Semantics-Based Process Support for Grid Applications

Gayathri Nadarajan, Areti Manataki, Yun-Heh Chen-Burger
Centre for Intelligent Systems and their Applications (CISA)
School of Informatics, University of Edinburgh, Scotland
gaya.n@ed.ac.uk, a.manataki@sms.ed.ac.uk, jessicac@inf.ed.ac.uk

Abstract

The infrastructure of Grid is approaching maturity and can be used to enable the utilisation and sharing of large scale, remote data storages through distributed computational capabilities and support collaborations and co-operations between different organisations. Grid can therefore be suitably used to support the creation and running of a Virtual Organisation (VO). However, to assist the smooth operation of VOs, robust computational and storage facilities alone are not sufficient. There must also exist appropriate rich business infrastructure.

In this work, we consider business process frameworks that utilise semantics-based Business Process Modelling (BPM) technologies, and we illustrate the multidisciplinary nature of our approach by applying them to three different fields: Supply Chain Management, Business Intelligence and Knowledge Management, and Intelligent Video Analysis. We aim to show that these three application areas that incorporate semantics-based BPM methods could be used to support developing Grid applications, and to subsequently support VOs.

Introduction

The Grid was envisioned to support collaboration between researchers who are not physically co-located. Thus, more sophisticated technologies and infrastructure are required to support flexible collaborations and computations between communicating e-Scientists. Current efforts which include the rapid development of Grid applications, however, lack to provide a high degree of easy-to-use and seamless automation that would allow such effective collaborations on a global scale (De Roure, Jennings & Shadbolt, 2005). Key requirements and technologies have been identified with respect to the pertaining gaps in existing Grid efforts. These include, among others, process description and enactment (e.g. workflows), information integration (e.g. Semantic Web technologies) and the use of ontologies and reasoning mechanisms. We attempt to address these key areas by incorporating traditional Business Process Modelling (BPM) methods which are rich in modelling aspects and in infrastructure.

We aim to show that BPM methods provide a suitable foundation to capture the process dynamics of an organisation, and when these methods are enhanced with semantics they could be used for the effective development of Grid applications that subsequently support VOs. The rest of the chapter is organised as follows. The *Background* section provides an overview of VO and Grid, enterprise and business process modelling and the different existing paradigms for activity modelling; the *Case Studies* section provides the context for the intelligent support in the fields of Supply Chain Management, Business Intelligence and Knowledge Management, and Automatic Video Processing with Grid Workflows using our proposed approach; the *Discussion* section discusses the relevance of our work to Grid applications and the last section concludes the chapter.

Background

In this section, we present an overview of Virtual Organisation and Grid, Enterprise and Business Process Modelling and the paradigms for activity modelling. This will provide a background and motivation for the work before the case studies are presented in the next section.

Virtual Organisation and Grid

A Virtual Organisation (VO) is defined as a set of individuals or institutions that wish to share resources in a controlled and coordinated way (Foster & Kesselman, 2004). The sharing not only involves files and data, but also software, equipment and human skills. In this way closer collaboration is enabled for the achievement of shared common goals. It has been argued that the VO paradigm will undoubtedly come to play a major role in the theory and practice of management (Mowshowitz, 2002, p.25). In fact, business experts state that currently competition is not between different enterprises but between different enterprise networks. Similarly, collaboration among different scientific groups and governmental organisations is crucial, thus signifying the era of VOs.

Issues in VO-related research include: i) trust and security, ii) the computing infrastructure, and iii) the capturing of the highly dynamic processes involved. Trust and cultural problems within the VO are studied mostly by business and information systems researchers. The Grid is suggested as the most suitable infrastructure for the emerging VOs; in fact, Foster and Kesselman adapt their first, computing-oriented definition of the Grid (Foster & Kesselman, 1998) to the era of VOs, as they consider relevant social and policy issues (Foster & Kesselman, 2004). Lately there is an increasing interest in a more encompassing view of VOs, including engineering management aspects such as knowledge management (Katzy & Löh, 2003). In this, the capturing of highly dynamic VO processes, the so-called "virtual operations", is considered to be crucial. We regard semantics-based BPM methods as a useful means for capturing these processes while allowing their use in the Semantic Web.

Enterprise and Business Process Modelling

The aim of applying an Enterprise Modelling (EM) method is to seek ways to improve an organisation's effectiveness, efficiency and profitability. EM methods are typically informal or semi-formal. They provide notations that enable entrepreneurs to describe aspects of their business operations. The notation is normally complemented with semi-formal or natural language descriptions which allow details of the business operations to be described.

Many EM methods have emerged to describe and redesign businesses, namely business process modelling, business system modelling and organisational context modelling methods. BPM methods are able to formally express informally practised procedures. More importantly, actions and effects of these processes can be demonstrated using simulation techniques. Some examples of BPM method representations include Process Specification Language (PSL) (Schlenoff et al., 1997), Integration DEFinition Language (IDEF3) (Mayer et al., 1995), extended UML's Activity Diagram (Rumbaugh et al., 2004) and Petri-Nets (Reisig, 1985). In the course of half a decade ago, new process languages and models have been developed to promote the understanding and interoperability of process semantics over the Web, with the extensibility of operating over the Semantic Web. They are characterised by XML and XML-based languages, such as the Resource Description Framework (RDF) (Klyne & Carroll, 2004) and the Web Ontology Language (OWL) (McGuiness & van Harmelen, 2004). Some of these languages include Business Process Execution Language for Web Services (BPEL4WS, 2003), BPML (Arkin, 2002), Web Service Ontology (OWL-S) (Martin, 2006), and more recently, Web Service Modeling Ontology (WSMO) (Roman et al., 2005).

BPM technologies are established enterprise modelling methods that can be used to describe complex, informal domains. Their role within an organisation is to seek ways to improve its effectiveness, efficiency and profitability. BPM methods provide notations which enable entrepreneurs to describe aspects of their business operations. The notation is normally complemented with semi-formal or natural language descriptions which allow details of the business operations to be described. More importantly, actions and effects of these processes can be demonstrated using simulation techniques. Utilising BPM methods within emerging fields such as the Grid would prove useful as they could assist in providing a more mature framework incorporating both business- and semantics-specific technologies. This will be demonstrated by the application of semantics-based BPM methods in three fields that will be introduced in the next section.

One example BPM language is Fundamental Business Process Modelling Language (FBPML) (Chen-Burger & Stader, 2003) which is an inherited, specialised and combined version of several standard modelling languages. FBPML has two sections to provide theories and formal representations for describing processes and data: the Data Language (FBPML-DL) and the Process Language (FBPML-PL). FBPML-DL is first-ordered and describes the data model or ontology in Prolog-like predicate syntax. FBPML-PL is both formal and visual. Its constructs are composed using FBPML-DL plus its specific vocabularies for describing processes, events, actions, conditions, life cycle status and communication facilities. Thus, although FBPML contains separate means to describe data and process, the data model may be used as a construct in the process language, as such it is a fully integrated BPM language.

The main aim of FBPML is to provide support for virtual organisations which are becoming more and more pervasive with the advancement of Web and Grid technologies. It ultimately seeks to provide distributed knowledge and semantics-based manipulation and collaboration. Most importantly, people with different responsibilities and capabilities could work together to accomplish tasks and goals without technological or communication barriers caused by the differences in their roles.

As Semantic Web-based languages such as OWL-S and BPEL4WS are procedural-based and cannot be directly used for logical reasoning, we wish to provide a neutral ground for such capability. One way of achieving this is by incorporating the more mature and established BPM methods in these emerging technologies. FBPML is goal-directed and exportable to Semantic Web-based languages, e.g. BPEL4WS (Guo et al., 2004) and OWL-S (Nadarajan & Chen-Burger, 2007). Hence, integrating FBPML with technologies such as ontologies would provide for rich semantics that is suitable to be used in a Grid context to support reasoning, analysis and flexibility of virtual organisations.

To this end, several efforts have been made to bridge the gap between EM methods and Semantic Web services. The notion of Semantic Business Process Management (SBPM) (Hepp et al., 2005) was proposed to utilise semantic capabilities within business process technologies in one consolidated technology. This is motivated by the limited degree of mechanisation within BPM methods which could be supplemented by Semantic Web service frameworks. The problem is understood at ontological level and a technology for querying and manipulating the process space is suggested. SBPM includes Semantics-Based Business Process Modelling, thus semantically annotated process models, which may facilitate semantics-based discovery of process fragments and auto-completion of process models (Wetzstein et al., 2007, p. 10). A different approach is adopted by Betz et al. (2006), who introduce a semantic description of Petri Nets with OWL DL. Finally, Tripathi & Hinkelmann (2007) introduce a change management system that, among others, will translate the OWL-S specification of business processes to BPEL4WS.

