An Organisation Ontology for Enterprise Modelling Preliminary Concepts for Linking Structure and Behaviour

Mark S. Fox, Mihai Barbuceanu, Michael Gruninger

Department of Industrial Engineering, University of Toronto, 4 Taddle Creek Road, Toronto, Ontario M5S 1A4 CANADA *tel:*+1-416-978-6823 *fax:*+1-416-971-2479

email: {msf, mihai, gruninger}@ie.utoronto.ca http://www.ie.utoronto.ca/EIL/eil.html

Abstract

The paper presents our preliminary exploration into an organisation ontology for the TOVE enterprise model. Its primary focus has been in linking structure and behaviour through the concept of empowerment. Empowerment is the right of an organisation agent to perform status changing actions. This linkage is critical to the unification of enterprise models and their executability.

1.0 Introduction

What is an organisation and how do we model it in an information system? Many disciplines have explored the former and every information system built has created a version of the latter. The purpose of this paper is to explore the latter from the perspective of Artificial Intelligence.

As information systems play a more active role in the management and operations of an enterprise, the demands on these systems have also increased. Departing from their traditional role as simple repositories of data, information systems must now provide more sophisticated support to manual and automated decision making; they must not only answer queries with what is explicitly represented in their Enterprise Model, but must be able to answer queries with what is *implied* by the model. The goal of the TOVE Enterprise Modelling project is to create the next generation Enterprise Model, a Common Sense Enterprise Model. By common sense we mean that an Enterprise Model has the ability to deduce answers to queries that require relatively shallow knowledge of the domain.

We are taking what can be viewed as a "second generation knowledge engineering" approach to constructing our Common Sense Enterprise Model. Rather than extracting rules from experts, we are "engineering ontologies." An ontology is a formal description of entities and their properties, relationships, constraints, behaviours. Our approach to engineering ontologies begins with defining an ontology's requirements; this is in the form of questions that an ontology must be able to answer. We call this the *competency* of the ontology. The second step is to define the terminology of the ontology - its objects, attributes, and relations. The third step is to specify the definitions and constraints on the terminology, where possible. The specifications are represented in First Order Logic and implemented in Prolog. Lastly, we test the competency of the ontology by "proving" the competency questions with the Prolog axioms.

Our initial efforts have focused on ontologies to support reasoning in industrial environments. The tasks that we have targeted to support are in "supply chain management" which extends MRP (Manufacturing Requirements Planning) to include logistics/distribution and "Concurrent Engineering" which looks at issues of coordination of engineering design. Much of our effort has been in creating representations of organisation behaviour: activity, state, causality and time, and the objects they manipulate: resources [Fadel 94, Fadel et al. 94], inventory, orders and products. We also have efforts underway in formalising knowledge of ISO 9000 quality [Kim & Fox 93], activity-based costing [Tham et al. 94] and organisation agility.

This paper describes the *organisation* ontology being developed as part of the TOVE Project. In particular it focuses on organisation structure, roles, authority and empowerment.

2.0 What is an Organisation?

We consider an organisation to be a set of constraints on the activities performed by agents. This view follows that of Weber [7?] who views the process of bureaucratization as a shift from management based on self-interest and personalities to one based on rules and procedures.

Mintzberg [1983] provides an early (and informal) analysis of organization structure distinguishing among five basic parts of an organization and five distinct organization configurations that are encountered in practice. This "ontology" includes several mechanisms that together achieve coordination, like goals, work processes, authority, positions and communication. The various parts of an organization are distinguished by the specific roles they play in achieving coordination with the above means.

The "language/action perspective" [Winograd 1987] on cooperative work in organizations provides an ontology that emphasizes the social activity by which "agents" generate the space of cooperative actions in which they work, rather than the mental state of individuals. The basic idea is that social activity is carried out by language and communication. The pragmatic nature of communication as the way of creating commitments among participants is exploited in the Coordinator system [Flores et al. 1988].

In the same vein, [Auramaki et al. 1988] present a method for modeling offices as systems of communicative action through which people engage in actions by creating, modifying and deleting commitments that bind their current and future behaviors.

