Strategy-driven Business Process Analysis

Carlos Pedrinaci[1], Ivan Markovic[2], Florian Hasibether[2], John Domingue[1]

[1] Knowledge Media Institute, The Open University, Milton Keynes, MK7 6AA, UK
{c.pedrinaci,j.b.domingue}@open.ac.uk
[2] SAP Research, Karlsruhe, Germany
{ivan.markovic,florian.hasibether}@sap.com

Abstract. Business Process Analysis (BPA) aims to verify, validate, and identify potential improvements for business processes. Despite the wide range of technologies developed so far, the large amount of information that needs to be integrated and processed, as well as the quantity of data that has to be produced and presented still poses important challenges both from a processing and presentation perspectives. We argue that to enhance BPA, semantics have to be the core backbone in order to better support the application of analysis techniques on the first hand, and to guide the computation and presentation of the results on the other hand. We propose a knowledge-based approach to supporting strategy-driven BPA by making use of a comprehensive and extensible ontological framework capturing from high-level strategic concerns down to lower-level monitoring information. We describe how corporate strategies can be operationalized into concrete analysis that can guide the evaluation of organisational processes, structure the presentation of results obtained and better help assess the well-being of corporate business processes.

Key words: Semantic Business Process Management, Business Process Analysis, Strategic Analysis

1 Introduction

In the business world the maxim "*if you can't measure it, you can't manage it*" is often used. Although it is too blunt a statement, it captures an important essence in current management approaches which try to maximise the aspects measured in order to evaluate, compare, and control the evolution of businesses. For instance the Balanced Scorecard is a popular "*set of measures that gives top managers a fast but comprehensive view of the business*" [1]. In a nutshell, the Balanced Scorecard defines four perspectives and suggests for each of them a set of aspects that managers should focus on. Assessing how well a company is doing is then a matter of calculating metrics and contrasting them with respect to pre-established goals for each of these key aspects.

Analysing business processes in an effective manner requires computing metrics that can help determining the health of business activities and thus the whole enterprise. However, this is not all there needs to be done. Aspects like client's satisfaction, whether a certain strategy will work out or not, how successful the research department is, or what would happen if we make a certain change in a process cannot "simply" be measured. Similarly, detecting, or better yet, anticipating process deviations with respect to expected behaviours can hardly be approached as a simple measurement problem. The closer we get to strategic analysis, the more impact analysis results are likely to have, but the more complex analysis techniques are required in order to deal with qualitative aspects, approximations, and uncertainty. In order to deal with these scenarios, BPA solutions need to apply advanced analysis techniques.

Semantic BPM, that is, the extension of Business Process Management (BPM) with Semantic Web and Semantic Web Services technologies has been proposed as a means for increasing the level of automation during the life-cycle of business processes [2]. This vision is pursued within the SUPER project[1] as part of a SBPM framework aiming at increasing the level of automation and enhancing the support for managing business processes by using comprehensive semantic models of the BPM domain. As part of this initiative, we advocate the use of a holistic integrated view of the enterprise spanning from high-level strategic concerns down to low-level operational details, in order to increase the level of automation in BPA and to better support the continuous improvement of strategic decision-making processes. In this paper, we focus on the conceptual models for capturing strategies and how they can be used for enhancing the analysis of business processes thanks to a formal and explicit modelling of the relations between strategies and operational aspects.

The remainder of the paper is organised as follows. We first introduce Core Ontology for Business Process Analysis (COBRA), which provides the core foundational conceptualisation for our work. Section 3 presents an extension of COBRA towards capturing different kinds of operational analyses that typically take place within BPM systems. Conversely, Section 4 describes Business Motivations Ontology (BMO) another extension of COBRA in this case focussed on strategic concerns and their operationalisation into outcomes that can be evaluated. Section 5 describes in more detail how our conceptual framework can support and guide the assessment of the well-being of enterprises by evaluating the desired outcomes that have strategic importance. Finally, we contrast our research with previous work in the area and we present our main conclusions and future work.

