Enterprise-standard ontology environments for intelligent e-business

Alan Flett **Mike Brown**
Alan.Flett@SemanticEdge.com Mike.Brown@SemanticEdge.com
SemanticEdge
www.SemanticEdge.com
Kaiserin Augusta Allee 10-11
Berlin, Germany

Abstract

This paper highlights SemanticEdge's use of ontologies within their much broader conversational e-commerce system. After sketching some of the problems in e-commerce and e-business, we introduce the conversational paradigm as applied to e-commerce. This paradigm requires the use of ontologies in many areas, and we go on to outline the major issues we face in applying ontologies, both from a technical and a methodological aspect. We then go on to outline more general issues facing ontologies which we believe will be crucial to ontology technology's acceptance within modern enterprise-standard information technology systems.

Keywords: Ontologies, e-business, e-commerce

1 Introduction

The aim of this paper is to provide a pragmatic perspective on the emerging information technology of ontologies; how it can help solve various information integration problems in electronic business, and how successfully it is being introduced into a leading-edge e-business company's business processes. How successfully this technology can be fostered from the research environment to becoming a useable commercial information technology is of the utmost importance for the development of network-based information access; that ontologies, in the form which we will discuss them here, have been studied in various esoteric fields within computer science, principally artificial intelligence, should alert the interested information technology professional to the potential weakness of such technology: namely lack of accepted standards, lack of methodological guidance and support, and lack of enterprise-standard environments and tools. These are, arguably, the crucial issues which govern the take up of any new technology. The fact that the technology works, or has a sound, well reasoned, argument for its existence and usage, is assumed as a given. Having been convinced of its usefulness, the interested information technology professional then starts asking the "mundane" questions which typically overlap little with research-oriented academia: How do I apply this to my business?; Are there existing standards?; Is there methodological support?; Is there a unified common accepted methodology?: Is the technology of a standard that I can trust in my enterprise information systems?; Are there people trained in these methodologies and technologies? Whilst the last issue is, admittedly, outside most people's control, especially in an emerging technology phase, the other issues are ones that should be addressed by parties interested in the successful fostering any new technology to maturity. These are the issues which we address in this paper. We seek to highlight those issues which we believe will affect the usability of ontology technology within enterprises (both at the technology level and business process level), and we make some suggestions as to what potential solutions might be. There are certainly more issues than those we highlight here which will affect the successful uptake of ontology technology (e.g., alternative technology, the current trials and tribulations of e-commerce business models), yet we believe those which we do expose are outstanding open issues for the field to address.

The outline of the papers is as follows. In Section Two, we outline some of the problems that e-business faces. In Section Three, we go on to introduce the conversational paradigm for e-business, and how ontologies help provide some of the functionality required by that it; ontologies are also introduced here. In Section Four we cover some problems of ontology technology and methodology we see as important, suggesting some solutions we believe may be beneficial. Section Six concludes our paper.

2 An introduction to e-business and its problems

People communicate imprecisely; computers take you literally. Language is rich and thoughts are multidimensional; computers have no room for variability or implied meanings. While this situation is tolerated in a research environment, it is intolerable and even disastrous in commerce or in real life. For this reason, the socially adept computer has become the holy grail of e-commerce, and, in fact, of human-computer interactions in general. This kind of computer must be mindful of who it is talking to, remembering what was already stated and who said it. It must react appropriately within the context of a situation, remaining flexible and self-adjusting, and understanding of changes in intention or direction. It must be flexible in detecting and offering appropriate alternatives if a specific question can't be answered precisely. It must accept alternative forms of statements and be conversationally comfortable to talk to, offering multiple modes of input and output, with an easy conversational style. It must be skilled at negotiating, eliciting goals, offering alternatives in price, size, style, and be informative by offering suggestions to the user based on its understanding of available alternatives and the context of the conversation. Above all, it must be *socially adept*—a social detective capable of interpreting intentions, background, level of user-need from how queries are phrased. It must match a social model to the appropriate conversation script. It must be able to infer a user's level of frustration and respond appropriately and be able to query and respond appropriately within the context of a given situation.

