

Towards a Managed Extensible Control Plane for Knowledge-Based Networking

David Lewis, John Keeney, Declan O'Sullivan, and Song Guo

Knowledge & Data Engineering Group (KDEG)
Centre for Telecommunications Value Chain Research (CTVR)
Department of Computer Science,
Trinity College Dublin,
Dublin, Ireland

{Dave.Lewis, John.Keeney, Declan.OSullivan, gsong}@cs.tcd.ie

Abstract. This paper proposes an open, extensible control plane for a global event service, based on semantically rich messages. This is based on the novel application of control plane separation and semantic-based matching to Content-Based Networks. Here we evaluate the performance issues involved in attempting to perform ontology-based reasoning for content-based routing. This provides us with the motivation to explore peer-clustering techniques to achieve efficient aggregation of semantic queries. The clustering of super-peers using decentralized policy engineering will deliver the incremental deployment of new peer-clustering strategies.

1 Introduction

As networks, distributed computing and pervasive computing converge, an increasing trend can be observed for making elements in a network more autonomous, be it an IP router, a wireless terminal, an application server or a sensor. This involves delegating more decision making capability to the elements and providing them information on their changing operational context, which can be used in making that decision. For example, much effort goes into the means for distributing information between elements when designing routing algorithms for ad hoc networks. From an operations perspective there is a need for such complex networks to become increasingly self-managed, or autonomic. Autonomic systems use knowledge of their operational state and operational context to self-manage, i.e. to self-configure, self-heal, self-optimize and self-protect, by monitoring state and context, planning and adapting. Though the need to self-manage was initially recognized as a challenge in dramatically reducing the operating costs of complex computing systems, increasingly complex networking systems comprising the Internet and telecommunications networks are also seen as needing to self-manage. Clark et al identified the central role of a knowledge-driven approach to support advanced AI techniques for monitoring and analyzing network conditions in order to drive the planning of optimization, protection or corrective strategies [clark].

The major challenges faced in distributing knowledge in converging networks are:

- Heterogeneity of contextual information: convergence means elements increasingly need to gather contextual information from peer elements of dissimilar types;
- Rapid evolution of applications: resulting in uncertainty about the type of contextual information that an element will need to gather in the future;
- Volatile peer set: with wireless and mobile communications playing an increasing role in integration, a given element will need to gather contextual information from a frequently changing set of peers.

This requires an autonomic knowledge delivery service that can inherently scale to the size of the network it supports. Equally, the recent interest in applying autonomic principles to communications networks has emphasized the need for access to network operational knowledge in an ontological form [mulvenna][stevenson].

In this paper we examine a common, scalable solution for gathering distributed context information that can be used by arbitrary peers in a future converged network. To address the volatility of peer sets we adopt the Publish-Subscribe (Pub-Sub) paradigm [meier] for the gathering of context information. The elements requiring contextual information express an interest through a subscription which is matched to messages published by other elements holding that type of information as it changes. Pub-Sub systems are already used for loosely coupled communication in a variety of applications. However, existing Pub-Sub systems require agreements on message types between the developers of publishing and subscribing applications. This places severe restrictions on the heterogeneity and dynamism of the information elements that can be exchanged. One solution to this 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 facilitate still looser coupling between producer and consumer applications than Pub-Sub. Several CBN solutions and prototypes exist, e.g. [carzaniga01][segall][pietzuch][chand][strom]. However, widespread CBN deployments have been slow to emerge. This is 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 for use in matching consumer subscriptions to message attributes, typically: Strings, Integers, Booleans and associated equality, greater then, less than, and regular expression matches on strings. This falls well short of supporting the heterogeneity and flexibility that elements in a converged network require to gather operational context. Selecting a more expressive language involves a difficult trade-off, since higher level features, e.g. set functions, introduce more complexity into a CBN node, and may only be of use to a subset of applications. We must aim therefore to have a CBN message and subscription language that can be expanded incrementally and autonomously to meet the requirements of emerging application domains without placing unnecessary overheads on the network as a whole.

