What Right Do You Have To Do That?

Infusing Adaptive Workflow Technology with Knowledge about the Organisational and Authority Context of a Task

Peter Jarvis[ϕ], Jussi Stader[ϕ], Ann Macintosh[ϕ], Jonathan Moore[Ψ], and Paul Chung[Ψ]

[ϕ]Artificial Intelligence Applications Institute,
The University of Edinburgh,
80 South Bridge, Edinburgh, EH1 1HN, UK.

[Ψ]Department Of Chemical Engineering,
Loughborough University, Loughborough,
Leicestershire, LE11 3TU, UK.

{Peter.Jarvis, Jussi.Stader, Ann.Macintosh}@ed.ac.uk
{J.P.Moore, P.W.H.Chung}@lboro.ac.uk

Abstract

Increasingly, workflow systems need to operate in dynamic environments where they are expected to ensure that users are supported in performing flexible and creative tasks while maintaining organisational norms. We argue that in order to cope with these demands, the systems must be provided with knowledge about the organisational structure and authority context of tasks. We support this argument by identifying a number of decision points that an adaptive workflow system must support, discussing how these decisions can be supported with technically oriented capability specifications, and describe how this support can be enhanced with the inclusion of knowledge about organisational structure and authority. We outline how such knowledge can be captured, structured, and represented in a workflow system. We then demonstrate the use of such knowledge by describing how the task initiation, task planning, activity scheduling, and agent interaction functions within a workflow system can be enhanced by it.

Key Words: Adaptive Workflow, Capability Matching, Organisational Modelling, Authority Modelling.

1. Introduction

There is a consensus in the workflow community that the application of systems based upon the unquestioned enactment of process models is inherently limited, particularly in dynamic working environments [Alonso et al. 1997] [Sheth 1997] [Sheth & Kochut 1997] [Klein 1996] [Georgakopoulos et al. 1995]. To move forward, the field needs to develop techniques that support people in making decisions about the activities required to achieve a task, the execution order of those activities, and the specific agents that will perform each activity. In order to provide such support, knowledge about tasks and agents is required. To provide a frame of reference we have identified three main stages of workflow in which the decisions above are made:

- **Task Initiation** is the stage where a user instructs a workflow system to begin a new process. There are choices of which task to initiate.
- **Task Planning** is the stage where the activities required to achieve a task are determined. There are choices of which method to use.
- **Activity Scheduling** is the stage where a specific agent is assigned the role of performing an activity. There are choices of who the activity should be assigned to.

We use these stages throughout the paper to show how different types of knowledge can provide different levels of decision support during the stages.

In this paper, we introduce our approach of “knowledge-based capability matching” which concentrates on the use of knowledge related to capabilities to support the decisions above. We first introduce the basic support that can be provided using knowledge about technical capabilities. However, organisations, are typically arranged into an organisational structure which defines the division of labour and communication channels that together provide a regulated co-operation between its members [Rupietta 1997]. Adding knowledge about this context of an activity in terms of organisational structure and authority, we can provide more sophisticated support. We show how this can be done by outlining a specification schema for describing the knowledge required and discussing how a workflow system can make use of it.

Throughout the paper, the term “task” corresponds to a goal or objective that is to be achieved. Typically to achieve a task, a number of “activities” have to be performed by “agents” (people or software systems). There may be different “methods” specified of how to achieve a task, each consisting of a number of activities. Once a “method” is selected, its constituent activities define the “process” for achieving a particular task.

The work reported here is part of an ongoing applied research programme at AIAI, University of Edinburgh, UK, investigating and developing techniques for supporting complex tasks within dynamic environments. Two projects in particular are relevant: the Enterprise project [Uschold et al. 1998] [Stader 1996] [Fraser & Tate 1995] which was completed in 1997 and the current Task Based Process Management project which is a collaboration between AIAI and Loughborough University, UK. Applications areas considered include the bid management process [Stader 1997] and the product innovation process. Our commercial partners include: IBM, Lloyd’s Register, Logica, Unilever, ICI, and BG.

