

# Applying Intelligent Workflow Management in the Chemicals Industries

Jussi Stader<sup>1</sup>, Jonathan Moore<sup>2</sup>, Paul Chung<sup>3</sup>, Ian MacBriar<sup>4</sup>, Mohan Ravinranathan<sup>5</sup>, Ann Macintosh<sup>6</sup>

## ABSTRACT

Many organizations are searching for ways to improve their management of new product development. Like most engineering activities, new product development involves dynamic, collaborative processes, undertaken in a context of competing requirements and influences [1]. 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].

There is increasing interest in using knowledge-based techniques to provide more flexible process management support than is possible using current workflow systems [2]. The Task-Based Process Management project (TBPM) investigates the provision of such knowledge-based support for workflow management in the area of new product development within the chemicals industries. It is intended that such support shall contribute to the speed and effectiveness of the product development process in the chemicals industries.

The empirical focus of the project is a case study addressing the scale-up process, which is part of new product development in the chemicals industries. Scale-up is a long-term process of experimentation and design, requiring a high degree of flexibility, and collaboration between specialists from many business and technical disciplines. The nature of the process leads to requirements that present significant challenges for process management. These requirements have led to the design of an architecture for a suitable process management system. The architecture consists of three central components:

1. A modeling system that allows the user to provide information about processes, agents and their organizational context;
2. A planning system that uses knowledge-based planning techniques to assist in the planning of processes, using a library of templates for common processes;
3. A process enactment system that uses a multi-agent paradigm for co-ordination of activities and dissemination of information.

In order to attain the degree of flexibility required, it must be possible to interleave the use of the three components. In addition, the components must be able to reason about the domain in which they are deployed. This is approached through knowledge-rich models, which are based on an ontology [21] that was developed during the project. The ontology provides explicit, well-defined terms covering the relevant aspects of the business domain (business processes, process support and enactment mechanisms, organizational context, and chemical engineering design).

---

<sup>1</sup> AIAI, the University of Edinburgh, UK.

<sup>2</sup> Department of Chemical Engineering, Loughborough University, UK.

<sup>3</sup> Department of Chemical Engineering, Loughborough University, UK

<sup>4</sup> Origin UK Ltd., UK.

<sup>5</sup> Unilever, UK

<sup>6</sup> International Teledemocracy Centre (ITC), UK..

The TBPM project is a collaborative project between AIAI, the University of Edinburgh and the Chemical Engineering Department at Loughborough University. ICI and Unilever are industrial partners on the project, providing real business requirements and in-depth knowledge of the application domain.

## INTRODUCTION

As the UK chemicals industry increasingly moves away from bulk commodity chemicals production and toward specialty chemicals there is an increasing need to rapidly develop and bring to market new products. The effective management of the new product development (NPD) process is essential for the survival of a business in today's marketplace. There is significant competitive advantage to be achieved from being first to market with a product, both from the opportunity to establish market share, and to make the most of today's limited product life span in the market.

However, like most engineering activities, new product development involves dynamic, collaborative processes, undertaken in a context of competing requirements and influences<sup>16</sup>. There is a very large number of interacting processes that must be completed before a product can be brought to market, and these are the province of a number of different disciplines and departments, including R&D, Engineering, Marketing, Finance, Safety Health and Environment and General Management. The effective management of the NPD process must achieve good collaboration and communication between disciplines and effective implementation of the necessary processes. Failure to manage NPD effectively can have a range of adverse consequences, from a failed product [2], with associated financial consequences for the business, to an incompletely studied chemical process, with potentially dire consequences for safety or for the environment.

It is the aim of workflow to support the management of processes and workflow management systems are widely used for the streamlined management of "administrative" processes. However, the NPD process and many other engineering activities have characteristics that pose significant challenges for workflow systems:

- It is a highly interdisciplinary process, requiring the co-ordination of individuals from many different engineering and business specialities.
- Many ad-hoc processes occur, which nonetheless are long-time activities and require specific technical and business skills and resources to perform, therefore needing careful management.
- The structure of the process followed is highly flexible—varying from one project to the next, so that each is unique.
- The process is information-intensive: a significant quantity of technical information of different types is generated and must be distributed to interested parties reliably and efficiently.

