**Proceedings of DETC99/CIE:**
**ASME Design Engineering Technical Conferences/ Computers in Engineering**
**Las Vegas, NV September 12-16, 1999**

# DETC99/CIE-9024

# AN ANALYSIS OF EXISTING ONTOLOGICAL SYSTEMS FOR APPLICATIONS IN MANUFACTURING

Craig Schlenoff
National Institute of Standards and Technology
Gaithersburg, MD

Peter Denno
National Institute of Standards and Technology
Gaithersburg, MD

Rob Ivester
National Institute of Standards and Technology
Gaithersburg, MD

Simon Szykman
National Institute of Standards and Technology
Gaithersburg, MD

Don Libes
National Institute of Standards and Technology
Gaithersburg, MD

Keywords: manufacturing, ontology, ontological engineering, semantics

## ABSTRACT

In all types of communication, the ability to share information is often hindered because the meaning of information can be drastically affected by the context in which it is viewed and interpreted. This is especially true in manufacturing because of the growing complexity of manufacturing information and the increasing need to exchange this information among various software applications. Different manufacturing functions may use different terms to mean the exact same concept or use the exact same term to mean very different concepts. Often, the loosely defined natural language definitions associated with the terms contain so much ambiguity that they do not make the differences evident and/or do not provide enough information to resolve the differences.

A solution to this problem is the development of a taxonomy, or ontology, of manufacturing concepts and terms along with their respective formal and unambiguous definitions. This paper focuses on an effort at the National Institute of Standards and Technology to identify, formally define, and structure the semantic concepts intrinsic to the capture and exchange of manufacturing information.

Specifically, this paper documents the results of the first phase of this project – that of analyzing existing ontological systems to determine which is most appropriate for the manufacturing domain. In particular, this phase involved the exploration of efforts that are studying both the uses of ontologies in the general sense and those that are using ontologies for domain-specific purposes.

## 1.0. INTRODUCTION

The objective of this work described in this paper is to move closer to the ultimate goal of seamless manufacturing systems integration using the principle behind ontological engineering to unambiguously define domain-specific concepts. A major challenge facing industry today is the lack of interoperability between heterogeneous systems. Current integration efforts are usually based solely on how information is represented (the syntax) without a description of what the information means (the semantics). With the growing complexity of information and the increasing need to completely and correctly exchange information among different systems, the need for precise and unambiguous capture of the meaning of concepts within a given system is becoming apparent.

The approach for the project described in this paper is to analyze current ontological systems to determine which is most suitable to model the concepts in the manufacturing domain. Examples of ontological systems include CYC[1,2], MikroKOSMOS, and the Ontolingua server. The project will move to formally identifying and modeling concepts (and definitions of those concepts) from various manufacturing domains and projects (e.g., process specification, product modeling, resource representations, etc.) in this ontological

[1] CYC is a registered trademark of Cycorp Inc.

[2] No approval or endorsement of any commercial product in this paper by the National Institute of Standards and technology is intended or implied. This paper was prepared by United States Government employees as part of their official duties and is, therefore, a work of the U.S. Government and not subject to copyright.

system. At this point, an analysis can help to identify inconsistencies in the use of terms among various domains as well as help to establish a means to generalize these terms to a level that is common among the domains in question.

The output of the work documented in this paper will be a taxonomy of terms and concepts along with formal definitions of exactly what each of those terms and concepts mean and how they interrelate. Although it would be impossible to create a complete taxonomy of every interpretation of every term, a high-level, extensible subset of this taxonomy will be created to serve as a basis for future, domain-specific additions and specializations. This shared understanding of concepts could then be used to integrate applications and systems that function towards a common goal.

## 2.0. WHY ONTOLOGIES FOR MANUFACTURING?

This section considers what value the investigated ontologies might provide to the area of information technology within manufacturing. Communication and context are important notions for understanding the role of the investigated ontologies vis a vis other technologies. This and the related concepts of formality, ground, and context availability are discussed. Three areas of potential benefit are then considered: unambiguous communication, standards making and semantic alignment efforts, and the future industrial information infrastructure.

## 2.1. COMMUNCATION, MEANING, AND CONTEXT INTRODUCTION

In this paper 'communication' has the following meaning: One communicates to another with the expectation that at some time thereafter the receiver will produce a behavior that is in some way consistent with the initiator's intention. For communication to succeed, the initiator must have (or the system design must reflect) some understanding of the context under which the receiver is operating and the relationship between the message it designs and the behavior it desires from the receiver.

