator service operations needed to complete a certain task.

In WebTransact, a composition and the sequence of operations invoked are described by WSTL. WSTL models compositions as composite tasks, which can consist of wither atomic tasks or recursively of other composite tasks. Tasks are identified by its signature, execution dependencies, data links, and rules. Tasks may also be further sub-divided into mandatory or desirable tasks.

5.2 *Model driven service composition*

The paper (Orriens et al., 2003a) introduces the approach of Model Driven Service Composition, which is based on dynamic service composition discussed in the previous section, since it should facilitate the management and development of dynamic service

compositions. UML (Unified Modelling Language) is used to provide a high level of abstraction, and to enable direct mapping to other standards, such as BPEL4WS. The OCL (Object Constraint Language) is used to express business rules and to describe the process flow. Business rules can be used to structure and schedule service composition, and to describe service selection and service bindings. The paper identifies two main use cases: the process of service composition development, which may be sub-divided into the four phases of service definition, scheduling, construction, and execution. In order to enable representation of all possible service compositions, the paper introduces an information model – an abstract meta-model – that models components and relationships between the components. For that purpose, service composition elements and service composition rules are defined.

Since service composition is closely related to workflow (Preuner and Schrefl, 2002; Piccinelli and Williams, 2003; Hansen et al., 2002; Georgakopoulos et al., 2002), service composition elements are based on the business process elements activity, representing a well-defined business function (condition), constraining the behaviour of the composition by adding pre- and post-conditions (event), describing occurrences during the process of service composition (flow), defining a block of activities, (message), containing input- or output-information (provider) describing a party offering concrete services, and the abstract class (role), describing a party participating in the service composition (process).

Service composition rules can be modelled using OCL and can be structured as follows: structural rules guide the process of structuring, scheduling and prioritising within service composition, data rules control the use of data and message relations, behavioural rules to control event occurrences and enforce integrity constraints, resource rules to guide the use of resources, such as service selection, providers, and event raisers, and, finally, exception rules to guide exceptional behaviour in the service composition process.

The business rule driven service composition introduced in Orriens et al. (2003b, 2003c) is very similar to the model driven service composition. The authors use the same service composition elements and define four components for composition architecture: the definer, the scheduler, the constructor, and the executor to map the four phases of the service composition process. The framework introduced, basically consists of a service composition manager SCM assisting the user in developing, executing, and managing service compositions, and a service composition repository maintaining composition elements and rules. The SCM determines user requirements and passes them to the definer to define an abstract service composition. In order to find out if any information is already available, the definer prompts the composition engine to query the composition rule repository, and in case a corresponding composition rule already exists, the composition elements repository returns the activities to the definer. New rules, on their part, are stored in the composition rule repository and composition elements repository respectively. The scheduler develops a set of concrete compositions and lets the user select an alternative, which is then passed to the constructor to generate the executable software. The executor monitors the service execution.

Paper (Orriens et al., 2003b) introduces a classification scheme for business rules: structure related rules facilitate the specification of the way in which service composition is to be carried out, role related rules govern the participants involved in the service composition process, message related rules regulate the use of information, event related rules govern the behaviour of service composition in reaction to expected or unexpected events and constraint related rules represent conditions in service composition.

5.3 *Declarative service composition*

The approach followed by enTish (Ambroszkiewicz, 2003) is somewhat different from typical composition platforms. Services are typically created on the fly to realise client requests. Anyway, most frameworks are based on the assumption that first the business process has to be created. For enTish, a different architecture is needed, since client requests are expressed in a declarative way using formal languages. The declarative approach consists of two phases: the first phase takes an initial situation and the desired goal as starting point, and constructs generic plans to reach the goal. The latter one chooses one generic plan, discovers appropriate services, and builds a workflow out of them. The first phase is realised using PDDL (Planning Domain Definition Language) and estimated-regression planning as used in XSRL (XML Web-services Request Language), which must provide machine-readable semantics and specify the abstract service behaviour. The second phase may be realised by using existing process modelling languages, such as BPEL.

enTish implements a layered architecture, containing a conversation layer, a functionality layer and a database management layer representing real world data. The functionality layer consists of two components, namely raw application and a filter associated with the raw application. The filter analyses input information of an operation, which may consist of several parameters, in order to produce the desired output by the raw application. Constraints are expressed in one common description language. The conversation layer implements a conversation protocol specifying the order for message exchange.

Since the architecture is quite different from the traditional web service architecture presented in the introduction, the concept of web service description and web service registry has to be revised. WSDL does not describe service attributes, and UDDI lacks of information about what the service does. The goal of the enTish approach is, thus, to create a description language that overcomes the drawbacks of currently existing standards, expresses client requests and is open, and is of distributed use to enable users to introduce new resource types, functions and relations with unique names, namely URIs. Furthermore, a coordination model and a universal protocol are required to realise the declarative approach.

In SELF-SERV (Sheng et al., 2002), web services are declaratively composed and then executed in a dynamic peer-to-peer environment. SELF-SERV defines three types of services: elementary and composite services, and service communities. Service communities can be seen as containers of alternative services. Service composition is based on state-charts, gluing together an operation's input- and output-parameters and produced events. Service execution is monitored by software components called coordinators, which initiate, control, and monitor the state of a composite service they are associated with. The coordinators retrieve the state relevant information from the service's state-chart and represent it in what is called a routing table containing pre-conditions and post-processings.

The SELF-SERV framework features a service manager and several pools of services, all implemented in Java and communicating via the exchange of XML documents. The service manager consists of a service discovery engine facilitating the advertisement and location of services, a service editor facilitating the definition of new services and the edition of existing ones, and a service deployer generating routing tables of every state of a service's state-chart and uploading these tables into the hosts of the

corresponding composite service. Input and output information are both structured as XML documents.

5.4 Automated vs. manual web services composition

In Section 2, when discussing service composition models, BPEL4WS were introduced as a manual composition model. In this sub-section automated service composition is dealt with in greater detail. In contrast to manually describing service composition, there is also a lot of effort done in the direction of automated or ontology based service composition (Tosic et al., 2002b; Fileto et al., 2003). The most frequently used definition of ontology in the literature on semantic web services (Motta et al., 2003; Medjahed et al., 2003; Horrocks and Tessaris, 2002) is that an ontology is a collection of web services that share the same domain of interest and describe how web services can be described and accessed. Ontology languages, such as RDF, DAML-S (DARPA Agent Markup Language for Web services), DAML+OIL, or OWL (Web Ontology Language) have formal specifications and, thus, can be queried and provide the means to define sophisticated class properties. A web service dealing with film equipment may belong to a film-production-ontology for example. DAML-S provides the mechanisms to organise web services into ontology.

