The Colloquium: An Ontological Characterization of Distributed Human-Machine System Supervision

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-95/18

[Submitted to the International Journal on Human-Computer Studies on August 4 1995]

The Colloquium: An Ontological Characterization of Distributed Human-Machine System Supervision

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

SUMMARY

The work of supervising complex networks and systems is not always a highly centralized and influential activity; sometimes, supervisory operations are broadly distributed among numerous autonomous organizations. In such cases, the administrative structures and goals may permit only quite loose coordination of their activities. Non-homogeneous methods and procedures for collecting and analyzing data are likely to be used in each operations center. Furthermore, each center may have quite limited direct influence over the actual events in the environment. This paper sets out a characterization of the analysis and problem-solving environment in which such operations centers function. The environment is referred to as the "colloquium"; it is distinguished from the "target domain", which is the common network or application system being administered. A description is also presented of some of the differences between the terms static and dynamic, when applied to these two environments. This establishes the background for a discussion of the theoretical basis underlying the design of the colloquium, which is presented in the form of a set of ontological entities. Intelligent agents that participate in the colloquium can improve their levels of mutual comprehension and cooperation by committing to this shared ontology. The paper culminates with some coded examples of how features of the colloquium can be structured in the ontology, and how it can be integrated with an ontology representing a target domain - one for the application area of shared road traffic information.

INTRODUCTION

The deployment of highly automated systems in a variety of complex industrial and administrative applications often involves the creation of sophisticated control centers where large amounts of real-time process information is accessible to the staffing personnel. Human-machine interaction researchers have noted that the typical role of such people has changed over time from that of direct control over the system to a more hands-off, supervisory role [Sheridan 1992; Liu, Fuld & Wickens 1993]. As a result, the operator's task in many such applications involves ensuring that the system remains in some nominal status condition - essentially an on-going vigilance that is intermittently broken by periods of fault identification, error analysis, problem resolution, and so on.

Research attention in this arena has traditionally been focused on work settings where operators have a significant level of control over the environment, such as flexible manufacturing plants, chemical processing, and aerospace control centers. However, such approaches do not accurately capture the activity in some operations environments - ones where there is quite limited direct control over the domain, and where operators must rely heavily upon encouragement and advisory communications to manage the system. We have discussed this phenomenon elsewhere [Murray & Liu n.d.], and have coined the term 'hortatory operations' for such environments.

Another shortcoming of existing research is that it often cannot accurately encapsulate an operations environment where there are several different organizations involved in administering the system, with the result that each one has only partial data directly available about the state of the domain. The various organizations may have different goals and priorities, but they recognize the need to pool their information in order to be effective. It is desirable therefore to try to characterize a knowledge system architecture which will assist the participating organizations to share all available information and to evaluate it in context - that is, as a common pool of loosely cooperating agents.

An ontology is an information structure used in knowledge engineering which permits the formal representation of a set of concepts, that is the entities and their inter-relationships which are assumed to exist in some domain of interest. Sharing some group of ontologies among several intelligent agents or other knowledge-based systems - or "committing" to them - is intended to harmonize the underlying agreements about the concepts among the participants. Such a commitment does not necessarily require any of them to support the ability to interpret and reason upon all facets of the shared concepts.

One of the limitations that is evident from published examples of knowledge-sharing in ontologies [Gruber 1992,1993] is their focus upon a static set of domain conditions - a static context as it were. Although they permit the sharing of generic knowledge about dynamic activity (such as spatial velocity primitives), the ontologies do not provide support for the updating and revision effects that such activity can have upon the information being shared among several intelligent agents. Agents dealing with a dynamic environment need to cope with on-going circumstances as they unfold over time. It therefore becomes necessary to introduce some concepts related to the cooperative reasoning and problem-solving context itself in order to deliberate upon the state of such a domain. These shared entities and their attributes need to be derived from the circumstances of the inter-agent communications, the underlying beliefs and assumptions they may have about each other, and so on. In this paper, we gather this set of concepts together and use them to structure the architecture of an autonomous reasoning and problem-solving environment - we refer to this domain as the "colloquium".

Researchers working on context-sensitive intelligent assistant systems have recognized the particular importance of viewing knowledge in context. While there may be disagreement between cognitive scientists and system engineers upon how to define context per se, there is general consensus on the existence and role of many of the elements of context. Furthermore, as Brezillion & Abu-Hakima [1995] have pointed out, there is also concurrence that the contextual component of knowledge is rarely accommodated with an explicit representation in intelligent system models. In describing situated agents and knowledge in context, the authors differentiate between the 'situation' - the overall physical environment including the agents - and the 'reasoning environment', which they describe as all of the 'mental stuff inside agent(s)'. In the present paper, we further break down that reasoning environment by distinguishing between entities that represent real-world concepts and those which pertain to the colloquium's internal inference processes themselves - the real-world components are referred to as relating to the "target domain".

The boundary between colloquium and target domain can get very blurred when agents participate directly in the target domain itself. The need for agents to model each other's behavior then becomes inextricably bound up with what goes on in the world [Nielsen 1995]. Focusing the modeling efforts upon hortatory operating environments offers a more tangible and definitive separation between the two domains. It also helps maintain a focus on the knowledge processes and cognitive aspects of the colloquium, precisely because the opportunity for exerting control over the target domain is limited.

The separation between the colloquium and the target domain is further refined in conditions where agents' subsystems for surveillance and influence are not closely merged with the target domain infrastructure itself. This can permit agents to differentiate more easily between the symptoms of target domain problems and those emanating from failures within each other's own equipment. It is also preferable that the communications between agents are not dependent upon the target domain infrastructure; again, this can improve the decoupling between the two environments.

We select road traffic management as the hortatory operations environment to be used as a sample application of our architecture. In this environment, knowledge about target domain conditions is distributed among several different entities such as freeway traffic managers, police, municipal authorities, transit operators, and commercial traffic information services. Access to this real-time knowledge is governed by the nature of each organization's information-gathering methods, which are primarily focused upon their own traffic-related responsibilities - typically involving monitoring traffic levels, identifying and resolving problems that cause congestion or danger, coordinating construction activity, and so on. Such organizations can help alert network users to current or expected conditions and can urge them to take certain actions but, unlike an air traffic control center say, they have little ability to control the actions of individual network users. Instead, road traffic operations rely heavily upon the use of information distribution and encouragement as a means of influencing the domain, rather than physical manipulation or administrative authority.

In conclusion, it may be noted that the operation of certain complex computer and telecommunications networks - such as the Internet - exhibits certain hortatory characteristics. The distributed form of such a system's locus of control, together with the limited level of influence available to any individual administrator, constitute the main features of a hortatory environment. However, the high degree of coupling between the infrastructure of the target domain network and the colloquium's internal communications system results in a much more detailed distinction separating the two environments.

The following is a summary of the remaining sections of this paper. The distinguishing features of the target domain and the colloquium are first explored in some more detail. Attention is also drawn to the variance between static and dynamic domains in the architecture. This is followed by an examination of the theoretical basis for articulating those attributes that are particular to the colloquium The subsequent section defines the primary ontological entities of the colloquium in detail, and that is followed by coded examples which demonstrate how the shared ontology is structured. In the final section, some of the potentials for further research work in this area are discussed.

ENVIRONMENT TYPES

The principal characteristic of the architecture discussed here is the division of the working arena or environment into two regions - the target domain and the colloquium. The target domain is specified as the application area or real-world system being monitored and analyzed. Ontological entities representing objects and states in the real world are said to be entities of the target domain. The target domain is sometimes referred to in the literature as the 'domain of discourse', since it is the source of most of the relevant data and the object of most of the administrative and supervisory actions.

The colloquium on the other hand is the environment in which the analysis and inference take place. It was mentioned earlier that a basic assumption of the architecture is that participating agents are not fundamentally integrated into the world being modeled. Rather, they stand off to the side in a monitoring or supervisory mode - this being particularly the case in hortatory environments where the extent of real control over the world is relatively limited. Ontological entities concerned with the reports, messages, and beliefs shared among the agents are said to be entities of the colloquium.

