WHO DOES WHAT?
MATCHING AGENTS TO TASKS IN ADAPTIVE WORKFLOW

Jonathan Moore†, Robert Inder‡, Paul Chung†, Ann Macintosh*, and Jussi Stader‡

† *Department of Computer Science,*
Loughborough University, Loughborough, UK

‡ *Artificial Intelligence Applications Institute*
Division of Informatics, The University of Edinburgh, Edinburgh, UK

* *International Teledemocracy Centre,*
Napier University, Edinburgh, UK

J.P.Moore@lboro.ac.uk,R.Inder@ed.ac.uk,P.W.H.Chung@lboro.ac.uk
A.Macintosh@napier.ac.uk, Jussi.Stader@aiai.ed.ac.uk

Abstract

The flexibility of current workflow systems needs to be extended to allow them to operate in more dynamic and uncertain environments. Adaptive workflow systems need to be able to provide intelligent support for the planning and enactment of complex processes; supporting their users in performing flexible and creative tasks, while respecting the norms of the organizations within which they are deployed. We discuss how knowledge about the dynamic context of a process may be represented in terms of roles within tasks, and the remits, authorities, and capability requirements associated with such roles. We outline how this knowledge may be used, together with information about the capabilities of available agents, to support the selection of an appropriate set of agents to fill the roles in a given task. An example is given based on an experimentation task within the chemical process industries.

Keywords: adaptive workflow, process management, capability, remit, authority.

1 Introduction

Workflow systems have proved successful for the management of "administrative" processes—characterized by clear, well-defined structure, and constant, predictable form—for some time. However, there is a consensus that current systems are insufficiently flexible to deal with complex, dynamic processes within a changing context (see Alonso et al, 1997, Sheth, 1997).

One approach to increasing the flexibility of workflow systems is to use knowledge-based techniques to represent the context in which a process is occurring to the system, and thus enable the system to use such knowledge to reason about the process, and to adapt to a changing context (e.g. Dellen et al, 1997).

Jarvis et al (1999) identify key points during the planning and enactment of processes at which an intelligent task management system needs to offer decision support. 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 positions within the organization. Such knowledge is not usually readily available, especially within today's increasingly dynamic and virtual organizations. To provide decision support in this area, we need to represent and capture the knowledge upon which such decisions are based.

We approach the selection of agents for a task from three complementary angles, characterized by the questions:

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

2 Overview of TBPM

The Task-Based Process Management (TBPM) project is investigating the provision of intelligent support for the management of complex, dynamic processes through the use of artificial intelligence techniques to represent, and reason with, knowledge about the domain in which a workflow system is deployed.

2.1 Intelligent Task Management

The approach taken centres around the flexible generation, specialization, and execution of hierarchically-structured process models, using an approach outlined in Jarvis (2000), and based on the provision of a library of possible process plans which may be used to achieve tasks which commonly arise within the domain.

2.2 Industrial Scenario

Alongside this domain-independent approach is a practical investigation of a complex, dynamic engineering process within the chemicals industry. This investigation serves two purposes:

1. To identify and refine requirements for the task management system within such a domain. See Moore et al, (1999a), for a summary.
2. To provide a realistic scenario against which to assess the task management approach taken.

The process investigated is the scale-up process—a combination of scientific experimentation and engineering design which occurs as part of product innovation within the chemical process industries.

2.3 Ontologies

Central to the knowledge-based approach of TBPM is the development of a number of interrelated *ontologies* for structuring knowledge of the domain and processes of interest.

Moore et al, (1999b), describe the different areas in which ontologies need to be developed to support intelligent process management. The most germane to this paper is the ontology of *tasks*, which identifies and describes the classes of task which commonly arise within the domain of interest.

The ontological modelling approach taken structures the important concepts in the domain into a generalization hierarchy. Each term in the hierarchy may specify a number of named parameters, which are take values which are terms from the ontology. This representation provides a simple yet flexible structure for ontological classification, from which it is straightforward to determine whether one ontological class encompasses, or subsumes, another.

The example in section 4 uses the term "experiment", which is a task within the ontology of the scale-up process and which is parameterized by the type of system investigated by the experiment, the scale at which the experiment is performed, and the fundamental technique used to carry out the experiment.

