Benchmarking Knowledge-based Context Delivery Systems

John Keeney, David Lewis, Declan O'Sullivan
Knowledge & Data Engineering Group and CTVR,
Department of Computer Science,
Trinity College Dublin,
Dublin, Ireland.
{ John.Keeney | Dave.Lewis | Declan.OSullivan }@cs.tcd.ie

Abstract

As context-aware systems become more widespread and mobile there is an increasing need for a common distributed event platform for gathering context information and delivering to context-aware applications. The likely heterogeneity across the body of context information can be addressed using runtime reasoning over ontology-based context models. However, existing knowledge-based reasoning is not typically optimised for real-time operation so its inclusion in any context delivery platform needs to be carefully evaluated from a performance perspective. In this paper we propose a benchmark for knowledge-based context delivery platforms and in particular examine suitable knowledge benchmarks for assessing the ability of platforms to deal with semantic interoperability.

1. Introduction

This paper is motivated by the challenge of establishing a common, scalable service for delivering context information to context-aware application central to pervasive computing. Pervasive computing promises to make available a vast volume of context messages from environmental sensors embedded in the fabric of everyday life reporting on user location, sound levels and temperature changes to name but a few. Any scalable context delivery system must ensure therefore the accurate delivery of context events to the consumers that require them. However, the wide range of sensors and sensed information, and the mobility of consuming clients, will present a level of heterogeneity that prevents consumers accurately forming queries to match possibly unknown forms of relevant context events.

In addition to the pervasive computing domain, a number of other application areas would benefit from scalable context delivery system. For example, autonomic systems gather knowledge of their operational state combined with operational context to self-manage, i.e. to self-configure, self-heal, self-optimise and self-protect. A key platform technology for the practical transition to autonomic management is a service that enables the efficient delivery of distributed operational knowledge to, and only to, nodes that have expressed an interest in that knowledge for use as the context information needed to make appropriate adaptation to their behaviour. The trends for pushing more operational intelligence towards each self managing element in order to achieve more context-awareness in self-managing behaviour often requires elements to gather knowledge without necessarily being explicitly aware of all of the potential distributed sources of that data.

Thus a knowledge-based approach to modelling and transmitting context has been proposed by a number of researchers. Typically these researchers are adopting ontology-based semantics [wang][masuoka] [belecheanu], spurred by the standardisation of ontology languages by the Semantic Web initiative at the World Wide Web Consortium (W3C) [berners-lee]. This approach promises loose semantic coupling between applications, which is vital as new waves of applications increasingly rely on using the information and services offered by existing heterogeneous distributed applications. To date, solutions have been only implemented in centralised or enterprise-scale middleware [stevenson][wang][chen], and not to a scale suitable for deployment over the Internet.

Meeting this challenge demands we address both the extreme heterogeneity and rapid evolution of context-aware applications and context information, in combination with the need for high event throughput, low-latency of event messages between large, volatile populations of event service clients.

In this paper we do not present a new solution to knowledge-based context delivery, but instead we examine a suitable benchmark for assessing and comparing such systems. This is based on our own experience and experimentation with various aspects of knowledge-based context delivery, specifically the routing of context information, semantic interoperability and caching [lewis05][keeney][power]. We assume that any scalable platform that can offer a common service at Internet-scales will be based on a decentralised asynchronous message-passing mechanism that supports the efficient notification of relevant changes of context to interested applications.

2. Context-Delivery Mechanisms

Publish-subscribe (Pub-Sub) event systems [meier] might be considered as the basis for such an event service as they avoid close coupling between producers of events and one or more event consumers that have expressed an interest in an event type. Currently, publish-subscribe systems, e.g. IBM MQSeries, are used widely as a communication bus to flexibly integrate business functions. However, such Pub-Sub systems require agreements on message types between the developers of producer and consumer applications. This places severe restrictions on the heterogeneity and dynamism of client applications.

In [crowcroft], Crowcroft et. al. envisage pervasive computing applications globally generating 10^9 events per second, which must be accurately routed to event consumer groups of widely varying size taken from a population of 10^8. The solution they propose is a Pub-Sub system that filters events based on matching client subscriptions to message attributes rather than the full message type, a technique known as content based networking. Content-Based Networks (CBN) thus facilitates still looser coupling between producer and consumer applications than Pub-Sub. Several CBN solutions and prototypes exist, e.g. [carzaniga01][segall][pietzuch][chand][strom], though they have not been demonstrated to operate anywhere near the aspirations for global pervasive computing set out above, which is approximately the equivalent to the packet per second throughput of today's Internet.