A second characteristic of both of the environments in this architecture is the extent to which dynamic behavior plays a role in the activity in the environment. When there is only limited need for updating, revisions, and other temporal elements in an environment, it may be regarded as static. Ontological entities commonly used in static environments tend to have attributes that remain constant or change infrequently in real-world time-scales and/or during the analysis process. In dynamic environments, those entities that affect much of the activity tend to change fairly frequently, so that updating and revisions play a salient role in reasoning and problem-solving activities. The distribution of the interaction of these environment types is summarized in Figure 1.

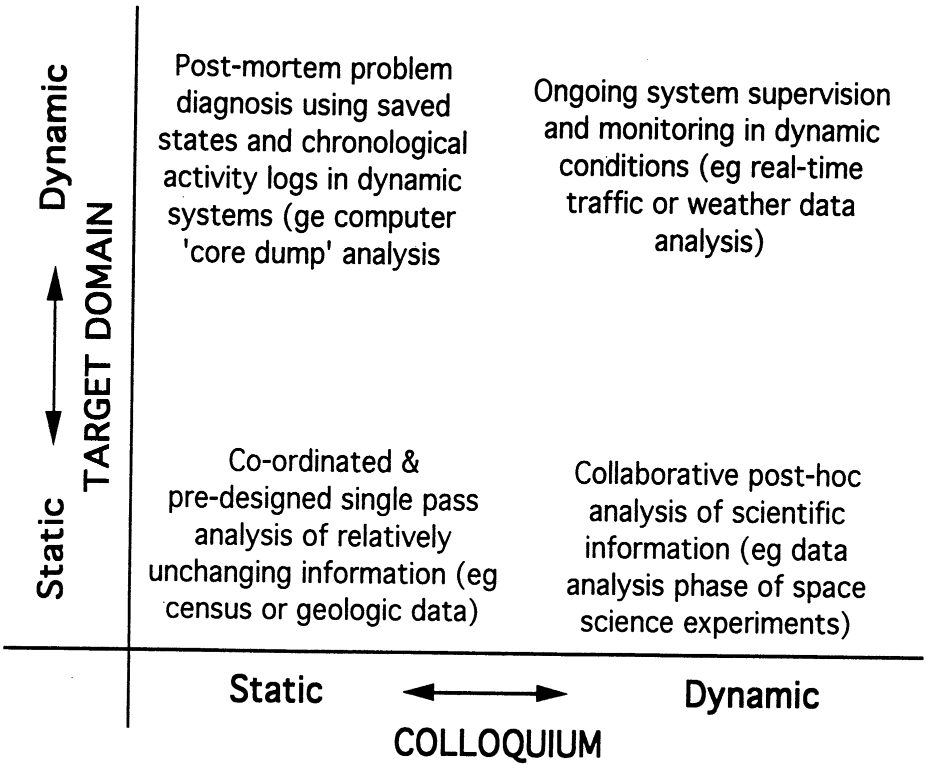

Figure 1 - Summary of Environment Type Interactions

The distinction between target domain and colloquium tends to be defined by the sharpness of the boundaries that delineate an intelligent agent's focus of attention. Thus, when an agent draws conclusions about what it observes in the real world (via its own sensors, access to databases, and so on), it is essentially working in the target domain. On the other hand, when it addresses issues concerning messages between agents, states of belief, etc. the agent is making inferences about the what's going on in the colloquium.

In contrast to this, the distinction between static and dynamic environments is more strongly accentuated by the type of real-world application or problem space being addressed. A static environment is one where most of the ontological entities have attributes that do not change substantially during the inference process; they may perhaps start out as unknown but once asserted they are not undergoing constant rework and revision. The distinguishing feature of a highly dynamic environment is the persistent need to deal with situations where time-varying issues and event sequencing have a strong effect upon the reasoning activity.

The following examples discuss the four combinations of environment types - i.e. target domain vs colloquium and static vs dynamic - and help demonstrate some of the basic differences between the environments.

Static Target, Static Colloquium:

A set of autonomous cooperating agents share the analysis of relatively static information in large databases (such as geological or census data), and each agent focuses on its own sub-specialization or unique analysis method. The shared results are compiled into overall conclusions, perhaps by additional consolidation agents, when each specialist agent's work has been completed.

Static Target, Dynamic Colloquium:

A greater degree of overlapping expertise may be found among the participating agents, so that higher levels of knowledge interchange are possible. As before, the target information is likely to be quite static and consists primarily of data whose prior states have little or no impact (of interest to the analysts anyway) upon present conditions. However, in this scenario the agents may publish interim reports of their results and beliefs asynchronously; thus, process management issues, like consistent spatial orientation and tracking of sequences such as "who reported what when", become important considerations.

Dynamic Target, Static Colloquium:

A variety of diagnostic analysis systems typify this category, especially those constructed from intelligent agents specializing in individual sub-systems. A dynamic process in the target domain encounters some error state or other out-of-nominal condition; this causes the relevant data to be saved which eventually becomes the subject of an analysis procedure. The specialist agents use the saved information, and possibly previous history of the target system, to reach individual conclusions that, as before, may be consolidated into an overall verdict. In these conditions, event sequencing and causal reasoning are important in the target domain, but have little importance specifically for knowledge interchange among the agents themselves.

Dynamic Target, Dynamic Colloquium:

The fourth category is comprised of dynamic target domains that are constantly being monitored by various intelligent agents; real-time analysis of weather data or road traffic conditions are examples. Episodic reasoning may become important, and participating agents often update their beliefs based on new information. The latest information can be shared among agents, but there is less likelihood of very close coordination, since each is primarily addressing its own goals. However, each other's methods and modes of operation often need to be taken into account, in particular to prevent confirmation biases, groupthink, and circular logic problems.

This paper focuses upon the colloquium and in particular on devising an appropriate ontological model of it. As mentioned before, the assumption is made that it is populated by intelligent agents that share knowledge about the target domain. This implies that the ontology of the colloquium should be shared by the participating agents, and that it will be able to accommodate the concepts associated with on-going changes and revisions of beliefs that occur in a dynamic inference and problem-solving process. The theoretical foundations for the architecture of the colloquium are the topic of the next section.

THEORETICAL BASIS FOR THE COLLOQUIUM

The notion of on-going change is endemic in a dynamic target domain; as we have noted, this means that the job of the human operator is more focused on monitoring the self-correcting and automated intelligence systems, tracking system changes and problem-solving activities, and handling the unanticipated and relatively unusual fail states and out-of-nominal situations. Hence, it is desirable to design a model that specifically incorporates an operating agent's behavior in addressing abnormalities. We are thus mandated to seek an architecture for the colloquium that helps integrate some important human qualities like coping, resilience to breakdown, etc. with the primary attributes of sophisticated intelligent systems.

Winograd and Flores [1987] have discussed the limitations of using a simple rationalistic approach to modeling information systems and cognition. They note that detached situational assessment is impossible in everyday activity, since the moment-to-moment realities around us have a persistent influence on our actions. Various scholars of cognition from the rationalistic tradition have developed models of the human mind which rely on an assumption of on-going conscious and deliberate thought. However, philosopher Martin Heidegger differs from many of them in designing his theories around a 'coping' strategy which determines everyday human activity [Heidegger 1962]. In his view, entities in the world don't come to a person's attention until they become unfit for their purpose. People spend much of their time absorbed in trying to achieve some goal, taking in their stride any intervening obstacles and overcoming them as they arise. The tools and procedures for reaching the objective are cognitively invisible until they need repair or adaptation in some way.

