National University of Ireland, Galway
Ollscoil na hÉireann, Gaillimh

Department of Information Technology

technical report NUIG-IT-101002

An Architectural Design for Multi-Agent Information Trading

P. McDermott (NUI, Galway)
C. O'Riordan (NUI, Galway)

itreports@nuigalway.ie http://www.it.nuigalway.ie/

An Architectural Design for Multi-Agent Information Trading

Paula Mc Dermott
Department of IT
NUI, Galway
paula.mc_dermott@geminga.nuigalway.ie

Colm O'Riordan
Department of IT
NUI, Galway
colmor@geminga.nuigalway.ie

Abstract

In this paper we discuss a multi-agent information trading system, which we argue provides a suitable and flexible design for distributed IR, and a suitable test-bed for studying agent cooperation, coordination and negotiation. Within the fields of Information Retrieval (IR) and Information Filtering (IF), agents represent a suitable paradigm to conceptualise, design and implement an IR system.

1 Introduction

Information overload is a term frequently used to describe the vast amount of information available to readers and the difficulty involved in identifying relevant material. Research into this problem has focused on improving retrieval techniques.

Many techniques have been developed to aid users in dealing with information overload. These techniques, and systems, attempt to analyse the content of documents and present only those that are relevant to a user. Techniques employed to date include the Boolean model, Vector Space model and AI-inspired approaches.

With the increased proliferation of data repositories, techniques are required to effectively query a number of distributed repositories. Problems in this domain include source selection (which repository is most appropriate for a given query) and result merging (how to combine results from different sites). Our system provides a model to incorporate these two problems within a MAS paradigm.

Multi-agent systems consist of a number of agents (autonomous, intelligent pieces of software [5]) that are able to interact with each other and the environment. These agents may possibly differ from each other in their skills and their knowledge [13]. This involves many problems including those of cooperation, coordination and negotiation.

The paper is laid out as follows. Section 2 outlines research in IR and distributed IR. Research and development in MAS is outlined in section 3. The following section discusses the design of our system. Section 5 describes current and future experiments. Section 6 summarises and offers some conclusions.

2 Information Retrieval

The International conference on scientific information in Washington 1958 is said to mark the beginnings of IR as it is now known. The task of an Information Retrieval system is to retrieve documents containing information that is relevant to a users information need [1]. This task is a difficult one as it usually deals with natural language which is not always well structured and may be semantically ambiguous [2].

2.1 Approaches

The problem of satisfying an information need can be approached in several ways. These approaches are termed models and range from simple string matching techniques to more complicated neural network based approaches. The task of IR can be broken down into three subtasks, namely, representation, comparison and feedback.

2.1.1 Representation

The user of a retrieval system has to translate their information need into a query. This normally implies specifying a set of words which convey the semantics of the information need [2]. The documents in the system must also be represented in a machine readable manner, such that the system can perform comparison techniques between the document set and the user information need. Representation techniques range from using indexes, vector and matrix representation to more modern representations, such as neural networks, connectionist networks and semantic networks [3].

2.1.2 Comparison

This process involves comparing the user query to the document set in order to obtain a measure of the similarity between documents and the query. The comparison mechanisms used are dependent on the underlying representations employed in the system. There are three classic representations (models) in IR, namely, Boolean, Vector and Probabilistic.

Boolean Model
In this model, the queries and documents are represented as sets of index terms. Index terms are either present or absent in a document, as a result the index term weights are assumed to be all binary i.e., $w_{i,j} = \{0, 1\}$. The user query is composed of keywords linked together using the Boolean primitives AND, OR and NOT.

Vector Model
The vector model, recognises the limitations of the Boolean model and proposes a framework in which partial matching can be performed. Non-binary weights are assigned to index terms in queries and in documents. The best known term-weighting formula was suggested by Salton:

$$w_t = f_{it} log \frac{N}{N_t}$$

N being the number of documents in the collection, N_t the number of documents in the collection containing the term t and f_{it} the number of times t occurs in the document i [3]. The vector space model evaluates the degree of similarity between d_j and q by computing the cosine of the angle between these two vectors. Extensions and variations have been proposed - these include alternative weight schemes [28], clustering of document vectors (self-organising maps [29], k-means clustering [30]) and LSI [6] (wherein the dimensions of the vectors are reduced to improve effectiveness and hopefully accuracy).