This definition of communication begs the question of what is 'meaning'. Along the lines of Bloomfield [BLO33] there is at the very least this sense of meaning to communication: it sets conditions for satisfaction. If I say "Pick up the book." and you do so, this is evidence of a relationship between a representation and a perceived reality. Philosophers might argue that there is more or less meaning in 'meaning', but computer scientists need not care; they are interested in the behavior of programs.

The phrase 'understanding of the context' in the above definition of communication is fraught with implication. Roughly, context is an environment, a 'place' where things occur or an utterance is made. The anticipation of, and preparation for, a particular environment is a basic design question.

In order to cut the notion of context down to a manageable size, we first make the distinction of three fundamentally different means by which software technology embodies context. Along the lines of Varala et al. [VAR91], these technologies are 'enactive', 'emergent', and 'symbolic'.

Neural nets are representative of enactive technology. In neural nets an associative memory embodies the relationship between environmental demands and behaviors in response to those demands, that is, context is reflected in the trained associative memory. Enactive technology is outside the scope of this paper.

Genetic algorithms are representative of emergent technology. In the emergent approach, context is reflected in the statistical distribution of individual designs within a spectrum of possible designs. Each individual is an attempt to encode (and cope with) the context. Emergent technology is outside the scope of this paper.

Symbolic technology includes procedural (e.g. object-oriented) constraint and logic programming, ontologies, and natural language. Unlike the others, symbolic technology is inextricably linked with language and hence meaning and context can be communicated among individuals, not just embodied by them. Roughly speaking, dictionaries, thesauri and computer-based ontology systems are each built upon networks of symbols related to each other by various notions of resemblance.

### 2.1.1. GROUND

Webster's defines 'ground' as “the foundation for an argument, belief or action.” Ontologies are one technology where the notion of ground is prominent. In some sense, unambiguous communication of information is enabled by relating consensus domain terminology to widely held ground terms. There are however, overriding issues in making this goal a reality. This is discussed in the section on unambiguous communication.

Some ontology systems (for example CYC) provide a mechanism to allow reasoning under multiple sets of ground terms ('assumptions'). [deK86] Intuitively this appears useful in manufacturing situations but more consideration is necessary.

### 2.1.2. CONTEXT AVAILABILITY IN INFORMATION EXCHANGE

In the design of communicating systems one may choose to resolve all questions of the context of information exchanged at design time. Alternatively, the exchange might include context-setting information, or 'meta-level' information.

Many high-performance systems do not communicate much context information between each other; the context of the information exchanged was agreed upon by the system designers a priori and is implicit in the software.

Within the investigated ontologies, the meta-level is available as references that could lead all the way to ground terms. What is not as clear at this point of our investigation is in what manner meta-level information is made available in information exchange. We may study the Knowledge Interchange Format (KIF) in our second year effort as a means towards this end.

### 2.1.3. SEAMLESSNESS

A context or meta-level may be available but only in a different technology or facility than the problem solving machinery. This is a common occurrence (e.g. the OMG implementation repository). Seamlessness refers to the continuity (or 'transparency') of the information content with its context-setting information. That is, in a seamless environment the problem solving machinery can get context setting information without appealing to another for service.

The principle question here with respect to the investigated ontologies is how to marry the technology to the 'problem solvers'. That is, how should ontologies be interfaced to applications such as schedulers and workflow systems? We did not attempt to answer this question in this year's investigation.

### 2.1.4. FORMALITY

At times, formality in ontologies seems to mean the degree to which the ontology resembles mathematical logic. Resemblance to mathematical logic in itself however does not suggest a purpose for formality. We suggest the following:

- Formality is about making valid inferences (and thus getting expected behaviors).
- Formality is about traceability to ground.
- Formality is about enabling computational tools (to manage complexity etc.).

There are many notions exchanged among engineering and manufacturing systems for which it might be quite inefficient to attempt to represent the ideas in anything but traditional mathematical terms. Computer-Aided Design (CAD) geometry is of this sort.

## 2.2. ANSWERS TO “WHY ONTOLOGIES?”

The previous section provides a foundation of ideas concerning communicating systems. In this investigation we have not yet resolved many design questions regarding how this technology may be employed. Assuming that there are reasonable answers to these questions, however, we may still ask what added value the investigated ontologies might provide to the development of a manufacturing information infrastructure. Here we consider the contribution ontologies might make to eliminate ambiguity in communication, and in the standards making process.