2 COBRA

Supporting the level of automation demanded by enterprises nowadays requires enhancing BPA with support for applying general purpose analysis techniques over specific domains in a way that allows analysts to use their particular terminology and existing knowledge about their domain. To this end we have defined Core Ontology for Business pRocess Analysis [3]. COBRA provides a core terminology for supporting BPA where analysts can map knowledge about some

[1] http://www.ip-super.org

particular domain of interest in order to carry out their analyses. It is worth noting that COBRA does not aim to provide a fully-comprehensive conceptualisation for supporting each and every kind of analysis since the scope would simply be too big to be tackled appropriately in one ontology. Instead COBRA, depicted in Figure 1, 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.

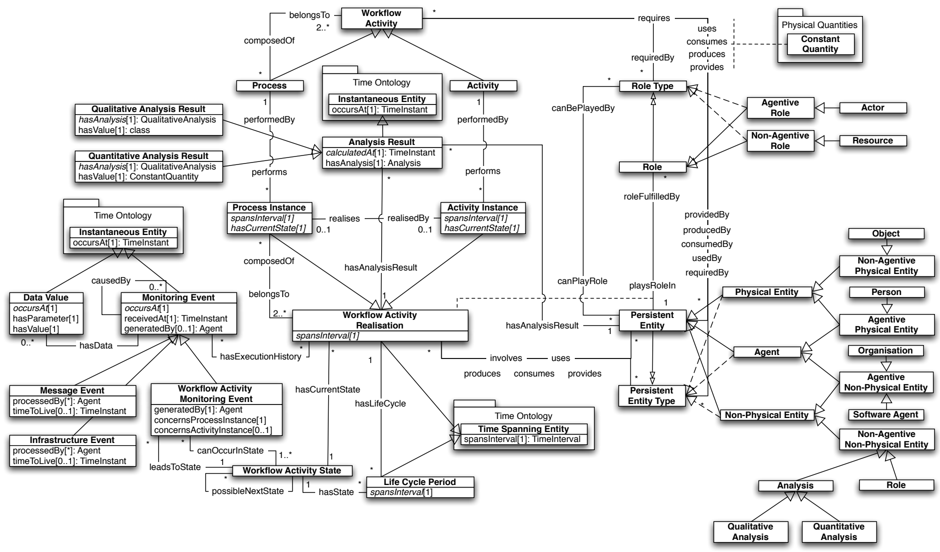

Fig. 1. Core Ontology for Business pRocess Analysis.

COBRA builds upon Time Ontology that provides a temporal reference by means of which one can determine temporal relations between elements. COBRA provides a lightweight foundational basis in order to ensure a coherent view among additional extensions. It defines concepts such as *Process*, *Activity*, *Process Instance*, *Role*, *Resource*, *Organisation* or *Person* which are to be refined within specific ontologies as defined within SUPER, or other approaches like the Enterprise Ontology [4] or TOVE [5].

COBRA has been extended with a reference Events Ontology (EVO) [3] that provides a set of definitions suitable to capture monitoring logs from a large variety of systems and ready to be integrated within our core ontology for analysing business processes. It is based on existing syntactic formats, e.g., MXML [6] or the Audit Trail Format by the Workflow Management Coalition [7] which therefore confers on it the ability to capture logs generated by a plethora of systems. As prescribed by COBRA, EVO is centred around a state model that accounts for the status of processes and activities. The state model has been captured ontologically and enhanced with additional relations, , see [3] for further details.

Concerning the analyses themselves such as metrics, the previously presented version of COBRA [3] solely captures the concept *Analysis Result*, a Temporal Entity, which has two disjoint sub-concepts: *Qualitative Analysis Result* and *Quantitative Analysis Result.* As part of our work on strategy-driven analysis as well as metrics definition and computation, we have slightly extended COBRA itself. First, COBRA includes support for units of measure and their manipulation. Secondly, we have introduced the concept *Analysis*, which is refined into *Qualitative Analysis* (e.g., "is the process critical?") and *Quantitative Analysis* (e.g., "process execution time") based on the type of Analysis Result they produce. This provides us the means for maintaining a library of Analysis specifications (e.g., metrics, time series, etc.), and it allows us to distinguish between the analysis themselves and the actual results. Indeed, the relationship between Analysis and Analysis Result has also been captured, in such a way that every Analysis Result is a result for a particular Analysis, and every Analysis may have several Analysis Results. Hence we can obtain all the results for a particular analysis, track its evolution over time, apply time series analysis, etc.