2. Knowledge-Based Technical Capability Matching

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

Clearly, capability matching requires specifications of capabilities both for activities and for agents. If capability specifications are to be matched, it is important that the specifications use common and well-defined terms which can be related to each other. Our approach during the Enterprise project was to develop a technical capability ontology with our industrial partners which is published as part of the Enterprise Ontology [Uschold et al. 1998]. This technical capability ontology provides a hierarchy of capabilities. For example, databases provide a *Store* capability and - more specifically - relational databases provide a *Store Relational* capability. This hierarchical structure can ease the task of specifying capabilities (required or available) because specifying a high-level capability implies that all its lower-levels are covered too.

However, in our experience organising all terms required for specifying capabilities of agents and activities into such a hierarchy is too big a task for any realistic application area. We decided to impose more structure by splitting the specifications into two parts: the technical capability itself and the area (or “knowledge space”) in which it can be applied. For example, if a specific database application can store data about skills, it can apply its *Store* capability to *Skills*. Each of the parts uses its own hierarchy of terms shown in part in Figure 1.

Capability Ontology

- Storage Capability
  - Store
    - Store Structural
      - Store Relational
      - Store Hierarchical
  - Access
    - Retrieve

Knowledge Space Ontology

- Legal Entity
  - Corporation
  - Shareholder
  - Partnership
- Document
  - Report
    - Technical report

Figure 1: Example Capability and Knowledge Space Elements from the Enterprise Ontology

Using this specification schema in a matching function, the workflow support we implemented during the Enterprise project can not only determine which agents match the capability requirements of an activity exactly, but it can rank all agents available at the time of execution according to how closely they match the capability requirements. Exact matches of a capability specification are best, but agents that can apply the required capability in a wider area than required are nearly as suitable. Similarly, agents that have a more general capability are suitable, although more specialised agents would be preferred because they are likely to perform the activity more effectively.

In summary, by providing a well-defined ontology of capability and knowledge space terms, statements about capabilities can be made and matched consistently. The use of a generalisation structure within the ontology simplifies the specification of capabilities. It can also be exploited by a workflow system to apply “generalist vs. specialist” heuristics and make the best use of the agents available during a task’s execution. Our experience with the Enterprise project shows that a technical capability matching function can underpin the three workflow stages as follows:

- **Task Initiation** Using specifications of technical capability requirements of tasks, users can be assisted in choosing tasks that they themselves can initiate and manage.
- **Task Planning** A technical capability matching function can help a user rank alternative methods by considering how much agent support is currently available for each approach.
- **Activity Scheduling** A technical capability matching function can help a user identify and rank potential performers of an activity so that the most suitable agent can be selected.

3. Motivation for Adding Knowledge about Organisation and Authority

The importance of at least sensitising a workflow system to the organisational structure and authority context within which it operates is well argued in the literature (cf. [Tate 1993], [Joosten 1996], [Kappel et al. 1995], [Dellen et al. 1997], and [Rupietta 1997]). If this context is ignored, the system will undoubtedly break organisational conventions. For example, activities may be assigned to people who are not related to the organisational unit responsible for the overall task. We argue that a workflow system can be much more than just sensitised to organisational structure and authority issues. With this knowledge it can proactively guide a user’s decision making by highlighting how an organisational structure can be navigated and authority constraints maintained. Using our three stages of workflow, the following additional support can be provided:

- **Task Initiation:** After determining that a user has the technical capability to initiate and manage a task, the system can use authority knowledge to determine if the agent has the authority to take this action. If an agent does not have the necessary authority, the system can advise the agent on who to ask for it.
- **Task Planning**: After determining if there are agents technically qualified to complete a task, the system can assess relationships between the agent managing the task and the agents that will be asked to perform activities. Methods for achieving a task can then be ranked using an “organisation and authority metric”. For example, a method where all the agents are employed in the same department as the agent managing the task may be considered preferable to one where the agents are distributed across departments.
- **Activity Scheduling:** After determining the set of agents available for performing a task through technical capability criteria, the set can be ranked according to the authority held by the agent requesting the assignment. For example, it may be considered better to assign activities to agents that one has direct authority over.

Activity scheduling requires agents to communicate not just about results of activities but also about the delegation of activities between agents. With knowledge about organisational structure and authority we argue that support for such communication can be enriched as follows:

- **Interaction Styles:** if a system understands the relative authorities between two individuals, it can assist in the selection of appropriate interaction styles for the communication between them. For example, the system can prioritise entries on a to-do list according to the seniority of the individual requesting the action, and arranging dialogues so that “accept” becomes the default option when responding to a superior’s request.