### 2.2.1. UNAMBIGUOUOS COMMUNICATION

We consider two questions with respect to the investigated ontologies and the goal of unambiguous communication:

1. Where does the problem of ambiguity reside?
2. Do communicating systems need access to ground terms?

In regards to the first question we note that generally speaking, manufacturing systems have lagged financial systems in their ability to exchange data. One explanation for this might be that less time and effort has been invested in the integration of manufacturing systems. There are other equally valid explanations:

- Manufacturing data and its interrelationships are complex, perhaps much more so than financial data.
- There is no universally accepted meaning to terms used in manufacturing.

It is important to note that these two points are statements about the nature of manufacturing information itself: their solution requires knowledge of manufacturing foremost. Information technology isn't the issue here. The problem is getting individuals in whole industries to agree on the meaning of perhaps thousands of terms such as "part version" and "part revision." Standards efforts such as IEC 61360-4, a dictionary of standard terminology for electronic components [ISO97] are representative of the sort of work that must be done.

With regards to the second question, (whether communicating systems need access to ground term), communicating systems rarely need access to ground terms. The exception to this is those few systems that mediate data. The problem of bringing meaning to the exchanged data always leads back to reference to shared understanding. For this reason, we conclude that ontologies do not offer significant benefit towards making current information exchange methods more reliable.

### 2.2.2. STANDARDS MAKING AND SEMANTIC ALIGNMENT EFFORTS

As suggested above, there is a need for comprehensible industrial terminology developed through consensus. As increasing amounts of the work involved with bringing a product to market are possible on the Internet and as increasing amounts of the supply chain become integrated, the audience for this terminology widens. A terminology shared among manufacturers of centrifugal pumps is most useful when it can also be understood by makers of bearings for these pumps and builders of process plants using the pumps.

Although this aspect of the studied ontologies has not been a central focus of this year's effort, we expect it to be an area where the investigated ontologies may provide great benefit.

### 2.2.3. INDUSTRIAL INFORMATION INFRASTRUCTURE

Distributed objects, agents, integrated workflow and supply chains are common themes emerging in the development of an industrial information infrastructure. This mode of operation emphasizes ad hoc access to objects. This is in contrast to the more traditional approach of organizing systems around the semantics of a shared database. The emerging architecture suggests that ad hoc access to shared meta-data and terminology might also prove useful in future information systems. If this turns out to be true, then technology such as the investigated ontologies may prove to be essential in supplying a terminology and meta-data in computational form.

### 2.3. SUMMARY

Assessing the value of ontologies (or any other technology) to the problem of communicating manufacturing information is, in part, a matter of determining whether it is the simplest possible means to establish a context on exchanged data under which valid inferences (and thus expected behaviors) can be achieved.

The investigated ontologies provide seamless access to its meta-data and ground terms in a computational, formal form. However, one cannot say absolutely whether seamless systems with access to meta-data and possessing a theory of ground are better or worse designs because they possess these attributes. Many very successful, high-performance systems will continue to possess seams, no theory of ground and no access to meta-level information. The genius of design is in part in making the right choice with respect to how context is established and used in communication.

The investigated ontologies can contribute significantly to the alignment of consensus domain terminology.

## 3.0. APPROACH AND MAJOR FINDINGS FOR MANUFACTURING ANALYSIS

A systematic approach was taken throughout this project to ensure that a proper cross-section of manufacturing-related ontological systems were chosen, appropriate analysis criteria were determined, and a proper analysis was performed. The project started by doing a literature survey to determine what appropriate ontological systems were available. This survey included a thorough search of the web and numerous interactions with colleagues in the ontology field. From this survey, the following ontological systems were identified:

- ANSI Ad Hoc Group on Ontology Standards Representation [ANSI98]
- CYC [CYC98]
- Enterprise Ontology [ENT98]
- LOOM [LOOM98]
- MikroKosmos [MIKRO95]
- Ontolingua [ONTO96]
- Sensus [SENS98]
- SPAR (Shared Planning and Activity Representation) [SPAR98]
- STEP (Standard for the Exchange of Product model data) [ISO94]
- TOVE (Toronto Virtual Enterprise) [TOVE98]
- Wordnet [WORD98]

A high-level analysis of each of the above ontological systems was performed and a few systems were eliminated due to their lack of appropriateness to this project. In general, the project analyzed these ontologies against the following three criteria:

- the ontology's ability to represent manufacturing information (e.g., time-varying concepts, flow of materials, constraints, etc.),
- the amount of manufacturing information that was already represented in the ontology,
- the ability for the ontology to inference over the information represented.