We now turn our attention to the different paradigms taken to model activities and show how the one that we have adopted is relevant for the provision of semantics-based support for Grid applications.

Overview of the Paradigms for Activity Modelling

The different types of activity modelling paradigms are often distinguished based on their modelling objectives in capturing and emphasising the different types of domain objects during their modelling procedure. The most popular ones are: process-, data- and agent-centric. A process-centric approach treats the process (how things are done) as the first class modelling primitive over other modelling objects. It describes processes and their relationships as their main focus and often in great detail. A data-centric approach treats data (information on what is being manipulated by processes) as the first class primitive over processing. In this approach, data items and their relationships are explicitly represented as main modelling objects. Similarly, an agent-centric model treats agents as the main modelling objects.

This is analogous to the imperative versus declarative approaches. The data-centric approach allows large collections of data to be processed and passed between different components of a system and is primarily concerned with optimising the data flow between these components, e.g. functions within a scientific workflow. The process-centric approach is more concerned with the low-level achievement of the functions or tasks within an activity. Both these approaches capture the mechanistic components of an activity. On the other hand, an agent-centric approach captures specific aspects of the human component or software agent within the system, thus focuses on understanding an agent's roles, involvements, competence and collaboration with other agents in their tasks. This approach is useful in a distributed environment where the aim is to achieve a flexible mechanism to allow collaborative decision-making or task-achievements. In the following sub-section, we provide an example to discuss the applicability of the first two paradigms in relation to our work.

<u>A Discussion: Process- versus Data-Centric Modelling Approaches</u>

When considering the approach taken by our framework, it is appropriate to discuss process- and data-oriented paradigms for modelling the activities in our applications, specifically in workflow systems. As mentioned before, the former treats processing (how things are done) as the first class primitive over data; while the latter treats data (what is being manipulated) as the first class primitive over processing. One may argue that scientific workflows are specialised for data-intensive applications where the flow of data takes precedence over the way the processes are manipulated. This implies that the sequence of steps taken to achieve the goal of the workflow execution is fixed. Figure 1 illustrates this.

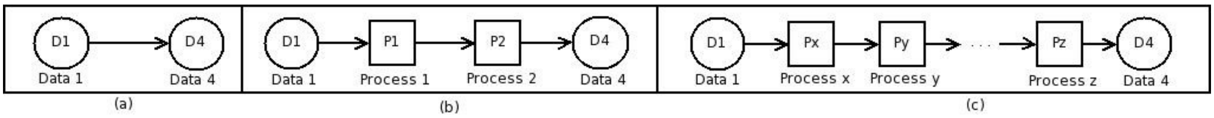

Figure 1: Data-centric (b) versus process-centric (c) approaches

Suppose the goal is to transform Data 1 to Data 4 (shown in Figure 1(a)). Figure 1(b) shows how this data transfer is achieved via two fixed processes, Process 2 and Process 3. If the workflow is fixed as such, then the main concern would be to optimise the data transfer between the two processes. Many data-intensive workflows operate in this manner, for instance workflows for Bioinformatics applications. Figure 1(c) shows how the goal is achieved via a combination of a variable number of processes. In this case, the workflow is not fixed.

Our approach has focused on this aspect of workflow construction due to the nature of the problem we are trying to solve, e.g. a video processing task may be achieved using several different combinations of steps. When the workflow is not fixed, the concern would be to optimise the sequence of processes involved, which is to select the set that minimises the cost of achieving the overall goal. Thus by manipulating the way things are done, even for a large dataset, one could also optimise the way the data transfer takes place. Prioritising processing over data does not imply that the data is not of concern. By focusing on the techniques applied to the data, one is also implicitly giving importance to the data because the methods extract useful information from the videos. Thus the process-oriented approach complements the data-oriented approach. Similarly, process-centric approaches may also be used to help optimise processes in a Virtual Organisation to enhance performance of internal and external organisational collaborations, or to provide additional information to help knowledge consolidation and extraction in a complex distributed organisational context. In our example, we demonstrate how traditional BPM methods may be used to enact scientific workflows, workflow automation and optimisation and knowledge management in a distributed environment such as the Internet or Grid.

In the next section we demonstrate how semantics-based BPM technologies could be used to provide intelligent support for Grid applications and VOs by presenting our work in three application areas.

Case Studies

We now present three case studies where semantics-based BPM technologies are utilised. For each case study, the introduction and motivations are first presented, followed by the content of the work done.

Case Study 1 – SCM and VO: The Need for a Knowledge-Based Approach

Supply Chain Management (SCM) is perceived as both an emergent field of practice and an emerging academic domain (Storey et al, 2006), and its role has growing importance in today's business world. From a purely operational approach of the 1960s we can now detect a more strategic one, where alignment, integration, collaboration and agility are seen as desirable characteristics of a supply chain (Storey et al, 2006; Chandra & Kumar, 2000; Tan, 2001). Hence, SCM is defined today as "the integration of key business processes from end user through original suppliers that provides products, services and information that add value for customers and other stakeholders" (Lambert & Cooper, 2000).

In the knowledge and globalisation era, enterprises are expected to collaborate in order to gain competitive advantage; it has been argued that companies today are competing not only through their product range but also through their supply chains (Harrison & van Hoek, 2008; Tan, 2001; Lambert & Cooper, 2000). Thus enterprises form supply chains (SC) that are tightly connected and flexible, where each member does not wish to satisfy only its goals, but also the whole SC's goals. This trend is defined as extended enterprise, which is "a kind of enterprise represented by all organisations or parts of organisations, customers, suppliers and sub-contractors, engaged collaboratively in the design development, production and delivery of a product to the end customer" (Browne & Zhang, 1999).

One can clearly see that SCM and extended enterprises are closely related to virtual organisations and virtual enterprises specifically. As Storey et al (2006) state, there is a shift of SCM to a collaborative model and virtuality, two basic characteristics of virtual organisations. Furthermore, Kangilaski (2005) views VOs as a SC strategy, while Browne & Zhang (1999) illustrate the similarities and differences of extended and virtual enterprises.

Our work on knowledge-based analysis of supply chain strategies is motivated by the increasing importance of SCM, and it intends to serve as a vehicle for insight into and comparison of different supply chain strategies. While most of the literature covers supply chain strategies and practices in a theoretical way, we suggest analysis at a lower level. We thus view semantics-based Business Process Modelling as a useful methodology for capturing SC strategies, and we suggest the use of a declarative workflow engine for the simulation of such business process models. The simulation results can be used when comparing different SC strategies, thus assisting further SCM analysis.

In our work we argue that knowledge-based techniques can be useful for the analysis of SC strategies. Given the strong relation between SCM and VO, we believe that the suggested framework presented in the section *Knowledge-Based Analysis for SCM* can also be applied in the case of a virtual organisation. Hence, BPM and simulation can assist in the decision-making procedure on the development of VO-related architecture and infrastructure, including Grid-based applications.

Knowledge-Based Analysis for SCM

Processes are perceived as the prevailing unit of organisational analysis (Adamides & Karacapilidis, 2006); similarly, business process models provide an abstraction of business strategies, and hence supply chain strategies. We recognise FBPML as a useful modelling language within our framework, as it has rich visual modelling methods and formal semantics, and it supports workflow system development. We also suggest a *cyclic modelling framework* for the development of a BPM that captures a firm's supply chain strategies. The framework is initiated with the identification of existing models and data from the organisational environment that are relevant to business and SC strategies. The second step involves the evaluation of the existing model and the detection of gaps. If no gaps are found, then the existing model is satisfactory and it constitutes the final BPM. Otherwise the modeller needs to deal with the gaps found and then to improve the existing BPM to create a new version, which is provided as input for the next cycle of the modelling procedure. The modelling framework is ended when no gaps are identified during the evaluation procedure. The relevant decision is made on the following four *evaluation criteria*: soundness, completeness, realism and level of detail of the model. Hence, the respective four questions are asked when evaluating a BPM: 1) Is it correct according to the modelling specification and does it behave correctly? 2) Does it cover all important SC strategies of the firm? 3) Does it correspond to correct business and SC strategies of the firm? 4) Is it abstract enough to provide an overview of the firm's SC strategies and detailed enough to provide interesting information?

