ICEBERG: Exploiting context in information brokering agents

Conference Paper *in* Lecture Notes in Computer Science · July 2000

DOI: 10.1007/978-3-540-45012-2_4 · Source: DBLP

ICEBERG: Exploiting Context in Information Brokering Agents©

Catholijn M. Jonker[1] and Arjen Vollebregt[2]

[1]Vrije Universiteit Amsterdam, Department of Artificial Intelligence,
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands, email: jonker@cs.vu.nl, URL: www.cs.vu.nl/~jonker
[2]KPN Research, St. Paulusstraat 4, Leidschendam, The Netherlands, email: A.M.Vollebregt@kpn.com.

Abstract. The research reported in this paper has both a scientific and a commercial aim. The scientific interest is to explore the use of contexts in order to improve the quality of information brokering. In this paper it is shown that contexts in information brokering can be exploited to enable four directions of query reformulation: up and down (standard query expansion) and sideway reformulations of the user request that can even involve going from one context to another.

1. Introduction

A lot of research is done on the subject of information retrieval, federated databases, search engines, and information brokering agents[1]. The basic problems are recall and precision (as formulated within multi-database theory). This paper focusses on precision. A well-known technique to improve precision is relevance feedback to enable query expansion; the user can indicate "more like this result", or is asked to choose from a number of concepts that are automatically distilled from the top ranking results. Even though query expansion approaches help in the sense that more items similar to those first retrieved are found, there is much room for exploring interactive methods for improving precision (Manflano et al., 1998).

The work reported here is based on the assumption that precision can be proved by exploiting the knowledge gained from good domain analysis. The notion of a multi-dimensional context is introduced and it is shown how this can be used to:

- disambiguate the user request by finding the context relevant for his request
- sharpening his request (the downward direction of query expansion)
- widen his scope xt (upward query expansion)

© Published in: Klusch, M. and Kerschberg, L., *Cooperative Information Agents IV, The Future of Information Agents in Cyberspace,* Proceedings of the 4th International Workshop, CIA 2000, Lecture Notes in Artificial Intelligence, Vol. 1860, Springer-Verlag, pp. 27-38.

[1] See any of the major conferences, shops, and journals on agent- or multi-agent systems, information retrieval, and multi-databases.

- sideways query reformulations within the same context or even by going from one context to a different context that is more fitting with respect to the actual wishes of the user. This is not covered by traditional query expansion.

In Section 2 the notion of a multi-dimensional structure of contexts is introduced. In the third section the information brokering process as performed by the ICEBERG broker is described from a global viewpoint. In Section 4, the internal processing of the ICEBERG information broker is explained, with an emphasis on the exploitation of contexts. It describes the process of reformulation of user request on the basis of the multi-dimensional structure of contexts (this includes standard forms of query expansion based on feedback mechanisms). In Section 5 it is explained how the search results are evaluated on the basis of the multi-dimensional structure of contexts. Section 6 compares the work reported in this paper to other work and contains the conclusions.

2. Notion of Context

The research has been performed within the context of the ICEBERG project of KPN Research (the research department of the royal Dutch telecom). The reference to the ICEBERG is made to emphasize that the average user when trying to find information has only the amount of knowledge on the topic that could be compared to the very tip of a huge ICEBERG. The aim of the ICEBERG project is to get that part of the ICEBERG above water in which the user is interested. Domain-specific knowledge is specified using ontologies. Ontology, as Gruber (cf. (Gruber, 1993)) describes it *"explicit specification of a conceptualization"*, provides a vocabulary for talking about a domain. In this paper a novel way to structure the domain of application is presented, i.e., the *contextual universe*. This notion is first defined and then explained.

Definition
A contextual dimension D for a domain of application P is a set of concepts relevant for P. A contextual universe U for a domain of application P is a tuple <D1, …, Dn> (n•1) of contextual dimensions. A context C within a contextual universe U is a tuple <I, T>, where I is the identifier of the context and T is a semantic network of concepts from the domain ontology of P. The identifier I of a context within a contextual universe U is defined as a tuple <d1, …, dn>, with di ∈ Di for all i•n.