In Tosic et al. (2002a) the need for requirements for ontology is discussed. The paper contains a list of some proposed ontology needed for the management of web services:

- ontology of QoS metrics
- ontology of measurement units
- ontology of currency units
- ontology of measured properties
- ontology of measurement methods.

Furthermore, in Tosic et al. (2002a) also an ontological definition of Quality of Service (QoS) metrics is provided and may be summarised as follows:

- metric name
- short human-readable textual description
- measured property (a link to the ontology of measured properties)
- formulae (from zero to many) how the given QoS metric can be computed from other QoS metrics, each accompanied by a unit conversion rule
- invariant relationships with other QoS metrics.

In Akram et al. (2003), the web service environment is characterised by the exploratory, volatile, and dynamic features. Exploratory refers to dynamic service invocation during runtime when the need for a certain service is established while the feature dynamic covers the service evolution aspect when the content of a web service changes over time. Volatile means that a web service handling a request may not be available at a later point of time. The approach taken here proposes the usage of ontology in an exploratory

service space and the employment of agents (Velasco and Castillo, 2003) to deal with web service volatility and dynamism.

Changes in web service execution may be internal, where referring to the change of information provided by the web service, or external nature, including the volatility of web services, which refers to a temporal or permanent unavailability of a web service due to network failure, service relocation, request overload, or operation rename. The problem of supporting dynamic service requests contains the two facets web service discovery and selection and adapting to web service changes. The authors define the term change management as generic term for detection, propagation, and reaction to internal and external changes as well.

A service request is defined as a list of atomic orders where each in turn consists of application specific attributes. For example, stock identifier or stock category in a stock market service environment. Formally, a request may be specified by the following grammar:

```
req ::= req ; req |
      req || req |
      order
```

This sample shows that a request may either consist of two sequential requests or of two requests executing concurrently or of an order. An order id defined as tuple containing a value for each attribute necessary for the request. An asterisk may be used as wildcard, if the value cannot be specified.

```
Req: { ("buyEquipment", Camera, *, *, *, 50, 00:15) }
```

The sample request shown in the grey box above describes an order to buy 50 cameras, in order to improve the film equipment of the crew. In this example, the last value represents a timestamp while the remaining three may take any value.

To overcome the problems that arise when it comes to dynamic service discovery and selection, Akram et al. (2003) proposes an approach to organise web services into ontology.

DAML-S divides service descriptions into profiles, providing a high level description of the service, models describing the execution flow of the service, and groundings providing a mapping between DAML-S and WSDL, in order to describe how the service has to be invoked. The profile should contain enough information to actually discover a web service including functional descriptions of the service, such as input or output parameters. The description part provides human understandable information, such as service name and some textual description. The following code snippet gives a brief idea of how a DAML-S description looks like:


```
<daml:Class rdf:ID="Camera">
  <rdfs:label>Equipment Item</rdfs:label>
  <rdfs:subClassOf rdf:resource="&service;" />
</daml:Class>
<rdf:Property rdf:ID="Equipment">
  <rdfs:label>Equipment</rdfs:label>
  <rdfs:subPropertyOf
    rdf:resource="&profile;serviceCategory" />
  <rdfs:domain rdf:resource="&service;serviceProfile" />
  <rdfs:range rdf:resource="&daml;#Thing" />
</rdf:Property>
```

The first four lines represent a human understandable description of the web service. They define the name and specify that the class is a service. The rest of the code defines a property of the service. If a search is performed, a request is sent to the DAML-S registry, which in turn looks for matching services. If a proper description is found, the properties of the service are checked to see if they fit the requirements.

Now, what if the price for a stock for example rises above a certain limit? If a film director wants to buy equipment for a certain amount of money – considering the sample request on the previous page – and the price rises during execution, the request has to be aborted and another service has to be found. To react to changes, Akram et al. (2003) suggests the usage of agents that play the roles of monitors and notifiers. Changes may occur before service invocation or during service execution. Soft states store a participant list that is used to maintain the participants in a web service environment interacting with the web services. An agent is assigned to each web service to monitor changes in the status of the service, to verify the availability and to notify the participants. Changes are detected by comparing a service's operations with the corresponding service description in a UDDI registry. If operations have changed, these changes are updated in the registry, and the participant list is refreshed by removing references to web services that have become invalid, due to changes of web service interfaces. The way a participant then reacts to a change depends on the type of change and the availability of alternate services. If no alternate service exists, the request must be cancelled. In case of an internal change, the previous data will be replaced by the current response (Akram et al., 2003).

The authors propose an architecture including a GUI to interact with the user and a request broker as intermediary between GUI and web services to support dynamic service requests. The GUI allows the user to specify service requests and validates user input. The request broker then handles decomposition of user requirements, service selection, service invocation, change management, and response consolidation. The request broker consists of several components, which are the Communication Interface namely: that is responsible for a secure communication between GUI and web services using the SOAP standard, the request decomposer that decomposes a request into a list of orders, the request manager that evaluates the request execution and determines whether a request has been executed safely or not and the notifying agents that monitor changes in the distributed environment and propagate changes to the entities that need to know about these changes.

Implementation of the system has been done using Java Technologies, such as RMI and JDBC, WSDL, and DAML-S to describe the web services, UDDI to publish service descriptions and SOAP to allow messaging among the entities involved.

Wu et al. (2003) proposes a DAML-S process model supporting AI planning techniques for automatic web service composition called SHOP2. DAML-S and OWL and DAML+OIL as well are used for semantic web service description. Web services are seen as actions with pre- and post-conditions. In DAML-S ontology services are modelled as processes: atomic, composite, and simple processes. Atomic processes can be directly executed by calling the corresponding web service. Composite processes may be decomposed into atomic or other composite processes. Simple processes are abstract representations of atomic or composite processes. They are not directly executable but provide an abstract view of the process.

