Towards a Distributed Intelligent Agent Architecture for Human-Machine Systems in Hortatory Operations

John Murray & Yili Liu, [jxm,yililiu]@umich.edu
Department of Industrial & Operations Engineering
Univ. of Michigan, Ann Arbor, MI 48109-2117, USA

Technical Report IOE-95/17

[Submitted to the IEEE Transactions on Systems,

Man, and Cybernetics on January 14 1995.]

Towards a Distributed Intelligent Agent Architecture for Human-Machine Systems in Hortatory Operations

John Murray and Yili Liu (jxm, yililiu@engin.umich.edu)
Department of Industrial & Operations Engineering
University of Michigan, Ann Arbor, MI 48109-2117

ABSTRACT

Traditional models of supervisory monitoring and control operations generally rely upon two underlying assumptions: (i) the operations center is a centralized entity which exerts a large amount of direct influence and physical control over the system under supervision (the target system or 'domain of discourse'), and (ii) the vast majority of state information about the target system comes to the operations center via systems which are under its direct control. However, these assumptions are not always valid. For example, an operations center concerned with road traffic has limited ability to monitor and control all traffic flow directly; we introduce the term *hortatory operations* to describe this situation. The traffic operations center is typically one of several autonomous entities who must share their information and coordinate their activities with each other. This paper discusses our research into developing a distributed intelligent agent architecture for such an environment. We cover the main components of the architecture, and then examine in detail the ontologies needed in modeling the environment. We also demonstrate how the ontological approach led us to identify an innovative construct which has the potential to improve existing algorithms for detecting traffic congestion. Finally, we show how the intelligent reasoning features of our architecture can be applied in analyzing real-world practice.

Keywords: Human-machine systems, intelligent agents, hortatory operations, knowledge sharing, ontology design, road traffic management.

Towards a Distributed Intelligent Agent Architecture for Human-Machine Systems in Hortatory Operations

I. INTRODUCTION

Researchers in various fields have discussed the changing role of many humans who operate complex machines, from that of direct manipulation and control over the system to a more hands-off, supervisory role. The typical role of these people involves on-going vigilance over the system's activity and status, punctuated by periods of problem diagnosis and resolution; Sheridan [1] has characterized such supervisory control as "non-anthropometric tele-robotics". Studies of control center activity have typically focused upon characterizing the work as being done in a single, relatively autonomous environment. Such centers are typically staffed by one or more individuals who have significant levels of influence over the domain, which may be a processing plant, communications network, or other dynamic environment [2,3].

Human-machine interaction researchers have been successful in promoting improvements in these types of supervisory control environments. However, such approaches often cannot not accurately encapsulate an operations environment which has quite limited direct control over the domain, such as those associated with routine freeway traffic activity. Organizations handling the management of a freeway network have responsibilities that typically involve monitoring traffic levels, identifying and resolving problems which cause congestion or danger, coordinating construction activity, and so on. They can help 'inform' network users of current or expected conditions and can urge them to take certain actions. However, unlike an air traffic control center, they have little ability to direct the actions of the network users. And, unlike a public utility network controller say, systems to physically alter flow directions and rates are almost non-existent. We therefore propose to characterize this form of activity as *hortatory operations*, a term which reflects the use of information distribution and encouragement as a means of influencing the domain, rather than physical manipulation or administrative authority.

A second characteristic also differentiates the freeway operations domain from many other fields. There are typically several distinct organizations involved, each having only partial data directly available about the domain. Thus, the overall real-time knowledge about the domain is essentially distributed among several different entities such as freeway traffic management, police, transit operators, and commercial traffic information services. The divided access to this real-time knowledge is primarily governed by the nature of each organization's information-gathering methods. For example, police control operators tend to rely on emergency calls and patrol car communications, whereas freeway traffic managers often use inductive loop sensors and video surveillance as their primary data sources. The various organizations may have different goals and priorities, but they recognize the need to pool their knowledge in order to be effective. From a system modeling point of view, it is desirable therefore to find an architecture which permits the integration of all information into an overall and seamless picture of the domain.

The combination of distributed information availability and the hortatory nature of most traffic operations suggests that the more popular research models of control room activity are not well-suited to a freeway operations environment. These usually focus upon centralized access to information and decision-support systems - for example, the discussion of NASA's Mission Control Center in Kearney [4] - and generally assume a domain populated by technically skilled and understanding individuals such as pilots or astronauts. Efficient road traffic management is a major goal of research in intelligent transportation systems (ITS) [5,6], and the integration of multiple information sources is of significant interest for the needs of advanced traffic operations [7,8].

Winograd and Flores [9] have described the overall problems involved when modeling information systems and cognition using simply a rationalistic tradition. In particular, they outline some of the limitations inherent in traditional decision support systems. Heidegger's concept of 'thrown-ness' is cited as an example, wherein nobody is capable of true, detached situational assessments in everyday life, since we all are influenced by the moment-to-moment realities around us. They suggest that an alternative direction for research in this area is to start modeling activity in the form of networks of conversations among entities in the domain. Using this approach, protocols such as those encoded in speech act theory (requests, offers, acceptances, etc.) might be used in a somewhat ad hoc manner to cope with the situation at hand.

Some degree of insight may thus be gained into how a person models their situation in an environment by examining the nature of these type of networks and the semantics of the interactions involved. There are significant real-time data interpretation challenges in this arena, and a basic architecture must encourage the process of combining information from different modalities and of synchronizing different interpretation techniques. The application of this research approach in transportation systems has recently received encouragement by ARPA policy-makers in the intelligent systems arena [10]. The particular point about integrating multiple sources of information in a timely fashion is mentioned by the authors of the policy. Moreover, they also call for further work in improvements in 'common-sense' knowledge representation and reasoning, especially in the temporal and spatial domains.

We mentioned earlier that a basic feature of the traffic environment is that composite knowledge of the full domain is not totally resident in any one administrative entity, but is distributed among a number of organizations. Consequently, it may be suggested that some of the research in computer-supported cooperative work (CSCW) could be applied to modeling the joint operations activity of the traffic domain. After all, CSCW researchers do indeed address the broad topic of how a computer-based environment can improve the effectiveness of people working together as a group.

However, much of the CSCW research has been focused upon coordinating meeting-based activities such as joint document development [11], or supporting long-distance close collaboration on design or visualization, for example by using video whiteboards [12]. Less emphasis has been placed by the CSCW community on examining looser levels of coordination among individuals who are distributed among several organizations but who have various levels and extents of influence over a common domain.