For the ICEBERG project the domain of courses has been chosen. The contextual universe of courses has been modelled using two dimensions: the background and the theme dimension, see Figure 1. The background dimension describes the reason the intended participant is interested in the course, and the theme dimension describes the theme of the course irrespective of the background of the intended participants. For the prototype a restriction was made to those courses that have something to do with communication, i.e., the first column of Figure 1.

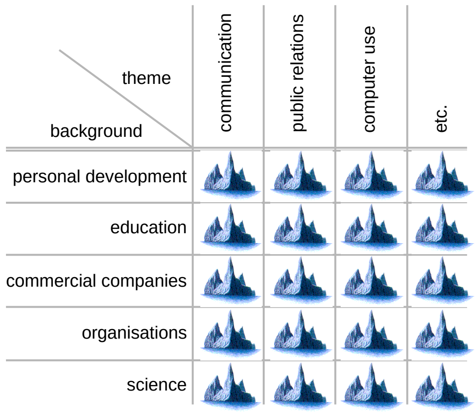

Fig. 1. Dimensions of Context

Each cell (iceberg) in the matrix corresponds to a context. In general a contextual dimension is a set of concepts with which the domain of application can be categorised.

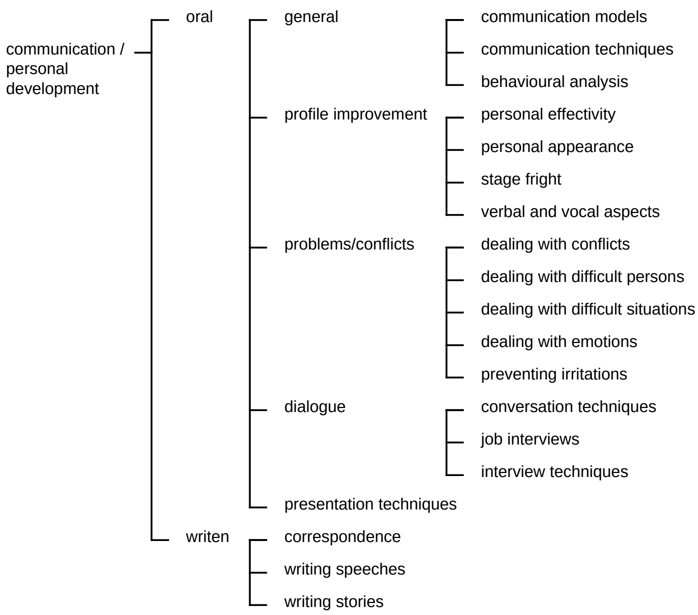

Fig. 2. Personal development x Communication

For example, by defining the theme dimension as the set {communication, public relations, computer use, others}, the domain of all courses can be categorised: each course can be associated to one or more themes. The topics in a context are also structured in some kind of semantic network, e.g., a taxonomy like that in Figure 2, for the context identified by <personal development, communication>. Of course also other techniques could be used to structure the concepts in each context, and the use of an additional dimension can reduce the size of the structures per concept.

Note that topics are not unique, they can occur in several contexts, and even a partial semantic network can occur in several contexts. For example, the topic presentation techniques occurs in the context <personal development, communication>, but also in all the other contexts underneath the column communication. Determination of the proper background is essential for finding the most appropriate courses on the theme communication. For example, a person might think he needs to follow a course on giving presentations, but in order to help him adequately, he might need help to realise that what he needs is a course on how to overcome stagefright.

3. The ICEBERG Broker Agent from a Global Perspective

The system developed is a multi-agent system, in which a broker agent, called the ICEBERG broker plays a central role. The ICEBERG broker mediates between information providers and users. The information providers have access to databases containing information on all available courses of the institute represented by the information provider. The ICEBERG broker agent that plays the central role in the system has been designed by reusing the generic broker agent model as presented in (Jonker, and Treur, 1998). The complete system has been modelled using DESIRE (see e.g., Jonker et al., 1998), the rules presented in the next sessions are part of the knowledge sources within the system, which executes them using chaining as the inference relation. The system has been tested with the software environment of DESIRE, which is based on Prolog.