The problem with designing models of many complex environments is that there is no single and complete set of constraint symbols which determine the operation of the overall system. Indeed, it often may not be very clear how to derive such a set from the (frequently ill-formed) goal statements for the system. As Dreyfus has noted, formal symbolic logic, although unambiguous and universal, is still limited in its ability to express day-to-day meaning and events [Dreyfus 1993]. One of the outcomes of this problem is the rapid disintegration of an expert system's performance that occurs when real-world conditions move it out of its area of competence. Enhanced knowledge-sharing among intelligent entities can help address this type of difficulty. Heidegger's concepts help underpin our approach by providing a useful paradigm for an agent's behavior in focusing on the unexpected and on the breakdown of routine activity in the target system.

Hence, we seek to characterize the colloquium in terms of the participation of a set of intelligent entities, each one of which is - in Nilsson's words [1995] - a habile system, meaning that its architecture incorporates commonsense reasoning (typically the prerogative of the human operator) and reflexive reaction to immediately-identifiable situations (such as automated knowledge-based expert system behavior). Minsky [1985] uses the term 'agency' to describe clusters of agents participating in a reasoning activity together - in our architecture an analogous definition is adopted and, like Minsky, the terms intelligent agent, agency, and participant are used interchangeably depending upon the context. Each agent in the colloquium needs to know about the existence of other participants; thus, the primary candidate for inclusion in the shared ontology is the concept of an intelligent agent itself.

A basic feature of the architecture is the interchange of information and coordination of activity among the participants. While the subjective or propositional content of such interchanges is the activity and events in the target domain, it is precisely the shared protocols, underlying assumptions, and abstracted formats involved in the conversations that are the purview of the colloquium. Agents are inter-linked by some communications infrastructure which affords certain 'networks of conversations', as Winograd and Flores use the term. They provide the guidance on how 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. Our shared ontology should also incorporate some simple features of computer communications protocols [Day & Zimmermann 1983]; these permit it to differentiate between concepts like one-to-one conversations vs general broadcast messages.

The focus of attention in the colloquium is the internal message or report of a target domain incident or event. This is directly in line with Heidegger's identification of malfunctions or 'unavailable-ness' being the driving force behind everyday behavior. While an individual operating agent may be absorbed with monitoring or managing its own part of the target domain, the basic reason for communication among operating agents is a condition in the target domain that is perceived to be out-of-nominal or otherwise worthy of attention. Thus, it is desirable that our ontology recognize a format for sharing such information.

The concept of communication by way of trouble tickets or error reports is common in general engineering practice, particularly in software engineering because of its adaptability and inherent potential for modification. Processes for characterizing change requests and version control processes are quite well developed and thus we can draw upon this work to provide the inter-agent message form. It should be noted that the report itself is distinct from any representation of the event itself - the latter is an entity in the target domain. The specific identification of report documents as the control vehicle of primary focus has been discussed elsewhere, both for electronic hardware problem diagnosis using an intelligent reasoning environment [Hahn 1991] and for coordinating the activities of multiple (human) agents upon a target domain [Murray 1988].

In expressing a report about the target domain, an agent is communicating its belief about some phenomenon or state of affairs there. Our ontology should therefore accommodate the possibility that reports address either an apparently new condition in the target domain, or provide additional information about a previously-reported condition. Furthermore, formal shared reports often include some means of status classification to indicate state changes between various nominal and exception conditions. It is desirable therefore to include an ontological commitment to the types of state change which can occur in a dynamic environment, since our coping-focused model for the colloquium is centered predominantly upon dealing with exception conditions. Once again, it is important to draw a distinction between the actual or plausible set of states for a given target domain entity and the *types* of state change, as set out in our ontology, which can affect the inference activity.

Some basic spatio-temporal frames of reference need to be available in any distributed reasoning and analysis of a dynamic environment. A specific set of ontological commitments is necessary to promote coordination and agreement, not simply upon the target domain activity, but also on the relationship between such events and the corresponding belief revision activity within the colloquium itself. In particular, it is insufficient merely to commit to an agreed spatial structure or other inter-linked layout of objects in the world. This only reflects what is sometimes referred to 'public spatiality' - that arena where entities show up for people. Which in our environment is of course no more than the generally-agreed topology of the target domain.

Overlaid upon the target domain topology is a set of location and orientation primitives that are the concern of the problem-solving process, and are inferred from attributes of the focus of attention, i.e. whatever an agent is currently dealing with. Dreyfus uses the term 'pragmatic perspective' for this concept, and is careful to point out that it is not an individual or private concept, but rather is accessible to anyone adopting the same focus of attention. When a participant makes statements like 'X is further away than Y', it is routinely understood that the relative distance and orientation is being determined by whatever the agent is currently attending to. This is not necessarily synonymous with the individual agent's physical location; rather, it is established by the context in which the agent is acting.

In a dynamic environment, where entities are frequently characterized by their attributes of motion or topological directedness, that context is strongly influenced by the actual or implied movements in the domain. Our ontology therefore should address the basic axioms and relationships implied by these spatial primitives. The dynamism of the environment also suggests the likelihood of change over time, and a parallel justification may be made for accommodating an agent's perspective in a temporal sense. The architecture of the colloquium thus should also incorporate a set of elements establishing the implications of relative positions in the time dimension. It may be suggested that spatial and temporal identifiers are more properly characteristics of the target domain. However, we note that the need to pay attention to a concept like 'before' or 'beyond' and its implications only arises in the context of an event or incident in the target domain which is worthy of note. Therefore, these abstract primitives, and the various axioms they carry with them, are more like mental 'tools' for the intelligent agent, and thus they reside in the colloquium.

It was noted earlier that a report may be expressed concerning an apparently-new incident or phenomenon in the target domain, or alternatively it may serve to revise or update an agent's belief about a previously-notified event. This provides the participating agents with a means of tracking the evolution of the event. However, a further set of requirements is needed in order to provide a more complete environment for inference and analysis - that is, a method of linking and de-linking events in a causative sense. The physical capability for event and symptom propagation resides in the target domain of course, and some aspects of it may therefore appear in models of the target domain or its common ontology. However, it is the specialized knowledge of individual agents which encapsulate the notions of cause and effect over space and time. Consequently, the abstract rudiments of event propagation need to be included in the shared ontology of the colloquium. This permits individual participants to utilize shared information in developing various types of episodes and scenarios, but without necessarily imposing on all participants some blanket definitions of what might constitute those concepts.

Although the focus of this discussion has been on the characteristics of the colloquium, it should be noted that practical reasoning seamlessly encompasses artifacts from the ontologies of both target domain and the colloquium. The coping strategy adopted by the agents participating in the colloquium is constantly being influenced by on-going events in both environments. But if that strategy specifically indulges in some form of introspection and starts addressing the distinction between the two environments, then a fundamental context switch has occurred - from reasoning about the target to meta-reasoning about the colloquium itself.

Therefore, the architecture imposes no requirements concerning the internal structures or processes of individual participating agents. In particular, it avoids articulating a set of intentional states as a basis for the expressing of reports or messages. (Since the underlying metaphor is that of cooperating participants and reliable communications, there is no prima facie necessity for agents to model each other's behavior.) However, the colloquium architecture incorporates two elements that concern how agents present information in reports, both of which emanate from a phenomenological approach rather than an intentional stance - they are data origination and incompleteness.

Agents in the colloquium typically have access to target sensing systems and/or data archiving facilities which are directly under their control and which provide data primarily intended for the agent's internal use. The architecture incorporates a recognition that, in passing on this data or beliefs derived from it, an agent is essentially introducing or 'originating' this piece of information into the problem-solving process. This is fundamentally different from a participant reiterating or passing on assertions which it received from other agents. By focusing upon the information itself, rather than the physical source of it, we avoid the need to incorporate additional entities such as sensing systems into the architecture; this makes for a clearer ontological boundary between the two.

The colloquium model recognizes that information presented in shared reports may not be complete - in other words, that a full set of attributes needed to uniquely locate and identify a target domain event may not be present. The ontology provides a bounded set of practical circumstances in which an item of data is omitted or left unstated in an agent's report. The use of shared definitions for these cases permits participants to make better inferences about data interchanged among them.

