Semantic Business Process Management: Scaling up the Management of Business Processes

Carlos Pedrinaci and John Domingue
Knowledge Media Institute
The Open University
Milton Keynes, MK7 6AA, UK.
c.pedrinaci, j.b.domingue@open.ac.uk

Christian Brelage
SAP Research
Vincenz-Priessnitz-Str. 1, 76131 Karlsruhe
christian.brelage@sap.com

Tammo van Lessen and Dimka Karastoyanova and Frank Leymann
Institute of Architecture of Application Systems
University of Stuttgart,
Universitätsstrae 38, 70569 Stuttgart, Germany
firstname.lastname@iaas.uni-stuttgart.de

Abstract

Business Process Management (BPM) aims at supporting the whole life-cycle necessary to deploy and maintain business processes in organisations. Despite its success however, BPM suffers from a lack of automation that would support a smooth transition between the business world and the IT world. We argue that Semantic BPM, that is, the enhancement of BPM with Semantic Web Services technologies, provides further scalability to BPM by increasing the level of automation that can be achieved. We describe the particular SBPM approach developed within the SUPER project and we illustrate how it contributes to enhancing existing BPM solutions in order to achieve more flexible, dynamic and manageable business processes.

1. Introduction

The pervasive use of IT systems is nowadays visible is most environments. Medical centres, book shops, logistic companies, even the small shop where you buy the newspaper, all have some sort of IT support for driving their businesses. Indeed, the degree of complexity of the IT infrastructure is closely related to the activity performed and its associated complexity, which itself depends on several aspects such as the nature of the business activity performed, its dynamism, and its scale. Historically, IT systems were mainly used for efficiently storing and retrieving information. Then, automation gradually shifted from a data-driven approach into more process-oriented support of business activities. This trend led to so-called Workflow Management Systems (WFMS) which support the automated enactment of processes. Still, despite the performance increase brought by WFMS, their lack of flexibility led to further evolution.

BPM intends to support "business processes using methods, techniques, and software to design, enact, control, and analyze operational processes involving humans, organizations, applications, documents and other sources of information" [25]. BPM results from the limitations of WFMS which mainly focus on the enactment of processes by generic engines and do not take into account the continuous adaptation and enhancement of existing processes. BPM acknowledges and aims to support the life-cycle of business processes which undoubtedly involves post-execution analysis and reengineering of process models. BPM has made more evident the difficulties for obtaining automated solutions from high-level business models, and for analyzing the execution of processes from both a technical and a business perspective [12].

The fundamental problem is that moving between the Business Level and the IT Level is hardly automated. Deriving an IT implementation from a business model is particularly challenging and it requires an important and ephemeral human effort which is expensive and prone to errors. Conversely analysing automated processes from a business perspective, e.g., calculating the economic impact of a process or the performance of departments within an organisation, is again an expensive and difficult procedure which typically requires a human in the loop. In short, the BPM life-cycle currently suffers from a lack of automation which is impeding the scalability of current businesses from a management perspective. In this paper we shall review some of the existing drawbacks by analysing each of the phases in the BPM life-cycle and we shall discuss the approach adopted within the SUPER project[1] to avoid or minimise their negative effects in an attempt to increase the applicability of BPM.

The remainder of this paper is organised as follows. Section 2 is devoted to presenting Semantic Business Process Management (SBPM) as an enhancement of BPM towards achieving further scalability. The section first introduces an overall vision explaining the current limitations of BPM, whereas subsequent subsections focus on particular issues within the phases of the BPM life-cycle. Finally, Section 3 presents the main conclusions reached so far and introduces future work that we will carried out.

2. Semantic Business Process Management

So far BPM has focussed mainly on supporting the graphical definition of business processes and on the derivation of skeletal executable definitions that could automate them. From the modelling perspective, notable examples are Event-driven Process Chains (EPC) [14] and the Business Process Modelling Notation (BPMN) [21]. On the technical side, the so-called Service-Oriented Architecture and related technologies such as Web Services, WS-BPEL [20] or Message-Oriented Middleware are perhaps the main enabling technologies [13].