There are three major phases in dealing with a user request: query (re)formulation, information resource discovery, and response construction.

If a new request from the user has been received by the broker, the broker first helps the user to (re)formulate the request: a precise answer needs a precise query. Given a "good" query the broker can solicit information form the appropriate providers (information resource discovery). If the ICEBERG broker received information on courses that might satisfy a user request, then it presents the most precise answers to the user. The broker also constructs additional propositions to the user, which may help the user to reformulate his request again (this is part of the response construction phase). The ICEBERG broker can thus be used in an iterative way.

4. Query Reformulation

Before the broker knows what to ask of the information providers, the broker first needs to analyse the question of the user. Most of the information in the user’s request is straightforward, e.g., in our opinion, the user knows quite well how much money he is willing to spend, how much time he is willing to invest, where (in which region, or city) he wants to follow the course. Therefore, most information can be used directly by the information providers to return courses that more or less fall into the required categories. The one piece of information that most often is the most problematic is the information about the course topic(s).

Self-evidently, the user is not yet an expert in the field he wants to follow a course about, therefore, he might not know the perfect topics that would result in the right course offers. Furthermore, a user might be interested in courses that cover a topic in a certain context. For example, a course on Java in the context of programming languages, is quite different from a course on Java in the context of studying various cultures, or following a course on presentations in the context of presenting business proposals does not help in the most ideal manner if you need the course to help you teach a group of students.

The broker analyses the topics of the course being searched. By making use of the domain-specific information, the broker is capable of determining the different contexts to which the course topic(s) can belong. The process consists of two parts:

1. request analysis
2. reformulation of user requests

4.1 Request Analysis

The request analysis is done in two stages: initialisation and determination of the set of contexts.

Initialisation

In the initialisation part of the first phase the first topic is taken from the user request and the contexts that contain that topic are identified as interesting contexts. Within the broker this knowledge is expressed declaratively as follows:

```
if      course_searched(partial_topic(T: TOPIC, 1), R: REQUEST_ID)
        and     occurs_in_context(T: TOPIC, C: CONTEXT)
then    interesting_context_for(CONTEXT, R: REQUEST_ID)
```

For example, the user entered the topic "presentation" which occurs in the context <communication, personal development>, but also some other contexts like <communication, science>.

If the user request contains 1 topic or if the set of interesting contexts is empty at the end of the initialisation, then the process continues with reformulation of the user requests. If the set is not empty, the complete set of contexts is determined.

Determination of the Set of Contexts

The set of contexts defined in the initialisation as interesting contexts are used as the basic set of contexts, from which the overall set of interesting contexts is determined iteratively. For each of the subsequent topics in the user request the set of interesting contexts is adapted. Each additional topic possibly reduces the set of interesting contexts. If the topic is not covered by one the contexts, this context is removed from the set of interesting contexts. The process continues until all topics in the user request have been analysed, or until the set of interesting contexts is empty. Consider the following definitions:

SAC is the set of all contexts in the multi-dimensional structure.
S is the set of interesting contexts produced in phase 1a
t_j is the j'th topic mentioned in the user request
n is the total number of topics in the user request
$SC(t_j) = \{ C \in SAC \mid t_j \text{ occurs in } C \}$

Formalised, the process is specified as follows:

```
j = 2;
Repeat
        PS := S;
        S := S ∩ SC(t_j) ;
        j := j + 1;
Until j = n or S = ∅;
```

Note that if the set S became empty in the above procedure, then PS is the last not empty set produced in this process. If, for example, the user entered also the topic "personal effectivity", then the set S will only contain the context <communication, personal development> after phase 1b, since this topic cannot be combined with presentations in the other contexts still available after phase 1a.