The following systems were excluded from the analysis, along with the respective reason:

- ANSI Ad Hoc Committee on Ontology Standards – at the time the analysis was performed, this ontology was not mature enough to properly analyze. In addition, since the upper level of CYC was to be merged with this ontology, an analysis of CYC would be sufficient to also analyze this ontology.
- Sensus – only a taxonomy of terms without definitions were provided and the concepts represented in this system had already been merged with CYC through the Ad Hoc Group on Ontology Standards work
- SPAR – at the time this analysis was performed, it was not mature enough to analyze.
- STEP – it was too limited in domain (only product data), there were no formal definitions of concepts, and from the project participants' previous work with STEP, we know it would not be appropriate.
- Wordnet – it is more of an on-line super-dictionary than a knowledge base.

Once the ontological systems to be analyzed were determined, we moved on to determining the appropriate analysis criteria. It was decided that the project would base our analysis on typical manufacturing scenarios. This would involve identifying appropriate manufacturing scenarios, extracting the concepts inherent to that scenario, grouping the concepts into appropriate categories, and developing inferencing questions that are based on those concepts. We would then see how well existing ontological systems could model those concepts and determine how well they could answer questions pertaining to those concepts.

The CAMILE "Factory from Hell" scenario [CAMI91] was identified as an appropriate scenario for our manufacturing analysis. This scenario was developed by Ken McKay as part of an assignment through CAM-I (Consortium for Advanced Manufacturing, International). The scenario details a fictitious factory (based heavily on knowledge gained through site visits to actual factories) including information on many departments and the decision making processes which occur throughout the development of a product. The concepts, which were detailed in the scenario, were extracted and grouped into manufacturing-related categories. The chosen categories were (in no particular order):

a) Penalties
b) Costs
c) Financials
d) Scheduling
e) Process Planning
f) Product Configuration
g) Resource Planning
h) Resources
i) Inventory
j) Batches/Lots
k) Orders
l) Customer/Vendor
m) Scrap/Rework
n) Manufacturing Execution

Using these categories and the concepts in each category, we initially examined each of the ontological systems to determine how well they could represent those concepts. Namely, we rated each ontology with respect to the following four categories:

1. Required concepts are not represented in ontology. Related information infrastructure is not available and must be modeled before concepts can be represented.
2. Required concepts are not represented in ontology. Related infrastructural concepts are available. Modeling of required concepts could take place primarily by combination of existing concepts.
3. Representation of required concepts could be achieved through specialization or minor modification of existing concepts.
4. Required concepts are available in ontology and would require either trivial modifications or none at all.

During the initial phases of this analysis, it was found that a few other ontologies were not appropriate for further analysis for the reasons described below.

- LOOM – it is a language and environment. It is not an ontology itself but is quite suitable for implementation of projects using ontologies. Therefore, LOOM would not be appropriate for the development and modeling of a manufacturing ontology.
- MikroKOSMOS – its purpose is to provide a general mechanism for mapping meaning between languages. As such, it has been developed with different capabilities and design structure than would be needed for a manufacturing ontology. Specifically, Mikrokosmos provides no inferencing capability to answer questions that are not explicitly answered in the knowledge base.
- Ontolingua – it is an ontology authoring tool and not an ontology itself. Since this body of work is a development environment, it is not appropriate to attempt to evaluate its direct applicability to manufacturing.

For the above reasons, these ontologies were not further analyzed.

The remaining three ontologies, CYC, Enterprise Ontology, and TOVE, were then analyzed in further detail. The results of this analysis showed that all three packages were approximately equally able to represent manufacturing information. However, the inferencing capabilities in CYC seemed a bit more mature than the other two packages analyzed. Also, the close relationship that NIST and the ATP Ontology project have with Cycorp would allow the project to more easily leverage Cycorp staff's expertise while modeling the manufacturing ontology. For these reasons, the project decided to proceed with CYC to model the manufacturing ontology.

## 4.0. MANUFACTURING ONTOLOGICAL SYSTEMS INVESTIGATED

Table 1 summarizes the major points related to the ontologies that were investigated. A more detailed description of each of these ontologies can be found in the subsections below.