The DAML-S ontology describes a set of classes and properties, specific to the description of web services. The upper ontology of DAML-S comprises the service profile for describing service advertisements, the process model for describing the actual program that realises the service, and the service grounding for describing the transport-level messaging information associated with execution of the program (Zeng et al., 2003). The code example below shows the definition of the process term in DAML-S:

```
<daml:Class rdf:ID="Process">
  <daml:unionOf rdf:parseType="daml:collection">
    <daml:Class rdf:about="#AtomicProcess"/>
    <daml:Class rdf:about="#SimpleProcess"/>
    <daml:Class rdf:about="#CompositeProcess"/>
  </daml:unionOf>
</daml:Class>
```

The atomic process on the other hand is described as follows:

```
<daml:Class rdf:ID="AtomicProcess">
  <daml:subClassOf rdf:resource="#Process"/>
</daml:Class>
```

To stick with the example described in our case study, we would like to provide a small code sample to give a basic idea of what a DAML-S description of a process definition looks like:

```
<daml:Class rdf:ID="ShootingLocator">
  <rdfs:subClassOf rdf:resource="Process.daml#AtomicProcess"/>
</daml:Class>

<rdf:Property rdf:ID="LocationPreference">
  <rdfs:subPropertyOf rdf:resource="..."/>
  <rdfs:domain rdf:resource="#ShootingLocator"/>
  <rdfs:range rdf:resource="concepts.daml#LocationPreference"/>
</rdf:Property>
```

The class `ShootingLocator` may be used to find proper locations for shooting a film. It is supposed to be an atomic process and is associated with a `LocationPreference` property.

SHOP2 (Wu et al., 2003) is a domain-independent HTN (Hierarchical Task Network) planning system that creates plans by task decomposition. SHOP2 decomposes a task into very small primitive tasks that can be directly executed in the order planned by SHOP2.

The approach taken in Zeng et al. (2003) describes a framework called OntoMat-Service, which allows for seamlessly browsing conventional web pages, including XHTML advertisements for web services, direct, manual invocation of an advertised web service as a one-off use of the service, tying web service advertisements to each other when browsing them, tying web service advertisements to one's own conceptualisation of the web space when browsing them and invoking such aggregated web services.

The WSDL definition listed below shows the typical WSDL elements types, message and portType with one or more operation tags enclosed by the definitions root element. The types element is enhanced by RDF class and properties definitions. The RDF part describes a class called Film crew in our example with a property element further specifying the service class.

```
<?xml version="1.0" encoding="UTF-8"?>
<definitions name="LaptopService"
  targetNamespace="http://filmmanagement.wsdl/"
  >
  <types>
    <rdf:RDF>
      <rdfs:Class rdf:ID="Filmcrew">
        <rdfs:label>Film crew</rdfs:label>
      </rdfs:Class>
      <rdf:Property rdf:ID="wage">
        <rdfs:label>Wage</rdfs:label>
        <rdfs:range rdf:resource="#rdfs:Literal"/>
        <rdfs:domain rdf:resource="#Filmcrew"/>
      </rdf:Property>
    </rdf:RDF>
  </types>
  <message name="totalFilmCost">
    <part name="cost" type="rdf:ID=costs"/>
  </message>
  ...
  <portType name="FilmcrewService">
    <operation name="getTotalFilmCost" />
    <output message="tns:totalFilmCost"
      name="totalFilmCost"/>
    </operation>
    ...
  </portType>
</definitions>
```

The OntoMat-Service starts with a common WSDL service description by the service provider. The WSDL description is then parsed and made human readable and formatted as an HTML document. This page may contain several service descriptions and cross-links to related services. In the next step, the user may either choose text-phrases that correspond to some underlying machine understandable semantics and map them to his own terminology or directly invoke the web service. The result of the mapping step is a set of mapping rules between the web service ontology and pre-loaded ontology. Before invoking the web service operations, a user might wish to choose the operations and the mappings he wants to use. The OntoMat-Service provides a web service-planning module, which computes logically possible service flows. Therefore, a knowledge base provides pre- and post-conditions, goals, etc., for web services.

The OntoMat-Service-Surfer described in the final section represents a graphical user interface, which implements a drag and drop mechanism enabling the user to highlight items on a web page, and pulling them into the ontology browser of the OntoMat-Service framework. The resulting ontology may in turn be published to enable third parties to execute services based on these semantics.

In Su et al. (2003), the authors focus on a composability model for web services, automatic generation of composite web services, and prototype implementation and experiments. They develop a framework called WebDG (Web Digital Government) to provide E-government services and to serve citizens in collecting benefits to which they are entitled, especially those facing low income, mental illness, addiction, or mental retardation.

Web service composition (Chen et al., 2003; Cardoso and Sheth, 2003) requires semantic service description, in order to enable services to interact with each other. WSDL does not provide any semantic description to web services. The figure below depicts a diagram developed by Su et al. (2003), which shows the key features of WSDL in ovals surrounded by dashed lines and those added by semantic descriptions in grey filled ellipses. Edges are unidirectional and add multiplicity information to describe the relations between the entities.

Operations, for example, consist of an input message, an output message, exactly one name, a description and an operation mode, which may either be notification, solicit response, one-way, or request-response. Furthermore, it can be described by purpose and quality, which are tuples of three attributes. Function, synonym, and specialisation define the operation's purpose and fees, security and privacy specify the operation's quality. Operations are categorised by domain, synonym, and specialisation. Domains are areas of interest where the operation might be used (for example an operation `calculateFilmCost` may only be interesting for a film production company). Synonyms are alternative function names and specialisation contains a set of characteristics of a certain function.

