The Knowledge Engineering Review, Vol. 18:3, 193–195. © 2004, Cambridge University Press
DOI: 10.1017/S0269888904000013 Printed in the United Kingdom

Foreword: ontologies for distributed systems

HEINER STUCKENSCHMIDT

1 Introduction

The benefits of using ontologies have been recognised in many areas such as knowledge and content management, electronic commerce and recently the emerging field of the Semantic Web. These new applications can be seen as a great success of research in ontologies. On the other hand, moving into real application comes with new challenges that need to be addressed on a principled level rather than for specific applications. This special issue will be devoted to less well-explored topics that have come into focus recently as a response to the new problems we face when trying to use ontologies in heterogeneous distributed environments. These environments include the use of ontologies in peer-to-peer and pervasive computing systems.

Peer-to-peer architectures have been proven useful for large-scale information sharing. Their benefits in terms of a low entrance barrier and a high degree of flexibility and fault tolerance have led to successful applications like Napster or Gnutella. One of the drawbacks of current peer-to-peer systems, however, is their limited search and querying capabilities. Most systems only support plain keyword queries that suffer from the lack of naming conventions and the ambiguity of natural language. Recently, different research projects attempt to overcome this limitation by integrating schematic and conceptual knowledge into peer-to-peer environments. Examples of such projects are Edutella,[1] EDAMOK,[2] Piazza[3] and our own project, SWAP.[4]

A similar situation can be observed for pervasive and ubiquitous computing. (see for instance `http://www.ebiquity.org/`). Here, however, the notion of context is even more essential as it becomes an inherent parameter for any kind of information sharing. When people move around, their environment changes, imposing new constraints and requirements on their information needs and communication capabilities. Modelling these constraints and using them to adapt and improve information exchange between users in the system is a very challenging task. It quickly turns out that besides the need to capture conceptual information, notions of space and time become more important. This requires not only advanced context-matching methods but also reasoning capabilities that go beyond typical inheritance reasoning normally associated with ontologies. Currently a number of projects on ubiquitous information access address these problems. Examples of such projects are CODA,[5] GAIA[6] and CoBrA.[7]

2 The content of this issue

The work contained in this special issue contributes to improving information processing in distributed systems such as peer-to-peer systems and pervasive computing environments using ontologies. The contributions can roughly be divided into two topic areas.

[1] `http://edutella.jxta.org/.`
[2] `http://edamok.itc.it/.`
[3] `http://data.cs.washington.edu/p2p/piazza/.`
[4] `http://swap.semanticweb.org/.`
[5] `http://www-2.cs.cmu.edu/coda/.`
[6] `http://choices.cs.uiuc.edu/gaia/.`
[7] `http://cobra.umbc.edu/.`

2.1 Ontological modelling for distributed systems

One of the central topics of this issue is the use of ontologies in distributed and decentralised systems. This class of system has gained a lot of attention recently. Concepts like the Semantic Web, grid computing, peer-to-peer systems or ubiquitous computing are topics of ongoing research. Ontologies can have different functions in these systems. They might define shared terminologies, explicate hidden assumptions of system components or become a vehicle to specify system architectures and information formats.

Undercoffer *et al.* address the problem of intrusion detection in networks and the use of ontologies and logical reasoning methods in order to solve this problem more effectively. Ranganathan *et al.* describe the enrichment of a CORBA-based pervasive computing environment with explicit ontologies in order to facilitate matchmaking and interoperability. Chen and Finin address the same problem, focusing on the recognition of the local context a user if currently in and the conclusions than can be drawn from this fact. Schlenoff *et al.* discuss the use of ontologies to capture world knowledge necessary to build autonomous vehicles by supporting reasoning about potential obstacles and their characteristics.

2.2 Meaning negotiation

One of the most pressing and challenging problems of decentralised semantic-based systems is the need to compare different representations of semantics. In the world of ontologies, this is known as mapping or integration. In the broader area of distributed systems the term ''meaning negotiation'' has been used to indicate that the problem is not restricted to the use of formal ontologies. New and unconventional methods are needed to achieve at least partial results on this difficult problem. Existing approaches range from theorem-proving methods to voting and bidding mechanisms each being applicable under certain assumptions. Up to now no general solution to this problem has been found.

Giunchiglia and Shvaiko discuss the general problem of matching graph-based representations such as thesauri, XML documents or ontologies on a semantic basis. Ge *et al.* address the special problem of finding a common terminology in peer-to-peer systems and proposes a method for agreeing on a shared model based on local definitions.

2.3 Overview of the content

Ontologies in distributed systems

''An ontology for context-aware pervasive computing environments''
H Chen and T Finin
''Use of ontologies in a pervasive computing environment''
A Ranganathan, RE McGrath, RH Campbell and MD Mickunas
''A target-centric ontology for intrusion detection''
J Undercoffer, A Joshi, T Finin and J Pinkston
''Using ontologies to aid navigation planning in autonomous vehicles''
Craig Schlenoff, Stephen Balikirsky, Mike Uschold, Ron Provine and Scott Smith

Meaning negotiation

''OntoVote: a scalable distributed vote-collecting mechanism for ontology drift on a P2P platform''
Yanfeng Ge, ong Yu, Xing Zhu, Shen Huang and Min Xu
''Semantic matching''
F Giunchiglia and P Shvaiko

3 Discussion

Looking at the papers in this issue it appears that there are promising approaches for using ontologies to model and reason about distributed systems. It seems a problem that developers of ontology-based distributed systems frequently hit the limits of existing ontology languages and tools. Although efforts in the area of the Semantic Web have developed languages and infrastructure for supporting the use of ontologies in a Web-based setting, many important aspects are not yet supported to the degree needed. I would like to stress the following two points that demand a fundamental treatment.

3.1 Context versus ontology

Ontologies are explicit models of a view on a domain that is *shared* by different parties. Its use enables members of the different parties to communicate and exchange information. Contexts represent the *personal* view of a party on a specific domain. In order to communicate with other parties, relations between their context definitions have to be found. Both of these views are valuable for the Semantic Web as the benefit of semantic models can be that they define a common understanding of specific terms. Thus reference to a term in a shared ontology makes it possible to communicate between systems on a semantic level. On the other hand semantic models can be used to explicate what a certain term means in the context of a specific information source. In order to make use of both of these benefits of semantic models the representational infrastructure of the Semantic Web has to consist of a mixture of ontologies and context definitions. In such a mixed infrastructure, on the one hand, the intended meaning of terms provided by individual persons can be captured more accurately by a context definition describing the personal view of the author of that term. On the other hand, different context descriptions can be aligned using shared ontologies that provide a shared terminology for comparing terms.

3.2 Spatio-temporal knowledge

The problem of representing space and time has also been addressed in the framework of description logics which currently is the dominant framework for encoding ontologies on the Semantic Web resulting in special logics for spatial and temporal reasoning. These approaches, however, suffer from a high complexity that often enforces some restrictions on the expressiveness of concept descriptions. Another approach often seen is to establish links to generic ontologies of space and time. These generic ontologies, again, often contain rich axiomatisations that are not really exploited by existing tools. The most important argument against using these description logic-based approaches is the kind of reasoning they support. In particular, the logics only support exact matches on the basis of logical deduction. With respect to the temporal and spatial context of statistics, however, we will almost never find exact matches in practice. Therefore we rather need representations of space and time that support approximate matching in terms of relevance to a given request. Approaches supporting this kind of relevance reasoning, however, often lack a strong formal foundation. Here, new ideas about a flexible but well-founded connection between models for dealing with conceptual, spatial and temporal knowledge is needed.