### 4.1. ANSI AD HOC COMMITTEE ON ONTOLOGY STANDARDS

The goal of the ANSI Ad Hoc Group [ANSI98] (associated with the ANSI X3T2 committee on Ontology Standards) is to merge the upper level ontologies of many of the well-known ontological systems (CYC, Pangloss, Penman, Wordnet, etc.). An "upper level ontology" is an ontology of the most general conceptual categories. There are a number of such ontologies out in the world that have proved very useful in natural language processing and other AI oriented applications, as well as in enterprise modeling and database integration. The challenge is that it is difficult to translate between these applications because of the differences in their upper level ontologies. The purpose of the standard will be to provide a sort of ontological baseline to support translation and integration between ontology-based applications, and hopefully also to serve as the starting point for future upper level ontologies.

At the time the analysis was performed, all that was available from this group was a high level taxonomy of terms without any definition of what the terms meant. It was assumed that the location of any term within the taxonomy was meant to serve as a loose definition of the term. However, because this ontology standard was being adopted by other systems that we were analyzing, such as CYC, the analysis of those other systems would indirectly allow us to analyze the ontology standard. In addition, because those systems provided additional capabilities that the ontology standard alone did not (e.g., inferencing capabilities, formal definitions of terms, user interfaces, etc.), the respective systems would be a more appropriate choice for use to model a manufacturing ontology. For these reasons, this Ontology Standard was not investigated any further.

**Table 1: Summary of Ontologies Investigated**

| Ontology | Domain | Purpose | Provides Inferencing? | Development Framework or Full Ontology |
|---|---|---|---|---|
| CYC | Generic | Enable common sense reasoning about the world | Yes | Full ontology |
| Enterprise Ontology | Business enterprise and organization modeling | Comprehensive ontology whose main groupings consist of activities, organization, strategy, marketing, and time. | No | Full Ontology |
| LOOM | Generic | A language and environment for constructing intelligent applications | Yes (forward, truth maintenance) | Development Framework |
| Mikro-KOSMOS | Knowledge-based translation of natural language | Translate natural language text from one language to another via a language-neutral text meaning representation | No | Full Ontology |
| Ontolingua | Generic | Development environment and authoring tool for the creation of modular, reusable ontologies. | No | Development Framework |
| TOVE | Enterprise integration | Provide a generic, reusable data model including shared terminology and meaning that each agent can jointly understand and use | Yes | Full Ontology |

### 4.2. CYC

CYC [CYC98] is a very large, multi-contextual knowledge base and inference engine developed by Cycorp. The goal of the CYC project is to construct a foundation of basic "common sense" knowledge--a semantic substrate of terms, rules, and relations--that will enable a variety of knowledge-intensive products and services. CYC is intended to provide a "deep" layer of understanding that can be used by other programs to make them more flexible.

A drawback to CYC is that its level of knowledge is so "deep" as to be unintuitive to all but CYC knowledge experts. Higher-level knowledge is left to application developers. Not surprisingly, there are large gaps in CYC 's higher-level KB as it has only been extended to support whatever application was required for its use. Only some aspects of these extensions are publicly available. Manufacturing is not well represented by the KB.

The CYC technology is composed of the knowledge base and inference engine, the CycL representation language, interface tools and application modules. Cycorp is currently working on tools to ease the difficulty of adding to the KB. At the present time, the CYC KB contains tens of thousands of terms and several dozen hand-entered assertions involving each term. CycL, the CYC representation language, is a large and flexible knowledge representation language. It is essentially an augmentation of first-order predicate calculus (FOPC), with extensions to handle equality, default reasoning, and some second-order features.

### 4.3. ENTERPRISE ONTOLOGY

The Enterprise Ontology [ENT98] was built as part of the large Enterprise Project at the Artificial Intelligence Applications Institute at the University of Edinburgh, in collaboration with industry partners. The focus of the project is to promote the use of knowledge-based systems in enterprise modeling and organizational support. The result of this initiative was an Enterprise Toolset, one component of which is the Enterprise Ontology.

The Enterprise Ontology is relatively comprehensive and includes over 90 different concept classes and over 60 relations between concepts. In order to represent concepts within the Enterprise Ontology itself, a *meta ontology* was developed, which includes more general modeling terms such as entities, relationships, roles, attributes, and so on. Building on these terms, the concepts in the Enterprise ontology are divided into five categories: activities, organization, strategy, marketing and time. Of course, there are interactions among the various categories of concepts. For example an activity may take place over an interval of time as part of a plan.