3 Operational Analysis Ontology

Operational Analysis Ontology is an extension of COBRA focussed on supporting operational analysis, that is the analysis over operational aspects that can directly be computed. More abstract concerns such as strategic analysis being captured in another ontology that will be described in Section 4.

Operational Analysis Ontology, depicted in Figure 2, is mainly based on the notions of Analysis and Analysis Result defined in COBRA. In particular it defines the main kinds of analyses over operational aspects and their respective results. Central to Operational Analysis Ontology is therefore the concept of Metric which supports the definition of operational measures used for monitoring the well-being of business processes and enterprises. The concrete definition of Metric as well as the means for computing them will be explained in more detail in Section 3.1. Additionally the ontology defines another Quantitative Analysis, namely *Quantified Quantitative Analysis Variation.* This kind of analysis basically captures the variation of the value of a certain Quantitative Analysis over a certain period (e.g., "the Process Execution Time increased by 5 ms"). Similarly, we also capture the variation qualitatively as *Quantitative Analysis Variation* for which the results are solely the *Variation Type*, e.g., *Increased*, *Strictly Decreased.*

Finally, the ontology captures two Qualitative Analyses, namely *Qualitative Analysis Comparison* and *Quantitative Analysis Comparison.* The former allows us to specify things like "is the Process a bottleneck?" whereby determining the kind of Process is a Qualitative Analysis and "bottleneck" is one of the kinds of results that can be obtained. Defining these kinds of Qualitative Analysis Comparison is based on the specification of a Qualitative Analysis, a *Logical*

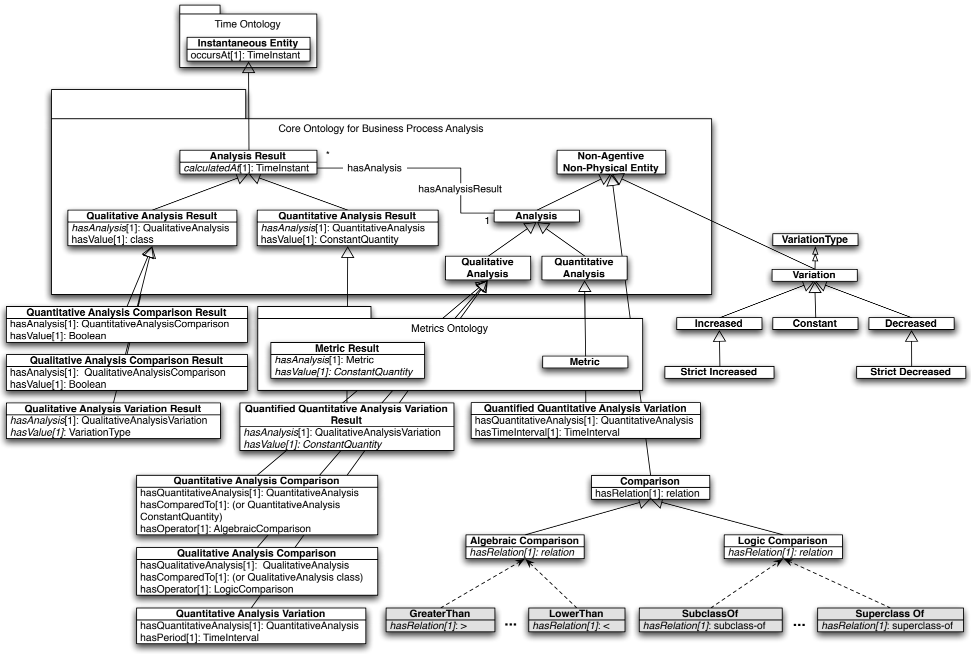

Fig. 2. Operational Analysis Ontology.

Comparison, and the value to compare to. Among the Logical Comparisons contemplated we include things like equality or subsumption relations. Conversely, Quantitative Analysis Comparison supports the definition of comparison between any Quantitative Analysis (e.g, a particular Metric) and a reference value which could itself be another Quantitative Analysis. In this way one can define things like "is the Process Execution Time > 10 ms?" or "is Process A Execution Time > Process B Execution Time?".