After the development of a BPM that provides insight into a firm's SC strategies, we suggest its simulation with an appropriate *workflow engine*. We have designed and implemented such a meta-interpreter adopting a declarative approach. The workflow engine is business context sensitive, calculating the total duration and cost of a BPM execution (in the user defined time and cost units). Its architecture is described by three main components: a process model specification, a world state description and the workflow engine algorithm, all specified using logics. When the workflow engine is fed with a BPM and an initial world state description, the algorithm is activated and execution begins; throughout execution the workflow and world states are updated, and a relevant record is kept, while real-time feedback is provided to the user. An example for a process model specification is provided at the end of this section for the case study of Dell, while the predicates describing the entities and data of the world state are the following:

```
entity_occ(EntityName, EntityId, EntityAttribute).
data(SubjectID, Subject, Attributes).
```

The algorithm of the workflow engine has been implemented in Prolog, and an abstraction is provided below in pseudo code. We should briefly mention that time is treated explicitly, and in each time point the following sub-steps take place: actions scheduled for that time point are applied, junctions are checked for execution and are processed, and processes are checked for and are executed. The BPM stops execution at the time point when the finish junction has been reached, and no process is in execution, and no event is scheduled to fire the execution of some process at some later time point. The `ActionsAgenda` is a list with pairs of actions and time points at which they are scheduled, while the `ProcessedList` is a repository of junctions and processes that have been executed, thus keeping track of the internal workflow state.

```
1. Read the process model and initial world state.
2. Initialize the internal state of the workflow engine:
   ➢ Set T=0.
   ➢ Instantiate the ActionsAgenda to an empty list.
   ➢ Instantiate the ProcessedList to an empty list.
3. Execute scheduled actions:
   ➢ Search ActionsAgenda for actions with firetime=T.
   ➢ Execute them.
4. Check the process model for junctions to execute:
   ➢ Search the process model for junctions whose execution semantics are
     satisfied.
   ➢ Execute them at once and move them to the ProcessedList.
5. Check the process model for processes to execute:
   ➢ Search the process model for processes whose triggers and
     preconditions are satisfied.
   ➢ Calculate their completion time and move them to the ProcessedList.
   ➢ Move their actions to the ActionsAgenda and schedule them to be fired
     at their completion time.
6. Check for completion of the process model execution:
   ➢ If the finish-junction is in the ProcessedList and there is no process
     in the ProcessedList with completion time bigger than T and no event
     is expected in the future to trigger some process instance, then
     return T and the total cost of the processes in the ProcessedList and
```

```
    halt.
  ➢ Else go to 7.
7. Set T=T+1. Go to 3.
```

The simulation of a firm's BPM that provides an insight into its SC strategies could serve the further analysis of its SCM practices. We suggest two types of experiments. The first involves the optimisation of the BPM by parallelising two or more sequenced processes, thus decreasing the duration of the BPM execution. The second deals with the comparison of the detected SC strategies with alternative ones. This can be done by creating a BPM for the alternative SC strategies, by simulating it and comparing the simulation results to the ones of the investigated firm. The advantage of this framework is that the difference between alternative SC strategies can be made clearer when comparing BPMs rather than when comparing theoretical issues.

We have applied the suggested framework for knowledge-based analysis and modelling of SC strategies on the *case study of Dell*, a PC company that owes a big part of its success to its supply chain strategies. The two most important points of its strategies are the direct model (i.e. direct sales to all customer segments) and the build-to-order strategy, which means that a computer is built only after a relevant order has been placed. In fact, Dell places high importance in its relation with its key suppliers and customers and aims to virtually integrate with the most important ones through ValueChain.Dell.com, Premier Pages and Platinum Councils (Dell & Fredman, 2006). So Dell adopts, to some extent, the VO approach, a fact that makes this successful case very interesting to investigate.

While most of the literature covers Dell's SC strategies from a theoretical viewpoint, we have applied the suggested modelling and analysis framework, satisfying the need for an analysis at a lower level. Hence, following the cyclic modelling framework we have developed a BPM for Dell that provides an insight into its SC strategies, and which is strategic, business-goal-oriented and executable (Manataki, 2007). The initial input for the cyclic modelling framework was the MIT Process Handbook entry on the decomposition of the process "Dell-Create Computers to Order" (MIT Process Handbook), as well as SCM-related literature that covers the successful case of Dell. The cyclic modelling framework has led to a three-phased modelling procedure for Dell: The first BPM was the one implied by the MIT Process Handbook; this was enriched with sequence details about the processes, providing us with the second BPM version, which was later enhanced by incorporating SCM-related information, leading to the final BPM for Dell. Two example processes of the final BPM that provide an insight into Dell's SC strategies can be seen in Figures 2 and 3.

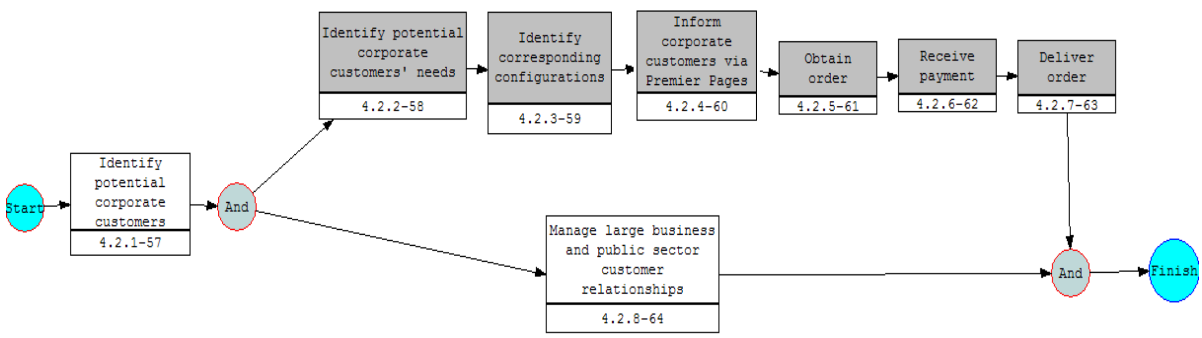

Figure 2: *Decomposition of "Sell directly to large business and public sector customers"*

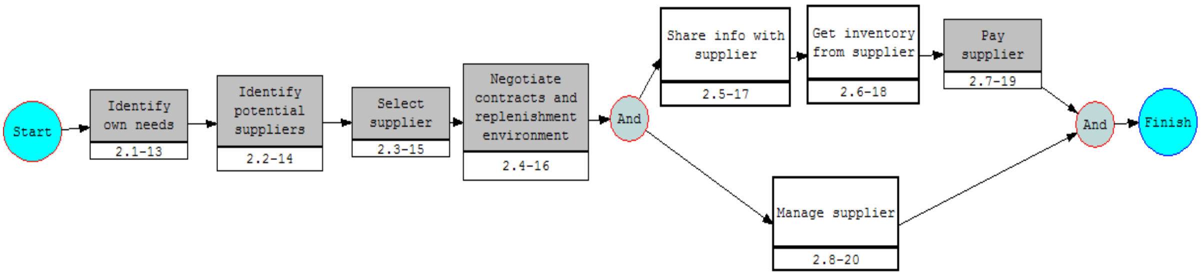

***Figure 3:** Decomposition of "Buy standard item to order"*

Let us now explain the rationale and the logical specification of Figure 2. The presented process decomposition illustrates Dell's direct sales strategy adapted to the corporate customer segment. As mentioned previously, it is strategic, thus it aims to provide us with structured information about Dell's direct sales strategy, rather than focus on operational details. We should mention that the first process 4.2.1 is further decomposed into two sequential processes, "Identify key personnel" and "Identify employee groups", and process 4.2.8 is decomposed into the parallel processes "Support large business and public sector customers" and "Get feedback from large business and public sector customers". The further decomposition of the latter two makes clear the use of Premier Pages, Account Teams and Platinum Councils for the desired virtual integration (Manataki, 2007). The logical specification of the examined process model above consists of the junctions' and processes' description. A formal representation that describes the control structure of the process model in Figure 2 is provided below using the junction/3 predicate (where process IDs are used instead of their full names). The predicate process/7 is used to give process descriptions as included in Figure 2.

```
junction(JunctionType, PreProcesses, PostProcesses).
junction(start, [], [p4_2_1]).
junction(and_split, [p4_2_1], [p4_2_2, p4_2_8]).
junction(link, [p4_2_2], [p4_2_3]).
junction(link, [p4_2_3], [p4_2_4]).
junction(link, [p4_2_4], [p4_2_5]).
junction(link, [p4_2_5], [p4_2_6]).
junction(link, [p4_2_6], [p4_2_7]).
junction(finish_and, [p4_2_7, p4_2_8], []).
```

```
process(ProcessID, ProcessName, Trigger, Precond, Action, Duration, Cost).
process(p4_2_6, receivePayment,
[exist(event_occ(customerOrder)), exist(event_occ(customerPayment))],
[exist(entity_occ(customerPremierPage))],
[create_data(customerOrderPaid, [orderID_tbre34, ccID_c3154a,
  amount_5000])], 1, 40).
```

Using the developed workflow engine we have run experiments on Dell's BPM improvement and on its comparison with a traditional PC company. The BPM optimisation experiments have shown that Dell's BPM cannot be improved, unless the semantics of the processes are changed. The SC strategies' comparison of Dell with a traditional PC company involved the creation of a BPM for the latter, and its simulation with the declarative workflow engine. The simulation results of the two BPMs have shown that: i) Dell's strategic choice of direct sales guarantees a faster and cheaper sales procedure ii) Dell's choice of buying to order and virtually integrating with its suppliers results in higher cost and time values in the case of a new supplier, but it leads to shorter duration than that of a traditional PC company on the long run (Manataki, 2007).