The work of Lee [1988] looks at language acts in the bureaucratic office, viewing language not as a mechanism for information transfer but as a mechanism for social interaction and control. He presents a logic-based representation of deontic notions - authorization, permission, prohibition and the like - and shows how this can be used to model cooperative work in the office.

More recently, Yu and Mylopoulos [1994] have proposed a framework for modeling organizations as being made of social actors that are intentional, having motivations, wants and beliefs and strategic, evaluating their opportunities and vulnerabilities with respect to each other. This formal model is used to explore alternative process designs in business reengineering.

3.0 Ontology Competence

A problem in the engineering of ontologies is their evaluation. A number of criteria have been proposed including [Fox et al. 93] [Gruber 93]:

- **Generality:** To what degree is the representation shared between diverse activities such as design and troubleshooting, or even design and marketing?
- **Competence:** How well does it support problem solving? That is, what questions can the representation answer or what tasks can it support?
- Efficiency: Space and inference. Does the representation support efficient reasoning, or does it require some type of transformation?
- **Perspicuity:** Is the representation easily understood by the users? Does the representation "document itself?"
- **Transformability:** Can the representation be easily transformed into another more appropriate for a particular decision problem?
- Extensibility: Is there a core set of ontological primitives that are partitionable or do they overlap in denotation? Can the representation be extended to encompass new concepts?

- **Granularity:** Does the representation support reasoning at various levels of abstraction and detail?
- **Scalability:** Does the representation scale to support large applications?
- Minimality: A minimal set of terms should be in the ontology.

But the criterion we have found most useful is *competence*. For any task in which the ontology is to be employed, the task imposes a set of requirements on the ontology. These requirements can best be specified as a set of queries that the ontology should be able to answer, if it contains the relevant information. These requirements, which we call competency questions, are the basis for a rigorous characterization of the information that the ontology is able to provide to the task. Competency questions are benchmarks in the sense that the ontology is necessary and sufficient to represent the tasks specified by the competency questions and their solution. They are also those tasks for which the ontology finds all and only the correct solutions. Tasks such as these can serve to drive the development of new ontologies and also to justify and characterize the capabilities of existing ontologies.

This characterization of competency raises an important issue: where does the representation end and inference begin? If no inference capability is to be assumed, then query processing is reducible to "looking up" an answer that is represented explicitly. In contrast, object/semantic network representations assume at least inheritance as a deduction mechanism. In defining an ontology a key question then becomes: should we be restricted to just a terminology? Should the terminology assume an inheritance mechanism, or some type of theorem proving capability as provided, say, in a logic programming language with axioms restricted to Horne clauses (i.e., Prolog)? What is the *deductive capability* that is to be assumed by an ontology? In the TOVE project we assume a theorem prover of the power of Prolog.

The basic entities in the TOVE ontology are represented as objects with specific properties and relations. Objects are structured into taxonomies and the definitions of objects, attributes and relations are specified in first-order logic. An ontology is defined in the following way. We first identify the objects in our domain of discourse; these will be represented by constants and variables in our language. We then identify the properties of these objects and the relations that exist over these objects; these will be represented by predicates in our language.

We next define a set of axioms in first-order logic to represent the constraints over the objects and predicates in the ontology. This set of axioms provides a declarative specification for the various definitions and constraints on the terminology. Further, we need to prove the competency of the ontology. The ontology must contain a necessary and sufficient set of axioms to represent and solve these questions, thus providing a declarative semantics for the system. It is in this sense that we can claim to have a competent ontology, and it is this rigor that is lacking in previous approaches to ontology engineering.

The competency questions are generated by requiring that the ontologies be necessary and sufficient to support the various tasks in which it is employed. Within our applications, these include:

• Planning and scheduling -- what sequence of activities must be completed to achieve some goal? At what times must these activities be initiated and terminated?

- Temporal projection -- Given a set of actions that occur at different points in the future, what are the properties of resources and activities at arbitrary points in time? This includes the management of resources and activitybased costing (where we are assigning costs to resources and activities).
- Execution monitoring and external events --What are the effects on the enterprise model of the occurrence of external and unexpected events (such as machine breakdown or the unavailability of resources)?
- Hypothetical reasoning -- what will happen if we move one task ahead of schedule and another task behind schedule? What are the effects on orders if we buy another machine?
- Time-based competition -- we want to design an enterprise that minimizes the cycle time for a product [Blackburn 91]. This is essentially the task of finding a minimum duration plan that minimizes action occurrence and maximizes concurrency of activities.

