PROCEEDINGS

SEKE 2009

The 21st International Conference on Software Engineering & Knowledge Engineering

Sponsored by

Knowledge Systems Institute Graduate School, USA

Technical Program

July 1-3, 2009

Hyatt Harborside Hotel, Boston, Massachusetts, USA

Organized by

Knowledge Systems Institute Graduate School

Huanjing Wang, *Western Kentucky University, USA*
Christiane Gresse von Wangenheim, *Universidate do Vale do Itaja, Brazil*
Tim Weitzel, *Bamberg University, Germany*
Victor Winter, *University of Nebraska at Omaha, USA*
Guido Wirtz, *Bamberg University, Germany*
Eric Wong, *University of Texas at Dallas, USA*
Franz Wotawa, *Technische Universitaet Graz, Austria*
Haiping Xu, *University of Massachusetts Dartmouth, USA*
Chi-Lu Yang, *Institute for Information Industry, Taiwan*
Hongji Yang, *De Montfort University, UK*
Ren-Dar Yang, *Institute for Information Industry, Taiwan*
Huiqun Yu, *East China University of Science and Technology, China*
Du Zhang, *California State University Sacramento, USA*
Jing Zhang, *Motorola Research, USA*
Min-Ling Zhang, *Hohai University, China*
Zhinan Zhou, *Samsung, USA*
Hong Zhu, *Oxford Brookes University, UK*
Xingquan Zhu, *Florida Atlantic University, USA*
Eugenio Zimeo, *University of Sannio, Italy*
Andrea Zisman, *City University, UK*

PUBLICITY CO-CHAIRS

Daniel Beimborn, *Bamberg University, Germany*
S. Masoud Sadjadi, *Florida International University, USA*

South America Liasion

Jose Carlos Maldonado, *University of Sao Paulo, Brazil*

Industry Advisory Committee

Yi Deng, *Dean, School of Computer Science, Florida International University, USA*
J. S. Ke, *Senior Fellow, Institute for Information Industry, Taiwan*
A. J. Rhem, *Senior Partner, A. J. Rhem and Associates Inc., USA*

Proceedings Cover Design

Gabriel Smith, *Knowledge Systems Institute Graduate School, USA*

Conference Secretariat

Judy Pan, *Chair, Knowledge Systems Institute Graduate School, USA*
Omasan Etuwewe, *Knowledge Systems Institute Graduate School, USA*
Chen-Cheang Huang, *Knowledge Systems Institute Graduate School, USA*
Daniel Li, *Knowledge Systems Institute Graduate School, USA*

Table of Contents

Semantic Service Matchmaking in the ATM Domain Considering Infrastructure Capability Constraints

Thomas Moser, Richard Mordinyi, Wikan Danar Sunindyo, Stefan Biffl

Institute of Software Technology and Interactive Systems, Vienna University of Technology
Favoritenstrasse 9-11/188, Vienna, Austria
{thomas.moser, richard.mordinyi, wikan.sunindyo, stefan.biffl}@tuwien.ac.at

***Abstract*—In a service-oriented environment business processes (BPs) flexibly build on software services (SSs) provided by systems in a network. A key design challenge is the semantic matchmaking of BPs and SSs in two steps: 1. Find for one BP the SSs that meet or exceed the BP requirements; 2. Find for all BPs the SSs that can be implemented within the capability constraints of the underlying network, which poses a major problem since even for small scenarios the solution space is typically very large. In this paper we analyze requirements from mission-critical BPs in the Air Traffic Management (ATM) domain and introduce an approach for semi-automatic semantic matchmaking for SSs, the "System-Wide Information Sharing" (SWIS) BP integration framework. A tool-supported semantic matchmaking process like SWIS can provide system designers and integrators with a set of promising SSs candidates and therefore strongly reduces the human matching effort by focusing on a much smaller space of matchmaking candidates. We evaluate the feasibility of the SWIS approach in an industry use case from the ATM domain.**