The main lesson learnt from the use of BPM and simulation on Dell's case is the fact that knowledge-based techniques can help extract essence of SCM's operations and assist the analysis and comparison of different supply chain strategies therefore constructing arguments to support or criticise certain SCM strategies. The presented BPM provides a logically structured view on Dell's SC strategies, and the conducted experiments with the use of the developed workflow engine allow us to structurally compare different SC strategies. Since VOs can be regarded as a specific case of SCM, we believe that the suggested framework for SCM could also be useful when analysing a VO constitution or when developing or improving Grid technologies for supporting VOs.

Case Study 2 – Business Intelligence and Knowledge Management

Large enterprises today are virtual organisations that their sub-organisations are distributed across different parts of the world. In addition, it is common practice for large organisations to contract out parts of its operations to different external organisations to enable the whole of its operations and services. This forces an organisation to communicate, collaborate and share knowledge with these external companies that again may not be co-located with each other. As a result, organisational information that is vital to enable the efficient and effective running of an organisation's operations and to sustain its continuous growth is inevitably distributed in all of the cooperated sub-organisations that may be both internal as well as external to the company.

To make the problem worse, important organisational information may be tacit knowledge, i.e. not described explicitly anywhere in the form of documentation, and it may only be known locally, making it difficult to be gathered accurately and in a timely fashion. This means that it will not be a trivial task to gain an accurate overview of an organisation when it is needed to support important decision-making processes. To overcome this problem, personal knowledge and past experiences, educational guess and ad-hoc approaches are often employed that may be error-prone and deployed in an inconsistent manner.

Such difficulties stem from the problem of not being able to gain quality information and turn them into appropriate knowledge when needed. To address these important issues, Knowledge Management (KM) methodologies suggest to use a set of conceptual models to help capture relevant information; while offering a structural process whose aim is to help turn relevant information into usable knowledge and provide it to the end user *at the right time, presented in the right form and provided to the right people,* therefore helping them in doing their tasks and achieving overall organisational goals (Schreiber et al. 2000).

Moreover, Business Intelligence (BI) techniques may be used in conjunction with KM by utilising a holistic reasoning framework that allows one to carry out business motivated analysis based on knowledge already stored in those models, by leveraging knowledge technologies. In this approach, important aspects of an organisation are identified and described in individual models using appropriate modelling methods. These models can also be modified over time as the organisation evolves. Information described in those models is often used together to create a coherent overview of the organisation and to support business-goal oriented enquiries.

In this paper, we describe a modelling framework that records three important aspects of an organisation: the Actor, Data and Process (ADP) aspects. These are the three corner stones of the domain knowledge of any (virtual) organisation. In this approach, each aspect is captured in a model using an appropriate modelling method and objects described in those models are formally defined in an underlying (formal) ontology. Because common concepts of different models are defined in the same ontology, relevant knowledge stored in different models can be related to each other, thus promotes knowledge sharing between these models. It also enables the derivation of new knowledge through cross-examination of different models.

This modelling framework therefore allows us to answer several types of business and knowledge management related questions that are frequently asked in an organisation, e.g. who created the data, who maintains the information, who authorised the operation, what is the impact of an operation, how critical is a dataset or an operation, etc. This line of work is therefore relevant to anyone who wishes to build a VO over the Grid or (a controlled) Internet, as it provides simple access to critical knowledge and an overview over the different components of a complex VO.

An Integrated Actor, Data and Process-oriented (ADP) Approach

To support today's knowledge economy, an appropriately designed Organisational Memory must closely support organisational operations and its business aims. To address these needs, we combine the use of the ONA (Ontology Network Analysis) method that provides an analytical framework for data analysis together with two other modelling methods (Role, Activity and Communication Diagram (RACD) modelling and a semantics-based BPM method – FBPML) to capture and analyse the three important aspects of an organisation – data, human and operational aspects. The closely combined use of these three methods allows us to create an integrated ADP modelling and reasoning framework.

The ONA approach deals mostly with the data aspects of the domain – it systematically works out the relationships and builds structure between domain concepts based on their attributes. The RACD is a role modelling method that captures the human aspect of a (virtual) organisation, including the hierarchical and functional relationships between them. A rich semantics-based BPM method using FBPML is also presented that captures the operational aspect of an organisation that is linked with the data and actor aspects. In this approach, an underlying ontology is utilised to allow data exchange and knowledge integration between the different aspects of an organisation, thus allowing its user to ask complex questions that needs knowledge existing in different parts of a VO (Chen-Burger & Kalfoglou, 2007).

Figure 4 gives example models that may be created and used as a part of modelling initiatives using the ADP framework. Figure 4(a) gives a part of the domain ontology created for the US Air Operations where concepts, terminologies and the relationships between them are defined. Common concepts and terminologies used in different models are defined in the ontology, so that different models may communicate with each other (knowledge sharing) through this ontology. To enable automated knowledge sharing and derivation, a common formal language is used to represent this ontology. User queries may also make use of terminologies defined in the ontology for communicating with the system; and answers to those queries may also be provided based on those concepts and vocabularies, although they may have been derived from a combined examination of several diagrams from different models.