The BPM life-cycle is composed of four phases [25]. The first phase called Design or Modelling is concerned with the (re)design of business processes. Within this phase process modelling tooling is used to produce a process specification in such a way that some business goal(s) can be achieved in, presumably, the most efficient manner. The second phase, which is often referred to as Configuration or Deployment, is in charge of mapping the process definition previously designed into the underlying IT infrastructure (e.g., WFMS, ERP, Web Services, etc) so that it can be automated to the biggest extent. After the process has been deployed, the Execution or Enactment phase is in charge of actually executing them. Finally, the execution is analysed during the Analysis or Diagnosis phase in order to assess the well-being of the business processes, detect deviations, or identify potential improvements.

Current approaches to BPM suffer from a lack of automation that would support a smooth transition between the business world and the IT world [12]. On the one hand, current technologies only support the derivation of partial definitions of executable processes and still require an important human effort in order to obtain robust deployable solutions. On the other hand, once deployed these automated processes need to be continuously monitored, analysed, enhanced and adapted to meet evolving (business or technical) requirements and to accommodate ever-changing (business or technical) environments.

We have previously argued that the difficulties for automating the transition between both worlds is due to a lack of machine processable semantics [12]. Often business modelling is in fact approached as process modelling [10], and process modelling mainly focusses on the graphical representation of processes using modelling languages, e.g., BPMN, which cannot capture domain specific semantics. As a result, process definitions do not provide machine processable semantics that could support business practitioners in the analysis and reengineering of processes, and executable processes definitions, like WS-BPEL, are bound to inflexible syntactic definitions which pose important technical difficulties.

SBPM that is, the combination of Semantic Web and Semantic Web Services technologies with BPM, has been proposed as a solution for overcoming these problems [12]. SBPM aims at accessing the process space of an enterprise at the knowledge level so as to support reasoning about business processes, process composition, process execution, etc. SBPM builds upon the use of ontologies as a core component providing the required semantic information and enhances the composition, mediation and discovery of Web Services by applying Semantic Web Services techniques. In the remainder of this section we shall cover each of these phases of the BPM life-cycle identifying the main bottlenecks and introducing how our SBPM approach complements current practices towards a greater scalability in the management of business processes.

2.1. Design

Business process (re)design is in charge of generating the (presumably) most efficient process for achieving a set of business goals, guided by overall business strategies and constrained by existing legal regulations and commercial agreements. It is therefore a particularly complex task that needs to harmonise a diversity of aspects both from a business and, often also, from a technical perspective spanning several levels of abstraction. In reality, business process modelling is even more complex since existing processes which might influence or directly interact with the one being modelled need also to be accounted for.

Usually, business process modeling issues are addressed on a very specific level of abstraction, a typical partitioning being as follows:

- **Business Level (strategic)**: this level can be seen as the "big picture". It contains a functional breakdown of a domain and is usually subdivided into multiple layers which allow navigating through the problem

[1]This work is supported by SUPER (IST-026850) www.ip-super.org.

space. The modeling techniques are informal or ad hoc. Consequently, the models do not have clear semantics. They are mainly used to communicate about the universe of discourse in a guided manner and to explain the "what" of a business problem. Typical artifacts on this level may include large scale reference models such as the supply chain reference model or solution maps.

- **Business Level (operational)**: this level usually provides multiple layers which depict concrete processes. Modeling techniques in this layer are better defined and formalized but do not necessarily have clear execution semantics. There are countless modeling techniques (business scenario maps, EPC, BPMN, flowchart-techniques) in this space. BPMN is however recently emerging as the de-facto standard in this area.

- **Technical Level (processes)**: this level is the first level for which well defined semantics are provided since the models are directly executable by workflow engines. The Business Process Execution Language (BPEL) [20] is currently the de-facto standard in this layer.

- **Technical Level (services)**: this layer is concerned with the actual implementation of processes decomposed into activities that can hopefully be automated. Within this technical layer the Service-Oriented Architecture and the so-called WS-* specifications are the most applied technologies [13].

- **Technical Level (implementation)**: this level comprises actual implementation artifacts, i.e., business functions.