In summary therefore, the colloquium may be thought of as a network of agents, each of which has individual access to certain information about the target domain, but which depends on interchanges with other agents to develop a fuller picture of the target domain activity. The on-going communication in the colloquium uses new and updating reports published by the agents; it relies on agreements among them about some fundamental communication protocols. The problem-solving and analysis process is facilitated when the participants commit to common understandings about the meanings of condition change indicators, basic spatio-temporal primitives, and event propagation characteristics.

Some of the practical limitations of distributed reasoning models are mitigated by agreements on sourcing of evidence, and on the implications of incomplete information. We should reiterate that, in the spirit of Winograd and Flores, our architecture seeks to provide some accommodation for reasoning outside the rationalistic tradition. The colloquium ontology helps participating agents to share information about activity in the target system, without requiring them to commit to completion and closure. It permits them to better make sense of how phenomena fit together, without relying on a centralized definition of 'how the target domain works'. And it opens opportunities for agents using substantially different inference processes to cooperate together in a more meaningful fashion.

Figure 2 sets out in diagrammatic form this set of generalized features which characterize the colloquium. The next section describes each of them in more detail.

COLLOQUIUM ENTITIES

Participants:

The colloquium is occupied by agents or agencies which control access to the sensing data, share the distributed knowledge, and do the intelligent work and inference procedures. They "participate" in the analysis and reasoning process as the members of the colloquium. Agreement on the identity or types of participating agents is particularly useful in environments where relatively few agents are cooperating in overseeing a fairly complex dynamic domain. In the static ontology example discussed by Gruber [1992], this may be analogous to associating the identity of the bibliography publisher (as distinct from the publisher cited in the bibliography entry) with each reference. While this may have some merit in certain circumstances, it is appreciably less necessary than in the case of a dynamic domain, where arbitration between on-going assertions and belief revisions by agents become more necessary.

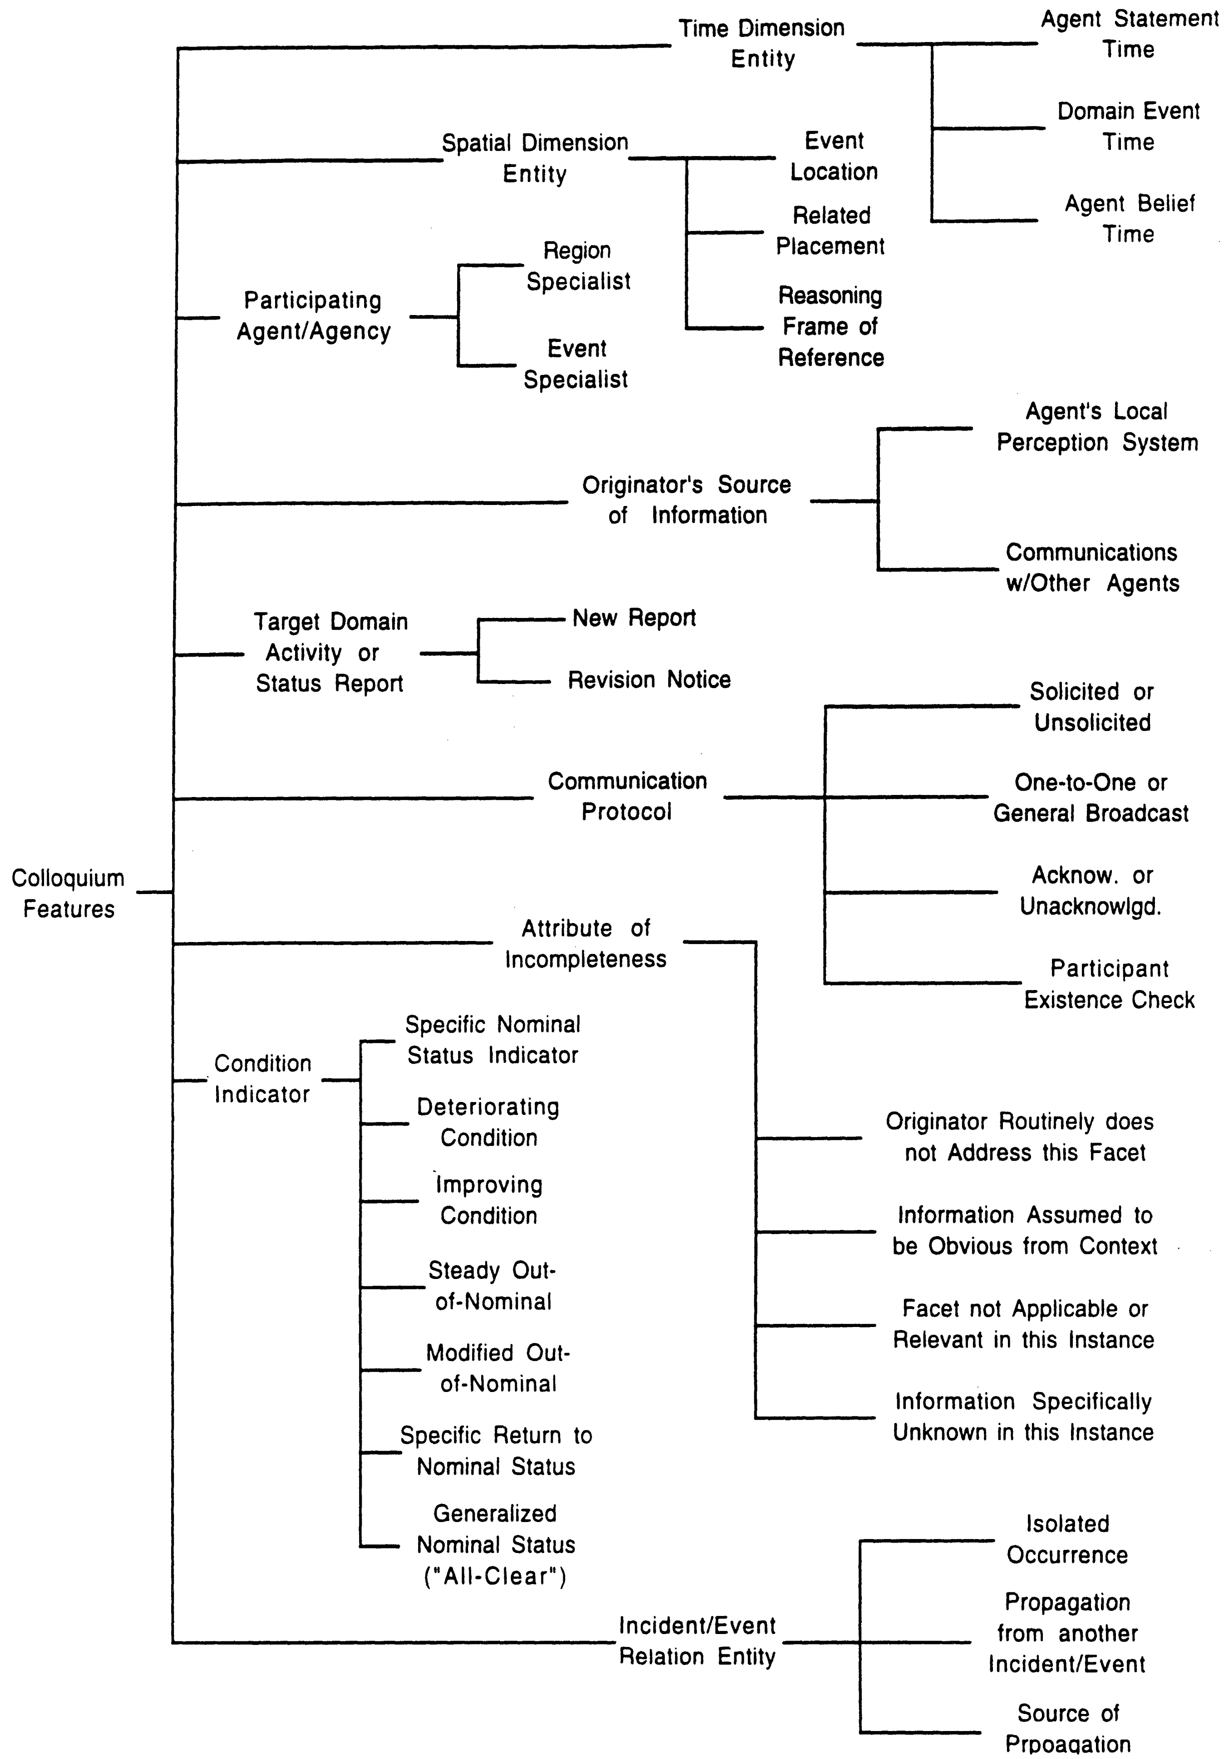
Colloquium Features
Time Dimension Entity
Agent Statement Time
Domain Event Time
Agent Belief Time
Spatial Dimension Entity
Event Location
Related Placement
Reasoning Frame of Reference
Participating Agent/Agency
Region Specialist
Event Specialist
Originator's Source of Information
Agent's Local Perception System
Communications w/Other Agents
Target Domain Activity or Status Report
New Report
Revision Notice
Communication Protocol
Solicited or Unsolicited
One-to-One or General Broadcast
Acknow. or Unacknowlgd.
Participant Existence Check
Attribute of Incompleteness
Originator Routinely does not Address this Facet
Information Assumed to be Obvious from Context
Facet not Applicable or Relevant in this Instance
Information Specifically Unknown in this Instance
Condition Indicator
Specific Nominal Status Indicator
Deteriorating Condition
Improving Condition
Steady Out-of-Nominal
Modified Out-of-Nominal
Specific Return to Nominal Status
Generalized Nominal Status ("All-Clear")
Incident/Event Relation Entity
Isolated Occurrence
Propagation from another Incident/Event
Source of Prpoagation