3 Selecting Agents

When a task needs to be performed, one of the crucial decisions to be made is the selection of the correct set of agents to participate in the task. Jarvis et al (1999) set out the motivations for knowledge-based capability matching between agents and task requirements.

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. We refer to this agent as having a *remit* which 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*.

3.1 Roles

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 current 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.

3.2 Remits

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.

3.3 Capabilities

There arise situations in which a task arises for which no encompassing remit exists, or where the agent possessing the encompassing remit cannot, for some reason (such as workload), accept the new task. In such cases, an alternative agent must be sought, and it is here that capabilities can be used to locate a suitable agent.

Like remits, a capability in TBPM identifies a class of tasks. The possession of a capability by an agent indicates that the agent is able to take on (i.e. fill the "responsible agent" role) tasks within that class. The attachment of a capability description to a role indicates a requirement: for an agent to be acceptable to fill the role, the set of capabilities the agent possesses must encompass those specified by the role.

In addition to the class of tasks, each of an agent's capabilities also specifies a level of competence, currently on a discrete scale: "none", "aware", "familiar", "skilled", and "expert". This enables a role's capability requirements to specify a minimum competence level required in addition to the class of tasks involved.

3.4 Authority

There are currently two distinct, though complementary, approaches being taken to modelling authority.

The first, which is not described in detail here, addresses the authority to consume or access a resource, based on a model of ownership of the resource, which may be partitioned and delegated.

The second approach addresses cases where authorization to proceed with a task is not needed because the task involves the use of a resource, but because of some other organizational constraint. A good example in the chemical process industries is safety authorization for a chemical experiment, which must be obtained before any experiment can be performed.

This type of authority is modelled, as with capabilities and remits, by the expression of a class of tasks. An agent filling a role with a particular authority is able to provide authorization for particular instances of that kind of task to proceed. In addition to the class of tasks, an authority specifies a "perspective", enabling the same task to require authorization from a number of distinct points-of-view. For example, a chemical experimentation task may require both "safety" and "business strategy" authorization—the authority to provide which will typically be vested in different roles.

Authority may be inherited down the hierarchical process structure in a similar manner to remits. However, authorities may not be arbitrarily created as remits may: a role's authority can only arise by delegation from an encompassing role, which must possess authority which encompasses that delegated.

4 Example

During the scale-up of a chemical reaction, many experimentation tasks typically arise. One possible such task might be ontologically expressed as:

```
experiment(
    system:    organic_reaction
    scale:     laboratory
    technique: distillation)
competence:    skilled
```

There will typically be several "research chemist" roles established in the top-level activity of the project. One such role, "organic research chemist" might specify a remit such as:

```
experiment(
    system:    organic_reaction
    scale:     any
    technique: standard_laboratory)
competence:    expert
```

Since the remit encompasses the experimentation task which has arisen, the agent filling this role may be identified as the one to whom the task *should* be assigned. A similar matching process between the role's capability requirements and the known capabilities of available agents may be used to find an agent to fill the role in the first place.

The overall process within which the experimentation task arises will almost certainly impose a requirement for safety authorization before the experiment can proceed. Again, there will usually be a "safety manager" role on the project, with an imputed authority such as:

```
experiment(
    system:    any
    scale:     any
    technique: any)
perspective:   safety
```

By searching the roles in the process hierarchy, the individual(s) able to give the required authorization may be identified.

5 Conclusions

The selection of the correct set of agents to participate in a given task is one of the major decisions for which an intelligent task management system needs to provide support. We have described how we may represent knowledge about the remits, capabilities, and authorities required and possessed by roles and agents within a process, and use such knowledge to support intelligent adaptive workflow.

An example has been outlined, drawn from the domain of the chemical process industries.

6 Acknowledgements

The TBPM project is a joint project between the Computer Science Department at Loughborough University, and the Artificial Intelligence Applications Institute at The University of Edinburgh. The project is funded under the EPSRC Systems Engineering for Business Process Change programme, and has ICI and Unilever as industrial partners.

7 References

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

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

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.*

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.

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.

Moore, J., Stader, J., Chung, P., Jarvis, P., and Macintosh, A., 1999b. Ontologies to Support the Management of New Product Development in the Chemical Process Industries. *12th International Conference on Engineering Design (ICED 99), Munich, 24-26 August 1999* 159–164.

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.*