3.1 Metrics Ontology

Operational Analysis Ontology has been defined in rather abstract terms without explaining the mechanisms by which one could take these definitions and apply within a concrete domain. In this section we focus on Metrics Ontology which allows us to define metrics in way that can support their automated computation. The reader is referred to [8] for the details of SENTINEL, a monitoring tool that makes use of Metrics Ontology for computing and presenting metrics about the execution of business processes.

Metrics Ontology provides us with the capacity for specifying and computing metrics, as necessary for analysing and managing business processes, in a domain-independent way. On the basis of our conceptualisation we can capture kinds of metrics, e.g,. "process instance execution time", as well as specific metrics to be computed, e.g., "process instance X execution time". The former are

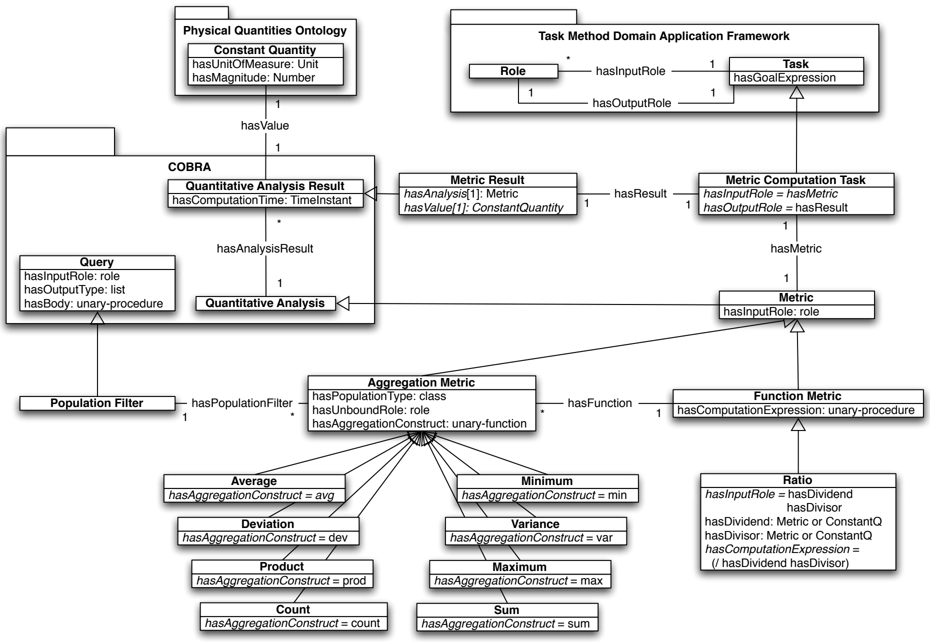

Fig. 3. Metrics Ontology.

defined as concepts, whereas the latter are modelled as instances. In this way we can provide libraries of metrics such as general purpose ones, or specific for some domain like Supply-Chain, and at analysis time the analyst can specify which of these metrics should be computed over which entities by instantiating them. This provides a convenient way for organising metric definitions and seamlessly supports the comparison of results by kind of metric, e.g., "which is the process which takes longer", as well as it allows tracking their evolution over time.

Central to Metrics Ontology is the concept *Metric* which is defined as a Quantitative Analysis (see Section 2). Metrics are specified by a set of input roles that point to domain-specific knowledge [9]. We refine Metrics into two disjoint kinds, Function Metrics and Aggregation Metrics. A Function Metric is a metric that can be evaluated over a fixed number of inputs. For example, the Metric *Process Instance Execution Time* is a Function Metric which takes as input one Process Instance. Conversely, Aggregation Metrics (e.g., "average process execution time") take an arbitrary number of individuals of the same kind as input (e.g., a set of Process Instances). Therefore, Aggregation Metrics are computed over a population in order to obtain an overall perception of some aspect of interest such as the average execution time of some particular process. The population to be processed can be defined intensionally as an ontological query so that the metric computation can focus on certain processes, or resources of interest. In this respect the use of semantic technologies plays a key role towards supporting business analysts in the analysis of processes, allowing them to use their domain-specific terminology and still use a generic machinery to process the information in a seamless way.

