This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Knowledge-based process management – an approach to handling adaptive workflow

PLEASE CITE THE PUBLISHED VERSION

PUBLISHER

© Elsevier

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Chung, Paul Wai Hing, Larry Y.C. Cheung, J. Stader, P. Jarvis, J. Moore, and A. Macintosh. 2019.
"Knowledge-based Process Management – an Approach to Handling Adaptive Workflow". figshare.
<https://hdl.handle.net/2134/2348>.

This item was submitted to Loughborough's Institutional Repository by the author and is made available under the following Creative Commons Licence conditions.

For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

Knowledge-Based Process Management – An Approach to Handling Adaptive Workflow

PWH Chung, L Cheung, J Stader*, P Jarvis*, J Moore and A Macintosh***

Department of Computer Science, Loughborough University

*AI Applications Institute, University of Edinburgh

**International Teledemocracy Centre , Napier University

Abstract

In recent years, many organisations have found enterprise modelling, especially business process modelling, to be an effective tool for managing organisational change. The application of business processing modelling has brought benefits to many organisations, but the models developed tend to be used for reference during business operations and re-engineering activities; they rarely play an active role in supporting the day-to-day execution of the processes.

While workflow management systems are widely used for the streamlined management of "administrative" business processes, current systems are unable to cope with the more dynamic situations encountered in ad-hoc and collaborative processes [1]. A system that supports complex and dynamically changing processes is required.

There is increasing interest in making workflow systems more adaptive [8][20] and using knowledge-based techniques to provide more flexible process management support than is possible

using current workflow systems [4][21].

This paper describes the results of a collaborative project between Loughborough University and the University of Edinburgh. ICI and Unilever were industrial partners on the project, providing real business requirements in the application domain. The project investigated the use of ontologies, agents and knowledge based planning techniques to provide support for adaptive workflow or flexible workflow management, especially in the area of new product development within the chemical industries.

Keywords: workflow management, adaptive workflow, knowledge-based process management, process modeling, agent selection

1. Introduction

Enterprise modelling, especially business process modelling, is an effective tool for managing organisational change. Business processing modelling has brought benefits to many organizations. Organisations and their processes undergo changes from time to time, and in some cases changes are continual. Organisations change either through proactive efforts to become more competitive or in response to a need to maintain competitiveness in a changing environment.

Workflow systems are designed to support business processes. These systems embody explicit process models that will need to be modified to reflect the changes in the organisation. A major limitation of workflow systems is that they can, typically, only support simple, predictable processes, but not the dynamically changing and complex processes that are present in many organisations.

van der Aalst et al [25] point out that existing tools that support collaborative work are typically in one of two extremes: unstructured, information centric approaches (CSCW) and structured, process centric ones (product workflow). “Adaptive workflow aims at providing process support like normal workflow systems do, but in such a way that the system is able to deal with certain changes. These changes may range from ad-hoc changes such as changing the order of two tasks for an individual case to the redesign of a workflow process ...”

The Task Based Process Management (TBPM) project aims to support the management of change in business organisations with the help of intelligent task management and coordination technologies. The area of new product development (NPD) within the chemicals industries was chosen as the application focus for generating and testing ideas. The reason is that the NPD process, like many other engineering activities, has characteristics that pose significant challenges for workflow systems:

- It is a highly interdisciplinary process, requiring the co-ordination of individuals from different engineering and business specialities.
- Many ad-hoc processes occur, which nonetheless are activities requiring a significant amount of time, specific technical and business skills and other resources to perform, thus needing careful management.
- The structure of the process is highly flexible, which vary from one project to the next.
- The process is information-intensive, where a significant amount of technical information of different types is generated and must be distributed to interested parties reliably and efficiently.

These characteristics make conventional workflow systems unsuitable for handling NPD processes ([1], [7], [11]). However, if such support could be provided, there are potential benefits to be gained by:

- providing a single computing framework allowing the planning, execution and monitoring of processes. This ensures that planned processes are followed faithfully and allows the inspection of information about the current status of the process.

- permitting flexibility in process modelling and planning, so that process plans may be revised in the light of events and experiences gained during the process.
- improving the quality of decision-making because of the effective management of information and its dissemination to interested parties as it becomes available (for example, technical difficulties or discoveries which may have an impact on the business-case for the product being developed).

Because of these and other related potential advantages, there is increasing interest in making workflow more adaptive [8][20] and in using knowledge-based techniques to allow workflow to cope with complex and dynamically changing processes [5][21].

The TBPM project extends the application scope of current workflow system by incorporating knowledge about processes in general, and about the domain in which the system is deployed in particular. Such knowledge enables the system to reason about processes within those domains, providing the necessary power and flexibility for computer support.

The industrial partners on the project provided the scale-up process within NPD as a test-case application to elucidate the requirements for an intelligent process management support tool from a particular real-world standpoint.

2. The Scale-Up Process

Scale-up typically occurs at a point during NPD when a promising product has been identified, preliminary marketing investigation has been done, and a potential chemical process for manufacturing the product has been proposed, but not yet fully investigated. There is a series of experiments to be performed at gradually increasing scale, starting in the laboratory and ending (if all proves satisfactory) with a working pilot plant in order to investigate the behaviour of the proposed chemical process, and the nature of the engineering necessary to implement it at the intended scale of production.