That is an ambitious vision for e–commerce to realise. At present, there are various problems impeding the realisation of such a vision. These impediments break down into what may be called information integration problems and human-computer interaction problems. Issues from both categories will now be outlined.

Query rather than search

E-marketplaces are more confusing than their real world counterparts. Not only must products from several vendors be collected in the right place, but these products must be matched to the requirements or queries of the end user. Simple search systems can answer direct questions and match products to those questions when the terminology of both the product description and the query match. That is, most search techniques are keyword-based (including the most WWW search engines). This kind of search paradigm not only retrieves information which merely has the same terminology, but it also misses information which uses different terminology.

Static product search

However, search is just the beginning of a broader negotiation process in purchasing products (or, in the generic sense, of retrieving pertinent information). In a real marketplace, the buyers and the sellers come to an accommodation; buyers enter with a set of requirements, including price; so do vendors. These requirements are rarely met by any product; the price may be too high, the features of the product may be missing some of the initial requirements, or there may be no one product that proves a good match. At this point, both the buyer and the seller must decide what they are prepared to negotiate: Can the buyer pay a slightly higher price?; Could the vendor add an additional feature affordably?

Unstructured text

In addition, most e-marketplace technologies rely on structured databases to store and retrieve data. However, structured databases do not have the capability to handle the unstructured, unpredictable and complex documents that are typical of both product descriptions and other related information such as product reviews. The issues then become those of extracting the salient information from such unstructured text sources, or of structuring such documents with meaningful, machine-processable, annotation—all complex problems.

Heterogeneously modelled data

In the web environment, users are further confounded by their lack of knowledge about how any system is structured. Interactions must be forgiving—capable of handling any form of query and responding with good matches to vague questions. Whilst most of today's e-marketplaces stop at search, some of the more advanced also categorise their product offerings so that roughly similar products are retrieved together. While categorisation improves the chances that the right product will be retrieved, it is seldom an easy process to automate. In reality, most categorised systems are a cobbled together collection of automatic features and manual labour. Semantically mapping between and among catalogues is a complex process; each vendor describes his wares in different terms. Equivalent products must be determined, and judging equivalence is usually beyond most automatic systems. Also, manual systems are not scalable; it is not possible to keep constantly changing systems up to date in a timely manner using manual labour.

Cultural and individual personalisation

Today's increasingly global market, especially as it is manifested in the Web, demands support for multilingual communication, information access and transactions. It is expected that the demographics of the Web will dramatically shift away from the US and away from English over the next few years. An e-commerce application that can only converse in English, is of no use to a native Spanish speaker. Multilingual support can be important to employees and business partners as well as potential customers. So, too, for someone searching a catalogue for a product or service to purchase. But language is only the "tip of the iceberg"; these systems must also be localised to the modes of data presentation (e.g., currency, dimensions), product and business regulations, business practices and cultural norms of the user's own country.

Telecoms infrastructure media

E-marketplaces also need to meet the needs of an increasingly mobile user population. The mobile user wants to purchase products on the run. This means that a wireless device such as a PDA or WAP telephone is the interaction venue of choice. Large text files or graphics are useless in this environment. Voice input and output would greatly enhance such mobile and wireless transactions. Speech recognition and speech generation are, however, only functionally useful in very limited ways. One of these limitations is the lack of strategic interactive knowledge, that is, the ability to hold dialogs with users interested in fulfilling some goal, be it access to information or to exercise a transaction for some desired product.

All these problems with current e-commerce/-business require solutions. One of the most talked about solutions to many of the information search and integration issues is ontologies. Ontologies promise a lot; a shared and common understanding of some domain between application systems, and between them and people. In terms of human-computer integration, conversational systems are also being talked about as the ideal way automatic e-business should be executed. It is to these conversational systems, and how ontologies will aid in providing these systems with various functionalities, which we turn our attention to next.