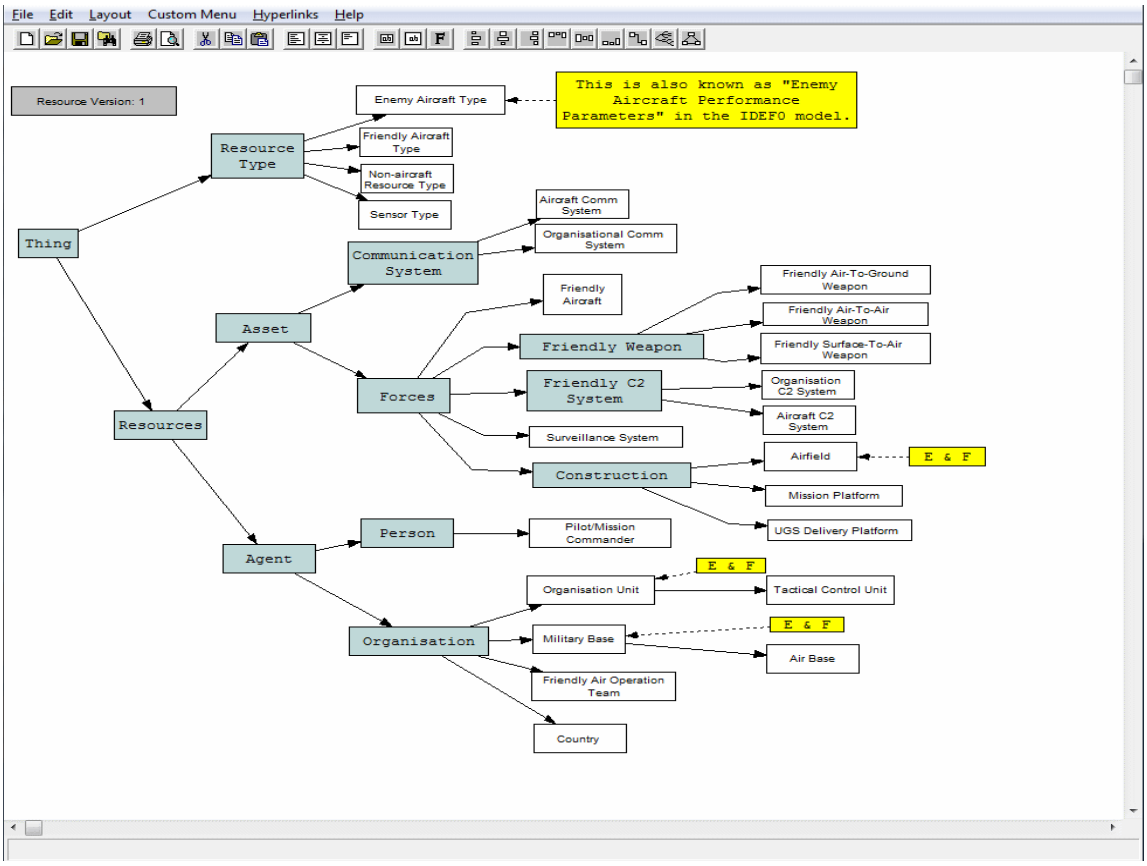

(a) Resource Ontology

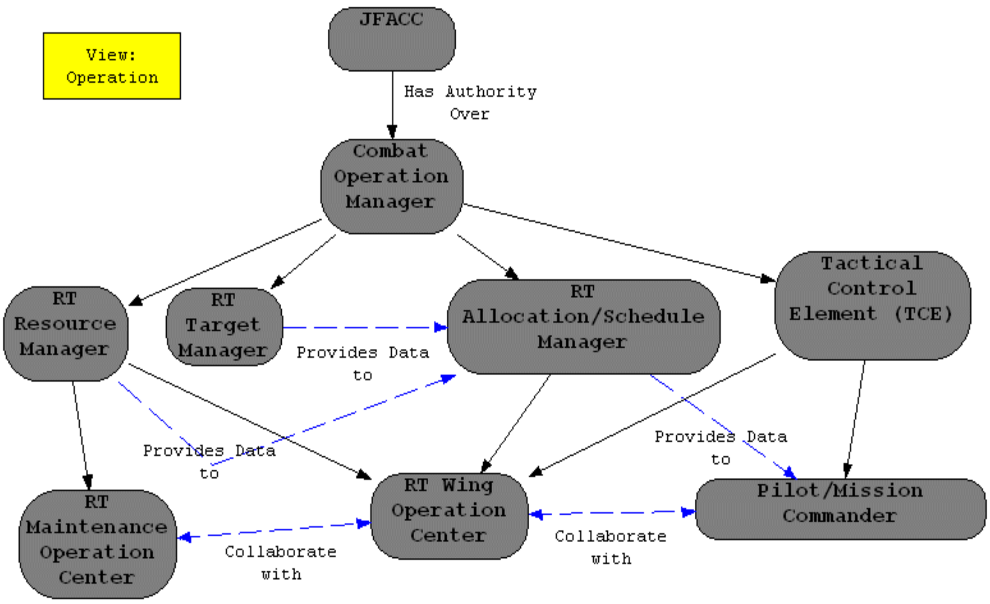

(b) RACD Role Model

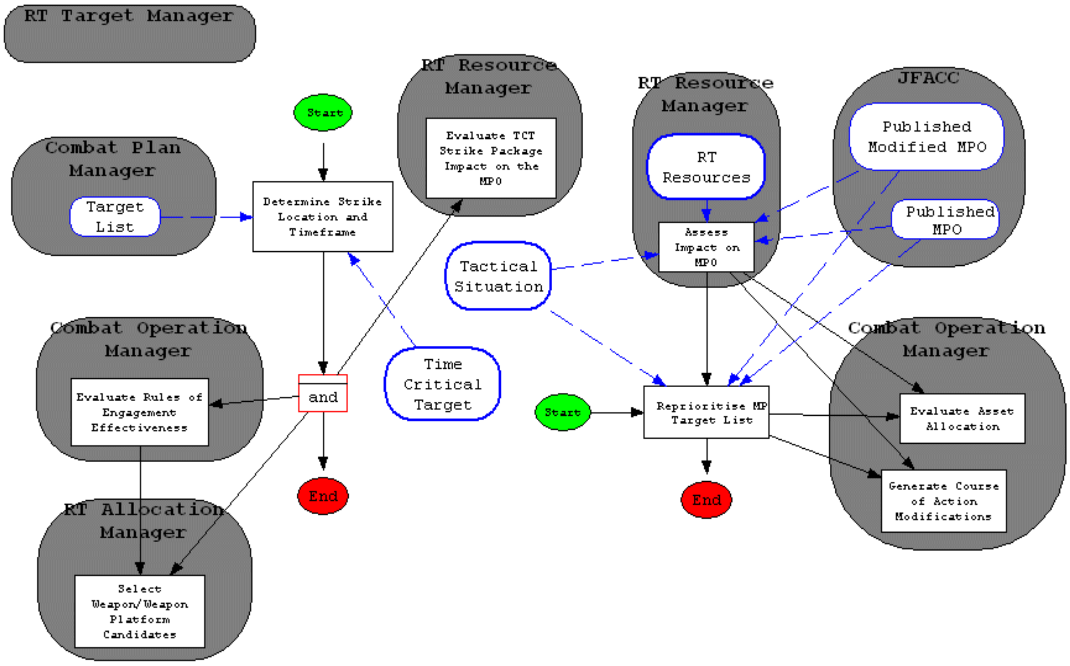

(c) Activity Model

Figure 4: *Example models used in the ADP framework – these are models developed for the domain of US Air Operations.*

Figure 4(b) is a diagram that is a part of the RACD Role Model where roles of personnel involved in the air operations as well as the relationships between these roles are defined. In this diagram, formal and informal "Influences" over roles are presented as links within the VO and the information flow between them are also defined. Note that those roles may be decomposed into more detailed roles, when appropriate. Roles that are used in RACD are also being used (consistently) in other models, so that different models overlap with each other, thus enable knowledge sharing and integration between them.

Figure 4(c) gives an activity model where personnel and data information are included in the context of organisational operations. This combined description includes information that already exists in other models, i.e. the ontology and RACD role model, enables one to draw more comprehensive conclusions regarding an organisation's operations: such as who carries out what operations, how they interact with each other and what information is used as a part of their activities. As data is included as a part of activity descriptions, new relationships between data can be derived via relationships between activities. At the same time, new relationships between processes may also be discovered via the relationships between data that are defined elsewhere (e.g. in an ontology). One can use these new linkages between data and activities to derive dependencies, relationships and detect inconsistencies between different data items, between activities and between data and activities - even where related data and processes may be originally defined in remote parts of different models (Chen-Burger & Robertson, 2005).

When adding typical business process analysis information to the model, such as the frequency of a process execution (as normally carried out in an organisation), the criticality of a process, the typical time and cost of a process, etc., one may use this information to infer properties about the data that is related to those processes. For example, from process analysis we may identify critical processes in a model. We may therefore infer that the fundamental data that is used as a main input by critical processes are also critically important. In addition, people who are responsible for generating main results for such processes may also be seen as critically important in their roles (called Actors in the framework). These are reasonably good quality approximate answers to queries that can be derived with relatively little effort. A formal representation to identify such critical data and actor information based on existing process knowledge is provided below:

$\forall$P, D. criticalProcess(P) $\wedge$ mainInput(D, P) $\rightarrow$ criticalData(D).
$\forall$P, A. criticalProcess(P) $\wedge$ mainActor(A, P) $\rightarrow$ criticalActor(A).

Other example queries that the ADP framework may help answer are: "Who created these data items?", "What process(es) created them?", "How are they being used?" and "Who uses them?", "Where are they being stored", "How are they being stored" and "What are the frequencies that those data items are being used and in what operational context", "How critical are those data items?" and ultimately "What are the impacts of those data and processes to the enterprise?". A carefully combined ADP approach can provide good approximate answers to most of these questions.

What is also interesting is that since a VO may include several sub- and external organisations, the roles and activities described in the models may therefore be cross-organisational. It may therefore be necessary to give a higher level view of the functions in the (virtual) organisation. In RACD, roles may be decomposed into sub-roles which are eventually mapped to individuals who carry out the tasks. This is similar to the decomposability of a process where process may be divided into sub-processes until they are leaf nodes. Such ability is useful as it enables one to give a higher level abstraction of the domain; while being able to answer lower level concrete questions such as those provided above. Thus this framework provides rich and flexible mechanism for capturing the complex context of a Virtual Organisation that is enabled via a distributed infrastructure such as the ones built upon the Internet and Grid.

Case Study 3 – Automatic Video Processing with Grid Workflows

Today, large scale data such as videos and images are accumulated on the Grid and other similar storages, but not analysed in an efficient manner. Manual processing would prove to be impractical, thus effective automated mechanisms should be in place to allow such analyses to be conducted more efficiently. With Grid technology, there should be provisions to allow image processing applications or programs across distributed locations to access these data via a Grid middleware. Grid workflows could provide automation to tackle this problem, with a centralised management system that could coordinate the execution of the applications. This would mean that a user in one location can invoke applications from remote machines in different locations. However, the user may not have enough expertise to determine which application this is, e.g. a marine biologist wishing to analyse videos of underwater marine life is not able to determine the image processing function calls to do so.

