INTELLIGENT TASK MANAGEMENT SUPPORT FOR NEW PRODUCT DEVELOPMENT IN THE CHEMICAL PROCESS INDUSTRIES

Jonathan Moore[†], Jussi Stader[‡], Ann Macintosh[‡],
Andrew Casson-du Mont[‡], Paul Chung[†]

[†] Chemical Engineering Department, Loughborough University, Loughborough, UK
{J.P.Moore, P.W.H.Chung}@lboro.ac.uk

[‡] Artificial Intelligence Applications Institute, University of Edinburgh,
80 South Bridge, Edinburgh, UK
{jussi.stader, ann.macintosh, andrew.casson-dumont}@aiai.ed.ac.uk

ABSTRACT

Current workflow management systems are not able effectively to manage complex, dynamic, collaborative processes such as those which characterise new product development (NPD). The TBPM project is investigating the extension of workflow approaches to support such processes using techniques from artificial intelligence. In this paper we describe the scale-up process which forms part of NPD within the chemicals industries, and the requirements which must be met by intelligent workflow systems if they are to manage this process successfully. We then outline the architecture of the prototype intelligent workflow system which is being developed as part of the project.

INTRODUCTION

The need for intelligent workflow

Workflow technology has been successful for supporting "administrative" processes (see Alonso et al [1]) for some time. Administrative processes are characterised by their clear, well-defined structure and consistent, predictable form, which enable workflow implementors to predict, and explicitly cater for, every possible route which may be taken through the process.

However, in the past workflow technology has not been successfully applied to complex and dynamic processes such as those typically involved in new product development (NPD), because current systems have not proven sufficiently flexible. Recently, techniques from artificial intelligence are being used to make workflow more adaptive and to allow it to cope with complex and dynamically changing processes (see, for example, [2] and [3]). Such work opens up scope for computer support for the management of processes such as NPD. The provision of such support may be expected to contribute to the speed and effectiveness of the NPD process in several ways:

- 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 makes information about the current status of the process available at all times.
- 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 these potential advantages, there is increasing interest in the use of techniques from artificial intelligence to endow workflow systems with explicitly represented knowledge of the task management process in general and 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.

The TBPM project

The TBPM (Task-Based Process Management) project investigates the provision of intelligent computer-based support for the scale-up process which forms part of NPD within the chemical process industries. TBPM is a joint project between the Chemical Engineering Department of Loughborough University and the Artificial Intelligence Applications Institute of the University of Edinburgh, with Unilever Research and ICI Research and Technology as industrial partners. There are two major themes to the project:

1. A theoretical and technical approach using techniques from artificial intelligence to extend the current state-of-the-art in workflow and develop a prototype workflow system capable of using knowledge-based techniques to manage complex, dynamic tasks. This approach is taking a generic view of task management, not tied to any single example or domain.
2. A practical investigation of the scale-up process which forms part of NPD within the chemicals industries, in order both to identify requirements for the workflow system and to provide a practical test-case.

This paper describes the project mainly from the point-of-view of the second of these themes. An outline of the scale-up process is given, and the requirements which it imposes on a task management system identified. The means by which these requirements are being addressed in the TBPM project are then discussed.

THE SCALE-UP PROCESS AND REQUIREMENTS FOR MANAGING IT

The scale-up process

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 which 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 process models and the knowledge-bases necessary to implement a prototype system for managing scale-up processes as a test-case and demonstration for the task management 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.

Scale-up involves investigating the behaviour of this proposed chemical process, and the nature of the engineering necessary to implement it at the intended scale of production. 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. Scale-up projects are typically long-term, involve co-operation amongst experts from a diverse range of scientific, engineering, and business disciplines, and give rise to many ad-hoc processes.

The overall process followed for NPD is reminiscent of that proposed by Cooper [4], 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

The major requirements for a task management system capable of handling the scale-up process are discussed below.

<u>adaptability</u>

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 scale-up itself. Thus, in order to be generally applicable across a range of different scale-up projects, process models cannot be expressed in sufficient detail that they can be slavishly followed by a workflow system in order to implement a scale-up process. However, they *can* be expressed at a more abstract level, and then be interpreted in an appropriate manner for the specific context in which they are being applied.

<u>mixed initiative</u>

The management of a scale-up project is an expert task, requiring experience, judgement, and knowledge of the domain in order to implement a given process model effectively in a particular situation. The range of expertise required is wide, and we cannot with the current state of technology hope to capture it all completely in a computer system. We must instead employ a mixed-initiative approach involving both the human manager and the system. We focus on providing as much support as possible for the process management decisions which must be made, providing sensible defaults and helpful suggestions, while allowing the human manager of the process full freedom to change these if necessary for the task in hand.