3 An introduction to an e-business solution: ontologies meet conversational systems

3.1 The requirements

All business transactions consist of a complex set of interactions among vendors, buyers, and the information that exists surrounding the product or service—evaluations of products, notions of reliability, stylishness, appropriate price, service, etc. The marketplace is heavily dependent on such information, and is, in fact, a specialised information system. It matches vendors' offerings with buyers' needs. It informs the consumer so that he understands his options. This process is ideally carried out through a dialogue that helps both sides adjust their offerings and requirements. That is, the technology has to be an adaptive system based on natural language understanding. It has to be able to discern the meaning and even the underlying intent of a question. It also has to adapt during the course of a negotiation with a customer, as the customer's requirements narrow or change in reaction to the availability or suitability of products in the marketplace. It has to have multilingual input and output as well as the capability for speech recognition and generation to make it appealing in an increasingly global marketplace.

Furthermore, any such system must be integrated and robust enough to sit on top of the complex functions that constitute modern information-mediated transactions, that is: context-dependent search; multilingual; product retrieval across multiple suppliers, each with different descriptions of analogous products; continuous updating of products, prices and other information; tracking of orders; feedback to suppliers on successful/unsuccessful products and the products that are requested but are unavailable; fast response time: fast updating; sticky features to keep people on the site; and an improvement on current systems, both in quality of service and in cost savings.

3.2 The solution

SemanticEdge has developed a state of the art multilingual natural language (text and voice) dialog system capable of handling dialogs with humans wanting to access information, for example, to purchase products and services. The technology extends naturally to Customer Relations Management (CRM) and other e-business functions. This technology depends on several distinct technology areas within Artificial Intelligence: natural language processing, including deep language processing and statistical analyses; machine learning, including inductive learning; speech recognition; automated dialog generation, both user and content specific; and knowledge representation and ontologies.

The system mediates between humans and information. That is, it mediates between an information space and a human's conceptualisation of that information space; for example, between a product space and a customer's conceptualisation of that product space, and how they will consequently go about searching and querying that product space. Users hold negotiations with the system, which is mediating access to the product spaces, and it will ask questions of them. This requires the system to have the ability to guide those dialogs according to a representation of that product space. This ability to a large extent is supported by ontologies. Not only does the technology model products objectively, as might be done with a sophisticated database system, but we also model subjective quality judgements that consumers tend to use when conceptualising the product space before them. These subjective, ad hoc, categories gives the system the ability to communicate to the consumer in a human friendly way, in a way that is, in terms of the ontological commitments made by the system, similar to those of the typical customer or user. These human-oriented aspects are further enhanced by other technologies within the system, such as user models of consumer reaction to the dialog process as it happens.

3.3 Ontologies' role in the solution

As noted, many of these information search, integration, and modelling, functionalities are supported by ontologies. SemanticEdge uses both cutting-edge ontology editors and environments from third parties and its own ontology technology to build both domain and user ontologies. By domain ontology we mean both objective and subjective ontologies. Objective ontologies model the standard product descriptions typically released by product manufactures. These include attributes such as weight, price, product features, and so on. By subjective ontologies we mean those attributes of a product which are somehow generated by consumers' (common) conceptualisation of the product, such as a "fast", "noisy", "family" printer, and so on. These subjective quality attributes and ad hoc categories are typically the main communication vocabulary that people use to conceptualise the product/information space and consequently the terminology they use in dialogs. Also, user ontologies play the role of modelling different types of user and mapping those different consumers, and their attendant different requirements, to different dialog strategies and ultimately to different products and services. Since these knowledge bases are concept based, they are language independent—an important feature in a multilingual system that must retrieve in any language; new languages need only to be mapped to concepts, not translated term by term.

This ontology building effort is large in scale and complexity and its management is non-trivial. It requires extensive automation and support from intelligent tools. It is our experience that at present, whilst competent inference engines and editors are available from third parties, large parts of what might be called a comprehensive ontology engineering workbench or environment are missing; worse, it is those parts which would allow for the efficient application of ontology technology in enterprises which are missing. Consequently, SemanticEdge has developed its own set of ontology building technologies. These technologies allow for the capture of large amounts of instance level information from unstructured through to highly structured sources; that is, to populate an already existing intentional structure. SemanticEdge is also developing ontology learning technology which will help in the burden of building the intentional structure; that is, acquiring the concepts and conceptual relations (e.g., subclass relations) from free text as well as from more structured sources. The technology being developed uses a mixture of adaptive pattern recognition, inductive learning, several machine learning technologies, and an extensive set of heuristics. SemanticEdge is also currently involved in several academic projects, such as OntoWeb [4], and industry consortia all with the common goal of working to develop more advanced shared ontology technologies, to allow for their efficient and efficacious application in enterprises.