Probabilistic Model
The classic probabilistic model was introduced in 1976 by Robertson and Sparck Jones [10]. The probabilities $P(d|R)$ and $P(d|\overline{R})$ (the probability that a document is relevant or irrelevant to a query) can be calculated based on the probabilities $P(t_i|R)$ and $P(t_i|\overline{R})$ (the probability that a term t_i describes a relevant or irrelevant document). Sample implementations of the probabilistic theory include inference networks [11] and belief networks [9].

2.1.3 Feedback

The third and final stage in an Information Retrieval system is that of feedback. Users of IR systems may have little or no knowledge of the operations of the system or the documents in the collection. This may lead to difficulties in formulating queries which allow the system to meet their information needs. Also, due to the nature of IR the returned results will be a partial match to the users query. Mechanisms are required to improve the accuracy of the returned set. This is achieved by query reformulation and may be performed by the user but is largely automated by IR systems. It involves two steps: expanding the original query with new terms and re-weighting the terms in the expanded query. This process allows the system to modify the search criteria and thus return more suitable documents to the user.

2.2 Distributed Information Retrieval

As document collections grow larger they become more costly to manage with an information retrieval system. Searching and indexing costs increase with the size of the underlying document collection. Larger document collections invariably result in longer response times [4]. To support the demanding requirements of modern search environments, a need for alternative architectures and algorithms has emerged. This has led to the increased adoption of both distributed and parallel architectures.

A more important factor within distributed information retrieval is the proliferation of distributed information sources. Modern information environments have become large, open and heterogeneous [12]. This has increased the need to develop distributed IR systems.

Commonly identified problems within distributed IR are those of source selection and result merging. With respect to source selection many alternatives exist. If the collection is semantically partitioned, rule-based approaches may be used to query only relevant sites. If, however, the document collection is not structured in this manner, as is becoming more common, we need means to select the relevant sources. Approaches to date have involved indexing centrally all document collections and at query time using this index to further query a subset of sites. Oftentimes, such centralised statistical information is not available and means are needed to identify suitable sources.

A related problem is that of result merging or fusion, where the ranked results from a distributed set of sites must be merged into a complete ranked list. Voorhees et al. [27] proposed strategies based on modelling sites using training queries and by query clustering. Callan et al. [26] provide a weighting scheme whereby the ranked list is weighted according to the belief that a given site is relevant.

In recent years, agents have been developed and 'released' into information management environments in an attempt to improve the mechanisms for finding, fusing, using, presenting, managing and updating information. Work by Sheth [32], Payne [31] and Kilander et al. [33] has directly applied agents to information retrieval, filtering and browsing. The following section details the main concepts within MAS research.

3 Multi-Agent Systems

Software agents date back to the early days of AI work and Carl Hewitt's concurrent actor model [7]. In this model, Hewitt proposed the concept of a self-contained interactive and concurrently executing object which he termed an actor. The object had some encapsulated internal state and could respond to messages from similar objects [5]. In recent years interest in agent technology has increased and as a result several definitions of an agent have been proposed. Wooldridge [24] states:

> An agent is a computer system that is situated in some environment, and that is capable of autonomous action in this environment in order to meet its objectives.

The interest in multi-agent systems is largely founded on the insight that many real world problems are best modelled using a set of agents instead of a single agent. In particular, multi-agent modelling makes it possible (i) to cope with natural constraints like the limitation of the processing power of a single agent or the physical distribution of the data to be processed and (ii) to profit from the inherent properties of distributed systems like robustness, fault tolerance, parallelism and scalability. Generally a multi-agent system is composed of a number of agents that are able to interact with each other and the environment and that differ from each other in their skills and their knowledge about the environment [13].