Widespread CBN deployments have been slow to emerge partly due to the difficulty in reaching a general compromise between the expressiveness of event types and subscription filters and the need both to match these efficiently at CBN nodes and to efficiently maintain forwarding tables by aggregating new subscriptions with any existing ones that cover a superset of matching messages [carzeniga99]. As a result current CBNs only support a very limited range of data types and operators that can be used in matching consumer subscriptions to message attributes, typically: Strings, Integers, Booleans and associated equality, greater than, less than, and regular expression matches on strings. This falls well short of supporting the heterogeneity and flexibility that clients of a global pervasive computing event service require. We assert, therefore, that the target for CBN expressiveness should be the subject-predicate-object structure of ontological knowledge representations, standardised as the W3C's Resource Description Framework (RDF) [rdf]. Thereby, subscription queries can contain arbitrary logic based on any binary predicate defined for message attributes. A CBN based on triple-structure messages and corresponding RDF queries is far more flexible, open and reusable to new applications. We call such a semantic-based CBN a Knowledge-Based Network (KBN).

3. Knowledge-Based Networking

Such a KBN presents major new and unconsidered challenges in how the messages and subscriptions rendezvous in the network, thereby establishing routes that event messages follow to subscribing clients with matching subscriptions. For instance, scalability of a KBN to Internet proportions requires a routing control plane that minimises both the size of routing state held in KBN nodes and the overhead of ontological reasoning in nodes. At the same time this control pane must itself auto-configure in response to topology changes, exhibit robustness to network failures and maximise reachability. The scalability of the routing control plane in the Internet relies, through the use of the Border Gateway Protocol [rekhter], on the natural administrative partitioning of the Internet into Internet Service Provider (ISPs) domains. However, in a KBN the efficient partitioning of routing space must be based on groupings related to the semantics of message contents rather than grouping within the hierarchies of network addresses.

Clearly, any software-based event forwarding algorithm will struggle to match the hardware optimised performance of packet forwarding in IP routers. Basing the forwarding algorithm on today's ontological reasoners would seem likely to incur a heavy computational load. However, a number of strategies can be adopted to minimise this load. For instance optimised context ontology reasoners may be developed for KBN routing purposes, or intelligent clustering in KBN routing algorithms that are cognisant of the performance profiles of existing reasoners and of the semantics being exchanged by client applications could be explored. Equally, the caching of context change events could produce efficiencies in a context delivery system with large numbers of subscribers. In addition, to fully exploit the potential benefits of ontologies in handling heterogeneity of context information, semantic interoperability functions driven by ontology-based mapping may be deployed.

To assess and compare different schemes and combinations of techniques it is necessary to establish a benchmark for the knowledge-based delivery of context-information. Measurements from these benchmarks may also benefit the solutions themselves, especially when addressing semantic clustering and loading across a large, decentralised network of KBN nodes forwarding context events.

KBNs represent a significant departure from the mainstream of knowledge-based systems research that has rarely extended beyond desktop applications and enterprise scope. Though the Semantic Web enables knowledge exchange at a global scale, communication is via HTTP or enterprise-scale middleware, thereby insulating the systems from the issues of Internet scaling. By addressing asynchronous messaging over a highly decentralised network we attempt to reconcile Internet engineering values and knowledge engineering solutions, thereby exploiting the new efficiencies yielded by clustering KBN nodes based on semantic measures. Thus, any useful benchmarking system for such a context delivery system needs to combine benchmarking techniques from both the networking community and the knowledge engineering community.

4. Benchmarking Contextual Knowledge Delivery

A benchmark for Pub-Sub systems is proposed in [carzaniga02], which we extend here to cover Knowledge Based Networking. It is proposed that for common Pub-Sub applications such as news broadcast, networked games, application management etc., that the application dependent behaviour of the system under observed traffic loads should form the basis of a benchmark. However, to test for uncommon workloads or the non-application dependent behaviours typical in a context delivery network, a synthetic benchmark is also presented. In addition to comparison based on the set of differentiating features of different systems, the workload behaviour of the system over a topology created using an extension of a random graph generation model are to be characterised using a number of key metrics. These metrics include: the publication rate, the parameters used to automatically generate subscriptions, the subscription / un-subscription cycle durations, and the parameters used to automatically generate the publications.

