Negotiating Domain Ontologies in Distributed Organizational Memories

Ludger van Elst and Andreas Abecker

German Research Center for Artificial Intelligence (DFKI)
Knowledge Management Department
Postfach 2080, D-67608 Kaiserslautern, Germany
{elst|aabecker}@dfki.uni-kl.de

Abstract

Organizational Memory Information Systems (OMIS) have a strong need to represent shared understanding of various actors in the information landscape. Ontologies are widely seen as an adequate means for this purpose. In the FRODO project we develop an agent-based framework for Distributed Organizational Memories, and one important type of agent in the framework are agents responsible for managing domain ontologies (Domain Ontology Agents, DOA). The process of establishing shared conceptualizations in our framework takes place on three levels: i) A DOA collects evidence that a portion of knowledge might be sharable among a group of actors. ii) The DOA coordinates a negotiation procedure between the relevant actors. iii) Ontology Societies explicitly reflect the sharing scope of the knowledge managed by a DOA. These societies are grounded on the rights and obligations of the actors with respect to a specific domain ontology.

The integration of all three levels is a cornerstone of FRODO's approach to support a full ontology lifecycle in a distributed environment. Our actual research focus is the elaboration of level i). In particular, we are working on an instance-based approach for finding ontology overlaps on the basis of text analysis techniques.

Domain Ontologies in Distributed Organizational Memories

An ***Organizational Memory Information System***, (OMIS, or shortly OM) is an intelligent information and assistant system that fosters creation, accumulation, sharing, reuse and further development of explicit knowledge in an organization which may be prevalent in manifold different forms, formats and systems. With the KnowMore architecture (Abecker et al., 1998) we presented a layered approach for an OM with context-specific, proactive information services offered to support the end user's actual task at hand. The KnowMore architecture was a basis for several further research activities in this area (Staab & Schnurr, 2000; Holz et al., 2001; Abecker et al., 2001; Jansweijer et al., 2000). A core element is the so-called ***knowledge-description layer*** which provides homogeneous access to heterogeneous information sources through ontology-based description of information objects and their relationships. Naturally one major element of such an approach is a suitable and up-to-date ***domain ontology*** used for describing the actual content of information objects.

In the FRODO project we argue that for pragmatic (dealing with legacy systems, methodologies for starting KM in initiatives) as well as for theoretical reasons (optimal strategies for knowledge creation and maintenance, enabling of virtual enterprises) it is required to proceed from the centralized single-OM approach of KnowMore to a ***distributed, multi-OM scenario*** with an arbitrary number of autonomous, but cooperating OMs (Abecker et al., 2002). Communication and collaboration between different OMs may (and normally must) occur in manifold forms and between different levels of the KnowMore layered OM architecture. Together with other authors with similar goals (see (Bergenti et al., 2000; Dignum et al., to appear)), we propose to implement the complex information ecology in such a scenario on the basis of software agents organized in an ***agent society*** by taking different roles and responsibilities, rights and obligations, as defined in a social layer of the agent system design.

In (van Elst & Abecker, 2001a; van Elst & Abecker 2001b) we argue that creation and maintenance of the *domain ontologies* in such a distributed OM scenario are of particular interest to start with such agent society considerations. We observe a trade-off between complexity of communication and quality of service, as well as a balance required between competing goals such as sharing scope of ontological commitments, degree of formalization of knowledge exchanged, and stability of the stored knowledge. In order to effectively manage the several "balancing acts"- which mainly come from the negotiation between individual concerns and organizational concerns – we propose to establish an ***ontology society*** as a set of software agents acting as stakeholders in different roles with respect to a given domain ontology, organized via the technical means of ***rights and obligations*** which constitute the specific roles in a given ontology society (van Elst & Abecker, 2001a; van Elst & Abecker 2001b).

A central mechanism in such a multi-OM scenario is the communication between different OMs which might refer to different domain ontologies. In the next section we briefly sketch three "levels of harmonization" which stepwise bridge understanding gaps and enable meaning negotiation. In the subsequent section we give a simple example which shows how to support the ongoing evolution of domain ontologies and mediation knowledge going through these three steps.