4.0 Activity/Time Ontology

In this section we define the ontology of time and action that is used to represent the behaviour of the organisation. An important component of representing behaviour is the ability to temporally project, that is, to determine the possible set of future states given a current state. Temporal projection induces the following set of requirements on the ontologies:

• Temporal projection requires the evaluation of the truth value of a proposition at some point in time in the future. We therefore need to define axioms that express how the truth of a proposition changes over time. In particular, we need to address the frame problem and express the properties and relations that change or do not change as the result of an activity.

- We must define the notion of a state of the world, that is, define what is true of the world before and after performing different activities. This is necessary to express the causal relationship between the preconditions and effects of an activity.
- The time interval over which the state has a certain status is bounded by the times at which the appropriate actions that change status occur. This interval defines the duration of a state if the status is enabled. This is essential for the construction of schedules.
- We want a uniform hierarchical representation for activities (aggregation). Plans and processes are constructed by combining activities. We must precisely define how activities are combined to form new ones. The representation of these combined activities should be the same as the representation of the subactivities. Thus aggregate activities (sets of activities or processes) should themselves be represented as activities.
- The causal and temporal structure of states and subactivities of an activity should be explicit in the representation of the activity.

4.1 Situation Calculus Specification

We represent time as a continuous line; on this line we define time points and time periods (intervals) as the domain of discourse. We define a relation < over time points with the intended interpretation that t < t' iff t is earlier than t'.

One important property that must be represented to define what holds in the world after performing some action in order to capture the notion of causality. How do we express these notions if we have a continuous time line? The extended situation calculus of [Pinto & Reiter 93] allows us to incorporate the notions of situations and a time line by assigning durations to situations.

The intuition behind the situation calculus is that there is an initial situation, and that the world changes from one situation to another when actions are performed. There is a predicate Pos $s(a,\sigma)$ that is true whenever an action a can be performed in a situation σ .

The structure of situations is that of a tree; two different sequences of actions lead to different situations. Thus, each branch that starts in the initial situation can be understood as a hypothetical future. The tree structure of the situation calculus shows all possible ways in which the events in the world can unfold. Therefore, any arbitrary sequence of actions identifies a branch in the tree of situations.

Further, we impose a structure over situations that is isomorphic to the natural numbers by introducing the notion of successor situation [Reiter 91]. The function $do(a,\sigma)$ is the name of situation that results from performing action *a* in situation σ . We also define an initial situation denoted by the constant σ_0 .

Situations are assigned different durations by defining the predicate start(s,t) [Pinto & Reiter 93]. Each situation has a unique start time; these times begin at 0 in σ_0 and increase monotonically away from the initial situation.

To define the evaluation of the truth value of a sentence at some point in time, we will use the predicate $holds(f,\sigma)$ to represent the fact that some ground literal *f* is true in situation σ . Using the assignment of time to situations, we define the predicate $holds_T(f, t)$ to represent the fact that



some ground literal f is true at time t. A fluent is a predicate or function whose value may change with time.

Another important notion is that actions occur at points in time. The work of [Pinto & Reiter 93] extends the situation calculus by selecting one branch of the situation tree to describe the evolution of the world as it actually unfolds. This is done using the predicate actual. To represent occurrences, we then introduce two predicates, $occurs(a,\sigma)$ and $occurs_T(a,t)$, defined as follows:

 $occurs(a,\sigma) \equiv actual(do(a,\sigma))$ (EQ 1)

$$occurs_T(a,t) \equiv occurs(a,\sigma) \land start(do(a, \sigma), t)$$
 (EQ 2)

We will now apply this formalism to the representation of activities in an enterprise.