Our goal in this paper is to identify some features of a human-machine system which can seamlessly link a control center operator with the information sources at their disposal, both those which are internal to the center and those from other organizations. We focus on using a distributed intelligent system as a metaphor for the traffic information-handling environment, and draw on some of the research in that arena to help characterize our architecture. Hortatory operation is the primary method in our domain for traffic controllers to influence the general freeway traffic conditions. The extent of an individual organization's actual control over the domain, as well the means of exercising it, varies considerably from one organization to another. Hence, we focus our architectural design upon the nature of the distributed management structure and the information exchange and analysis among the various supervisory entities rather than examining in detail their individual methods of influencing the domain.

The next section outlines the salient components and features of our distributed agents architecture. We also set out several of the major assumptions, and provide examples how the architecture conforms with some existing approaches to freeway traffic management. Section three focuses on the ontology of traffic information, and discusses how the concepts and relationships derived from a preliminary knowledge elicitation process map into the ontological structures in our architecture.

A sample encoded entry from the ontology is presented in section four, one which demonstrates how certain aspects of existing traffic analysis systems can be enhanced by using our approach. Section five shows how the intelligent reasoning mechanism in our architecture can be used to add some encoding structure to the processes and activity found in an operational environment. The final section notes that a distributed agent architecture can help provide a framework for future research in distributed system management and hortatory operations.

II. DISTRIBUTED AGENT ARCHITECTURE

This study draws on the computing concept of a system of distributed intelligent agents and uses it as a work context model for traffic control activity. In a freeway traffic management environment, a relatively loosely-coupled network of autonomous cooperating organizations share knowledge and real-time needs in a manner which enables each entity to achieve its own goals. These organizations include police, transit authorities, freeway traffic controllers and administrators, and so on.

The major components of the distributed agent architecture may be outlined as follows:

(i) The domain of discourse or 'target system', i.e. the freeway network and traffic conditions in it which are being supervised,

(ii) A distributed set of agents, each one of which represents a participating organization,

(iii) The communications infrastructure which provides the means for data to be passed among the agents,

(iv) A set of concept specifications (i.e. ontologies), commonly agreed upon and sufficiently seamless to enable agents to exchange meaningful information with each other,

(iv) a reasoning framework that structures an agent's inferences about conditions in the target system and in other agents.

Figure 1 shows an example of a target system and a set of five agents which are linked by a communications infrastructure. The Freeway Operations Center agent has been expanded to show internal structural details.

The distributed agent architecture relies on several basic assumptions about behaviors in the target system and among the agents themselves. We now describe the various components and the assumptions in more detail.