Frequently, each layer contains artifacts which do not have representations on lower levels of abstraction. For instance, BPMN models may contain purely conceptual parts that are not implemented by the corresponding workflow. This is typically the case when manual processes are integrated in the model. Thus, there is no direct or complete translation from models on a higher level to a lower one. Moreover, it is important to note that the separation of these layers is neither strict nor fixed. Instead, the extremes "business strategy" and "implementation" represent the boundaries of a modeling continuum. The amount and characteristics of layers in the continuum depends on concrete needs within a particular BPM project which may often vary. Further complications arise in some domains like telecommunications where the technical details regarding a service execution (e.g., Quality of Service) are of particular relevance at the process level and even at the business levels. Hence, being able to properly correlate the data across layers can be of crucial importance.

While the current limited view is sufficient to address specific problems, a complete and holistic view of the BPM modeling space is indeed useful and ultimately required in order to avoid isolated solutions and to provide an overall view over the whole enterprise. A holistic view not only provides the means for dealing with domains like telecommunications where the layers are strongly intertwined, it also allows companies to better understand and manage their business processes. Supporting the modelling continuum, requires the use of semantic information that spans these layers of abstraction in order to be able to move back and forth across the vertical layers and even horizontally in order manage the whole enterprise. The fundamental approach we have adopted in SUPER is to represent both the business perspective and the systems perspective of enterprises using a stack of ontologies, and to use machine reasoning for navigating across the layers.

The stack of ontologies, depicted in Figure 1, builds upon the use of Web Service Modelling Ontology (WSMO) [7] as the core Semantic Web Services conceptualisation and Web Service Modelling Language (WSML) [7] as the representation language supporting the specification of Ontologies, Goals, Web Services and Mediators. The integration between the different conceptualisations is provided by the Upper-Level Process Ontology (UPO) which captures general concepts such as Process, Activity, Agent or Role which are extensively reused across the ontologies. UPO plays an integrating role by linking high-level business conceptualisations to lower level ones concerning both technical and business details. High-level business aspects are conceptualised in a set of ontologies represented in Figure 1 as the Organisational Ontologies cloud. It consists of ontologies capturing resources, organisational structure, business functions, business policies, business strategies, etc. This covers what in the literature is often referred to as enterprise modelling and in fact it builds upon previous research in this area such as TOVE [8], the Enterprise Ontology [24], and REA [9].

On lower levels of abstraction we have defined the Semantic EPC (sEPC) and Semantic BPMN (sBPMN) ontologies, which conceptualise EPCs and BPMN respectively incorporating the appropriate links to WSMO concepts. These ontologies therefore provide support for two of the main modelling notations currently used in BPM. The Business Process Modelling Ontology (BPMO) provides a common layer over both sEPC and sBPMN and links them to the rest of the ontologies from the SUPER stack. BPMO links process models to organisational information as conceptualised in the Organisational Ontologies. It is also linked to the Behavioural Reasoning Ontology (BRO) whose aim is to support the composition of processes by reasoning about their behaviour. Finally, BPMO enables the transformation of business processes modelled using different nota-

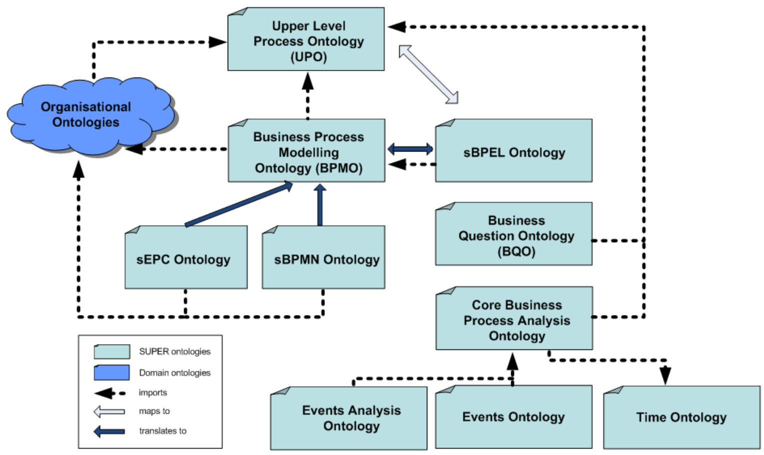

Figure 1. SUPER Ontology Stack.

tions into their executable form.