Agents may be specialists of some kind. For example, one may specialize particularly in information about a given region of the target domain, or it may focus on specific types of events or activity in the target domain. Specialization may also be delimited on a temporal scale, such as when the duties of one agent are subsumed by another during certain periods such as nighttime or weekends. The existence of such specializations is of importance to other agents in the colloquium, since it provides a more clear-cut means to seek verification of suspected conditions in the target domain. This is particularly important in dynamic environments, because of the on-going need to deal with the 'situation-at-hand'. It is desirable therefore that sharing an ontology incorporates a commitment to the notion of a specialization link between an agent and a target domain entity.

Messages & Reports:

The colloquium includes some support for sharing information about the target domain among two or more participating agents. This support may be comprised of a combination of database entries, frames, inter-agent messages of some type, etc. It is desirable to incorporate some ontological commitment to basic categorization of these knowledge-sharing vehicles, which are referred to here as reports. For example, a report may refer to a "new" condition in the target domain - i.e. one that is new in the belief of the issuing agent. Alternatively, the report may be a specific clarification or update of some type to an episode which has been previously reported upon. A further possibility is the case where the reports may not refer to each other directly, but implicitly make reference by using the same domain objects in multiple reports. A shared ontology should provide agents with the opportunity to agree upon these types of concepts.

Reports are typically comprised of a group of references to entities in the target domain. Although they are application-specific, these references can also be characterized in terms of one or other of the colloquium entities. For example, a report in a diagnostic system may include information about condition changes in the target domain; generic characteristics of such changes are described elsewhere in this section. Committing to a shared ontology of simple, condition-changing raises the abstraction level at which agents can communicate and reason.

Communications Protocol:

There are a number of facets of inter-agent communication which can affect inference processes in the colloquium. Firstly, a participating agent may routinely initiate reports of its own volition, it may only issue a report when requested to do so, or it may operate in a way that combines both policies. In the second place, it may publish reports only in relation to exception conditions in the target domain, it may explicitly report that conditions are believed to be normal, or some combination of both. Thirdly, some communications may occur simply between a limited group of two or more agents, while others are broadcast or otherwise made available to any interested party. Finally, some agents may communicate using the assumption that reports they issue will be acknowledged by recipients, while others have no such expectation.

The ontological model is not intended to incorporate a rigorous presentation of system reliability and tolerance of faults within the colloquium itself, consigning most of those issues to the design of the underlying infrastructure instead. However, the shared ontology can benefit greatly by making one exception to this exclusionary rule on meta-reasoning. The distributed nature of the colloquium, when taken in combination with the hortatory character of the environment, means that a participating agent's presence may not necessarily be evident to others in the sensor data they receive from the target domain. Therefore, the communications protocol needs to include a representation of an existence check, whereby one agent can verify that another is actually present and active. In many practical situations, this purpose may be served by some part of the protocol described earlier. However, the inclusion of basic presence-checking as an explicit ontological concept enables better abstraction in an agent's reasoning.

Condition Indication:

In many environments, the measurements received from the target domain have some nominal structure or behavior pattern which can be handled by some well-defined background process or self-regulating system. The real work in the colloquium only gets going in earnest when abnormal patterns or exception conditions are encountered. This necessitates ontological support for the various types of condition changes that correspond to broad classification expressions in English like deteriorating condition, improving condition, reverted to nominal, etc. Typically, the reports published in the colloquium by the participating agents either implicitly or explicitly specify one of the condition indicators, the full list of which is shown in the following table:

Condition Indicator Type	Description
Specific Nominal	A single target domain node or region of nodes is believed to be in a nominal or non-problem state.
Deteriorating	A node or region has changed to a state that is more problematic or serious than the earlier state, and thus is likely to merit more attention.
Improving	A node or region has changed to a state that is less problematic or serious than the prior state.
Steady Out-of-Nominal	A node or region continues to remain in an out-of-nominal state.
Modified Out-of-Nominal	A node or region has changed its state from one out-of-nominal one to another, with no apparent overall deterioration or improvement.
Reverted-to-Nominal	A previously problematic node or region is believed to have reverted back to a nominal or non-problem state.
Generic Nominal (All-Clear)	All nodes and regions in the target domain are believed to be in nominal or non-problem states.

It is important not to over-specify the implications of a condition indication entity by including 'value' judgments in the shared ontology; for example, participating agents should not be expected to commit to a shared priority scale. The ontology is not designed to make agents agree on the criticality level of a specific deteriorating condition in the target domain, or even that it *is* deteriorating at all. Rather, the agreement is on what the concept of deteriorating implies - i.e. that more attention and/or resources are probably needed, that adverse effects on other entities may occur (if only because attention means lack of resources available elsewhere), etc.

Certain types of colloquium and communications protocols may use an 'all-clear' form of report category, indicating an agent's belief that no reportable episodes currently exist. In other circumstances, a time-out period may apply, after which the target domain condition may be assumed to have reverted to normal. [In the colloquium, this corresponds to an agent holding a belief for a specific duration of time.] The shared ontology should also provide support for characterizing two versions of the out-of-nominal condition, (i) steady out-of-nominal, meaning 'condition unchanged', and (ii) modified out-of-nominal, meaning 'neither better nor worse, but different'.

Temporal Entities:

The colloquium can provide some participants with the ability to publish information concerning a dynamic domain for unspecified and asynchronous access by the general body of agents. In other words, some knowledge may be shared by blackboarding or broadcasting in some form, rather than as responses to specific requests. The time of publishing can become relevant in the overall reasoning process - it may be important to be able to express the notion "At time t1, agent A published a statement that at time t2 agent A [did/does/will] hold belief X about the state of the target domain at time t3".

Therefore, for practical purposes it is postulated that information in the colloquium may have up to three different times related to it. For example, a problem diagnosis report, published at t1 about activity in a dynamic target domain, might note that a chronological log entry made at t2 provides evidence about a target domain event that occurred at t3. This is about the limit of belief indirection that is needed for many practical human reasoning purposes. In common business parlance, it roughly corresponds to the sentiment "She gave a date-for-a-date at the project scheduling meeting on Monday". Authors in the cognitive science and AI literature have pointed out that most people encounter difficulty dealing with higher levels of nesting, such as "A believes at t1 that B believes at t2 that C believes at t3 that D believes at...etc".