Increasingly, researchers are turning to the use of ontology-based semantics to address this issue. The standardization of ontology languages by the Semantic Web initiative at the World Wide Web Consortium (W3C) [berners-lee] has spurred an increasing number of researchers to use ontology-based semantics to support interoperability in evolving systems [wang][masuoka][belecheanu]. However, these solutions have either been centralized or based on middleware scaled only to enterprise networks.

Therefore, we assert that the target for CBN expressiveness should be the subject-predicate-object structure of ontological knowledge representations, standardized 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 such 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)*. Previous work by authors detail a number of prototype systems that have been implemented to partially achieve this [lynch05][lynch06][keeney05][keeney06].

The scalability of a KBN to Internet proportions requires a routing control plane that minimizes 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 maximize 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 addressing schemes play only a small part in how the knowledge is organised and partitioned in the network. 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.

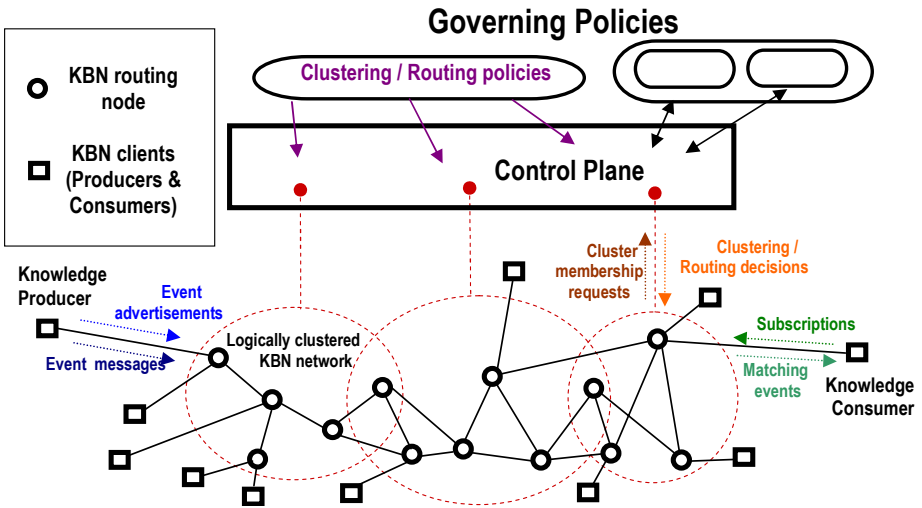


Fig. 1. A control plane for a Knowledge Based Network

Clearly, any software-based event forwarding algorithm will struggle to match the hardware optimized performance of packet forwarding in IP routers. Basing the forwarding algorithm on today's ontological reasoners will incur a particularly heavy computational load. However we do not seek to develop optimized reasoners for KBNs. Instead we are exploring the extent to which intelligent clustering in KBN routing algorithms that are cognizant of the performance profiles of existing reasoners and of the semantics being exchanged by client applications, can off-set this relatively poor forwarding algorithm performance.

In the remainder of this paper we propose an architecture for the KBN that separates the routing protocol from the matching and forwarding aspects (Figure 1). This split aims to minimize routing table size, as well as minimizing matching latency due to ontological reasoning. The routing protocols will exploit the robust behaviour and efficient look-up of peer-to-peer networks, and will use hierarchical clustering to minimize the ontological load on individual routers. To correctly dimension such a clustering scheme we must first understand the performance of existing reasoners. To this end we present in section 5 performance measures of such reasoners for use in KBN routers. We also discuss the impact this may have on the design of the routing plane. In addition we discuss related work (section 3) and design issues (section 4).

2 Architectural Principles

To reach its potential in flexibly supporting heterogeneous autonomic systems at Internet scales, the KBN must exhibit the following characteristics:

- Scalability in terms of latency and message throughput by minimizing the reasoning overhead at each router while minimizing routing configuration overhead.
- Robustness to the node and link failures that are inevitable at Internet scales.
- Self-configuration in response to the addition or removal of KBN nodes.
- Governance through policies, so that operational goals and constraints of different administrative domains can be applied in terms of high level rules to the relevant portions of the global KBN.

The use of RDF will support extensibility and applicability to new applications, while allowing the system to be extensible and modular in its routing and self-management functions by allowing co-existence of different routing and management schemes based on different sets of application semantics with configurable reasoning logic. It will support incremental deployment by different administrative domains, initially edge-based overlays, and later in the ISP-operated network core.