For the benchmark presented in this paper the synthetic benchmark proposal presented in [carzaniga02] is extended. The benchmark presented by Carzaniga et. al. is specifically designed for CBNs, where the set of operations and message and subscription attributes are taken from a limited set of predefined types and values. The benchmark extends this in a manner whereby the information used in automatically generated subscriptions and publications is drawn from ontological knowledge bases rather than from a given set of attribute types and values. We specifically focus on the semantic nature of the knowledge being distributed, the key difference between a CBN and a KBN. In this extended benchmark the sets of publishers and subscribers are again initially defined according to a randomly generated topology. More complicated topologies based on clustering can also be added.

- *Publication rate:* as in the original benchmark this remains a Poisson distribution of publications generated over time.
- *Active / inactive subscription cycle durations:* this remains as a timer / counter to define the intervals during which a subscriber has no subscriptions (inactive) or one or more subscriptions (active).
- *Subscribers' / publishers' ontologies:* the ontologies used to model the knowledge base of the subscribers and producers may be based on existing ontologies for specific application domains, or for more general cases, the ontologies may be automatically created. Initially, the ontologies should be created with a Zipf-like distribution of ontological terms, given that recent analysis of the usage of ontological terms in semantic web documents have shown Zipf-like distribution characteristics [swoogle].
- *Publication generation parameters:*
    - *Number of fields in publication:* the number of fields in a generated publication remains random according to a Poisson distribution.
    - *Names of attributes in publication:* instead of from a dictionary, the names of the attributes will be randomly drawn from the publisher's ontology. The prevalence of these names in each publication will need to be defined according to a domain specific probability function as before.
    - *Type of each attribute:* for ontologically based publications, the set of types is extended beyond

the simple data types described in section 2 to include those supported by ontologies: class, property, and instance. Each generated attribute must be randomly assigned one type according to a domain specific probability function.

  - *Value space for each attribute:* a set of values will be defined for each possible attribute, with the prevalence of each value defined according to a domain specific probability function in a manner similar to the original benchmark.

- *Subscription generation parameters:*
  - *Number of subscriptions per subscriber:* this defines the distribution of active subscriptions across subscribers since each subscriber can have more than one active subscription. The length each subscription remains active is described by the active / inactive subscription cycle durations described above.
  - *Number of filters per subscription:* this remains randomly distributed, again according to a Poisson distribution.
  - *Names of attributes used in each filter:* instead of from a dictionary, the names of the attributes used in the subscription filters will be randomly drawn from both the ontologies used by the publishers (if known) and from the subscriber's own ontological knowledge base. The percentage drawn from the publishers' ontologies must also be defined. The prevalence of each attribute name used will again be defined according to a domain specific probability function.
  - *Type of each attribute used in the filter:* the type of each attribute used in the filter must also be defined in a manner similar to that of the attribute types in publications. For attributes drawn from the publishers' ontologies the attribute types should remain the same. The prevalence of each type should again be defined according to a domain specific probability function.
  - *Attribute values used in filters:* the value space of each attribute used in filters must be defined as above with the attribute values used in publications. The overlap of these value spaces for each attribute that exists in both the subscribers' and publishers' ontologies must also be defined.
  - *Operators used in filters:* in the original benchmark the set of operators supported by the subscription filter mechanism was limited to the simple comparison and equivalence ones described in section 2, with the set of operators fixed by the implementation of the system. This needs to be extended to include the operators to describe subsumption, reverse subsumption, and equivalence of ontological classes, properties and instances. In a fully ontologically enabled system, every property can be considered an operator on classes or instances. This greatly extends the expressiveness of the subscription filter mechanism but complicates the routing and covering algorithms used by the KBN. The prevalence of each operator must also be defined according to which attributes can support which operators, and also according to domain specific requirements.

These parameters will also need to be expanded to take into account the size, complexity, expressiveness, bushiness, and information content [blanchard] of the publishers' and subscribers' ontological knowledge base. Each of these ontology features will also affect the performance of any reasoner used.

5. Applying the Benchmark and Further Work

As discussed, in order to undertake comparable analysis of the KBN systems, there will be a requirement to quickly auto-generate large ontologically based data sets, publications, and subscriptions in line with the different application domains and loading patterns. The auto-generation capability should initially be built upon existing work for the automatic generation of ontologically based data [guo05][patel-schneider]. For example, the Lehigh University Benchmark package developed in Java [guo04], currently consists of a university domain ontology, with auto-generated synthetic data and a set of test queries upon that test data. This can be extended to use the presented benchmark to allow for the generation of test data and test events/subscriptions based on other domain ontologies.

