Enabling Intelligence with Ontology

Richard A. Martin, Tinwisle Corporation
RichardM@tinwisle.com

I expect some feel that “intelligent enterprise” is an oxymoron, that intelligence and enterprise occupy extremes of a socially organized dimension of human endeavor, or that intellect is an individual quality not attainable in aggregate. Yet, the extent to which an enterprise survives in time attests to the intellectual capacity necessary to craft something distinct from any or all of its constituent parts. Rather than antithetical in their relation, “intelligent enterprise” states synonymy – an intersection of dimensions in the human endeavor. The intelligence of an enterprise is reflected in the emergence of robust operations.

Two distinct notions of enterprise should be considered when conjuring the intelligent creation. The preceding paragraph considers enterprise in its most expansive organizational meaning for which there is indeed some question regarding synonymy upon occasion. For systems of immense size, where mere existence rather than intellectual capability may drive pragmatic behaviors, the identification and quantification of intelligence, quite independent of our ability to influence behavior, is generally considered to be beyond the scope of the engineering disciplines. A second notion of enterprise, as an industrious, systematic activity (1) of moderate scale, seems more appropriate for our consideration.

The preceding paragraphs present a natural language description of concerns for the conjunction of the symbols “intelligent” and “enterprise” in the context of a systems engineering discipline. Just as we bound systems for the purpose of examining and constructing their features, so too do we bind symbols of that system to a context and thereby fix the relationship of symbols to meaning in that context. In this manner we establish an ontological context for a system. (In this paper, the symbol “context” refers to the set of propositions held as true by the system and hence “ontological context” refers to the definition and relationship of symbols as constrained by the propositions assigned to them.)

Unfortunately, while we have achieved considerable success in binding components to a system, our success in binding meaning to symbols in a context is highly stratified. The symbolic means of communication, i.e., the assignment of equivalent meaning to a symbol at both extremes of a communication channel, improves in efficiency as the context becomes more precise and less ambiguous. The propagation of binary state along a communications buss inside this machine I’m now using is highly efficient. But then the context for that transfer of meaning is very narrow indeed. As the context broadens to include me as a composer of symbols, the efficiency remains high because the machine is transcribing rather than interpreting my keystrokes.

Now you enter the context of this presentation by participating in its systemic operations. Assuming that the transcription and publishing process are effective, you can read the words and the sentences they compose. But, do those sentences have the same meaning to you as they do to me? Is your interpretation that which I intend or is your contextual perception distinct enough to impart a different meaning? For example, my use of term context is likely different from many who attach a strictly external rather than largely internal meaning of the term with respect to a system boundary. The most significant characteristic I can ascribe to natural language is the diversity in interpretation it allows. So by which means can I assure the extension of my intent if natural language does not suffice - if the words as symbols and sentences as composition lack the clarity of meaning that I give them as I write. Is it as simple as binding those symbols to an ontological context?

In establishing a system solution for an enterprise we have become accustomed to partitioning the problem space through hierarchical decomposition and composing a solution space by aggregation of parts. As the decomposition/composition is occurring, we attach symbolic labels to the various parts and write narrative, although often terse, descriptions of the resulting components. On rare occasions we even present narrative arguments for the existence of such components. The labels and narratives are drawn from natural language often “enhanced” by domain jargon. The labels form a term taxonomy for the problem/solution space and the narratives serve to refine the distinctions attached to the terms. Is this a suitable ontological context?

Because we desire the automated abstraction of meaning for efficient communication, the ontological context should be cast in a form amenable to such automation. This requirement most often results in determining the truth value of a proposition judged according to the rules of construction and constraints embodied in the ontological referent.

When that referent consists of terms supported by formal axioms bound as properties that constrain the truthful use of the term, it is possible to use an inference engine to determine the truth value of a proposition. This automated recognition of truthful statements is the driving force behind efforts to standardize both the development and use of ontology.

After centuries of debate and practical experience, it is now clear, despite some persistent attempts to the contrary, that no single ontology will serve adequately for all interpretations of symbolic meaning. I can even argue that each enterprise context has its own unique ontology that is bound to its distinctive identity. Thus the emphasis is shifting from the appropriate ontology to the appropriate methods and structures to support the interoperation of distinct ontology. Unlike the notion of natural language translation where interpretation is tempered by dialect and culture, the formal description of ontology, i.e., its precise mathematical construction, can result in accurate correspondence of intent from one end of the communication channel to the other. We can raise the binary channel efficiency to more complex semantic content.