These characteristics make NPD unsuitable for management by conventional workflow systems, which are unable to cope with the more dynamic situations encountered in ad-hoc and collaborative processes (1, 6, 10). However, if such support could be provided for the NPD process, there are potential benefits to the speed and effectiveness of the process:

- By providing a single framework of computer tools allowing the planning, execution and monitoring of processes. This ensures that the process followed faithfully reflects the process planned, and allows the inspection of information about the current status of the process.
- By 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.
- By improving the quality of decision-making by 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 such potential advantages, there is increasing interest in making workflow more adaptive (7, 18) and in using knowledge-based techniques to allow workflow to cope with complex and dynamically changing processes (4, 19).

The Task Based Process Management (TBPM) project is one of these investigations. Its knowledge-based techniques are based on explicit representations of knowledge about the task management process 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 for the management of such complex processes.

#### SCENARIO AND REQUIREMENTS

The specific example considered in the TBPM project is the scale-up process that forms part of NPD in the chemical process industries. The scale-up process has been investigated by a process of knowledge acquisition, lead by members of the TBPM project team and involving people (engineers and scientists) who are routinely involved in the process. This knowledge acquisition investigation has produced a composite description of the process and its context that is realistic in scope and complexity, without reflecting in its precise detail the practice of a single business. The purpose of the investigation was twofold:

1. To identify requirements which must be satisfied by the computer support system if it is to be able to manage the scale-up process satisfactorily;
2. To develop the knowledge and models necessary to implement a prototype system for managing scale-up processes as a test case and demonstration for the TBPM system.

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. During scale-up, a chemical manufacturing process that has shown promise in the laboratory is developed to a scale suitable for the industrial manufacture of the product in suitably large quantities for commercial exploitation.

Scale-up involves investigating the behavior of this proposed chemical process, and the nature of the engineering necessary to implement it at the intended scale of production. Many complex technical issues must be investigated during scale-up, since the reaction at industrial scale will typically differ markedly from that at laboratory scale, and the engineering issues related to the full-scale process need to be identified and resolved. This investigation is performed through a series of experiments at gradually increasing scale, starting in the laboratory and ending (if all proves satisfactory) with a working pilot plant.

Concurrently, the strategic and marketing concerns of the business need to be kept in mind, updated, and refined. Scale-up projects are typically long-term, involve cooperation among experts from a diverse range of scientific, engineering, and business disciplines, and give rise to many ad-hoc processes. As with all new product development processes, they are conducted under a high degree of time-pressure, caused by the desire to bring new products to market quickly.

The overall process of which scale-up forms a part is reminiscent of that proposed by Cooper<sup>3</sup>, in that 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.

#### *Requirements*

In developing the scale-up scenario, 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 dynamic processes such as scale-up. These requirements are outlined below.

#### *Flexibility*

There are almost no fixed rules for process management in scale-up. The complexity of the domain, and the unpredictability of the types of projects and difficulties that may arise from time to time preclude the *a priori* specification of fixed, unbreakable 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. About the only absolute, inviolable rules are those touching on safety: if a manager allowed an experiment to be conducted without its having first undergone the required safety assessment and approval, he could expect to be severely disciplined.

To take account of this, 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. A project manager (or manager at another level) must thus be able to re-specify the process to be followed at any stage, according to the prevailing conditions. By ensuring that it requires a positive decision to countermand organizational norms, the system can ensure that the manager is fully aware of, and takes responsibility for, such a decision.