It will also be necessary to establish a validation testbed for the practical evaluation of the KBN function. This will provide a platform for evaluating different RDF matching functions in the KBN routers, including the performance of RDF reasoners with differing profiles for relevant prevalence of classes, instances, predicates, and class hierarchies, including the time to load new semantic models, and the use of pre-calculated vs. run-time inference. This will provide base line reasoning performance figures as well as addressing the practical issues of running multiple routing strategies over a KBN. A large number of ontology reasoners are available, including: KAON2 [motik], Pellet [parsia], Racer [haarslev], FaCT [tsarkov], F-OWL [zou]. Any choice of reasoner must be based on examination of performance evaluations in the literature, such as [pan][guo04][pelletperf][motik], as well as separate benchmarking. These evaluations must also be compared to the performance characteristics of domain specific reasoners, or existing reasoners cut-down to give reduced but sufficient results in return for enhanced performance. This trade-off of reasoning performance versus expressiveness and accuracy of the model after the inference cycle is of particular importance where such reasoning may be required for efficient and correct routing in the network. The differing performance characteristics of different reasoners under different conditions, such as, the impact of the ratio of concepts to relationships or of subsumption relationships to user defined predicates, must also be evaluated. The performance of different reasoners, and the reasoning load, will also change in a non-linear fashion depending on the size and expressiveness of the ontologies used and the level of ontology language used (e.g. OWL-Lite vs. OWL-DL) [pan][guo04][guo05][pelletperf][motik]. Of particular importance is the amount of reasoning that can be performed at ontology load time versus when the first or subsequent queries are submitted to the ontology. This becomes particularly important if ontologies are added or removed dynamically, as would be typical in a network of knowledge producers and consumers where joins and departs from the network will occur.

As described in the benchmark metrics above, the KBN should initially be optimised to cope with a Zipf-like distribution of ontological terms. However, the accuracy of the Zipf distribution model for ontologies in applications domains other than semantic web documents should be considered. In [keeney] the authors have generated ontologies for the network management domain based on the different standard models used to represent network management information. A comparison of these ontologies with ontologies for different application domains, especially context representation (e.g. SOUPA [chen]), should prove beneficial.

To describe the characteristics of different ontologies, a range of semantic distance measures have been proposed, such as [blanchard][rada][jiang][sussna] [kashyap], and initial investigation will be required to determine the most suitable ones for use for different networking characteristics and different application domains. Another key research topic arises when the ontologies used by the subscribers and producers are different but where ontology mappings between them exist [osullivan04][lewis05]. In this case the semantic interoperability of the concepts used in each ontology will lead to complications in describing concepts such as semantic distance and information content of the ontologies.

Initial research has also begun on how nodes in a KBN could be logically clustered to improve the routing performance of the KBN to achieve efficient aggregation of semantic queries. Clustering both increases the scalability of RDF-based routing and supports the deployment of routing schemes tailored to specific application domains, thus allowing a wide range of strategies to co-exist in rendezvousing events advertisements and subscriptions via the KBN control plane. A number of different clustering schemes are being proposed. In addition to clustering based on dynamic subscription and publication loads, nodes with domain specific reasoning capabilities or resources for reasoning can be clustered together so that queries on ontologies requiring such reasoning or dealing with specific domains can be routed directly and efficiently towards the correct cluster. Clustering based on non-functional concerns such as security and trustworthiness are also being investigated [quinn06]. Clustering can also be driven by the semantic distance between knowledge sets of different nodes, so that queries can be initially routed towards nodes that may be able to answer that query, and then routed directly to the correct nodes, as seen in [loser] for example. However, each of these strategies need to be evaluated and compared in an objective manner, ideally using a comprehensive but flexible benchmark such as the one described here.

6. Conclusions

The synthetic benchmark described in this paper, extended from [carzaniga02], is the first such benchmark for knowledge based networking. It will demonstrate the feasibility and utility of a KBN and allow similar systems to be evaluated, compared, and improved in an objective and scientific manner. We believe that such a benchmark is critical to the development of knowledge distribution mechanisms for loosely coupled environments, as typically seen where context must be shared to drive autonomicity. While the promise of such a benchmark is obvious, it is also obvious that many of the additional issues discussed in the previous section must also be addressed. However, we assert that this benchmark is a useful first step.

Acknowledgements

This work was partially funded by Science Foundation Ireland under the CTVR project, and by the Irish Higher Education Authority under the M-Zones programme.