In order to support the automated computation of metrics, which is indeed metric dependent, each metric has a computation expression which is defined as a unary procedure. In this respect it is worth noting that the language used to define the metrics themselves as well as to develop the metrics computation engine is Operational Conceptual Modelling Language (OCML) [9]. OCML seamlessly supports the integration of static and dynamic knowledge paving the way for a rapid prototyping of a fully operational solution[2].

4 Business Motivation Ontology

The purpose of the Business Motivation Ontology (BMO) is to allow for representation of the notions important for the strategic aspect of business process analysis. BMO is inspired by the Business Motivation Model [10], a business modeling standard published by the Object Management Group (OMG). BMO is a result of conceptual reengineering, refinement and formalization of concepts described in the standard specification. There are four top level concepts that are core to the BMO: Means, End, Metric and Influencer. We visualize BMO concepts and relations in Figure 4, where the imported concepts from COBRA are marked with gray boxes.

Ends refer to any aspiration concept (cf. Figure 4, upper right). They state, *what* an enterprise wants to be. This could be about changing what the enterprise is, as well as maintaining the actual position. Ends subsume the concepts Vision (abstract End) and Desired Outcome (concrete End). A Vision describes a future, ultimate, possibly unattainable state, an enterprise as a whole wants to achieve. It is made operative by Mission and amplified by Goals. A Desired Outcome is an attainable state or target that the enterprise, or some part of it, intends to achieve. Desired Outcomes are established for a certain time interval spanning the period of time for which the Desired Outcome holds. We distinguish between outcomes that need to be achieved at the end of the period, and those that have a periodic check by means of which one can define what we refer to as *continuous* Desired Outcomes. Continuous Desired Outcomes have the particularity that they specify outcomes that are continuously desired during a given interval and that will be checked periodically. This allows to express things like “sales should increase by 1 percent per month for the next year” as well as “increase sales by 5 percent by the end of the year” in a simple and concise way.

Given the high-level of abstraction Desired Outcomes can have, it is particularly difficult to assess the level of achievement currently attained, and decide how to map these desires into concrete enterprise-wide actions and decisions. In consequence, we support the decomposition of Desired Outcomes into finer grain ones, as a gradual process that takes us from purely strategic outcomes to the operational level where one can perform measures and contrast the achieved

[2] The ontologies described herein can be found at http://www.cpedrinaci.net

results with those desired. This gradual refinement of Desired Outcomes is supported through the refinement of Desired Outcomes into *Goals* and *Objectives*, and by using what we refer to as *Logical Decomposition* constructs. Goals tend to be longer term, qualitative, and ongoing in comparison to Objectives. An Objective is a measurable, time-targeted Goal. It is a step along the way towards achieving a Goal.

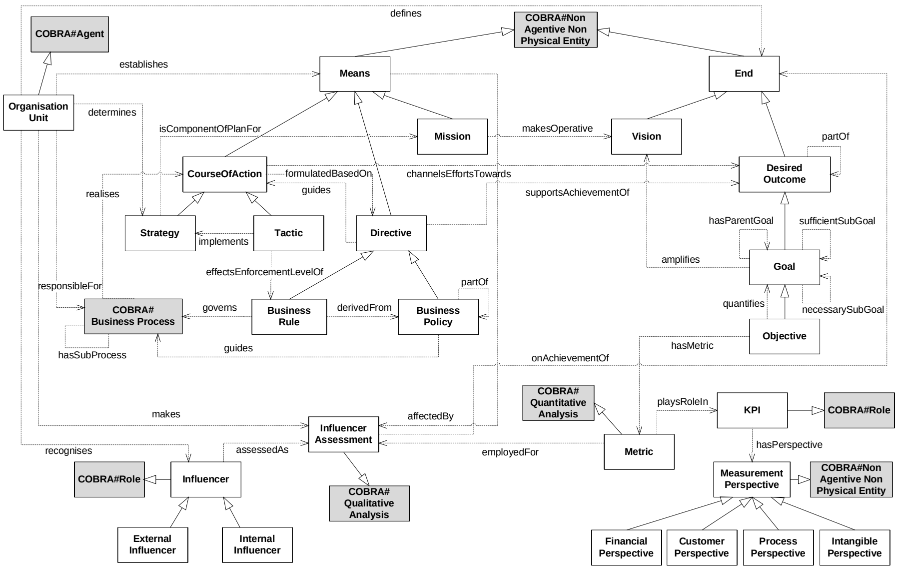

Fig. 4. Business Motivation Ontology.