2.2. Configuration

The outcome of the design phase are typically business process models that define a control flow between activities but usually lack of specific data flow and references to particular implementations of business functions. The configuration phase is in charge of transforming this high-level model of the business process into its executable form. In traditional BPM the binding to real services and the mapping of variables to specific data types typically involves qualified IT staff to translate the conceptual model into an executable model. This task is difficult, expensive and not always long-lasting as organisations, market needs, etc. are nowadays likely to change more frequently.

The incorporation of external changes once a business process has been deployed most often requires a new iteration within the BPM life-cycle. That is, it requires redesigning the process, configuring the newly obtained design, and ensuring the execution fulfils the business requirements. Introducing changes within processes, may this be due requirements evolution or simply external changes, is therefore a relatively slow and particularly expensive process. This is mainly due to the lack of meaningful descriptions of services and business process models so that the configuration phase can be automated to a bigger extent. Defining the semantics of services and processes in terms of functional and non-functional properties can support an automated processing for a simplified translation process.

In SUPER we use WSMO as the ontological framework to describe the semantics of services in terms of their capabilities, their data models and their behavioural interface, i.e. how they need to be invoked. Such descriptions can be used for both defining services as well as for retrieving implementations of a service that match the requirements formulated by a client as a WSMO goal. Similar concepts can be used to describe the functionality required by the tasks composing a process and the functionality services provide. Finally, in order to to support the execution of business processes, we complement WSMO with BPEL [20] given its extensive support and use within the industry.

There exist several approaches within the literature for translating from conceptual models expressed in some notation like BPMN to an executable model such as BPEL. In SUPER we complement these approaches with ontologies. In particular, we support the definition of business process models by the previously introduced sEPC, sBPMN, and BPMO ontologies. The Semantic BPEL (sBPEL) ontology [18] formalises BPEL and includes additional constructs linked to WSMO so as to support the mediation between heterogeneous data or processes, or the invocation of Goals as opposed to explicitly specified Web Services. Different transformations have been defined between these different conceptualisations, see blue arrows in Figure 1. An additional transformation, although not shown in the figure, has been defined for transforming sBPEL into a serialisation format, BPEL4SWS [17], for executing processes on extensions of existing workflow engines. BPEL4SWS allows using both Web Service Description Language (WSDL) Web services and Semantic Web Service frameworks like WSMO to describe the requirements imposed on activity implementations.

We contemplate two different strategies for binding implementations to executable process models during configuration:

- **WSMO goals as activity implementations**: Using WSMO goals as activity implementations requires a middleware that implements the WSMO model following the SEE [19] reference architecture. Goals can be used to query for suitable service implementations, similar to query by example (QBE) in databases. Depending on the completeness of the formulated goal the result of such a query can be either one matching service (static binding) or a ranked list of services implementing the requested functionality. In the latter case the actual binding is shifted to the execution phase where the execution engine picks the best matching service implementation during runtime (late binding).

- **WSDL services as activity implementations**: Instead of connecting activities with WSMO Goals BPEL4SWS supports also static binding to conventional Web services. However, as a result, the flexibility of the executable process model is limited because using WSDL for describing activity implementations hampers using services that are functionally equal but implement different WSDL interfaces.

Using semantic service descriptions eases the mapping of conceptual models to their executable form. Reacting to changes in the data model or in conceptual model itself provokes still changes on the underlying executable model, however the implications of such changes are less important in SBPM as most of the changes can be derived and solved automatically, either during runtime by pushing the service selection completely to the execution phase or by guiding the implementor in the configuration phase. This way the gap between business people and IT staff can be reduced at a significant scale and can finally result in a faster time to market.

2.3. Execution

In traditional BPM the execution phase relies on syntactic and rigid process models which interact with a fixed and predefined set of partner services. This rigidity impedes very desirable features like the replacement of services based on their current state, the selection of those that better fit a certain context, etc. A typical approach is modifying the process models with somewhat artificial branches. Unfortunately with this approach, the resulting models are more complex, and adapting them to changing conditions harder.

The use of Semantic Web Services within our framework provides the appropriate flexibility in this respect. At runtime, Goals can be bound to specific Semantic Web Services selected on the basis of the existing conditions and informed by contextual knowledge which includes monitoring data. Furthermore, since services are described semantically, both functional and non-functional properties have clear semantics. This enhances the interpretation of services by humans, and more importantly, it allows data mismatches to be resolved at runtime as supported by the SEE middleware [19].