Automatic workflow composition would be ideal for this purpose, i.e. one that doesn't require the end-user to know which image processing application to invoke but that does it automatically for them. This would support collaboration between different e-Scientists, such as domain experts, computer scientists and non-technical users. This is challenging because it is inherently difficult to capture the functionality of the workflow components and their data types based on user input alone. Existing Grid workflows such as Pegasus (Deelman et al., 2005), Kepler (Ludäscher et al., 2006), Taverna (Oinn et al., 2005) and Triana (Taylor et al., 2003) play a major force in realising Grid applications. However, they are still largely composed manually, that is the user will need to describe the components of the workflow either graphically or by other means. Triana is a visual workflow system that allows the user to create workflows by the use of drag-and-drop. Taverna is another system that helps biologists and bioinformaticians to compose workflows manually in a graphical manner. Kepler is a highly reliable system that is particularly useful for workflows with complex tasks, however it does not do this automatically. Pegasus is a configurable system that can map and execute complex workflows on the Grid, used in a range of applications. It has the capability of composing workflows automatically from a user- or automatically-generated abstract workflow. However, this provision is still limited because Pegasus does not support all the modelling primitives required for vision tasks (e.g. loops) and does not manipulate the extensive use of semantics-based technologies. Lack of automated workflow composition is an obstacle that needs to be addressed. We propose a semantics-based workflow enactment framework to support naive users who need to use specialised or technically-difficult-to-understand tools.

BPM and Intelligent Video Analysis

Semantics-based BPM workflow enactment could provide a suitable support to overcome some of the gaps pertaining to Grid-based applications. In this scenario, we investigate its application in providing intelligent video analysis. Data originating from a Grid source are often huge in quantity and will need to be processed in an efficient manner by applications distributed across the Grid. In the Taiwanese EcoGrid project (EcoGrid, 2006), videos of underwater marine life have been acquired using wireless sensor nets and stored for analysis. One minute of video clip typically takes 1829 frames and is stored in 3.72 MB. That translates into 223.2 MB per hour, 5356.8 MB per day and 1.86 Terabytes per year for one operational camera. Due to the unpredictability of nature, one may not easily skip frames as they may contain vital information. Based on our own experience, one minute's clip will on average cost manual processing time of 15 minutes. This means that one year's recording of a camera would cost human experts 15 years' effort just to perform basic analysing tasks. Currently there are three under water cameras in operation and this will cost a human expert 45 years just to do basic processing task. This is clearly an impractical situation and more appropriate automation methods must be deployed.

In order to provide an intelligent BPM-based image processing (IP) system to tackle this problem, we have proposed a solution enriched with planning, process and case libraries and ontologies within a three-layered framework (Nadarajan, Chen-Burger & Malone, 2006). The system is grounded by ontologies that provide an agreed vocabulary between the Grid resources. The ontologies are contained within the design layer of the framework. We have provided three ontologies to describe the goals (IP tasks), domain description and capabilities (IP tools). A workflow system with full ontological integration has several advantages. It allows for cross-checking between ontologies, addition of new concepts into the workflow system and discovery of new knowledge within the system.

Initial work in the development of an image processing ontology was conducted through the Hermès project[1], whereby an extensive vocabulary was developed to describe user-friendly terminologies for image processing problems. The user ontology was built based on the fact that non-image processing experts are able to provide a priori information based on their domain expertise. The image processing ontology is used to formulate the image processing task and constraints. Extending this approach, we modularised these ontologies by separating the relevant components into goals, domain descriptions and capabilities. In particular, the

[1]GREYC Laboratory, Image Team, Caen, France (http://www.greyc.ensicaen.fr/~arenouf/Hermes/index-en.html)

capabilities were added as a new component of the ontology development process. This would contribute towards the use of ontologies within the *problem-solving* aspect as well, whereby a performance-based planning approach is adopted.

The goal ontology contains concepts related to the high-level IP task and the constraints on these tasks; it captures the user requirements in as much detail as possible. As a first step, the user interacts with the system by providing input values for the goal and constraints. The goal can be one of those specified in the ontology, e.g. Compression, Detection, etc. The constraints are additional parameters to specify rules or restrictions that apply to the goal. These include qualifiers e.g. Acceptable Errors, Optimisation Criteria, Detail Level and Quality Criteria contained within the goal ontology. For instance, the user may have a high level goal *"Detect all fish in video."* This is interpreted as having the following meaning in the Goal ontology.

(Goal: Detection) [Occurrence = all occurrences]
[Performance Criteria: Processing Time = real time]

Due to space limitation, only the relevant parts of the ontology are illustrated and highlighted in Figure 5. A more detailed version of the ontology is described in Nadarajan & Renouf (2007).

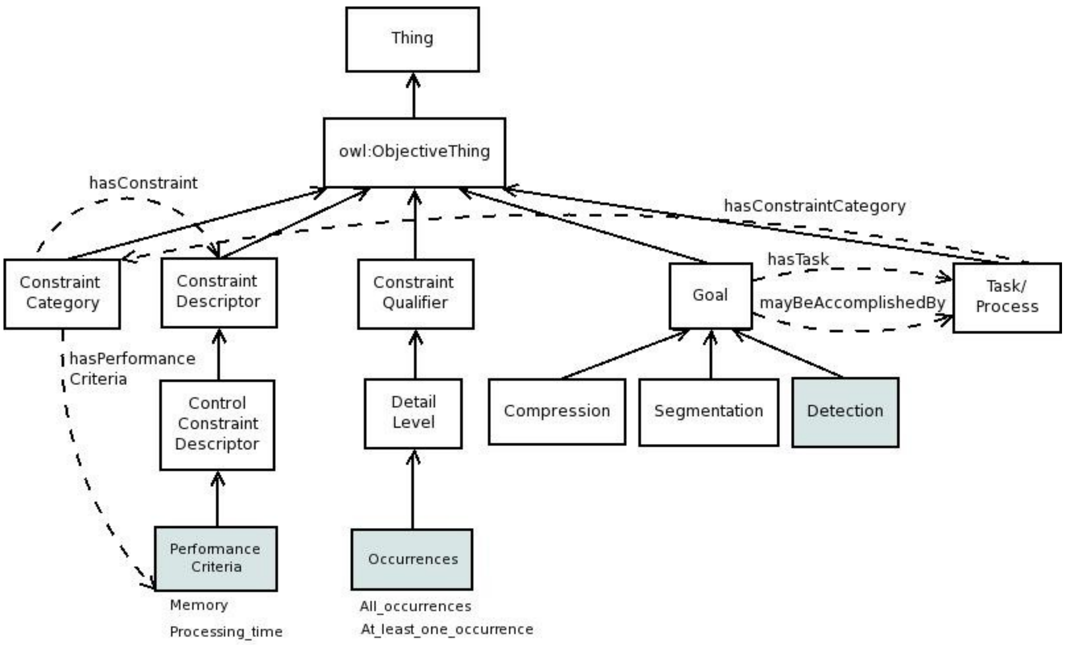

Figure 5: *Goal ontology containing concepts related to the user goal "Detect All Fish in Video".*

We have used FBPML to provide a basis for the data and process models. For a given video processing task such as 'Detection', a planning mechanism is used to obtain the sequence of sub-tasks involved to perform it. This is achieved by viewing it as a process modelling task using rich BPM constructs, where complex image processing tasks can be represented. Furthermore, the user can provide domain description such as the brightness level of the video and constraints that would restrict the task. Thus using the domain information and constraints, the system is assisted to provide a more accurate solution based on the user requirements as these details provide a basis for the preconditions and effects of the actions in the plan. For the detection task above, the high-level breakdown is given by *Pre-processing, Segmentation* and *Recognition*.

For the sub-task *Segmentation*, it can be hierarchically decomposed as shown in Figure 6.

We can model this using FBPML process model constructs. One of the methods to perform segmentation is 'Background Subtraction'. In this method, first a background model is constructed by averaging a series of consecutive images. Once this is done, the difference between the current frame and the background model is obtained. This would result in non-background objects, i.e. the fish. Then the background model is updated accordingly, that is by taking into account the new frame. This process is repeated until the last frame in the video sequence is obtained. The process model is described in Figure 7. Note that this process model includes a looping construct represented by the 'OR' control construct and arrow looping back.