I. INTRODUCTION

Safety-critical systems and business processes, e.g., in the Air Traffic Management (ATM) domain, have to become more flexible to implement changes due to new business environments (e.g., mergers and acquisitions), new standards and regulations. A promising approach follows the service-oriented architecture (SOA) paradigm that builds flexible new systems for business processes (BPs) based on a set of software services (SSs) provided by system nodes in a network. A key design challenge is the matchmaking of BPs and SSs, i.e., finding the SSs that a) best meet the requirements of the BPs under consideration and b) can be implemented with the available network capabilities. The solution space is typically large even for small problems and a general semantic solution to enable comprehensive tool support seems infeasible.

To provide a SOA solution for a set of BPs, meaning to identify suitable SSs for BPs, designers and system integrators need to overcome 3 integration challenges that build on each other:

1. Technical integration connects networked systems that use heterogeneous technologies, i.e., different protocols, operational platforms, etc. Current technical integration approaches like Enterprise Service Bus (ESB) [2] or Service Oriented Architecture (SOA) [14] need manual configuration on the technical detail level and tool support is typically focused on a single technology or vendor.

2. Semantic integration translates data content and format between systems that use heterogeneous semantics, i.e., different terminologies for service names, data formats, etc. For semantic integration, there is no standard or framework available, making the semantic transformations between multiple services inefficient and expensive.

3. Business process support builds on technically and semantically integrated systems that provide SSs the BP needs to fulfil its goal. The system integrator has to select SSs that really match the requirements of the BP, and check whether the infrastructure capabilities can support the communication requirements of the chosen solution.

Large BP and SS integration networks consist of hundreds of integration nodes; changes of SS properties and network capabilities make the correct and efficient identification of feasible BP and SS pairs a recurring complex and error-pro task. Current service matchmaking approaches focus on eit technical or semantic integration issues [20], while busin process support is, to our knowledge, missing. Tool sup for matchmaking of BPs and SSs need to make the req ments of BPs and SSs as well as the capabilities of SSs the underlying infrastructure understandable for machine

In previous work, we introduced a systems integratio proach, the "system-wide information sharing" (SWIS proach. The SWIS framework explicitly models the se of integration requirements and capabilities in m understandable form (semantic integration) [17]; and nectors and transformations between heterogeneous systems (technical integration) to simplify systems in (business process support) [16].

In this paper, we describe the semantic matchm BPs and SSs and the optimization of the integratio with respect to available network capabilities. Sem chmaking uses the machine-understandable SWIS describe BP and SS requirements and SS and netw ities to derive 2 results: 1. Provide sets of po each BP; 2. Optimize the set of selected SSs objectives (e.g., costs, delay) while observing of the underlying network infrastructure, a knapsack problem [11]. We evaluate the fe SWIS approach in a use case from the ATM d

The remainder of this paper is structured tion 2 summarizes related work, Section 3 search issues, while Section 4 describes the 5 elaborates the semantic service matchma the optimization of the integration solutio ates the approach and Section 7discusses gard to the research issues. Finally, Sect identifies further work.

II. RELATED WORK

This section summarizes related work on technical integration, semantic integration with semantic web services, and service matchmaking with multi-objective optimization.

A. Technical Integration

Technical system integration is the task to combine networked systems that use heterogeneous technologies to appear one big system. There are several levels at which system integration could be performed [1], but there is so far no standardized integration process that explains how to integrate in general.

need for integration over heterogeneous middleware gies with different APIs, transportation capabilities, rk architecture styles implies either solutions like nd SOA [14] or the development of static and thereible wrappers between each combination of midchnologies, and thus increases the complexity of tion.