A scale-up project tends to be long-term, typically 6 to 24 months, and involves a very large number of interacting and ad-hoc processes that are the province of a number of different disciplines and departments, including R&D, Engineering, Marketing, Finance, Safety Health and Environment and General Management. Failure to achieve good communication, collaboration between disciplines and effective implementation of the necessary processes can have a range of undesirable consequences, from a failed product [2], with associated financial consequences for the business, to an incompletely studied chemical process, with potentially disastrous consequences for health, safety or for the environment.

Cooper [3] suggests that improper management of the NPD process contributes to a large proportion of new product failures. He presents a canonical process model for effective NPD, in which it consists of a series of “go/no-go” decisions, separated by activities intended both to develop the design of the product itself and to improve the business and marketing case for its introduction. The scale-up process straddles one of these decision points – the decision of whether or not to proceed by sanctioning the construction of a pilot plant, which usually represents a major capital outlay.

3. Requirements

In considering the scale-up process, extensive interviews were carried out with our industrial partners to identify the key characteristics of the process and to elucidate the key requirements of a support system.

A number of challenges were identified which are not addressed by current commercial workflow systems, and which must be addressed by a system if it is to succeed in supporting complex and dynamic scale-up processes. These requirements are discussed below.

3.1 Flexibility

There are almost no fixed rules for process management in scale-up even though a number of canonical process models exist. The complexity of the domain, the unpredictability of the types of projects and difficulties, which may arise from time to time, preclude the a priori specification of fixed rules. Instead, guidelines and norms are established, which are open to interpretation by the project’s management

team: people trusted with managing projects are experienced, and expert in the field, and much is left to their discretion. While expected norms can be stated, e.g. “never do any work without an approved budget for it”, these are open to interpretation and variation in particular circumstances. For example, a project manager might choose to undertake a few days of work without a specific budget, in an initial feasibility study for a good, known customer where it seems likely that the work will lead to the placing of a valuable contract. To take account of this situation, a process management system must not impose fixed constraints. Instead, constraints should be advisory, so that the manager is aware of situations when he is breaking them, but is still at liberty to do so.

The only inviolable rules are those touching on safety: if a manager allowed an experiment to be conducted without having it first undergone the required safety assessment and approval, he could expect to be severely disciplined.

3.2 Adaptability

Each scale-up project typically takes a unique form, depending on many influences, including the nature of the product being developed, the target market for the product, and the knowledge developed about the product during the scale-up process. While certain characteristics and activities can be predicted, much of the process cannot be fully specified at the start, since it requires information that only becomes available some way into the project. For example, while it is known that considerable

engineering analysis and design work will need to be undertaken for most large scale-up projects, the type of specialist engineering activities which will be required can not be known with certainty until a considerable amount of experimental work has been done to identify the problems associated with the specific chemical process in question. Therefore, the process models should be expressed at an abstract level in the beginning, and then be interpreted in an appropriate manner for the specific context during the scale-up process. This means that it must be possible to interleave the planning and execution of tasks, deferring the complete specification of later tasks while earlier tasks are being executed.

3.3 Common Processes

There are commonalities between parts of the process at different levels of detail. Most scale-up processes look very similar at a high level, with a fairly consistent breakdown into management, engineering, scientific and commercial activities, which is a collection of best practice and experience within the industry for that specific type of project [6]. There is a need to provide expression of such common project structures as a resource for those setting up a specific project. However, such common structures should not be rigidly enforced, a manager should be allowed to select a template that best matches the current situation and then specialise it where necessary.

3.4 Capability Matching

Once the process has been decided, the individual tasks that comprise that process must be performed.

In scale-up, many of these tasks are highly technical, and require specific skills and experience of the agent chosen to perform them (the term “agent” denotes activity performers, including humans or software systems). In order to be able to assist with the selection of appropriate agents and delegation of tasks to them, the system must have access to knowledge about capabilities – both those required by the tasks and those possessed by the available agents.

3.5 Organisational Context and Authority

A process must be performed in compliance with organisational norms. For example, a task within the process should not be delegated to a specific person without the agreement of that person’s manager, if such agreement would normally be regarded as necessary. Similarly, if a person has the responsibility for performing a particular type of task on a project, the workflow system should not delegate that task to someone else, even if they are technically capable of performing it. In addition, different stakeholders in the process may work at different levels of detail, both for setting up a process and during its execution. For example, managers may further sub-divide their respective parts of the project according to technical discipline, delegating each sub-part to a relevant specialist. In order to achieve this, the system must be able to reason with knowledge describing the organisation structure and how lines of authority and management are implemented within that structure.

3.6 Information Management

Different technical disciplines are involved in different parts of the process and the people involved must communicate effectively with one another because the work of one discipline may impact on the work of another. However, a commonly encountered problem during scale-up is that crucial information is passed on too late (or not passed on at all), so that work may be wasted or delayed unnecessarily. This is because it is often not clear who needs the information. For example, engineers and chemists are not always good at understanding the consequences that their design choices or experimental results have on marketing concerns for the product. Routes for the transfer of information cannot be “hard-wired” into the process models, because of the flexible and ad-hoc nature of the processes and their models. Instead, the system must be able to match information that becomes available against process requirements and participants’ interests. The system should be able to distinguish, for example, between a “specification for a pump” and a “costing for a pilot plant”, and know that the former may be of more interest to a process engineer than to the project budget controller, who will in turn be more interested in the latter.