4.2 Terminology

At the heart of the TOVE Enterprise Model lies the representation of an *activity* and its corresponding enabling and caused *states* ([Sathi et al. 85], [Fox et al 93]). In this section we examine the notion of states and define how properties of activities are defined in terms of these states. An activity is the basic transformational action primitive with which processes and operations can be represented; it specifies how the world is changed. An enabling state defines what has to be true of the world in order for the activity to be performed. A caused state defines what is true of the world once the activity has been completed.

An activity, along with its enabling and caused states, is called an *activity cluster*. The state tree linked by an *enables* relation to an activity specifies what has to be true in order for the activity to be performed. The state tree linked to an activity by a *causes* relation defines what is true of the world once the activity has been completed. Intermediate states of an activity can be defined by elaborating the aggregate activity into an activity network (see Figure 0).

There are two types of states: *terminal* and *non-terminal*. In Figure 0, *es_fabricate_plug_on_-wire* is the nonterminal enabling state for the activity *fabricate_plug_on_wire* and *pro_fabricate_plug_on_wire* is the caused state for the activity. The terminal conjunct substates of *es_fabricate_plug_on_wire* are *consume_wire*, *consume_plug*, and *use_inject_mold* since all three resources must be present for the activity

to occur; the terminal states of *pro_fabricate_plug_on_wire* are *produce_plug_on_wire* and *release_inject_mold*.

In TOVE there are four terminal states represented by the following predicates: use(s,a), consume(s,a), release(s,a), produce(s,a). These predicates relate the state with the resource required by the activity. Intuitively, a resource is used and released by an activity if none of the properties of a resource are changed when the activity is successfully terminated and the resource is released. A resource is consumed or produced if some property of the resource is changed after termination of the activity; this includes the existence and quantity of the resource, or some arbitrary property such as color. Thus consume(s,a) signifies that a resource is to be used up by the activity and will not exist once the activity is completed, and produce(s,a) signifies that a resource, that did not exist prior to the performance of the activity, has been created by the activity. We define use and consume states to be enabling states since the preconditions for activities refer to the properties of these states, while we define release and produce states to be caused states, since their properties are the result of the activity.

Terminal states are also used to represent the amount of a resource that is required for a state to be enabled. For this purpose, the predicate *quantity*(*s*,*r*,*q*) is introduced, where *s* is a state, *r* is the associated resource, and *q* is the amount of resource r that is required. Thus if *s* is a consume state, then *q* is the amount of resource consumed by the activity, if *s* is a use state, then *q* is the amount of resource used by the activity, and if *s* is a produce state, then *q* is the amount of resource produced.

A state may have a status whose value is one of the following constants: {*possible, committed,*

enabled, *completed*, *disenabled*, *reenabled*}. The status of a state is changed by one of the following actions:*commit(s,a)*, *enable(s,a)*, *complete(s,a)*, *disenable(s,a)*, *reenable(s,a)*. Note that these actions are parametrized by the state and the associated activity.

Similarly, activities have a status whose value is one of the following constants: {dormant, executing, suspended, completed}. The status of an activity is changed by one of the following actions: execute(a), suspend(a), complete(a).

As part of our logical specification of the activity ontology, we define the successor axioms that specify how the above actions change the status of a state. These axioms provide a complete characterization of the value of a fluent after performing any action, so that we can use the solution to the frame problem in [Reiter 91]. Thus if we are given a set of action occurrences, we can solve the temporal projection problem (determining the value of a fluent at any point in time) by first finding the situation containing that time point, and then using the successor axioms to evaluate the status of the state in that situation. We present one of the successor axioms in the ontology:

The status of a state is committed in a situation iff either a commit action occurred in the preceding situation, or the state was already committed and an enable action did not occur.

A more complete specification can be found in [Gruninger & Fox 94].

5.0 Competency

In linking the structure of an organisation with the behaviour of agents within the organisation, we must define how the organisation ontology is integrated with the activity ontology.

If we consider organisation to be a set of constraints on the activities performed by agents, then the competency questions for the organisation ontology are extensions of the temporal projection and plan existence problems to incorporate the abilities and obligations of agents. The temporal projection problem is used to characterize the constraints that agents must satisfy to be able to perform activities, and plan existence characterizes the set of achievable goals. We can then propose the following set of competency questions for the organisation ontology.