ic Integration with Semantic Web Services

integration is solving problems originating from are data across disparate and semantically het- sources [6]. These problems include the tologies or schemas, the detection of duplicate nciliation of inconsistencies, and the model- lations in different sources [18]. One of the nd most actively studied problems in seman- establishing semantic correspondences (or vocabularies of different data sources. [3] ogies as a solution option to semantic inte- rability problems has been studied over Ontologies promise to provide machine- ntation of knowledge, while allowing certain facts as well as the derivation of ng approaches based on the modeled neral domain, semantic integration has f not unsolvable. However, in a specia- ATM domain, semantic integration identified three major dimensions of es for supporting semantic integra- mappings (semi-)automatically, the tation of these mappings, and rea-

Web Services and the need for enabling them are by now well time, recognition is growing of pecifications of Web Services, flexible automation of service construction of more powerful romote the use of semantically services. Furthermore, richer r automation of activities as tion, supply chain manage- of services. [12]

have been developing lan- pproaches for so called Semantic Web services [13]. The Ontology Web Language for Services (OWL-S), which seeks to provide the building blocks for encoding rich semantic service descriptions in a way that builds naturally upon the Web Ontology Language (OWL), supplies Web Service providers with a core set of markup language-constructs for describing the properties and capabilities of their Web Services in unambiguous, computer-interpretable form [4].. WSDL-S [15] is another approach for annotating current Web Service standards with semantic descriptions. The Web Service Modeling Ontology (WSMO) [10] is a framework for Semantic Web Services which defines a rich conceptual model for the development and the description of Web Services based on two main requirements: maximal decoupling and strong mediation.

All three approaches, OWL-S, WSDL-S and WSMO, provide mechanism for semantically describing Web Services, with the major goal of allowing generic description of service functionality as well adding semantics to general service descriptions like provided/consumed messages or service bindings. This ambitious goal seems very useful for generic service descriptions; however its usage is limited in specific domains like in the ATM domain, since too specific features would complicate a generic approach too much. Therefore, we defined our own ontology-based architecture for describing the properties and features of the ATM services [17].

C. Service Matchmaking Approaches

Semantic matchmaking can be seen as major feature of semantic integration which supports designers and system integrators by providing sets of possible integration partners regarding both structural and semantic attributes. However, the relevant semantic concepts are hard to define unambiguously for general domains, thus the focus on a well-defined domain like ATM provides semantic clarity.

Kolovski et al. [8] provide a mapping of WS-Policy to OWL. WS-Policy provides a general purpose model and syntax to describe the policies of a Web service. It specifies a base set of constructs that can be used and extended by other Web service specifications to describe a broad range of service requirements and capabilities. The main advantage of representing Web Service policies using OWL is that OWL is much more expressive than WS-Policy and thus provides a framework for exploring richer policy languages. Verma et al. [20] present an approach for matching the non-functional properties of Web Services represented using WS-Policy. Oldham et al. [19] present a framework for the matching of providers and consumers based on WS-Agreements. The WS-Agreement specification defines a language and protocol for capturing relationships with agreements between two parties.

Both WS-Policy and WS-Agreement define a generic framework for the representation of standard Web Service policies, however both frameworks seem too generic to be effectively used in a concrete scenario from a specialized domain like the ATM domain is. Therefore, we used the concept of describing Service policies using a knowledge representation language like OWL, but defined our own extendable policy representation language which is better suitable for the ATM domain [17].

III. Research Issues

Recent projects with industry partners from the ATM domain raised the need for semi-automated BP integration support in technology-driven integration environments. Recently, we developed a data-driven approach [16] that explicitly models the semantics of the problem space, i.e., BP integration requirements and network infrastructure capabilities [17]; the solution space, i.e., the connectors, and data transformations between SSs. Finally, we provide a process to bridge problem and solution spaces, i.e., identify feasible BP and SSs pairs while fulfilling business requirements and optimizing the chosen integration solution according to multiple objectives.

