Scalable Computing: Practice and Experience
Volume 12, Number 3, pp. 293–305. http://www.scpe.org
ISSN 1895-1767

SUPPORT OF SEMANTIC INTEROPERABILITY IN A SERVICE-BASED BUSINESS COLLABORATION PLATFORM

KAROL FURDÍK[†,‡], PETER BEDNÁR[†], GABRIEL LUKÁČ[†], AND CHRISTOPH FRITSCH[§]

Abstract. The paper describes a system prototype that is aiming at the provision of an environment for flexible project-oriented collaboration of networked enterprises, so-called virtual business alliances. The system, which was developed within the FP7 project SPIKE, employs the technology of semantically enhanced service bus, supported by underlying semantic structures such as ontologies and abstract business process models. The focus of this approach is to achieve effective interoperability of possibly heterogeneous services provided and consumed by members of the collaboration environment. The architecture of the main system modules is presented and the processing of semantically annotated services orchestrated in a collaboration workflow is explained in more details.

Key words: semantic annotation of services, business process ontologies, semantic interoperability, networked enterprises

AMS subject classifications. 94-04, 94B99

1. Introduction. Recent achievements in information and communication technologies (ICT) towards cross-network communication, cloud computing, service-based architectures and standardised interfaces bring new opportunities and challenges in many areas. When applied to the e-Business domain, these advanced technologies can improve flexibility and adaptability of business collaboration in so-called networked or cloud enterprises [26], which can be established as a temporal inter-organisational business alliance of enterprises and organisations co-operating on a well-defined project. Such a collaboration is usually characterised by rapidly varying business environments and requires proper technological background enabling interoperable provision and consumption of services between alliance members. Moreover, business processes in this type of networked enterprises need to be defined, structured and maintained on the alliance level as an inter-company composition of particular processes and services provided or consumed by the alliance participants in an interoperable manner.

The concept of interoperability in its technical, organisational, and semantic aspects was identified and emphasised in numerous initiatives and reports as a crucial, cross-cutting task in e-Business field [25] and other related areas. Semantic interoperability, which will be specifically addressed in next sections of this paper, refers to seamless service invocation, communication, and information exchange in an ICT environment of e-Business solutions based on principles of Service Oriented Architectures (SOA). According to this approach, services, which are formally described (i.e. annotated) by concepts of a standardised and shared knowledge base, can be provided, accessed, orchestrated, invoked, executed, and used in a flexible manner. Inputs, outputs, and other characteristics of possibly heterogeneous services can be semantically matched and integrated, enabling composition of services into customisable workflow structures. Moreover, the advanced technology of *Enterprise Service Bus* (ESB) [9] can be combined with the underlying semantic infrastructure and be employed to mediate potential incompatibilities of communicating services and applications, orchestrate their interactions, and make the integrated services available for broad access and re-use [20].

Various approaches for ICT solutions that enable networked enterprises by means of semantically enhanced ESB were designed and proposed, namely in several European research and development projects integrated in the FInES cluster [26]. Some of the most relevant projects, together with related approaches and technology frameworks, are discussed in the following subsection. In this context, we will present a specific approach for achieving semantically interoperable services, which was designed and adopted in the European project SPIKE (*Secure Process-oriented Integrative Service Infrastructure for Networked Enterprises*, FP7-217098, http://www.spike-project.eu). The project lasted from January 2008 till March 2011 and was co-founded by the European Commission within the 7th Framework Programme. The project consortium consisted of eight partners from five European countries and was coordinated by the University of Regensburg, Germany.

[†]Technical University of Košice, Faculty of Electrical Engineering and Informatics, Letná 9, 042 00 Košice, Slovakia. (Karol.Furdik@tuke.sk, Peter.Bednar@tuke.sk, Gabriel.Lukac@tuke.sk).

[‡]InterSoft, a.s., Floriánska 19, 040 01 Košice, Slovakia. (Karol.Furdik@intersoft.sk).

[§]University of Regensburg, Department of Information Systems, Universitätsstraße 31, 93053 Regensburg, Germany. (Christoph.Fritsch@wiwi.uni-regensburg.de).

SPIKE targets the interoperability of services and processes in networked enterprises, focusing on the design and development of a software platform that enables easy, fast, and secure start-up of virtual business alliances. Particular emphasis is given on semantically enhanced business processes that are capable to integrate heterogeneous services provided and consumed by alliance members, security aspects of service access, single-sign-on principles and identity federation, which is supported by the ESB-based semantic infrastructure of ontologies, services, and business processes [15]. The description of the SPIKE system, together with its principles and main outcomes, is presented in this paper. The remainder of this paper is organised as follows. Section 1.1 brings an overview of related research, including related projects and supportive technologies that were employed in SPIKE. The vision, objectives, and general approach of SPIKE are presented in Section 1.2. The following sections specifically focus on individual aspects of our designed and developed solution. Section 2 describes an overall system architecture, particular functional components and their interactions. The semantic infrastructure for service annotation and workflow specification is presented in Section 3. It includes a description of semantically enhanced ESB, means of dynamic service selection, composition, and execution in the environment of business alliances. Section 4 provides an overview of pilot applications and the results achieved during the testing of the system prototype. Finally, Section 5 summarises the project outcomes and identifies possible directions of future research and development.

1.1. Related Approaches and Technologies. The central concept of the SPIKE solution (cf. Section 1.2 and Section 3) is a semantically enhanced ESB that allows interoperability of services provided or consumed by participants of business alliance. It is based on general ESB technology, which can be seen as a messaging and communication middleware that defines standardised service interfaces and message routing for possibly heterogeneous and incompatible applications or services [9]. The implementation of the ESB technology in an organisation typically requires a specification of workflow structures that model particular processes, tasks, actions, and flow of information artefacts and messages between communicating applications. Construction of these workflow-based models is known as business process modelling and is nowadays mostly handled by the standardised BPMN (*Business Process Modelling Notation*, http://www.bpmn.org) and BPEL (*Business Process Execution Language*, [2]) notations. However, the ESB infrastructure itself and the notations for process modelling are both focusing on syntactic specification of interfaces, message exchange, and workflow structures. An effective integration of services and processes, which would be based on the meaning expressed in a machine-readable way, was identified as one of key challenges in ESB-related research [16].

The vision of semantic business process modelling, formulated in [17] and further elaborated, for example, in the European FP6 integrated project SUPER [4], aims at achieving a higher degree of automation in discovery and mediation of co-operating services [27]. The use of semantic technologies, namely Semantic Web services and underlying ontologies, for process modelling, service configuration, execution, monitoring, and analysis is envisioned as a method that can overcome the heterogeneity and incompatibility problems towards the semantically interoperable services. It may also help to reduce the human intervention throughout the life cycle of business process modelling [21]. The *Semantic Service Bus*, which can be seen as an enhancement of general ESB, makes use of semantic description of service capabilities, properties, and exchanged information artefacts, which then enables the service integration by means of automated service discovery, routing, composition and data mediation [20].