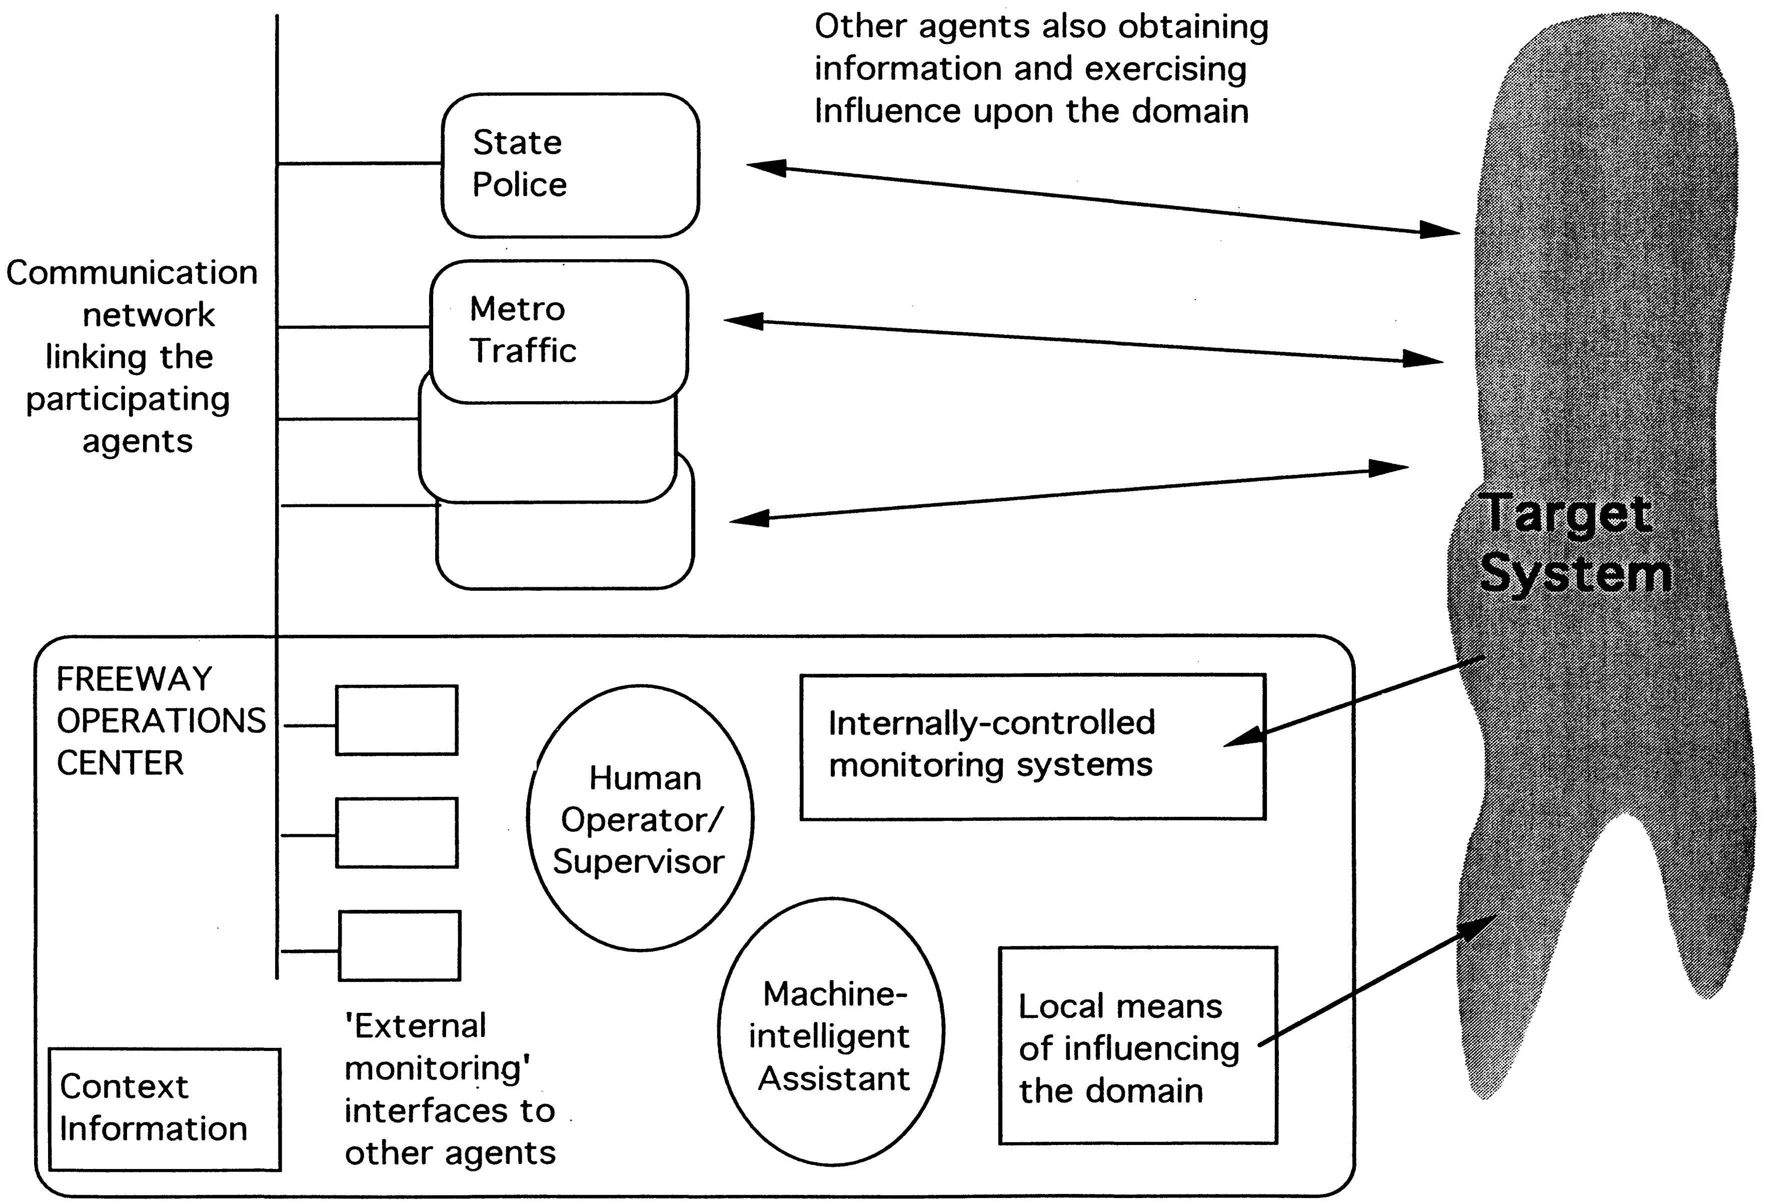

Overall Traffic Information Environment

DOMAIN OF DISCOURSE

The regional or metropolitan freeway network being supervised, and in particular the traffic conditions within it, constitute the target system or domain of discourse. A major operational goal of freeway traffic managers is to 'keep things moving'; thus, the ability to identify and locate obstructions quickly is of great importance. ITS researchers have devised a variety of incident detection algorithms to discern the difference between congestion resulting from accidents or other blockages and that which develops from capacity limitation (for example, the recurrent excessive demand on urban freeway systems during commuting hours) [13].

The principal source of data for such purposes is typically inductive loop sensors in the roadway pavement which detect the presence of passing vehicles. Information about average traffic speed, and its flow rate and spatial occupancy can be derived from the loop sensor outputs. Incident detection algorithms rely on an assumption that the measurement data exhibits different dynamic behavior patterns when an incident has occurred, and researchers have taken various approaches to characterizing these properties. Most algorithms can be classified into one of four types: pattern recognition, statistical models, neural networks, and catastrophe theory. However, many of them are not very effective in moderate or heavy traffic conditions, and may often require considerable retuning and adjustment for local conditions.

Many knowledge systems or methods suffer from the problem of trying to force a structure onto what is essentially incomplete, vague, and possibly inconsistent knowledge. This constitutes part of the basis upon which the need for bounding the domain of discourse rests. In the 'real' world however, people more or less manage to muddle along, as has been discussed by Norman [14] among others. A disadvantage common to many of the incident detection algorithms is that they are ill-suited to adaptation to incorporate other sources of information which are typically used in traffic operations environments. Therefore, an information modeling process which identifies an incident solely from some observed sequence of data values may not really reflect the actual process which a human operator uses.

It should be noted that the freeway network itself is simply the domain of discourse in this architecture, and therefore we do not attempt to characterize individual vehicles in the domain as separate entities. This is unlike some other models which use intelligent agents in transport control situations, such as air traffic control environments and cooperating robots [15,16]. The essential issue differentiating those domains from ours is that they permit appreciably greater levels of control over the individual participating vehicles than in the freeway traffic case.

DISTRIBUTED AGENTS

The architecture incorporates a set of agents, each one of which represents the unique characteristics of a individual participating organization such as a traffic information service or a transit operator. Each agent is comprised of three major elements, as follows:

(i) Some means of directly obtaining particular types of information about the state of the domain using its own monitoring or surveillance sub-systems,

(ii) Some locally-controllable methods of influencing specific aspects of traffic behavior in the target system,

(iii) Some intelligent entities which dispatch, supervise, monitor, or otherwise administer the organization's participation in the overall management activity.

For example, a freeway operations center might receive continuous telemetry about the current average speeds and densities of traffic in the network. It may provide some mechanisms for exerting influence over the domain, such as a system to display traffic status information and warnings on changeable message signs. And it probably has some human operators or machine intelligence which can reason and make inferences about general levels of congestion, etc. Similarly, a commercial traffic information service may have some surveillance aircraft, a set of contracted broadcast timeslots, and some dispatchers and announcers who control the operation.

An agent also has access, either implicitly or explicitly, to background or circumstantial information which establishes the context for its activities. The sources for such data in the traffic environment include event schedules, construction plans, and weather reports, as well as general knowledge about human behavior, workday patterns, and so on.

COMMUNICATIONS INFRASTRUCTURE

Agents typically have access to a set of communications links which interconnects them. In some cases, different pairs of agents may communicate using different methods, while in other cases a group of agents may utilize a common infrastructure to share information. Architecturally, the communications infrastructure can be regarded as a series of layers, such as the OSI Reference Model [17], which addresses the underlying technological support needed, as well as the various interaction protocols and methods. Inter-agent communication may often supported by ancillary equipment or links, rather than being deeply incorporated into an agent's own instrumentation. This typically occurs because the overall set of agents is comprised of a variety of private and public entities with limited centralized coordination, and hence the interrelationships have evolved over time in a relatively ad hoc manner.