Figure 1 provides an overview on the integration layers, data flows between the integration layers, and the steps of the semantic service matchmaking process: SM1: For each BP, identify the suitable SSs sets, which fulfil all BP service and data requirements. From these possible BP and SSs sets, the system integrators choose the most promising sets, the so-called collaboration sets. SM2: The selected collaboration sets are then optimized regarding the original infrastructure requirements of both the business BPs and the SSs, as well as the available limited capabilities of the infrastructure's nodes and links. The outcome of SM2 is an optimized configuration of the integration solution, consisting of the selected collaboration sets as well as their grounding to the underlying integration network infrastructure.

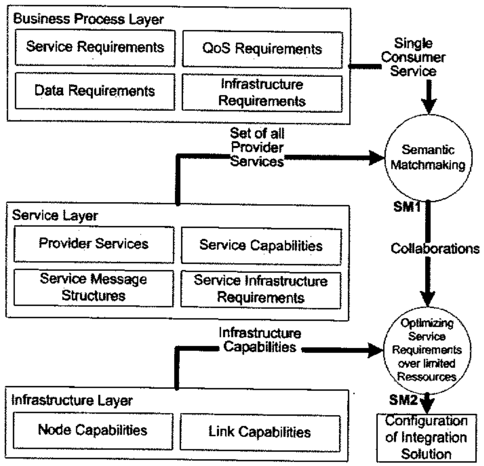

Figure 1: Semantic Service Matchmaking Process Steps.

Based on this, we derive the following research issues:

RI-1: Semantic Matchmaking of SS candidates for one BP (*SM1*). Provide machine-understandable descriptions for BP and SSs requirements as well as SS and network capabilities to provide tool support for SM1 to make the search space reduction effective (low number of false negatives and false positives) and efficient (less human effort required) co[...] to the current human-based approach.

RI-2: Resource Feasibility Check and Optimiz[...] all Collaborations (*SM2*). Provide a framework to [...] checking the validity of a set of BPs and SSs with the [...] structure capability constraints and b) ranking valid [...] by multiple optimization criteria like network cost and [...] delay.

IV. ATM Scenario Description

This section describes the integration scenario [...] ATM domain used throughout this paper. The ATM [...] (Figure 1) represents information that is typically [...] from customers/participants in workshops on requi[...] and capabilities elicitation for information systems in [...] tion domain. In safety-critical domains like ATM BP[...] tion solutions have to pass certifications before depl[...] which typical dynamic SOA solutions [2, 14] cann[...] regarding the current rigid integration network in the [...] domain designed to guarantee integration requirement[...] in case of partial failure.

In the ATM domain semantic matchmaking is an [...] scarce human experts who have to cope with a huge [...] space and often miss better solutions due to their sim[...] ristic search strategies. Tool-supported semantic match[...] provides designers and system integrators with a set of [...] ising integration partner candidates and therefore [...] reduces the human matching effort by focusing on [...] smaller space of feasible matchmaking candidates that [...] rated according to relevant optimization criteria.

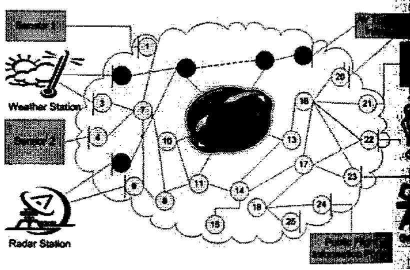

Figure 2: A Typical ATM Domain Integration Ne[...]

As shown in Figure 2, the integration network co[...] business services connected to integration network [...] Between these nodes, there may exist different kinds [...] work links using different transmission technologies [...] radio or wired transmission) as well as different mid[...] technologies for communication purposes. The capabi[...] nodes and links, like throughput, availability, reliab[...] security are explicitly modelled in order to be capabl[...] lecting suitable communication paths for particular [...] requirements, e.g., the communication link between [...] ATMIS Node and the red Node 12 represents a relia[...] secured communication path which may be requested [...] the ATMIS business service.

V. Semantic Service Matchmaking

[...]ion describes the semantic service matchmaking [...] well as the multi-objective optimization of the [...]ration services candidates.