References

[belecheanu] Belecheanu, R., Jawaheer, G., Hoskins, A., McCann, J.A., Payne, T., "Semantic web meetings autonomic ubicomp" in proc of the Workshop on Semantic Web Technology for Mobile and Ubiquitous Applications, Hiroshima, Japan, 2004

[berners-lee] Berners-Lee, T., Hendler, J., Lassila, O., "The Semantic Web", Scientific American, May 2001

[blanchard] Blanchard, E., Harzallah, M., Briand, H., Kuntz, P., "A Typology Of Ontology-Based Semantic Measures", in proc of EMOI-INTEROP'05 workshop, at CAiSE'05, Porto, Portugal, 13-14 Jun 2005

[carzaniga99] Carzaniga, A., Rosenblum, D., Wolf, A.L., "Challenges for Distributed Event Services: Scalability vs. Expressiveness" in proc of Engineering Distributed Objects (EDO '99), ICSE 99 Workshop, Los Angeles CA. May 1999.

[carzaniga01] Carzaniga, A., Rosenblum, D. S., and Wolf, A. L., "The Design and Evaluation of a Wide-Area Event Notification Service", ACM Transactions on Computer Systems, Vol. 19, Issue 3, Aug 2001.

[carzaniga02] Carzaniga, A., Wolf, A. L., "A Benchmark Suite for Distributed Publish/Subscribe Systems", Technical Report CU-CS-927-02, Dept. of Computer Science, University of Colorado. Apr 2002, http://serl.cs.colorado.edu/~carzanig/papers/

[chand] Chand, R., Felber, P.A., "A Scalable Protocol for Content-Based Routing in Overlay Networks", Second IEEE Int'l Symposium on Network Computing and Applications, Apr 2003, Cambridge, MA.

[chen] Chen, H., Perich, F., Finin, T., Joshi, A., "SOUPA: Standard Ontology for Ubiquitous and Pervasive Applications", in proc of Int'l Conf. on Mobile and Ubiquitous Systems: Networking and Services 2004, Aug 22, 2004, Boston, MA.

[crowcroft] Crowcroft, J., Bacon, J., Pietzuch, P., Coulouris, G., Naguib. H., "Channel Islands in a Reflective Ocean: Large-Scale Event Distribution in Heterogeneous Networks", IEEE Communications, Vol 40 No. 9, Sept 2002.

[guo04] Guo Y., Heflin J., Pan Z., "An Evaluation of Knowledge Base Systems for Large OWL Datasets", Technical Report, CSE Dept., Leigh University, 2004.

[guo05] Guo Y., Heflin J., "LUBM: A Benchmark for OWL Knowledge Base Systems", Journal of Web Semantics, Vol. 3 Issue 2, 2005.

[haarslev] Haarslev, V., Moller, R. 2001. "RACER System Description", in proc of IJCAR 2001, LNAI Vol 2083, 701–706. Siena, Italy, Springer.

[jiang] Jiang, J., Conrath, D., "Semantic Similarity based on corpus statistics and lexical taxonomy", in proc of International Conference on Research in Computational Linguistics, 1997.

[kashyap] Presentation, NIST Invitational Workshop on Semantic Distance, Gaithersburg, MD, Nov 2003.

[keeney] Keeney, J., Lewis, D., O'Sullivan, D., Roelens, A., Boran, A., Richardson, R., "Runtime Semantic Interoperability for Gathering Ontology-based Network Context", to appear in Proc of IEEE/IFIP Network Operations and Management Symposium (NOMS 2006), Apr 3-7, 2006, Vancouver, Canada.

[lewis05] Lewis, D., O´Sullivan, D., Power, R., Keeney, J., "Semantic Interoperability for an Autonomic Knowledge Delivery Service", in proc of 2nd IFIP TC6 International Workshop on Autonomic Communication (WAC 2005), Oct. 2005, Vouliagmeni, Athens, Greece.

[loser] Loser, A., Naumann, F., Siberski, W., Nejdl, W., Thaden, U., "Semantic overlay clusters within super-peer networks", in proc of Int'l Workshop on Databases, Information Systems and Peer-to-Peer Computing in Conjunction with the VLDB 2003.

[masuoka] Masuoka, R., Labrou, Yannis, Parsia, B., Sirin, E., "Ontology-Enabled Pervasive Computing Applications", IEEE Intelligent Systems, Sep-Oct 2004, pp 68-72.