#### *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, reflecting a corpus of best practice and experience within the industry for that specific type of project. At a more detailed level, many of the engineering and scientific activities follow some identifiable type of basic process, again reflecting common best practice and similar technical training of participants: see, for example<sup>5</sup>, which provides recommended methods for various types of experimentation encountered during scale-up. There is a need to provide expressions of such common process structures as a resource for those setting up a specific process. However, such common structures should not be rigidly enforced: a manager should be able to select the template process that best matches the current situation, then change it in order to specialize it to match even more closely.

#### *Delayed Planning*

For all but trivial tasks, the process to be followed cannot be completely specified *a priori*. While certain characteristics and activities can be predicted, much of the process cannot be fully specified at the start, since it requires information that becomes available only 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 that will be required are typically not 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. This means that it must be possible to interleave the planning and execution of processes, deferring the complete specification of later stages while earlier stages are being executed.

#### *Agent Representation and Selection*

Many individuals involved in the scale-up process are highly skilled and have very specific capabilities. "Mechanical engineers," for example, cannot simply be lumped

together as a category, and treated as interchangeable. Each engineer is highly skilled, and has an individual mix of qualifications, qualities and experience that may be suitable to perform specific parts of the process. This means that a system that is to assist with the selection of the most suitable person to perform a task must support expressive representations of the capabilities of agents, and their comparison against similarly expressive representations of the capabilities that a task requires.

#### *Management at Different Levels*

Different stakeholders in the process may work at different levels of detail, both for setting up a process and during its execution. For example, the project manager is interested in the overall project structure, but may delegate management of the details of the engineering and research aspects of the project to other managers who are experienced in each of these areas. Each of these managers may then further sub-divide their respective portions of the project according to technical discipline, delegating each section to a relevant specialist. An intelligent process management system must support such hierarchically-structured distribution of management responsibility.

#### *Information Management*

Different technical disciplines are involved in different parts of the process (the engineers do the plant design, the chemists do the experimentation). The different disciplines must communicate effectively because the work of one discipline may impact on the work of another (engineering decisions may have implications for packaging which in turn affect the marketing of the product and potentially the financial viability of the product).

Relevant information should be made available to those who need it as soon as it is produced so that there are no unnecessary delays, or unpleasant surprises discovered late into the project. However, it is often not clear who needs to be made aware of information that is becoming available, particularly in situations where information must be passed between disciplines (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.

#### *A Working Scenario*

To illustrate some of the requirements, consider a practical example from the scale-up process: an overall high-level plan for the process can be specified, but at various points during the process a requirement may be identified for previously unplanned experimentation to be performed. Planning such experiments would typically fall to the research manager on the project. The manager decides on the amount and complexity of engineering activities needed: a straightforward experiment at the laboratory bench would require essentially no engineering input, while an experiment requiring specialized laboratory-scale test apparatus (a “lab-rig”) or even a pilot plant would entail inclusion of some engineering design and construction activities.

Having selected an appropriate strategy for the task in hand, the manager needs to adapt it to the current situation. There may be activities that are normally carried out but that are obsolete in the current situation. For example, if an appropriate lab-rig design already existed (perhaps having been generated for a previous project which did not progress to construction of the rig), then the existing design can be reused rather than designing a new lab-rig from scratch.

All experimentation processes share one important characteristic: the experiment to be performed must be specified, and safety approval obtained for the experiment to be run as described on the specified equipment, before the experiment may be performed. However, there are different forms of experiment safety approval, depending on the nature of the experiment and the equipment used. A simple experiment at the laboratory bench may require only a relatively simple assessment of the dangers inherent in the use of the intended chemicals, while a more complex experiment (conducted at elevated temperatures and/or pressures, for example) might require a more in-depth experimental safety assessment. At the top end of the scale, some lab-rigs and all pilot plants would require a full hazard and operability study (HAZOP). So, each activity in the experimentation process may require further elaboration and specialization before it can be performed.

When the plans for the scale-up project are put into action, suitable people must be found to perform the different parts of the project. In scale-up, the relevant staff are mostly highly skilled individuals who are specialists in their particular area of chemistry. In addition, in order to provide continuity, there are roles that have to be filled for the lifetime of the project, like “project manager” and “project safety officer,” many of which require the person to have a certain amount of authority.