4. Organisational Structure Modelling Language

A modelling language for describing organisational structure must contain constructs that can be used to model a wide variety of organisations. Our language is based upon the one published as part of the Enterprise Ontology [Uschold et al. 1998] [Uschold 1998]. We are confident in the generality and adequacy of this ontology as it was developed by a working group that included representatives from three international organisations and it is similar to others that have been developed, independently, for similar purposes (cf. [Hoog et al. 1994] [Rupietta 1997]). The language is centred around the *organisational unit* concept which can be used to describe departments, divisions, projects, working groups etc. The definitions in Figure 2 outline the central concepts within our organisational modelling language.

Organisational Unit: An entity responsible for managing the performance of activities to achieve one or more purposes. An organisational unit can be used to describe departments, working groups, project etc.

Agent: An entity that can perform an activity.

Person: A human being.

Machine: A non-human entity that has the capacity to carry out functions. A machine is similar to a person. However, it is anticipated that some functions and roles are exclusive to one or the other. For example a machine cannot be held responsible for anything.

Manages - Organisational Unit to Organisational Unit: An *organisational unit* can *manage* an *organisational unit*. With the relationship, one *organisational unit* takes on the role of the *manger* and the second organisational unit the role of *managed by.*

Figure 2: Definition of the Concepts in the Organisational Modelling Language

Organisational units can be connected by a number of relationships (Figure 3). The *manages* relationship can be used to represent the subdivision of *organisational units*; a committee into working groups, for example. Both *machine* and *person* are *agents* and can be linked to *organisational units* through relationships. A *person,* for example, may be related to an *organisational unit* through the manages relationship, taking the role of a *manager*.

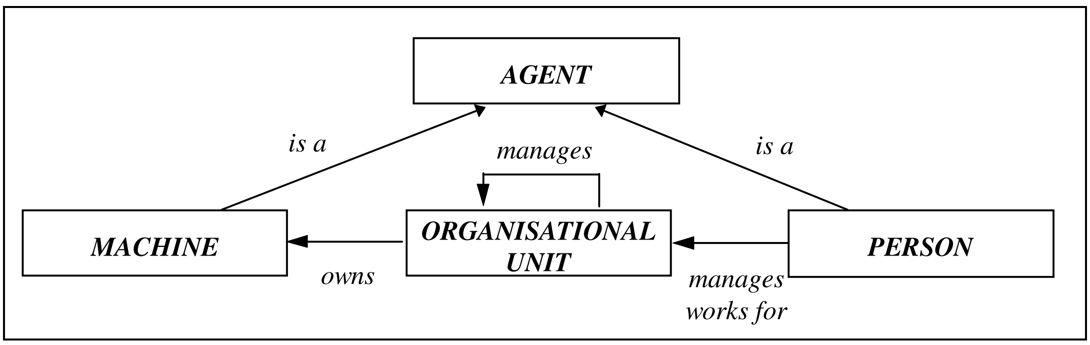

Figure 3: Graphical Representation of the Organisational Structure Modelling Language

5. Authority Modelling Language Proposal

Our authority modelling language aims to provide constructs that can be used in conjunction with an organisational model to define the authority relationships within an organisation. The model is based upon the following authority primitives:

- **Obliged:** an agent is obliged to perform an action.
- **Permitted:** the agent may decide itself whether or not to perform an action.
- **Forbidden:** an agent must not perform an action.

These primitives are taken from the Deontology[1] field; a field aiming to describe the duties and responsibilities of individuals that has been used in a number of areas within computer science [Wieringa 93]. The primitives have been successfully applied to authority modelling issues in other areas of Computer Supported Cooperative Work. Liu and Dix [1997], for example, use them to endow computer agents with knowledge about their responsibilities.

Consider a sales department that has a manager, a secretary, and a meeting room. This situation is shown using our organisational modelling language in Figure 4.