Semantically enhanced business process modelling, workflow management, and design of semantic ESB is in focus of research and standardisation organisations such as the Object Management Group (OMG, http://www.omg.org), W3C consortium (http://www.w3c.org), OASIS group (http://www.oasis-open.org), or Workflow Management Coalition (WfMC, http://www.wfmc.org). In the European context, particular solutions were provided as outcomes of several FP6 and FP7 research projects, mostly integrated in the FInES cluster initiative [26]. Some of the projects, identified as relevant and related to the SPIKE approach, are:

- STASIS (FP6-034980, http://www.stasis-project.net) provides an infrastructure for semantic mapping of services by means of a distributed peer-to-peer repository and shareable ontology structure [5];
- TrustCom (FP6-001945, http://www.eu-trustcom.com) has designed a framework for trust, security and contract management for service-based collaboration of networked enterprises;
- SUPER (FP6-026850, http://www.ip-super.org) provides a framework for semantic business process management, including generic formal languages, process models, and shareable ontology resources [4];
- COIN (FP7-216256, http://www.coin-ip.eu) targets the long-lasting enterprise collaboration, networking, and interoperability by integrating services and business processes in a generic service platform [19];

- NisB (FP7-256955, http://www.nisb-project.eu) is aiming at a provision of user-centric tools for hierarchical interoperability of enterprises, by means of designing and applying various business model archetypes and principles of dynamic business ecosystems.

The novelty of the SPIKE approach, in comparison to the above-mentioned projects, lays in the design and implementation of a light-weighted ESB framework, which is strongly supported by underlying semantic structures. The holistic and robust infrastructure, proposed in projects such as STASIS, SUPER, or COIN, was reduced in SPIKE to a rather simple and straightforward mechanism of semantic service bus, presented in more details in Section 3. From the technological perspective, the SPIKE solution is built on the WSMO framework (*Web Service Modelling Ontology*, http://www.wsmo.org), namely on its simplified variant WSMO-Lite [30]. The underlying semantic knowledge base was created in SPIKE following a generic and reusable methodology [15], which transforms user requirements, provided as textual descriptions of application cases, into a ready-to-use implementation of ontologies and business process models (cf. Section 3). In addition, a part of the knowledge base, namely the ontologies for modelling service properties and business process characteristics, was created by reusing ontologies of the SUPER project and other well-established semantic structures such as Dublin Core (http://dublincore.org), SKOS [24], etc.

Similarly as in the NisB project, a set of abstract models of sub-processes were developed in the BPMO notation [4]. WSMO Studio (http://www.wsmostudio.org) was used in SPIKE as a general toolkit for design and implementation of all the semantic representations, including ontologies, abstract business process models, and semantic annotations of services. Finally, the ESB runtime environment was designed to employ the semantic knowledge base for service mediation and orchestration in workflow sequences, directly focusing on information exchange in short-term business alliances. The results of TrustCom and STASIS projects were considered during the design of SPIKE architecture, namely in the specification of modules for business alliance maintenance and security [12].

1.2. SPIKE Objectives and General Approach. The SPIKE project primarily focused on the design, implementation, and testing of a software service platform that is capable to support a transparent, easy-to-use, and effective collaboration of organisations within an environment of virtual business alliances. At the organisational level, the project objectives were defined with respect to simplify the collaboration between members of a business alliance and to enable outsourcing of parts of the value chain to business partners (and vice versa, offering such parts in a form of services) in short-term business alliances [15]. Special focus was given on security and trust during all phases of alliance life cycle. The scientific and technology objectives of SPIKE targeted research, development, implementation and validation of software components for semantically enhanced business process management environment, which are capable to handle customised reference processes, ad-hoc defined workflow structures and distributed processes built from generic process fragments.

With respect to the defined objectives, three levels of collaboration were identified and designed as operational modes for the overall platform as follows:

- *Collaborative processes* that enable to produce physical or intangible artefacts and are modelled by means of complex workflow patterns;
- *Sharing services*, where alliance partners can offer their services in the scope of a given business process. Offered services can be retrieved, negotiated, contracted, and finally be used (i.e. invoked and consumed) by other alliance members according to the conditions specified in the service contract;
- *Identity federation*, enabling and mediating access of alliance members to the internal resources or services of other partners.

The approach adopted in the SPIKE project, aiming at a support of all defined collaboration levels, is schematically depicted in Figure 1.1. Three phases of the alliance life cycle, i.e., setting-up, running, and closing down, were addressed together with proper security settings and federation of identities enabling to access internal services or information resources of an alliance partner in a single-sign-on mode [7]. Business organisations, presented in the top bar of the Figure 1.1, may decide to form an alliance focused on a production of an artefact. A collaborative value chain, which determines particular steps and the main target of the short-time alliance, is defined as a first step. The value chain is then modelled and expressed in the Conceptual Layer in the standardised BPMN notation and can be further semantically enriched by means of the concepts from a shared knowledge model [17]. Resources and services of participating organisations can then be mediated and integrated according to known and formalised meaning. On the Service Layer, particular tasks in the process model are grounded to executable services provided by the alliance partners. It allows sharing and using the