#### TBPM'S APPROACH

The approach adopted by the TBPM project is based on work carried in the Enterprise project<sup>19</sup>, and centers around an intelligent workflow engine, the Task Manager. The Task Manager includes a process modeling tool and an interactive process planner. The planner uses artificial intelligence techniques to assist in the planning of tasks, 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. All system components operate on the basis of knowledge-rich models of the processes and the organizational context within which they are carried out.

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 software), who carry out tasks specified in the models.
- TBPM software components; in particular the Task Manager, 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*. Uschold and Gruninger<sup>21</sup> 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.

#### TBPM'S ONTOLOGIES AND MODELS

There are several key features of an ontology that are relevant in this context:

- 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 reuse 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 accruing 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 (17, 20) and previous work is being built on (22, 19). 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. This 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 that specialise the general concepts of the process ontology for the particular application domain in question (here, scale-up).

Both sets of ontologies are described in more detail below. The TBPM Task Manager depends only on the process ontology, which makes the system independent of the particular 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 13 and 14.

The process ontology defines the concepts central to the handling of all business processes. Key concepts are:

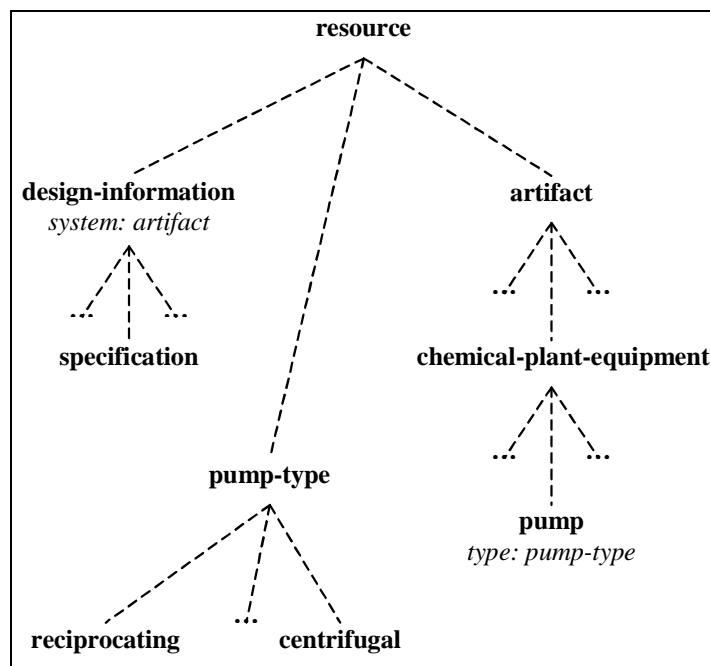
- **Task Type:** A task type is a specification of the basic nature of a task, in terms of the purpose for which the task is carried out.
- **Role:** A role represents a coherent body of work within a task, all of which is expected to be carried out by the same agent. Key aspects of a role are its remits (a set of task types, specifying the types of task which an agent is expected to carry out by virtue of its filling the role) and its authorities (again a collection of task types, in this case specifying the types of task for which the agent can authorise the execution by virtue of its filling the role).
- **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 which that agent can perform, each 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.
- **Organizational Unit:** TBPM borrows heavily from the Enterprise Ontology<sup>22</sup> for its organizational modelling primitives, which are based on the concept of organizational units, between which management and ownership relationships may exist, and which may own resources.

Specializing 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 specializations of a single term from the process ontology:

- **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 objects (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.

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 specialisations. Figure 1 shows some terms from the resources section of the hierarchy for scale-up.



**Figure 1.** Some terms from the resources section of the ontology, illustrating the hierarchical structure. "Resource" is a term from the process ontology, the other terms are specializations of "resource" defined in the domain ontology for scale-up.