What are ontologies anyway?

The standard working definition to which most ontologists refer to is that of Gruber [1]. Others have defined ontologies in largely similar ways [2]. The common characteristics an ontology is supposed to have are that it should be a *formal, explicit specification of a shared conceptualisation*. That is, ontologies (in the sense propagated here, and the one recognised in artificial intelligence) should be a conceptualisation of some phenomena. How complex this conceptualisation is up to the conceptualisor, and with regard to the formal specification, also depends on how expressive the specification language is (humans having very expressive specification languages being able to have very complex ontologies of their world). Formality means that that we may be able to automatically map and reason with our specification, typically with the aim of having machines reason with the various ontological knowledge. The notion of the ontology being shared means that the specifications made are to some extent common throughout some group of members. This conceptualisation may have been arrived at by common consent or not, but the members of the group are said to have *committed* themselves to the ontology.

The elements of ontologies

Ontology languages are the formalising structure which represent the domain/universe of discourse or world we are interested in. To accomplish this, editors are typically used which bypass many of the textual characteristics of specification languages and provide more perspicuous and efficient means for developing, maintaining, and modifying ontologies. Taking the ontology editor metaphor to another higher level, we end up with an ontology environment metaphor, where various potentially time saving and quality increasing features are to be had (e.g., modularisation, versioning, and reuse mechanisms in general, verification and validation). Once the ontology has been developed, we can think of performing reasoning with the ontology, requiring inference engines; which may also be part of the development environment, thereby supporting intelligent editing, and perhaps enabling debugging and tuning of the ontology from reasoning efficiency and competency viewpoints. The interested reader might want to read [2] for a detailed exposition of ontology technology.

This brief overview of how ontology technology fits into SemanticEdge's conversational system motivates some important issues in the application of ontology technology and methodology. One of the most important is that it requires the ability to manage large ontologies. This scalability requirement motivates several other ontology engineering issues, including: acquisition; visualisation; modularization and versioning; reasoning transparency; multitasking; competency; and methodology. In the remainder of this paper, we will explore all of these issues more fully.

4 Issues in enterprise-standard ontology environments

Many of the above elements are now discussed, but some are not. Notably, languages are not discussed, as these are, arguably, where academia can be and has been of greatest input; these languages are complex, the semantics issues complex, and not many people outside academia are competent to go around designing semantically well designed ontology languages. The other elements are largely those of engineering application; they require manpower, capital, and management.

4.1 Acquisition

The acquisition of ontological knowledge and the instantiation of such ontologies is a prime issue in the building of large ontologies. There are several distinct knowledge acquisition phases in ontology building. There is the initial acquisition phase, where the structure and terms of the domain are acquired and represented; this is also part of the conceptualisation phase mentioned before. There is then the acquisition of rules and axioms of the domain; this is also part of conceptualisation. Lastly there is the acquisition of instances, the instantiation of the ontology with facts to be reasoned with. The aim in all these phases of acquisition should be to acquire what needs to be acquired quickly, efficiently, correctly, and with as little effort as possible.

The different phases impose different requirements, and the goal of making them automatic are more realisable in certain phases than in others. The least amenable to our wishes are the first and second, that of acquiring conceptual knowledge of the terms, structure, rules and axioms, of the domain; that is, the conceptualisation of a domain is a skilled process. It is in the last phase, the acquisition of instances, that most can be accomplished, though this a non-trivial task when the source data is informal, such as text in web pages.