We therefore place an emphasis on Internet engineering values, (i.e. incremental deployment of new network features, support for application innovation, minimal standardization of core network features) in addition to the measurable properties of network scalability and performance. Overall, these issues 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 infrastructures, and are thereby insulated from the issues

of Internet scaling. By addressing asynchronous messaging over a highly decentralized network we uniquely attempt to reconcile Internet engineering values and knowledge engineering solutions, thereby exploiting the new efficiencies yielded by clustering KBN nodes based on semantic concerns.

The significance of the approach taken is that it exploits the ontological knowledge to construct efficient routing schemes through semantic clustering so that routing table sizes can be minimized and the knowledge base at each router can be kept small to minimize reasoning overhead. Clustering thereby 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.

3 Related Work

There has been little examination of the use of ontology-based semantics in content-based networking in the scientific literature. In [petrovic], an extension to the Toronto Publish/Subscribe System (ToPSS) is described that proposes extending the event/subscription matching function of this CBN to include class equivalence, ontological sub-class and super-class relationships (i.e. subsumption) and semantic mapping based relationships, which is equivalent the CBN extensions we have implemented as described in [keeney06][lynch05][lynch06]. However, [petrovic] does not address scalability issues of including ontology-based reasoning into the CBN, and no proposal is made to integrate this with the P2P routing extension for ToPSS. More significantly, however, no report of an implementation or evaluation of this proposal has yet emerged. In [li] a semantic publish/subscribe system is presented, but it is based on a centralized pub/sub bus implementation and thus is limited to enterprise scale and does not offer true CBN capabilities.

Much work to date on content-based networks has focused on how efficiency in routing can be gained through subscription aggregation and merging, where routes to subscribing clients are multiplexed with ones with covering subscriptions, i.e. broader subscriptions that will match all the event messages that would match the covered subscription. Recent progress with the XNET CBN has shown that perfect routing can be achieved in a scaleable manner independently of subscriber joins and leave rates though subscription aggregation [chand]. XNET does not, however, address how possible routes from KBN nodes to message producers are established. The default approach is flooding, where a node requests all other nodes for relevant routes, but this is not scaleable to large numbers of nodes [muhl]. This is addressed in the SIENA CBN through the static set up of spanning trees [carzeniga03] from producers to all possible consumers. However, these are then costly to recalculate in the event of configuration change or failures, thus failing our requirements for robustness and self-configuration.

The HERMES CBN [pietzuch], ToPSS [muthasumy] and the REBECCA CBN [tempstra] have all addressed these issues by applying peer to peer distributed hash table (P2P DHT) mechanisms to the formation of routing tables in CBN nodes. P2P DHTs such as CAN, CHORD and Pastry have well known properties of scalability, robustness and self-organization. It should be noted that though P2P system are concerned with efficiently routing queries to matching information sources, they do not address the CBN concern of optimally routing a sequence of asynchronous replies

back to the set of querying, or in CBN terms, subscribing clients. P2P DHTs provide efficient routing by using a cost metric keyed to the physical topology of the network resulting in average hop-counts for a route in the order of the log of the number of nodes in the network i.e. $O(\log(N))$. It is the demonstrated strengths of DHT-based routing protocols for CBNs that indicate the appropriateness of peer-to-peer Semantic Overlay Networks as a routing mechanism that meets our requirements for an Internet-scale KBN.

There are several attempts at applying P2P DHT techniques to the retrieval of distributed ontology encoded knowledge information, e.g. in RDF, in semantic overlay networks [tempich][cai][loser]. In supporting an ontology-driven DHT-based P2P routing mechanism for the KBN, the approach outlined in [loser] seems most promising due to its support for peer clustering. Used in this way, peer clustering introduces a hierarchy of peer groups based on policies. Such policies can admit nodes based on semantic closeness, recorded performance, administrative domains or indeed reasoning capabilities. It therefore provides a mechanism for these different routing configuration strategies to co-exist, serving different application domains or user communities in a way that supports incremental deployment and innovation.