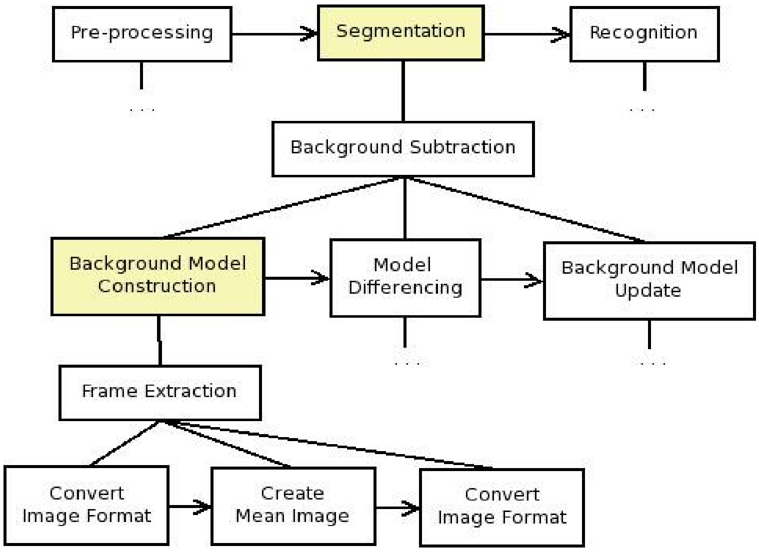

Figure 6: *Hierarchical decomposition for sub-task Segmentation within overall task Detection.*

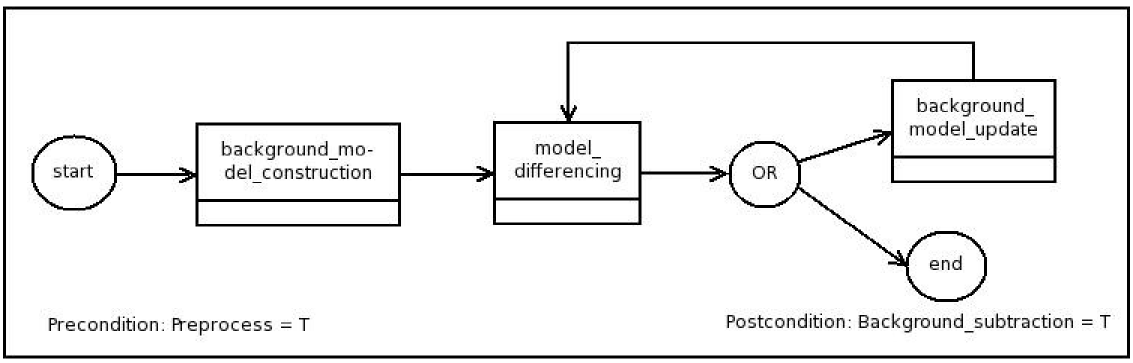

Figure 7: *An FBPML representation of the method 'Background Subtraction'.*

Each of the processes may be decomposed into other sub-processes, for example one way of performing *Background Model Construction* is *Frame Extraction* which involves 3 steps: i) Conversion of the original image to a suitable format for the proceeding computation, ii) Creation of the mean image and iii) Conversion of the image to its original format again. All these processes are primitive, that is they are not decomposable and correspond to a function call in a vision library. These process model representations can be encoded as methods in the planning component of the workflow engine. A walkthrough of the mechanism of the workflow engine is provided below.

Walkthrough
The goal and constraints for the 'Detection' task above is fed by the user and read by the workflow. The planner interacts with the goal and domain ontologies to obtain the implementation-specific values for the user input via an extension of FBPML-DL. The goal, constraints and initial domain state are determined and fed to the Planner. Then the decomposition is performed to obtain a sequence of sub-tasks (primitive tasks) involved to perform the overall goal. The process library which contains the instances of the processes available is consulted to check for the available operator(s) to perform each sub-task. The capability ontology is then consulted to check for the performance level of the selected operator. During this phase, it may be the case that a complete sequence is not obtained, or that more than one sequence is found to solve the goal. When this happens, the case library is consulted to retrieve the most *similar* solution based on past solutions. This adapted solution that corresponds to a set of function calls in an IP library is obtained and finally deployed for execution.

By grounding the workflow system with ontologies and incorporating rich BPM constructs to model complex IP tasks, we have provided a semantics-based workflow composition method that would enable e-Scientists to perform complex video analyses.

Relevance to Grid Applications and VO

We have demonstrated how semantics-based BPM technologies are used in three application areas; Supply Chain Management, Business Intelligence and Knowledge Management, and Intelligent Video Analysis. We have used FBPML as our formal representation for process modelling and ontologies for modelling the data. Such a representation is ideal for information sharing and exchange across the Grid.

In Dell's SCM, we have shown that BPM can give a good insight into supply chain strategies, and, with the use of an appropriate workflow engine, we can simulate and check on different such strategies. Since VOs can be regarded as an extension of specific supply chain strategies, the framework that we have suggested in our work can be useful for the analysis and improvement of a VO as well. So, if one wants to capture how a VO works and collaborates in order to satisfy its goals (i.e. when developing Grid technologies for a VO), BPM is highly recommended. Also, if one considers improving the function of a VO through business process reengineering, then our workflow engine could be a useful tool and help toward this direction.

Using the ADP approach for Business Intelligence and Knowledge Management, each aspect is represented by an appropriate process modelling method communicated with an underlying ontology. We have shown that appropriate diagnosis, analysis and inference could be performed within an organisation. This approach is ideal for any VO as it captures the human, business and operational aims of the organisation.

For more pervasive problem domains such as distributed image processing, the approach of incorporating BPM with the use of ontologies has proved useful in a Grid context due to the complexities of image processing itself. The EcoGrid is one of such examples where automatic support for non-IP users is essential. The framework we have proposed is designed to be used in the Grid as we have data originating from one source (Taiwan), while the image processing algorithms are developed in Italy and the workflow management is controlled in Edinburgh.

This is a typical Grid scenario and much work is under way to realise this framework. Once the whole system is in place, it would operate as a VO.

Conclusion

Semantics-based BPM could be used effectively to support Grid applications which provide a basis for the creation and running of a VO. BPM methods such as FBPML introduce an established formal representation to Grid applications that lack mature formalisms at present. This paper gives three relevant applications where different semantics- and process-based techniques have been applied. In this paper, knowledge-based techniques have been demonstrated to assist the analysis of Supply Chain strategies, and the simulation of corresponding semantics-based business process models can provide us with tangible numerical results to help gauge SCM (Supply Chain Management) performance.

The ADP approach uses and extends existing knowledge based modelling approaches to provide a useful framework to help answer queries that are frequently asked within a virtual organisation. It combines standard practices already existing in the different domains of knowledge management, ontology and business process modelling to provide a rich semantics-based support for organisational memory management. It is suitable for storing and presenting data and process information that is cross-organisational and can present organisational information at different levels of abstraction to suit different viewing and analytical needs. Based on information stored in these models, one is able to derive good quality approximate answers using logical inference techniques with relatively little effort.

In the field of Intelligent Video Analysis, BPMs provide a foundation for encoding methods and a language of communication between the workflow layer and other components. This provides a rich context-sensitive framework to allow e-Scientists to perform complex video processing tasks without possessing any technical expertise. Communicating the process models with underlying ontologies further enhances the support as ontologies are also formal, which enable machine processability. More importantly, they promote sharing and reuse of data, resources and applications across the Grid. To summarise, the benefits of utilising semantics-based BPM within Grid applications include support for reasoning, analysis and flexibility.

References

Adamides, E., & Karacapilidis, N. (2006). A Knowledge Centred Framework for Collaborative Business Process Modelling. *Business Process Management Journal*, 12(5), 557-575.

Arkin, A. (2002). Business Process Modeling Language (BPML), Version 1.0.

Betz, S., Klink, S., Koschmider, A. & Oberweis, A. (2006). Automatic User Support for Business Process Modeling, In K. Hinkelmann, D. Karagiannis, N. Stojanovic, G. Wagner, *Proceedings of the Workshop on Semantics for Business Process Management at the 3rd European Semantic Web Conference 2006*. Budva, Montenegro, 1-12.

BPEL4WS (2003). Business Process Execution Language for Web Services Version 1.1 IBM, BEA Systems, Microsoft, SAP AG, Siebel Systems. http://www128.ibm.com/developerworks/library/specification/ws-bpel/

Browne, J., & Zhang, J. (1999). Extended and virtual enterprises - similarities and differences. *International Journal of Agile Management Systems*, 1(1), 30-36.