The automated acquisition of such instances is something to which SemanticEdge has invested considerable resources in. Proprietary information extraction technology has been developed to support the acquisition of information from unstructured to more structured sources. A screenshot of the GUI can be seen in Figure 1. SemanticEdge is among a growing number of companies that offer specialised technology for carrying out this information extraction task. A number of trainable and self-learning Artificial Intelligence (AI) technologies are encapsulated inside a single Information Extraction Engine. These AI technologies can be configured to map any number of different product catalogue formats onto a single intermediate, predefined product schema. From this schema, information can be exported into one or more formalised representations (including ontology languages).

Export involves two basic steps:

- **Normalisation**: This can simply involve mapping one of a number of synonyms for a given piece of product information onto a single predefined symbol. It can also involve more complex normalisation rules such as converting numeric attribute values that can be given in one of a number of units onto a single standard unit.

- **Generation of Export Syntax**: Through the attachment of formatting rules to the intermediate product schema, high flexibility in the export format can be achieved, and as noted, the information can be output to ontology languages, such as, for example, F-logic.

One further requirement is that the acquisition technology should be useable by non-knowledge engineers. This again is a distinguishing feature between the three phases, with again, the first two being more knowledge engineer heavy, whereas the third, if done properly, can be accomplished by non-experts in conceptual modelling. Of course, there is more to acquisition, especially when considering the first two phases where the process is more one conceptualisation. Here the progressing conceptualisation and formalisation would be partly helped by perspicuous visualisation, and it is that which we will now look at in more detail.