5.1 Structure

- What is the structure of the organisation? How is the organisation decomposed into units?
- What are the members of a particular unit of the organisation?
- What positions exist in the unit?
- What position does person X occupy?
- Who must person X communicate with?
- What kinds of information does person X communicate?
- Who does X report to?

5.2 Behaviour

- What are the goals of the unit?
- What are the goals of the position?

- What are the goals of person X?
- What activities must a particular position perform?
- What activities must person X perform?
- Is it possible for an agent to perform an activity in some situation? That is, does the agent have the ability to perform the activity?

5.3 Authority, Empowerment and Commitment

- What resources does the person have authority to assign?
- What activities may a person execute without explicit permission?
- In order to perform a particular activity, whose permission is needed?
- Is an agent allowed to perform an activity in some situation?
- What goals is person X committed to achieving?
- Is a goal achievable by an agent given its current commitments and the commitments of other agents?
- If a goal is unachievable for a given set of agents, how can they be empowered to be capable of performing the activities to achieve the goal? That is, how can the constraints defining empowerment for the agents be modified so as to be able to achieve the goal?
- What authority constraints are necessary among a set of agents in order to achieve a goal?

5.4 Goal Achievement

- What goals are solitarily unachievable for a given agent? That is, what goals are unachievable using a plan that contains only activities that the agent is capable of performing? Such goals require the assistance of other agents to achieve them.
- What goals are achievable by an agent given the effects of activities that other agents are capable of performing?
- If a goal is solitarily unachievable for a given agent, what agents are required to assist the agent in achieving the goal?

6.0 Organisation Terminology

In this section we introduce the basic terminology, i.e., ground terms, of our organisation ontology.

Figure 1 shows the basic elements of our organisation ontology.





6.1 Organisation Agent

A concept found in almost all of the literature is that of an agent. An agent performs activities in order to achieve one or more goals. An agent can be a human being, a computer program, or a group of people and/or programs.

Organisation-Agents (OAs) are the "active" entities in an organisation.

Organisation-Agent(OA)	(EQ 4)
Individual-Agent(OA)	(EQ 5)
Group-Agent(OA)	(EQ 6)

Individual-Agent and **Group-Agent** are subclasses of **Organisation-Agent**. They represent either individuals, like employees and contractors, or groups like departments, divisions, boards of directors, etc. (figure 2).

The basic structural relationship of the organisation is the **member-of** relation.

member-of(
$$OA_1, OA_2$$
) (EQ 7)

In this example, organisation agent OA_1 is a member of group agent OA_2 .

OAs play various **Organisation-Role**s, they have **Goals** to achieve, fill **Oganisation-Positions**, and communicate with other OAs using **Information-links**.





6.2 Organisation-Role

An **Organisation-Role** defines a prototypical function of an agent in an organisation.

Organisation-Role(OR) (EQ 8)

A particular agent can assume several roles at the same time. For an individual agent, examples of organisational roles include "project manager", "reviewer", "troubleshooter", etc. Once an agent is assigned to a role, that creates a commitment (more on commitments later) on the agent's part to act to achieve the goal(s) of the role.

Each organisational-Role has:

• *Goals*: one or several **goals** which the agent playing the role is responsible for.

has-goal(OR, G) (EQ 9)

• *Skills*: one or more skills required to achieve the goals.

has-skill(OR, S) (EQ 10)

• *Processes*: activity networks that have been defined to achieve the goals.

• *Policies*: constraints on the performance of the role's processes. These constraints are unique to the organisation role.

has-policy(OR, Po) (EQ 12)

• *Information-Link*: these are communication links to other agents in specified roles. Communication consists of exchanging speech acts according to specific conversation structures that are also formally represented [Barbuceanu & Fox 95].

```
has-information-link(OR, IL) (EQ 13)
```

6.3 Organisation Position

An **Organisation-Position** defines a formal position that can be filled by an OA in the organisation.

(OP is the unique identifier for a specific organisation position.) Examples of positions include "president", "laboratory director", "senior researcher", "sales-representative", etc. Any position essentially consists of a set of roles the OA filling it will have to carry out. For each positions we specify:

• *Roles*: the roles to be assumed in the position.