Whether an agent is collaborating or competing with another agent, interaction between the agents will occur. The agents will need to communicate with each other, to be aware of each other and to reason about each other. It is clear that the concept of a multi-agent system is far more complex than that of an individual agent. The complexity introduced by agent-agent interactions has created, among others, the research areas of agent cooperation, coordination, coalition formation and negotiation.

3.1 Cooperation

Cooperation is often presented as one of the key concepts which differentiates multi-agent systems from other related disciplines such as distributed computing, object-oriented systems and expert systems [14]. Given a multi-agent system wherein the individual and the subgroups are assigned one or more (possibly implicit) goals, cooperation occurs when the actions of each agent satisfy either or both of the following conditions:

- the agents have a (possibly implicit) goal in common, (which no agent could achieve in isolation) and their actions tend to achieve that goal.
- the agents perform actions which enable or achieve not only their own goals, but also the goals of agents other than themselves [14].

Notions of shared or joint goals and notions of complimentary individual action have been posited. Theories from formal logics [15], philosophy [23] and game theory [22] have been adopted to reason about cooperation and to design and develop models whereby agents cooperate.

3.2 Coordination

Coordination is a central issue in software agent systems and in Distributed Artificial Intelligence (DAI) [16, 17]. Coordination is required in a multi–agent system for many reasons, including: preventing anarchy or chaos, meeting global constraints, efficiency, distributed expertise resources or information and dependencies between agents' actions. It is important to note, at this stage, that coordination may require cooperation but cooperation will not necessarily result in coordination.

Coordination Strategies
Many coordination strategies have been developed, Nwana, Lee and Jennings [16] categorise them into four broad categories:

- Organisational Structuring: this is the simplest coordination scenario which exploits the *a priori* organisational structure [16]. The organisation defines implicitly the agents responsibilities, capabilities and control flow.
- Contracting: The contract net protocol for decentralised task allocation is one of the most important paradigms developed in DAI [18]. When an agent needs to find another agent to deal with a problem that the former cannot solve, it broadcasts a contract announcement for that problem. The broadcasting agent then awards the contract to the agent with the best bid.
- Multi-Agent Planning: Within this approach agents build a multi-agent plan that details all their future actions and interactions required to achieve their goals. The system operates in an iterative manner, with task execution, planning and re-planning. Multi-agent planning can be broken into centralised (one agent gathers all partial plans and analyses them) and decentralised (each agent has a model of all other agents plans) multi-agent planning.
- Negotiation: Negotiation is involved in most of the other coordination strategies in some way. Much of the work on negotiation strategies has its roots in game theory: utility values for each outcome of an interaction are inserted into a payoff matrix which can (usually) be viewed by both parties. An interactive process of offers and counter-offers then ensues, in which each agent chooses a deal which maximises its expected utility value.

3.3 Coalition Formation

Work on negotiation ranges from bilateral negotiation (agreements reached between two agents) to n-agent agreement. In the case of bilateral negotiation there are two options: the agents reach an agreement, or they do not. On the other hand, n-agent agreements result in a more complicated scenario. It is these n-agent scenarios which will be the focus of this section.

A coalition can be defined as a group of agents who have decided to cooperate in order to perform a common task [19]. Shehory, Sycara and Jha [19] believe that the incorporation of a coalition formation mechanism will increase the efficiency of group-wise task execution, resulting in near-optimal task performance. In general agents will only form a coalition if each member of the coalition gains more by joining the coalition than by working alone.

The work in this domain ranges from sub-additive coalitions (Zlotkin & Rosenschein [21]) to greedy strategy coalitions (Shehory, Sycara & Jha [19]) and those that are based on agent negotiation (Reusable Task-based System of Intelligent Networked Agents).