It is interesting to observe that scenarios can be devised where the three time concepts occur in different real-time sequences; this may be done by mixing various types of predictions and post-hoc analyses. Also, it seems likely that, given a sufficiently large timegrain, many environment activities may be represented adequately using a default axiom that the report is valid now [i.e. all three times are equal to 'now'] unless some other time is specifically stated or is declared to be unknown.

Spatial Entities:

A sense of placement is an important facet of most colloquium activity because of the need to address the effects of behavior propagation among entities in a dynamic target domain. Furthermore, positional commitments should be included in a shared ontology to facilitate the resolution of everyday prepositions like "beyond", "downstream from", "before", and so on. Positional information is predicated upon some agreed or implied frame of reference, and this 2D or 3D frame of reference may not be the same for all participants or from the context of one message to the next. Thus, intelligent reasoning may require modification of the basis reference frame. It should be noted that spatial entities in the target domain do not necessarily have to refer to 2D or 3D physical space. Positional statements can be made about locations where motion is implied - "beyond" a switch or fuse in an electrical circuit for example - or in hierarchic structures that are conceptual, such as the use of "above" and "below" in administrative organizations.

Reconciling interactions between reported events in the target domain may require merging of information relying on differing reference frames. As with the temporal elements, up to three spatial entities may typically be associated with a message or report. For example, consider a statement of the form "The region of interest P1 is located beyond position P2 [when viewed from location P3]". In this case, P1 can be thought of as the location of the subject of the message while P2 is essentially the domain reference point. In a dynamic domain, P3 frequently emerges from the context of the message rather than being explicitly stated. It may become evident because of assumed motion in the domain - in ATMS for example, a positional statement like "beyond exit 15 on I-80 westbound" implies "to the west of". In aerospace applications, the problem arises in three dimensions, and there is a need to resolve expressions using above, forward of, etc. from differing frames of reference.

A dynamic reasoning ontology therefore should support axioms such as "If the region has an attribute of 'westbound', then 'beyond' means 'to the west of'". It is proposed to incorporate a fundamental set of shared commitments to positional primitives, such as those defined for Basic English [Ogden 1932]. As with the temporal commitments, the goal is to facilitate practical work in the colloquium rather attempt a total codification of the nature of space-time.

Propagation Entities:

Error propagation can occur in the target domain when one abnormal state or event causes a second one to happen, the latter being either temporally or spatially separated from the first, or both. It is also conceivable that the second state or event actually occurs independently of the first, but that its effect or outcome is altered as a result of the first. A further possibility is that incomplete recovery or clean-up procedures in the target domain resulted in later failing sequences. In some environments, it may be desirable to identify sets of linked events and characterize them as episodes.

In the colloquium, knowledge is shared among the participants using reports. The incident or event addressed by a report can be characterized as either the source of a propagation effect or the result of one, or both. Alternatively, the incident or event may be regarded as having no salient propagation effect and therefore the report addresses an isolated occurrence. Propagation is therefore a report-linking process in the colloquium. An agent may reason internally upon the possibility of propagation links between reports. It may include in a report some propagation information that links the report with other reports. Or the agent may perhaps publish an additional report specifically to establish propagation links between some other reports.

It is important for the purposes of the shared ontology to differentiate carefully between the definitions of propagation in the target domain and in the colloquium. In the target domain, the dynamics of the environment may result in an effect propagating among entities there. For example, a valve failure in a drainage pipe may cause an overflow and consequent flooding further back in the system. In the colloquium, a propagation attribute would only be asserted if the flooding event had been independently identified, had become the topic of a separate report, and some inference process was required to link the two events. Otherwise, every state change and belief revision becomes characterized as a propagation from the previous state and the explicative power of propagation becomes lost.

Propagation also has one or more types associated with it. For example, a failure in an electromechanical system might have electrical and thermal effects as well as mechanical ones. In the colloquium therefore, it may be necessary to characterize the propagation links among reports by type. It is also plausible that several threads of different types might link reports in differing sequences.

Information Origination:

The fact that agents are cooperating in their reasoning can often lead to a tacit acceptance concerning the origin of information about dynamic domains; this is the issue of initiation of knowledge versus hearsay. An agent can be thought of as *originating* an item of knowledge - either by publishing it or providing it in response to a specific request - when the item has been identified directly by the agent using its own sensing system. On the other hand, the agent may be involved in hearsay if the reported item is primarily being inferred using similar or related knowledge received from other participating agents.

The level of originality may be a significant attribute of the shared knowledge item; it is desirable to expose this differentiation explicitly by incorporating an assessment of the knowledge origination within the shared ontology. Naturally, the intent here is to assist with the practical sharing of information, not to provide a complete and rigorous solution to the AI problems of common and distributed knowledge. The provision of an ontological entity defining knowledge origination helps system designers accommodate direct versus indirect knowledge, and thus discourages the circular logic and reasoning inaccuracies that arise from cognitive biases like repetition and groupthink. It is recognized that an ontology of cognitive biases could possibly be developed from a comprehensive list [e.g. Sage 1991]. However, although it might be valuable for use in some colloquium activities, such an ontology specifically is not considered architecturally part of the colloquium.

Incompleteness Attributes:

It was stated earlier that the intent of a shared ontology is to harmonize underlying agreements among participating agents about target domain entities. One accepts of course that participating agents are not committed to support the ability to interpret and reason upon all facets of the shared concepts in the target environment. However, in the colloquium, this has some implications for knowledge incompleteness. The unstated/unknown issue is particularly relevant to a dynamic environment because of the need to identify some default course of action when incomplete information is encountered.

Various researchers [e.g. Cheeseman, Kelly, Self, Stutz, Taylor & Freeman 1988; Gennari, Langley & Fisher 1989], while noting that this is a rich topic for further investigation, have also pointed out that reasoning systems that simply ignore a fact or interpolate a value because of missing data may destroy potentially valuable knowledge. It is preferable to articulate the potential situations implied by the absence of information. For example, in the colloquium architecture, several different nuances may be suggested by an item of knowledge which is left unstated in a message or report. An agent may leave something unstated because (i) the agent has no information in this particular instance, (ii) it routinely does not address this facet of the shared ontology, (iii) it assumes that the recipient will be able to infer the facet from context, or (iv) the facet is fundamentally inapplicable or not relevant in this case, (which also should be inferable from the context). Ontological support for these varying interpretations of incompleteness is necessary in the colloquium.

TRAFFIC CONGESTION DOMAIN - AN EXAMPLE ONTOLOGY

We now discuss some of the practical aspects of implementing and using shared ontologies for the target domain and the colloquium. The selected operating environment is road traffic management, specifically the sharing of information on freeway conditions among traffic agencies in south-east Michigan. The development environment chosen for constructing the colloquium and target domain ontologies is the Ontology Editor [KSL Network Services 1995]. The combination of the two ontologies - plus any other more fundamental ones that they incorporate - constitutes the set of entity specifications to be shared among the participants. This section discusses how the ontological entities would be combined in the construction of a knowledge-based intelligent agent tool. It also provides two examples of how axioms from the colloquium ontology can implement architectural facets like spatial reasoning and incompleteness.

The entity specifications that are to be shared using the colloquium ontology implement the artifacts discussed in the previous section. Figure 3 shows an example of how the overall colloquium ontology hierarchy is structured. The set of classes described earlier are shown with some further sub-classes also detailed. (The titles of colloquium entities in these examples are identified with the prefix Co-.) Definitions of entities directly referring to the real-world environment are shared using the target domain ontology; figure 4 shows portion of the hierarchy for the shared ontology for road traffic. Terminology for many types of traffic infrastructure locations and entities can be included in the example ontology - points of interest, road types, bridges, tunnels, and so on. Classes are also provided to codify the participating agents and the shared information. Besides the concept of a shared message (`Traffic-Condition-Report`), the ontology includes definitions for the principal items of content in those messages. Entities are provided to represent the name or number of the affected freeway(s), the traffic direction, and positional location information.