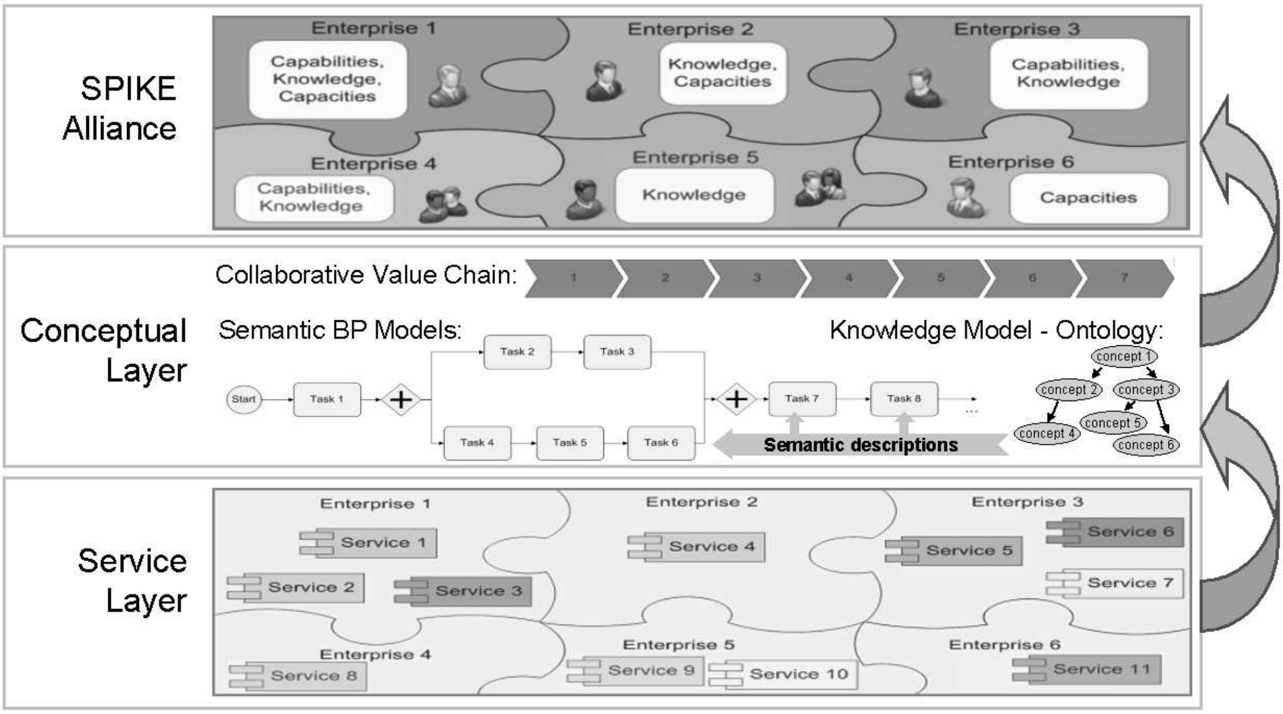

Fig. 1.1: Basic schema of the SPIKE approach for networked enterprises

services in the scope of defined processes by authorised organisations. Identities and credentials necessary for the service invocation are distributed to the authorised users in a secure way. The alliance can then operate according to a dynamic process model, which, if needed, may be modified and adapted during run time.

2. Architecture of the SPIKE Business Collaboration System. The described high-level functionality of a collaborative system for networked enterprises led to the design of the SPIKE system architecture. In accordance with the SOA principles, the architecture was proposed as highly modular and extensible [22]. The methodology of [28] was adopted to identify the viewpoints, perspectives, and stakeholders of the envisioned system. User partners of the SPIKE project, responsible for particular pilot applications (cf. Section 4), provided initial descriptions of required functionality from their perspective [31]. These descriptions have subsequently been used as a background for the specification of system views and perspectives as well as a platform for the validation of the system design.

2.1. Information View, Data Elements. The information view, as an initial step during architecture design, defines a structure of data elements and information resources that are stored and manipulated by the system. The design of data structures was accomplished by analysing the descriptions and requirements of user partners on information and data types that may be exchanged within the business alliance environment of the envisioned functionality [15]. The analysis resulted in a design of the main data elements, as presented in Figure 2.1.

The *Process*, *Workflow*, and *Task* elements are basic building blocks for modelling an alliance of collaborative business processes. The Task element, representing particular workflow actions, is further specified by the parameters such as inputs, transformations, and outputs. These parameters, consumed and produced by a task in a workflow, are represented by a set of sub-types of the generic and abstract *Resource* data element. Resource defines a common set of properties inherited by all its child data elements, in particular by the resource types such as *Document*, *Service*, *Report*, *Message*, etc. Properties of these information resources are provided as semantic metadata defined in an ontology schema (cf. Section 3). This solution enables to combine the standardised business process modelling with semantic descriptions created according to the Semantic Web principles [17]. The semantic knowledge base is represented by the *Ontology* and *Metadata* elements. These

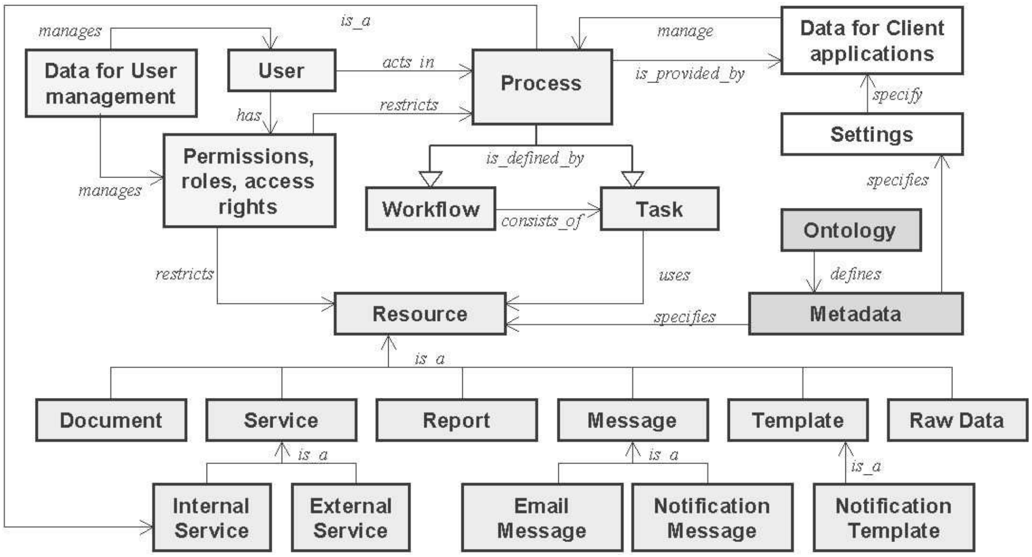

Fig. 2.1: Data elements and their structural relationships

elements store and provide both the metadata schema as well as instantiated data that specify and semantically describe the elements of other information resources. Finally, the information view contains the data elements for user management, security, authentication, and system settings needed for configuration of the client-side tools.

2.2. Functional View, Structure of System Components. The functional view of the architecture, which defines particular system modules, their inner components and interactions between them, was designed as a structure of user interfaces, data modules, and business logic of the semantic service bus [22]. Four main functional subsystems, schematically depicted in Figure 2.2, were specified as follows:

- The *SPIKE System Core* (SSC), a back-end providing functions for accessing and processing all the system data, namely the data storage, security, and maintenance of semantically enhanced business processes, workflow sequences, and services;
- The *SPIKE Portal Instance* (SPI), a web-based user interface that acts as a front-end to the SSC functionality;
- The *SPIKE Administration, Monitoring and Reporting* (SAMR), a subsystem that provides tools for overall system maintenance and day-to-day operation;
- The *SPIKE Service Bus* (SSB), an infrastructure that handles the communication between other SPIKE subsystems and external entities.

Each subsystem is further divided into a set of loosely coupled components, so-called managers, which provide autonomous and elementary functionality. The components of SSB and SSC subsystems are responsible for semantic workflow maintenance, including mediation, orchestration and execution of services in a pre-defined business process that corresponds to the alliance value chain.

Managers integrated into the SSC provide the core functionality of the overall system and are responsible for manipulations on the service level. The *Content Manager*, a very central component of the SSC, provides means for storage, retrieval and update of all data presented in and brokered through the SPIKE infrastructure, namely the ontologies, service registrations and descriptions, user sessions, and metadata of identity management. The *Service Manager*, which implements a built-in Web Service engine, provides discovery and execution capabilities for the services integrated into a process workflow. The *Semantic Manager* handles all functionality involved in dealing with semantic information, namely the semantic search, matching, mediation, mapping, and reasoning over semantically described data. Semantic metadata descriptions of services, sub-processes, and artefacts, which are specified as input and/or output parameters of services, are maintained and provided by the Semantic Manager by means of metadata mapping.