4 Design - Information Trading

The primary concern of this paper is the application of agent-based systems to the domain of information management. Weiss [13] describes information agents as: agents that have access to multiple, potentially heterogeneous and geographically distributed information sources. Information agents have to cope with the increasing complexity of modern information environments, ranging from relatively simple in-house information systems, through large-scale multidatabase systems, to distributed information and knowledge repositories in the Internet. One of the main tasks of the agents is an active search for relevant information in non-local domains on behalf of their users or other agents. This includes retrieving, analysing, manipulating and integrating information available from different information sources.

The purpose of the system described and proposed in this paper is twofold. Firstly, the system acts as an agent based information trading system and secondly it forms a test bed and framework for testing various cooperation, coordination and coalition strategies. Within this framework, agents may have the potential to learn about the information needs, about the other agents abilities etc. As sources become more distributed, we have a set of information providers, with their set of repositories and an associated cost to access this information. We envisage a set of agents satisfying a set of needs by trading information.

4.1 Agent Community

The concept of an agent as an autonomous entity and as part of a larger team has been discussed in a previous section. Agents have been defined as having knowledge, beliefs and capabilities.

4.1.1 Agent Architecture

The agents in this system will have the following architecture:

Knowledge
The agents' knowledge base includes the concepts of users and their information needs. Upon registration the user is assigned an agent and is expected to supply information needs to the agent. The information need is essentially the task that the agent is obliged to carry out for its user. The utility is a measure of whether or not the agent can 'afford' to employ another agent to help with the task. Another way to look at the utility is as a measure of how satisfied the user is with the agent. If the user is happy with the agents work he/she will have greater trust in the agent and will increase its utility. This increase will allow the agent more freedom to employ and 'trade' with other agents thus increasing their knowledge base and their ability to satisfy the users information need.

Beliefs
This part of the agent's architecture details the agent's opinions of one another. The belief rating refers to one agent's level of belief in another agents ability to perform a task. The trust rating, on the

other hand, relates to an agents belief that the agent will actually carry out the task. This section is an important part of the agent architecture. The belief and trust levels that an agent has in another agent will be taken into account when an agent is considering who is most suitable for a particular task.

Capabilities

The capabilities section refers to the tasks that the agents will be able to perform. All agents will be capable of carrying out the same tasks (including retrieval, learning and trading) but they will differ in their ability to perform them.

4.2 Interactions

The system is relatively straight-forward if the agents are capable of performing the task specified by their users. They simply query their own repositories, using their IR module, and return the retrieved documents to their user. This acts as a set of separate centralised IR systems. The more complicated and interesting aspects of the system can be seen if it is cheaper or easier for another agent to perform the task or if an agent is unable to perform the task required (due to poor capabilities or repositories). If, for example, the agent in question does not have access to the information needed to fulfill its task(s) then it must employ the help of another agent. This is quite a complex task and there are many possibilities that could occur.

4.2.1 1-1 Agent Interactions

In this scenario there are two possibilities that must be considered, namely, the agents can cooperate with each other or not. In order to cooperate the agents must be able to communicate their actions to one another. They must be able to:

- Propose a course of action
- Reject a course of action
- Accept a course of action
- Counter propose a course of action

The propose and accept performatives will allow the agents to cooperate with one another and the addition of the counter propose performative allows the agents to negotiate with one another. The propose and counter propose performatives are essentially doing the same tasks.

4.2.2 N-Agent Interactions

This scenario introduces many more possibilities. The agents must be able to communicate with one another, the performatives described in the last section still hold for this case. The agents will again be capable of cooperation and negotiation. The possibilities of the agents coordinating their actions and forming coalitions must also be considered.

When an agent performs a task it expects its utility to increase. When a group of agents work together on a task they still expect to be paid for their part in the operation. This introduces the notion of a payoff scheme. The bonus available for the task is still the same regardless of how many agents work on the task. The bonus must be split in some way between the agents. Several coalitions may form within one agent community. There can be coalitions within coalitions. It should be noted that the agents are under no obligation to join in any coalition or even to cooperate with other agents. It is possible that an agent may not want to share the payoff with other agents and therefore undertakes to do the task by him/herself (the concept of a 'selfish' agent). Although the agent will get the full payoff from this task it may take him longer to complete as he works alone.