4.2 Reformulation of User Requests

The main goal of this process is to determine whether or not adaptations to the user question must be created (and which adaptations if so) and what the expected accuracy of the results for that user request is and to adapt the user request if necessary. The reformulation process operates according to three cases: S is empty after initialisation of the query analysis, S is empty after query analysis, or S contains more than 1 element. The set S would for example become empty if the user entered "presentations" as a first topic, and then "airplane", since the topic airplane is not contained in the contexts that are related to "presentations". In the prototype, the topic "airplane" is not supported at all, because there are no courses on airplanes within the theme communication. So, if airplane had been the first topic, then S would have been empty after the initialisation of the query analysis.

Case 1

If the set S of interesting contexts is empty after the initialisation of the query analysis, then the broker initiates an interaction with the user. The broker explains that it did not recognise the topics presented by the user (the expected accuracy of the search results is then 0) and offers him the choice:

- search for it anyway.
- adapt his request; for this the broker immediately provides the different contextual dimensions from which the user can choose. After a context has been selected by the user, the broker presents the topics in that context, enabling the user to focus his request. Of course, the user can retract his steps and choose a different context.

Case 2

Suppose that during query analysis, the set S of interesting contexts became empty. The expected accuracy of the search results for the request as formulated by the user is the fraction of recognised topics with respect to the total number of topics in the user request:

Definition

Let PS be the last not empty set of interesting contexts identified in phase 1. A topic t is recognised iff there is a context $C \in PS$, such that t occurs in C. Let NRT be the number of recognised topics in request RID, let TNT be the total number of topics in request RID. The expected accuracy (EA) of the search results for RID is defined by

$$EA(RID) = NRT(RID) / TNT(RID)$$

Given that the set of interesting contexts became empty in phase 1b, this fraction will be less than 1. Therefore, it is to be recommended that the user adapts his request. The proposed adaptations are based on the set PS, i.e., the last set of interesting contexts found in phase 1 before this set became empty. The contexts in PS cover the first topics in the user request and thereby determine possible contexts of the user interests. The broker presents his analysis to the user (the set of contexts based on the first topics, the conflicting topic that is not covered by those contexts, and the expected accuracy of the search results if the request remains as it is) and offers him the choice:

- search for it anyway, then the information providers have to receive the current user request.
- adapt his request. The user can choose one or more alternative topics from the contexts of the set PS. If necessary, the broker can also help him to reformulate his request more drastically, see the case that the set S is empty after the initialisation of the query analysis.

Case 3

Suppose that after query analysis, the set S of interesting contexts still contains more than one element. The request is now a bit ambiguous. The broker presents the user with the contexts in S and asks the user whether he wants to focus his request to one (or a few) of those contexts. The user is free to stick to the whole set S, but is cautioned that the search will then probably give too many results and that it will be hard to pick out the useful ones.

5 Information Resource Discovery

In the previous phase the broker has analysed the request of the user and on the basis of that analysis has formulated one or more questions to be posed to the information providers. In the second phase the broker determines which information provider is best suited to answer which question. The broker maintains profiles of the information providers. This profile contains the contexts for which the information providers has offered courses in the past.

Each question is matched to the profile of the different information providers. The match is made by checking the following items in the question against the corresponding information in the profile of the information providers:

- topics
- region
- price
- level (beginner, .., expert)
- language
- duration
- dates of the meetings

The user is not obliged to fill in all these items, but the items he did fill in are checked. If the match is favourable, the question is sent to that information provider.

6. Response Construction

The broker has the task of composing an appropriate report of the search results and of creating suggestions to the user with which the user can adapt his request if necessary. For this, the component needs information about the (reformulated) request from the user, and, of course, the search results as reported by the different information providers.

The broker has to match each found course against the user request, collect those courses that have a sufficiently high match, propose these best courses to the user, but also create other search possibilities that may help the user to improve his requests. If no courses have been found that match the request sufficiently, a failure report is constructed for the user, and also other search possibilities are formed and presented to the user. Matching is discussed a lot in literature, and nothing special needs to be said about the ICEBERG broker on this count.