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. In the example in Figure 1, each term, which is a specialisation of “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))

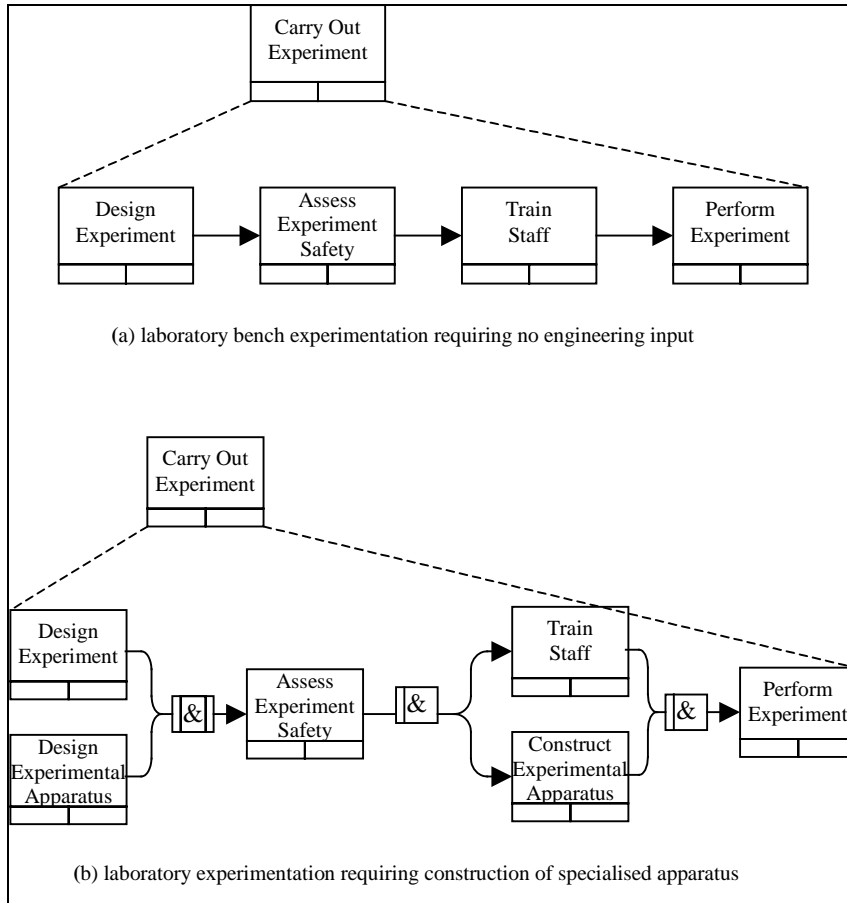
can be generated to represent specifications for different types of pump. This structuring of the ontology represents a compromise between expressive power and likely acceptance by users, and is discussed further in 12.

The process and domain ontologies define a framework in which TBPM’s models can be expressed. TBPM maintains several main types of model that reflect the state of the organization and the processes within it. Models of the *organizational structure* specify the division of the organization into organizational units, in terms of “manages” relationships between units, and “owns” relationships between organizational units and resources. An *Agent Registry* is used to model and manage the agents (both human and software) that are relevant to the system, maintaining information about their capabilities.

The most complex of the TBPM models are task specifications and plan specifications, both held in the *Plan Library*. A task specification in the library consists of a task type, specifying the basic purpose to be achieved, together with a number of resource specifications for both the inputs and outputs of the task, and a number of roles that are expected to be filled for the task to be accomplished successfully.

Each task has at least one role: the “owner” role, which indicates responsibility for the management of the task. A task may have any number of additional roles, each specifying a coherent body of types of work that are expected to be involved in the conduct of the task, and which should be performed by a single agent. A typical scale-up project, for example, will involve a project manager role (the owner role for the overall project), together with roles such as engineering manager, research manager, project accountant, project safety manager and project planner. Each represents a distinct body of work on the project that is expected to be carried out by a single agent (for small projects, a single agent may fill several of these roles, but the roles still remain distinct).