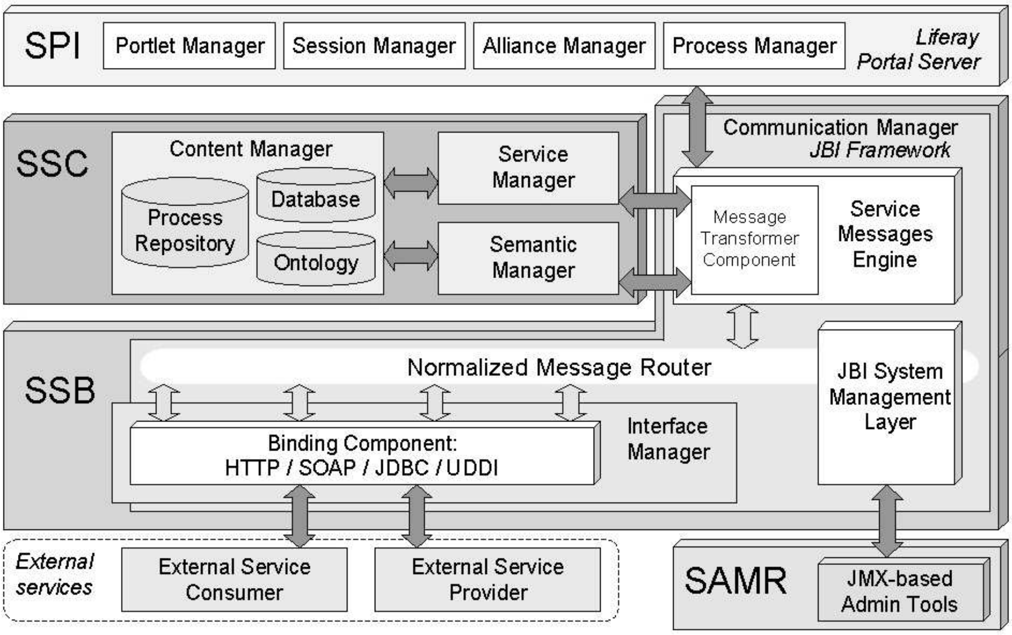

Fig. 2.2: Functional subsystems and components of the SPIKE platform

The business logic of the SSC managers is transmitted to the rest of the system by means of the SSB functionality. SSB serves as a central communication channel that handles messaging and data exchange between the system core, user interface, and administration parts of the SPIKE infrastructure. The *Interface Manager* is the only component employed by other managers to interact with external services. It provides basic capabilities for service usage, i.e. for connecting the services and transmitting service requests and responses to other components of the platform. The Binding Component acts as a proxy to remote services. It makes the services available to the service bus in a form of normalised messages, independently of the service's actual transport protocol and data format. The *Communication Manager* is an implementation of a semantically enhanced ESB, specifically designed for dynamic service selection and mediation (cf. Section 3). The service bus component is built on the Java Business Integration (JBI) specification JSR 208, using the Normalised Message Router as a central messaging backbone. The Service Messages Engine provides the lifting-lowering semantic transformations and the related business logic during the processing of external services (see also Figure 3.3 in Section 3.2). The grounded services are then orchestrated into a process workflow, exposed to the portal interface of SPI subsystem, and provided to the authorised alliance members.

3. Semantic Structures and Enhancements for Services and Business Processes. To support semantic message routing, service annotation and mediation in the environment of semantically enhanced ESB, the underlying semantic structures for the SPIKE system were built on the WSMO framework [23]. The value chain of a business alliance is semantically represented by an abstract business process model, which is implemented by the *Business Process Modelling Ontology* (BPMO) representation of WSMO [4]. The advantage of this approach is that BPMO is compatible with the standard notation of BPMN and, in addition, the underlying ontology format allows seamless and straightforward integration with other semantic elements. BPMO, as well as sBPMN, and sBPEL ontologies [18], published as outcomes of the project SUPER [4], were adopted as basic semantic structures for representing the elements of business processes in SPIKE.

The resource ontologies, which are capable to serve as a semantic base for annotating service inputs, outputs, non-functional properties, and various information artefacts exchanged between services in a workflow, were created in the format of *Web Service Modelling Language* (WSML, http://www.wsml.org). The developed resource ontologies can be divided into three logical groups [15]:

- *Process-related ontologies* provide conceptual models for semantic description of business processes and their elements such as *Process*, *Task*, *Service*, etc. Existing ontologies have been reused or new ones derived from the above-mentioned BPMO, sBPMN, and sBPEL ontologies where necessary. Moreover, this ontology group contains the concepts enabling semantic annotation of services included in a SPIKE collaboration process and referenced by the Task data elements. It covers the WSMO-Lite ontology for semantic description of Web Services and a specific SPIKE service ontology that interconnects the service-related concepts with collaboration processes by means of human tasks and available types of online services.

- *System-related ontologies* semantically describe the platform environment. It includes concepts describing non-functional properties of WSMO services, SKOS classification schemes [24], concepts representing user profiles, and general concepts such as *CollaborationObject*, *Alliance*, *Contract*, *Organisation*, *Address*, *Person*, *Actor*, etc.
- *Domain ontologies* extend the conceptual models towards particular pilot applications of the SPIKE project (cf. Section 4). The domain ontologies, created from background materials and textual descriptions provided by the SPIKE user partners [15], cover areas such as identity federation, service contracting, authorisation and secure access to distributed legacy applications, and collaboration environment of documentation services.

Both the developed abstract business process models and resource ontologies are publicly available at http://www.spike-project.eu/BPmodels/ and http://www.spike-project.eu/ontologies/, respectively.

3.1. Semantic Annotation of Services and Processes. Developed semantic structures of ontologies and abstract process models can be seen as initial steps towards an orchestrated workflow of interoperable services. To anchor an abstract process model into real services and artefacts, its activity elements such as WSMO Goal tasks, Web Service tasks, and manual tasks need to be grounded to a concrete WSDL representation of executable services [10]. The semantic interoperability can then be achieved by associating the WSDL elements with proper ontology concepts that express the meaning of inputs, outputs, and characteristics of a service in a machine-readable way.

SPIKE adopts the specification of *Semantic Annotations for WSDL and XML Schema* (SA-WSDL) [11], which is probably the best known mechanism for semantic annotation of Web Services. It defines XML attributes for linking WSDL elements to the respective ontology concepts that may semantically specify service inputs, outputs, and types. The advantage of the SPIKE approach is that the WSMO framework directly supports the SA-WSDL annotation mechanism and the respective user interface is included as part of the WSMO Studio toolkit. It thereby enables seamless integration of WSMO ontologies, BPMO models of processes, and semantically described services. Additionally, SA-WSDL attributes can specify transformations between XML messages and related ontology instances, enabling semantic data mediation by lifting and lowering procedures from XML descriptions to ontologies and vice versa. This feature was employed in SPIKE for dynamic service selection and routing as described in Section 3.2.