[meier] Meier, R., Cahill, V., "Taxonomy of Distributed Event-Based Programming Systems", The Computer Journal, vol 48, no 5, pp 602-626, 2005

[motik] Motik B., Sattler U., "Practical DL Reasoning over Large Aboxes with KAON2", found at: http://www.fzi.de/KCMS/kcms_file.php?action=link&id=580, expected publication 2006.

[osullivan04] O'Sullivan, D., Lewis, D., Wade, V., "Enabling Adaptive Semantic Interoperability for Pervasive Computing", in proc of 2nd International Workshop on Managing Ubiquitous Communications and Services (MUCS), Dublin, Ireland, Dec 13-14, 2004

[parsia] Parsia, B., Sirin, E. 2004., "Pellet: An OWL-DL Reasoner", Poster at ISWC 2004, Hiroshima, Japan, 2004.

[pan] Pan Z., "Benchmarking DL Reasoners Using Realistic Ontologies", in proc of the International workshop on OWL: Experience and Directions (OWL-ED2005). Galway, Ireland. 2005

[patel-schneider] Patel-Schneider, P.F., Sebastiani, R., "A new general method to generate random modal formulae for testing decision procedures", Journal of Artificial Intelligence Research. (JAIR) Vol 18, 2003.

[pelletperf] "Pellet Performance", www.mindswap.org/2003/pellet/performance.shtml

[pietzuch] Pietzuch, P., Bacon, J., "Peer-to-Peer Overlay Broker Networks in an Event-Based Middleware". 2nd International Workshop on Distributed Event-Based Systems (DEBS'03) at SIGMOD/PODS 2003, San Diego, CA, Jun 2003

[power] R. Power, D. O'Sullivan, O. Conlan, D. Lewis, V. Wade, "Resolving Queries in a Heterogeneous Context Rich Environment", in proc of the Information and Telecommunication Technologies conference, 20-21st Oct 2004, Limerick, Ireland.

[quinn06] Quinn, K., Kenny, A., Feeney, K., Lewis, D., O'Sullivan, D., Wade, V., "A Framework for the Decentralisation and Management of Collaborative Applications in Ubiquitous Computing Environments" to appear in proc 10th IEEE/IFIP Network Operations and Management Symposium (NOMS 2006), Apr 3-7, 2006, Vancouver, Canada.

[rada] Rada R., Mili H., Bicknell E., Blettner M., "Development and application of a metric on semantic nets", IEEE Transactions on Systems, Man, and Cybernetics 19, pp 17-30, 1989.

[rdf] World Wide Web Consortium (W3C), Resource Description Framework (RDF) Model and Syntax Specification, W3C Recommendation, 1999, http://www.w3c.org/

[rekhter] Rekhter, Y., Li, T., (eds) "A Border Gateway Protocol 4 (BGP-4)", IETF RFC 1771, Mar 1995.

[segall] Segall, B. et al, "Content-Based Routing in Elvin4", in proc of AUUG2K, Canberra 2000.

[stevenson] Stevenson, G., Nixon, P., Dobson, S, "Towards reliable wide-area infrastructure for context-based self-management of communications", in proc of 2nd IFIP WG6.3, WG6.6 Workshop on Autonomic Communications (WAC), Athens Greece, Oct 2005.

[strom] Strom et al., "Gryphon: An Information Flow Based Approach to Message Brokering", in proc of Int'l Symp. on Software Reliability Engineering, 1998.

[sussna] Sussna M., "Word sense disambiguation for free-text indexing using a massive semantic network", in proc of the 2nd Int'l Conf. on Information and Knowledge Management, pp 67-74, 1993.

[swoogle] "Swoogle statistics: Figure 6. Cumulative Term/Namespace usage distribution", http://swoogle.umbc.edu/modules.php?name=Swoogle_Statistics&file=figure&figurename=figure6

[tsarkov] Tsarkov, D., Horrocks, I., "Ordering Heuristics for Description Logic Reasoning", in proc of IJCAI 2005. Jul 30-Aug 5, 2005, Edinburgh, UK.

[wang] Wang, X., Dong, J.S., Chin, C.Y., Hettiarachchi, S. R., Zhang, D., "Semantic Space: An Infrastructure for Smart Spaces", IEEE Pervasive Computing Magazine, Jul-Sep 2004, pp 32-39

[zou] Zou Y., Finin T., Chen H., "F-OWL: an Inference Engine for the Semantic Web", Proceedings of 3rd Int'l Workshop on Formal Approaches to Agent Based Systems, Apr 2004, MD, USA.