COOPERATIONS AND FEDERATIONS OF TRADERS IN AN INFORMATION MARKET

Michael Christoffel
Institute for Program Structures and Data Organisation
Universität Karlsruhe
Am Fasanengarten 5
76149 Karlsruhe, Germany
christof@ira.uka.de

Abstract

We observe the increasing value of information in our society. A consequence is the development of an information market, where the "product" information is dealt under the law of supply and demand. In this paper, we present a trading service for an information market, which assists in the location and selection of appropriate providers for the special demand of a customer. One main idea described in this paper is to increase the efficiency of the single traders by cooperations and federations of traders. We show the realisation of our approach in an agent platform for a market of scientific literature.

1 Introduction

Together with the development of an information society, we observe the increasing value of information. Not only science, but also many other professions depend on a steady supply with newest information. The Internet has accelerated this development, since it allows the fast creation, duplication, and distribution of information world-wide.

The consequence is the development of an information market, where the “product” information is dealt under the law of supply and demand (Shapiro and Varian, 1999).

In opposite to material goods, information can be digitised and edited very easily. This fact favours electronic commerce, because the "product" information can be delivered to the customer via the Internet directly.

There are other significant differences between information markets and traditional markets, such as relatively high production and low distribution costs, and the obstacle to assess the value of information (one can determine the value of a piece of information only if one has read this information). A major topic is information overload (Varian, 2000).

A customer in this market is often unable to locate providers for his/her special demand and to decide which provider is the most appropriate. This selection considers not only the product itself and the price; quality of service attributes are equally important. For instance, a cheep providers who supplies tomorrow is worthless if the information is needed today.

In this paper, we present a trading service for an information market. This service deals exactly with the before mentioned problem: how to select a provider, how to bring together supply and demand. We will discuss a definition of this service and the most important challenges in section 2. A more general description of mediation services within electronic markets can be found in Guttman et al (1999).

We will emphasise the idea of increasing the potential of individual traders by cooperation with other traders. This cooperation is described in section 3.

A special kind of cooperation is the installation of fixed relations between traders, forming a trader federation. This approach is discussed in section 4.

In section 5, we present two different direction of a realisation of our ideas. The first direction is a simulation study, the second direction is the application in a platform for special kind of information market. The scenario we will present is search and delivery of scientific literature under market aspects. This work is part of the UniCats project, where we develop an agent framework capable the requirement of this information market. The UniCats project is funded by the German Research Foundation (DFG) as a part of the national strategic research offensive “Distributed Processing and Delivery of Digital Documents (V^3D^2)”.

Section 6 gives an overview on related approaches concerning trader cooperations and federations. Finally, in section 7 we give a conclusion of this paper.

2 Traders in an Information Market

In this section we will discuss the requirements for a trading service in a open information market. We state three qualities of this market:

- Distribution: Market participants can be located anywhere in the Internet.
- Heterogeneity: Market participants may be different in their properties, behaviour, and expectations.
- Dynamism: The market may change at any time. Especially, market participants are free to enter and to leave.

A simple and general definition of a trader can be found in Bearman (1993): “A trader is a third party object that enables the linking of clients and servers in a distributed systems.” Transferred to the market scenario, we can define: A trader is a market instance that brings together supply and demand. It does this by the selection of those providers appropriate for a given customer demand. So the main task of a trader is to remedy the lack of market transparency. A theoretical discussion of trading services in information markets can be found in Rose (1999).

In order to select one provider or a couple of providers from all providers in the market, the trader has to hold information about these providers. For each provider, it holds a profile that contains collected metadata. The selection of appropriate providers is done in a query handling process, where the available profiles are tested in which degree they match the customer query. The trader answers with a ranked list of provider addresses enriched with additional information about the providers, which is generated out of the profiles.

When handling a query received from a customer, the trader calculates for each available provider the conformity of its profile with the query. The degree of conformity is the basis for the ranking of the providers. The structure of the profiles and the process of query handling is described in Christoffel (1999).