The FRODO Ontology Evolution Approach

Evidences for Ontology Matches

The most basic activities for implementing ontology negotiation processes concern noticing an ontology mismatch and finding potential ontology mapping operations. In the literature mainly ***evidence for ontology merging operations*** are mainly based on some combination of evidence from two sources (see, e.g. (McGuiness et al., 2000a; McGuiness et al., 2000b; Dieng & Hug, 1998)):

a) *Term-based evidence* considers the textual description (i.e., the "name") of a concept in an ontology.

b) *Topology-based evidence* considers the structure of the concept graph representing (mostly) the is-a hierarchy of the ontology.

In the case of an OM application where ontology concepts are used for indexing and retrieving (in the general case, multimedia, but today normally mainly text) documents there is at least one further possibility, namely gathering *instance-based evidence* (see also (Lacher & Groh, 2001)). Here, the idea is to compare documents (or, better document sets) indexed with certain ontology concepts using text analysis and document understanding techniques. If, for instance, the set of documents indexed with concept A of one ontology is considered by some document comparison algorithms very similar to a set of documents indexed with concept B in another ontology, it could be a good idea to consider A and B to represent the same concept.

For experimenting with this idea, we implemented a domain ontology agent by adding agent communication facilities with the JADE[1] agent platform to the Protégé[2] tool for ontology development, and coupled it with the mindAccess[3] document analysis software. First results will be presented at the workshop.

A topic which still seems to be a bit neglected is the ***aggregation of evidences*** for determination of candidate ontology-merging operations. Obviously concepts in two ontologies under consideration are merging candidates if we find arguments for a sufficient degree of *similarity* between them. The notion and use of similarity in (Dieng & Hug, 1998), for example, suggests to think about the notion of similarity in Case-Based Reasoning which has been developed up to a considerably sophisticated stage (cp. (Wess, 1996)). Taking into account that an ideal aggregation mechanism would have to deal with different kinds of evidence (term-based, structure-based, instance-based), coming with different degrees of confidence, maybe with conflicting evidences, maybe with some degree of ignorance, it stands to reason whether such a mechanism should not be based on a strong mathematical foundation, as e.g. suggested by Richter's grounding of similarity in CBR on *Dempster and Shafer's theory of evidence* (Richter, 1995; Kohlhas & Monney, 1994). We did not yet work into this direction, but it is straightforward how to apply this theory in the given scenario, and it seems an interesting question for the workshop discussion whether the participants think such a step relevant and necessary or over-formalized.

Communication for Ontology Negotiation

If two agents detect a negotiation need regarding their respective ontologies, two ingredients are required to enable successful communication:

- *Speech acts* for meaning negotiation provide the appropriate primitives for talking about suggestions, questions, answers, and decisions about negotiating meaning between two agents. In (van Elst & Abecker, 2001a; van Elst & Abecker, 2001b) we started with a top-level analysis of the required expressiveness, which is mainly determined by the representation means of the underlying ontology formalism and the wished merging operations. In the workshop we will present a FIPA-compliant description of our proposed performatives for negotiating about a TRIPLE ontology (http://www.dfki.uni-kl.de/frodo/triple/).
- *A negotiation protocol* describes the process how two agents can come to an agreement about a meaning negotiation question. In FRODO we did not yet investigate this topic. However, (Bailin & Truszkowski, 2001) made a proposal which seems a good starting point to work from.

Ontology Society Framework

As already mentioned above, we see ontology lifecycle management as an integral part of a comprehensive agent society for managing distributed OMs. Other kinds of agents comprise, e.g. user agents, information retrieval agents, process and task agents, etc. With the FRODO agent platform (Abecker et al., 2002) we provide the basis for setting up such complex systems and defining their social layer with the concepts of actor roles defined via agents' rights and obligations which are a subset of their competencies. The FRODO agent platform is currently being implemented on top of JADE. In (van Elst &

[1] http://sharon.cselt.it/projects/jade/
[2] http://protege.stanford.edu/index.shtml
[3] http://www.im-insiders.de/download/mindaccess_engl.pdf

Abecker, 2001a; van Elst & Abecker, 2001b) we introduce six kinds of ontology agents described below in terms of the speech acts they are allowed, or obliged, respectively, to use with respect to a specific ontology in quest:

- *Passive Users* are members of an Ontology Society who intend just to *use* the respective domain ontology. Because they are only consumers, they are not obliged to any services. On a voluntary basis they can suggest ontology updates.
- The role of *Associate Users* is similar to the consumer role of Passive Users. However, Associates have the additional right to receive updates of the ontology whenever they want (e.g., when the ontology has changed), i.e., this role defines a kind of subscription service for a domain ontology.
- *Partner Users* have not only the right to use a domain ontology, but also committed themselves to help improving it by suggesting ontology updates, e.g. on the basis of quality feedback wrt. ontology exploitation.
- Ontology *Experts* are obliged to answer queries about a domain ontology.
- In addition to Experts, an *Editor* of a domain ontology is also responsible for ontology maintenance. The editor accepts update suggestions from the various users, negotiates possible updates with the partners and tries to guarantee global quality criteria of the domain ontology (e.g., soundness). The editor also realizes the service for all ontology society members with the right to receive update notifications. Furthermore, the editor manages the social model of "its" ontology society. This means, the editor knows the associations between roles and rights/obligations as well as which roles are enacted by the various actors. In particular, the editor grants roles (e.g., negotiates rights and obligations) to the applicants that want to join the ontology society.

	Passive User	Associate User	Partner User	Expert	Editor
Query	**R**	**R**	**R**	**R**	**R**
Receive Update		**R**	**R**	**R**	**R**
Suggest Update	**R**	**R**	**R/O**	**R**	**R/O**
Answer Queries				**R/O**	**R**
Edit					**R**
Send Upd. Notif.					**R/O**
Grant Guarantees					**R**
Guarantee Quality					**O**

Table 1: Example Rights (R) and Obligations (O) of Ontology Agents

Based upon these specifications we define ***Domain Ontology Agents (DOA)*** and ***Distributed Domain Ontology Agents (D^{2}OA)*** which are responsible for ontologies *within one* OM, or are located *between several OMs* to facilitate cross-OM communication, respectively. Typical questions to DOAs are "What are the subconcepts of concept A?" whereas D^{2}OAs answer questions like "Which OM contains concepts like A and B?" or "What does A mean in OM$_y$?". In terms of the roles mentioned above, both DOAs and D^{2}OAs typically enact the role of *Ontology Experts* and *Ontology Editors*.

Imagine for example two groups of experts, one for domain D1, one for domain D2. Each group negotiates its own domain ontology managed by DOA$_{D1}$ and DOA$_{D2}$, respectively. D^{2}OA has knowledge what these ontologies are about and tries to identify points of contact or overlaps between them. Then, D^{2}OA initiates a negotiation procedure between DOA$_{D1}$ and DOA$_{D2}$. The result might be a common upper level ontology or a mapping for some parts of the ontologies. So, in general, the architecture can handle three levels of negotiated meaning:

a) *Level 1* („No shared conceptualization"): Neither the local ontology agents (DOA$_{D1}$, DOA$_{D2}$) nor the D^{2}OA have an explicit agreement on the ontological level. In this case, to enable some level of communication techniques are applied which do not rely on a shared vocabulary (e.g., text retrieval). Communication on this level is monitored by D^{2}OA.

b) *Level 2* („Ontology Mapping"): D^{2}OA maintains mapping rules between the ontologies of the two local Domain Ontology Agents. These rules are based on the evidence gained by low-level communication on level 1.

c) *Level 3* („Ontology Negotiation"): In case of high evidence for overlap, a negotiation process is initiated by D^{2}OA. Possible outcomes are, e.g., that the two local DOAs do not want to establish are shared conceptualization or that they define a common top-level ontology. If the two ontologies agree to have some part of their domain ontologies in common, this results in a change of the underlying role model. This would have an effect on further developments of this part of the ontology, because both are now responsible, i.e., each change has to be preceded by a negotiation step as defined by the agent's roles (cf. Table 1).

Summary

Three major points describe our research approach for meaning negotiation in distributed OM applications:

- The use of text analysis techniques for instance-based evidence finding (which is possible because our ontologies are embedded into the distributed OM usage scenario) offers new, rarely investigated, but practical possibilities.
- The formally sound examination of evidence aggregation seems a pretty new topic.

- The management of meaning negotiation questions as an integral part of a comprehensive ontology society approach allows full lifecycle support for ontologies in OMs.