[...]ation of Possible Collaboration Candidate Sets

[...]fication of possible collaboration candidate sets is [...] as a heuristic algorithm. Step by step, the possi[...]ation candidate sets are reduced by applying the [...]ed to the possible collaboration candidate sets. [...]c rules that are applied during the source/sink [...] described in the following paragraphs.

[...]apping. During the description of the SS mes[...] SS message segment was mapped to a domain [...]ich has been specified in the common domain [...]erefore, for all segments of the message required [...] BP, it is searched for messages of the SSs that [...]ents, which are mapped to the same domain con[...]ossible, to the same message format.

[...]licies. In addition, SSs can define requirements [...]arding preferred or unwanted SS partners, as [...] non-functional requirements, e.g., QoS require[...]ing the underlying integration network. A policy [...]n or a condition for a single collaboration or a set [...]ons, in order to allow the communication via the [...]egration network. In SWIS-based applications, [...] kinds of policies. On the one hand, there are [...] are valid for all collaborations. They specify [...]ons that need to be fulfilled by all collaborations, [...]m time for the delivery of messages. On the [...]ere are policies which are required only for a [...] of collaborations. These policies specify condi[...] to be fulfilled by the collaborations containing [...] e.g., the communication has to use only secure [...] specified set of other SSs is allowed to partici[...]boration. The SS policies that regard other SSs [...] checking whether the attributes and tags of [...] particular collaboration candidate meet the [...] defined by the BP.

[...]lation. If a message segment is mapped to [...] concept as the required message segment, [...] of the two segments differ, check whether [...] defined for the two formats. A converter is [...] format of message segments from one ba[...] different one. An explicit identifier is de[...] search for the converter at runtime (e.g., by [...]ion).

[...] Transformation. If the message segments [...] concept they are mapped to, check if a [...] consumes a segment mapped to the same [...] the segment of the message of the SS and [...] with a segment mapped to the same do[...] segment of the message of the BP.

[...] As last rule it is checked whether there is [...] between the nodes connecting the SSs and [...] the BP.

If all the rules mentioned above are successfully applied to a set of one or more SSs and a BP, then the particular set is accepted as collaboration candidate. If any of the rules cannot be met, the particular set is discarded as collaboration candidate.

B. Validity-Check and Optimization of Collaborations

Once all collaborations have been specified a Scenario is derived. A Scenario contains beside all collaborations a specification detailing how to configure the network infrastructure, so that the integration solution is optimized according to the given objectives. In the following the process steps needed to optimize the scenario is explained.

Preliminary Checks. The process step checks whether there is at least one single network route for each collaboration satisfying all global and collaboration specific policies. If this step cannot be completely satisfied the process raises an exception. The system integrator either updates or removes the collaborations which cannot be mapped to a network route, and restart the process step, or adapts the semantic infrastructure model, by adding additional nodes and links.

Route Derivation. Once it has been verified that each collaboration can be mapped to at least one route in the network, the process step derives every possible route for each collaboration. The only restrictions are that no node is allowed to appear twice within the same route and all policies have to be satisfied. The valid ones are retained; the ones violating the restrictions are removed. At the end of this process step, each collaboration will have either a single route or a set of valid routes to choose from.

Creating Scenarios. The processing step combines each route of each collaboration with each other. This means that a scenario consists of a set of collaborations where each collaboration represents exactly one route. The more scenarios are created, the higher the probability to find a scenario that is well suited for achieving the stated optimization objectives.

Evaluation. The process iterates through all scenarios and calculates their fitness according to the optimization objectives. The fitness of a scenario is the fitness of all its containing collaborations, and represents the real values (e.g. the time a message needs and the costs along the chosen route) of the objectives. The fitness represents the trade-off of the configuration, the routes of each collaboration predetermine. The set of fitness values is then analyzed according to the Pareto Front approach [5]. The Pareto Front contains either a single Scenario or a set of Scenarios. In the latter case there may be several "nearly equivalent" configurations as integration solutions. Thus, the system integrator has to decide which one to pick for practical deployment.