The engineering of policies for a potentially complicated network of overlapping clusters, however, manifests as further technical policy engineering challenges. However, we have developed a novel policy engineering platform which uses the concept of communities as a resource and policy grouping abstraction [feeney]. This avoids the need for centralized engineering required in role-based policy schemes, and is therefore well suited to the decentralized evolution of policies by the interconnected network of autonomous decision makers that would characterize the operators of KBN clusters in a global setting. The core technical significance of this architecture is in applying semantic clustering based on reasoned knowledge bases at KBN routers through policy-driven super-peer hierarchies to KBN route management.

4 Design Issues

The fully decentralized approach of peer-to-peer systems results in highly scalable systems that can operate on an Internet-scale, prompting the proposed use of P2P techniques in the KBN control plane. By identifying the more significant concepts used to determine where messages should be routed around the network, it is envisioned that a set of determining keys can be used to semantically cluster peers. By supporting such clustering of producers and consumers, message forwarding can be largely restricted within clusters or between super peers, thereby improving scalability and load balancing. Further, these semantic clusters will be based on application semantics and usage, rather than static address hierarchies, so the management of KBN routing domains will need to respond to the needs and activities of application user communities as much as to ISP administrators, especially when operated as an overlay. As such, application-based grouping semantics will be very heterogeneous and so the control plane in which routing information is calculated and exchanged must be inherently more flexible. It must support the operation of multiple routing domains to reflect application heterogeneity while still supporting the sub-grouping needed to enable efficient dissemination of routing information. In addition, as shown in [carzaniga03], such use of determining keys in the node-level forwarding algorithm

of CBNs leads to substantial improvements in performance and scalability. Route recalculation and dynamic clustering of peers will be a relatively expensive operation, and will be effected by the amount of churn in the network as new nodes arrive and depart or new subscriptions and advertisements are published. This will therefore be a key performance indicator for the efficiency and scalability of the proposed scheme. Well optimized routes within and between clusters will still be required to minimize the number of hops required between producers and consumers.

The performance of the dynamic clustering and route determination algorithms will be of importance in terms of the overhead in the network and its ability to stabilize during and after churn. However a key overhead (see next section) is the loading of new ontology classes into a KBN node's reasoner, so the ability of the routing scheme to limit this is a key metric to be observed during simulation. This can be evaluated in terms of how the load of the network is balanced across clusters, but in a manner where traffic can be largely contained within the clusters and short routing paths are maintained between producers and consumers. The performance of the forwarding algorithms will depend on the size and complexity of the routing information stored at each node. This can be evaluated in terms of the average time taken to determine the forwarding strategies for individual messages in terms of the semantics of their content. It will be necessary to evaluate the degree to which new or cancelled subscriptions are either covered by another subscription or cause churn within the network, and contrast this with the heterogeneity of the semantics of the subscriptions present in the network. To further compare performances of the KBN a benchmark is presented in [keeney06b].

5 Initial Evaluation of Knowledge Model Reasoning

Recent exploratory experimentation by the authors has evaluated the integration of ontological equivalence and subsumption into SIENA CBN event/subscription matching [lynch05][lynch06][keeney06] yielding the following results. Firstly, the loading of new ontologies into a reasoner embedded in a KBN node is computationally expensive due to load-time inference, so the frequency of changes to the ontological base of a given KBN node must be minimised since changes will need to be distributed to each of the nodes in the network. Secondly, ontological reasoning is memory intensive and memory usage is proportional to the number of concepts and relationships loaded into the reasoner so reasoning latency can be controlled by limiting this number in any given KBN node. However, once loaded and reasoned over, the querying of such an ontological base is relatively efficient, with performance relative to size of the ontological base. Based on these initial observations, that the (re-)loading and (re-)reasoning of ontologies is expensive and that such operations will greatly affect the scalability of the network, it was considered possible that such axioms could form the basis of semantic clustering policies for partitioning the routing mechanism, thereby localising the effect of such operations and minimising the ontological base at each node and so improving routing performance.