• *Agent*: the organisation agent filling the position. In general we assume that positions are filled by individual agents. Note that a group agent may fill a position.

• *Policies*: constraints on the performance of position's processes (inherited from the required roles). These constraints are unique to the organisation position.

6.4 Information-Link

The Information-Link relation

is a unidirectional link used to communicate information from one agent to another. It describes, for an agent in a given organisational role, the information it is interested in receiving from another agent. Information links are formal in the sense that they are specified for every organisational role, and they are informal in that they can be associated with an pair of agents.

The Information-Link specifies:

• *Sending-Agent*: the agent sending information along the link.

has-sending-agent(IL, OA) (EQ 19)

• *Receiving-Agent*: the agent receiving information from the link.

has-receiving-agent(IL, OA) (EQ 20)

• *Sending-Role*: the organisation role played by the sending agent.

has-sending-role(IL, R) (EQ 21)

• *Receiving-Role*: the organisation role played by the receiving agent.

has-receiving-role(IL, R) (EQ 22)

• *Interests*: the information interests of the receiving agent.

• *Volunteers*: the information the sending agent can supply to other agent.

It is understood that information distribution in the above case is non-committal, in the sense that it does not create obligations for either the sender or the receiver.

6.5 Authority and Commitment

We introduce the concept of an OA's commitment to achieving a goal [Jensen 93]. The predicate

signifies that **Organisation-Agent** OA is committed to the achievement of **Goal** G. Conse-

quently, the totality of activities performed by OA must include the achievement of G. Prioritisation of goals, etc. are not considered here.

We use *authority* to refer to the control relationship that exists between two organisational agents. For OA_1 to have *authority* over OA_2 implies that OA_1 is able to extract a commitment from OA_2 to achieve a **goal** that is defined as part of OA_2 's **organisation-roles**. In order to extract that commitment, OA_1 has to be related directly or indirectly by a **authority-link** relation.

authority-link(AL)	(EQ 26)
--------------------	---------

has-supervisor(AL, OA_1) (EQ 27)

has-supervisee(AL, OA₂) (EQ 28)

can-assign-goal(AL, G) (EQ 29)

The **authority-link** relation entails a number of constraints such as the ability of the supervisor to extract commitments from the supervisee. The **can-assign-goal** relation explicitly states what goals the supervisor may assign to the supervisee.

7.0 Empowerment: Linking Structure and Behaviour

With the introduction of organisational knowledge, we now have to address the problem of how to specify "who can do what". That is, what is the set of activities that an OA is allowed to perform as a member of the Organisation. It would appear that by associating processes with OAs via the **has-process** property, we have solved the problem. That is, an OA can perform any activities specified by its processes. But consider the following situation: "Jill, in her role as a CNC machine operator, has a process she must perform in order to achieve the goal of producing an order. The process is composed of three activities: 1) machine-setup, 2) machine-run and 3) machine-teardown. But before the machine-run activity can commence, she must receive permission from her supervisor."

The problem is that Jill has a process that specifies a sequence of activities that she must perform. But she cannot perform the second activity, machine-run, without permission. The implication is that within our Activity ontology, she is not allowed to change the state of the machine-run activity to "execute".

An obvious way to solve the problem is to insert a fourth activity between machine-setup and machine-run where she seeks approval from her supervisor. If approval is obtained, then she can commence the subsequent machine-run activity. Again we have a problem. Who is allowed to change the status of this new approval activity to completed? If Jill is allowed to make any status changes she wants to the activities in her process, she can change the status of the approval activity regardless of whether she obtained approval or not.

The problem lies with who is allowed to make status changes to states and activities. When Jill goes to her supervisor for permission, is it Jill who changes the status of the approval activity to completed or her supervisor? It is not clear. Therefore the only solution to the problem of permission to perform an action lies in precisely stating who is allowed to change the status of the activity, e.g., from dormant to executing. We introduce the concept of Empowerment as a means of specifying the status changing rights of an OA. *Empowerment is the right of an OA to perform status changing actions*, such as commit, enable, suspend, etc. Empowerment naturally falls into two classes: state and activity empowerment.