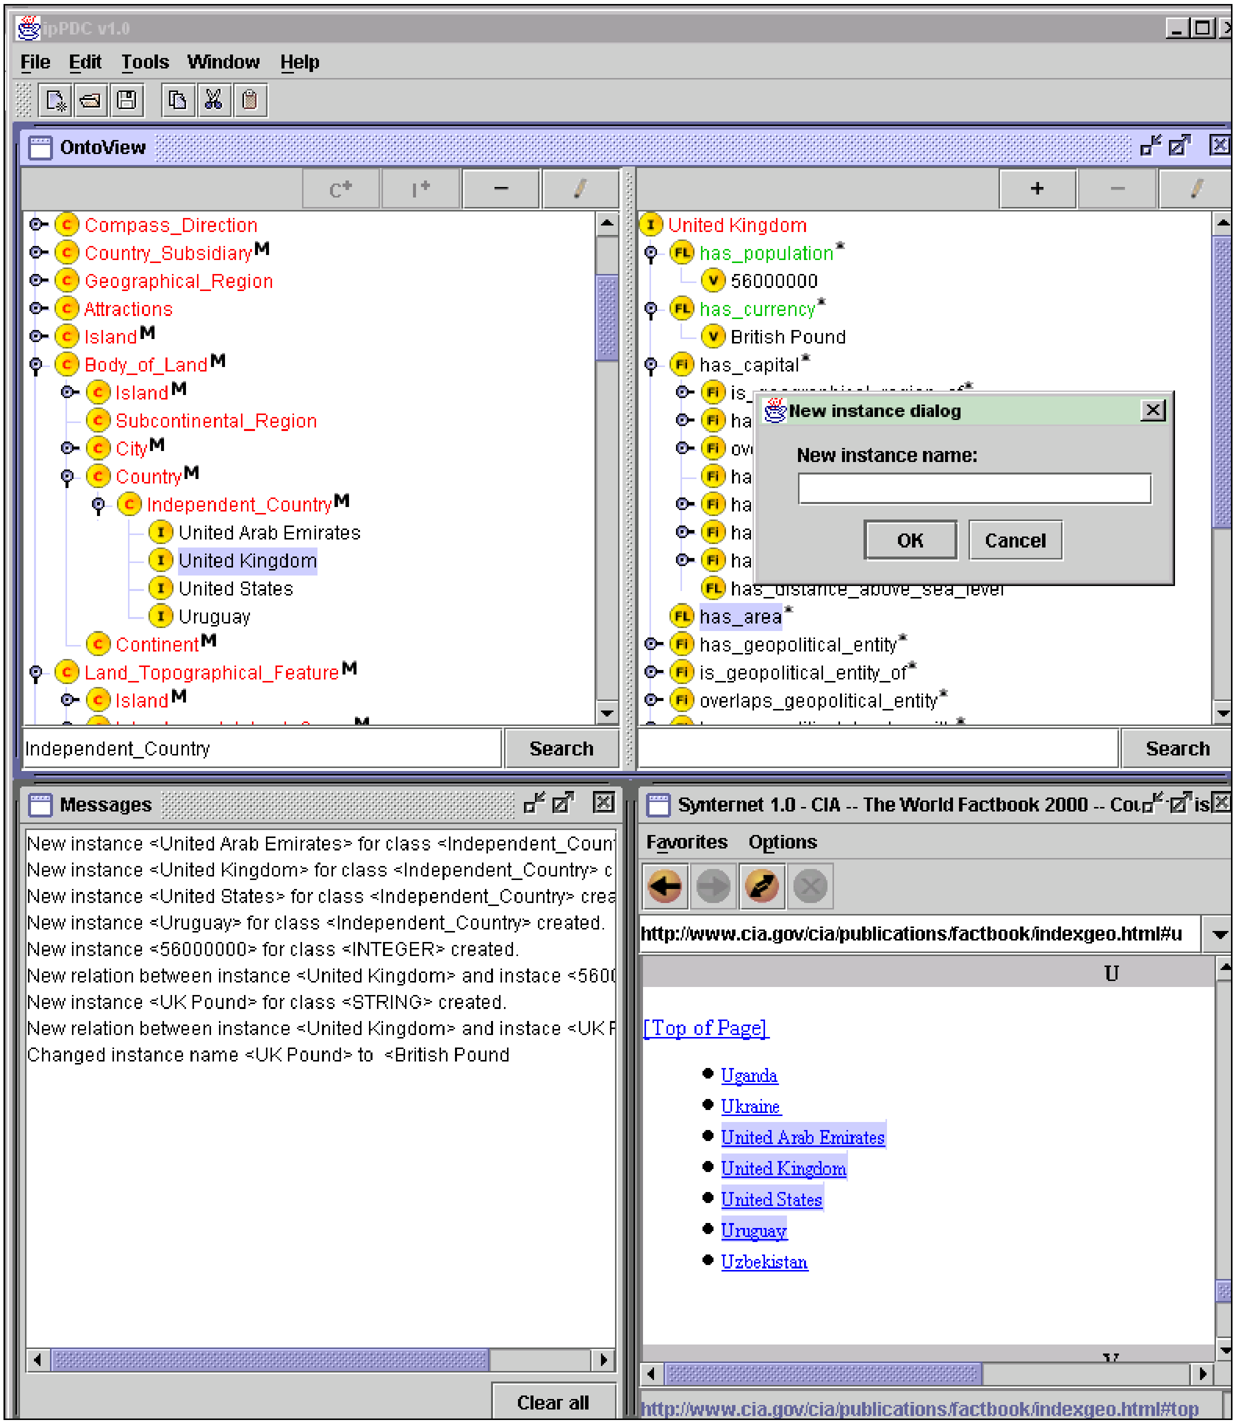

Figure 1: SemanticEdge has developed the sePDC to enable the acquisition of product instances. This extensional information has to conform to the imported ontology, and a number of ontology formats, including F-logic, can be accommodated. Here we are capturing some new instances of the Country concept from the CIA World Fact-book.

4.2 Visualisation

Visualisation has dogged conceptual modelling for years, with the taxonomic metaphor often blinding all to any other conceptualisation of how ontological knowledge might be presented. While a high quality implementation of a taxonomic presentation is both important and useful, it is not the whole story. Certainly, for acquiring taxonomic subclass relations, it is probably a fine paradigm to use. The issue become more complex when one considers that the visualisation should somehow convey other, more complex, aspects of the conceptualisation. This visualisation of a conceptualisation has to include not only the terms but also the axioms in relation to the terms they relate to, and so on; although the visualisation of axioms themselves is a tricky issue. Further, editors which somehow support the user in defining axioms must be welcomed.

Visualisation should aid in the conceptualisation of the domain, as well as being some kind of first-pass validation of the ontology's commitments.