3.7 Domain Knowledge

The specification of capabilities and information will both need to refer to concepts in the engineering domain. Therefore, the system should be provided with the knowledge about the domain of chemical process engineering in order, for example, to distinguish an engineer who can design control systems from one who can design mechanical systems, and a specification of a pump from one of a voltmeter.

4. Task Based Process Management System Architecture

The approach adopted by the project is common in workflow systems in which process models are actively used for supporting business processes. However, the project extends the scope of current workflow systems by using not only simple models of processes themselves, but also models incorporating knowledge about the application domain and organisational context in which the processes are being enacted.

These knowledge-rich models of business processes and their organisational context are used by TBPM's workflow engine, the Task Manager, to support heterogeneous agents (humans and software agents) in working together to achieve their tasks. The Task Manager includes a process modelling tool and an interactive process planner. The planner uses artificial intelligence techniques to assist in the planning of tasks [4], while permitting the user to participate in planning decisions. An agent-based architecture supports the execution and coordination of the planned process among multiple participants distributed across a computer network.

TBPM is not a single, fully integrated, monolithic IT system. Rather, it provides a process management framework that integrates existing stand-alone tools such as word processors, spreadsheets, etc. This is particularly important in a highly technical field such as chemical engineering, where the use of specialist third-party software, such as simulation and analysis tools, is essential. The structure of the system is illustrated in figure 1.

Figure 1: The architecture of the TBPM system

4.1 Process Models

TBPM relies heavily on the use of knowledge about business processes and their context. This knowledge is provided to the system in the form of models that give information about processes, organisational structure, and agents.

There are three different classes of “user” interacting with the process models in TBPM:

- Managers, who select and/or generate the process models to be enacted.
- User agents (both human and non-TBPM software), who carry out tasks specified in the models.
- TBPM software components; in particular the Task Manger, which enacts the overall process based on the models.

It is essential that the models used are unambiguous, carrying the same semantics for each of these classes of user, and also that the TBPM software is able to interpret and manipulate the models. To this end, models in TBPM are formally expressed using terms drawn from a small number of related ontologies.

4.2 Ontologies

Uschold and Gruninger [23] define an ontology as “... a shared understanding of some domain of interest ...”, and the TBPM ontologies have been developed to provide just such a shared understanding, in the form of unambiguous semantics for the models used in the system. The key features of an ontology include:

- An ontology of some domain of interest identifies and precisely describes the important concepts in the domain and the valid relationships between these concepts.
- The set of important terms and their definitions are agreed between all participants within the domain, and thus form a basis for communication about the domain.
- An ontology can be specified independently from the intended application for which it is developed. This enables its re-use for other purposes and applications touching the same domain. By separating the ontology (i.e. the language and concepts) of a domain from the uses to which it is put, different applications are enabled to use the same domain ontology, and thus to communicate in terms of this shared ontology.
- An ontology can be formalised and thus support communication between IT systems as well as between humans.

The main benefits derived from the use of ontologies are:

- Ontologies support communication by providing a shared vocabulary with well-defined meaning, thus avoiding ambiguities and misunderstandings. They can support communication between any agents whether they are human agents or software agents. This is particularly useful in situations where experts from different fields need to work together, as is the case in NPD.
- In order to provide flexible support in a non-trivial business situation (such as NPD), models of different aspects of the domain have to be interrelated to make best use of them. However, it is difficult to capture different domain models in a way that takes their relationships into account. Using ontologies, it is possible to specify related models independently, thus reducing the difficulty of capturing domain models.

For TBPM's ontologies, wherever possible, existing ontologies or standards are being used ([19], [22]) and previous work is being built on ([24], [21]). There are two distinct sets of ontologies employed in TBPM:

- The process ontology: a single ontology defining the concepts central to the management of business processes in general. The ontology specifies characteristics of the concepts and thus defines the components of the models used within the system.
- The domain ontologies: A small number of interrelated ontologies with specialise the general concepts of the process ontology for the particular application domain in question.

The TBPM Task Manager depends only on the process ontology, which makes the system independent of the particular application domain being studied. To deploy TBPM in another domain requires the replacement of the domain ontologies, but no changes to the Task Manager software or the process ontology. More details of the relationships between the ontologies are available in [14] and [15].

The process ontology defines the concepts central to the handling of all business processes. Some of the key concepts are shown in table 1.