The specific protocols used for individual knowledge transfer on inter-agent links can vary depending on the participants and the message, but broadly they will exhibit the following attributes:

(i) One-to-one transfers vs. broadcast messages. There may be direct one- or two-way communication between two specific agents, or one agent may distribute information which is accessible by all other agents.

(ii) Solicited vs. unsolicited knowledge. One agent may request another to provide some specific piece of information. Alternately, an agent may discover or 'overhear' relevant information when monitoring another agent's communications.

(iii) Acknowledged vs. unacknowledged messages. The issuance and receipt of acknowledgments can influence one agent's understanding about the state of another agent's information.

For example, a freeway operations staff member may receive routine traffic status reports via fax from state police controllers, overhear CB radio conversations about congestion among truck-drivers, request (and acknowledge) clarification of an incident from a transit dispatcher, and routinely monitor one or more commercial traffic reporting services.

SET OF ONTOLOGIES

All knowledge-based systems rely either explicitly or implicitly upon certain underlying assumptions about the world. In order to permit communication and knowledge interchange between intelligent systems, there must be agreements about these underlying assumptions. Formal ontologies have been presented as a specification mechanism for this purpose, and a specialized discipline of ontology design, akin to software engineering, has been postulated [18]. In practical terms, common ontologies determine the vocabulary to be used for conducting interactions between agents. Unlike a taxonomy, which refers simply to a classification of domain objects, an ontology permits a wider universe of discourse by encompassing the relevant axioms and attributes of an object which may be reasoned upon.

Sharing ontologies does not correspond to the sharing of knowledge per se; each agent maintains knowledge about things which the others do not. Also, an agent which commits to supporting a particular ontology is not necessarily capable of addressing all queries which are posable using the ontology. Incorporating shared ontologies into our architecture helps reflect the general consistency among traffic operations agents about concepts in the domain, as well as admitting incomplete knowledge about them. The set of ontologies needed to characterize the traffic operations environment has four broad sub-sections or components:

(i) A "traffic-information" specific ontology which is shared among the agents, and to which they commit in the sense of supporting basic axioms about it.

(ii) A "general" or "background" ontology, addressing contextual concepts such as weather conditions and time-of-day, as well as day-to-day knowledge such as spatial and temporal primitives like "north of" or "happens during".

(iii) A "local" or "group" ontology, which is concerned with the knowledge and procedures used predominantly within each organizational agent. This ontology is unique to each agent, but may also be partially understood by other agents.

(iv) An "individual" ontology which reflects the concepts used as part of the internal expertise or reasoning process of a human or an intelligent machine which is part of one of the agents.

The first two components constitute the shared ontologies to which all the participating agents are committed, while the latter two are primarily unique to the individual agents. The domain-specific ontology is the set of concepts that all agents use to characterize entities and events that are specific to the freeway traffic domain and are of common interest. The context ontology is essentially a representation of the 'common-sense' concepts which are beyond the scope of the domain-specific one; this is the field of foundational knowledge, and such an ontology has been under development for some time [19].

The group-level ontology represents the set of unique entities, relationships, and events which characterize each agent organization's methods and activities. Such an ontology corresponds directly to the results which one might be expect from a comprehensive object-oriented analysis of an organization's internal procedures and systems. The fourth "individual" ontology is in essence the design rationale for a particular intelligent machine or the components of human's mental model [20].

Figure 2 provides an example of how three agents utilize the various types of ontology for sharing knowledge and for their own internal purposes.

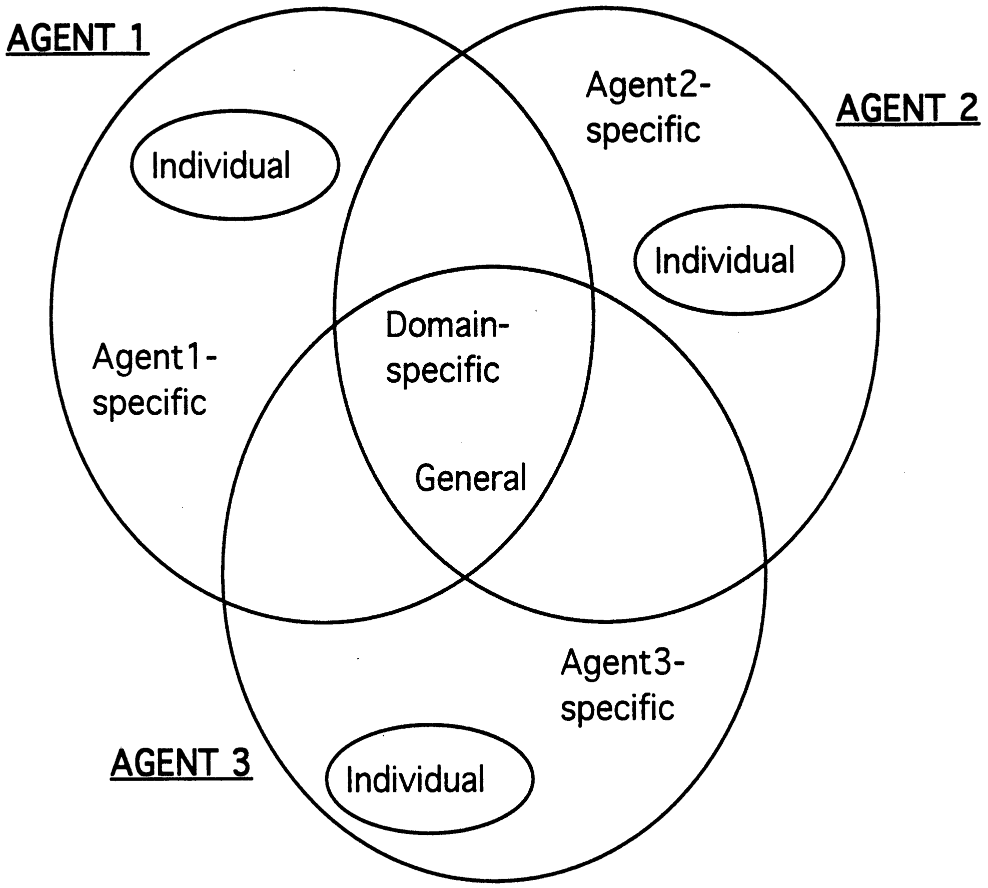

Figure 2: Ontologies shared among agents

REASONING MECHANISM

Our distributed agent architecture requires a schema for analyzing causal relationships and reasoning about the concepts and entities identified in the set of ontologies. Specifically, we need to identify a suitable means of representing the internal information-handling and knowledge processing of the various participating agents. A schema which uses a strict hierarchical representation of the physical domain is unlikely to be accurate, given the nature of the operating processes, in particular because the channels by which information becomes shared among the agents are not totally formalized. For example, an individual police officer or transit driver may receive information which is more pertinent to their own immediate situation via other agencies than they do from their own organization's hierarchy. Therefore, agents can influence each other's internal actions and information flows in ways which are not easily accommodated by utilizing a hierarchy consisting purely of systems, sub-systems, and components.