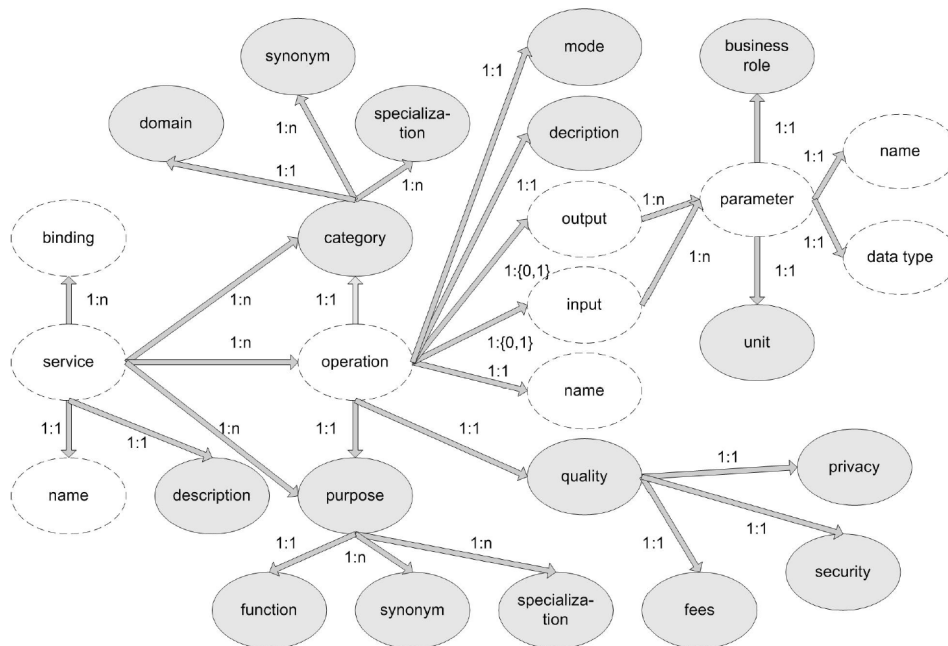
After extending WSDL with semantic capabilities, it is now possible to deal with a composability model for web services. Su et al. (2003) identifies two different composability rules: syntactic and semantic rules. Syntactic rules include mode and service composability. Assuming that two services are mode composable – a one-way operation, at either the client or server side, must be connected to a notification mode at the other side and solicit response maps to the request-response messaging pattern – both services must also agree to compatible transport protocols in order to be able to exchange messages.

Semantic rules comprise message composability, operation semantics composability, qualitative composability, and composition soundness. Message composability is based on the definition that message types, which in fact are parameter data types, used by the operation to be invoked, must be compatible with the ones passed by the client call. Messages are directly compatible if they have the same data type, and indirectly compatible if the type used by the client call is derived from the one used by the web service operation. Operations are semantically composable if the domains of interest are similar or synonymous, and if both operations provide the same characteristics as well. Operations are qualitative composable if both match the same qualitative properties, such as security or privacy. To compose services in a way that provides an added value (Su et al., 2003) introduces the notion of composition templates to provide composability soundness. For example, two independent services such as `engageFilmCrew` and

bookFlight may be composed in order to engage actors and actresses for shooting a film and flying them to the shooting location. Composition templates are associated with each composite service and are used to compare values added by different compositions. To check whether a composition is sound or not, stored templates are used to store a composite service's template in the service repository. A service composition is considered to be sound if its template is the subset of a stored template.

With all these rules in mind, Su et al. (2003) proposes the architecture for automatic service composition. The approach consists of four large modules, called specification, matchmaking, selection, and generation. For the service selection phase, Su et al. (2003) introduces a high level language – CSSL (Composite Service Specification Language) – which is a semantic description language for web service compositions. CSSL extends WSDL and represents the ontology-based graph presented in Figure 3. In the matchmaking phase, a special matchmaking algorithm, integrating the composability rules described above, is applied to generate composition plans that meet the composer's requirements. QoS constraints are used in the selection phase to select the most proper composition plan from the output. Finally, the resulting composite service is generated using a composition language such as WSFL, XLANG or BPEL4WS as specified in the CSSL description.

Figure3 Ontology based description of web services



5.5 Context based Web Service Discovery and Composition

Services become available through different channels, meaning that e-services may be accessed using different devices, such as PCs, palmtops, cell-phones or TV sets, and also different network technologies and protocols (Baresi et al., 2003). The goal should be to deliver the same service via web, SMS, or call centres, for example. The Multi-channel

Adaptive Information Systems project MAIS proposed in Baresi et al. (2003) aims to create a platform, a methodology and design tools for the development of distributed information systems, based on e-services. Baresi et al. (2003) defines a modelling framework for adaptive information systems separating the application and technological levels and providing formalised contracts between the Service Provider and the Service Consumer.

E-services in this context are defined by a request perspective and a provisioning perspective, both dealing with QoS issues in order to allow service negotiation based on user requirements, channels, and provider constraints. Negotiation is encoded in adaptation rules, which allow the dynamic adaptation of the execution flow of the services, described as workflows.

The provisioning perspective specifies the service provides, the service functionality, and how to invoke the service. The request perspective specifies the service consumer, the QoS level and a certain context. In Baresi et al. (2003), the authors use UML to describe the entities involved in both perspectives and show its associations.

In the provisioning perspective shown in Figure 4, an e-service is described by a unique identifier, a textual description, and a category. An e-service is associated with one or more channels and an association class called C-Quality Dimension to further specify QoS attributes, such as response time, availability, usability, accessibility, integrity, bandwidth, reliability and price. An e-service is also associated with exactly one service provider and an association class called P-Quality Dimension to describe, similar to C-Quality Dimension, quality parameters guaranteed by the provider. The functional description of an e-service contains pre-conditions, post-conditions, operations with inputs and outputs, and events that may occur during the execution of an operation. Finally, an e-service is grouped into compatibility classes associated with an abstract e-service to express several services with the same functionality in terms of the process execution, to enable service substitution during execution if a service fails. Therefore, an abstract e-service has an association to the functional description of an e-service. Of course, e-services may be aggregated and composed to a composite e-service, which extends the e-service class.