Chandra, C., & Kumar, S. (2000). Supply chain management in theory and practice: a passing fad or a fundamental change?. *Industrial Management & Data Systems*, 100(3), 100-113.

Chen-Burger, Y.-H., & Robertson, D. (2005) Automating Business Modelling: A Guide to Using Logic to Represent Informal Methods and Support Reasoning Book Series of Advanced Information and Knowledge Processing, Springer-Verlag.

Chen-Burger Y.-H., & Stader J. (2003). Formal Support for Adaptive Workflow Systems in a Distributed Environment. In Layna Fischer (Ed.)*, Workflow Handbook* (pp. 93-118). Florida, USA: Future Strategies Inc..

Chen-Burger, Y.-H., & Kalfoglou, Y. (2007). Knowledge Management Support for Enterprise Distributed Systems. In Peter Rittgen (Ed.), *Handbook of Ontologies for Business Interactions* (pp. 294-310). IGI Global (former Idea Group Inc.).

De Roure, D., Jennings, N.R. & Shadbolt, N.R. (2005). The Semantic Grid: Past, Present, and Future, *Proceedings of the IEEE*, 93(3), 669-681.

Deelman E., Singh G., Su M.-H., Blythe J., Gil Y., Kesselman C., Mehta G., Vahi K., Berriman G. B., Good J., Laity A., Jacob J. C., & Katz D. S. (2005). Pegasus: a Framework for Mapping Complex Scientific Workflows onto Distributed Systems. *Scientific Programming Journal*, 13(3), 219-237.

Dell, M., & Fredman, C. (2006). *Direct from Dell: Strategies that Revolutionized an Industry.* New York, NY: Harper-Collins Publishers.

EcoGrid. (2006). National Center for High Performance Computing, Taiwan. Available as: http://ecogrid.nchc.org.tw/.

Foster, I., & Kesselman, C. (1998). *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers Inc., San Francisco, CA.

Foster, I., & Kesselman, C. (2004) *The Grid 2: Blueprint for a New Computing Infrastructure.* Morgan Kaufmann Publishers Inc., San Francisco, CA.

Guo, L., Chen-Burger, Y.-H., and Robertson, D. (2004). Mapping a Business Process Model to a Semantic Web Service model. *3rd IEEE International Conference on Web Services (ICWS'04).*

Harrison, A., & van Hoek, R. (2008). *Logistics Management and Strategy: Competing Through the Supply Chain.* FT: Prentice Hall.

Hepp, M., Leymann, F., Domingue, J., Wahler, A. & Fensel, D. (2005). Semantic Business Process Management: A Vision Towards Using Semantic Web Services for Business Process Management. In *Proceedings of the IEEE International Conference on e-Business Engineering*. IEEE Computer Society.

Kangilaski, T. (2005). Virtual organization and supply chain management. In *10th IEEE International Conference on Emerging Technologies and Factory Automation, Vol. 1* (pp. 705-712). IEEE.

Katzy, B., & Löh, H. (2003). Virtual Enterprise Research State of the Art and Ways Forward, In *9th International Conference on Concurrent Enterprising*. Centre for Concurrent Engineering, University of Nottingham, UK.

Klyne, G. & Carroll, J. (Eds.) (2004). Resource Description Framework (RDF): Concepts and Abstract Syntax. W3C Recommendation. http://www.w3.org/TR/rdf-concepts/

Kraemer, K., & Deddrick, J. (2001). Dell Computer: Using E-commerce to Support the Virtual Company, Center for Research on Information Technology and Organizations, Globalization of I.T., paper 236.

Lambert, D. M., & Cooper, M. C. (2000). Issues in Supply Chain Management, *Industrial Marketing Management,* 29(1), 65-83.

Ludäscher B., Altintas I., Berkley C., Higgins D., Jaeger-Frank E., Jones M., Lee E., Tao J., & Zhao Y. (2006). Scientific Workflow Management and the Kepler System. *Concurrency and Computation: Practice & Experience*, 18(10), 1039-1065.

Manataki, A. (2007). *A Knowledge-Based Analysis and Modelling of Dell's Supply Chain Strategies.* MSc. Dissertation, School of Informatics, University of Edinburgh, UK.

McGuinness, D., & van Harmelen, F. (2004). *OWL Web Ontology Language.* World Wide Web Consortium (W3C). http://www.w3.org/TR/owl-features/

Martin, D. (Ed.) (2006). OWL-S Semantic Markup for Web Services, Pre-Release 1.2. World Wide Web Consortium (W3C). http://www.ai.sri.com/daml/services/owl-s/1.2/ (Temporary location at SRI).

Mayer, R., Menzel, C., Painter, M., Witte, P., Blinn, T., & Perakath, B. (1995). *Information Integration for Concurrent Engineering (IICE) IDEF3 Process Description Capture Method Report.* Knowledge Based Systems Inc.

MIT Process Handbook, Phios Repository, Case Examples, Dell. Retrieved June 12, 2008, from http://process.mit.edu/Info/CaseLinks.asp

Mowshowitz, A. (2002) *Virtual Organization: Toward a Theory of Societal Transformation Stimulated by Information Technology.* Westport, CT, USA: Greenwood Press.

Nadarajan, G., Chen-Burger, Y.-H., Malone, J. (2006). Semantic-Based Workflow Composition for Video Processing in the Grid. *IEEE/WIC/ACM International Conference on Web Intelligence (WI-2006)*, IEEE, 161-165.

Nadarajan, G., & Chen-Burger, Y.-H. (2007). Translating a Typical Business Process Modelling Language to a Web Services Ontology through Lightweight Mapping. *IET Software* (Formerly IEE Proceedings Software), 1(1), 1-17.

Nadarajan, G., & Renouf, A. (2007). A Modular Approach for Automating Video Analysis. In *Proceedings of the 12th International Conference on Computer Analysis of Images and Patterns (CAIP'07).* Springer-Verlag, 133-140.

Oinn T., Greenwood M., Addis M., Alpdemir M. N., Ferris J., Glover K., Goble C., Goderis A., Hull D., Marvin D., Li P., Lord P., Pocock M. R., Senger M., Stevens R., Wipat A. & Wroe C. (2005). Taverna: Lessons in creating a workflow environment for the life sciences. *Concurrency and Computation: Practice and Experience*, 18(10), 1067-1100.

Reisig, W. (1985). Petri Nets, an Introduction. *Eatcs: Monographs on Theoretical Computer Science, Vol 4.* Springer-Verlag.

Roman, D., Keller, U., Lausen, H., de Bruijn, J., Lara, R., Stollberg, M., Polleres, A., Feier, C., Bussler, C., & Fensel, D. (2005). Web Service Modeling Ontology. *Applied Ontology*, 1(1), 77-106.

Rumbaugh, J., Jacobson, I., & Booch, G. (2004). *The Unified Modeling Language Reference Manual, 2nd Edition.* Addison-Wesley.

Schlenoff, C., Knutila, A. and Ray, S. (Eds.) (1997). *1st Process Specification Language (PSL) Roundtable.* Sponsored by National Institute of Standards and Technology, Gaithersburg, MD. http://www.nist.gov/psl/

Schreiber G., Akkermans H., Anjewierden A., de Hoog R., Shadbolt N., van de Velde W., & Wielinga B. (2000). *Knowledge Engineering and Management: The CommonKADS Methodology*. Cambridge, MA: The MIT Press.

Storey, J., Emberson, C., Godsell, J., & Harrison, A. (2006). Supply chain management: theory, practice and future challenges. *International Journal of Operations & Production Management*, 26(7), 754-774.

Tan, K. C. (2001). A framework of supply chain management literature. *European Journal of Purchasing & Supply Management*, **7**(1), 39-48.

Taylor I., Shields M., Wang I., & Rana O. (2003). Triana Applications within Grid Computing and Peer to Peer Environments. *Journal of Grid Computing*, 1(2), 199-217.

Tripathi, U. K. & Hinkelmann, K. (2007). Change Management in Semantic Business Processes Modeling. In *Proceedings of the Eighth international Symposium on Autonomous Decentralized Systems*. IEEE Computer Society, 155-162.

Wetzstein, B., Ma, Z., Filipowska, A., Kaczmarek, M., Bhiri, S., Losada, S., Lopez-Cobo, J.M. & Cicurel, L. (2007). Semantic Business Process Management: A Lifecycle Based Requirements Analysis. In M. Hepp, K. Hinkelmann, D. Karagiannis, R. Klein and N. Stojanovic (eds.): *Semantic Business Process and Product Lifecycle Management. Proceedings of the Workshop SBPM 2007*, CEUR Workshop Proceedings.