Web Services, and online services in general, can obtain the WSDL descriptions inherently. However, in the case of SPIKE pilot applications, most services were of off-line type, where a human interaction was required. For this type of services, referenced in SPIKE as "human tasks", the description of properties can be modelled by means of standardised *XForms* format [6], while the *BPEL4people* extension [1] can be used to model these tasks in the executable process. In accordance with this technology background, inputs and outputs of all services that represent human tasks were enhanced in SPIKE pilot applications by respective SA-WSDL references to the semantic representations of artefacts required by service inputs and/or provided on the service output. Figure 3.1 presents sample XForms representation of input entry fields, which are required for a human task that initiates a collaboration process. The entry fields of the form are associated with proper ontology concepts in advance. During the workflow run time, a human actor in the process is asked to fill in the form fields with proper values, which are then automatically linked to the respective ontology concepts.

Abstract business process models of BPMO format specify an alliance workflow, consisting of a sequence of semantically annotated tasks. Tasks can be grounded to particular services of various types, including Web Services, electronic web forms, or offline services represented as human tasks. In principle, such an abstract process model, properly grounded and semantically described, can be semi-automatically transformed into its corresponding executable BPEL form. For such a transformation, SPIKE combines the *BPMO-to-sBPEL* translation mechanism [8] with the *Eclipse BPEL Designer* toolkit (http://www.eclipse.org/bpel/).

To process the executable workflow, a specific JBI runtime environment is created in the service bus for each of the SPIKE collaboration processes. Orchestration of services into a complex workflow is handled by

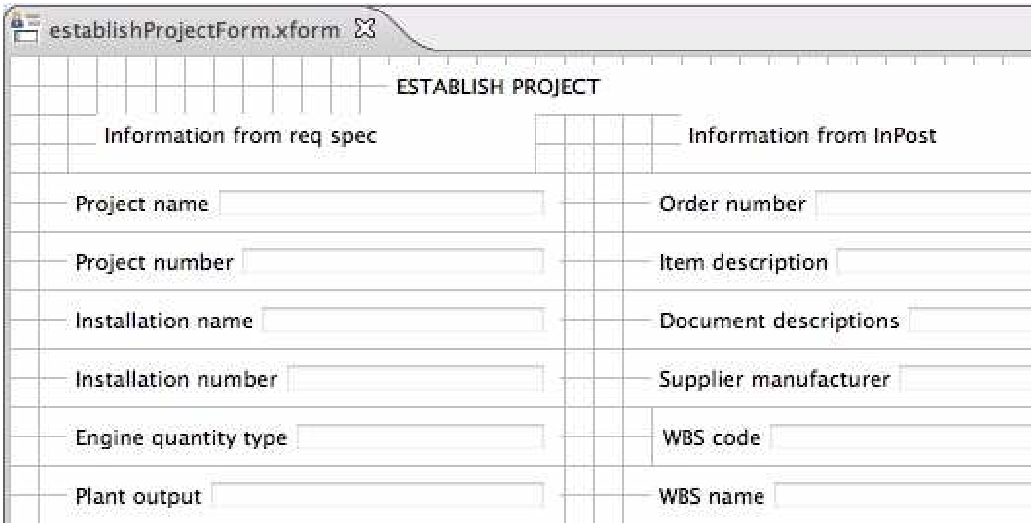

Fig. 3.1: XForms interface of input properties for a human task

the *Apache Orchestration Director Engine* (ODE, http://ode.apache.org), where a customised service resolving mechanism was implemented as an iterative selection of the best candidates for service execution, according to semantic Quality of Service properties.

3.2. Dynamic Selection and Mediation of Services in Semantic Service Bus. The presence of semantic annotations on all components of abstract business process models, i.e. on sub-processes, tasks, services, and exchanged artefacts, can facilitate the transformation to the executable workflow, as well as service interoperability, basically in two opposite modes. First, services may be linked to particular process tasks during the design time, using the correspondence between semantic descriptions of tasks and SA-WSDL attributes of services. This so-called "static service allocation", which is schematically presented on the left side of Figure 3.2, may be useful and advantageous to overcome heterogeneity problems between communicated services. The transformation of abstract models to executable BPEL processes is rather straightforward. The Invoke operation of BPEL is static; the service has to be furnished with a concrete WSDL of a service instance at the design time already. However, the role of semantic descriptions at run time is ignored or reduced to the matching of service inputs or outputs with provided or consumed artefacts. The drawback of this approach is that if such a statically allocated service is corrupted or not accessible for whatever reason at invocation time, then the running workflow is interrupted and may cause a failure of the whole process. Moreover, newly published services cannot be considered in the process model without altering and redeploying it.

The second option is so-called "dynamic service binding" [14, 13] and its schema is depicted on the right side of Figure 3.2. The *Process Layer*, which in SPIKE corresponds to the Process Manager component of SPI (see Figure 3 in Section 2.2), handles abstract process models and the respective semantic annotations of workflow tasks. Furthermore, this layer conducts the deployment of executable process representations to a workflow engine, as well as the execution and monitoring of running workflow. The *Mediation Layer*, i.e. an implementation of SSB, provides virtual interfaces (IF1-4 in Figure 3.2) for semantic descriptions of workflow tasks, which can be mapped to ontology instances and used for lifting and lowering transformations of SA-WSDL service descriptions. On the opposite side of the layer, there are binding components (BC1-4) to all available service instances that form the pool of service candidates. As a result, the Mediation Layer mediates between virtualised interfaces of workflow tasks and concrete instances of executable and available services. The *Service Implementation Layer* consists of the executable service instances, which are registered in the SPIKE service repository and properly contracted for usage within an alliance.

Dynamic service binding is implemented in SPIKE by means of the semantically enhanced ESB, as it is depicted in Figure 3.3. The process of service mediation is initiated by a service requester, which can be any stand-alone client or workflow engine. The requester sends a SOAP message containing a semantic description of required service to the message router. SSB then acts as a communication and messaging infrastructure