In a nutshell, the use of Semantic Web Services provides the following benefits from a process execution perspective:

- Process models are independent of the used partner services. If a particular partner is not available the SEE middleware chooses another functional equivalent service without the need of changing the process model or its deployment information.
- Process models are independent of the partner's internal data model. The process model has its own semantically annotated data model. If the partner's data model differs from that, semantic models help to bridge this gap.
- Partner services can be selected based on business aspects. Non-functional information about cost, quality of service, trust, legal constraints etc. can be taken into account so that the selected service is most suitable from a business perspective.

In SUPER, executable business process models are represented in BPEL4SWS. As it is an extension of BPEL 2.0, data flow is implicitly modelled using shared variables, declared by means of XML Schema data types. In order to take the advantages of ontologically described data models, SAWSDL [6] annotations link XML data to their semantic counterpart. This is the key enabler for semantic data mediation in data manipulation tasks. Additionally, BPEL4SWS supports evaluating logical expressions at runtime and branching conditions can be expressed in ontological terms which are presumably closer to human understanding.

In order to execute BPEL4SWS process models they are deployed to an execution engine [27] which is able to interpret them, to consume incoming messages (both XML-based and ontological instances) from services, to invoke partner services (both traditional and semantic), to perform data mediation, to evaluate logical expressions in control flow conditions and finally to emit monitoring events for each step in the execution. As we shall see in the next section, these monitoring events are also expressed in terms on an ontology so that they can later support advanced semantic analysis techniques.

2.4. Analysis

Experience shows that many factors can alter the ideal evolution of business processes (e.g., human intervention, mechanical problems, meteorological adversities, etc) and the quick adoption of special measures can mitigate to an important extent the eventual consequences. Furthermore, the competitive world we live in requires companies to adapt their processes in a faster pace. Costs pressure, offer and demand evolution, and market globalisation are examples of relevant changing conditions that force companies to have a continuous and insightful feedback on how business processes are actually being executed. Additionally, legal regulations like the Sarbanes-Oxley Act, require enterprises to show their compliance to standards. In short, there is an urgent need for deploying solutions able to adapt business processes, in a more or less timely manner, guided by the existing business objectives and organisational strategies.

One of the distinguishing characteristics of BPM solutions with respect to traditional WFMS is commonly referred to as Business Process Analysis (BPA) [25]. The main goals pursued by BPA are on the one hand the verification or validation of the execution with respect to prescribed or expected processes, and on the other hand the identification of potential improvements of business processes. The knowledge gained in this phase is thus employed for reengineering and fine tuning existing process definitions. This area therefore comprises a wide-range of fields such as Business Activity Monitoring (BAM), Business Intelligence (BI), Business Process Mining and Reverse Business Engineering. The importance of BPA is widely acknowledged and in fact all the main vendors provide their own solutions [28]. The quality and level of automation provided by these tools are rather similar and not surprisingly major efforts are devoted to presenting the information in a simple yet meaningful way better supporting humans in the analysis phase. As a consequence the state of the art in BPA represents yet another bottleneck in the management scalability of business processes.

BPA uses the logs captured by the underlying IT infrastructure such as Enterprise Resource Planning, Customer Relationship Management and WFMS to derive information concerning the well-being of business activities. Common practice within the industry is to build a Data Warehouse which consolidates all sorts of corporate information and enriches it with derived statistical data [3]. Constructing a Data Warehouse is however an expensive, delicate, and somewhat brittle process which is indeed particularly sensible to changes on the underlying IT infrastructure. Current approaches are based on a so-called Extract-Transform-Load phase which takes asynchronously data from a myriad of systems in, typically, highly heterogeneous formats and loads them into a data warehouse for further analysis. Not surprisingly one main challenge envisaged by BPA solutions regards gathering and integrating large amounts of heterogeneous yet interrelated data within a coherent whole.

Gathering and integrating corporate information in a meaningful and extensible way needs to be based on semantics as opposed to current practices based on fixed XML-based formats of relational database schemas [1]. In this respect, ontologies are particularly well-suited for defining sharable models supporting the integration of heterogeneous systems by providing a sharable conceptualisation that abstracts away from system specific details [11, 15]. Last but not least, the use ontologies brings corporate information to the knowledge level which is indeed closer to human understanding and therefore supports better business analysts in understanding and interpreting it.