Rasmussen [21] presents an abstraction model which avoids simply decomposing the system into its sub-parts and components. It is composed of five hierarchical layers based on different types of system principles instead of using levels of physical structure. Thus, the characteristics of individual components in the domain are addressed at a different layer than the more abstract concepts like functional behavior models or issues related to strategic policy. (Of course, it may also be the case that concepts at the higher levels of abstraction are more applicable to the larger sub-systems in the domain.)

Additional work on this schema presented by Goodstein & Rasmussen [22] discusses a format for knowledge representation which overlays a three-layer reasoning structure onto the five-level hierarchy. The composite schema is illustrated in Figure 3. Goodstein & Rasmussen's structure is designed to address three basic issues concerning direct control over a target system - (i) what entity should be controlled, (ii) the means with which it is controlled, and (iii) the rationale for controlling it.. In the information-handling activity of our hortatory operations environment, we re-interpret these for the purpose of our architecture as (i) what is happening, (ii) how has it come to attention, and (iii) why is it of importance. The reasoning structure can be applied at various levels in the means-ends hierarchy so that, if a target system entity requiring attention exists at one of the means-ends level, the rationale behind that need is formulated at a higher level. Conversely, the methods by which it comes to attention exist at a lower level in the means-ends hierarchy.

This adapted schema provides our architecture with a suitable method for characterizing the reasoning mechanisms used in participating agents.

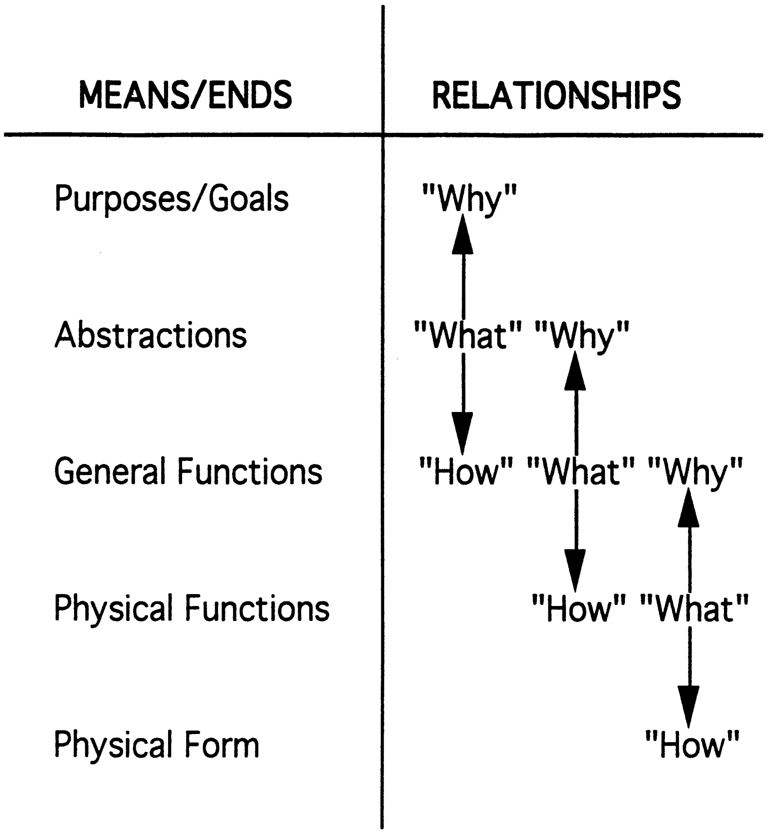

Figure 3: Decision-making space (from Goodstein & Rasmussen 1988)

ASSUMPTIONS

A basic requirement of the distributed agent architecture is that all agents are in agreement about a basic set of tenets concerning the domain; this is inherent in the commitment to sharing the domain-specific ontology. A further requirement is that the agents are mutually trusting, in other words that they do not deliberately try to mislead or outmaneuver one another. The architecture represents the domain of traffic control with the following characteristics:

(i) The set of participating agents consists of a variety of public and private information-gathering and administrative organizations, such as municipal roads and freeway administrators, transit system operators, police and emergency service authorities, and commercial traffic information services.

(ii) The environment in which the agents work essentially consists of a very large number of independently-operated vehicles using the road and freeway network. Only a insignificant percentage of these are directly controllable or dispatched by any of the participating agents.

(iii) None of the agents has central authority or control of the domain, either the target system or the information gathering and monitoring system.

(iv) Each agent has some means of obtaining certain information about traffic conditions, which helps add to the overall knowledge about the state of the domain.

(v) Agents may be able to take some action which influences the domain, and may be able to adapt or coordinate those actions to facilitate the other agents.

(vi) An agent has reasonable knowledge and understanding of its own monitoring and control systems, their status and limitations, but may not have a very accurate or complete picture of how the other agents' internal systems and procedures operate.

III. TRAFFIC INFORMATION ONTOLOGY

We now examine part of this architecture in some further detail by looking at how the major concepts and entities used by one agent organization map into the ontological structure outlined above. Operators at a freeway operations management center participated in knowledge elicitation sessions where they were interviewed about how they gathered and interpreted traffic information in their normal working environment. (We recognize of course that this process gives only a 'keyhole view' - and through the eyes of just one agent - into the overall architecture.)

Figure 4 summarizes the static, IS-A relationships among the main concepts which arose during the knowledge elicitation sessions. The main categories which emerged from an analysis of the discussions were the relevant physical entities, either in the target system or in the monitoring system, the background or task context entities, and the messages and variables which influence actions - these are referred to as evidential messages. The details of these concepts and their mapping into the set of ontologies in our architecture are discussed below.