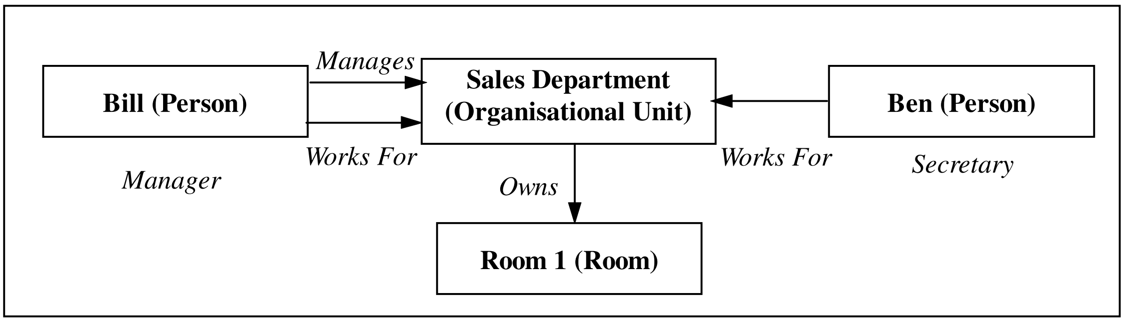

Figure 4: Example Organisational Model

Consider two capabilities that a secretary can provide: the retrieving of salary details, and the booking of rooms. These capabilities are specified below using our technical capability and knowledge space approach.

Retrieve(Person's Salary)
Book(Room)

Figure 5: Example Capability Specifications

The example shown in Figure 6 adds the authority context to the secretarial capabilities introduced in Figure 5. Considering first the *Retrieve* capability, the "obliged" slot states that the secretary must provide the capability if the person requesting it manages organisational unit for which he or she is a secretary. The "permitted" slot states that the secretary is free to choose between answering or not answering a request made by an individual who is asking for their own salary details. The "forbidden" attribute states that the secretary must not return salary details to anybody who does not meet the criteria stated in the "obliged" or "permitted" attributes. The authority statements added to the *Book* capability state that the secretary must book the room for people working in the organisational unit that owns the room. The secretary is permitted to accept bookings from individuals who work for the overall organisation that the secretary's organisational unit is a part but is forbidden to accept bookings from individuals who do not meet the "obliged" or "permitted" criteria.

[1] *Greek: deon "duty", and logos "science".*

Authority Specification		
For Capability:	*Retrieve(Person's Salary)*	
	Obliged:	If the person requesting the salary details is the manager of the organisational unit that the secretary works in.
	Permitted:	If a person is requesting his or her own salary details.
	Forbidden:	To the rest of the world.
For Capability:	*Book(Room)*	
	Obliged:	If the person requesting the room works in the organisational unit that owns the room.
	Permitted:	If the person requesting the room works in the company that owns the organisational unit that owns the room.
	Forbidden:	To the rest of the world.

Figure 6: Authority Context of the Secretary's Capabilities

5.1 The "Culture" Perspective

The authority specifications in Figure 6 are from what we term the "Capability" perceptive as they each apply to a single specified capability. Our "Culture" Perspective is designed to complement the "Capability" Perspective with the specification of general authority statements that reflect the authority culture within an organisation. For example, a manager may have authority over all the agents working in or owned by the department(s) he or she manages and thus the capabilities provided by these individuals. Figure 7 outlines such a specification. The essential difference between this culture related specification than those in Figure 6 is the inclusion of "for all" type statements that are applicable over a set of capabilities. In the detailed perspective above, authority statements are specified in terms of a single and explicitly stated capability.

For Capability:	*For all capabilities possessed by people who work in the organisational unit that is managed by a particular manager.*	
	Obliged:	to accept requests from that manager

Figure 7: Example Capability Specifications for a Manager Role

Whilst the culture perspective allows a single authority statement to be applied to many technical capabilities with the benefit of reducing modelling effort, the perspective is less precise than the detailed perspective. The two specifications given in Figure 6 and Figure 7 are currently consistent. Figure 6 states that the secretary is obliged to answer salary retrieval requests from his or her manager. Figure 7 implies that the secretary is obliged to answer all requests from his or her manager. If the statement in Figure 6 is modified to state that the secretary is forbidden to return salary details to his or her manager, then the two specifications contradict each other. The capability perspective is now stating that the request cannot be answered whilst the cultural perspective is stating that the request must be answered. To resolve such conflicts, a system working with authority specifications must employ some conflict resolution strategy. A general approach can be to always favour the capability oriented statement as this perspective is more detailed. However, the system may simply identify the conflict and allow the users to resolve it.