After the acceptable results of the search are ordered and the best n selected, the end report of the search is constructed and presented to the user. For this the broker has to determine the overall successfulness of the search.

If the search did not fail, then the search will be reported to have been successful and other search possibilities are offered by presenting topics to the user that come from the same context as the topics requested by the user. These topics can focus the mind of the user on either more specific topics, or on topics with a slightly different focus (topics of the same specificity, for example, siblings in the taxonomy, if a taxonomy is used to model the contents of a context).

If no courses are found that could be of interest to the user, then the search is labelled to be failed and the user is to be helped formulating a more successful query. If the search failed and furthermore, the broker had made this prediction before the search was initiated, then the broker can explain to the user that according to the domain knowledge of courses available to the broker it was highly unlikely that the topics provided by the user would be the topics of an existing course, and that in this case the search did indeed fail. The broker then continues with a list of other search possibilities with the aim to make the user aware of the context of each topic in his request. If a topic occurring in a context is related to other topics in that context (some more specific, some more general), then presenting these related topics as possible search topics to the user, gives the user insight in the context determined by his request, and he can use this knowledge to adjust his request.

7. Other Work

We compared the IceBerg approach with some approaches in the field of Information Retrieval and Multi-Databases. We will first cover the literature on Information Retrieval systems.

It turns out that there are several differences but also some similarities. The difference are that IR (Information Retrieval) usually operates on generic domains and are not domain specific like IceBerg (Rijsbergen, 1979, Bodner et al., 1996).Furthermore IR systems employ a different search strategy. IR systems only search up and down in the search tree. IceBerg also searches between the siblings in the search tree. Another difference is that most IR systems automatically generate the ontologies with automated methods, like LSI (Deerwester et al., 1990) and WAIS (Kahle et al., 1991). IceBerg in contrast uses a human constructed ontology. Both IR systems and IceBerg use user feedback to facilitate query expansion. IR systems often use relevance feedback (Harman, 1992) to enhance the retrieval this contrary to IceBerg which uses direct ontology information to reformulate the query.

Like with the IR systems there are also differences and similarities between IceBerg and Multi-Databases. We will first point out the differences. Multi-Databases presume that the user knows or has domain ontologies available. Also as with IR systems Multi-Databases only employ a search strategy in the semantic network that searches up and down the network but not sideways. For Multi-Database access, the OBSERVER system (Mena et al., 1999), the user first selects the appropriate ontology (assumed to be known or available), and then a GUI is used to formulate the exact query (which is then assumed to represent exactly what the user needs). The IceBerg broker does not assume that the user knows the ontology, but assumes that the initial query contains enough information to find the relevant parts of the ontology. Then the IceBerg broker uses the different contexts related to that part of the ontology to help the user formulate a request corresponding to his needs. In the work of Kashyap and Sheth (e.g., Kashyap et al., 1995, Mena et al., 1999), the query has its own context which determines the semantics of the query and which is used directly to match with possible contexts of the different databases. In IceBerg, the interactive query formulation leeds to a query with a unique context as well. Although the role of a context in multi-databases is not exactly the same as that in IceBerg enough similarities exist to allow for a coupling of the multi-database techniques to IceBerg. The relationships between contexts in IceBerg can be compared to the semantic interontology relationships presented in (Mena et al., 1999).

Next to the comparison with work on Information Retrieval, Query Expansion, and Multi-Databases in general, the IceBerg broker is also compared to some of the recently developed and operational models of virtual market places, information brokering on the Web, and Web commerce based applications. For a good overview of recent developments in intelligent information agents see (Klusch, 1999).

The broker for environmental regulations (Stasiak, Garrett, and Fenves, 1999) aims at finding the right regulation from the user request by searching the regulation agency in its own database. It searches on topics mentioned in the user request. The broker presents the regulations found and the user selects one of them. The broker queries the database of the regulation agency with the selected regulation to search for more details. The information found is send back to the user. The topics used by the broker to search its database are defined by the classification on author based on the document specific classifications. Automatically generated classifications have to be made if the system is to be used commercially.