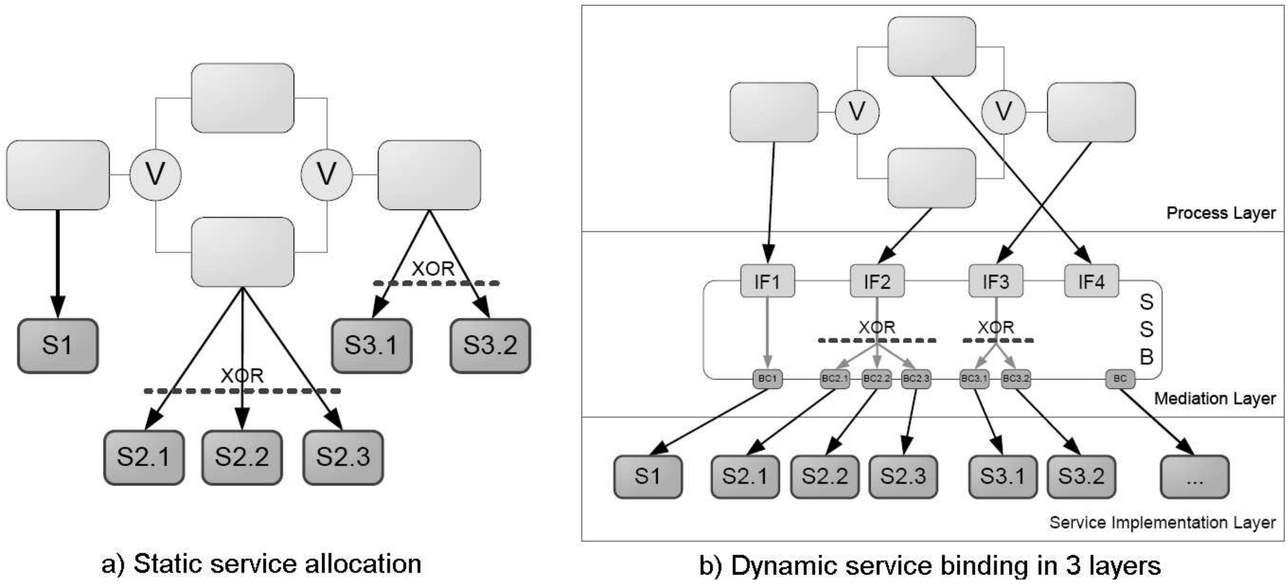

Fig. 3.2: Distinction between the static and dynamic service binding

and provides the JBI binding components for message sinks. The message router delivers the SA-WSDL virtual interface of the requested service to the Message Transformer, which forwards the interface to the lifting procedure. The virtual SA-WSDL description of the requested service contains definitions of required service properties, which are annotated with the *sawsdl:liftingSchemaMapping* attributes pointing to the instances of ontology concepts. These semantic instances are retrieved from the ontology and are used for semantic mediation and resolving of candidate services.

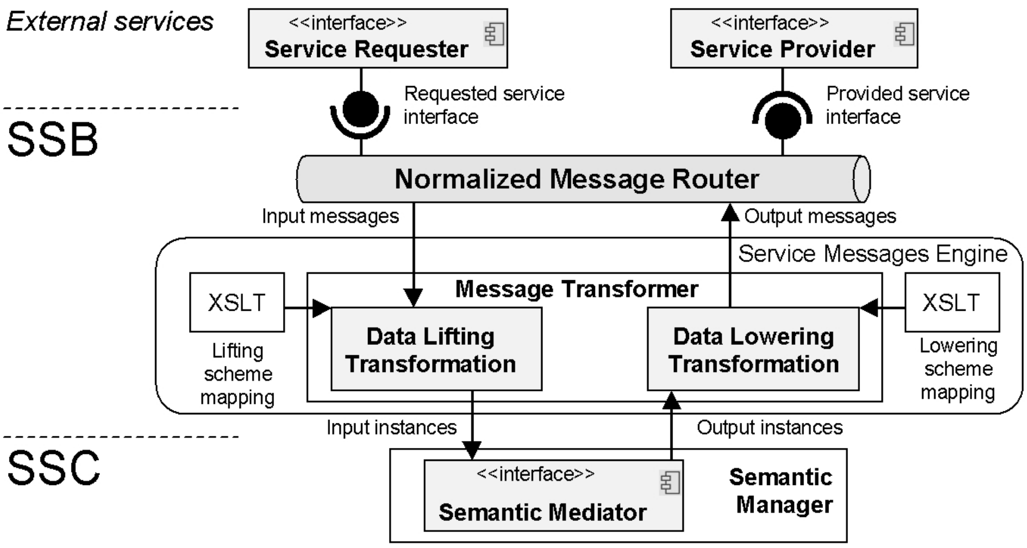

Fig. 3.3: Interactions during semantic transformation of service messages

The Semantic Manager component of SSC contains run-time implementations for ontology storage, maintenance, semantic mediation, validation and querying/reasoning over semantically described data on services and messages. In combination, it provides all required capabilities of the Mediation Layer for dynamic service binding, as presented in Figure 3.3. The service matching and resolving capability is based on semantic annotations that may be assigned to virtual service interface using the *sawsdl:modelReference* SA-WSDL attribute. The annotations can be specified by concepts of the SKOS or domain ontologies. The semantic matching procedure enables an arbitrary combination of both types of annotation. During the subsequent reasoning, the hierarchical structure of categories and subclass/superclass relations of input/output types can be recursively expanded.

The process of semantic mediation is employed to overcome possible differences between the domain ontology that annotates the virtual service interface and the ontology of resolved service candidates. As a first step of the mediation, the lifting schema is used to transform the input semantic data from the virtual interface into a set of semantic instances. If there is a difference between the ontologies annotating input and resolved services, then the instances can be transformed from the source ontology to the target ontology by means of pre-defined semantic axioms and transformation rules (instance-to-instance transformation). Instances are transformed back to normalised messages using the lowering schema specified for the resolved service.

The developed solution of semantic ESB is based on the WSMO Lite framework [30]. The *Goal* and *Web Service* objects of general WSMO, as well as its choreography and orchestration mechanisms are not used. In-memory representation of ontology elements, together with facilities for ontology validation, parsing, and serialisation, is handled by the *wsmo4j* library (http://wsmo4j.sourceforge.net). Due to performance reasons, the WSML ontologies of SPIKE knowledge base were converted to the RDF format. Physical storage of RDF ontologies, implemented as a part of the Content Manager component, uses the *Sesame repository* (http://www.openrdf.org). Mapping of top WSMO Lite ontology elements into the RDF model is based on the *ORDI framework* (http://ordi.sourceforge.net), which allows integrating various data sources and provides a common RDF API for data access. Infrastructure components such as external service interfaces, runtime environment for service selection and mediation, were built on the JBI-compliant ESB *Apache ServiceMix* (http://servicemix.apache.org).

4. Prototype Evaluation on Pilot Applications. The SPIKE system was implemented as a prototype, which is now available under the LGPL / Apache 2.0 / X11 open source license. The prototype system was tested on pilot applications run by corporate project members in Finland and Austria. The individual application cases (AC) of the pilots focused on particular aspects of collaboration in a business alliance and can be summarized as follows:

- AC1: *Information Hotel* was dealing with scenarios in collaborative manufacturing and documentation development. This AC aimed at demonstrating the intra-enterprise collaboration on a process of creating a complete set of documentation materials for an industrial system. SPIKE provided an environment for maintenance of distributed documentation services by means of secure knowledge sharing and content management. Beyond that, the overall documentation development process and its involved actors have been represented and controlled by the SPIKE system.
- AC2: *Legacy Applications*, focused on the creation and management of business alliances, including maintenance of service providers (namely, the *Siemens DAMEX*©toolkit for maintenance of business contracts), location and configuration of services, integration into a workflow, as well as tracking, contracting, and ordering of services. A particular focus of this AC was the seamless and flexible integration of legacy applications and services into the virtual business alliance.
- AC3: *Identity Federation* demonstrated the management of user identities in a networked enterprise, namely the maintenance of access rights, credentials, roles, and resources within a collaborative environment. It particularly considered aspects of cross-organisational identity federation and the involved technical and organisational procedures.

The architecture and functionality of the SPIKE platform was designed as a generic solution, which is capable to support business alliances of any domain. However, the defined ACs were taken as a framework determining the scope and focus of the sample data for the system prototype. Namely, the semantic structures of ontologies and abstract process models, presented in Section 3, were created specifically for the mentioned ACs. Nevertheless, the solution of a semantically enhanced ESB, together with related mechanisms for semantic annotation and dynamic service binding, is rather general and can be adapted to a SPIKE application in any domain.