Within SUPER, we have defined a stack of ontologies that capture log information and connect it to higher-level conceptual models about business processes, organisational structures, business goals and even strategic aspects [22], see Figure 1. In particular Events Ontology provides a sharable conceptual model for capturing business processes logs, whereas Core Ontology for Business pRocess Analysis (COBRA) connects it to the business level. COBRA provides a pluggable framework based on the core conceptualisations required for supporting BPA and defines the appropriate hooks for further extensions in order to cope with the wide-range of aspects involved in analysing business processes. Some extensions have already been developed for actively processing monitoring information when it is generated as well as for defining and computing business metrics automatically.

Once a Data Warehouse has been built and populated, Online Analytical Processing (OLAP) and Data Mining tools enable sophisticated data analysis that can help business analysts understand their businesses and even predict future trends. However, the semantics of the data being implicit, both OLAP and Data Mining techniques can hardly benefit from contextual knowledge about the organisation at analysis time, and strictly rely on human interpretation of the results [2]. This not only brings additional manual labour to an already complex and time consuming task, but it also prevents to an important extent the automation of certain decision making procedures and often forces developing expensive domain-specific solutions which become an additional management overhead when changes within the enterprise need to be implemented.

Important efforts have been devoted to enhancing mining techniques within SUPER [1, 2] and the use of semantic technologies have already proven their benefits. So far, the techniques developed as part of the ProM framework [26] have focussed on enhancing existing ones with semantics. This includes, a semantically enhanced Linear Temporal Logic (LTL) plugin which supports checking whether some LTL formulae hold. The enhanced version of the plugin supports benefiting from contextual knowledge within these formulae and therefore provides a more robust and stable technique since it not based any more on plain labels within the logs. For instance, a formula like "*when Activity A and B happen, C should be performed by a Manager*" can be defined independently from the existing managers or the actual configuration of the organisational structure.

Further work has been devoted to enhancing control-flow mining with abstraction and drill-down capabilities based on task hierarchies, organisational mining based on tasks similarity and performance mining which includes support for dealing with organisational structures. Although, the enhancement of the results obtained may not be groundbreaking, the integration of semantic technologies provides an outstanding improvement in terms of management scalability. Indeed, it enables the application of general purpose solutions over domain-specific data in a seamless manner, thus reducing human labour when performing analyses and minimising to an important extent the management tasks to be performed in order to cater for contextual changes.

So far we have focussed on the static aspects of BPA, but the main bottleneck in current techniques is arguably most visible within dynamic "real-time" analysis. In fact, as we discussed earlier, increasing the IT support within businesses brought on the one hand the capacity for dealing with further complexity, and on the other hand an increased automation in the activities performed. In current settings, communication is quicker, activities take less time, money flows faster and therefore decisions have to be adopted in a faster pace. Although, many decisions cannot be automated nor would entrepreneurs accept delegating them to a machine, many others can, and probably should, be addressed automatically. However, current BPA solutions hardly support this since the semantics of the data manipulated are implicit and cannot therefore support further machine processing unless domain-specific solutions are developed [28, 1, 22].

Reaching the level of automation demanded by businesses requires enhancing current analysis techniques with the capacity for reasoning over the knowledge gained by applying monitoring and mining techniques combined with pre-existing contextual knowledge about business processes and domains. At the core of our approach lays again our extensive conceptualisation of the BPM domain. Still, ontologies is not all there needs to be. In fact, reusing words from Musen "to build systems that solve real-world tasks, however, we must not only specify our conceptualizations, but also clarify how problem solving ideally will occur" [16].

Our approach to BPA [1] builds upon the research previously carried out in the context of Problem-Solving Methods (PSM) [23, 5]. PSM are intelligent software components that capture the expertise for solving knowledge intensive tasks in a domain independent manner by using ontologies as their *lingua franca*. They support reusing highly complex problem-solving expertise across domains and therefore represent an appropriate means for enhancing the state of the art in BPA while maintaining the genericity and thus the scalability of the solutions provided. So far, we have developed a generic metrics computation engine which supports the automated computation of general purpose as well as user-specified metrics in a domain independent manner. Our current research is focussing on the application of Heuristic Classification [4] to detecting and diagnosing process deviations. By doing so we expect to provide further automation within the analysis of business processes, reducing the time required for obtaining some analysis results and therefore increasing the overall manageability of business processes within entreprises by delegating more tasks for automated machine interpretation.