More important for our considerations is the question which metadata have to be selected to answer the customer queries. As shown in section 1, metadata have to cover different aspects, such as content, cost, and quality of service. There are two ways to gain metadata: from customers or from providers. Metadata can be gained from customers by following the customer-provider communication or by analysing customer feedback. For instance, the average cost for a query to a provider can be estimated out of past queries' costs. Metadata can also be gained from providers, either directly or by analysing test queries. For example, the size of the offered stock (e.g., the total number of offered books in a bookstore) can usually be determined directly, and the availability of supplies can be determined by test queries.

Because of the dynamism of the market, conditions and supply of the providers can change, and providers can vanish without notification, or new providers can appear. A trader in such a market must be able to react on the market dynamism and correct the collected profiles. Moreover, it cannot rely on the friendly behaviour and the honesty of the market participants. So we supposed that the quality of future trading services is determined by their capacities for metadata validation and generation. A more detailed discussion of this subject can be found in Christoffel, Pulkowski et al. (2000).

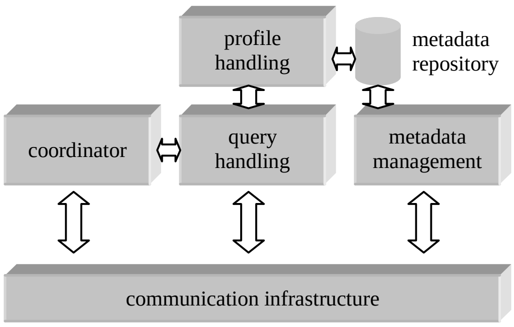

Figure 1 : simplified structure of a trader

The principle components of a trader (see Figure 1) are coordinator, query handling, profile handling, metadata repository, and metadata management. The figure shows that query handling not only founds on the profiles (via profile handling) but also on additional queries to providers and other traders.

3 Cooperating Traders

It is quite obvious that in an adequately big market system more than one trader will be established (unless one does enforce an artificial monopoly). In this section we suppose a community of traders that are independent from each other. These traders can share providers, but often a provider is only mediated by one trader. Also for scalability reasons, a division of the total market may be necessary. Of course, traders compete for the best customers and providers.

Nevertheless, in some cases a cooperation between traders can be useful, most of all when a trader cannot answer a customer query in a sufficient way (if he returns bad results, the customer would fairly ask him a second time). So we need a way that the trader may ask another trader for assistance without having the need for fixed relations. Customers and providers will also benefit from this situation, since there will be more successful mediations.

The solution is to use the same mechanism already used in customer queries. Assisting queries received from another trader will be handled like customer queries (if the trader decides to cooperate), and the results will be sent back to the asking trader, which merges the results together with its own results. Of course, cascades of assisting queries are thinkable; a control path mechanism is needed to prevent loops. Figure 2 shows the cooperation of independent traders.

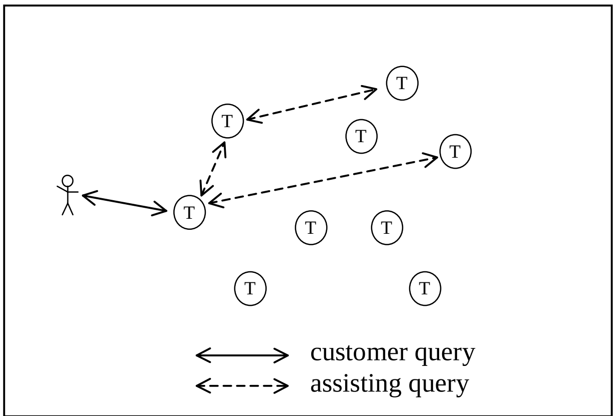

Figure 2: trader cooperation

An important observation is that there are no contracts between the cooperating traders. A trader who uses the assistance of another trader has no guarantee about the quality of the recommendations of this trader. Moreover, it is likely that a trader has to pay immediately for the delivered results. As a consequence, a trader will use cooperation only if the total estimated profit will exceed the additional cost; the use of assisting queries at random does not make sense. The decision whether to include another trader in the process of query handling (and which trader) is based on collected profiles for these traders. This decision is done dynamically in query handling. In order to collect trader profiles, a trader may use corresponding methods like it uses for provider profiles.