<u>capability matching</u>

Once the process to be followed for a particular task has been decided, the individual activities which comprise that process must be performed. In scale-up, many of these activities are highly technical, and require specific skills and experience of the agent chosen to perform them (we use the term “agent” to denote activity performers, be they 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.

organisational context, authority, and responsibility

The system must respect organisational norms while handling a process. For example, it should not delegate a task to a specific person without the assent of that person's manager, if such assent would normally be regarded as necessary. Similarly, if a person has 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 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.

information handling

A commonly encountered problem during scale-up is that people are given information that is crucial for their work too late (or even not at all), so that work may be wasted or delayed unnecessarily. For example, experimentation in the laboratory may reveal that a different grade of raw material is needed than originally thought (maybe of a higher purity). This may appear a straightforward change to the scientists involved, and not be explicitly brought to anyone’s notice. However, the effects of such a change on those considering the sourcing of raw materials and establishing a supply chain may be profound, and they need to know of such discoveries at the earliest opportunity. In order to support more effective communication, the system needs to keep track of the information generated during a process, and have some representation of the nature of that information so that it can be distributed to the agents who need it.

domain knowledge

The specifications of capabilities and of information will both need to refer to concepts from the engineering domain—to distinguish, for example, 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*. For managing scale-up, therefore, we need explicit representation of knowledge about the domain of chemical process engineering.

THE TBPM TASK MANAGEMENT SYSTEM

In this section we outline the architecture of the intelligent adaptive workflow system being developed by the TBPM project and indicate how it addresses the requirements listed above. The structure of the system is illustrated in Figure 1.

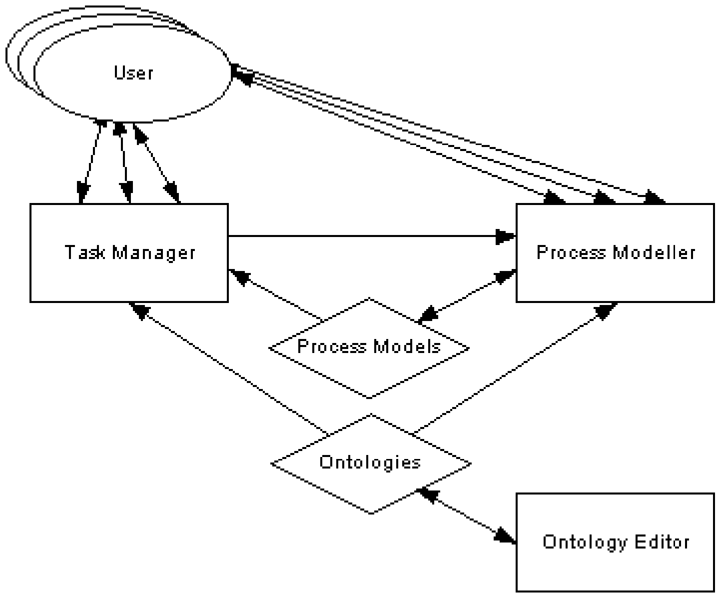

Figure 1. the architecture of the TBPM task management system

Ontologies

Apart from the requirements for adaptability and mixed initiative, the remaining requirements identified above specify the types of knowledge with which the system needs to be able to reason. Jarvis et al. [5] outline how such knowledge can be structured and represented in a workflow system. These requirements are being addressed by the development of a number of related *ontologies*. An ontology is defined by Uschold & Gruninger [6] as a "…shared understanding of some domain of interest…" The TBPM system uses two types of ontology: domain-independent ontologies that describe concepts related to task management in particular, and domain-dependent ontologies which describe concepts specific to scale-up, building upon the domain-independent ontologies.

<u>domain-independent ontologies</u>

A core set of high-level ontologies provides the knowledge necessary for the system to reason about processes in a generic way. Wherever possible, existing ontologies or standards are being used. The ontologies include descriptions of the nature of:

- Processes (building on emerging standards such as [7] and [8] for the representation of processes).
- Organisational structure (we adopt the *Enterprise Ontology*—see [9]).

- Capabilities (building on work arising from the Enterprise project—see [3] for details).
- Authority (see Jarvis et al. [5] for a description of the approach adopted to the modelling of authority within an organisation).

domain-specific ontologies

The core ontologies described above provide the fundamental knowledge which a workflow system needs in order to be able to reason about processes in the required manner. For the system to be useful within a specific process domain, each of these ontologies must be specialised for the domain in question. For our scale-up case study:

- We build on the core organisational and authority ontologies to describe a hierarchical, corporate organisation and authority structure.
- The process ontology is used to define a library of commonly occurring activities—for scale-up, largely activities associated with experimentation and engineering design—with which the task management components of the system may work.
- The capabilities ontology is extended to describe the capabilities of participants in the scale-up process, such as engineers (with capabilities such as *design* and *analysis*).