Multi-Objective Optimization. We have accomplished the process of optimizing collaborations by implementing a Java version of the mPOEMS approach into the SWIS framework. mPoems is an evolutionary algorithm using the concept of dominance for multi-objective optimization. The results and explanations of the approach can be found at [9].

VI. Evaluation

In this section, we evaluate the SWIS framework using a clear and comprehensible example to show the principles of our approach.

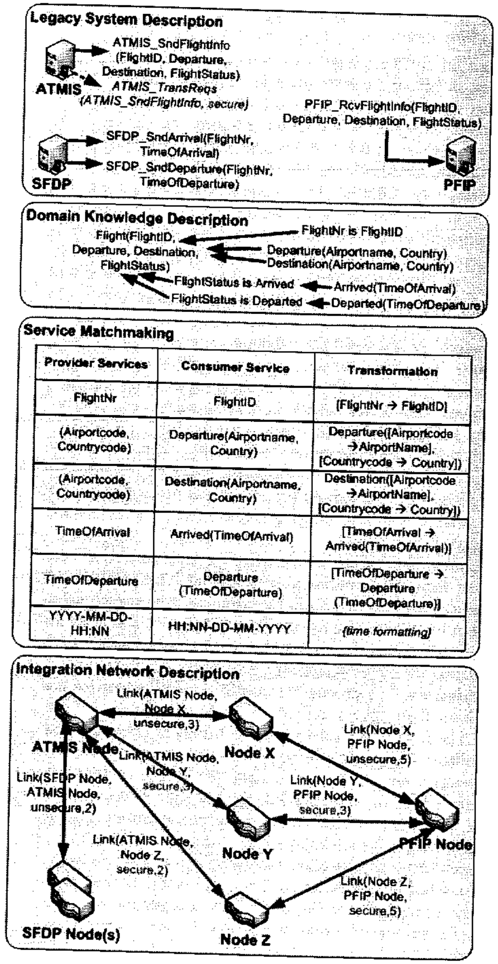

Figure 3: Service Matchmaking Example.

An example for semantic service matchmaking in the SWIS framework is shown in Figure 3. There are three services of provided by legacy systems, two provider services (*ATMIS* and *SFDP*) and one consumer service (*PFIP*). The consumer service needs information that can be obtained from the provider services, i.e. *FlightID, Departure, Destination and FlightStatus*. This needed information is provided separately by the two provider services, so the system has to find the suitable information that match with the consumer service's needs. Additionally, the service *ATMIS_TransReqs* defines a policy for the underlying integration network, stating that only secure network links may be used for the communication.

From the domain knowledge description, we know that *Flight ID* is a synonym for *Flight Number*, that *Departure* and *Arrival* are combinations of the airport code and country code of departure/arrival, and that the *FlightStatus* arrived or departed, can be derived by checking the occurrence of either *TimeOfArrival* or *TimeOfDeparture*.

Next, we calculate the network resources needed for sending messages from the *SFDP* Node to the *PFIP* Node with less capacity. From the integration network description, we can see several nodes connected by links. Each link contains information regarding source node and target node, support for secure transmissions and the transmission delay. The communication between *ATMIS* to *PFIP* needs to be done using secure connections only. There are two possible connections, either via *Node Y* or via *Node Z*. The system will choose connection via *Node Y* because it has less delay (6) than connection via *Node Z* (7).

VII. Discussion

The example shows that even for small problems the solution space is typically large. However, large BP and SS integration networks consist of hundreds of integration nodes; and changes of SS properties and network capabilities make the correct and efficient identification of feasible BP and SS pairs a recurring complex and error-prone task. By providing only sets of feasible/promising service provider and consumer candidates, semantic matchmaking supports designers and system integrators by providing sets of possible integration partners regarding both structural and semantic attributes. However, the relevant semantic concepts are hard to define unambiguously for general domains, thus the focus on a well-defined domain like ATM provides semantic clarity.