The intent of the Enterprise Ontology is not to model specific types of enterprises, but to provide a general model that is oriented more towards business and organization than towards a specific domain. From the perspective of the evaluation being performed in this paper, the Enterprise Ontology is greatly lacking. Virtually all concepts and terms that are specific to manufacturing enterprises are missing from this enterprise model. However, the Enterprise Ontology is still viewed as a valuable resource because of the infrastructure it provides. The *meta ontology* provides a flexible set of primitives for building concepts, and since manufacturing enterprises are a subset of business enterprises in general, many of those aspects of a manufacturing enterprise that are not manufacturing-specific are present in the existing ontology. For instance, concepts such as resources, people, machines, and plans will have direct applicability within a manufacturing enterprise model. It should be noted that that in most cases, for application to a manufacturing enterprise further specification of concepts existing within the current ontology will be necessary.

### 4.4. LOOM

"Loom is a language and environment for constructing intelligent applications. The heart of Loom is a knowledge representation system that is used to provide deductive support for the declarative portion of the Loom language. Declarative knowledge in Loom consists of definitions, rules, facts, and default rules. A deductive engine called a classifier utilizes forward-chaining, semantic unification and object-oriented truth maintenance technologies in order to compile the declarative knowledge into a network designed to efficiently support on-line deductive query processing." [LOOM98]

As this quote makes clear, Loom is a language and environment. It is not an ontology itself but is quite suitable for implementation of projects using ontologies. Loom is written in Common Lisp and the Common Lisp Object System (CLOS) and is easily integrated into Common Lisp programs. The importance of Loom in this study is that it exemplifies the sort of infrastructure that exists to enable development of high-quality knowledge-based systems. Because Loom is not a commercial product (it is the intellectual property of the University of Southern California) there are fewer barriers to its use.

Our exploratory work with Loom suggests that it is easy to use. Although there may be concerns among some about Common Lisp not being a mainstream programming language, the development of a robust Common Lisp-based HTTP server and a CORBA binding to Common Lisp has eased this problem somewhat.

### 4.5. MIKROKOSMOS

The ultimate objective of the Mikrokosmos [MIKRO95] research project is to define a methodology for representing the meaning of text in a language-neutral format called a text meaning representation (TMR). This would provide a mechanism for Knowledge-Based Machine Translation (KBMT) of natural language text from one language to another (via an intermediary translation into a TMR). In pursuit of this goal, researchers at New Mexico State University have conducted a comprehensive study of linguistic and language use phenomena. These phenomena have been encapsulated in various "microtheories" which are united through the control architecture of the KBMT system.

The principle objective of the Mikrokosmos project is, unfortunately, not directed at arbitrary queries of a specific knowledge base, but rather, a general mechanism for mapping meaning between languages. As such, it has been developed with different capabilities and design structure than would be needed for a manufacturing ontology. Specifically, Mikrokosmos provides no inferencing capability answering questions that are not explicitly answered in the knowledge base. This capability is vital for providing useful information in a Manufacturing context. The Mikrokosmos ontology contains a wide variety of basic concepts related to manufacturing (e.g., drill, cut, and make), but it has very few detailed concepts that would be helpful for manufacturing. As such, implementing a manufacturing ontology using Mikrokosmos would require the development of tools for inferencing capabilities and general querying of the knowledge base, as well as adding a tremendous number of detailed concepts to the knowledge base.

### 4.6. ONTOLINGUA

The Ontolingua [ONTO96] ontology development environment, developed at the Stanford University Knowledge Systems Laboratory, consists of a suite of authoring tools for creating and browsing modular, reusable ontologies. The set of tools provides a World Wide Web-based interface for ontology creation, allowing remote ontology creation or browsing of existing ontologies, many of which are available through the server Ontolingua Server at Stanford University.

The Ontolingua ontology development environment models information using the Ontolingua language [GRU93], a language based on the Knowledge Interchange Format (KIF) [GEN92]. Ontolingua expands the basic first-order predicate logic formalism provided by KIF to also include syntax for an object-oriented representation (classes, instances, slots, relations, etc.) In addition to the web-based authoring interfaces, the development environment also provides translation into other knowledge representation languages, including Loom [MAC91], Epikit [GEN90], Generic-Frame [CHA97] and pure KIF.

The purpose of this paper is to evaluate ontologies and not ontology authoring tools. Since this body of work is a development environment, it is not appropriate to attempt to evaluate its direct applicability to manufacturing. However, because of its advantages (ease of use, availability of existing modular ontologies to leverage from, ties to KIF and translator facilities to interface with other knowledge representation languages), this environment would be a strong candidate for consideration if a new manufacturing-related ontology were to be built from scratch. Indeed, this development environment was used to model the Enterprise Ontology, which is one of the ontologies evaluated in this paper.