Concept	Description
Task type	A task type is a specification of the basic nature of a task, in terms of the purpose for which the task is carried out.
Role	A role represents a coherent body of work within a task, all of which is expected to be carried out by the same agent. Key aspects of a role are its remits and its authorities.
Agent	An agent is an entity (person or software) that is capable of carrying out some set of tasks. Each agent has an associated set of capabilities, which specify the task types that agent can perform, together with a level of competence at which the agent performs that type of task.
Resource	A resource is any entity that may be produced, consumed, or used during the conduct of a task. Resources include both information and the physical objects of the application domain.
Task	A task in TBPM consists of a task type, specifying the basic purpose to be achieved, together with a number of roles that are expected to be required to be filled for the task to be accomplished successfully, and a number of resource specifications for both the inputs and outputs of the task.
Plan	A plan is a set of tasks, and associated sequencing constraints and resource flows between them, intended as a structural breakdown of one possible way of accomplishing another task type.
Organisational Unit	TBPM borrows heavily from the Enterprise Ontology [24] for its organisational modelling primitives, which are based on the concept of organisational units, between which management and ownership relationships may exist, and which may own resources.

Table 1: Key concepts of process ontology

Specialising TBPM to handle processes within a particular domain involves the development of a number of distinct, though interacting domain ontologies. Each such ontology represents the set of domain-specific specialisations of a single term from the process ontology. For example:

Task Types: Defines domain-specific types of task, such as “design”, “conduct experiment” and “purchase”. This enables the specification of domain-specific capabilities, authorities and remits.

Resources: Defines the main object (physical and information) associated with the domain. For example, “heat exchanger”, “hydrocarbon”, “specification” and “risk analysis”.

Tasks: Defines common tasks in the domain, in terms of their types, roles, and resources. For example, an experimentation task might be specified as being of type “conduct experiment”, with inputs “experiment design” and “apparatus”, output “experiment results”, and at least one role specifying a requirement for a research chemist capable of carrying out chemical experiments.

4.3 Ontology Implementation and Use

Both process and domain ontologies are structured as a single generalisation hierarchy, with the more abstract concepts of the process ontology forming the root terms of which other terms are specified.

Figure 2 illustrates part of the hierarchical structure. *Resource* is a term from the process ontology, the other terms are specialization of *resource* defined in the domain ontology for scale-up.

Figure 2 Part of the domain ontology

In addition to the hierarchical structure, each term may have associated with it a number of named parameters, each of which takes as its values an expression from any section of the ontology. A particular term's parameters, and their associated constraints, are inherited by all that term's children, which may add further specialisation to the constraints. For example the term *design-information* has an associated parameter *system*, which describes the type of system to which the design information in question refers.

Thus, expressions such as:

Specification (system: pump (type: reciprocating))

and:

Specification (system: pump (type: centrifugal))

can be generated to present specifications for different types of pump. This structuring of the ontology represents a compromise between expressive power and likely acceptance by users, and further discussion can be found in [13].

There are several ways in which ontologies are used in supporting the scale-up process at different phases. During the specification phase the following benefits can be achieved with the help of the ontologies:

- a user can search for ontology terms and their relationships. For example, given an appropriate resource ontology, a system can work out that an *equipment capital cost estimate* is of interest to somebody who has expressed an interest in *project financial information*.
- The agent ontology can help to infer an agent's capabilities and interests by looking at the agent's type. For example, a human agent who is a mechanical engineer can be assumed to be able to analyse technical information relevant to mechanical equipment.
- The agent, task, and resource ontologies all support the process of turning software systems into agents by providing a well-defined, common interface for communication.

During active process support, the ontologies are used in the following ways:

- To determine which process models can be used to perform a given task, the system can use the task ontology to match user task requirements with available process fragments.
- To ensure that the most suitable agent is selected to perform a task, the system can use the agent and organisational unit ontologies to check that the agent has the capabilities and authority required to perform the task. The system uses the ontologies to match agent specifications with process specifications.
- To determine a suitable style of negotiation during delegations the system can use the organisational unit ontology to determine the relative status of the agents involved.

- To pass information along the process and between agents, the system can use the resource ontology to match between a task's information requirements and information generated by other tasks in the process.
- To ensure information is passed to interested agents, the system can use the resource ontology to match between agents' declared interests and information generated by the process. The roles defined within the organisational unit ontology can be used to infer some of the interest of agents.
- The system can use the resource ontology to discover the format in which information is available in and determine whether and how agents can communicate using that format.
- During both process enactment and specification, human users can use all the ontologies to ensure that they understand the relevant concepts (i.e. processes, information, etc.)

4.4 Plan Library

Even within a fairly tightly specified domain, there is no single process model that can be applied to all projects. Many factors conspire to make the process followed unique to each project: the nature of the product, the customer, the history and structure of the business, the availability of people and resources, etc. Even for a single project, the nature of the process followed cannot be known a priori: it emerges over time, being affected by events and discoveries, which occur during the project's lifetime.

On the other hand, many recognizable similarities do exist between different projects. The broad structure of most product innovation processes in any one business may be essentially the same (Cooper [3] suggests an outline process intended to be broadly applicable); many engineering design projects contain an identifiable progression from requirements, through specification, to design; many engineers, when presented with similar tasks, will conduct them in essentially the same way.

The availability and applicability of such process patterns is one of the key aspects of domain knowledge that must be available to an adaptive workflow system if it is to offer effective process management support. When planning a project, a user must be able to identify suitable compositions of canonical processes for achieving the tasks, select the ones most appropriate to the current situation, and, if necessary, adapt them for that situation.