Design, implementation, and testing of the SPIKE system were performed in two rounds. The first phase of pilots testing mostly focused on proof-of-concepts implementations of research ideas and resulting functionality of particular components. The first trial phase was then accomplished in autumn 2009. This version was built upon the *Intalio BPMS Community Edition solution* (http://community.intalio.com), which served as an environment for maintenance of executable processes. However, compatibility and licensing problems were encountered during the evaluation of this approach in the first round of pilots. For that reason we decided to substitute the Intalio environment in SPIKE by our own implementation of the Human Task Environment (cf. Section 3.1). The resulting solution enables an integrated management of processes, services, actors, and resources of virtual business alliances and was developed by employing the Eclipse BPEL Designer tool and Liferay portlet environment (http://www.liferay.com). This user side tool for maintenance of business alliances, referenced as Alliance Manager of SPI (see Figure 2.2 in Section 2.2), was functionally connected with the semantic ESB and other inner SSB/SSC components. In addition, the security environment, developed upon the Shibboleth framework (http://shibboleth.internet2.edu), was integrated to the SPIKE system to enable cross-company single sign-on and federation of identities between various external online services [12]. The second phase of SPIKE implementation was completed in winter 2010. The resulting system prototype was tested in January and February 2011 in the second trial of pilot applications AC1-AC3. Testing results, which are layed out in detail in [3], indicate the suitability of the overall SPIKE approach and prove the prototype functionality for service interoperability, secure and flexible collaboration in an environment of business alliances.

In more detail, we have learned during the trial runs that technological answers alone, even if they prove to be very well applicable, only slightly advance virtual business alliances. Additional aspects such as "ease of setup" and, even more important, "ease of use" will be among the non-technical key factors for the success of any kind of software that aims at supporting virtual business alliances. Similarly important is the consideration of organisational and social aspects (our insights can be looked up in [3] and several other SPIKE publications). Nevertheless, the trial results regarding the technology produced some meaningful findings. One major goal was to investigate the maturity and possibilities of semantic technology combined with portal and service bus architecture to resolve issues in managing technical information and technical information flow in networked enterprise. Our results show that there are plenty of possibilities in software for collaboration platform, in underlying technology and especially in semantic modelling of business processes. However, commercial phase and maturity of this technology is still quite far away from the level that it could be taken into use without extensive support and error-situation and reliability improvements. This is especially due to large number of involved users and their low IT competence. According to our experiences, the SPIKE platform has satisfactory potential to improve information management efficiency in networked enterprises even if final estimations are difficult. A tentative return on investment (ROI) calculation for AC1 indicates that in the current business process 20% cost savings could be made with 45% deployment rate in the customer base of the industry partner for AC1.

Beyond that, the trial results show that grasping the concept and benefits resulting from semantic annotations is still quite difficult even for inexperienced developers, let alone regular fellow employees. It still requires major efforts and knowledge engineers. However, if a 'semantic basis' is created, the resulting gains are considerable. Therefore, our approach of designing the system in a 'semantic agnostic' manner, meaning that in case services and processes are not semantically annotated, components can be still used (with some limitations regarding flexibility) proved to be very convenient. This allows that services can be just gradually annotated in the cases when needed. The same proved right for the consideration of legacy services and the integration of human tasks. While 'real' web services are currently still emerging in many companies, a huge amount of legacy services and manual tasks prove their capability in day-to-day business and are therefore worth integrating in collaboration platforms as the one presented.

5. Conclusions. The presented system prototype, designed and implemented within the European FP7 project SPIKE, provides an environment for service-based and process-oriented collaboration in virtual business alliances. The solution aims at effective support of service interoperability, which was achieved by employing semantically enhanced service bus that enables both static and dynamic binding of concrete executable services to the representations of tasks in a pre-defined abstract process model. Semantic structures employed for the solution were built upon the WSMO Lite framework [30], enhanced by customised service resolving, mediation, and orchestration mechanism. This light-weighted approach brings more flexibility to the composition and maintenance of applications employing ESB for intra- or inter-enterprise service-based collaboration.

The prototype of the SPIKE system was tested in application cases covering the main aspects of collaborative processes in business alliances, namely the collaborative manufacturing, inclusion of external legacy applications, and identity federation. Achieved testing results proved the usability of the developed solution [3], namely the capability to integrate various heterogeneous services in an interoperable manner. However, further enhancements could be considered, for example, towards a more advanced support of service contracts between the alliance partners. It can be achieved by providing a support for formal operational level and service level agreements, as well as by adopting other standardised processes of service management and operation [29]. More information about the SPIKE system, including methodology materials, deliverables, and other project outcomes, is available at the project web site, http://www.spike-project.eu.

Acknowledgments. The SPIKE project was co-funded by the European Commission within the contract No. 217098. The work presented in the paper was also supported by the Slovak Grant Agency of the Ministry of Education and Academy of Science of the Slovak Republic within the 1/0042/10 Project "Methods for identification, annotation, search, access and composition of services using semantic metadata in support of selected process types".

REFERENCES

[1] A. Agraval et al, *WS-BPEL Extension for People (BPEL4People)*, Version 1.0, Active Endpoints Inc., Adobe Systems Inc., BEA Systems Inc., IBM Corporation, Oracle Inc., and SAP AG., 2007.

[2] T. Andrews, F. Curbera, H. Dholakia, Y. Goland, J. Klein, F. Leymann, K. Liu, D. Roller, D. Smith, S. Thatte, I. Trickovic, and S. Weerawarana, *Business Process Execution Language for Web Services. Ver. 1.1*, IBM developers' library, 2003.

[3] P. Bednar et al, *Deliverable 9.4: Evaluation of Trial 2 and specification for revision of the SPIKE system*, Project Deliverable, Consortium of the FP7 project SPIKE, ICT 217098, 2011.

[4] R. Belecheanu, L. Cabral, J. Domingue, W. Gaaloul, M. Hepp, A. Filipowska, M. Kaczmarek, T. Kaczmarek, J. Nitzsche, B. Norton, C. Pedrinaci, D. Roman, M. Stollberg, and S. Stein, *Deliverable 1.1: Business Process Ontology Framework*, Project Deliverable, Consortium of the FP6 project SUPER, IST 026850, 2007.

[5] D. Beneventano, N. Dahlem, S. El Haoum, A. Hahn, D. Montanari, and M. Reinelt, *Ontology-driven Semantic Mapping*, in Enterprise Interoperability III, Springer London, 2008, pp. 329–341.

[6] J. M. Boyer, *XForms 1.1*, W3C Recommendation, 20 October 2009, [on line], [cit. July 28, 2011], available at http://www.w3.org/TR/2009/REC-xforms-20091020/.

[7] C. Broser, C. Fritsch, O. Gmelch, *Towards Information Security Management in Collaborative Networks*, in Proceedings of 2010 Workshops on Database and Expert Systems Applications, DEXA conference, Bilbao, Spain, IEEE Press, 2010, pp. 359–363.