Moore et al. [10] give an overview of the domain-specific ontologies used and the relationships between them.

TBPM task management

At the heart of the system is the *task manager* component. This is the intelligent workflow engine which drives and manages the processes. It is integrated with a process editing tool which allows users to access a library of "process fragments"—descriptions of common activities which recur frequently in the domain, derived from the domain-specific process ontology and expressed in terms of how each breaks down into lower-level activities. An activity may have more than one description in the library, each representing a different way of carrying it out, and the lower-level activities in each fragment may themselves also have multiple expansions. Thus, the user is able to plan a task by assembling it from "canned" fragments, and then to refine it further by elaborating on the sub-activities in the same manner, generating a hierarchical model of the process to be performed. For maximum flexibility, however, the user is also able to edit the assembled process model both before and during its enactment, in order to specialise it for the particular case in hand.

In order to enact a process, it is necessary to select agents to perform each activity and to co-ordinate the execution of different activities. Specifications of the capabilities, authorities, and organisational roles (expressed in terms of the domain-specific ontologies) that an agent must possess or fulfil in order to perform the activity form part of the definition of each activity in the library of process fragments. A registry of the capabilities of all available agents enables the identification of agents suitable for the performance of the activity. A model of the organisational structure is also available, together with rules for how to infer the authorities and responsibilities possessed by agents from that structure.

A task management example

<u>planning the task</u>

Consider a practical example from the scale-up process: 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, who would select an experimental process from the library. Several such fragments would typically be available, differing mainly in 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 specialised laboratory-scale test apparatus (a "lab-rig") would entail inclusion of some engineering design and construction activities, and an experiment at pilot plant scale even more so.

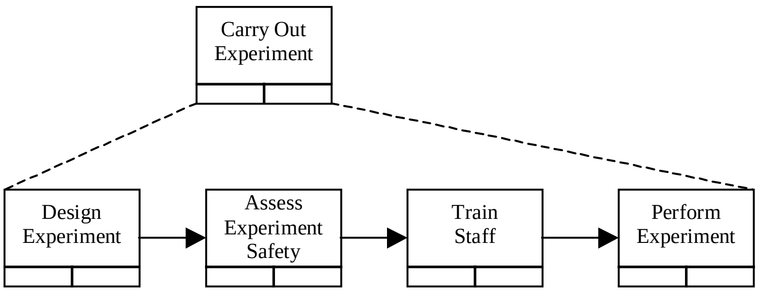

(a) laboratory bench experimentation requiring no engineering input

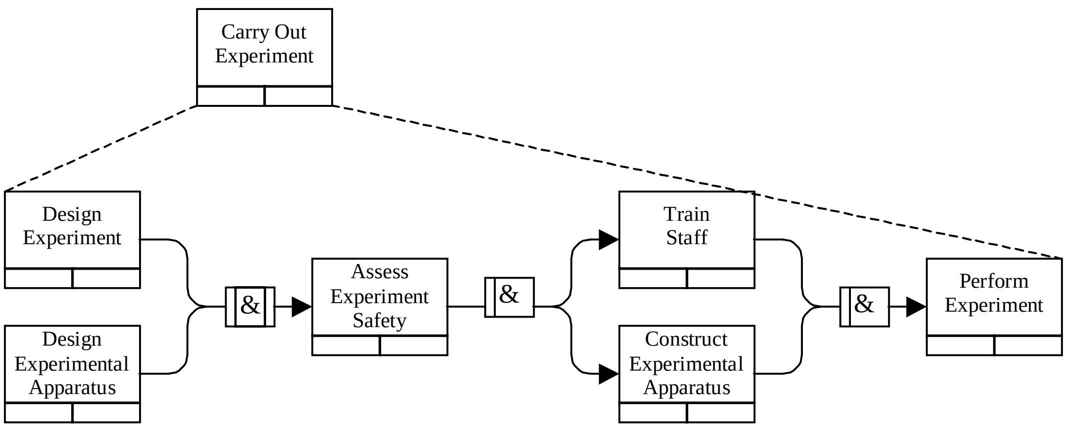

(b) laboratory experimentation requiring construction of specialised apparatus

Figure 2. two alternative process fragments for experimentation activities

Figure 2 shows a somewhat simplified illustration of the "laboratory bench" scale and "lab-rig" scale process fragments for experimentation, using the IDEF3 process description notation [11]. Only the structural breakdown is shown—each activity specification would also include details of the capabilities, authorities, and organisational roles required of an agent in order to carry out the activity.