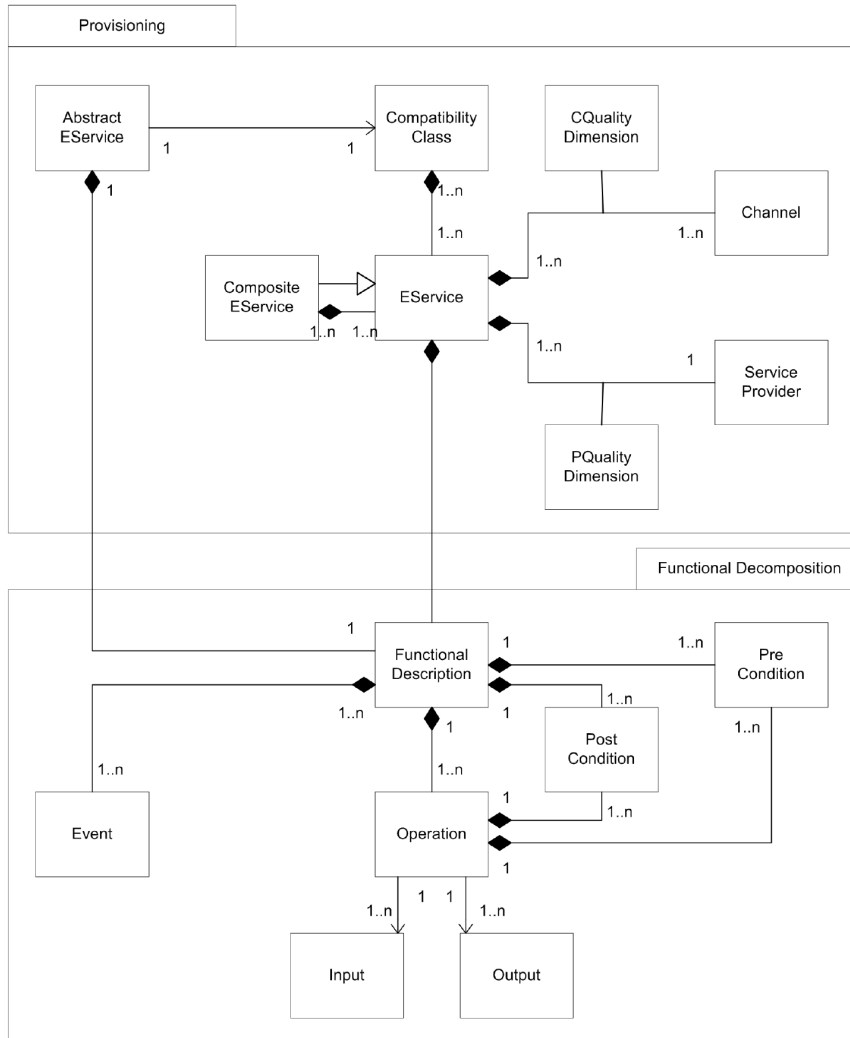
The class structure depicted in Figure 4 can be translated into a XML structure as depicted below:

```

<Service>
  <Provider>
    <PQualityDimension>
      <ProvisioningTime>10</ProvisioningTime>
      <Availability>90%</Availability>
      <Price>100</Price>
    </PQualityDimension>
  </Provider>
  <Channels>
    <Channel name="...">
      <Device>PDA</Device>
      <Network>GPRS</Network>
      <CQualityDimensions>
        <Availability>99%</Availability>
      </CQualityDimensions>
    </Channel>
  </Channels>
</Service>

```

Figure 4 The provisioning perspective



Services may be accessed by different types of devices using different technological channels. Information about networks, devices, and service providers is supplied by a context description that is used by the context manager in the MAIS architecture. The context manager is accessed both by service providers and service consumers as well. The context manager is linked to an interaction-enabling platform via an adaptive channel. The interaction-enabling platform passes the QoS requests to the above e-service composition platform, which retrieves the services matching user requirements out of an e-service directory and performs the composition based on the adaptation rules.

Heterogeneous client capabilities and the increasing number of connected devices require information services to be accessible from all devices in a similar fashion (Keidl and Kemper, 2004b). The authors define the word ‘context’ as the kind of information that makes information services aware of their current context.

They propose a framework in which the context information is passed within a SOAP header between web services. The header may contain several context blocks, each associated with one context type. During execution a web service may change its context information by inserting different context information into the SOAP header and sending the SOAP message to another web service, for example. Totally, there are four components that may process context information: the web service, context plug-ins, context services and also clients. Web services always have full control over the context information. They decide how the information influences their execution and their replies. Context plug-ins are programmed in Java and installed at each local host. Each plug-in is associated with one context type. Context services are also associated with one context type and must be available over the internet. It is not necessary to install them locally. Context information is pre-processed and post-processed before being actually sent by the web service. In the framework, the context block could be processed by several components. Thus, processing instructions are needed to specify rules of precedence.

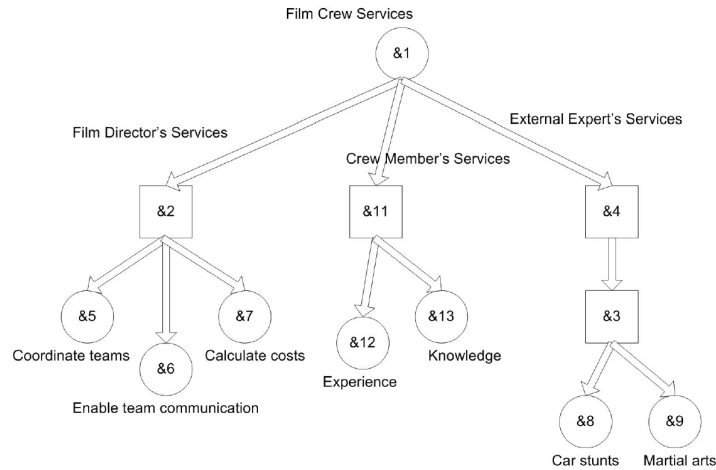
The context framework proposed here provides several context types, such as location (GPS coordinates, country, local time and time zone), information about the consumer invoking service (name, e-mail address, preferences) and client context information, such as hardware or software. The authors demonstrate the functionality of the framework by developing an information service, which is invoked by several different clients, such as PCs or PDAs. In case of a PDA invoking the service, unnecessary representation information is cut out of the SOAP message, since the display size would not meet the graphical representation requirements.

Accessing data via mobile or pervasive systems through web services has the advantages of hiding the heterogeneous nature of data and providing a well-defined interface for data access (Doulkeridis et al., 2003). In such a dynamic environment, the service discovery plays a crucial role. The approach taken in this paper is to describe a web service entry in a web service registry in a context-based way that answers the following questions:

- which services (or devices) are available in location L at time T?
- which services return results that the requestor device can represent?
- which services were available at timestamp T?
- which services are published by user U and his device D?

Since nowadays applications rarely use context information, the aim of the paper is to provide a context representation model to handle any type of context information as long as it can be represented as key-value pairs. In order to develop a context-aware service directory, the authors introduce a data model called multi-dimensional OEM (MOEM) graph, which may hold information presenting different facets under different contexts.

Figure 5 shows the graphical representation of a service directory as a MOEM graph. Each service category may be represented as an individual MOEM graph. Rectangular nodes represent multi-dimensional entities that present different facets under different contexts. Circular nodes may either contain data or hold the reference to a service. Bold lines are called context edges and define the context under which the services hold. The plain entity edges represent relations between the entities.

Figure 5 A MOEM graph representing a service directory

To search and find services matching a certain search criteria, a breadth-first search algorithm is developed. It takes the user defined context description as input and performs the search starting at the root. If a context edge matches the context specifier, the node is kept for further processing or else it is ignored. The result of the search is a list of all services, which meet the search criteria.