Fig 4: Traffic Information Relationships

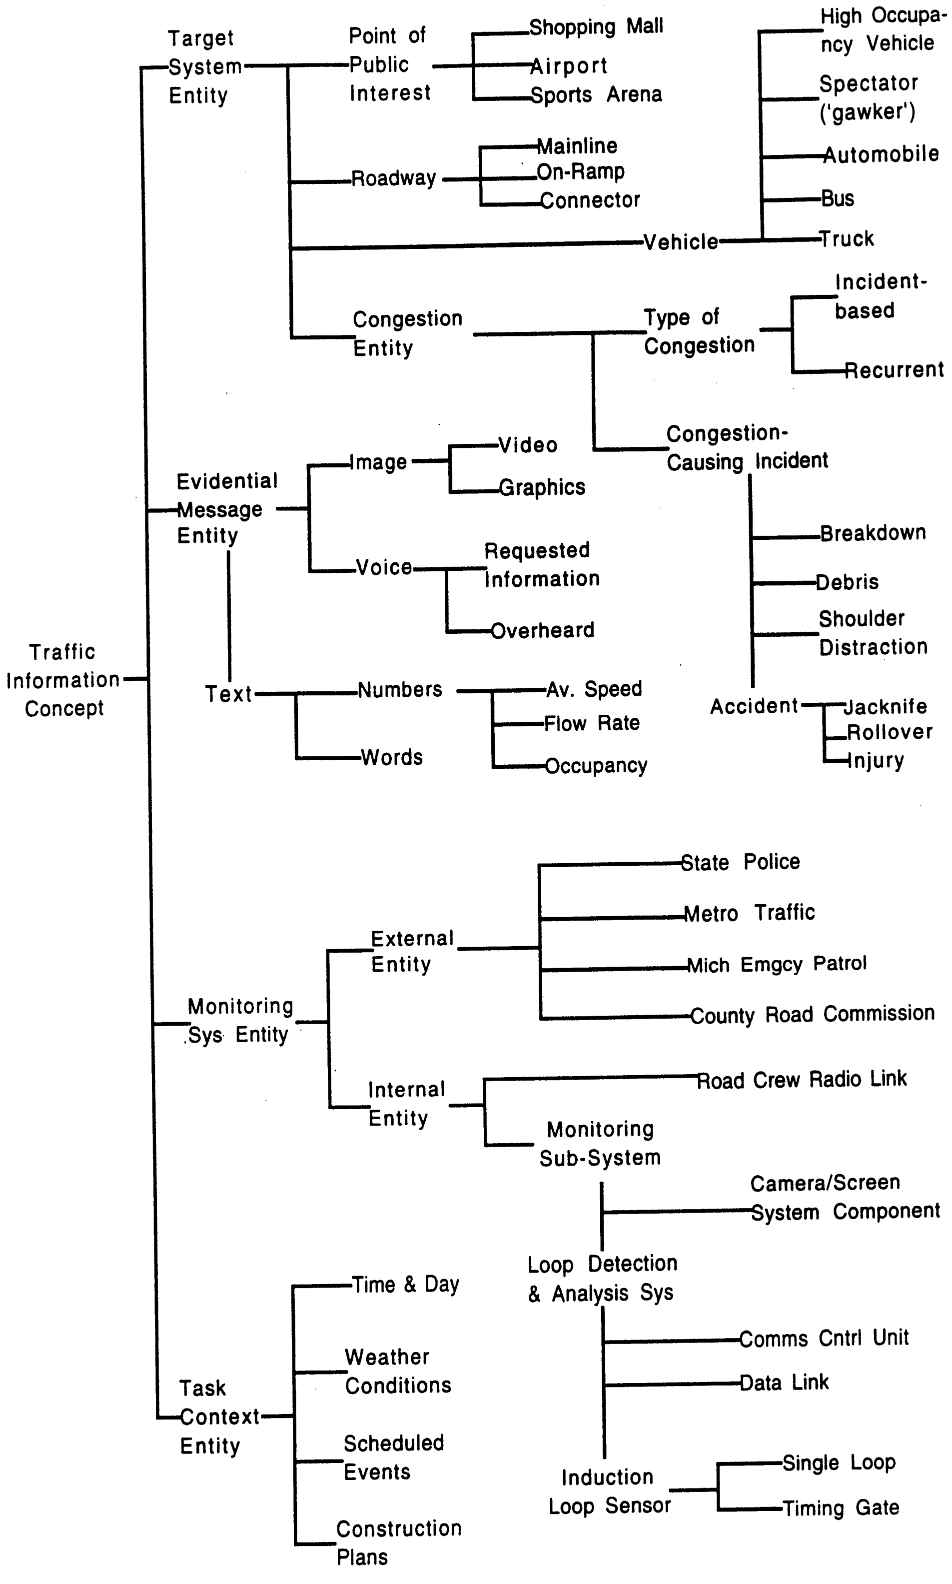

TARGET SYSTEM ENTITIES

These are terms which describe physical entities in the real world; examples from the traffic domain are roadways, vehicles, incidents, and places of interest. These items are further divided into narrower type specifications; for example roadways might be classified into on-ramps, interchanges, connectors, and so on. The specification is only made to the extent that the differentiation has significant relevance to traffic activities, demand on capacity, etc. Thus, while an airport might be a place of interest which would figure strongly in traveler navigational needs and roadside signs, it would only be required in a traffic operations ontology if its presence has an impact on ground traffic levels.

It is important in designing an usable ontology to clarify the sometimes subtle difference between domain entity types and the roles they play. For example, the terms 'automobile', 'truck', and 'bus' refer to types of the more general concept 'vehicle'. But the term 'HOV' (high-occupancy vehicle) can be used both as a concept (of which by definition all buses are types) or as a category which is entered for particular journeys by certain vehicles like automobiles. And an expression like 'spectator' (also 'gawker' or 'rubber-necker') exclusively describes a role which can on occasion be played by virtually any vehicle.

The target system entities map predominantly into the shared, domain-specific ontology in our architecture, insofar as these entities form the common subjects of communication between the various participating agents.

MONITORING SYSTEM ENTITIES

These terms represent the physical objects and organizational entities which form the primary sources of information about the target system. The basic sub-categories are internal sources and external sources. Internal information sources include the monitoring and surveillance equipment, analysis computers and visualization systems, and other instrumentation which the operator manipulates and uses as tools. These sources can also include two-way communications links to field personnel for whom the operator may act as a primary home base contact or dispatcher.

An external source of information is another agency or organization, together with its interfacing equipment and communications links, for which the operator has limited responsibility or control. In the traffic domain, these sources may include links to police dispatch messages, commercial traffic information services, municipal transit reports, and so on. Although operators may have some means to exchange information with the external sources, the opportunity to access and manipulate the other entity's equipment and tools is very limited or non-existent.

Most of the internal sources of information map into the group-level or agent-specific ontology, since internally-controlled monitoring systems are essentially unique to each agent. The monitoring entities defined as external sources constitute the ontological representations of the communications links which one agent has with another in the distributed network. These entities are likely to map into both agent-specific ontologies.

TASK CONTEXT ENTITIES

The background information environment within which actions are selected is captured in the task context. These entities represent factors which influence the target system conditions and may modify the actions of the operator - things they must cope with but over which they have essentially no control. Context entities may be implicit - for a traffic operator, entities like weather conditions or time-of-day are implicit - or explicit like roadway repair and maintenance activity, sports events, and so on.