Having selected an appropriate process fragment for the task in hand, the manager might wish to specialise it for the current situation. If the process fragment for experimentation using a lab-rig were selected, but 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 "design experimental apparatus" activity might be removed or curtailed, and the existing design used.

All experimentation process fragments 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 (i.e. different process fragments), 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, or HAZOP. So, each sub-activity in the experimentation process may require further elaboration and specialisation by the selection of an appropriate model fragment and its editing for the current situation.

The use of a library of process fragments and the ability to specialise a process address the requirements identified above for *adaptability* and *mixed initiative*. Rather than provide a monolithic process model to be followed (and probably fought against) in every situation, the system gives a high level of support for the generation of a model appropriate to the situation by the assembly of normative process models which are provided for common activities. The system retains flexibility by allowing those models to be edited for each situation in which they are applied.

<u>enacting the process</u>

Once the process to follow for the experimentation has been defined in enough detail for it to start, it must then be enacted (note, however, that planning and enactment can be interleaved, so that the process need not be completely specified before it can start). Process enactment involves the selection of appropriate agents to perform each activity, and the performance of the activities in the correct sequence. Precedence links and information flows between activities are a standard part of the process models, and enable the system to sequence activities so that they take place in the correct order, and to co-ordinate the transfer of information between activities.

As indicated above, activity specifications can include details of the capabilities, authorities, and organisational roles which an agent should possess in order to perform a task. Together with a registry of the capabilities, authorities, and roles actually possessed by the available agents, this information enables the system to assist with the selection of appropriate agents to perform particular tasks. 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.

As with the support for planning discussed above, the emphasis in the use of this information is on *support*. The human planner is not forced to accept the recommendations or options presented, but is explicitly made aware of situations when he or she is violating the normal rules for managing scale-up, encouraging the decision to do so to be fully thought through.

CONCLUSION

Although there are potential advantages to be gained by providing computer support for the management of complex processes such as new product development (NPD), current workflow systems are neither powerful enough nor flexible enough for the task. The TBPM project is focusing on enhancing the capabilities of workflow systems using techniques from artificial intelligence to endow the system with knowledge about the domain in which it is operating and the capacity to reason with that knowledge in order to support process management decisions.

The TBPM project approaches the problem through a specific example, i.e. the scale-up process which forms part of new product development (NPD) in the chemical process industries. However, the dependency on this particular process is restricted to the domain ontologies only. The task management and process specification components, and the core ontologies on which they depend, are generic, and we expect the approach to be straightforwardly generalisable (although requiring the development of domain-specific ontologies in each case) to other processes, both within and without NPD.

ACKNOWLEDGEMENTS

This work has been carried out jointly between Loughborough University Chemical Engineering Department and the Artificial Intelligence Applications Institute at the University of Edinburgh, under EPSRC research grants nos. GR/L42346 and GR/L42179 on the "Systems Engineering for Business Process Change" programme. We would like to thank Ian McBriar of ICI Research & Technology, Stuart Moralee and Mohan Ravindranathan of Unilever Research, and Lee Young of BG Technology for their assistance.

REFERENCES

[1] Alonso G, Agrawal D, El Abbadi A, & Mohan C, Functionality and Limitations of Current Workflow Management Systems, *IEEE Expert*, **12**(5), 1997.

[2] Dellen B, Maurer F, & Pews G, Knowledge-Based Techniques to Increase the Flexibility of Workflow Management, *Data and Knowledge Engineering*, North Holland, 1997.

[3] 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.

[4] Cooper R G, A Process Model for Industrial New Product Development, *IEEE Transactions on Engineering Management*, **30**(1), February 1983.

[5] Jarvis P, Stader J, Macintosh A, Moore J, & Chung P, What Right Do You Have To Do That? Infusing Adaptive Workflow Technology with Knowledge about the Organisational and Authority Context of a Task, to appear in *Proceedings of the First International Conference on Enterprise Information Systems (ICEIS-99)*, Setubal, Portugal, 1999.

[6] Uschold M & Gruninger M, Ontologies: Principles, Methods and Applications, *The Knowledge Engineering Review*, **11**(2), 1996, 93-136.

[7] 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.

[8] Tate A, Roots of SPAR—Shared Planning and Activity Representation, *The Knowledge Engineering Review*, **13**(1), 1998, 121-128.

[9] Uschold, M., King, M., Moralee, S., & Zorgios, Y., The Enterprise Ontology, *The Knowledge Engineering Review*, **13**, 1998, 31-89.

[10] Moore J, Stader J, Chung P, Jarvis P, & Macintosh A, Ontologies to Support the Management of New Product Development in the Chemical Process Industries, to appear in *Proceedings of the International Conference on Engineering Design (ICED)*, Munich, August 1999.

[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/