3. Conclusions and Future Work

The outstanding evolution of IT has influenced many activities in our quotidian life. Most business activities have seen an enormous increase in their automation by means of a more pervasive use of IT support. This phenomenon has given rise to trends like WFMS and more recently BPM which is seen as the silver-bullet for appropriately and efficiently managing business processes. However, this continuous seek for automation has recently found certain limitations in existing approaches. We are currently experiencing a limitation in the scalability for managing processes due to the need for human labour in several steps within the life-cycle of business processes. The reason for this limited scalability is mainly attributed to the lack of machine processable semantics that could support further automation of the tasks necessary to manage business processes.

Semantic BPM, that is, the extension of BPM with Semantic Web and Semantic Web Services technologies has been proposed as a means for increasing the automation of tasks through the provisioning of semantic descriptions of the artefacts involved in the life-cycle of business processes. This vision is pursued within the SUPER project which has already produced an extensive set of ontologies and tools within an overall framework that spans from methodological aspects of SBPM to the deep technical details required to orchestrate a set of Web Services which allows a business process to achieve the desired business goals. Although the project is still ongoing, very promising results have been obtained towards increasing the scalability of BPM thanks to an extensive use of semantic technologies to support the overall approach. It is worth noting however, that our aim has not been to provide scalable reasoning technologies but instead to use semantic technologies to support BPM so that business processes can be managed in a more comprehensive, simple, and thus scalable manner.

Future research will be devoted to enhancing the ontology stack, completing the integration of the tools developed so far, and refining the translations across the layers. This last task is currently being enhanced and generalised by treating the navigation through the layers as either Abstraction or Refinement as implemented within a generic heuristic classification Problem-Solving Method [4].

References

[1] A. K. Alves de Medeiros, C. Pedrinaci, W. van der Aalst, J. Domingue, M. Song, A. Rozinat, B. Norton, and L. Cabral. An Outlook on Semantic Business Process Mining and Monitoring. In *Proceedings of International IFIP Workshop On Semantic Web & Web Semantics (SWWS 2007)*, 2007.

[2] A. K. Alves de Medeiros, W. Van der Aalst, and C. Pedrinaci. Semantic Process Mining Tools: Core Building Blocks. In *16th European Conference on Information Systems*, Galway, Ireland, June 2008.

[3] S. Chaudhuri, U. Dayal, and V. Ganti. Database Technology for Decision Support Systems. *Computer*, 34(12):48–55, December 2001.

[4] W. J. Clancey. Heuristic classification. *Artificial Intelligence*, 27(3):289–350, 1985.

[5] J. Domingue, L. Cabral, S. Galizia, V. Tanasescu, A. Gugliotta, B. Norton, and C. Pedrinaci. IRS-III: A broker-based approach to semantic Web services. *Web Semantics: Science, Services and Agents on the World Wide Web*, 6(2):109–132, 2008.

[6] J. Farrell and H. Lausen. Semantic Annotations for WSDL and XML Schema. http://www.w3.org/TR/sawsdl/, January 2007. W3C Candidate Recommendation 26 January 2007.

[7] D. Fensel, H. Lausen, A. Polleres, J. de Bruijn, M. Stollberg, D. Roman, and J. Domingue. *Enabling Semantic Web Services: The Web Service Modeling Ontology*. Springer, 2007.

[8] M. S. Fox. The TOVE Project Towards a Common-Sense Model of the Enterprise. In *IEA/AIE '92: Proceedings of the 5th international conference on Industrial and engineering applications of artificial intelligence and expert systems*, pages 25–34, London, UK, 1992. Springer-Verlag.

[9] G. L. Geerts and W. E. McCarthy. An accounting object infrastructure for knowledge-based enterprise models. *IEEE Intelligent Systems*, 14(4):89–94, 1999.