A possible consideration (and first experiments confirmed this consideration) is that the missing reliability will be an obstacle for the efficient use of trader cooperations. We expect the development of longer-term relations between traders in the sense of business associates. Here the traders win in reliability but beware their independence, and the dynamism of the cooperations is not touched.

The existence of several independent traders raises a problem for the customer. Since the quality of results differs from one trader to the other, he/she has to decide which trader to consult. A possible solution would be the introduction of a metatrader whose task is the mediation of an appropriate trader. There are two disadvantages of this solution: First, the metatrader acts as a bottleneck in the market structure, since every query has to be led via the metatrader. Second, it is not very probable that all traders in the system will agree in the use of a common metatrader.

4 Trader Federations

The problems mentioned in the last paragraph can be avoided if there are fixed relation between traders. These traders form a trader federation. The structure of a trader federation is not fix. Traders can leave the federation, and new traders can join; relations can be disbanded, and new relation can be set.

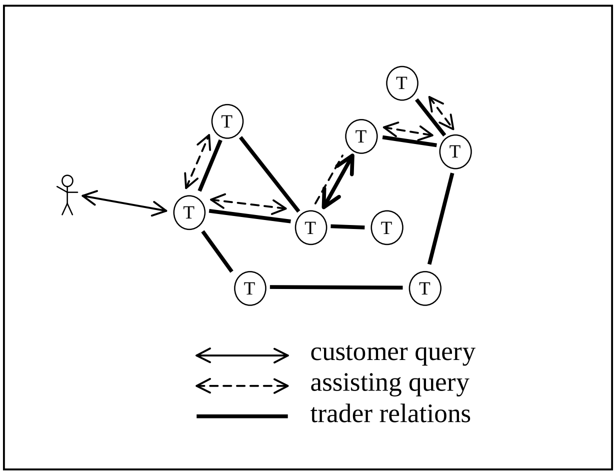

Figure 3: decentralised non-hieratchical trader federation

4.1 Non-hierarchical trader federations

The easiest way to form a trader federation is when the federated traders form a graph where trader relations are modelled as bi-directional edges (see Figure 3). Traders may send assisting queries to related traders, and the related traders may send the queries to further traders related to them. It is clear that it is possible that a query may reach every trader in the federation, although this is rational only in few cases. The difference to the previously discussed model of loose cooperation lies in the fixed relations between the agents, which act in the sense of contracts. So each trader has an amount of traders from those it will receive reliable assistance. Of course, since a federations is treated to be dynamic, the relations do not necessarily have to be set for lifetime.

There is also an effort needed to install and administer a federation. These tasks should be done by the federation itself. There are two different approaches: The decentralised approach foresees no administration instance other than the traders joined in the federation, while the centralised approach contains a special federation leader.

In a decentralised federation (Figure 3), there is no agent other than the traders that takes part in the federation. So each trader should be able to decide itself whether to create, build or join a federation. A new federation will be created when two previously independent traders join together. Other traders can join with one of the federated traders and by doing this join the whole federation. The same way two federations can merge together. The trader basis the decision whether to join with a trader (or a federation) is based on profiles about the other trader. So before two traders join in a federation, there must have been a cooperation and an exchange of profiles. The structure (the trader relations) of the federation can change during the time. Trader may leave the federation when they receive no benefit from the membership. A trader may also be expelled out of the federation, e.g., if it produces poor results or does not follow the rules. A continued idea is the ability of the federation to create and terminate traders dynamically.