Context entities are rarely the central focus of the operator's attention, but can affect their level of situation awareness. For example, rush hour congestion can degrade the monitoring system's ability to identify incidents like accidents or breakdowns. Similarly, the manner in which the background information is made available to the operator is generally of little relevance to their task effectiveness. Provided the wall-clock is legible and reasonable notice of sports events is available, context entities form the backdrop to normal work routines rather than taking center stage. Context entities map directly into the shared, common-sense or background ontology.

EVIDENTIAL MESSAGE ENTITIES

These concepts concern the primary items which influence human operators or their machine-intelligent assistants in their day-to-day actions. There are two subtypes to these entities: the codes - words, numbers, images, and so on - which are used as justification for selecting a particular action (or inaction), and the modal form - video/graphical/text and aural/speech/alarm signals - in which the codes are presented. Evidential message entities are the central items of focus with which the operator or assistant builds a coherent picture of the world.

Many of the codes are abstract or measured values like 'average speed' or 'occupancy' which are obtained using information processing in the monitoring system. Others like 'southbound' are used to modify or identify target or monitoring system entities in different ways - such as 'roadway' or 'vehicle'.

In some cases, there may be agreement among agents about the actual domain concept which is being represented in a shared ontology, but the function which the code for a concept serves internally for each agent may be different. For example, the traffic operator's classification of incidents in the target system such as breakdowns and accidents is based primarily on likely durations, number of lanes blocked, and so on. Accident characteristics such as the type of vehicles involved or the number of persons injured is useful to the traffic operator only for assessing the impact on the extent and severity of congestion. But on the other hand, these particular codes will have a more salient role for emergency service personnel. In this situation, there is a variation in function which a common ontological concept provides for different agents.

The modal forms in which the codes are presented are determined by the monitoring system user interface hardware, but in ontological terms are fundamentally different to it. The operator uses their understanding of the monitoring system and its user interfaces to help explain possible anomalies in the coded information; for example, they can use test procedures and verification methods to resolve problems and inconsistencies. On the other hand, the modal form of code presentation influences cognitive processing of the information, and affects the rate of interpretation and efficiency at handling simultaneous inputs from multiple sources.

Thus, we note that the evidential messages are precisely the entities which provoke specific actions (or inactions) on the part of intelligent individuals. They form the bridge between the physical world and the world models used by each intelligent individual, and therefore map into the individual-level ontology in our architecture.

IV. SAMPLE ONTOLOGY ENTRY

The focus upon developing an ontology of the environment, rather than simply encoding the expertise into a knowledge base, encourages the development of more faithful models of tasks undertaken by agents. It also assists the design of an underlying knowledge system architecture which can better be integrated into the work environment.

For example, we mentioned in the last section that, in the right circumstances, any vehicle can be described as a spectator or gawker. Significant accidents and breakdowns tend to cause secondary traffic congestion on adjacent roadways - typically those traveling in the opposite direction - because of spectating. Delays due to such secondary congestion are rarely as severe as on the primary roadway, but traffic professionals report that they use such patterns as additional evidence in their search to identify and confirm incidents. It is therefore desirable that the concept of spectator be incorporated into our ontology.

One method for specifying ontological entities is provided by the Knowledge Interchange Format or KIF [23], which emphasizes the representation of epistemological knowledge, rather than the expression of heuristic information. It is a format which can be translated into various implementation languages, which thus permits the sharing of ontologies among different users and systems. KIF provides the mechanism for linking the entities for roadways, vehicles, types of incidents, etc. and their attributes together to represent a given conceptualization.

Since the domain of discourse is traffic congestion, the defining criteria for spectators are related to the behavior of vehicles rather than the behavior of their occupants. Thus, spectator is a class of vehicle, membership of which is determined by some characteristic behavior patterns. In particular, a spectator is located on a stretch of road where the average speed is low, and downstream of which the average speed is not low. Furthermore, a reason for spectating must exist. This could be a suspected incident on another road stretch, in which case that road stretch must be visible from the spectator's location. Alternatively, a shoulder distraction condition must exist for the spectator's own location. The KIF expression of this conceptualization is as follows:

```
(define-class SPECTATOR (?s)
  ; Spectator is a role class played by vehicles
  ; which can be distinguished as follows. . .
   :def (and (vehicle ?s)
          (roadway ?r1)
          (located ?s ?r1)
          (average-speed ?r1 low)
          (roadway ?r2)
```

```
(downstream ?r2 ?r1)
(not (average-speed ?r2 low))
(or (and (shoulder-distraction ?d)
         (located ?d ?r1))
    (and (roadway ?r3)
         (suspect-incident ?i)
         (located ?i ?r2)
         (visible-from ?r2 ?r1))))
```

The rationale behind selecting particular conceptual structures can be examined using Gruber's five design criteria for analyzing an ontology - clarity, coherence, extendibility, encoding bias, and commitment.

Clarity focuses on objective and preferably complete definitions. The spectator class is restricted to those vehicles (not people) which are in a region of discontinuous traffic speeds. If there is no effect on general traffic flow, then the existence of spectators is irrelevant. Coherence refers to the maintenance of logical consistency - for example, membership of the spectator class of vehicle does not preclude membership of another vehicle class such as truck or auto.

Extendibility concerns the promotion of sharing by the affordance of specialized usage. A more technical traffic measure such as occupancy can be utilized in conjunction with the more everyday concept of average-speed by interested agents without requiring other, less-specialized agents to adopt it also. Encoding bias should be minimized by avoiding specific representation requirements. The use of 'low' and 'not low' as thresholds in the definition permits a multiplicity of representation schemata.

Ontological commitment should also be minimized by making as few claims as possible about the world. This is the criterion behind the existence requirement in the spectator class for a shoulder-distraction or suspected-incident. Since the domain of discourse is highway traffic conditions, an off-highway or distraction on the shoulder such as a fire should be confirmed before assignments to the class are established. On the other hand, we select the claim of suspected-incident, which is weaker than an incident which has already been confirmed, as sufficient for membership. This permits the use of spectating behavior as evidence to help identify actual on-highway traffic incidents, which more accurately reflects the real traffic operator's reasoning method in particular circumstances.

V. ANALYSIS OF REASONING PROCESS

It is desirable to examine how effective the reasoning mechanism part of our architecture is in formatting the causal and heuristic part of the traffic information domain. We noted earlier that concepts used at higher levels of abstraction may be more applicable to the larger physical sub-systems in the domain, but that there was not necessarily a direct mapping between them. Certain features and attributes of the physical structure of the system may therefore fall out of this type of analysis, but they will occur in the context of a person's reasoning about the domain rather than in the system designer's specifications, engineering drawings, and so on. Hence, an analysis using this architecture should provide a better insight into the perceived and salient functioning of the domain, one which will be more focused on information and less on structure.