4.3 Modularization and versioning

One issue which we have noted as being of importance to the ontology engineering issue is a very important lesson learnt in software engineering, that is, of reuse and modularization. The benefits of reuse and modularization have been often repeated, and we will do so here again in the light of ontologies.

Reuse is a good thing. It reduces effort through not requiring basic (or, nowadays, more complex) components to be built from scratch, increases quality through the reuse of quality components, and thereby reduces effort and costs. But reuse also brings problems of its own making such as building the most reusable components possible and finding and integrating the most up to date versions of these components into a working and consistent whole. In the ontology world, the merging, alignment, version control, and so on, of ontologies, is only now beginning to make it into tools, and has a long way to go. These are, however, going to be crucial issues which will become very important in years to come, as more and more ontologies are built and available on the WWW, where different components might be inconsistent with each other, and out of date versions abound. We can imagine a time when we will be able to select from a library of well designed, up to date, mutually consistent components which can be easily, even automatically, integrated, all done in some kind of developer studio environment.

4.4 Reasoning transparency

The are several issues regarding reasoning which require close attention. Reasoning in the small and in the large are totally different propositions. Issues such as inference profiling, debugging, inference efficiency concerns, how the modelling affects the reasoning to be performed, what the competency of the ontology implies for the completeness of the reasoning, logically erroneous axioms causing problems (circular axioms), and so on, all require attention, especially when the ontology is scaled up. For instance, let us assume we have a slow query. Is this slowness an implementation problem such as a bug in the reasoner, is it a problem of having too complex a model, is the model okay but the reasoner is providing us with too much reasoning, for which we have no use as regards the competency we have decided the ontology should support, is the reasoner sound and complete but not up to the job?, or is it simply a slow machine on which we are running our system? To answer these questions requires the ability to observe the reasoner in action and to have access to some kind of statistics and summary information. This aspect of ontologies is crucial, as the model-axiom-reasoner interaction is a non-trivial one and the space of possible interactions is immense and totally unpredictable in terms of the efficiency of the concerted artefact. This is something, that, for instance, databases do not suffer from to anything like the same extent. Indeed, this issue goes to the heart of much knowledge representation research in that it encompasses completeness, soundness, tractability, and so on.

It is our opinion that reasoners cannot be black boxes in to which no one may look. For the engineering of large ontologies, it is necessary to be able to at least appreciate the problems and to try and solve them however we can or is allowed by the reasoner paradigm in question. Altering the completeness and soundness characteristics is probably an extreme way of solving any potential problem, but at the very least, one should be aware of what inferences are being computed and why, so that, for example, over burdensome axioms may be modified if the competency specification allows it. These issues of reasoning efficiency become even more important when considering that these reasoners may very well be on-line knowledge servers to which multiple users (e.g., hundreds) may be accessing. It is to the issue of multitasking and multiuser access that we turn to next.

4.5 Multitasking

If ontologies, the models and reasoners, are to become trusted on-line servers of knowledge, then they must have certain features developed over the years in the database community. For example, the capability to handle multiuser access would seem to be the minimum. Multiuser access may be querying, browsing, or editing. There is much demand for all at SemanticEdge, where ontological tools and the conceptual models they support are required to be browsed by several people, who have to keep their work somehow consistent with the ontological commitments already made. And, indeed, many people have concepts which they might want to add, if suitably qualified to do so. And there are other issues, such as stability and reliability. This implies some professional engineering support and development to take ontologies to the next level of enterprise integration.

4.6 Competency

The notion of competency is one of the most important issues in conceptual modelling. It is important in that when we model a domain, we expect the resultant concerted artefact of model, axioms and inference engine to be competent with regards to the queries we should wish to ask of it. It is in the interaction of the these three components of any ontology that the complexity arises, and we require someway of assessing whether or not we have been successful in building the ontology artefact which satisfies our requirements. For example large amounts of global terminological assertions such as partitions and so on can have quite large effects on the kind of deductions possible (as well as reasoning efficiency, see earlier section on reasoning transparency).