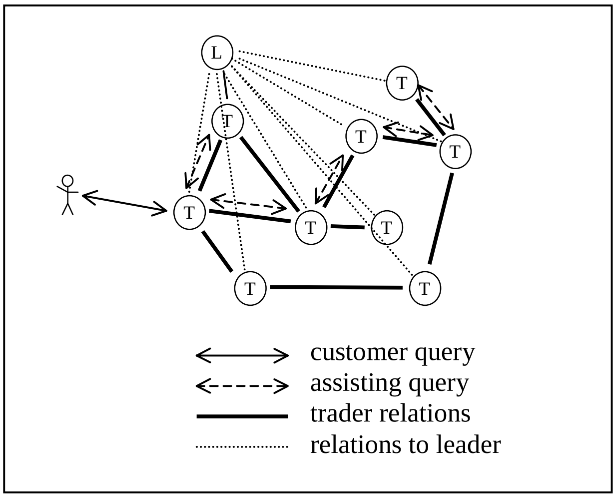

Figure 4: centralised non-hierarchical trader federation

A major drawback of the decentralised approach is that there is no agent in a federation which knows the federation in its entire dimension. The organisation of centralised federation is easier in that sense that there is a federation leader which knows every trader of the federation and the relations between them. The leader can be a trader or a separate type of agent. The question that arises is how the leader will be determined. An easy solution is that the leader will be determined when the federation is created and stays the leader as long as the federation exists. A “more democratic” idea would be the introduction of a kind of elections among the federated trader. However, as long as there are no results that point out the qualities of a “good” federation leader, there are no criteria for an election. In any case must be stated clearly that the centralised approach brings in a new bottleneck situation. Figure 4 shows a centralised federation with a separate federation leader.

Another consideration is the procedure to distribute profits among the federated traders. There are three different approaches:

- Traders keep there own profit (received from customers and/or providers) and pay the other traders for their assistance (same like in a loose cooperation);
- traders keep their own profits but don't pay for assistance;
- all profit will be shared among the federated traders.

The advantage of a federation of traders for customers and providers is that they do not have to decide among a multitude of different traders. Trader federations can appear as a unity that can be addressed by customers directly (e.g., by addressing the federation leader), or customers can access a single trader of the federation, but know that they will be forwarded to another trader, if the trader is not the best available for their special problem. Both approaches allow traders to specialise: Being an expert in one area (e.g., thematic or geographical), and passing on other queries. Same as customer queries, traders can also forward provider registrations and provider profiles

It must be clear that there must be trust and reliability among the traders joined in a federation. The traders must be willing to give away potential customers (in the hope of getting other opportunities from federated traders and profit in the end). If the traders are controlled by independent parties, this level of trust is not self-evident.

4.2 Hierarchical trader federation

One disadvantage of a non-hierarchical trader federation is that it is difficult (not to say impossible) to reach all those traders who can contribute to the desired result and skip the rest. A trader knows those federated traders with which it is related directly, but does not know the rest of the federation. A trader may include the profiles of the traders related to it into its own profile, but there is no obvious algorithm how to define a clear profile this way without integrating the profiles of all traders of the federation. It may be rational to declare parts of the graph for traders with similar specialisation (in this case a common profile makes sense), but the general and liberal structure of the non-hierarchical federation does not enforce this.

So a query may be forwarded to a much wider range of the federation than necessary to answer the query. But it is also possible that the most relevant trader for a given customer query (perhaps this trader is specialised in exactly the subject the customer is looking for) will not receive the query (unless each query is forwarded to all traders in the federation).

This is the reason for the introduction of hierarchical trader federations, where there is a hierarchy of superior and subordinated traders. The traders and trader relations in such a federations build a tree or a directed anti-cyclical graph. Figure 5 shows a decentralised hierarchical trader federation in the form of a tree. Of course, a centralised hierarchical trader federation may also be possible, where a separate agent or one of the traders of the federation (e.g., the root of the tree) acts as federation leader.

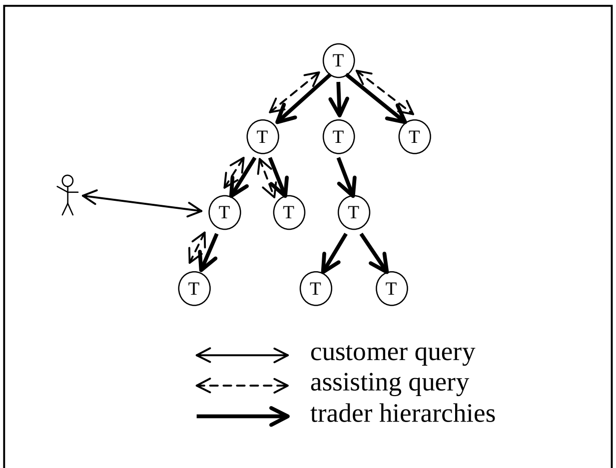

Figure 5: decentralised hierarchical trader federation

An important observation is that the a definition of a common profile for a subtree is well-defined. When a trader receives a query, it regards also the profiles of the subordinated traders, and sends an assisting query if the degree of conformity is high. However, it sends an assisting query to a superior trader only if it cannot handle the customer query even with the assistance of the superior traders. This procedure allows to limit the distribution of a query in the federation, but also allows to reach every trader if it is necessary for the special query.