### 4.7. TOVE (TORONTO VIRTUAL ENTERPRISE)

In order to support enterprise integration, it is necessary that shareable representation of knowledge be available that minimizes ambiguity and maximizes understanding and precision in communication. Secondly, the creation of such a representation should eliminate much of the programming required to answer "simple" common sense questions about the enterprise. The goal of the TOVE [TOVE98] project is to create a generic, reusable data model that has the following characteristics:

- provides a shared terminology for the enterprise that each agent can jointly understand and use,
- defines the meaning of each term (a.k.a. semantics) in a precise and as unambiguous manner as possible,
- implements the semantics in a set of axioms that will enable TOVE to automatically deduce the answer to many "common sense" questions about the enterprise, and
- defines a symbology for depicting a term or the concept constructed thereof in a graphical context.

The TOVE reusable representation represents a significant ontological engineering of industrial concepts. All axioms and definition are specified natively in the Knowledge Interchange Format (KIF) [GEN92]. It also has presentations using the Frame Ontology from the Knowledge Systems Laboratory (KSL) (http://www.ksl.stanford.edu/) from Stanford and will shortly have a presentation in XML (eXtensible Markup Language) [XML98].

The work began by translating the ontologies developed at Carnegie Mellon from LISP into a C++ environment. The ontology was then modified and extended. Currently, the ontology spans: activities, state, causality, time, resources, inventory, order requirements, and parts. There has also been work to axiomatize the definitions for portions of our knowledge of activity, state, time, and resources. The axioms are implemented in Prolog and provide for common-sense question answering via deductive query processing. Future work focuses on the development of ontologies and axioms for quality, activity-based costing, and organization structures.

## 5.0. DISCUSSION AND CONCLUSION

Apart from the major findings described in the above two sections, there were a couple of other interesting findings, discussed below.

### 5.1. CONTEXT IN ONTOLOGIES

The main objective of ontology development is to develop a standard vocabulary or to predefine terminology in order to facilitate exchange of information. Ontologies help create a uniform basis for information exchange by enabling the representation and communication of the meaning of a given term. However, a secondary issue that must be addressed arises when a term has multiple valid definitions. Being able to represent these definitions formally does not solve the problem of knowing which definition to use in a given circumstance.

This problem is being addressed in several ontology development efforts through the use of contexts in ontologies (e.g. [CYC98], [MIKRO95], and [TOVE98]). Context, also referred to in some efforts as microtheories, allows additional information beyond specific formal term definitions to be incorporated into an ontology. This contextual information may be represented implicitly or explicitly within an ontology. In the case of the Ontolingua Ontology Development Environment, modular ontologies are created and combined or included as components of larger ontologies. In one sense, this can be thought of as an implicit representation of context, since a term may be defined one way in one ontology and differently in another. MikroKosmos uses context to help resolve the meaning of words that could have multiple meanings. Although the way they do this is vague (possibly because it provides them with part of their proprietary advantage), it partially involves grammatical rule (e.g., adjectives follow nouns in Spanish, adjectives precede nouns in English). The placement of these words in a sentence provides the context to help to define what the words mean. CYC represents context using "microtheories", each of which is essentially a bundle of assertions that share a common set of assumptions; typically microtheories are focused on a particular domain of knowledge, a particular level of detail, a particular interval in time, etc.

### 5.2. INFERENCING IN ONTOLOGIES

Inferencing, in general terms, is the ability for a system to deduce new information that is not explicitly represented from concepts that are currently represented in a knowledge base. For example, let's assume that a particular manufacturing process (Process B) must be performed within 24 hours of the completion of another manufacturing process (Process A). In order for a scheduling program to decide when to schedule Process A, it must have access to certain information. Some of this information would be explicitly represented, such as the expected durations of Process A and B, the current time, and the standard hours that the factory is open. However, some of the information necessary is unlikely to be explicitly represented, such as, whether or not the factory is open "tomorrow". This type of information would need to be deduced from information that is explicitly represented, such as, today's date, today's day of the week, scheduling holidays, and factory hours. An inference engine could provide this deductive capability to determine the information needed but not explicitly represented.