The MOEM graph is similar to a tree with the exception that a leaf can belong to one or more sub-trees. Subsequently, an alternate data representation is introduced by translating the MOEM graph into an equivalent graph representation, and by means of a state table containing the device-type, user, return type, and service identifier in order to improve the search.

6 Summary

So far, we have dealt with many different approaches and frameworks that have been developed in order to provide widely usable web service composition platforms. The aim of this section is to give a summary of what we have presented so far. This is achieved by creating a table, listing most of the frameworks that have been discussed in this paper, mapping them to the corresponding composition approach, and listing some of the most important features that were identified to categorise the composition frameworks. Some approaches, such as enTish do not meet the requirements listed in the Table 1, since they do not provide a development platform for composite web services, but rather propose language or protocol specifications.

By the analysis of the composition frameworks, some modules and support features that are of great importance for designing and developing composite services have been identified. Those include an execution monitor to monitor and trace service execution, dynamic service selection, which is a profound part of service composition and may be used to choose appropriate service conversation protocols from a conversation repository, QoS modelling and evaluation, transaction support, and the possibility to graphically modelling service flows.

Table 1 Web services composition models

<i>Framework</i>	<i>Composition strategy</i>	<i>Execution monitor</i>	<i>Dynamic conversation selection</i>	<i>QoS Modelling</i>	<i>WSDL Language extension</i>	<i>Transaction support</i>	<i>Graph support</i>
E-flow	Dynamic composition	Yes	Yes	No	No	No	Yes
MAIS	Context based	No	No	Yes	Yes	No	No
MOEM	Context based	No	No	No	No	No	Yes
SELF-SERV	Declarative composition	No	No	No	No	No	Yes
OntoMat-Service	Semantic composition	No	No	No	Yes	No	No
SHOP2	Semantic composition	No	No	No	Yes	No	No
WebTransact	Dynamic composition	No	No	No	Yes	Yes	No
StarWSCoP	Dynamic composition	Yes	No	Yes	Yes	No	No

Table 1 provides an overview on current web service composition frameworks. The reader will also notice that there still has to be done a lot of development effort until a service composition platform will provide all those features listed above. As far as we are able to judge, service composition is a technology that will work only conditionally using the frameworks listed above (Srivastava and Köhler, 2003). We think that the current state of the development efforts can be seen as first attempts to create service composition platforms, since the current state of research lacks additional specifications concerning QoS extensions in WSDL, to name a very simple but crucial example, and some platforms listed in Table 1 have developed their own data structure for extending WSDL. Research results in the area of service composition show how limited the current web service technology is since it is purely based on standards, such as SOAP, WSDL or UDDI.

Unfortunately, most papers do not provide information about how well the described approach already works. Many of the systems and approaches discussed in this paper do not provide the actual system. For those systems, which are available, we are currently in the process of applying the frameworks to our film crew management case study. In our view, platforms, such as e-flow, WebTransact, and StarWSCoP sound very promising; unfortunately, they lack semantic description, and thus, it would also be interesting to follow the progress that will be made in the area of semantic web services, unless service composition works without semantic service description.

7 Conclusion

Many different standards have been proposed and various approaches have been taken to create a widely accepted and usable service composition platform for developing composite web services. web services composition seems to have higher chances of success compared to traditional composition middleware, due to the standardisation efforts that have taken place already. With WSDL, web services may be described in a

consistent way according to their functionalities. Of course, QoS and semantic descriptions have been proposed to extend the current WSDL standard, but have not yet found overall acceptance. UDDI and other registry based data models have been implemented, but are not widely used, and in the case of dynamic service discovery, it does not yet meet the requirements.

In this paper, different service composition strategies, introduced some existing composition platforms and frameworks have been presented and an attempt has been made to compare existing research approaches with each other by finding common characteristics and features.

New software is increasingly built by web services. Hence, service composition is a crucial topic, which has a high impact on many different areas of research. The main ambition is to enforce our research interest in the service composition area and also service evolution, since it is very likely that aggregated services may be changed over time.

Specifications, interactions, non-functional attributes, and internal changes invisible to the outside world have to be considered. Different views on service changes, who is managing the changes, and who propagates them, issue topics, such as service metrics and monitoring of service execution and service environment must be provided. Additionally, disruptive changes when one wants a web service to cooperate with others, constructive changes when one wants to provide better services than those offered so far must be thought of. The latter topic is closely related to planned changes while unplanned changes must be made whenever a service fails. All these considerations affect issues such as compatibility and versioning of web services. In the ongoing research work in service composition and evolution, a reduction of development time of services, integration effort and testing, provisioning an integration platform for management, and monitoring of aggregated services and their environments, and methodologies for testing dynamic service interactions and attributes is expected.

References

- Agarwal, S., Handschuh, S. and Staab, S. (2003) Fensel, D. *et al.* (Eds.): *Surfing the Service Web*, ISWC, LNCS 2870, Springer-Verlag Berlin Heidelberg, pp.211–226.
- Akram, M.S, Medjahed, B. and Bouguettaya, A. (2003) *Supporting Dynamic Changes in Web Service Environments*, ICSOC, LNCS 2910, Springer-Verlag Berlin Heidelberg, pp.319–334.
- Alonso, G., Casati, F., Kuno, H. and Machiraju, V. (2004) *Web Services. Concepts, Architectures and Applications*, Springer-Verlag Berlin Heidelberg.
- Álvarez, P. *et al.* (2003) Guelfi, N. *et al.* (Eds.): *A Java Coordination Tool for Web-service Architectures: The Location-Based Service Context*, FIDJI, LNCS 2604, Springer-Verlag Berlin Heidelberg, pp.1–14.
- Ambroszkiewicz, S. (2003) Benatallah, B. and Shaw, M.C. (Eds.): *Entish: An Approach to Service Composition*, TES 2003, LNCS 2819, Springer-Verlag Berlin Heidelberg, pp.168–178.
- Andrews, T., Curbera, F., Dholakia, H., Golland, Y., Klein, J., Leymann, F., Liu, K., Roller, D., Smith, D., Thatte, S., Trickovic, I. and Weerawarana, S. (2003) *Specification: Business Process Execution Language for Web Services Version 1.1*, <http://www-106.ibm.com/developerworks/library/ws-bpel/>.
- Benatallah, B. *et al.* (2004a) Atluri, V. (Ed.): *On Automating Web services Discovery*, Received: December 15, 2002/Accepted: September 15, Published online: February 6, Springer-Verlag.
- Benatallah, B., Casati, F. and Toumani, F. (2004b) *Web Service Conversation Modeling. A Cornerstone for E-business Automation*, IEEE Internet Computing, January–February.