The colloquium ontology contains generic specifications that do not rely upon specific application definitions. Thus, there will be no inheritance of attributes from any target domain ontology into the colloquium one. On the other hand, the integration of the two ontologies is achieved precisely because of the transfer of attributes in the opposite direction - from generic to specific. For example, a traffic administration organization such as a public transit authority would have an ontological representation as a target domain entity. This entity would have domain-specific attributes concerning service areas, types of transit systems operated, etc. It would also inherit from the colloquium ontology those generic attributes that characterize it as a participating agency. This is similar to the fashion in which many application-specific ontologies draw upon other, more generic ontologies for the inherent assumptions and characteristics of entities such as time units or mathematical relationships.

The traffic ontology makes a distinction between the type of incident or event (`Reported-Event-Type`) and the outcome or impact of such an event (`Effect-Of-Event-On-Traffic`). This reflects the common usage difference between causes like accidents, construction, debris, etc. on the one hand, and effects such as congestion or lane closures on the other. (This part of the design is based upon an analysis of actual written reports on freeway conditions shared among traffic agencies.) Sometimes messages explicitly mention just one of these, and the other is either unknown or assumed to be implicit in the context - an example of the need for some simple colloquium support for incompleteness. The ontology also permits the distinction between the time of an event in the target domain and the time of information availability; this supports reports of the form "Revised bulletins on tomorrow's construction closures will be issued at noon".

Figure 3 - Extract from top level of Colloquium ontology

```
Colloquium-Entity
   Co-Participating-Agency
   Co-Report
   Co-Comms-Protocol
   Co-Condition-Entity
      Co-Status-Change-Type
   Co-Time-Entity
      Co-Report-Issue-Time
      Co-Agent-Belief-Time
      Co-Target-Domain-Event-Time
   Co-Spatial-Entity
      Co-Spatial-Point
      Co-Spatial-Region
      Co-Spatial-Direction
         Co-Before
         Co-Between
         Co-Beyond
   Co-Propagation-Entity
   Co-Information-Origin
   Co-Incompletion
```

Figure 4 - Extract from top-level of Road Traffic Domain ontology

```
Traffic-Entity
   Traffic-Agency
      ....
   Traffic-Information-Entity
      Traffic-Condition-Report
      Roadway-Name/Number-Reported
      Roadway-Direction-Reported
      Traffic-Event-Locator
      Reported-Event-Type
      Effect-Of-Event-On-Traffic
      Condition-Change-State
      Time-Agent-Has-Information
      Time-Of-Traffic-Event
      ....
Domain-Place-Name-Entity
   Domain-Freeway/Street-Name
   Domain-Business-Location
   Domain-Entertainment-Facility
   .....
Roadway-Type-Entity
   Surface-Street-Type
   Freeway-System-Road-Entity
      Frontage/Service-Road
      Distributor/Collector
      Express-Lanes/Local-Lanes
      On/Off-Ramp
```

Figure 5 - Colloquium Report frame definition.

```
(define-frame Co-Report
   :own-slots
   ((Documentation
          "Co-Report is a class used to represent
          statements made by colloquium agents
          about the state of the target domain.")
    (Instance-Of Class) (Subclass-Of Co-Entity))

   :template-slots
   ((Co-Originating-Agent
          (Slot-Value-Type Co-Participating-Agency))
    (Co-Report-Issue-Time
          (Slot-Value-Type Co-Time-Entity))))
```

Figure 5 shows a small part of the colloquium entry for the report entity. In the shared ontologies for traffic, the entity (`Traffic-Condition-Report`) inherits these generic characteristics - by virtue of its being a type of report that is used in the colloquium (`Co-Report`) - as well as certain application-specific features, since it is also defined as a type of (`Traffic-Entity`) in the target domain. Thus, a builder of an intelligent system who designs a traffic report entity using the application-specific ontology also inherits representations for the issuing agent and the time at which it was issued, i.e. (`Co-Originating-Agent`) and (`Co-Report-Issue-Time`).

Similarly, the two target domain time entities discussed earlier - (`Time-Agent-Has-Information`) and (`Time-Of-Traffic-Event`) - automatically incorporate the characteristics of the colloquium generics (`Co-Agent-Belief-Time`) and (`Co-Target-Domain-Event-Time`). By making a commitment to the existence of these temporal distinctions, the system builder enables a greater sharing of information among quite diverse participants, even though his or her individual system may not require such detail internally.

The colloquium ontology provides for conceptual spatial entities to be represented as individual points (`Co-Spatial-Point`), or as regions (`Co-Spatial-Region`) which may be groups of points or groups of lower-level regions. There is no commitment to any particular dimensions or hierarchical structure - these are very much in the preserve of the relevant target domain ontology. As mentioned earlier, no particular commitment is made to represent actual 2D or 3D physical space; the target domain could be an organizational structure for example.

The concept of flow or motion through the system is accommodated by the ontological entity (`Co-Spatial-Direction`) which is the class used to characterize routing attributes like downstream, westbound, backwards, etc. Spatial direction is applied as an attribute of spatial region entities, and for each instantiation a pair of 'conceptual end-points' are defined; these are (`Co-Origin`) and (`Co-Goal`). They represent the abstract sources and destinations of the flow or motion, and are used to support qualitative reasoning about position in the target network. Figures 6 and 7 present portions of the definition of these three entities, and the following discussion shows how the ontological axioms can be applied in practice.

Figure 6 - Extracts from Colloquium ontology, showing part of the Spatial Direction entry and the positional axioms.

```
(define-frame Co-Spatial-Direction
   :own-slots
   ((Documentation
         "This class is used to identify 'classical'
         directions of common-sense notions of linear
         movement.  Examples of these are expressions
         like downstream, eastbound, upwards, outbound, etc.
         Each instantiation of a spatial direction include
         two generic spatial point definitions - the Origin
         and the Goal - which are used in the accompanying
         axiomatic definitions of positional relationships.")
    (Instance-Of Class) (Subclass-Of Co-Spatial-Entity))

   :template-slots
   ((Co-Goal (Slot-Value-Type Co-Spatial-Point))
    (Co-Origin (Slot-Value-Type Co-Spatial-Point)))

   :axioms
    ((=> (Co-Before ?locn ?anchor)
        (and
            (< (Distance ?anchor ?co-goal)
               (Distance ?locn ?co-goal))
            (> (Distance ?anchor ?co-origin)
               (Distance ?locn ?co-origin))))
    (=> (Co-Beyond ?locn ?anchor)
        (and
            (> (Distance ?anchor ?co-goal)
               (Distance ?locn ?co-goal))
            (< (Distance ?anchor ?co-origin)
               (Distance ?locn ?co-origin))))
    (=> (Co-Between ?locn ?anchor1 ?anchor2)
        (or
           (and (Co-Before ?anchor1 ?locn)
                (Co-Beyond ?anchor2 ?locn))
           (and (Co-Before ?locn ?anchor1)
                (Co-Beyond ?locn ?anchor2))))))
```

Figure 7 - Extracts from Colloquium ontology, showing the Goal & Origin relations.

```
(define-relation Co-Origin
   (?frame ?value)
       "This is the slot which is used as the starting point
       in Co-Spatial-Direction from which the directional
       vector originates."

   :def
       (and (Co-Spatial-Direction ?frame)
            (Co-Spatial-Point ?value)))

(define-relation Co-Goal
   (?frame ?value)
       "The is the ultimate destination slot for use
       in Co-Spatial-Direction - it is the final goal
       point for vectors in spatially directed regions."

   :def
       (and (Co-Spatial-Direction ?frame)
            (Co-Spatial-Point ?value)))
```

A target domain entity representing a single section of freeway might be defined to have the characteristics of the colloquium entity (`Co-Spatial-Region`) linked to the spatial direction attribute (`Co-Northbound`). That attribute would be instantiated using abstract extremities like "North-Pole" as destination and "South-Pole" as origin; these points essentially define locations well beyond the edges of the territory covered by the target domain. The extremities provide a means of doing qualitative reasoning upon natural language positional locator statements like "Interstate 75 northbound beyond Exit 10" by establishing a generic outer edge. Figure 6 includes spatial direction axioms for the prepositional expressions Before and Beyond; having established them, the expression Between is defined in terms of the other two.