In summary, we have a set of repositories, agents and users. These can be easily expanded if required. Primitives and protocols exist which allow retrieval, trading and coordination schemes (coalition formation). We envisage that this architecture can be reused with different mechanisms in place in order to compare approaches.

5 Experimental Setup

5.1 Information Sources

The information collections used in the system are standard IR test collections. A collection of abstracts from the communications of the association of computing machinery (CACM), the CISI collection and the Cranfield collection. These collections provide the agents with over five thousand documents to trade.

5.2 Information Retrieval Systems

Each agent in the system has the ability to search its own document collection. We require that different agents have different capabilities so we have several IR engines. The IR modules available in the system are: the Boolean model, the Vector Space model and the INFOrmer system [25]. The Boolean model and the Vector Space model have been discussed in a previous section. The third and final model available to the agent community is the INFOrmer system which uses a semantic network representation and spreading activation techniques to effect the comparison [25].

5.3 Multi-Agent System Protocols

As mentioned the system has been designed with its use as a test bed in mind. The main motivation behind this approach was to allow for the implementation of various agent cooperation, coordination and coalition formation strategies. It was decided that a possible manner in which to test these strategies would be to slowly increase the number of agents in the system and monitor the changes that occur. The different scenarios that will be examined are 2-agent scenarios. 3-agent scenarios and n-agent scenarios. We hope to measure complexity (messages passed), fitness and user satisfaction.

5.3.1 2−agent Scenarios

Within the 2-agent scenario we can consider the possibility of the agents cooperating with one another or working in isolation. The agents have the ability to communicate, propose, reject, accept and counter propose courses of action. These performatives allow the agents to cooperate with one another by proposing a course of action and accepting a course of action. We can test various cooperation strategies and examine the behaviour of the system to gain an indication of how the agents react to the differing strategies. These results can then be compared and the 'best' strategy discovered.

5.3.2 3−agent Scenarios

If we consider a situation where there are three agents in a system more complexities will come into play. As there are three agents in the system their actions will need to be coordinated in order for the community to maintain a stable state. We will be able to introduce the possibility of coalition formation to the system. It is possible that two of the agents will form a group (coalition) and work together to achieve their goal(s), thus excluding the other agent from the coalition. The remaining agent is then left to work in isolation to achieve its goal.

An agent may want to form a coalition with another agent or set of agents if it is unable to perform the task required of it. In order to form a coalition the agent has to carry out the following steps:

- Decide what agent(s) should be in the coalition. This decision will be made based on the belief and trust ratings described in the agents architecture.
- The agent will then broadcast a message to the selected agents offering them the opportunity to join the coalition
- The agent then waits for responses from those agents. These will either be accept or reject performatives
- The next step may involve negotiation. If the accepting agents are happy with the terms of the coalition then there will not be any need for negotiation. On the other hand if the agents are not satisfied with the terms of the contract then they will try to re-negotiate with the offering agent. This cycle can have several iterations

until the agents decide to join or to finally reject the coalition

- The contract is implemented. Each agent in the coalition carries out the task that they were employed to do and the overall task is completed.
- The final task to be carried out involves the modification of the belief and trust ratings. The agent that proposed the formation of the coalition must now examine his belief and trust ratings in the agents and adjust these based on the agents performance in the coalition.

The n-agent scenario allows a more complex community to develop. It is possible that several coalitions develop and the agents could move in and out of coalitions. This scenario is mainly a larger version of the 3-agent one. As there are more agents in the system there will be more opportunities available. The agents can be more selective in the coalitions that they join. It will also introduce the concept of a 'price war'. As there are more agents it should be easier to find one that is capable of performing the task required at the 'right price'.

5.4 Summary

5.4.1 Experiments

Initial experiments involve the simulation of a distributed multi-user IR system wherein users information needs are satisfied via a set of information agents. A number of information repositories will be created and will be accessible by some or all of the agents.