- Benatallah, B., Dumas, M., Fauvet, M-C., Rabhi, F.A. and Sheng, Q.Z. (2001) *Towards Patterns of Web Services Composition*, November.
- Benatallah, B., Dumas, M., Fauvet, M-C., Rabhi, F.A. and Sheng, Q.Z. (2002) *Overview of Some Patterns for Architecting and Managing Composite Web Services*, ACM SIGecom Exchanges, August, Vol. 3, No. 3, pp.9–16.
- Booth, D., Haas, H., McCabe, F., Newcomer, E., Champion, M., Ferris, C. and Orchard, D. (2004) *Web Services Architecture*, W3C Working Group Note 11, February, W3C Technical Reports and Publications, <http://www.w3.org/TR/ws-arch/>.
- Bultan, T., Fu, X., Hull, R. and Su, J. (2003) *Conversation Specification: A New Approach to Design and Analysis of E-service Composition*, ACM 1-58113-680-3/03/0005, WWW, May 20–24, Budapest, Hungary.
- Bunting, D., Hurley, M.C.O., Little, M., Mischkinsky, J., Newcomer, E., Webber, J. and Swenson, K. (2003a) *Web Services Coordination Framework (WS-CF) Ver1.0*, <http://developers.sun.com/techtopics/webservices/wscaf/wscf.pdf>, July 28.
- Bunting, D., Hurley, M.C.O., Little, M., Mischkinsky, J., Newcomer, E., Webber, J. and Swenson, K. (2003b) *Web Services Composite Application Framework (WS-CAF) Ver1.0*, <http://developers.sun.com/techtopics/webservices/wscaf/primer.pdf>, July 28.
- Bunting, D., Hurley, M.C.O., Little, M., Mischkinsky, J., Newcomer, E., Webber, J. and Swenson, K. (2003c) *Web Services Transaction Management (WS-TXM) Ver1.0*, <http://developers.sun.com/techtopics/webservices/wscaf/wstxm.pdf>, July 28.
- Bunting, D., Hurley, M.C.O., Little, M., Mischkinsky, J., Newcomer, E., Webber, J. and Swenson, K. (2003d) *Web Services Context (WS-Context) Ver1.0*, <http://developers.sun.com/techtopics/webservices/wscaf/wsctx.pdf>, July 28.
- Cabrera, F. *et al.* (2002) *Web Services Coordination (WS-coordination)*, <ftp://www6.software.ibm.com/software/developer/library/ws-coordination.pdf>.
- Cabrera, F., Copeland, G., Cox, B., Freund, T., Klein, J., Storey, T. and Thatte, S. (2001) *Specification: Web Services Transaction (WS-Transaction)*, <http://www-106.ibm.com/developerworks/webservices/library/ws-transpec/>.
- Cardoso, J. and Sheth, A. (2003) ‘Semantic e-workflow composition’, *Journal of Intelligent Information Systems*, Kluwer Academic Publishers, Vol. 21, No. 3, pp.191–225.
- Casati, F. and Shan, M-C. (2001) *Dynamic and Adaptive Composition of E-services*, Published by Elsevier Science Ltd., 12th International Conference on Advanced Information Systems Engineering (CAiSE 00).
- Casati, F., Ilnicki, S., Jin, L-J. and Shan, M-C. (2000) *An Open, Flexible, and Configurable System for Service Composition.*, HPL technical report HPL-2000-41.
- Chen, L., Shadbolt, N.R., Goble, C., Tao, F., Cox, S.J., Puleston, C. and Smart, P.R. (2003) Fensel, D. *et al.* (Eds.): *Towards a Knowledge-based Approach to Semantic Service Composition*, ISCW, LNCS 2870, Springer-Verlag Berlin Heidelberg, pp.319–334.
- Doulkeridis, C., Valavanis, E. and Vazirgiannis, M. (2003) Benatallah, B. and Shan, M-C. (Eds.): *Towards a Context-Aware Service Directory*, TES, LNCS 2819, Springer-Verlag Berlin Heidelberg, pp.54–65.
- Dustdar, S. and Treiber, M. (2005) *A View Based Analysis on Web Service Registries*, Distributed and Parallel Databases, Springer, forthcoming.
- Fileto, R., Liu, L., Pu, C., Assad, E.D. and Medeiros, C.B. (2003) ‘POESIA: an ontological workflow approach for composing web services in agriculture’, *The VLDB Journal*, Published online: September 30, Springer-Verlag, Vol. 12, pp.352–367.
- Freund, T. and Storey, T. (2002a) *Transactions in the World of Web Services, Part 1. An Overview of WS-transaction and WS-coordination*, <http://www-106.ibm.com/developerworks/webservices/library/ws-wstx1/>.
- Freund, T. and Storey, T. (2002b) *Transactions in the World of Web Services, Part 2. An Overview of WS-transaction and WS-coordination*, <http://www-106.ibm.com/developerworks/webservices/library/ws-wstx2/>.