Means (cf. Figure 4, upper left) are counterparts of Ends - they state *how* an enterprise intends to achieve its Ends. A Means may be either a Mission, a Course of Action or a Directive. A Mission describes a long-term approach which is focused on achieving the Vision. Its definition should be broad enough to cover all possible strategies and the complete area of enterprise operations. A Course of Action is an approach (plan) for configuring some aspect of the enterprise, undertaken to achieve its Desired Outcomes. It includes the subconcepts Strategy and Tactic. A Strategy is accepted by the enterprise as the right approach to achieve its Goals. Each Strategy is implemented by Tactics, which tend to be shorter term and narrower in scope in comparison to Strategies. A Strategy may be implemented by several Tactics and a Tactic may contribute to the implementation of more than one Strategy. To clarify the distinction, we say Strategies channel efforts towards Goals and Tactics channel efforts towards Objectives. Thus, the two concepts establish an analogy for concepts on different level of abstraction on the complementary sides of Ends (the *what*) and Means (the *how*). A Course of Action is realized by Business Processes.

Directives are set up to guide the Courses of Action. A Directive defines or constrains some aspect of an enterprise. A Business Policy is a non-actionable Directive whose purpose is to guide or govern the Business Processes within an enterprise. It is not focused on a single aspect of governance or guidance. Business Policies provide the basis for Business Rules. A Business Rule is a Directive that is directly actionable. It is often derived from (implements) Business Policy. Business Rules govern Business Processes and their enforcement level is effected by chosen Tactics.

Besides the fact that Metrics play an important role in business process analysis, setting targets is crucial to the motivational perspective of an enterprise's process space. BMO follows the Balanced Scorecard approach [1] here, and assigns performance measures with the company's aspirations. A Metric takes the role of a Key Performance Indicator (KPI) for an Objective, if it is applied to indicate the performance for a particular Objective. A KPI can take several perspectives; Financial, Customer, Process and Intangible (sometimes called Innovation and Learning).

An Influencer is something that can cause changes that affect the enterprise in the employment of its Means or achievement of its Ends. Almost anything within or outside a company could act as an Influencer in some situation. An Influencer Assessment is the procedure of judging the impact of an Influencer on an enterprise. Various methods can be used for performing this assessment. Consider an example method of SWOT analysis [11], where the result of Influencer Assessment can fall in the following categories: Strength, Weakness, Opportunity and Threat.

5 Strategy-Driven Analysis

The *what* vs. *how* duality of BMO allows us to contrast the strategies implemented with respect to established outcomes, i.e., "is the strategy contributing adequately to our ends?" Furthermore, it supports the analysis phase by guiding the computation and presentation of information based on what are the most important outcomes that need to be assessed. Evaluating the strategies themselves and determining whether the established Ends and their decomposition were reasonable, despite being very interesting research topics, fall outside of the scope of this paper.

As we indicated previously, Objectives are low-level Desired Outcomes that specify a desired condition over directly measurable aspects. These conditions are algebraic comparisons (e.g., greater than)[3] between any two Quantitative Analyses or between a Quantitative Analysis and a Constant Quantity (e.g., 5 ms). The specification of these conditions follows the approach established by COBRA in the sense that it distinguishes the analysis specification from the actual results obtained at certain points on time (e.g., see Quantitative Analysis Variation and Quantitative Analysis Variation Result). Objectives therefore allow us

[3] Figures do not include all the operators for the sake of clarity.

to define conditions over Metrics or their evolution over time (see Quantitative Analysis Variation) which can be computed directly from monitoring data (e.g., "execution time > 10 days"). This allows us to bring the results to a higher-level of abstraction where one can talk about Objectives in terms of whether they have been met or not.