In the ontologies investigated, the tools were designed to work with specific representations; namely: 1) inference engines developed by Cycorp Inc. to work on their CycL representation, 2) a deductive engine developed with LOOM to specifically work on the LOOM knowledge representation, and 3) a set of tools developed all around the world to work on information represented in the Knowledge Interchange Format (KIF). The LOOM deductive engine is discussed in Section 4.4.

## 6.0. REFERENCES

[deK86] deKleer, Johan, "An Assumption-Based Truth Maintenance System" Artificial Intelligence Journal 28 (163-196), Deerborn, MI, 1986.

[ANSI98] ANSI Ad Hoc Group on Ontology Standards, http://www.ksl.stanford.edu/, August 20, 1998.

[BLO33] Bloomfield, Leonard, "Language" University of Chicago Press, 1933.

[CAMI91] McKay, Ken, et. al., "Report: Intelligent Manufacturing Management Program, State of the Art Scheduling Survey," CAM-I Report R-91-IMM-01, Waterloo, Ontario, Canada, 1991.

[CHA97] Chaudhri, V., A. Farquhar, R. Fikes, P. Karp and J. Rice, *The Generic Frame Protocol 2.0*, KSL-97-05 (technical report), Stanford Knowledge Systems Laboratory, Stanford University, Stanford, CA, 1997.

[CYC98] Cycorp Inc, The CYC Technology, http://www.cyc.com/tech.html, July 31, 1998.

[ENT98] Uschold, Mike, et. al., "The Enterprise Ontology," The Knowledge Engineering Review, Vol. 13, Special Issue on Putting Ontologies to Use (eds. Mike Uschold and Austin Tate), 1998.

[FAR96] Farquhar, A., R. Fikes and J. Rice, "The Ontolingua Server: A Tool for Collaborative Ontology Construction," *Proceedings of the Tenth Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, Canada, November 1996.

[GEN90] Genesereth, M. R., *The Epikit Manual*, Epistemics, Inc., Palo Alto, CA, 1990.

[GEN92] Genesereth, M. R. and R. E. Fikes, *Knowledge Interchange Format, Version 3.0 Reference Manual*, KSL-92-86 (technical report), Stanford Knowledge Systems Laboratory, Stanford University, Stanford, CA, 1992.

[ISO94] ISO, 10303-1:1994, Product data representation and exchange: Part 1: overview and fundamental principles.

[ISO97] ISO/IEC 61360-4:1997 Standard data element types with associated classification scheme for electric components. http://www.codus.co.uk/intstds.htm

[LOOM98] "Loom Homepage" http://www.isi.edu/isd/LOOM/LOOM-HOME.html, August 20, 1998.

[MAC91] MacGregor, R., "The Evolving Technology of Classification-based Knowledge Representation Systems," *Principles of Semantic Networks: Explorations in the Representation of Knowledge*, J. Sowa (ed.), Morgan Kaufmann, San Mateo, CA, 1991.

[MIKRO95] Mahesh, K., and Nirenburg, S., A Situated Ontology for Practical NLP. In Proc. Workshop on Basic Ontological Issues in Knowledge Sharing, International Joint Conference on Artificial Intelligence (IJCAI-95), Montreal, Canada, Aug. 19-20, 1995.

[ONTO96] A. Farquhar, R. Fikes, & J. Rice. The Ontolingua Server: A Tool for Collaborative Ontology Construction. Knowledge Systems Laboratory, KSL-96-26, Stanford, CA, September 1996.

[SENS98] "Ontology Creation and Use – SENSUS," Natural Language Group at USC/ISI, http://www.isi.edu/natural-langauge/resources/sensus.html, November 30, 1998.

[SPAR98] Tate, Austin, "Planning Initiative: Shared Planning and Activity Representation: SPAR Homepage", http://www.aiai.ed.ac.uk/~arpi/spar/, November 30, 1998.

[TOVE98] Gruninger, Michael, "Enterprise Modeling," http://www.ie.utoronto.ca/EIL/comsen.html, August 19, 1998.

[VAR91] Varela, Francisco J., Thompson, Evan and Rosch, Elanor, "The Embodied Mind: Cognitive Science and Human Experience" The MIT Press, Cambridge, Massachusetts, 1991.

[WORD98] "Wordnet -- A Lexical Database for English", Cognitive Science Laboratory, Princeton University, http://www.cogsci.princeton.edu/~wn/, November 30, 1998.

[XML98] Extensible Markup Language (XML) 1.0: W3C Recommendation from 10-February-1998, http://www.w3.org/TR/REC-xml, August 28, 1998.