In order to determine the extent to which ontology reasoner implementations differed with respect to initial loading and subsequent querying times, it was decided to undertake an experiment to compare OWL Description Logic (DL) implementations. Three typical DL implementations are measured in our experiment: Racer

[haarslev], the Jena framework [jena], and Pellet [parsia]. Racer implements a highly optimized tableau calculus for very expressive description logic and offers reasoning service for multiple TBoxes and for multiple ABoxes. Pellet is an open-source Java based OWL DL reasoner and supports the full expressivity of OWL DL including reasoning about nominals (enumerated classes). Two Pellet structural reasoners were tested: one is the “OWLReasoner” designed to provide an Jena-like interface to applications using Jena to access ontologies by translating the results of common calls to provide Jena data structures; the second “OWLAPIReasoner” reasoner is designed to support applications using the OWL API framework [owlapi] to query and manipulate ontologies. Both Pellet and Racer implement tableaux optimization techniques to improve their description logic reasoning performance. Finally, the Jena framework supports a number of different inbuilt reasoning modes, four of which are discussed here: The first, the “OWL_MEM_RDFS_INT” reasoner, incorporated directly into Jena, performs RDFS entailment reasoning. The second “OWL_MEM_TRANS_INF” reasoner, also supplied with Jena, performs transitive reasoning. The third reasoner configuration is where Pellet is plugged directly into Jena to perform full OWL DL reasoning and is accessed directly from within Jena like one of Jena’s own reasoners. The fourth mode “OWL_MEM”, is where Jena undertakes no reasoning, and is included for comparison.

A thorough and fair comparison of all aspects of these implementations is difficult, because they are implemented in different programming languages with the reasoning algorithms varying considerably. However for our purposes this is not important, rather we were interested in the performance of the implementations given a certain expected pattern of usage. This pattern of usage for knowledge routing is where the ontology is loaded and reasoned over initially at KBN start up time (which we call the loadtime stage) and then the ontology would be subsequently queried on a repetitive and ongoing basis (which we call the runtime stage) to facilitate the KBN routing. Thus in the experiment, for a number of ontologies, loadtime performance is given as the times taken for the different reasoners to load, parse and check the ontology, combined with the time taken to perform TBox classification, perform ABox realisation and an initial query of all concepts. The time taken to perform subsequent queries for the set of concepts in those ontologies was measured as the runtime performance of the reasoners.

For the purpose of investigating the impact of the complexity of ontologies upon the reasoner implementations, five ontologies were chosen for test. These were chosen in order to reflect a range of ontology complexity, from simple ontologies with instances through to complex ontologies with no instances. First, a large but relatively simple subset of DMTF’s CIM (Common Information Model) ontology with 167 classes and 733 individuals was used in the test. This ontology contains the set of CIM concepts required to manage printing devices. (A plug-in for the Protégé editor from the Universidad Politécnica de Madrid (UPM) [Lópezdevergara02] was used for the conversion of the CIM object-oriented model of classes and properties to an ontological model.) Secondly, we tested a more extensive ontology that represents the entire DMTF CIM Core Model, with 1215 classes and 5734 individuals. This ontology is very large by comparison to most ontologies, but relatively simple in terms of content. Next, the wine ontology [wine] from the W3C community contains a classification of wines, with only 138 classes of which 61 are imported from a food ontology, and 206 individuals of which 45 are also imported. The nominals

(individuals in class expression) and advanced DL constructors (e.g. disjunctions and equality) used in wine ontology makes it very complex to reason over, and is used as the de-facto stress test for OWL DL tools. The fourth test data set is the Galen ontology [galen]. It is a medical terminology ontology with a very large and complex class and property model (2749 classes) but without instances. It has traditionally been used as a benchmark for terminological reasoning. Finally, the Mindswap ontology [mindswap] is a relatively small but detailed ontology containing the student and staff details and listings of the Semantic Web Research Group in the University of Maryland. It contains 49 classes, of which 37 are imported, and 122 individuals, of which only 6 are imported.

All tests were performed on a desktop computer with Dual 3.2 GHz Intel Xeon processors, 2GB of RAM, running Windows XP Service Pack 2. For Java-based tools, Sun's JDK 1.5.0 was used. All the timings presented in this section are computed as the average of numerous independent timings.