Conversely, Goals are qualitative Desired Outcomes defined at a higher-level of abstraction, thus more closely related to the final strategical outcomes wanted. Goals can be decomposed into subgoals by using Logical Decompositions. We currently contemplate two kinds of Logical Decompositions namely *conjunctions* and *disjunctions*. Conjunctions allow us to express that a certain Goal is achieved when all its subgoals are achieved. Disjunctions on the other hand specify that a Goal is achieved when at least one of its subgoals has been achieved.

To support the definition of Goal graphs of an arbitrary complexity and still be able to determine their assessability, we distinguish two kinds of Goals: *Operational Goals* and *Complex Goals*. Operational Goals are those that are defined as a composition of Objectives or Qualitative Analysis Comparisons (e.g., "is the QoS increasing?"). Operational Goals are therefore low-level Goals which can be directly assessed based on existing measures and by performing the appropriate metrics computation and/or comparisons. Complex Goals on the other hand are composed of Operational Goals and therefore support defining higher-level Goals while still retaining their assessability given that they are defined in terms of Goals which are directly assessable. The notion of assessability introduced above allows us distinguish Goals that can be evaluated automatically–*Assessable Desired Outcomes*–from those that can not–*Non Assessable Desired Outcomes*. Whereby Assessable Desired Outcomes are either Objectives, Operational Goals, Complex Goals, or Goals solely composed of assessable Goals. Non Assessable Desired Outcomes are those that do not meet the restrictions above.

Assessing Desired Outcomes is envisaged as a process like the one formalised in [12]. In their paper, Giorgini et al. describe both a qualitative and a quantitative model for evaluating goal models defined as AND/OR decompositions. Their model includes the formalisation of a set of propagation rules and a label propagation algorithm which is sound, complete, and scalable. They also include an extensive set of relations between goals that can capture typical situations such as the fact that Goals contribute or hinder each other, or even mutually impede or ensure the achievement of each other. The reader is referred to [12] for concrete details of their approach. The assessment of Desired Outcomes in our model can straightforwardly be implemented in the same way. More importantly, our concept of Objective being directly linked to measurable aspects, can be used as what the Giorgini et al. refer to as initial nodes and allow us to directly apply the evaluation of goals over monitored data without the need to introduce any data manually. As a consequence, our conceptual model i) supports the operationalisation of desired outcomes into a graph that can automatically be evaluated; and ii) allows for directing the analysis of processes and organisations based on those aspects which are known to be of strategic importance.

6 Related Work

With Balanced Scorecard [1], Kaplan and Norton devised an approach to strategy implementation based on the specification of strategic objectives and assignment of respective measures used for assessing the achievement of these objectives, according to the Scorecard perspectives. In contrast to our work, the Balanced Scorecard does not provide an explicit linkage between the strategic and operational aspects of a process-centric enterprise neither does it specify the evaluation mechanisms necessary to automate the assessment of objectives within the Scorecard.

In more general work on BPA, we find techniques based on plain statistical analysis, data and process mining techniques, neural networks, case-based reasoning, or heuristic rules [13, 7, 14, 15, 16, 17, 18]. Some techniques focus on automating analysis to the greatest extent whereas others pay particular attention to obtaining results that can easily be explained and presented to the user. What can be distilled is that researchers have so far focused on operational aspects, leaving strategic concerns up to analysts interpretation. The work described herein is therefore, to the best of our knowledge, the first attempt that provides a comprehensive conceptualisation ranging from strategic aspects to operational details in a way that is amenable to performing automated analysis driven by high-level strategic concerns.

7 Conclusions

Current practices in Business Process Management try to maximise what can be measured in order to better support the decision-making process. We have argued that in order to reach the level of automation demanded by businesses nowadays, we need a holistic integrated conceptual model of the enterprise spanning from high-level strategic concerns, down to low-level operational concerns. We have described our model based on a set of modular ontologies which are currently being developed an enhanced within the SUPER project as part of a wider conceptualisation of the BPM domain. The ontologies we have defined are aligned with a formal model for evaluating Goals which supports the development of sound, complete and scalable algorithms for the evaluation of strategic outcomes out of low-level monitoring details. Our conceptual framework represents, to the best of our knowledge, the most comprehensive approach developed for supporting automated strategic analysis within BPM .