In many ways, crafting a specific ontology is far easier than making available the means by which formal ontological components can interoperate. Two major issues have risen as impediments. The first is agreement on the very top-most terms of the ontology languages and the second is agreement on the way in which we account for context. Unfortunately, progress on either of these fronts is difficult since many existing ontological commitments must also be accommodated – a legacy problem.

For an enterprise to achieve robust intelligence it should commit to an ontological mechanism for context management that is much larger than itself precisely because the enterprise boundary is constantly shifting as it adapts to the forces of change. This realization has resulted in many different approaches to the construction of a well-defined ontology. A simple search using the term “ontology” yields both explicitly defined ontology targeting specific kinds of enterprises and discussions of the ways in which ontology is created. The advent of RDF and XML, with a self definition of category capability, is pushing the issue of ontological precision and commitment into renewed urgency (2). Whereas, individual dominant partners and dedicated consortia have in the past been able to dictate ontological constructs to achieve homogeneity of meaning, the wider base of consumers and producers attending to WWW activity result in much less organizational control for term use and meaning. To address this potential for misunderstanding among users, several notable efforts are underway and some of the historical efforts have new vigor.

The largest, and perhaps best known effort to codify symbolic meaning is that of Cycorp. Recently they have released OpenCyc as an open source version of their Cyc technology with a target of 6,000 concepts and 60,000 assertions together with an inference engine, knowledge browser and knowledge capture tools (3). *Assertions* in CycL, the language in which Cyc is written, consist of a *formula*, *microtheory*, *truth value*, *direction* and *support*. *Formulas* are well-formed expressions whose context is restricted by its *microtheory* containment. Five possible *truth values*, of which “default true” and “monotonically true” are the most common, resolve an *assertion*. *Direction* controls the inference mechanisms when applicable and *support* refers to the formal arguments that justify the *assertion*.

Even with its relatively large base of knowledge, the primary use of Cyc is now, as it has been for many years, the production of specialized ontology for particular applications. In this sense, OpenCyc is a foundational ontology to be shaped by selection and enhancement for a particular purpose. The microtheory feature both constrains the source selections and allows extension to add contextual content.

A more recent entrant is the SUMO (Suggested Upper Merged Ontology) effort of Teknowledge Corporation (4) in response to a request for candidates from the IEEE P1600.1 project, Standard Upper Ontology Working Group (5). This effort uses the KIF (Knowledge Interchange Format) language to specify ontology assembled from many sources and validated by mapping to known ontological references. Figure A shows the first term definition found in the Ecommerce Services Ontology (6) that is a middle level ontology formed by extending SUMO with particular domain content (=> denotes logical implication and ? prefixes individual variables).

```
;; Upper Ontology of Services

;; definition of WebService
(subclass WebService ComputerProgram)
(documentation WebService "Web sites that do not merely provide static
information but allow one to effect some action or change in the world,
such as the sale of a product or the control of a physical device.")

(=>
     (instance ?Service WebService)
     (exists (?Resource)
            (providesAService ?Service ?Resource)))

(=>
     (and
            (instance ?Service WebService)
            (instance ?Using UsingAService)
            (patient ?Using ?Service)
            (agent ?Using ?User)
            (providesAService ?Service ?Provider))
     (exists (?Contract)
            (and
                   (instance ?Contract ServiceContract)
                   (agreementMember ?Contract ?User)
                   (agreementMember ?Contract ?Provider))))
```

Figure A: Ecommerce Services Ontology definition of WebService (KIF version)

Also associated with the IEEE SUO effort is the IFF Foundation Ontology based upon the Information Flow Framework that "is designed to support the semantic interoperability among various object-level ontologies."(7) This approach is considered complementary to SUMO in that it will support many SUMO-like ontology by creating a meta-level framework based upon category theory. As such, it goes beyond the first and second order logic found in other formal ontology efforts. Unfortunately, this venture into theories of Information Flow and Formal Concept Analysis yield an artifact unfamiliar to most professionals in the information sciences and engineering disciplines. Therefore, this most promising approach using a formal meta-level must be made more accessible to practitioners before its potential can be realized. Figure B presents the architectural structure emphasizing the meta-level components.

Perhaps the most challenging issue regarding the top-level distinctions for ontological interoperability is the approach taken to time. Whether time is an identifying characteristic for entities or simply a clock for the system to use can have a profound impact upon the correspondence of entities represented by terms of the ontology that must interoperate. Note that while the meta-level is neutral in this respect it must nonetheless mediate the correspondence between distinct notions of time as held by the instance ontology for which it provides the means of interoperation. The recent WonderWeb Foundational Ontologies Library effort to establish a Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) (8) identify time as a fundamental category and would thus seem to support the PSL (Process Specification Language) effort at NIST(9) that emphasizes the temporal relationships among processes. In contract, BPML (Business Process Modeling Language)(10) treats time as a clock event for process use rather than a defining characteristic of the process itself and therefore emphasizes the functional dependency between processes. Each approach has a sound theoretical basis that should facilitate efforts for interoperability.