Fig. 2. Loadtime performance of reasoners




Fig. 3. Runtime performance of reasoners

As can be seen from the loadtime performance (Figure 2) of the different reasoners, all of the reasoners perform similarly with the small ontologies, with the ones accessed through Jena performing a little worse than Racer and Pellet when accessed directly. However for the large and complex ontologies the Jena based reasoners perform substantially better than the others. When the operation of the other Jena based reasoners is compared to the Jena configuration performing no reasoning it is clear that much of the Jena loadtime overhead is independent of the level of reasoning required. Of all of the reasoners, only Pellet, whether or not through Jena, performed an adequate level of reasoning on the complex Wine ontology. For this reason, of the reasoners tested, the Pellet reasoner configurations seems to perform best, perhaps using its own Jena style interface or its OWL API interface for small ontologies, but definitely embedded in Jena for large and complex ontologies.

From the runtime performance graph (Figure 3) it is clear that all of the reasoners accessed using the Jena framework perform far superior to the others. This seems to be achieved by the clever use of model caching for second and subsequent queries within Jena. When compared to the other reasoner configurations, the use of Pellet through the Jena framework would appear to be the best overall choice for runtime performance. However, further investigation may be required to explain the counter-intuitive poor performance of Jena's own reasoners with the smaller ontologies.

In the event that an off-the-shelf generic reasoner is required, taken together, the results point to the adoption of the Pellet reasoner accessed through the Jena framework as the default reasoning configuration. However, in situations where very good runtime performance is required, such as in network routing, and this performance requirement outweighs the need for complete and correct reasoning, the use of cut-down or application specific reasoners may prove beneficial. It is also envisioned that in a large scale network different reasoners will be used by different nodes in accordance with their role in the clustered network's routing scheme.

6 Further Work and Conclusions

In this paper we propose an approach to managing the distribution of routing information between clusters of peers using super-peer communication in order to minimize the impact of ontology-based reasoning in the matching function of a Knowledge Based Network router. We propose that the policies for managing such clustering must therefore be based on the closeness of concepts in a web of ontologies and the likelihood of semantic similarity between subscriptions and advertisement. This work builds on experience gained in extending the basic SIENA CBN router with RDF sub-class and super-class operators, described in detail in [keeney06][lynch05][lynch06]. We do not aim to build a specialized RDF reasoner for this purpose, but to select from the various ones available. Thus in this paper we present the results from comparing several reasoners in order to both select an initial candidate and to understand further the performance trade offs between the addition of new concepts to a reasoners loaded knowledge-based and matching involving those concepts. This latter understanding may be particularly important if semantic clustering prompts certain nodes to specialize in certain groups of types and predicates by integrating reasoners optimized for those semantics and resulting in

faster messages matching. This work will continue by evaluating further extensions to the KBN router incorporating the selected Pellet reasoner, both through simulation and through a limited testbed.

Future work will involve integrating an existing community based policy management system [feeney] with the KBN control plane as the mechanism by which clustering policies are defined and potential conflicts identified and resolved. Policy-driven clustering enables the size of the super-peer network and the size and granularity of peer clusters to reflect different application domains needs. This will support overlapping clusters and hierarchies of clusters under separate administrative control. Different policy combinations will be applied to the testbed implementation of the KBN to verify stable behaviour, before those policies are included in the simulated KBN to evaluate the scalability of the resulting clusters. For example, the clustering policy may be specified in terms of accuracy and latency as well as the semantic spread of the query-able knowledge-base, or in terms of queries across a peer population and of the querying load across that population. Though the practicalities and performance impact of policy enforcement will be verified in the testbed, the validation of policy sets and the communication between super-peers in different cluster will be assessed through simulation.

The scalability and flexibility of the KBN under high load of heterogeneity will be evaluated using an OPNET based simulation, in order to test the message overhead involved in P2P clustering and the effectiveness of semantic load sharing. The key semantic distance measures used in clustering will be recorded for each node such that an analysis of the entire network of nodes will be possible in order to evaluate the efficiency of the KBN node clustering scheme that was simulated. A range of semantic distance measures have been proposed, such as [rada][jiang] [sussna][kashyap], and initial investigation will be required to determine the most suitable ones for use in different sizes of clusters or in ontology profiles associated with application domain clustering. The clustering schemes will be initially optimized to cope 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].

Ultimately we aim to design and validate differing clustering policies that can be used to tune and compliment semantic distance calculations. We will also assess the impact of policies on the coexistence of different (simulated) reasoning capabilities in KBN nodes. It is also planned to extend our work on incorporating semantic interoperability in node matching functions [keeney05][lewis] and in inter-cluster communications.