[10] J. Gordijn, H. Akkermans, and H. v. Vliet. Business Modelling is not Process Modelling. In *Conceptual Modeling for E-Business and the Web, LNCS 1921 (ECOMO 2000)*, pages 40–51, Salt Lake City, USA, 2000. Springer-Verlag.

[11] T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.

[12] 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 F. C. M. Lau, H. Lei, X. Meng, and M. Wang, editors, *ICEBE*, pages 535–540. IEEE Computer Society, 2005.

[13] G. Hohpe and B. Woolf. *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2003.

[14] G. Keller, M. Nüttgens, and A.-W. Scheer. Semantische Prozeßmodellierung auf der Grundlage "Ereignisgesteuerter Prozeßketten (EPK)". Arbeitsbericht Heft 89, Institut für Wirtschaftsinformatik Universität Saarbrücken, 1992.

[15] E. Mena, A. Illarramendi, V. Kashyap, and A. Sheth. OBSERVER: An approach for query processing in global information systems based on interoperation across pre-existing ontologies. *International journal on Distributed And Parallel Databases (DAPD)*, 8(2):223–272, April 2000.

[16] M. A. Musen. Ontologies: Necessary–Indeed Essential–but Not Sufficient. *IEEE Intelligent Systems*, 19(1):77–79, January/February 2004.

[17] J. Nitzsche, T. van Lessen, D. Karastoyanova, and F. Leymann. BPEL for Semantic Web Services (BPEL4SWS). *On the Move to Meaningful Internet Systems 2007: OTM 2007 Workshops*, pages 179–188, 2007.

[18] J. Nitzsche, D. Wutke, and T. van Lessen. An Ontology for Executable Business Processes. In *Proceedings of SBPM 2007 Semantic Business Process and Product Lifecycle Management in conjunction with the 3rd European Semantic Web Conference (ESWC 2007)*, Innsbruck, Austria, June 2007.

[19] B. Norton, C. Pedrinaci, J. Domingue, and M. Zaremba. Semantic Execution Environments for Semantics-Enabled SOA. *it - Methods and Applications of Informatics and Information Technology*, Special Issue in Service-Oriented Architectures:118–121, 2008.

[20] OASIS Web Services Business Process Execution Language (WSBPEL) TC. Web services business process execution language version 2.0 committee specification. http://docs.oasis-open.org/wsbpel/2.0/CS01/wsbpel-v2.0-CS01.pdf, January 2007.

[21] Object Management Group. Business process modeling notation specification - final adopted specification. Available at http://www.bpmn.org, February 2006.

[22] C. Pedrinaci, J. Domingue, and A. K. Alves de Medeiros. A Core Ontology for Business Process Analysis. In *5th European Semantic Web Conference*, 2008.

[23] R. Studer, R. Benjamins, and D. Fensel. Knowledge Engineering: Principles and Methods. *Data Knowledge Engineering*, 25(1-2):161–197, 1998.

[24] M. Uschold, M. King, S. Moralee, and Y. Zorgios. The Enterprise Ontology. *Knowledge Engineering Review*, 13(1):31–89, 1998.

[25] W. M. P. van der Aalst, A. H. M. ter Hofstede, and M. Weske. Business Process Management: A Survey. In W. M. P. van der Aalst, A. H. M. ter Hofstede, and M. Weske, editors, *Business Process Management*, volume 2678 of *LNCS*, pages 1–12. Springer, 2003.

[26] B. van Dongen, A. K. Alves de Medeiros, H. Verbeek, A. Weijters, and W. van der Aalst. The ProM Framework: A New Era in Process Mining Tool Support. In *Applications and Theory of Petri Nets 2005. 26th International Conference, ICATPN 2005*, pages 444–454, Miami, USA, June 2005. Springer-Verlag.

[27] T. van Lessen, J. Nitzsche, M. Dimitrov, D. Karastoyanova, M. Konstantinov, and L. Cekov. An Execution engine for BPEL4SWS. In *2nd Workshop on Business Oriented Aspects concerning Semantics and Methodologies in Service-oriented Computing (SeMSoc) in conjunction with ICSOC, To appear*, Sept. 2007. Vienna, Austria.

[28] H. J. Watson and B. H. Wixom. The Current State of Business Intelligence. *Computer*, 40(9):96–99, 2007.