We can easily implement a naive round robin merging mechanism. Based on relevance feedback (automated using available human judgements) we can learn to, not only modify the user query as in traditional IR, but can also learn to attach more trust in certain sites and their associated agents. Similarly we can, over time, learn to identify suitable sources to query.

Irrespective of the learning algorithm employed by the individual agents and or the society of agents we have an empirically sound means of evaluation.

This setting allows us to categorically quantify the suitability of agent interaction protocols such as coalition formation schemes and trust and reputation models commonly adopted within the agent community. Initial agent experiments will involve local updatings of belief and trust models possessed by individual agents, experiments regarding coalitions and normative constraints can be easily incorporated into the existing system.

In summary we propose initially to build a simplistic IR system to use for baseline comparisons. Subsequent incorporation of agent interaction and associated protocols will lead to the production of results regarding the suitability of these approaches.

Motivations behind the development of such a system include the creation of a test-bed for the examination of agent interactions and the provision of a distributed information retrieval architecture which provides an alternative to and possibly better approach than, existing designs.

The system, as it is designed, allows users, n agents, n information needs and m IR modules to interact within the agent community. The n agents satisfy the users by providing information, through the use of m different IR modules.

As mentioned, we propose that the architecture developed in this system may subsume other distributed IR models. Problems associated with existing distributed approaches include source selection and result merging. The application of the multi-agent paradigm provides a means of overcoming these difficulties. Over time the agents will learn about the quality of information and the ability of an agent to retrieve this information. Before an agent decides to 'employ' another agent it will examine its belief and trust rating in that agent. If the agent is deemed to provide accurate information in an appropriate manner it will be selected to aid in the task. The belief and trust that each agent has in each other agents information collection and ability provides a means of automatic source selection. Agents with access to information of a poor quality or those using simplistic retrieval methods will not be invited to perform tasks with, or for, other agents.

6 Summary & Conclusion

The areas of IR and multi-agent systems represent different areas of study but can be combined in order to provide a robust, extendable, intelligent, potentially distributed information management system. We believe that this approach provides a unique method of dealing with the traditional IR problem. The agents provide an intelligent module, in which they have the ability to communicate with, learn from, trust and distrust each other. We believe that this process of information trading provides a better quality of service to the end user. The multi-agent paradigm undoubtedly extends the capabilities of IR modules in a distributed environment. Immediate future work will involve the completion of experimentation using existing coordination and coalition schemes.

7 References

[1] Sparck Jones, K., Willett, P., (Eds.) 1997. Readings in Information Retrieval. Morgan Kauffman Publishers Inc. San Francisco, California.

[2] Baeza-Yates, R., Ribeiro-Neto, B., 1999. Modern Information Retrieval. Addison-Wesley ACM Press, New York.

[3] O' Riordan, C., Sorensen, H., 1999. Information Retrieval and Filtering An Overview, in Dept of IT Journal, NUI, Galway.

[4] Brown, E., 1997. Parallel and Distributed Information Retrieval. in Baeza-Yates, R., Ribeiro-Neto, B., 1999. Modern Information Retrieval. Addison-Wesley ACM Press, New York.

[5] Nwana, H., Ndumu, D., 1995 An Introduction to Agent Technology. BT Labs.

[6] Dumais, S., 1993. Latent Semantic Indexing (LSI) and TREC 2. Proceedings of the 2nd Text Retrieval Conference

[7] Hewitt, C., 1997. Viewing Control Structures as Patterns of Passing Messages, in Artificial Intelligence.

[8] Salton, G., A Vector Space Model for Automatic Indexing in Readings in Information Retrieval. Sparck Jones, K., Willett, P., (Eds.) 1997. Morgan Kauffman Publishers Inc. San Francisco, California.

[9] Riberio-Netto, B., Muntz, R., 1996. A Belief Network Model for Information Retrieval, in Proceedings of the 19th ACM SIGIR Conference on Research and Development in IR.