The potential of trader federations lies in the specialisation of traders. So a trader specialised in biology may have a subordinated trader specialised in zoology. This way customer queries (and also provider registrations) can be forwarded easily to most suitable trader, even if the structure of the entire federation is not known.

There is also a big disadvantage of an hierarchical federation. We already mentioned that it may be problematic to join traders that are controlled by different organisations. Now a hierarchical federation does not only force them to give away customers and provider, it does also force them to accept the hierarchy, e.g., being the subordinate of another trader.

5 Realisation

In order to test our assumptions and ideas in a concrete market scenario, we have continued research in two separate directions: by simulation and by the development of an integration platform for a special information market.

5.1 Simulation Study

In our simulation scenario, we have modelled a simplified, but in its basics realistic fictive information market, containing various customers and providers. In our simulation model, we have taken concern on the qualities of an open market. Providers offer different kinds of product groups, which may be demanded by customers. The benefit of a provider is determined by its profit; the benefit of a customer by several factors, such as cost, time expended, and quality of service. We have also implemented learning behaviour and special influence factors such as migration and immigration of market participants and advertisement.

A special consideration in the simulation study was the influence of the traders. Traders act as service mediators as described in this paper. The benefit of a trader is also determined by its profit; there are different charging models available: fixed registration fees, rates per time, and bonuses. There is also the possibility of cooperations among traders, although not all features described in this paper are already implemented.

Until now, we have performed 16 controlled experiments within this simulation. The experiments verified and explained general market mechanisms such as competition and niches, which are forecasted by market theory (Varian, 1999) and previously observed in other experiments (Kephart et al., 1998). The experiments also covered continuing experiments, investigating the influence of traders.

One of the most fundamental results that is related with the subject of this paper is shown in Figure 6. The figure shows the average utility of the customer over the time in two markets with each 10 customers and 20 providers. The two markets differ only that in the first market there are no traders (the light graph), and in the other market there are 10 traders (the dark graph). As we can see, the average utility of a customer in the market with traders is approximately 19% higher than in the market without traders.

We plan to continue this study as soon as we can confirm the results in a real world scenario.

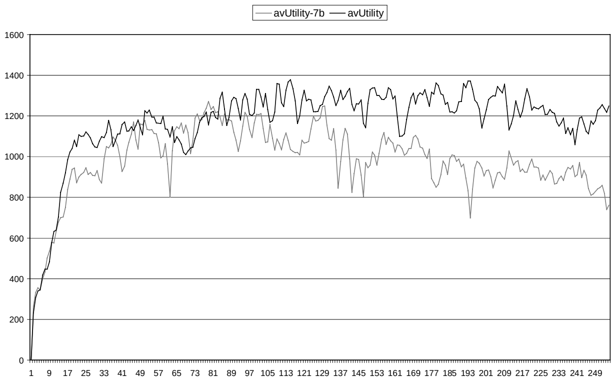

Figure 6: an experiment investigating the benefit of the traders for the customer

5.2 The UniCats project

The second direction is the practical application of our approach in a real market, which is part of the UniCats project (Christoffel et al, 1998, Christoffel et al. 1999). Aim of the UniCats project is to develop an infrastructure for a market of scientific literature. This market used to be in the hand of relatively few libraries and bookstores, but now the success of the Internet opened the way for competitive and new kinds of providers: delivery services, bibliographic databases, publishing houses, and Internet bookstores among others. Even traditional institutions like university libraries have to follow and offer their services online.

Core of the developed system is the UniCats environment, a framework for independent and communicative UniCats agents (Christoffel, Nimis et al., 2000). The environment offers two different kinds of communication: directed communication between two agents, and group agents between a group of agents or all agents. Communication is based on the exchange of XML messages.