References

Abecker, A., Bernardi, A., van Elst, L., Lauer, A., Maus, H., Schwarz, S., and Sintek, M. 2002. *FRODO: A Framework for Distributed Organizations - Milestone M1: Requirements Analysis and System Architecture.* DFKI Document D-01-01, Kaiserslautern.

Abecker, A., Bernardi, A., Hinkelmann, K., Kühn, O., and Sintek, M. 1998. Towards a Technology for Organizational Memories. *IEEE Intelligent Systems*, May/June 1998.

Abecker, A., Bernardi, A., Ntioudis, S., Herterich, R., Houy, C., Legal, M., Mentzas, G., and Müller, S. The DECOR Toolbox for Workflow-Embedded Organizational Memory Access. 2001. In: *ICEIS 2001, 3rd Int. Conf. on Enterprise Information Systems*, Setúbal, Portugal.

Bailin, S.C., and Truszkowski, W. 2001. Ontology negotiation using JESS. In *Proc. 3rd Int. Conference on Enterprise Information Systems, ICEIS-2001, Setúbal, Portuga*l.

Bergenti, F., Poggi, A., and Rimassa, G. 2000. Agent Architectures and Interaction Protocols for Corporate Management Systems. In *Proc. ECAI Workshop on Knowledge Management and Organisational Memories*, pp. 39-47, Berlin, Germany.

Dieng, R., and Hug, S. 1998. Comparison of "Personal Ontologies" Represented through Conceptual Graphs. In: *13th European Conference on Artifical Intelligence (ECAI'98)*, Brighton, UK.

Dignum, V., Meyer, J.J., and Weigand, H. *to appear*. Towards an Organizational Model for Agent Societies Using Contracts. To appear in: *Proc. of AAMAS'02, First Int. Joint Conf. on Autonomous Agents and Multi-agent Systems*, Bologna, Italy, July 15 – 19, 2002.

van Elst, L., and Abecker, A. 2001a. Ontology-Related Services in Agent-Based Distributed Information Infrastructures. In: *Proc. 13th Int. Conf. on Software Engineering & Knowledge Engineering*, Buenos Aires, Argentina, pp. 79-85.

van Elst, L., and Abecker, A. 2001b. Domain Ontology Agents in Distributed Organizational Memories. In: *IJCAI'2001 Working Notes of the Workshop on Knowledge Management and Organizational Memories*, Seattle, Washington, USA, pp. 39-48.

Holz, H., Könnecker, A., and Maurer, F. 2001. Task-Specific Knowledge Management in a Process-Centred SEE, In: *Workshop on Learning Software Organizations 2001, (LSO 2001),* Kaiserslautern, Springer LNCS 2176.

Jansweijer, W., van de Stadt, E., van Lieshout, J., and Breuker, J. 2000. Knowledgeable Information Brokering. In: *European eBusiness and eWork Conference 2000,* Madrid.

Kohlhas, J., and Monney, P.A. 1994. Theory of Evidence - A Survey of its mathematical Foundations, Applications and Foundational Aspects. ZOR 39, 1994, pp. 35-68.

Lacher, M., and Groh, G. 2001. Facilitating the exchange of explixit knowledge through ontology mappings. In: *14th Int. FLAIRS conference*, AAAI Press, Key West.

McGuinness, D.L., Fikes, R., Rice, J., and Wilder, S. 2000a. An Environment for Merging and Testing Large Ontologies. *Proc. of the Seventh Int. Conf. on Principles of Knowledge Representation and Reasoning (KR2000)*. Breckenridge, Colorado, USA.

McGuinness, D.L., Fikes, R., Rice, J., and Wilder, S. The Chimaera Ontology Environment. 2000b. *Proc. of the 17th National Conf. on Artificial Intelligence (AAAI 2000)*. Austin, Texas.

Richter, M.M. 1995. On the notion of similarity in case-based reasoning. In: *Mathematical and Statistical Methods in Artificial Intelligence* (ed. G. della Riccia et al). Springer Verlag, pp. 171-184.

Staab, S., and Schnurr, H.-P. 2000. Smart Task Support Through Proactive Access to Organizational Memory. *Knowledge-Based Systems,* Elsevier, 13(5), pp. 251-260.

Wess, S. 1996. *Fallbasiertes Problemlösen in wissensbasierten Systemen zur Entscheidungsunterstützung und Diagnostik*. Infix Verlag.