The use of an ontology entry for incompletion permits a level of agreement among agent designers, which would otherwise be hard to coordinate, about the meaning of data left unstated in reports. Various processes, such as the use of default values, interpolation or extrapolation from other data, methods for seeking out information, and so on can be grouped in this ontological class. The entry for (`Co-Incompletion`) which is shown in Figure 8 includes a sample axiom that covers two of the incompleteness attributes discussed in an earlier section. The English interpretation of the axiom is as follows: "If a report, which has been issued by a particular agent, contains a named slot which is empty, then either (i) that slot is to be treated as undefined in this report, or (ii) the corresponding slot in all reports originated by this particular agent are assumed to be empty".

Figure 8 - Extract from Incompletion entry showing Empty/Undefined axiom.

```
(define-frame Co-Incompletion
   :own-slots
   ((Documentation
          "This class provides a means to group commonsense
          ideas about colloquium entities which are
          incomplete or were left unstated for some reason.
          These notions help support definitions of the
          'appropriate' or 'most likely' defaults to use.")
    (Instance-Of Class) (Subclass-Of Co-Entity))

   :axioms
   ((=>
       (and
          (Has-Slot-Value@Ol-User%Slot-Constraint-Sugar
              Co-Report
              Co-Originating-Agent
              ?originator)
          (Has-Slot-Value@Ol-User%Slot-Constraint-Sugar
              Co-Report
              ?slotname
              ?slotvalue)
          (Empty ?slotvalue))
       (or
          (Undefined ?slotvalue)
          (forall
             (?report)
             (and
                (Has-Slot-Value@Ol-User%Slot-Constraint-Sugar
                    ?report
                    Co-Originating-Agent
                    ?originator)
                (Has-Slot-Value@Ol-User%Slot-Constraint-Sugar
                    ?report
                    ?slotname
                    (Empty))))))))
```

This example shows how a simple level of agreement can be achieved among independent agents about a colloquium entity without requiring explicit accommodations to be made by them. This provides the developers of intelligent systems relying upon shared knowledge with a greater degree of flexibility in their designs than was previously available. In the traffic engineering field, the search for designs which aid in sharing information among various organizations has been noted by several researchers. For example, Richie [1990] and Gilmore et al [1994] both describe architectures that envision separate intelligent entities sharing information using a distributed blackboard model. Although a design may permit the various participating agents to execute independently of each other, the blackboard architecture relies heavily upon very close coordination and high levels of communication among all agents - this may be difficult to achieve in practice.

Another limitation of such designs is that all relevant information must be published on the blackboard and direct inter-agent conversations are forbidden. This could well pose an implementation problem since individual organizational goals tend to determine the primary context of intelligent systems usage. For example, a police agency may merge its traffic-related information with data on its other activities - data which could well be inappropriate to publish - and separating such items could prove structurally very difficult to accomplish. On the other hand, obtaining commitments to a shared set of ontologies is likely to be less troublesome and permits the participating agents greater flexibility and autonomy in how they share information. In this respect, it is maintained that a focus upon agreed ontological supports more accurately reflects real world processes.

CONCLUSIONS

The use of automated intelligent assistants in support of human operators is becoming more widespread, particularly when an organization undertakes the integration of multiple sources of data on its activities. The difficulties encountered with such systems can become exacerbated when several organizations try to gain synergy by linking up their knowledge-based systems. However, such problems are ameliorated if the participants can commit to shared ontologies that pertain to their domains of responsibility. The process of developing such ontologies is made somewhat easier when dealing with a hortatory operations environment. This is because a knowledge system designer's conceptualization focus is more upon analysis and interpretation processes at a higher level of abstraction, rather than having to fit more closely into a direct feedback/control loop structure, which might be the case in a more traditional supervisory operations model.

It should be noted that, by their very nature, distributed operations environments - and hortatory ones in particular - are likely to be administered with the aid of a loose-knit network of intelligent agents, precisely because of the lack of opportunity for central control. Hence, it is contended that accurate characterizing of a colloquium architecture is particularly necessary, since its existence is directly implied by the loose-knit concept. The use of shared ontologies provides for the broadest degree of implementation variations, while maintaining coordination at the upper conceptual levels where it matters. Thus, organizations can maintain their own sensing systems, databases, network models, etc. and agree precisely on the concepts and objects of mutual concern - viz. the entities in the target domain, and the shared knowledge of the colloquium.

ACKNOWLEDGMENT

The authors would like to acknowledge the support provided for this work by a US Federal Highway Administration grant to the Intelligent Transportation Systems Research Center of Excellence at the University of Michigan.

REFERENCES

BREZILLION P & ABU-HAKIMA S (1995). "Using Knowledge in Its Context: IJCAI 1993 Workshop", AI Magazine, Vol 16/1.

CHEESEMAN P, KELLY J, SELF M, STUTZ J, TAYLOR W & FREEMAN D (1988) "AutoClass: A Baysian Classification System", Proc. of Fifth Intl Conf on Machine Learning, Ann Arbor MI.

DAY J & ZIMMERMANN H (1983). "The OSI Reference Model", Proceedings of the IEEE, Vol 71/12.

DREYFUS H (1993). "Being-in-the-World; A Commentary on Heidegger's Being and Time", MIT Press, Cambridge MA.

GENNARI J, LANGLEY P & FISHER D (1989). "Models of Incremental Concept Formation", Artificial Intelligence, 40/1.

GILMORE J, ELIBIARY K & FORBES H (1994). "Intelligent Control in Traffic Management", Proc .of Vehicle Navigation & Information Systems Conference.

GRUBER T (1992). "A Translation Approach to Portable Ontology Specifications", Technical Report KSL 92-71, Stanford.

GRUBER T (1993). "Towards Principles for the Design of Ontologies Used for Knowledge Sharing", Technical Report KSL 93-04, Stanford.

HAHN M (1991). "Watson: An Expert System for Mainframe Diagnosis", Stanford CS Report 226.

HEIDEGGER M (1962). "Being And Time", Harper Row, New York.

KSL NETWORK SERVICES (1995). Knowledge Systems Laboratory, Stanford University, URL http://www-ksl.stanford.edu.

LIU Y, FULD R, & WICKENS C (1993). "Monitoring system behavior in manual and automated scheduling systems", Int. J. of Man-Machine Studies, Vol 39.

MINSKY M (1985). "The Society of Mind", Simon & Schuster, New York.

MURRAY J (1988). "Problem Reporting And Development Tracking Using VM/SP", Software Engineering Notes, ACM SIGSOFT, Vol 13/4

MURRAY J & LIU Y (n.d). "Towards a Distributed Intelligent Agent Architecture for Human-Machine Systems in Hortatory Operations", IEEE TransSMC, in review.

NIELSEN P (1995). "Intelligent Computer Generated Forces for Command and Control", Proc. 5th Conf. on Computer Generated Forces and Behavioral Representation , Orlando, FL.

NILSSON N (1995). "Eye on the Prize", AI Magazine, Vol 16/2.

OGDEN C K (1932). "The Basic Dictionary", Kegan Paul, London.

RICHIE S (1990). "A Knowledge-based Decision Support Architecture for Advanced Traffic Management", Transportation Research A, Vol 24/1.

SAGE A (1991). "Decision Support Systems Engineering", John Wiley, New York.

SHERIDAN T (1992). "Telerobotics, automation, and human supervisory control", MIT Press, Cambridge MA.

WINOGRAD T & FLORES F (1987). "Understanding Computers and Cognition", Addison Wesley.