To support this TBPM incorporates a “Plan Library” which maintains a database of process structures, relating each structure to the types of tasks for which it is a suitable method. Each plan specifies a set of tasks, together with the ordering constraints and object flows between them. Thus, a plan represents one possible way of achieving a given type of task by breaking it down into a particular structure of sub-tasks.

Each plan specifies only a single level of structural decomposition. However, the decomposition is into a further set of tasks, for each of which further plans may exist in the library. These plans may in turn be selected to specialize the sub-tasks, and so a multi-level hierarchical process structure may be generated by composition of many plans.

Figure 3: Two possible alternative plans are shown for achieving the “obtain” task.

For any given task, there may be multiple possible plans, expressing different ways of breaking the task down for different situations. Figure 3 shows a simple plan for the task of acquiring some artifact. Tasks are shown as boxes with square corners and object flow by boxes with rounded corners. The plan has two sub-tasks: *specify* the artifact required, then *obtain* it. Two additional plans are shown, each a possible method for achieving the *obtain* task: one by designing and constructing the artifact in question, the other by purchasing it.

A manager of a process can put together such plans dynamically, decisions being based on the current situation and which of the alternative breakdowns is more appropriate. A similar process of specialization may then be applied to structure each of the sub-tasks within the chosen breakdown. For example, the "agree purchase" sub-task may have several different available plans; one might cover putting large orders out to tender, another simply ordering a part from a manufacturer's catalogue.

The plan library provides good support for carrying out common tasks in any one of a number of standard ways. Collaboration in process management is enhanced by the provision of plans at different levels of abstraction, so that each user can work with plans expressed at the necessary level, without having to commit unnecessarily to particular lower-level details. A project manager can outline the structure of the overall project, leaving the details of the experimentation and engineering involved to be fleshed out by the relevant scientists and engineers.

More details on TBPM's models can be found in [9][10][16]. In practice, models are kept small, and they are specified independently without direct reference to each other. The models are related to each other dynamically by the Task Manager at the time of process execution. This gives a just-in-time feel: many decisions can be delayed until the last moment, so that the latest state of the organization can be taken into account.

5. TBPM's Task Manager

The general cycle of workflow can be divided into three phases:

- Specification: determine what it is that needs to be done.
- Planning: decide how it is to be done and who is going to do it.
- Execution: do it.

For any single activity, these three phases must be performed in sequence. In conventional workflow systems, specification and planning for the whole process which is to be supported are completed before execution commences, and the model produced is used many times for different examples of the same process. In most organisations there are many business processes for which this approach is not feasible. Some processes may be too complex or unpredictable to fully specify in advance, others have to be performed in dynamic environments which frequently lead to changes to the original specifications and plans being required.

The task management approach of TBPM allows the three stages to be interleaved: as long as for any individual task the stages are conducted in the correct order, different tasks in the system may be at different stages. Specification of a project's specialist engineering activities may be postponed until after most of the experimentation has been executed; planning whether to purchase or construct specialist equipment can wait until an appropriate stage in the project (over the course of the project, new suppliers may emerge, old ones go out of business, and the organisation's own manufacturing facilities undergo significant changes).

TBPM's support for each of the three stages is outlined below, before some discussion of the types of inter-agent communication that are supported by the system.

5.1 Specification

Support for the specification of tasks to be performed is provided in the ontology of domain tasks. This ontology defines common forms of task which occur frequently in the domain. For each task, the basic type of the task is specified, along with the set of roles normally involved in carrying out the task, and the types of resources required as inputs and produced as outputs. TBPM incorporates a process modelling tool which is used to construct the plans in the Plan Library. This tool has access to all ontologies and to previously specified plans in the Plan Library.

Task specifications are brought into the system either as a new top-level activity (for example, a completely new project to be run), or as part of a plan used to implement another task specification.

5.2 Planning

The planning phase is divided into two parts. The first is process planning, which decides what tasks are to be carried out and in what order. The second is agent selection, which decides which agents (human or software) are going to carry out the tasks.

5.2.1 Process Planning

Initial planning support is provided by the integration of the process modelling tool with the Task Manager, and the provision of the Plan Library described above. The process modeller can be used to edit the structure and detail of specific plans, enabling a human manager of a task to adapt the plan to current (possibly changing) circumstances. Such changes can be made in advance of execution, but they can also be made when the process has already started, to take account of information about the process and its context as it becomes available.

(a) laboratory bench experimentation requiring no engineering input

(b) laboratory experimentation requiring construction of specialised apparatus

Figure 4: two alternative process fragments for experimentation activities

Figure 4 shows a somewhat simplified illustration using the IDEF3 process description notation [12].

The figure shows two alternative plans for the task “Carry Out Experiment” which differ in that (b) allows for the need for special apparatus to be constructed for the experiment, while (a) does not (it assumes that the necessary apparatus already exists). Only the structural breakdown is shown; each activity specification would also include details of the abilities, authorities, and organizational roles required of an agent in order to carry out the activity.