Figure 7 shows the layered structure of a UniCats agent. The communication layer includes the functionality for both directed and group communication (on the basis of TCP/IP). It also has a repository of all known groups and agents. The intermediate layer encompasses message encoding and decoding (and also encryption and decryption) and also methods for primitive transactions such as acknowledgements and status reports. The agent layer contains all the communication transactions that are available for an agent type (so the agent layer may differ from one type to the other). There is also a registry of all running jobs and queries. The application layer contains all the algorithms and logic of an agent. Therefore, it may be different for each individual agent. The application layer also contains the interface of the agent to a human user. In addition to that, every agent has a graphical observation interface that is used to observe an agent's behaviour and to set up parameters and direct commands. Figure 8 shows such an interface.

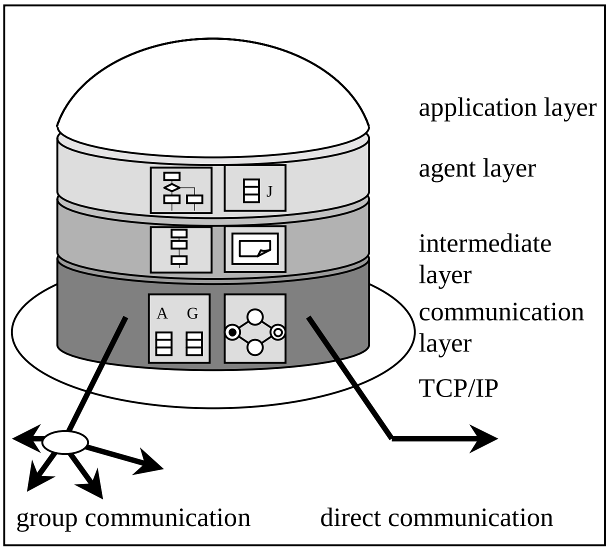

Figure 7: structure of a UniCats agent

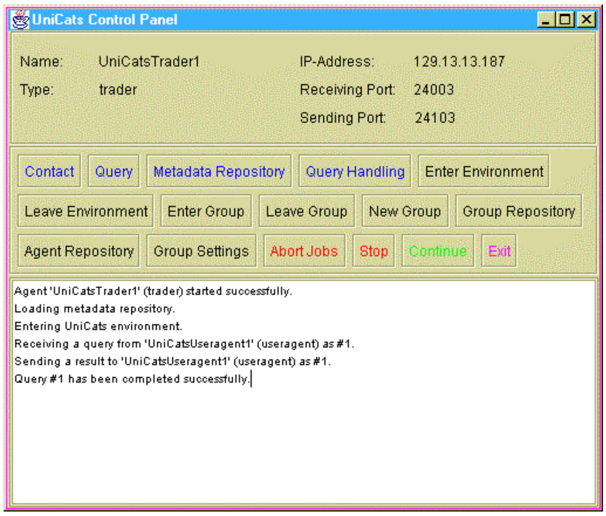

Figure 8: observation interface of a UniCats agent

At the moment, we have implemented prototypes of three agent types (see Figure 9):

- User agents (U) that establish the connection between customers and the environment,
- wrappers (W) that establish the connection between providers and the environment, and
- traders (T) that select appropriate providers for a given query.

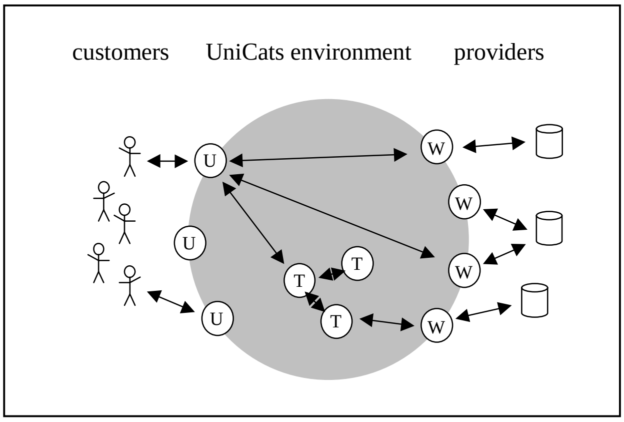

Figure 9: the UniCats environment

In a simple use case, a query is initiated by a customer with the assistance of a user agent. The user agents asks a trader to recommend a provider. Then the user agent sends the query to the provider via the wrapper and presents the results sent back. More complex use cases encompass sequences of queries, ranked recommendations for a series of providers, and additional communications between the participating agents.