State empowerment specifies the range of stati through which an OA may take a state by performing the appropriate actions, such as commit. State empowerment not only specifies allowable status changes but may be used to restrict the set of resources an OA is empowered to commit to a use/consume state. An OA may be empowered for any type of resource, including other OAs. The implication being the first OA may commit the second to a state.

Activity empowerment specifies the range of stati through which an OA may take an activity by performing the appropriate actions, such as execute and suspend. Even though an activity may be enabled, the OA whose role contains the plan which contains the activity may not be empowered to start its execution.

With the addition of empowerment, a second type of authority arises. That is, the supervising agent may alter what a supervise is empowered to do.

8.0 Axioms

The competency questions defined earlier drive the creation of the terminology and its axiomatisation. The structural and behavioural questions can be answered directly with the predicates introduced in the previous two sections with some simple axioms. The first axioms we introduce are the obvious taxonomic constraints. If an agent is an individual or group then it is constrained to being an Organisation-Agent.

 $Individual - Agent(OA) \supset Organisation - Agent(OA)$

 $Group - Agent(OA) \supset Organisation - Agent(OA)$

The **member-o**f relation imposes a constraint on the second parameter, that is, OA2 has to be a **Group-Agent**.

 $member - of(OA1, OA2) \supset Group - Agent(OA2)$

The next axiom states that any agent that fills an organisation role is committed to the goals associated with the role.

The remainder of the axioms presented here focus on empowerment. For any activity *a* that requires that the agent be empowered, the status changing actions for the activity require *holds(-activity_empowered(agent,a),* σ) as a precondition.

For any state *s* that requires that the agent be empowered, the status changing actions for the activity require *holds*(*state_empowered* (*agent*,*s*), σ) as a precondition.

It is possible to for one agent to empower another agent for an activity if the first agent supervises the second, and the supervisor is empowered for that activity.

 $\begin{array}{l} Poss(activity_empowers(agent, agent', a), \sigma) \equiv \\ holds(supervises(agent, agent'), \sigma) \land \\ holds(activity_empowered(agent, a), \sigma) \qquad (EQ 30) \end{array}$

It is possible to for one agent to disempower another agent for an activity if the first agent supervises the second, and the supervisor is empowered for that activity.

 $\begin{array}{l} Poss(activity_disempowers(agent, agent', a), \sigma) \equiv \\ holds(supervises(agent, agent'), \sigma) \land \\ holds(activity_empowered(agent, a), \sigma) \qquad (EQ 31) \end{array}$

An agent is empowered for an activity only as a result of the action *activity_empowers*, and is no longer empowered only as a result of the action *activity_disempowers*.

 $\begin{array}{l} holds(activity_empowered(agent, a), do(a', \sigma)) \equiv (\exists agent') \\ a' = activity_empowers(agent', agent, a) \lor \\ holds(activity_empowered(agent, a), \sigma) \land \neg(\exists agent') a' = \\ activity_disempowers(agent', agent', a) \\ \end{array}$

It is possible to for one agent to empower another agent for changing the status of states if the first agent supervises the second, and the supervisor is empowered for changing the status of that state.

 $\begin{array}{l} Poss(state_empowers(agent, agent', s), \sigma) \equiv \\ holds(supervises(agent, agent'), \sigma) \land \\ holds(state_empowered(agent, s), \sigma) \qquad (EQ 33) \end{array}$

It is possible to for one agent to disempower another agent for changing the status of that state if the first agent supervises the second, and the supervisor is empowered to change the status of that state.

 $\begin{array}{l} Poss(state_disempowers(agent, agent', s), \sigma) \equiv \\ holds(supervises(agent, agent'), \sigma) \land \\ holds(state_empowered(agent, s), \sigma) \end{array} (EQ 34) \end{array}$

An agent is empowered for changing the status of a state only as a result of the action *state_empowers*, and is no longer empowered only as a result of the action *state_disempowers*.