[8] L. Cabral, B. Norton, J. Nitzsche, T. van Lessen, *Deliverable 4.6: BPMO to sBPEL Two Way Translation*, Project Deliverable, Consortium of the FP6 project SUPER, IST 026850, 2008.

[9] D. A. Chappell, *Enterprise Service Bus*, O'Reilly Media, Inc., Sebastopol, CA, 2004.

[10] E. Christensen, F. Curbera, G. Meredith, S. Weerawarana, *Web Services Description Language (WSDL) 1.1*, W3C Note, 15 March 2001, [on line], [cit. July 28, 2011], available at http://www.w3.org/TR/wsdl/.

[11] J. Farrell, H. Lausen, *Semantic Annotations for WSDL and XML Schema*, W3C recommendation, 28 August 2007, [on line], [cit. July 28, 2011], available at http://www.w3.org/TR/ sawsdl/.

[12] C. Fernandez, G. Fernandez, M. A. Ramirez, J. M. Troya, *Deliverable 7.3: Implementation of Security Components for Service Bus Sub-System*, Project Deliverable, Consortium of the FP7 project SPIKE, ICT 217098, 2009.

[13] C. Fritsch, P. Bednar, G. Pernul, *DS^3I – A Dynamic Semantically Enhanced Service Selection Infrastructure*, in Proceedings of the 12th International Conference on E-Commerce and Web Technologies, EC-Web 2011, Toulouse, France, August/September 2011, Lecture Notes in Business Information Processing, Volume 85, Springer-Verlag Berlin Heidelberg, 2011, pp. 13–24.

[14] C. Fritsch, G. Pernul, *Security for Dynamic Service-Oriented eCollaboration - Architectural Alternatives and Proposed Solution*, in Proceedings of the 7th International Conference on Trust, Privacy and Security in Digital Business (TrustBus 2010), Springer-Verlag Berlin Heidelberg, 2010, pp. 214–226.

[15] K. Furdik, M. Mach, T. Sabol, *Towards semantic modelling of business processes for networked enterprises*, in Proceedings of the 10th International Conference on E-Commerce and Web Technologies, EC-Web 2009, Linz, Austria, September 2009, Lecture Notes in Computer Science, Volume 5692/2009, Springer-Verlag Berlin Heidelberg, 2009, pp. 96–107.

[16] P. Giangarra, J. DeMeester, *Enabling Network-Centric Operations with Semantic Web Technologies*, W3C submission, April 2005, [on line], [cit. July 28, 2011], available at http://www.w3.org/2005/04/FSWS/Submissions/14/Paper.pdf.

[17] M. Hepp, F. Leymann, J. Domingue, A. Wahler, and D. Fensel, *Semantic Business Process Management: A Vision Towards Using Semantic Web Services for Business Process Management*, in Proceedings of the IEEE International Conference on e-Business Engineering (ICEBE 2005), Beijing, China, IEEE Press, 2005, pp. 535–540.

[18] M. Hepp, D. Roman, *An Ontology Framework for Semantic Business Process Management*, in Proceedings of the conference Wirtschaftsinformatik 2007, Karlsruhe, Germany, 2007, pp. 423–440.

[19] S. Huber, C. Carrez, and H. Suttner, *Development of Innovative Services Enhancing Interoperability in Cross-organizational Business Processes*, in Enterprise Interoperability, Lecture Notes in Business Information Processing, Volume 76, Part 2, Springer-Verlag Berlin Heidelberg, 2011, pp. 75–88.

[20] D. Karastoyanova, B. Wetzstein, T. van Lessen, D. Wutke, J. Nitzsche, and F. Leymann, *Semantic Service Bus: Architecture and Implementation of a Next Generation Middleware*, in Proceedings of the 2nd International Workshop on Services Engineering 2007 (SEIW), ICDE Workshops, IEEE Press, 2007, pp. 347–354.

[21] D. Karastoyanova, T. van Lessen, F. Leymann, Z. Ma, J. Nitzsche, B. Wetzstein, S. Bhiri, M. Hauswirth, and M. Zaremba, *A Reference Architecture for Semantic Business Process Management Systems*, in Multikonferenz

Wirtschaftsinformatik 2008, Berlin, GITO-Verlag, 2008, pp. 1727–1738.

[22] M. Mach, P. Bednar, K. Furdik, *Support for Forming Temporal Business Alliances as Networked Enterprises*, in Proceedings of the Central European Conference on Information and Intelligent Systems (CECIIS 2009), Varazdin, Croatia, University of Zagreb, Faculty of Organisation and Informatics, 2009, pp. 179–186.

[23] M. Mach, J. Hreno, K. Furdik, *Prototype of a Platform for Business Collaboration*, in Proceedings of the Central European Conference on Information and Intelligent Systems (CECIIS 2010), Varazdin, Croatia, University of Zagreb, Faculty of Organisation and Informatics, 2010, pp. 347–354.

[24] A. Miles, S. Bechhofer, *SKOS Simple Knowledge Organization System Reference*, W3C Recommendation 18, August 2009, [on line], [cit. July 28, 2011], available at http://www.w3.org/TR/skos-reference/.

[25] Lios Geal Consultants Ltd., *E-Business Interoperability and Standards: A Cross-Sector Perspective and Outlook. Special e-Business W@tch report*, Enterprise and Industry Directorate General, European Commission, 2005.

[26] M. Missikoff, S. Drissi. R. Giesecke, A. Grilo, M. Li, N. Mehandjiev, and D. Werth, *Future Internet Enterprise Systems (FInES). Research Roadmap. Final report, ver. 3.0, June 2010*, European Communities, 2010.

[27] C. Pedrinaci, C. Brelage, T. van Lessen, J. Domingue, D. Karastoyanova, and F. Leymann, *Semantic Business Process Management: Scaling up the Management of Business Processes*, in Proceedings of the 2nd IEEE International Conference on Semantic Computing (ICSC 2008), Santa Clara, CA, IEEE Press, 2008, pp. 546–553.

[28] N. Rozanski, E. Woods, *Software Systems Architecture: Working with Stakeholders Using Viewpoints and Perspectives*, Addison Wesley, 2005.

[29] M. Sarnovsky, K. Furdik, *IT service management supported by semantic technologies*, in Proceedings of the 6th IEEE International Symposium on Applied Computational Intelligence and Informatics (SACI 2011), Timisoara, Romania, IEEE Press, 2011, pp. 205–208.

[30] T. Vitvar, J. Kopecky, D. Fensel, *WSMO-Lite: Lightweight Semantic Descriptions for Services on the Web*, WSMO Deliverable D11, Ver.0.2, DERI, 2008.

[31] S. Wiesbeck et al, *Deliverable 2.2: User requirements analysis and development/test recommendations*, Project Deliverable, Consortium of the FP7 project SPIKE, ICT 217098, 2008.

Edited by: Dana Petcu and Jose Luis Vazquez-Poletti
Received: August 1, 2011
Accepted: August 31, 2011