- Georgakopoulos, D., Schuster, H., Cichocki, A. and Baker, D. (2002) *Process-based e-service composition for modeling and automating zero latency supply chains*, Information System Frontiers, Kluwer Academic Publishers, Vol. 4, No. 1, pp.33–54.
- Hansen, M., Madnick, S. and Siegel, M. (2002) Bussler, C. *et al.* (Eds.): *Process Aggregation Using Web Services*, WES, LNCS 2512, Springer-Verlag Berlin Heidelberg, pp.12–27.
- Horrocks, I. and Tessaris, S. (2002) Horrocks, I. and Hendler, J. (Eds.): *Querying the Semantic Web: A Formal Approach*, ISCW, LNCS 2342, Springer-Verlag Berlin Heidelberg, pp.177–191.
- Keidl, M. and Kemper, A. (2004a) Bertino, E. *et al.* (Eds.): *A Framework for Context-aware Adaptable Web-services*, EDBT, LNCS 2992, Springer-Verlag Berlin Heidelberg, pp.826–829.
- Keidl, M., and Kemper, A. (2004b) *Towards Context-Aware Adaptable Web Services*, ACM 1-58113-912-8/04/0005, WWW2004, May 17–22, New York, USA.
- Khalaf, R. and Leymann, F. (2003) Benatallah, B. and Shan, M-C. (Eds.): *On Web Services Aggregation*, TES, LNCS 2819, Springer-Verlag Berlin Heidelberg, pp.1–13.
- Medjahed, B., Bouguettaya, A. and Elmagarmid, A.K. (2003) *Composing web services on the semantic web*, The VLDB Journal, Published online: September 23, Springer-Verlag, Vol. 12, pp.333–351.
- Motta, E., Domingue, J., Cabral, L. and Gaspari, M. (2003) Fensel, D. *et al.* (Eds.): *IRS-II: A Framework and Infrastructure for Semantic Web Services*, ISCW, LNCS 2870, Springer-Verlag Berlin Heidelberg, pp.306–318.
- Narayanan, S. and McIlraith, S.A. (2002) *Simulation, Verification and Automated Composition of Web Services*, ACM 1-58113-449-5/02/0005, WWW, May 7–11, Honolulu, Hawaii, USA.
- Orriens, B., Yang, J. and Papazoglou, M.P. (2003a) Orłowska, M.E. *et al.* (Eds.): *Model Driven Service Composition*, ICSOC, LNCS 2910, Springer-Verlag Berlin Heidelberg, pp.75–90.
- Orriens, B., Yang, J. and Papazoglou, M.P. (2003b) Jeusfeld, M.A. and Pastor, Ó. (Eds.): *A Framework for Business Rule Driven Web Service Composition*, ER 2003 Workshops, LNCS 2814, Springer-Verlag Berlin Heidelberg 2003, pp.52–64.
- Orriens, B., Yang, J. and Papazoglou, M.P. (2003c) Benatallah, B. and Shan, M-C. (Eds.): *A Framework for Business Rule Driven Service Composition*, TES, LNCS 2819, Springer-Verlag Berlin Heidelberg, pp.14–27.
- Papazoglou, M.P. (2003) *Web Services and Business Transactions, World Wide Web: Internet and Web Information Systems*, Kluwer Academic Publishers, Vol. 6, pp.49–91.
- Piccinelli, G. and Williams, S.C. (2003) van der Aalst, W.M.P. *et al.* (Eds.): *Workflow: A Language for Composing Web Services*, BPM, LNCS 2678, Springer-Verlag Berlin Heidelberg, pp.13–24.
- Pires, P.F., Benevides, M.R.F. and Mattoso, M. (2003) *Building Reliable Web Services Compositions*, Web Databases and Web Services 2002, LNCS 2593, Springer-Verlag Berlin Heidelberg, pp.59–72.
- Preuner, G. and Schrefl, M. (2002) ‘Integration of web services into workflows through a multilevel schema architecture’, *Proceedings of the 4th IEEE Int’l Workshop on Advanced Issues of E-commerce and Web-based Information Systems (WECWIS 2002)*, IEEE.
- Ran, S. (2003) ‘A model for web services discovery with QoS’, *ACM SIGecom Exchanges*, Vol. 4, No. 1, pp.1–10.
- Rao, J. and Su, X. (2004) Li, M. *et al.* (Eds.): *Toward the Composition of Semantic Web Services*, GCC 2003, LNCS 3033, Springer-Verlag Berlin Heidelberg, pp.760–767.
- Sheng, Q.Z., Benatallah, B., Dumas, M. and Mak, E.O-Y. (2002) ‘SELF-SERV: a platform for rapid composition of web services in a peer-to-peer environment’, *Proceedings of the 28th VLDB Conference*, Hong Kong, China.
- Srivastava, B. and Köhler, J. (2003) *Web Service Composition – Current Solutions and Open Problems*, <http://www.zurich.ibm.com/pdf/ebizz/icaps-ws.pdf>.

- Su, S.Y.W., Meng, J., Krithivasan, R., Degwekar, S. and Helal, S. (2003) 'Dynamic inter-enterprise workflow management in a constraint-based e-service infrastructure', *Electronic Commerce Research*, Kluwer Academic Publishers, Vol. 3, pp.9–24.
- Sun, H., Wang, X., Zhou, B. and Zou, P. (2003) *Research and Implementation of Dynamic Web Services Composition*, APPT 2003, LNCS 2834, Springer-Verlag Berlin Heidelberg, pp.457–466.
- Terai, K., Izumi, N. and Yamaguchi, T. (2003) *Coordinating Web Services based on Business Models*, ACM 1-58113-788-5/03/09, ICEC, Pittsburgh, PA.
- Tosic, V., Esfandiari, B., Pagurek, B. and Patel, K. (2002b) Bussler, C. *et al.* (Eds.): *On Requirements for Ontologies in Management of Web Services*, WES, LNCS 2512, Springer-Verlag Berlin Heidelberg, pp.237–247.
- Tosic, V., Patel, K. and Pagurek, B. (2002a) Bussler, C. *et al.* (Eds.): *WSOL – Web Service Offerings Language*, WES, LNCS 2512, Springer-Verlag Berlin Heidelberg, pp.57–67.
- Velasco, J.R. and Castillo, S.F. (2003) Bauknecht, K., Tjoa, A.M. and Quirchmayr, G. (Eds.): *Mobile Agents for Web Service Composition*, EC-Web, LNCS 2738, Springer-Verlag Berlin Heidelberg, pp.135–144.
- Virdell, M. (2003) *Business Processes and Workflow in the Web Services World*, <http://www-106.ibm.com/developerworks/webservices/library/ws-work.html>, January.
- Wang, Q., Yuan, Y., Zhou, J. and Zhou, A. (2003) Dong, G. *et al.* (Eds.): *Peer-serv: A Framework of Web Service in Peer-to-Peer Environment*, WAIM, LNCS 2762, Springer-Verlag Berlin Heidelberg, pp.298–305.
- Wohed, P., van der Aalst, W.M.P., Dumas, M. and ter Hofstede, A.H.M. (2003) Song, L.Y. *et al.* (Eds.): *Analysis of Web Services Composition Languages: The Case of BPEL4WS*, ER, LNCS 2813, Springer-Verlag Berlin Heidelberg, pp.200–215.
- Wu, D., Parsia, B., Sirin, E., Hendler, J. and Nau, D. (2003) *Automating DAML-S Web Services Composition Using SHOP-2*.
- Zeng, L., Benatallah, B., Dumas, M., Kalagnanam, J. and Sheng, Q.Z. (2003) *Quality Driven Web Services Composition*, ACM 1-58113-680-3/03/0005, WWW, May 20–24, Budapest, Hungary.