Figure 10: the 3D virtual reality interface of the UniCats system

We cannot set out the functionality of the user agent and the wrapper in this paper. Please refer to Schmitt and Schmitt (1999) and Pulkowski (1999) for more information.

We laid stress on low requirements of our system, which is a precondition for a system that satisfies an open market. Agents can be installed easily anywhere in the Internet, providers can be connected without having to install a new interface, and the connection to the customers is eased by user-adapted interfaces. Figure 10 shows one of the user interfaces we are experimenting with.

An extension of the UniCats environment may include new agent types such as certification centres and servers for digital payment.

6 Related Work

In this section we take a look on some other approaches dealing with interworking traders in distributed systems.

The use of cooperations between autonomous instances to enlarge the potential of the single traders has been suggested quite early. In an approach discussed in Canberra and Queensland (Bearman and Raymond, 1991), cooperation does not base on assisting queries but on sharing knowledge. Against the background of the idea of open distributed processing (ODP), they discuss decentralised federations of traders, where each trader holds its own set of federated traders. The connections between traders are defined by contracts. Bearman et al. (1997) gives a clearer description on the realisation of trader federations. Main instrument of the organisation of these federation is a global asymmetric graph of federated traders. An approach for the realisation of dynamism within federation or loose cooperations is not given.

The DRYAD project in Helsinki (Kutvoven, 1995) picked up the idea of federation contracts in their co-operations between trading domains. Here a trader may relay queries to federated traders, which pass the results back to the first trader. Traders store knowledge about all federated traders. The decision which federated traders are to be asked is determined by the global search policy, which takes into account also the dynamism of the federations.

The RHODOS project in Deakin university (Ni and Goscinski, 1994) deals with trading between several distributed systems. They focus on the naming problem which arises in the trading of objects and resources. They provide a model for assisting queries between trader. However, they give no approach for the organisation of trader cooperations and federations.

The university of Nottingsham (Lee and Benford, 1998) suggests non-hierarchical trader federations and introduces an explorative model, where queries can be passed to adjacent traders based on local decisions. They suggest affinity between traders for measure for these decisions. Links between traders are treated as dynamic and can be created and destroyed randomly.

While the previous mentioned approaches deal with object trading only and because of that cannot applied to our problem directly, Stanford suggests a model for federated trading for literature sources (Gravano and García-Molina, 1995). Their main idea is the installation of glossary services (GlOSS) that summarise the content of text databases. Accordingly, they suggest hierarchies of glossary services by summarising the content of several glossaries and selecting appropriate glossary services to forward queries. However, the need for a direct access to the document databases and the missing consideration of dynamism makes this approach not suitable for an open market.

Service mediation is considered in Hamburg (Müller et al., 1996). Cooperations are viewed as a key mechanism to enhance local traders. Connections between traders base on dynamic links. Cooperations can be established directly or with the help of a middleware layer. However, the missing consideration of an open system makes this approach not directly suitable for our problem.

7 Conclusions

In this paper, we have discussed the challenges for a trading service in an open information market. We have pointed out the importance of market survey and metadata management. We paid special attention to the idea of trader cooperations and federations, from which we expect an increased efficiency of the trading service. We have shown that the dynamic behaviour of the relation between the traders is an important issue for the efficiency. We have introduced several models of the structure and organisation of trader federations and cooperations.

Our approach can be enlarged when we consider that our trader community may not evolve into one big trader federation but into a group of competitive trader federations (and some isolated traders that work on their own). In this case, we need the definition and implementation of cooperations between federations.

We realised our approach in the simulation of the behaviour of an information market and in the implementation of an infrastructure for a special information market, namely the market of scientific literature. We plan to extend this application. We also want to perform experiments and measurements on this platform and expect to gain new ideas and results.

References

M. Bearman and K. Raymond. Federating Traders: An ODP Adventure. In: *Proceedings of the Workshop on Open Distributed Processing*, Berlin, Germany, 1991.