Deciding when an ontology should be declared competent is not easy. One solution is to have a test harness where several representative and/or complex queries can be entered and run on the ontology. If these cases give satisfactory responses then one might conclude that the ontology is competent and leave it at that. Alternatively, one could have a comprehensive set of queries somehow automatically generated which comprehensively exercises the ontology's competency (one may even want such functionality so as to cache the results of our ontology if we decide that its performance is not good enough to have as part of an online, live, system). The returned answers would then have to be somehow assessed and a judgement made on the ontology's competency.

4.7 Methodology and Ontology

One of the most challenging aspects of ontology work is developing new ontology structures, capable of representing what is intended. This is typically not easy when one wanders, even slightly, from the well-trodden path of EER modelling commonly practised in the database (and OO software engineering) world. Representing other kinds of world phenomena, well axiomatised, is not a trivial task. Upper level ontologies, where a solid, well thought out, conceptual structure, offers the benefits of providing methodological support for modelling and Ontology. For example, part-whole knowledge is one of the most common ontological structures humans use to think of the world, and it is also one of the more complex representation and reasoning paradigms to get right. Methodological support in this area, as well as the definition of various well-conceptualised upper-level concepts, is crucial for the development of ontologies that will talk to each other. Without such support, it is quite possible that ontologies unable to be integrated will find their way into various resources, and significantly harm the information integration dream that many believe ontologies offer. After all, the word Ontology, as per our working definition given earlier, apparently for many people has a notion of a common and shared meaning of terms, and it therefore seems fairly reasonable to hope for at least a common methodological upper level Ontology which people can use and extend.

5 Conclusion

E-marketplaces accentuate several dimensions of the product buying process. They bring incomprehensible scale, and thus the consequent problem of enabling a customer to, indeed, comprehend and interact with this space. This large space brings problems of heterogeneity (how to find all similar products), search (how to find your products amongst the hundreds or thousands available), choice (is there too much choice now available for the average customer), and optimality (is the customers choice the optimal one in the space of his product options). These problems of scale are not the only problems in allowing human customers to interface optimally with their chosen product space; there are conceptual problems independent of scale. These problems lie in allowing a human customer to conceptualise and communicate in as natural a way as possible with the buying process. This consequently means depending on complex natural-language and knowledge-based systems (including ontologies) supporting a dialog or negotiation. SemanticEdge is currently developing such cutting-edge technology (including ontology technology) which will allow it to so facilitate this natural transactional process between consumers and the products they wish to buy.

However, it is our opinion that ontology engineering has a long way to go before becoming truly enterprise-standard. It must embrace many of the engineering paradigms of object-oriented (OO) software engineering and database engineering that have become de facto over the last few years. Owing to the added complexity of ontologies in that they have inference engines, knowledge models, axioms and rules, all interacting in a non-trivial way, there are extra problems and complexity to be managed. At present, this additional complexity is not managed at all and is largely hidden from the user. Only symptoms such as slow queries, strange deductions, hanging inference engines and incompetent answers surface to alert the ontology engineer to problems. For the successful engineering of ontologies, it is essential that this complexity be better managed and thoughtfully exposed, so that when problems do occur there are recourses to take and information to be examined. This along with the substantial beefing-up of ontology environments will allow them to become the next essential components in future distributed ontology-enabled information systems (e.g., resources on the WWW). Such information environments are beginning to be seriously discussed by many as the next stage in the evolution of, for example, the WWW—the Semantic Web—with such initiatives as DAML [3], where ontology languages such as DAML+OIL are being designed to support this, reinforcing this belief.

6 Acknowledgements

We acknowledge that part of this paper was based on the SemanticEdge white paper as written by Susan Feldman of IDC.

7 References

[1] T. R. Gruber: A Translation Approach to Portable Ontology Specifications, *Knowledge Acquisition*, 5:199—220, 1993.

[2] D. Fensel: Ontologies: Silver Bullet for Knowledge Management and Electronic Commerce, Springer-Verlag, Berlin, 2001.

[3] DAML: http://www.daml.org

[4] OntoWeb: http://www.ontoweb.org