5.2 Using Organisational Structure and Authority in Workflow

The specifications of organisational structure and authority described above can be used to support the desired functionality outlined in section 3:

- **Task Initiation:** if task specifications include statements of authority, a workflow system can ensure that only agents who meet the requirements specified in those statements can initiate a task. If an agent does not have the necessary authority, the system can identify agents that do and suggest that the initiating agent discusses the task with one of these agents.
- **Task Planning:** considering the availability of agents for performing each method a workflow system can rank methods for achieving a task using an "authority metric". A method where

activities can be assigned to agents with a high level of obligation to perform their activities may be considered preferable to a method where agents are only "permitted" to perform the activities. Methods for which the only agents available are forbidden to perform their activities can be discounted.

- **Activity Scheduling:** the "authority metric" of task planning can be adapted to rank potential agents during scheduling. The set of available agents can be ranked according to the authority held by the agent who wants the activity to be scheduled. Agents "obliged" to perform the activity may be considered preferable to those only "permitted" to perform it.
- **Interaction Styles:** interaction styles could be modified by examining the relative authorities between two agents. If during an activity delegation, the receiving agent is "obliged" to perform the activity then a workflow system could offer a different communication style to the case where the agent is "permitted" to perform the activity.

5.3 Further Authority Modelling Issues

To enable a clear introduction, the description of our authority language so far has not made clear the relative "viewpoints" from which authority can be specified. Authority can be specified and viewed from either the viewpoint of the agent providing a capability or that of an agent requesting a capability. Considering the example in Figure 6, the secretary's authority is specified from the "provisions" viewpoint as the authority statements are attached to each of the capabilities that the secretary provides. Considering the example in Figure 7, the manager's authority is specified in terms of the capabilities that he or she can request and is therefore specified from the "requesting" viewpoint. Both the "provisions" and "requesting" cases are simply points from which a single authority model can be described and viewed. We anticipate that the perspectives will provide two useful viewpoints form which to present authority models within tools that support people in understanding, critiquing, and maintaining such models.

6. Conclusion

If workflow technology is to move forward into applications that are not based upon prescriptive process models, then techniques must be developed for supporting users in making decisions about the activities required to achieve a task and the agents that can perform those activities. In this paper, we have argued that knowledge about technical capabilities and the organisational structure and authority context within which they exist are important factors used by people when making such decisions. Our introduction to technical capability matching outlined how such a function can support important workflow decisions. Complementing this support with knowledge about an organisation's structure and authority context allows a workflow system to take organisational norms into account and to help users maintain these norms. The following points summarise how a number of workflow decisions can be supported by technical capability matching and how this support can be enhanced with knowledge about organisational structure and authority:

- **Task Initiation:**

  Technical Only: ensure that users can only initiate tasks that they have the technical skills to manage.

  With Authority: ensure that tasks are started only by users with the appropriate authority. If a user does not have the authority to initiate a task, the workflow system can suggest a user who does so that initiation can be discussed.
- **Task Planning:**

  Technical Only: rank the methods available for achieving a task using the availability of agents that are technically qualified to perform the constituent activities.

  With Authority: rank the methods available for achieving a task using the authority and level of obligation of available agents that perform the constituent activities.
- **Activity Scheduling:**

  Technical Only: identify the agents that have the technical capability to perform an activity.

  With Authority: identify the agents that have the authority to perform an activity and rank them according to their level of obligation to perform that activity.
- **Interaction Styles:**

Technical Only: no support

With Authority: modify the interaction support for two agents in accordance with the authority relationship between them.

We have noted how authority can be defined from both “capability” and “cultural” perspectives. The capability perspective lends itself to precise authority specifications that require significant modelling effort. The cultural perceptive lends itself to more general authority specifications that reduce the modelling effort with an (often acceptable) loss in precision.

We plan to implement the organisational and authority modelling languages outlined in this paper and experiment with them within an adaptive workflow system. This work should produce precise modelling languages and empirical evidence of their usefulness.