Kasbah (cf. (Chavez, and Maes, 1996), (Chavez, Dreilinger, Gutman, and Maes, 1997)) is a web-based multi-agent system using agents interacting with each other within the virtual market domain. The agents act on behalf of their users (CHAVEZ, AND MAES, 1996). Price Negotiation is one of the interesting features applied within Kasbah (Chavez, Dreilinger, Gutman, and Maes, 1997).

Market Space is an open agent-based market infrastructure. It is based on a decentralised infrastructure model in which both the humans and the machines can read information about the products and services, and everyone is able to announce interests to one another (Erikson and Finn, 1997). The aim in designing Market Space is to design a market place where searching, negotiation and deal settlement, e.g. interaction with users is done using agents.

The MeBroker (Doles, Dreger, Großjohann, Lohrum, and Menke, 1999) is an Information Broker project of the Freie Universität Berlin. The architecture of the system is resembles the one of ICEBERG. The difference is that the emphasis of MeBroker lies in the selection of the information provider. It uses meta-data on the information providers to make the best choice.

The Search Broker (Manber, and Bigot, 1999) performs its search in two stages. First the request is analysed to extract the topic from the request. The broker searches which server is the best to answer the request. The request is then passed to that server to hopefully find the right answer. New in this approach (in contrast to e.g. Yahoo) is that this happens in one regular search request. The index of the servers and topics is made by hand because the mapping is then more accurate.

Most brokers found on the Internet focus on the selection of the right information provider based on the given user request with an exception to a certain extent of the broker for environmental regulations. In contrast, the ICEBERG broker agent focuses on aiding the user in formulating the right request in an interactive process. The ICEBERG broker not only helps to sharpen the focus of the request, it also helps in disambiguation of the queries, widening the scope of the user request, and in other reformulations.

An additional difference with our approach is that these approaches have been implemented without using a principled design method, and do not use components as building blocks that are (formally) specified at a conceptual level. This is also a difference with the work described in (Martin, Moran, Oohama, and Cheyer, 1997), and (Tsvetovatyy, and Gini, 1996).

In this paper a multi-dimensional notion of context is defined to serve as necessary background knowledge for the broker. The multi-dimensional structure of contexts is used twice by the ICEBERG broker: for reformulation of the user request, and for evaluation of search results.

The clear structure in dimensions provided by the broker makes it easy for the user to find the best matching context for his request. Furthermore, each context is further elaborated by a semantical structure containing the topics relevant for that context. The ICEBERG broker exploits this structure to help the user find the best topics for his request. The broker does not only presents a relevant part of the multi-dimensional contexts if the user formulated an ambiguous request, but also if the user issued a request that points to a unique context. By presenting the relevant contexts to the user in all cases the broker ensures that the user and broker share the domain model. A shared domain model enables the broker to find the information that the user qualifies as good.

The multi-dimensional structure of contexts is also used to evaluate the search results delivered by the information providers. Given that the broker agent ensured a shared domain model of the context relevant for the user by exploiting the multi-dimensional structure of contexts, using that same structure to evaluate the results ensures that evaluation made by the broker matches to high degree the evaluation of the results by the user, and can be offered to the user within (and with indication of) the right context.

Further research can be done in automatically building and / or maintaining a multi-dimension structure of

Acknowledgements

This project was conducted for KPN Research. We like to thank Clair Moore, Ron van Riet, Louis Wennekes and Bertjan Teunissen from KPN Research from the project team for their input during the project and Jan Treur for his comments on this paper. We also like to thank the anonymous reviewers for pointing out a wealth of literature that we were unaware of. We learnt a lot.

References

Bodner, R., and Song, F. (1996), Knowledge-based approaches to query expansion in information retrieval. In McCalla, G. (Ed.), Advances in Artificial Intelligence (pp. 146-150), New York: Springer.