We used the concept of describing Service policies using [illegible] knowledge representation language like OWL, but defined ou[illegible] own extendable policy representation language which is bet[illegible] suitable for the ATM domain. We do not use standardiz[illegible] Web Service description frameworks because, since [illegible] strengths of Web Service description frameworks lies [illegible] generality of the approach, however their weakness is [illegible] may become complicated to describe domain-specific [illegible] For specific domains, it may be useful to use the princ[illegible] web service descriptions but tailor them to the domai[illegible] tionally, we defined our own ontology-based archit[illegible] describing the properties and features of the ATM [illegible]

We have developed a data-driven approach [16] [illegible] itly models the semantics of the problem space, [illegible] gration requirements and network infrastructure [illegible] [17]; the solution space, i.e., the connectors, an[illegible] formations between SSs. In this paper, we descri[illegible] to bridge problem and solution spaces, i.e., id[illegible] BP and SSs pairs while fulfilling business req[illegible] optimizing the chosen integration solution acco[illegible]

ple objectives. In order to evaluate the proposed process, we have derived two major research issues that will be discussed in the following paragraphs.

Semantic Matchmaking of SS candidates for one BP. Current service matchmaking approaches focus on either technical or semantic integration issues [20], while business process support is, to our knowledge, missing. In the SWIS framework, we presented a combined service matchmaking approach that performs matching based on the data of the services and available service policies regarding other services. The SWIS framework's semantic service matchmaking enables an effective search space reduction and poses lower risk and effort compared to the current human-based approaches.

Resource Feasibility Check and Optimization for all Collaborations. The optimization process steps allow using existing resources efficiently. Out of all possible collaborations for a single business process which are creatable by means of the proposed semantic matchmaking approach, only those are desirable to be deployed in the integration solution which fulfills certain criteria. Those criteria are set up by the integration expert so that existing collaborations use the un[illegible] integration network infrastructure with its limited [illegible] as efficient as possible.

[illegible]I. Conclusion And Further Work

[illegible] we presented an approach for semi-automatic [illegible]chmaking for software services (SSs), the "Sys[illegible] Information Sharing" (SWIS) Business Process (BP) [illegible]ramework. The SWIS BP integration frameworks [illegible]ne-understandable SWIS models to describe BP [illegible]ments as well as SS and network capabilities to [illegible] possible SSs for each BP. Out of these possible [illegible] integrators choose the wanted sets. These [illegible] then optimized with multiple objectives (e.g., [illegible] observing the capabilities of the underlying [illegible]

[illegible] feasibility of the SWIS approach in an [illegible] the ATM domain. The example shows [illegible] problems the solution space is typically [illegible] for large BP and SS integration net[illegible] hundreds of integration nodes. A tool[illegible] [illegible]chmaking process like SWIS can pro[illegible] [illegible] and integrators with a set of promising [illegible] therefore strongly reduces the human [illegible] on a much smaller space of mat[illegible]

[illegible] work will include a detailed de[illegible] [illegible] design to translate between matched [illegible] measuring the effectiveness and [illegible] semantic transformation with tool[illegible] [illegible] approach.

[illegible]owledgment

[illegible] acknowledge all project members [illegible] Information Sharing) project per[illegible] [illegible] University of Technology [illegible] Austro Control GmbH.

References

[1] K. Balasubramanian, A. Gokhale, G. Karsai, J. Sztipanovits, and S. Neema, "Developing Applications Using Model-Driven Design Environments," *COMPUTER*, 2006, pp. 33-40.

[2] D.A. Chappel, *Enterprise Service Bus*, O'Reilly Media, 2004.

[3] A. Doan, N.F. Noy, and A.Y. Halevy, "Introduction to the special issue on semantic integration," *SIGMOD Rec.*, vol. 33, no. 4, 2004, pp. 11-13.