7. Acknowledgements

The work described in this paper can be accredited to two main projects. The Enterprise project[2] was funded under the UK government’s *Intelligent Systems Integration Programme*. It was carried out at AIAI, the University of Edinburgh, UK, with input from industrial partners IBM UK, Lloyd's Register, Logica, and Unilever. The Task-Based Process Management project[3] is funded under the UK Engineering and Physical Sciences Research Council’s *Systems Engineering for Business Process Change* programme. It is a collaboration between AIAI and the Chemical Engineering Department at Loughborough University, UK. Industrial partners on this project are BG Technology, ICI, and Unilever.

8. References

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

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

Fraser, J. and Tate, A., 1995, The Enterprise Tool Set -- An Open Enterprise Architecture. In: Proceedings of the Workshop on Intelligent Manufacturing Systems, International Joint Conference on Artificial Intelligence (IJCAI-95), Montreal, Canada,

Georgakopoulos, D., Hornick, M. and Sheth, A., 1995, An Overview of Workflow Management: From Process Modelling to Workflow Automation Infrastructure. *Distributed and Parallel Databases*, 3, 119-153, Kluwer.

Hoog, R., Benus, B., Metselaar, C., Vogler, M. and Menezes, W., 1994, Organisation Model: Model Definition Document. University of Amsterdam Technical Report, KADSII/M6/UvA/041/3.0, Netherlands.

Joosten, S., 1996, Workflow Management Research Area Overview. *Proceedings of 2nd Americas Conference on Information Systems,* Phoenix, Arizona.

Kappel, G., Lang, P., Rausch-Schott, S. and Retschitzegger, R., 1995, Workflow Management Based on Objects, Rules, and Roles. IEEE Bulletin of the Technical Committee on Data Engineering, 18(1), 11-17.

Klein, M., 1996, Challenges and Directions for Coordination Science. *Proceedings of the Second International Conference on the Design of Cooperative Systems (COOP’96),* June 1996, France.

Liu. K. and Dix, A., 1997, Norm Governed Agents in CSCW. *Proceedings of the First international Conference on Computational Semiotics*, Paris, France.

Rupietta, W., 1997, Organisation and Role Models for Workflow Processes, in Lawrence. P, (ed.) Workflow Handbook, Wiley.

[2] Details available at http://www.aiai.ed.ac.uk/~entprise/

[3] Details available at http://www.aiai.ed.ac.uk/project/tbpm/tbpm.html

Stader, J., 1996, Results of the Enterprise Project. *Proceedings of Expert Systems '96, the 16th Annual Conference of the British Computer Society Specialist Group on Expert Systems*, Cambridge, UK.

Stader, J., 1997, An Intelligent System For Bid Management, *The International Journal of Project and Business Risk Management*, Project Manager Today Publications, 1(3) 299 - 314.

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, Toulouse, France, September.

Sheth, A. and Kochut. K., 1997, Workflow Applications to Research Agenda: Scaleable and Dynamic Work Coordination and Collaboration Systems. *Proceedings of the NATO Workshop on Advances in Workflow Management Systems and Interoperability*, Turkey.

Tate, A., 1993, Authority Management - Coordination between Planning, Scheduling and Control. Workshop on Knowledge-based Production Planning, Scheduling and Control at the International Joint Conference on Artificial Intelligence (IJCAI-93), Chambery, France.

Uschold, M., King, M., Moralee, S. and Zorgios, Y., 1998, The Enterprise Ontology. *The Knowledge Engineering Review*, 13, Special Issue on Putting Ontologies to Use (eds. M. Uschold and A. Tate). [Available From] ftp://ftp.aiai.ed.ac.uk/pub/documents/1998/98-ker-ent-ontology.ps.

Uschold, M., 1998, The Enterprise Ontology Ontolingua Encoding. [Available From] http://www-ksl-svc.stanford.edu:5915/FRAME-EDITOR/&sid=ANONYMOUS&user-id=ALIEN

Wieringa, R., and Meyer, J., 1993, Applications of Denotic Logic in Computer Science: A Concise Overview. In R. Wieringa and J. Meyer (eds.), *Deontic Logic in Computer Science*. Wiley.