A task always has an owner assigned to it and changes to the task are always made by its owner. The owner of the overall project, the project manager, is assigned when the project starts being supported by the Task Manager. That project manager can delegate ownership of specific tasks to others during the course of the project. In the scenario above, the owner of the task “Carry Out Experiment” will be the research manager. This person will be presented with the alternative plans in Figure 4 as options for different ways of achieving the task. Having selected an appropriate process fragment for the task in hand, the manager might wish to specialise it for the current situation: a suitable lab-rig design may already be available, so "Design Experimental Apparatus" can be deleted from the plan.

Further support is provided for an important class of process management decisions: that of locating suitable agents to carry out or participate in a task. Agents participate in tasks by filling roles in the task specification. For instance, in the experimentation process example above:

- The “perform experiment” activity might require an agent with an experimentation capability in the appropriate area of chemistry.
- The “assess experiment safety” activity may include an approval stage which should always be performed by someone in the role of “project safety officer”.
- The “train staff” activity may involve the approval of expenditure on such training, which would have to be performed by someone with the authority to do so.

The problem of finding agents arises when a role must be filled. The system provides support by considering two issues:

- First, is there someone who should fill the role? Is there a role defined at a higher level of abstraction in the process whose remit subsume the work type of the role in question? If so, the default is to unify the two roles. For example, if an agent has overall responsibility for all mechanical engineering design on the project, he would be expected also to take responsibility for the design of the mechanical handling systems for the pilot plant stage of the project.
- Second, is there someone who could fill the role? In some cases there will be no such subsuming remit, or the agent with such a remit will be unable to fill the role for some reason (perhaps having no more available time, given current commitments). In these cases, an agent is sought to fill the role on the basis of capabilities: i.e. an agent is sought with capabilities that encompass the work type defined for the role.

For any given role, the system will, at the request of the task owner, search the agent registry and process structure to determine the sets of agents with encompassing remits or suitable capabilities, and present the results in order, with the most closely matching remit given precedence. The task owner is then free to accept the system's top recommendation, select another suitable agent from the list, or even elect to use an agent which the system does not regard as suitable (the available models may not cover everything, and there are likely to be cases where the task owner knows more than the system about who should be chosen).

5.2.2 Agent Selection

When a task is planned, one of the crucial decisions to be made is the selection of the correct set of agents to participate in the task. The selection of the correct set of agents for a task is an important decision, made by managers on the basis of knowledge about different agents' abilities and roles within the organisation. The approach adopted in the project in the selection of agents for a task is from three complementary angles, characterised by the questions:

- Who should do it?
- Who can do it?
- Who may do it?

In a capability matching function, the capabilities required by a task are matched against the capabilities

held by available agents in order to identify suitable agents for the task. Capability matching here refers to the more sophisticated matching that takes into account knowledge about capabilities themselves and relationships between them. The reason for using any such matching function in a workflow context is that it is impossible to predict the exact environment in which a task is executed. Specific agents may not be available at the time of execution (people take holidays or leave the organization), or more suitable agents may have become available (people are hired and new software systems are developed). Similarly, activities may not be required in the specific context of task execution. Availability of agents not only has an impact on assigning activities to agents, but also on the decision of which method is chosen to achieve a given task. If a method for carrying out a task requires a particular capability but there are currently no agents available with that capability, then the task may be achieved using an alternative method.

Clearly, capability matching requires specifications of capabilities both for activities and for agents. If capability specifications are to be matched, it is important that the specifications use common and well-defined terms that can be related to each other.

Capabilities are not the whole story, however. It is typically not appropriate simply to select any agent capable of performing the task: frequently there will be a specific agent who is supposed to take the task on. This agent is referred to as having a remit that encompasses the task in question.

Remits and capabilities can assist, respectively, with the determination of who should take part in a given task, and of who can take part. The remaining issue mentioned above is: who may do it? This is a question of establishing authority for a task before it is allowed to proceed. The organizing structure for sets of capabilities, remits and authorities is a role. A role within a given task ties together a coherent set of capabilities, remits, and authorities into an entity describing an agent's relationship to the task and its component sub-tasks.

Many types of role are readily identifiable within organizations: e.g. "chief executive officer", "mechanical designer", "project manager". These generally relate to fairly high-level tasks, such as complete projects, or the day-to-day running of the business. Similar roles can be identified for tasks of much finer granularity, however. For example, you are currently fulfilling the "reader" role in a task of "read academic paper". To fill this role successfully, you require the capability to read the English language. However, the reason why you (as opposed to a colleague with equivalent capability) are doing the task may be that you have a particular remit to keep up to date with research in enterprise information systems, and this remit encompasses the task of reading papers published in the area.

Within TBPM, a given task may involve multiple roles, but must always have at least one: the "responsible agent" role, which is filled by the agent taking responsibility for the management and progress of the task.

A remit specifies a class of tasks within the domain of interest. Any particular remit is attached to a role, which is in turn attached to a task within the hierarchical process structure. The meaning of a remit is that the agent filling the role to which it is attached is expected to take on the management of tasks (i.e. fill the "responsible agent" role) which occur as part of the finer process structure of the task to which the role relates, and which are encompassed by the remit. So remits may in some sense be regarded as being "inherited" down the process hierarchy. When a task arises for which an encompassing remit exists, then:

- The agent filling the role is under some obligation (though this may be negotiable) to take on the task.
- It would violate the organization's rules for any other agent to be asked to take on the task.