A plan specification details one possible way of breaking down a single task into sub-tasks. Such a plan specifies precedence relationships and flow of resources between tasks. Figure 2 shows a somewhat simplified illustration using the IDEF3 process description notation<sup>11</sup>. 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).



**Figure 2. Two alternative process fragments for experimentation activities**

Each plan provides only a single level of breakdown, but each constituent task of the plan is in itself a task specification, for which a number of plans may exist in the library. In this way, a hierarchical process structure may be generated by successive refinement of tasks at increasing levels of detail.

The plan library can be updated by the addition or alteration of plans at any time. As soon as they are made public, the changes made will show up in the options presented to a task owner seeking to locate a suitable plan for the task.

More details on TBPM's models can be found in [8]. 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.

#### TBPM'S TASK MANAGEMENT

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.
- *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 that 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 organizations, there are many business processes for which this approach is not feasible. Some processes may be too complex or unpredictable to fully specify in ad-

vance, others have to be performed in dynamic environments that 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 organization'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.

### *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 that 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. There is a process-modeling tool that 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.

### *Planning*

Initial planning support is provided by the integration of the process-modeling tool with the Task Manager, and the provision of the Plan Library described above. The process modeler 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 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 2 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 specialize 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 decision: that of locating suitable agents to carry out or participate in a task. Agents participate in tasks by virtue of their 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 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, she would be expected also to take responsibility for the design of the mechanical handling systems for the pilot plant stage of the project.
- 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 remit 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 that the system does not regard as suitable (the available models may not cover everything. There are likely to be cases where the task owner knows more than the system about who should be chosen).

See 15 for a more detailed description of the selection of agents for roles within tasks.

#### *Execution*

The TBPM Task Manager supports the coordination of the execution of the overall process by tracking the status of tasks within it, maintaining sequencing 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 coordinated 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 recognized as computer-based operations (edit, view, etc.), and the input is data of a known format. A registry of tools that can be invoked to support users in carrying out such operations is maintained, and the appropriate tool is invoked automatically in such cases.

#### *Communication*

It is not the intention of TBPM to attempt to capture all the communication that occurs during the planning and execution of a process. In a process such as scale-up, there is much communication of an informal nature, between participants who are used to working closely with each other. However, there are two types of more formal communication that naturally form part of the system:

##### *Task management communication*

Agents communicate with each other in the process of setting up and executing a process. Since the nature and outcome of such communication is critical to the conduct of the process, it is conducted through, and recorded within the task management system.

A task's owner may wish to delegate the task to another agent, or may want to assign other agents to fill the non-ownership roles in the task. Once an appropriate agent for such a role has been located with the help of the system, the role is offered to the agent through the system. Negotiation about whether the agent accepts the role, and on what basis, may then occur: TBPM does not attempt to capture the details of the negotiation, but records the resulting acceptance or rejection of the offered role. Ideally, TBPM would provide more support for negotiation, and should support different default styles of negotiation, depending on the nature and/or relative positions within the organization of the agents (this issue is discussed further in 9).

Agents also communicate indirectly through the completion of tasks and provision of resources on which other agent's tasks depend. This is all handled by the Task Manager, and the owner of a task informed when the status of some important part of the task or its resources changes.

*Domain-specific communication*

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 that depend on it.

DISCUSSION AND CONCLUSION

TBPM's approach is to use ontology-based knowledge-rich models to improve a workflow system's understanding of processes and their contexts. Together with its strategy to interleave the three phases of workflow (specification, planning, execution), this makes the TBPM system powerful and flexible enough to provide effective support for dynamic and complex processes. In brief, the requirements identified in the scale-up scenario are covered by the TBPM approach in the following ways:

- **Flexibility:** the combination of explicit, independent representations, adaptable process templates, and the general just-in-time feel of the TBPM approach provide an extent of flexibility that is suitable for processes such as NPD (see also below).
- **Common Processes:** the Plan Library provides support in the form of pre-specified plans that can be assembled and adapted to suit the specific needs of the current situation
- **Delayed Planning:** TBPM's interleaved Specification-Planning-Execution cycle allows parts of the plans to be specified while the overall process is in progress. TBPM places only the minimum constraints on the cycle, i.e. that any individual task must be specified before it is planned and planned before it is executed.
- **Agent Representation:** TBPM uses a powerful, finely-grained representation of capabilities that is based on a capabilities ontology. The hierarchical structure of this and all ontologies allows the specification of capabilities at any desired level of detail, still being able to take advantage of their higher-level categorisation and more general characteristics.
- **Management at different levels:** TBPM's plan specifications are hierarchical and tasks can be viewed, managed, and delegated at any level of detail, providing full support for this requirement.
- **Information Management:** TBPM has explicit models of information in the form of resource specifications. Agents can register interest in information, and tasks can specify both information they require and information they

produce. Like all models, these are based on an ontology of resources which means that during process execution information requirements can readily be matched to the available information in an intelligent way.

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 organization and its 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 organization. For example, if an organization 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 organizational structure. Such changes can be actively supported by TBPM because the structures and processes involved in the change are modelled explicitly. If the organization wants to change a business process, TBPM can be used to model the new process and changes can take immediate effect in the organization.

Changes are localized 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 organization'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. The use of ontologies in combination of intelligent matching algorithms makes this possible. Note, however, that TBPM would not attempt to actively support such change: it is not aimed at managing the availability of agents (e.g. working out holiday schedules). The system simply ensures that it can operate in the presence of such change, thus accommodating it.

New product development is not the only cross-functional process that could benefit from the TBPM approach. Ongoing 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.

#### ACKNOWLEDGEMENTS

This work has been carried out jointly between the Artificial Intelligence Applications Institute at the University of Edinburgh and Loughborough University Chemical Engineering Department, under EPSRC research grants numbers GR/L42179 and GR/L42346 on the "Systems Engineering for Business Process Change" program. 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. Dellen, B., Maurer, F., and Pews, G., 1997. Knowledge Based Techniques to Increase the Flexibility of Workflow Management. Data and Knowledge Engineering, North Holland.
5. Euzen J-P, Trambouze P, and Wauquier J-P, "Scale-up Methodology for Chemical Processes, Editions Technip, Paris, 1993.
6. 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.
7. 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,
8. 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 organizational and authority context of a task. First International Conference on Enterprise Information Systems (ICEIS-99), Setubal, Portugal.
9. Jarvis, P., Moore, J., Macintosh, A., Stader, J., and Chung, P., 2000. Exploiting Organizational Knowledge in Adaptive Workflow Systems. In: Bustard, D., Kawalek, P., and Norris, M., (eds) *Systems Modelling for Business Process Improvement*, Artech House, 81-94.
10. Klein, M., 1996, Challenges and Directions for Coordination Science. Proc. 2<sup>nd</sup> International Conference on the Design of Cooperative Systems (COOP'96), France.
11. 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/>
12. 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
13. 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.
14. 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.
15. Moore J, Inder R, Chung P, Macintosh A, and Stader J, "Who Does What? Matching Agents to Tasks in Adaptive Workflow," submitted to: International Conference on Enterprise Information Systems, Stafford, UK, July 2000.
16. Rzevski G, "The Integrated Product Development Process: Issues and Methods," *Proceedings ICED 93*, vol 1, pp. 493-498, The Hague, 1993.
17. 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.
18. 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.

19. Stader J, Results of the Enterprise Project, Proceedings of the 16<sup>th</sup> International Conference of the British Computer Society Specialist Group on Expert Systems, Cambridge, UK, 1996.
20. Tate A, Roots of SPAR—Shared Planning and Activity Representation, *The Knowledge Engineering Review*, **13**(1), 1998, 121-128.
21. Uschold M, and Gruninger M, “Ontologies: Principles, Methods and Applications,” *The Knowledge Engineering Review*, 11(2), 1996, pp. 93-136.
22. 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>