[4] J. Dong, Y. Sun, and S. Yang, "OWL-S Ontology Framework Extension for Dynamic Web Service Composition," *18th Intl Conf on SE & Knowledge Engineering (SEKE'2006)*, 2006, pp. 544-549.

[5] M. Ehrgott, *Multicriteria Optimization*, Springer, 2005.

[6] A. Halevy, "Why your data won't mix," *Queue*, vol. 3, no. 8, 2005, pp. 50-58.

[7] S. Herr, K. Läufer, J. Shafaee, G.K. Thiruvathukal, and G. Wirtz, "Combining SOA and BPM Technologies for Cross-System Process Automation," 20th *Intl Conf on SE & Knowledge Engineering (SEKE'2008)*, 2008, pp. 339-344.

[8] V. Kolovski, B. Parsia, Y. Katz, and J. Hendler, "Representing Web Service Policies in OWL-DL," *4th International Semantic Web Conference (ISWC 2005)*, Springer, 2005, pp. 461-475.

[9] J. Kubalík, R. Mordinyi, and S. Biffl, "Multiobjective Prototype Optimization with Evolved Improvement Steps," *Evolutionary Computation in Combinatorial Optimization*, 2008.

[10] H. Lausen, A. Polleres, and D. Roman, "Web Service Modeling Ontology (WSMO)," *W3C Member Submission*, vol. 3, 2005.

[11] S. Martello, and P. Toth, *Knapsack problems: algorithms and computer implementations*, John Wiley & Sons, 1990.

[12] D. Martin, M. Paolucci, S. McIlraith, M. Burstein, D. McDermott, D. McGuinness, B. Parsia, T. Payne, M. Sabou, and M. Solanki, "Bringing Semantics to Web Services: The OWL-S Approach," *First International Workshop on Semantic Web Services and Web Process Composition*, Springer, 2005, pp. 26-42.

[13] S.A. McIlraith, T.C. Son, and H. Zeng, "Semantic Web Services," *IEEE INTEL. SYSTEMS*, vol. 16, no. 2, 2001, pp. 46-53.

[14] P.P. Mike, and H. Willem-Jan, "Service oriented architectures: approaches, technologies and research issues," *The VLDB Journal*, vol. 16, no. 3, 2007, pp. 389-415.

[15] J. Miller, K. Verma, P. Rajasekaran, A. Sheth, R. Aggarwal, and K. Sivashanmugam, "WSDL-S: Adding Semantics to WSDL-White Paper," 2004.

[16] T. Moser, R. Mordinyi, A. Mikula, and S. Biffl, "Efficient System Integration Using Semantic Requirements and Capability Models: An approach for integrating heterogeneous Business Services," *11th International Conference on Enterprise Information Systems (ICEIS 2009)*, 2009, accepted for publication.

[17] T. Moser, R. Mordinyi, A. Mikula, and S. Biffl, "Making Expert Knowledge Explicit to Facilitate Tool Support for Integrating Complex Information Systems in the ATM Domain," *International Conference on Complex, Intelligent and Software Intensive Systems (CISIS 2009)*, 2009, accepted for publication.

[18] N.F. Noy, "Semantic integration: a survey of ontology-based approaches," *SIGMOD Rec.*, vol. 33, no. 4, 2004, pp. 65-70.

[19] N. Oldham, K. Verma, A. Sheth, and F. Hakimpour, "Semantic WS-agreement partner selection," *15th International World Wide Web Conference*, ACM, 2006, pp. 697-706.

[20] K. Verma, R. Akkiraju, and R. Goodwin, "Semantic Matching of Web Service Policies," *2nd International Workshop on Semantic and Dynamic Web Process (SDWP 2005)*, 2005.

[21] H. Wache, T. Vögele, U. Visser, H. Stuckenschmidt, G. Schuster, H. Neumann, and S. Hübner, "Ontology-based integration of information-a survey of existing approaches," *Workshop on Ontologies and Information Sharing (IJCAI-01)*, 2001, pp. 108-117.