Thus, remits capture the essence of the reasoning necessary to decide who should perform a task. For more detailed description of the selection of agents for roles within tasks can be found in [16].

5.3 Execution

The TBPM Task Manager supports the co-ordination of the execution of the overall process by tracking the status of tasks within it, maintaining constraints between tasks, and managing the flow of resources between tasks.

Execution of individual tasks is not directly the responsibility of the Task Manager. When the owner of a task wishes to execute it, there are two options:

- Use the planning support to break down the task into finer structure, which will be co-ordinated by the Task Manager (of course, each sub-task generated will require execution).
- Execute the task. The system provides access to the required inputs, and in general, the executing agent is responsible for providing the required outputs to the system when they are available.

However, some tasks may be broken down to a level of detail where the individual tasks being executed are recognised as computer-based operations (edit, view, etc.), and the input is data of a known format. A registry of tools which can be invoked to support users in carrying out such operations is maintained, and the appropriate tool is invoked automatically in such cases.

5.4 Information Management

Much of the scale-up process is defined in terms of the flow of technical information: design documents, experiment designs and results, costings, business cases, etc. . All information-related specifications (agent interests, activities' inputs and outputs, etc.) are expressed in terms of the ontology, which classifies the information. This means that the information can be managed by the system: agents' expressed interests can be matched against the types of information available in the system, and the formal information flows captured in the plans being executed can be used to ensure that information generated by one activity is passed to other activities which depend on it.

6. Evaluation

A TBPM prototype was implemented to evaluate the proposed framework and a detailed scale up process and associated domain knowledge was input to the system for testing purposes. The flexibility requirement is the most demanding and the most important of the identified requirements because it reflects the need to cope with and accommodate change in an organisation and its processes. The strategy to interleave the three phases of workflow (specification, planning, execution) makes the TBPM system powerful and flexible enough to provide effective support for dynamic and complex processes.

There are different types of change and TBPM deals with these in different ways. Some changes can be supported in TBPM by changing models, ensuring that the models always reflect the organisation. For example, if an organisation decides to introduce safety auditing, related activities must be added to the process models and safety-related capabilities, authorities, and remits must be included in the models of agents and the organisational structure. Such changes can be actively supported by TBPM because the structures and processes involved in the change are modelled explicitly. If the organisation wants to change a business process, TBPM can be used to model the new process and changes can take immediate effect in the organisation.

Changes are localised by keeping the models small, and relating them dynamically: a change to a method of achieving a task may need only a single change to a plan in the plan library, while a change to an agent's capabilities requires a single change to the relevant entry in the agent registry. These changes thenceforth automatically propagate throughout the organisation's processes as the new models are related dynamically to the rest of the models in the system.

Other changes can be accommodated without any change to the models themselves, simply by relating independently specified models as late as possible. For example, fluctuations in agent availability are communicated to the system simply by logging on and off; using the capability and authority specifications of the process models to match against those of agents that are currently available ensures that at the time of execution the most suitable agents are chosen.

Intelligent matching algorithm will need to be designed and implemented in TBPM for agent selection, as it is not always possible to find an agent whose capabilities match exactly the requirements of a task.

7. Conclusions and Future Work

By carrying out an in-depth analysis of the NPD process with the help of industrial input, many requirements for supporting adaptive workflow were identified. To address the requirements, a novel framework was proposed and a prototype was implemented to demonstrate the feasibility of the approach. Although NPD was used as a test case application in this project, it is not the only cross-functional process that could benefit from the TBPM approach. Ongoing design and production involves sales, marketing, purchasing, production, delivery, engineering and maintenance functions, which must all work together and in parallel for the company to be a quick and efficient provider. As competition increases, speed of response and flexibility are becoming more and more important. This is accelerating particularly with the current focus on e-commerce and e-business. This revolution in the way business is done will require ever more flexible processes, which this technology will have great value in enabling.

With the increasing tighter legislation on health and safety, certain design and manufacturing processes have to conform to related national and international standards. There is on the one hand, the demand for flexibility and, on the other hand, the need to conform to standards. Therefore, there is the need for support for process compliance checking as well as support for change. A research project at Loughborough University is being carried out to pursue this line of investigation.

Acknowledgements

This work was carried out jointly between Loughborough University and the Artificial Intelligence Applications Institute at the University of Edinburgh, under EPSRC research grants numbers GR/L42346 and GR/L42179 on the “Systems Engineering for Business Process Change” programme. The authors would like to thank former members of the project team for their contributions to this work.

References

- [1] Alonso G, Agrawal D, El Abbadi A, and Mohan C, “Functionality and Limitations of Current Workflow Management Systems”, IEEE Expert, 12(5), 1997.
- [2] Cooper R G, Winning at New Products: Accelerating the Process from Idea to Launch (2nd Ed), Addison-Wesley, Reading MA, 1996.
- [3] Cooper R G, “A Process Model for Industrial New Product Development”, IEEE Transactions on Engineering Management, 30(1), 1983.
- [4] Currie, K., & Tate, A., “O-Plan: the Open Planning Architecture”, Artificial Intelligence, 52(), 1991, 49-86.
- [5] Dellen, B., Maurer, F., and Pews, G.. Knowledge Based Techniques to Increase the Flexibility of Workflow Management. Data and Knowledge Engineering, North Holland, 1997.
- [6] Euzen J-P, Trambouze P, and Wauquier J-P, Scale-up Methodology for Chemical Processes, Editions Technip, Paris, 1993.