One final comment regarding the structure of concepts is crucial. Because of the manner in which we tend to decompose/compose there is often an expectation that the meaning of the labels we attach will somehow mirror this hierarchic path structure. That is seldom the case. Concepts found in most ontology form a lattice or directed acyclic graph (DAG) that reflects both the divergence and convergence of term meaning. It is for this reason that the formal axioms associated with the term taxonomy are necessary to provide the insight for automated interoperation.

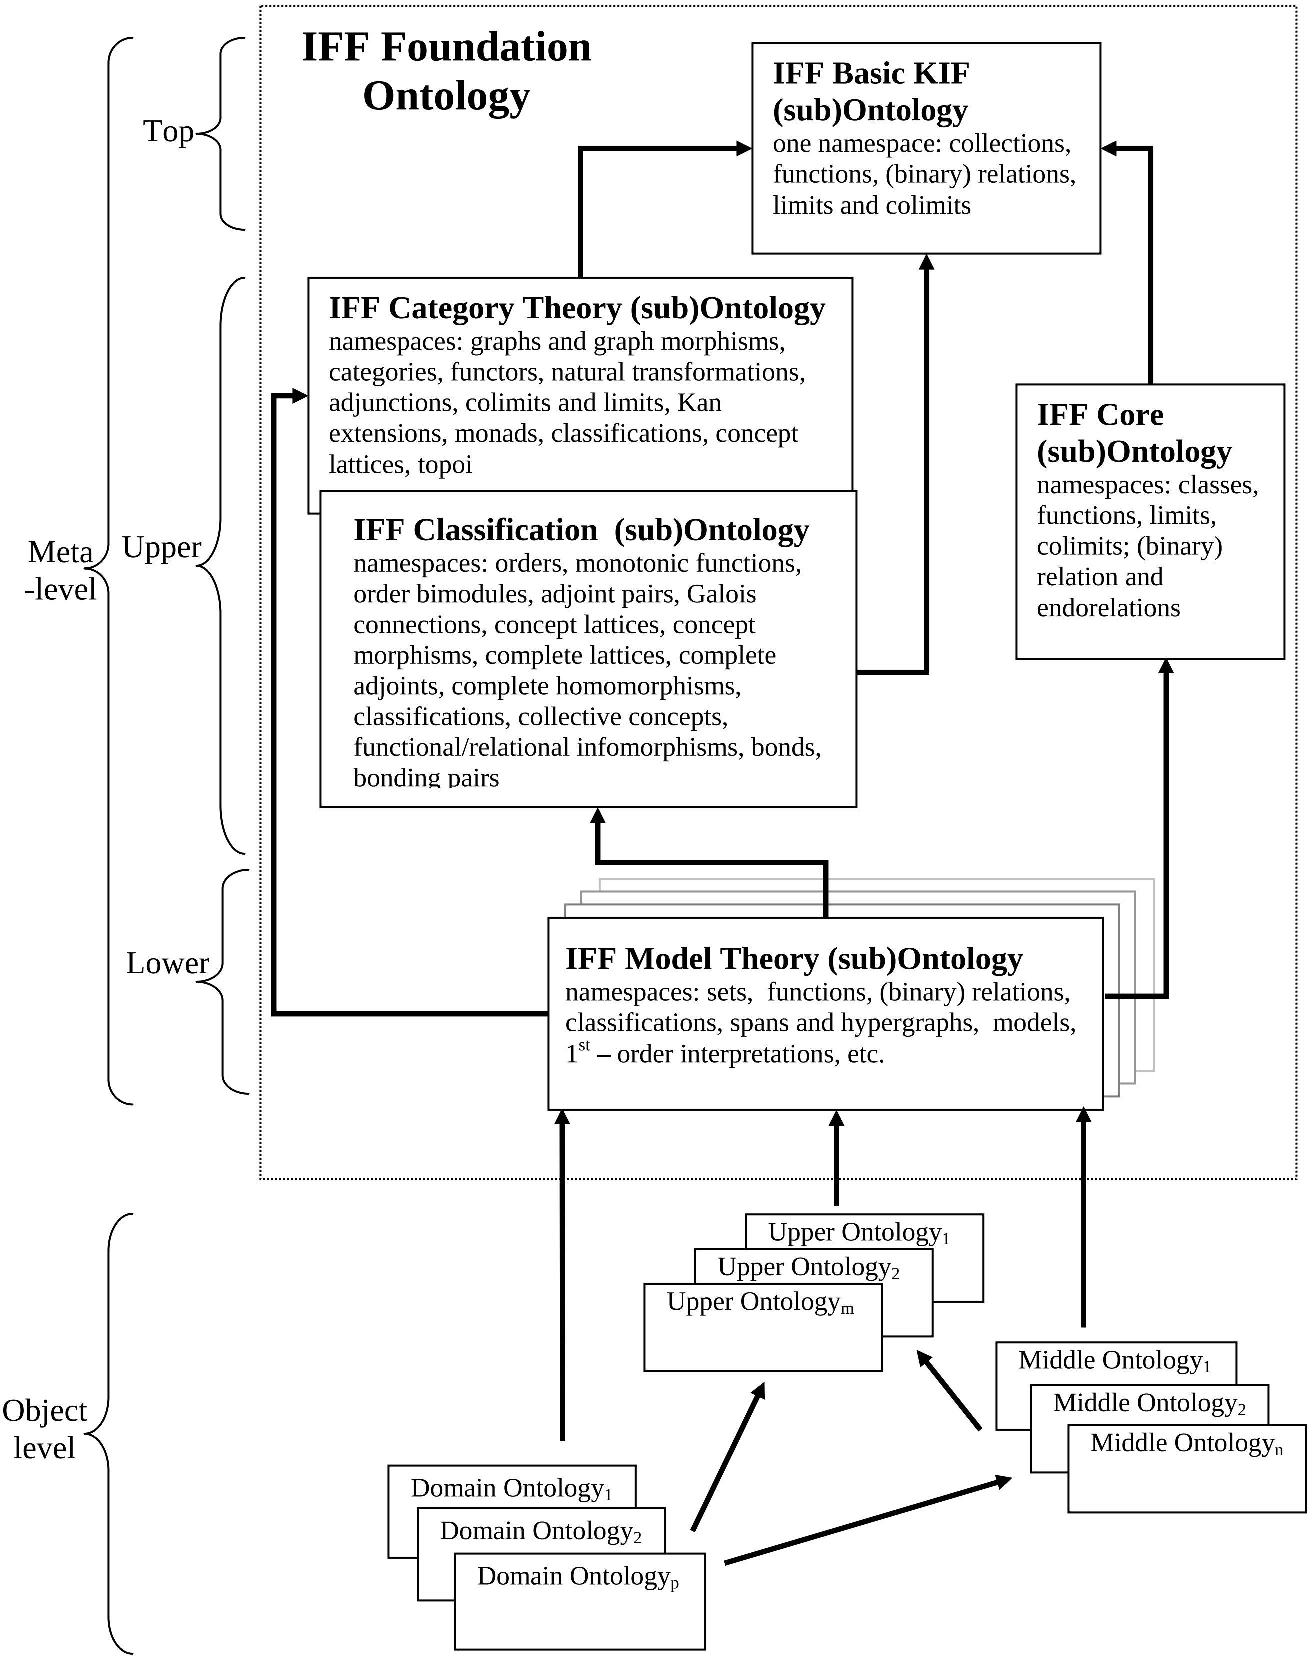

Figure B. The IFF Foundation Ontology Architecture (with dependencies)

A formal ontology should enable an enterprise to better comprehend its situation and should facilitate the emergence of intelligent actions. The ability to manage multiple ontological contexts will be essential for efficient and effective communications in support of goals as the intelligent enterprise pursues alignment and adaptation in response to change.

References:

(1) American Heritage Collegiate Dictionary @2002 Houghton-Mifflin Company accessed via http://www.yourdictionary.com/ (accessed November 11, 2002)
(2) see http://www.w3c.org/Consortium/ (accessed November 11, 2002) as a point of departure
(3) http://www.opencyc.org/ (accessed November 11, 2002)
(4) Niles, I., & Pease, A., (2001), Toward a Standard Upper Ontology, in Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001). Available at http://ontology.teknowledge.com
(5) http://suo.ieee.org/ (accessed November 11, 2002)
(6) Ecommerce Services Ontology. Available as an open source download at http://ontology.teknowledge.com
(7) http://suo.ieee.org/IFF/versions/20020515/IFFFoundationOntology.htm (accessed November 11, 2002)
(8) Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A., Schneider, L., (August 2002) WonderWeb Deliverable D17 The WonderWeb Library of Foundational Ontologies Preliminary Report. Available online at http://wonderweb.semanticweb.org/deliverables/D17.shtml (accessed November 11, 2002)
(9) http://www.mel.nist.gov/psl/psl-ontology (accessed November 11, 2002)
(10) http://www.bpmi.org (accessed November 11, 2002)