One should note that the layer boundaries in our architecture are not clean-cut and, in using it to help with knowledge structuring, one needs to recognize that people's intuitive feel for the domain is properly represented. This implies that the outward character of the knowledge structuring is in tune with the real-world concepts and procedures of the working environment - even though the actual underlying automated implementation may be relying on different paradigms. This is a point which parallels the encoding bias issue concerning ontologies which was made earlier.

The record of the knowledge elicitation sessions with the traffic management operators was analyzed to investigate the extent to which the participants' descriptions of their activities ranged over the levels in the hierarchy. The main topic of interest was the process of identifying traffic congestion patterns which result from accidents or other incidents. A preliminary assignment of some topics and concepts from the traffic ontology to the means/ends levels is as follows:

(i) Purpose/Goals Level: Organizational policies & strategy, values such as accuracy in reporting congestion and maintenance of public confidence.

(ii) Abstractions Level: Classical models of typical traffic behavior, such as forced vs. free speed, propagation & backwave phenomena.

(iii) General Functions: Variables and measures of influence, such as average speeds or flow rates, and routine operational procedures for handling incidents, interaction with other agents, etc.

(iv) Physical Functions: Typical vehicle behaviors, general characteristics of surveillance instrumentation, data display facilities, etc.

(v) Physical Form Level: Specific features and problems with monitoring units, unique features of certain locations, administrative issues and paperwork.

It was noted that the more experienced individuals tended to move up and down through these levels more fluidly and with greater rapidity. They also were more inclined to explain domain activity using more abstract and theoretical terms like flow-speed curves and free mean speed. On the other hand, the more junior operators focused their remarks in a tighter range levels and were more inclined to characterize the target system in more concrete terms.

For example, in one piece of conversation, an experienced operator moved very quickly between a mid-level assessment of the congestion likely to result from a particular obstruction, a low-level concern over potential faults in the sensing equipment, and some high-level strategic comments about control center actions might impact public credibility of the changeable message system. The corresponding conversation of a less experienced operator focused on the impact of the obstruction and the appropriate procedure for addressing it, without mention of monitoring system details or of global strategy.

The point to be made here is not that the variation in operator experience results in significantly differing ability to do the job (which in itself may or may not be true). Rather, our analysis suggests that an individual's experience may affect their likelihood to move as fluidly among the five means-ends levels when describing the target system and their own work procedures in a detached setting.

VI. CONCLUSION

We have discussed the development of a distributed agent architecture which can be applied to operational environments where direct control over a target system is both limited and decentralized. The architecture addresses the following features; domain of discourse, participating agent organizations, inter-agent communications means, ontological structuring for sharing knowledge, and intelligent reasoning processes. A static set of IS-A relationships among concepts were derived from knowledge elicitation sessions pertaining to one of the agent organizations, and the mapping of these to our ontology architecture was described. We showed how a focus upon the design of ontological entities which reflect real-world concepts can help improve existing, more theoretically-based algorithms and methods. An example was also given of how our architecture's intelligent reasoning features can be applied to the analysis real-world operational activities.

Decentralized system management and hortatory operations are both common characteristics of present-day work environments; the distributed agent architecture set out in this paper provides researchers in this arena with a preliminary formalism for structuring their investigations.

REFERENCES

[1] Sheridan T, "Telerobotics, Automation, and Human Supervisory Control", MIT Press, 1992.

[2] Mitchell C, "GT-MSOCC: A research domain for modeling human-computer interaction and aiding decision making in supervisory control systems". IEEE Trans. on Systems, Man, and Cybernetics -SMC-17, pp553-570, 1987.

[3] Hopkin V, "Air traffic control" in E. Wiener and D. Nagel (Eds), Human Factors in Aviation, Academic Press, pp.639-663, 1988.

[4] Kearney M, "The Evolution of the Mission Control Center.", Proceedings of the IEEE, Vol 75, pp399-416, March 1987.

[5] Saxton L (Ed), Special Issue on Intelligent Vehicle Highway Systems (IVHS), IEEE Trans. on Vehicular Technology, February 1991.

[6] Collier C & Weiland R. "Smart cars, smart highways", IEEE Spectrum, Vol 31 No 4, 1994.

[7] Morris J & Marber J, "Virginia Traffic Management". ITE Journal, Vol 62, July 1992.

[8] Pierce V et al. "The San Antonio Advanced Traffic Management System", ITE Journal, June Vol 64. 1994.

[9] Winograd T & Flores F, "Understanding Computers and Cognition", Addison Wesley, 1987.

[10] Grosz B & Davis R, "A Report to ARPA on Twenty-First Century Intelligent Systems", AI Magazine, Vol 13 No 3, Fall 1994.

[11] McGuffin L & Olson G, "ShrEdit: A Shared Electronic Workspace", CSMIL Technical Report 45, University of Michigan, 1992.

[12] Tang J and Minneman S, "VideoDraw: A Video Interface for Collaborative Drawing", Proceedings of CHI'90, ACM New York, 1990.

[13] Chen C & Chang G, "A Review of Recent Freeway Incident Detection Algorithms", Technical Report DTFH61-92-R-00122, University of Maryland, 1992.

[14] Norman D, "The Psychology of Everyday Things", Doubleday New York, 1988.

[15] Cammarata S et al, "Strategies of Cooperation in Distributed Problem Solving", Proc. of the Eighth IJCAI, Karlsruhe, West Germany, 1983.

[16] Steeb R et al, "Cooperative Intelligence for Remotely Piloted Vehicle Fleet", Technical Report R-3408-ARPA, Rand Corporation, 1986.

[17] Day J & Zimmermann H, "The OSI Reference Model", Proceedings of the IEEE, Vol 71, December 1983.

[18] Gruber T, "Towards Principles for the Design of Ontologies used for Knowledge Sharing", Technical Report KSL 93-04, Stanford Knowledge Systems Laboratory, 1993.

[19] Guha R & Lenat D, "Enabling Agents to Work Together", Communications of the ACM, Vol 37 No. 7, July 1994.

[20] Johnson-Laird P, "Mental Models", Harvard University Press, Cambridge MA, 1983.

[21] Rasmussen J, "A cognitive engineering approach to the modelling of decision making and its organization", Technical Report Riso-M-2589, Riso National Laboratory, Denmark, 1986.

[22] Goodstein L & Rasmussen J, "Representation of Process State, Structure and Control", Le Travail Humain, Vol 51 No 1, 1988.

[23] Genesereth M & Fikes R, "Knowledge Interchange Format Reference Manual", Technical Report Logic-92-1, Stanford Computer Science Department, 1992.