- [7] Georgakopoulos, D., Hornick, M., and Sheth, A., 1995, An Overview of Workflow Management: From Process Modelling to Workflow Automation Infrastructure. *Distributed and Parallel Databases*, Vol 3, pp 119-153, 1995.
- [8] Han, Y., Sheith, A., and Bussler, C., 1998, A Taxonomy of Adaptive Workflow Management. Proc. CSCW-98 Workshop Towards Adaptive Workflow System, held during the 1998 Conference on Computer-Supported Cooperative Work in Seattle, USA,
- [9] Jarvis, P., Stader, J., Macintosh, A., Moore, J., Chung P., 1999. What Right Do You Have To Do That?: Infusing adaptive workflow technology with knowledge about the organisational and authority context of a task. First International Conference on Enterprise Information Systems (ICEIS-99), Setubal, Portugal.
- [10] Jarvis, P., Moore, J., Macintosh, A., Stader, J., and Chung, P., 2000. Exploiting Organisational Knowledge in Adaptive Workflow Systems. In: Bustard, D., Kawalek, P., and Norris, M., (eds) *Systems Modelling for Business Process Improvement*, Artech House, 81–94.
- [11] Klein, M., 1996, Challenges and Directions for Coordination Science. Proc. 2nd International Conference on the Design of Cooperative Systems (COOP'96), France.
- [12] Mayer R, Cullinane T, deWitte P, Knappenberger W, Parakath B, & Wells S, IICE IDEF3 process description capture method report (al/tr-1992-0057). Technical Report, Air Force Systems Command, Wright-Patterson Air Force Base, Ohio, 1992. See also <http://www.idef.com/>
- [13] Moore, J., Inder, R., Chung, P., Macintosh, A., Stader, J, 2000. Combining and Adapting Process

- Patterns for Flexible Workflow. International Workshop on Enterprise and Domain Engineering, in conjunction with the 11th International Conference on Database and Expert Systems Applications (DEXA 2000). Greenwich, London, 4-8 September 2000
- [14] Moore, J., Stader, J., Macintosh, A., Casson-du Mont, A., and Chung, P., 1999a. Intelligent Task Management Support for New Product Development in the Chemical Process Industries. 6th International Product Development Management Conference (PDM 99), Cambridge, UK, 5-6 July 1999, 787-796.
- [15] Moore J, Stader J, Chung P, Jarvis P, and Macintosh A, "Ontologies to Support the Management of New Product Development in the Chemical Process Industries", In: Lindeman U, Birkhofer H, Meerkamm H, and Vajna S (eds), Proceedings of the International Conference on Engineering Design, ICED 99. Munich, August 24-26, 1999, pp. 159-164.
- [16] Moore J, Inder R, Chung P, Macintosh A, and Stader J, Who Does What? Matching Agents to Tasks in Adaptive Workflow, Proceedings of the 2nd International Conference on Enterprise Information Systems (ICEIS 2000), Stafford, UK, July 2000, pp181-185.
- [17] Rupietta, W., Organisation and Role Models for Workflow Processes, in Lawrence. P, (ed.) Workflow Handbook, Wiley, 1997.
- [18] Rzevski G, The Integrated Product Development Process: Issues and Methods, Proceedings ICED 93, vol 1, pp. 493-498, The Hague, 1993.
- [19] Schlenoff C (ed), Knutilla A, & Ray S, Unified Process Specification Language: Functional

- Requirements for Modeling Processes, National Institute of Standards and Technology, Gaithersburg, Maryland, 1996.
- [20] Sheth A., 1997. From Contemporary Workflow Process Automation to Adaptive and Dynamic Work Activity Coordination and Collaboration. Workshop on Workflows in Scientific and Engineering Applications, France, September 1997.
- [21] Stader J, Results of the Enterprise Project, Proceedings of the 16th International Conference of the British Computer Society Specialist Group on Expert Systems, Cambridge, UK, 1996.
- [22] Tate A, Roots of SPAR—Shared Planning and Activity Representation, *The Knowledge Engineering Review*, **13**(1), 1998, 121-128.
- [23] Uschold M, and Gruninger M, “Ontologies: Principles, Methods and Applications”, *The Knowledge Engineering Review*, 11(2), 1996, pp. 93–136.
- [24] Uschold, M., King, M., Moralee, S., and Zorgios, Y., 1998, The Enterprise Ontology. *The Knowledge Engineering Review*, Vol. 13. , p 31-89. Also available on World Wide Web at URL <http://www.aiai.ed.ac.uk/~entprise/enterprise/ontology.html>
- [25] van der Aalst, W.M.P., Basten, T., Verbeek, H.M.V., Verkoulen, P.A.C. and Voorhoeve, M. Adaptive Workflow, In *Enterprise Information Systems*, J. Filipe (Ed.), pp63-70, Kluwer Academic Publishers, 2000.