Further work will also focus on the use of further clustering policies. Such policies can admit nodes based not only on semantic closeness, but also in terms of recorded performance, availability of resources, trustability of nodes, reasoners used and reasoning capabilities, or administrative domains. It therefore provides a mechanism for these different routing configuration strategies to co-exist, serving different application domains or user communities in a way that supports incremental deployment and innovation. This use of clustering policies also supports innovation in clustering strategies by allowing peers to introduce new policy elements and the supporting super-peer matching capabilities.

Acknowledgements

This research was partly supported by Science Foundation Ireland under grant no. 03/CE3/I405, 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
- [cai] Cai, M., Frank, M., "RDFPeers: A scalable distributed RDF repository based on a structured peer-to-peer network", in proc of WWW 2004, May 2004, New York, USA.
- [carzaniga99] Carzaniga, A., Rosenblum, D., Wolf, A.L., "Challenges for Distributed Event Services: Scaleability 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, August 2001
- [carzaniga03] Carzaniga, A., Wolf, A. L., "Forwarding in a Content-Based Network" in proc SIGCOMM'03, August 25-29 2003, Karlsruhe, Germany, ACM Press
- [chand] Chand, R., Felber, P.A., "A Scalable Protocol for Content-Based Routing in Overlay Networks", IEEE International Symposium on Network Computing and Applications, April 2003, Cambridge, MA
- [clark] Clark, D., Partridge, C., Ramming, J.C., Wroclawski, J.T. "A Knowledge Plane for the Internet", in proc of SIGCOMM'03, 25-29 August 2003, Karlsruhe, Germany
- [feeney] Feeney, K. Quinn, K. Lewis, D. O'Sullivan, D. Wade, V., "Relationship-Driven Policy Engineering for Autonomic Organisations", in proc 6th IEEE Int'l Workshop on Policies for Distributed Systems, Stockholm, Sweden, 6-8 June 2005, pp 89-98
- [galen] The Galen Ontology, <http://www.cs.man.ac.uk/~horrocks/OWL/Ontologies/galen.owl>
- [haarslev] Haarslev, V., Moller, R. 2001. "RACER System Description", In proc IJCAR 2001, volume 2083 of LNAI, 701-706. Siena, Italy, Springer.
- [jena] Carroll, J., Dickinson, I., Dollin, C., "Jena: Implementing the Semantic Web Recommendations", in proc of WWW 2004, New York. May 2004. <http://jena.sourceforge.net/>
- [jiang] Jiang J. Conrath D., "Semantic Similarity based on corpus statistics and lexical taxonomy", in proc of Intl Conference on Research in Computational Linguistics, 1997.
- [kashyap] Presentation, NIST Invitational Workshop on Semantic Distance, Gaithersburg, MD, November 2003
- [keeney05] Keeney, J., Lewis, D., O'Sullivan, D., Roelens, A., Boran, A., Richardson, R., "Runtime Semantic Interoperability for Gathering Ontology-based Network Context", in proc of NOMS 2006, Vancouver, Canada. 3-7 April 2006
- [keeney06] Keeney, J., Lynch, D., Lewis, D., O'Sullivan, D., "On the Role of Ontological Semantics in Routing Contextual Knowledge in Highly Distributed Autonomic Systems", Tech. Report (TCD-CS-2006-15), Dept of Computer Science, Trinity College Dublin. 2006.
- [keeney06b] Keeney, J., Lewis, D., O'Sullivan, D., "Benchmarking Knowledge-based Context Delivery Systems", in proc of ICAS 06, Silicon Valley, USA, July 19-21, 2006.