Despite the promising conceptual results obtained and the firm foundations upon which our work is based, a fully automated system based on these notions needs to be implemented and thoroughly tested within real settings. Additionally, the establishment of methodologies and techniques for mapping concrete ends into particular means and their ulterior validation when deployed in concrete settings are indeed very interesting areas for further research.

References

1. Kaplan, R.S., Norton, D.P.: The Balanced Scorecard - Measures that Drive Performance. Harvard Business Review (January/February 1992)
2. Hepp, M., Leymann, F., Domingue, J., Wahler, A., Fensel, D.: Semantic Business Process Management: A Vision Towards Using Semantic Web Services for Business Process Management. In Lau, F.C.M., Lei, H., Meng, X., Wang, M., eds.: ICEBE, IEEE Computer Society (2005) 535–540
3. Pedrinaci, C., Domingue, J., Alves de Medeiros, A.K.: A Core Ontology for Business Process Analysis. In: 5th European Semantic Web Conference. (2008)
4. Uschold, M., King, M., Moralee, S., Zorgios, Y.: The Enterprise Ontology. Knowledge Engineering Review **13**(1) (1998) 31–89
5. Fox, M.S.: 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, London, UK, Springer-Verlag (1992) 25–34
6. van Dongen, B., Alves de Medeiros, A.K., Verbeek, H., Weijters, A., van der Aalst, W.: The ProM Framework: A New Era in Process Mining Tool Support. In: Applications and Theory of Petri Nets 2005. 26th International Conference, ICATPN 2005, Miami, USA, Springer-Verlag (June 2005) 444–454
7. Mühlen, M.: Workflow-based Process Controlling. Foundation, Design, and Implementation of Workflow-driven Process Information Systems. Volume 6 of Advances in Information Systems and Management Science. Logos, Berlin (2004)
8. Pedrinaci, C., Lambert, D., Wetzstein, B., van Lessen, T., Cekov, L., Dimitrov, M.: SENTINEL: A Semantic Business Process Monitoring Tool. In: Ontology-supported Business Intelligence (OBI2008) at 7th International Semantic Web Conference (ISWC2008), Karlsruhe, Germany (2008)
9. Motta, E.: Reusable Components for Knowledge Modelling. Case Studies in Parametric Design Problem Solving. Volume 53 of Frontiers in Artificial Intelligence and Applications. IOS Press (1999)
10. Object Management Group: Business motivation model. `http://www.omg.org/docs/dtc/06-08-03.pdf` (2006)
11. Ansoff, H.I.: Corporate Strategy. McGraw-Hill, New York (1965)
12. Giorgini, P., Mylopoulos, J., Nicchiarelli, E., Sebastiani, R.: Formal reasoning techniques for goal models. Journal of Data Semantics (2004)
13. Castellanos, M., Casati, F., Dayal, U., Shan, M.C.: A comprehensive and automated approach to intelligent business processes execution analysis. Distributed and Parallel Databases **16**(3) (2004) 239–273
14. Grigori, D., Casati, F., Castellanos, M., Dayal, U., Sayal, M., Shan, M.C.: Business Process Intelligence. Computers in Industry **53**(3) (April 2004) 321–343
15. Azvine, B., Cui, Z., Nauck, D.D., Majeed, B.: Real Time Business Intelligence for the Adaptive Enterprise. In: The 8th IEEE International Conference on E-Commerce Technology, Los Alamitos, CA, USA, IEEE Computer Society (2006)
16. Watson, H.J., Wixom, B.H.: The Current State of Business Intelligence. Computer **40**(9) (2007) 96–99
17. March, S.T., Hevner, A.R.: Integrated decision support systems: A data warehousing perspective. Decision Support Systems **43**(3) (2007) 1031–1043
18. Alves de Medeiros, A.K., Van der Aalst, W., Pedrinaci, C.: Semantic Process Mining Tools: Core Building Blocks. In: 16th European Conference on Information Systems, Galway, Ireland (June 2008)