 $\begin{aligned} holds(state_empowered(agent, a), do(a', \sigma)) &\equiv (\exists agent') a' \\ &= activity_empowers(agent', agent, a) \lor \\ holds(state_empowered(agent, a), \sigma) \land \neg(\exists agent') a' = \\ state_disempowers(agent', agent', a) \end{aligned}$

9.0 Conclusions

The paper presents our preliminary exploration into an organisation ontology for the TOVE enterprise model. Its primary focus has been in linking structure and behaviour through the concept of empowerment. Empowerment is the right of an organisation agent to perform status changing actions. This linkage is critical to the unification of enterprise models and their executability. Much work remains in the development of our ontology and especially its axiomatisation.

10.0 Acknowledgments

This research is supported, in part, by the Natural Science and Engineering Research Council, Manufacturing Research Corporation of Ontario, Digital Equipment Corp., Micro Electronics and Computer Research Corp., Spar Aerospace, Carnegie Group and Quintus Corp.

11.0 References

[Auramaki et al. 88] Auramaki, E., Lehtinen, E. and Lyytinen, K. A speech-act-based office modelling approach, *ACM Transactions on Office Information Systems*, Vol. 6, No. 2, april 1988, 126-152.

[Blackburn 91] Blackburn J. *Time-based Competition*. Business One Irwin, 1991.

[Fadel et al 94] Fadel, F., Fox, M.S., and Gruninger, M. A resource ontology for enterprise modelling. *Third Workshop on Enabling Technologies-Infrastructures for Collaborative Enterprises*,(West Virginia University 1994). [Flores et al. 88] Flores, F., Graves, M., Hartfield, B. and Wionograd, T. Computer systems and the design of organizational interaction, *ACM Transactions on Office Information Systems*, Vol. 6, No. 2, april 1988, 153-172.

[Fox et al. 93] Fox, M.S., Chionglo, J., Fadel, F. A Common-Sense Model of the Enterprise, *Proceedings of the Industrial Engineering Research Conference 1993.*

[Gruber 93] Gruber, T. R.,(1993): Toward principles for the design of ontologies used for knowledge sharing, Report KSL 93-04, Stanford University, august 1993.

[Gruninger & Fox 94] Gruninger, M., and Fox, M.S., (1994), "The Role of Competency Questions in Enterprise Engineering", *Proceedings of the IFIP WG5.7 Workshop on Benchmarking -Theory and Practice*, Trondheim, Norway. June 94.

[Jennings 93] Jennings, N., R., (1993): Commitments and conventions: The foundation of coordination in multi-agent systems, The Knowledge Engineering Review, vol. 8:3, pp 223-250, 1993.

[Kim & Fox 93] Kim, H. and Fox, M.S. Quality Systems Modelling: A Prospective for Enterprise Integration, *Fourth Annual Meeting of the Production and Operations Management Society.* 1993.

[Lee 88] Lee, R. M. Bureaucracies as deontic systems, *ACM Transactions on Office Information Systems*, Vol. 6, No. 2, april 1988, 87-108. [Mintzberg 83] Mintzberg, H. *Structure in Fives* - *Designing Effective Organizations*, Prentice Hall Inc., 1983

[Pinto & Reiter 93] Pinto, J. and Reiter, R. Temporal reasoning in logic programming: A case for the situation calculus. In *Proceedings of the Tenth International Conference on Logic Programming* (Budapest, June 1993).

[Reiter 91] Reiter, R. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. *Ar*-*tificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCar*-*thy*. Academic Press, San Diego, 1991.

[Sathi et al 85] Sathi, A., Fox, M.S., and Greenberg, M. Representation of activity knowledge for project management. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. PAMI-7:531-552, September, 1985.

[Tham et al. 94] Tham, D., Fox, M.S., and Gruninger, M., A cost ontology for enterprise modelling *Third Workshop on Enabling Technol*ogies-Infrastructures for Collaborative Enterprises, (West Virginia University 1994).

[Winograd 87] Winograd, T. A language/action perspective on the design of cooperative work, *Human Computer Interaction* 3, 1 (1987-1988), 3-30

[Yu & Mylopoulos 94] Yu, E. S. K. and Mylopoulos, J. From E-R to "A-R" - Modelling strategic actor relationships for business process reengineering, *13-th Int. Conf. on the Entity-relationship Approach*, Dec. 13-16 1994, Manchester, UK.