Boles, D, Dreger, M, Großjohann, K, Lohrum, S, and Menke, D, (1999), MeDoc Architektur und Funktionalität des MeDoc-Dienstes, Frei Universität Berlin. http://www.inf.fu-berlin.de/~medoc3/

Chavez, A., and Maes, P., (1996), Kasbah: An Agent Marketplace for Buying and Selling goods. In: *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology, PAAM'96*, The Practical Application Company Ltd, Blackpool, pp. 75-90.

Chavez, A., Dreilinger, D., Gutman, R., and Maes, P., (1997), A Real-Life Experiment in Creating an Agent Market Place. In: *Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology, PAAM'97*, The Practical Application Company Ltd, Blackpool, pp. 159-178.

Deerwester, S., Dumais, S.T., Furrnas, G.W., Landauer, T.K., and Hashman, R, (1990), Indexing by Latent Semantic Indexing. In: Journal of the American Society for Information Science, 41(6).

Erikson, J. and Finn, N., (1997), Market Space: an open agent-based market infrastructure. Master's Thesis, Computer Science Department, Uppsala University, Sweden.

Gruber, T., (1993), What is an Ontology? Knowledge Systems Laboratory, Computer Systems Dept., Stanford University, Stanford, CA94305, USA.

Gruber, T., (1997), Toward principles for the design of ontologies used for knowledge sharing. In: *Formal Ontology in Conceptual Analysis and Knowledge Representation*, N. Guarino and R. Poli, editors, Kluwer Academic Publishers.

Harman, D. (1992), Relevant Feedback Revisited, In: Proceedings of the 15th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval, Copenhagen, pp. 1-10.

Jonker, C. M., and Treur, J., (1998), Compositional Design and Maintenance of Broker Agents. In: J. Cuena (ed.), *Proceedings of the 15th IFIP World Computer Congress, WCC'98, Conference on Information Technology and Knowledge Systems, IT&KNOWS'98*, pp. 319-332.

Kahle, B, and Medlar, A., (1991), An Information System for Corporate Users: Wide Area Information Servers. In: Connexions - The Interoperability Report, 5(11).

Kashyap, V., and Sheth, A. (1995), Schematic and Semantic Similarities between Database Objects: A Context based Approach, Technical Report TR-CS-95-001, LSDIS Lab, Department of Computer Science, University of Georgia.

Klusch, M., (Ed.), (1999), Intelligent Information Agents Agent-Based Information Discovery and Management on the Internet. Springer-Verlag Berlin/Heidelberg. 498 pp.

Manber, U., and Bigot, P. A., (1999), The Search Broker, Department of Computer Science, University of Arizona. http://sb.cs.arizona.edu/sb/paper.html

Manglano, V., Beaulieu, M., and Robertson, S., (1998), Evaluation of Interfaces for IRS: modelling end-user searching behaviour. In: 20th Collogquium on Information Retrieval, Grenble, http://www.weic.org.uk/.

Martin, D., Moran, D., Oohama, H., and Cheyer, A., (1997), Information Brokering in an Agent Architecture. In: *Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology, PAAM'97*, The Practical Application Company Ltd, Blackpool, pp. 467-486.

Magennis, M. (1995), Expert rule-based query expansion, Presented at the British Computer Society Information Retrieval Specialist Interest Group Colloquium. Crewe, England. April 1995. http://ir.dcs.gla.ac.uk/publications/papers/Year/1995.html#magennis95a

Mena, E., Illarramendi, A., Kashyap, V., and Sheth, A. (1999), OBSERVER: An Approach for Query Processing in Global Information Systems based on Interoperation across Pre-existing Ontologies, Distributed and Parallel Databases Journal.

Rijsbergen, C.J. van, (1979), Information Retrieval, second edition, London: Butterworths.

Stasiak, D., Garrett, J., Jr., and Fenves, S., (1999), A Broker for Environmental Regulations, Carnegie Mellon University. http://www.ce.cmu.edu/GreenDesign/ research/broker.html

Tsvetovatyy, M., and Gini, M., (1996), Toward a Virtual Marketplace: Architectures and Strategies. In: Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology, PAAM'96, The Practical Application Company Ltd, Blackpool, pp. 597-613.