M. Bearman: ODP-Trader. In: *Proceedings of the International Conference on Open Distributed Processing*, Berlin, Germany, 1993.

M. Bearman, K. Duddy, K. Raymond, and Andreas Vogel. Trader Down Under: Upside Down and Inside Out. In*: Theory and Practice of Object Systems*, 3(1): 15-29, 1997.

M. Christoffel, S. Pulkowski, B. Schmitt, and P. Lockemann. Electronic Market: The Roadmap for University Libraries and their Members to survive in the Information Jungle. In: *ACM Sigmod Record Special*, 27(4): 68-73, 1998.

M. Christoffel, S. Pulkowski, B. Schmitt, P. Lockemann, and C. Schütte. The UniCats Approach - New Management for Books in the Information Market in: *Proceedings of the International Conference IuK99 - Dynamic Documents*, Jena, Germany, 1999.

M. Christoffel. A Trader for Services in a Scientific Literature Market. In: *Proceedings of the 2nd International Workshop on Engineering Federated Information Systems*, Kühlungsborn, Germany, 1999.

M. Christoffel, S. Pulkowski, and P. Lockemann. Integration and Mediation of Information Sources in an Open Market Economy. In: *Proceedings of the 4th International Conference Business Information Systems*, Poznan, Poland, 2000.

M. Christoffel, J. Nimis, S. Pulkowski, B. Schmitt, and P. Lockemann. An Infrastructure for an Electronic Market of Scientific Literature. In: *Proceedings of the 4th International Baltic Workshop on Databases and Information Systems*, Vilnius, Lithuania, 2000.

L. Gravano and H. García-Molina. Generalizing GlOSS to Vector-Space Databases and Broker Hierarchies. In: *Proceedings of the 21st Very Large Data Bases Conference*, Zürich, Switzerland, 1995.

R. Guttman, A. Mouskas, and P. Maes. Agents as Mediators in Electronic Commerce. In: Matthias

Klusch (Ed.). *Intelligent Information Agents*. Springer, Berlin - Heidelberg, Germany, 1999.

J. Kephart, J. Hanson, J. Sairamesh. Price War Dynamics in a Free-Market Economy of Software Agents. In: *Proceedings of the 6th International Conference of Artificial Life*, Los Angeles, USA, 1998.

L. Kutvonen. Achieving Interoperability through ODP Trading Function. In: P*roceedings of the 2nd International Symposium on Autonomous Decentralized Systems*, Meza, USA, 1995.

O. Lee and S. Benford. An explorative approach to federated trading. In: *Computer Communications*, 21(2), 1998.

S. Müller, K. Müller-Jones, W. Lamersdorf, and T. Tu. Global Trader Cooperation in Open Service Markets. In: *Proceedings of the International Workshop Trends in Distributes Systems*, Aachen, Germany, 1996.

Y. Ni and A. Goscinski. Trader Cooperation to Enable Object Sharing among Users of Homogenous Distributed Systems. In: *Computer Communications,* 17(3), 1994.

S. Pulkowski. Making Information Sources Available for a New Market in an Electronic Commerce Environment. In: *Proceedings of the International Conference on Management of Information and Communication Technology*, Copenhagen, Denmark, 1999.

S. Pulkowski. Intelligent Wrapping of Information Sources: Getting Ready for the Electronic Market. In: *Proceedings of the 10th VALA Conference on Technologies for the Hybrid Library*, Melbourne, Australia, 2000.

F. Rose. *The Economics, Concept, and Design of Information Intermediaries*. Physica, Heidelberg, Germany, 1999.

B. Schmitt and A. Schmidt. METALICA: An Enhanced Meta Search Engine for Literature Catalogs. In: *Proceedings of the 2nd Asian Digital Library Conference,* Taipei, Taiwan, 1999.

C. Shapiro and H. Varian. *Information Rules: a strategic guide to the network economy*. Harvard Business School Press, Boston, USA, 1999.

H. Varian. *Intermediate Microeconomics: A Modern Approach*. W. W. Norton & Company, New York, USA, 1999.

H. Varian. Markets for Information Goods. In: *Proceedings of Monetary Polity in a World of Knowledge-Based Growth*, Quality Change and Uncertain Measurements, 2000.