[10] Robertson, S., Sparck Jones, K., 1976. Relevance Weighting of Search Terms. Journal of the American Society for Information Sciences

[11] Turtle, H., 1990. Inference Networks for Document Retrieval. PhD Thesis.

[12] Huhns, M., Singh, M., (Eds.) 1998. Agents and Multi-Agent Systems: Themes, Approaches and Challenges in, Huhns, M., et al. Readings in Agents.

[13] Weiss, G., 1998 Learning to Coordinate Actions in Multi-Agent Systems, in Huhns, et al. Readings in Agents.

[14] Doran, J., Franklin, S., Jennings, N., Norman, T., 1997. On Cooperation in multi-agent systems. The Knowledge Engineering Review, 12(3).

[15] Wooldridge, M., Jennings, N., 1996. Towards a Theory of Cooperative Problem Solving. Proceedings of Modelling Autonomous Agents in a Multi-Agent Worls (MAAMAW-94).

[16] Nwana, H., Lee, L., Jennings, N., Coordination in Software Agent systems, Intelligent systems

Research Group, Applied Research and Technology, BT Labs.

[17] Jennings, N., 1996. Coordination Techniques for Distributed Artificial Intelligence. in O'Hare, G., Jennings, N., (Eds.) Foundations of Distributed Artificial Intelligence, Wiley.

[18] Sandholm, T., 1993. An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations. Proceedings of the National Conference on Artificial Intelligence, (AAAI).

[19] Shehory, O., Sycara, K., Jha, S., 1998. Multi-Agent Coordination through Coalition Formation. Proceedings of Agent Theories Agent Languages and Mechanisms.

[20] Shehory, O., Kraus, S., 1995. Coalition Formation Among Autonomous Agents: Strategies and Complexities. Proceedings of the 14th International Conference on Artificial Intelligence.

[21] Zlotkin, G., Rosenschein, J., 1994. Coalition, Cryptography and Stability: Mechanisms for Coalition Formation in Task Oriented Domains. Proceedings of the 12th National Conference on Artificial Intelligence.

[22] Axelrod, R., 1984. The Evolution of Cooperation. Basic Books, New York.

[23] Tuomela, R., 2000. Cooperation: A Philosophical Study. Philosophical Studies Series, Kluwer Academic Publishers.

[24] Wooldridge, M., Intelligent Agents, in Multiagent Systems A Modern Approach to Distributed Artificial Intelligence. Weiss, G., (Ed.) 2000. MIT Press, Cambridge Massachusetts.

[25] O'Riordan, C., Sorenson, H., 1999. An Agent-Based System for Intelligent Collaborative Filtering in Proc. of the third Int. Workshop on Co-operative Information Agents, Uppsala, Sweden.

[26] Callan, James P., Zhihonh, Lu W., Croft, B., 1995. Searching Distributed Collections with Inference Networks. Proceedings of the 18th Int. ACM SIGIR Conference on Research and Development in Information Retrieval

[27] Voorhees, E., Siemens Trec-4 Report. Further Experiments with Database Merging, in Harman, D K., Proceedings of the 4th Text Retrieval Conference, 1998. pages 121-130.

[28] Salton, G., Buckley, C., 1998. Term Weighting Approaches in Automatic Text Retrieval. Information Processing and Management, 24(5):513-523.

[29] Kohonen, T., 1990. The Self-Organising Map. Proceedings of the IEEE, 9, 1464-1479.

[30] Hartigan, J., Wong, M., 1979. A K-Means Clustering Algorithm.

[31] Payne, T., 1994. Learning E-mail Filtering Rules with Magi, A Mail Agent Interface. MSc. Thesis, Dept. of Computing Science, University of Aberdeen, Scotland.

[32] Sheth, B., Maes, P., 1993 Evolving Agents for Personalised Information Filtering. in Proceedings of the 9th IEEE Conference on Artificial Intelligence for Applications.

[33] Kilander, F., Fahraeus, E., Palme, J., 1997. PEFNA, The Private Filtering News Agent. Dept. of Computers and Systems Sciences, Stockholm University, Stockholm.