- [lewis] Lewis, D., O'Sullivan, D., Power, R., Keeney, J., "Semantic Interoperability for an Autonomic Knowledge Delivery Service", in proc of WAC 2005, Athens Greece, Oct 2005
- [li] Li, H., Jiang, G., "Semantic Message Oriented Middleware for Publish/Subscribe Networks", in proc of SPIE, Volume 5403, pp 124-133, 2004
- [loser] Loser, A., Naumann, F., Siberski, W., Nejd, W., Thaden, U., "Semantic overlay clusters within super-peer networks", in proc Int'l Workshop on Databases, Information Systems and Peer-to-Peer Computing in Conjunction with the VLDB 2003
- [lópezdevergara02] López de Vergara, J.E., Villagrà, V.A., Berrocal, J., "Semantic Management: advantages of using an ontology-based management information meta-model", in proc of HP-OVUA'2002, distributed videoconference, 11-13 June 2002
- [lynch05] Lynch, D., "A Proactive approach to Semantically Oriented Service Discovery", MSc dissertation, Dept of Computer Science, Trinity College Dublin, 2005
- [lynch06] Lynch, D., Keeney, J., Lewis, D., O'Sullivan, D., "A Proactive approach to Semantically Oriented Service Discovery", in proc of IWT'06, at WWW'06, Scotland. 2006.
- [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
- [mindswap] The Mindswappers Ontology, <http://www.mindswap.org/2004/owl/mindswappers>
- [muhl] Muhl, G., Fiege, L., Gartner, F., Buchman, A., "Evaluating Advanced Routing Algorithms for Content-Based Publish/Subscribe Systems", in proc of MASCOT 2002
- [muthusamy] Muthusamy, V., Jacobsen, H.A., "Small-scale peer-to-peer publish/subscribe" in proc Workshop on Peer-to-Peer Knowledge Management, San Diego, USA, July 2005
- [mulvenna] Mulvenna, M., Zambonelli, F., "Knowledge Networks: the nervous system of an autonomic communication infrastructure", in proc of WAC 2005, Athens Greece, Oct 2005
- [owlapi] Sean Bechhofer, Phillip Lord, Raphael Volz. Cooking the Semantic Web with the OWL API. In proc of ISWC2003, Sanibel Island, Florida, October 2003.
- [parsia] Parsia, B., Sirin, E. 2004., "Pellet: An OWL-DL Reasoner", Poster at ISWC 2004, Hiroshima, Japan, 2004.
- [petrovic] Petrovic, M., Burceaa, I., Jacobsen, H.A. "S-ToPSS – a semantic publish/subscribe system" in proc VLDB, Berlin, Germany, September 2003
- [pietzych] Pietzych, P., Bacon, J., "Peer-to-Peer Overlay Broker Networks in an Event-Based Middleware". in proc DEBS'03 at ACM SIGMOD/PODS Conference. California, June 2003
- [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] W3C (1999) World Wide Web Consortium, Resource Description Framework (RDF) Model and Syntax Specification, W3C Recommendation, <http://www.w3c.org/>
- [rekhter] Rekhter, Y., Li, T., (eds) "A Border Gateway Protocol 4 (BGP-4)", IETF RFC 1771, March 1995
- [segall] Segall, B. et al, "Content-Based Routing in Elvin4", In proc 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 WAC 2005, Athens, 2005
- [strom] Strom et al., "Gryphon: An Information Flow Based Approach to Message Brokering", In Intl. Symp. on Software Reliability Engineering 1998
- [sussna] Sussna M., "Word sense disambiguation for free-text indexing using a massive semantic network", proc of Conference on Information and Knowledge Management, 1993.
- [swoogle] "Swoogle statistics: Figure 6. Cumulative Term/Namespace usage distribution", http://swoogle.umbc.edu/modules.php?name=Swoogle_Statistics&file=figure&figurename=figure6

- [tempich] Tempich, C., Staab, S., Wranik, A., "REMINDIN': semantic query routing in peer-to-peer networks based on social metaphors" WWW 2004, New York, USA, 2004.
- [terpstra] Terpstra, W.W., Behnel, S., Fiege, L., Zeidler, A., Buchmann, A.P., "A peer-to-peer approach to content-based publish/subscribe", in proc of DEBS 2003, ACM Press 2003
- [wang] wang, X., Dong, J.S., Chin, C.Y., Hettiarachchi, R., Zhang, Daqing, d., "Semantic Space: An Infrastructure for Smart Spaces", IEEE Pervasive Computing Magazine, Jul-Sep 2004, pp 32-39
- [wine] W3C: The Wine Ontology, <http://www.w3.org/TR/2003/PR-owl-guide-20031209/wine>