

# The multi-agent system architecture in SEWASIE

Pablo R. Fillottrani

Faculty of Computer Science

Free University of Bozen/Bolzano

Piazza Domenicani 3, Bozen/Bolzano 39100, Italy

and

Departamento de Ciencias e Ingeniería de la Computación

Universidad Nacional del Sur

Av. Alem 1253, Bahía Blanca 8000, Argentina

## ABSTRACT

We describe the design, implementation and deployment of the multi-level agent-based system architecture developed for the SEWASIE project. The aim of the system is to help the user in querying heterogeneous data sources which are integrated by means of ontologies. The agent architecture is based on a two level data integration scheme supported by mediators and brokers, connected by a peer to peer mechanism. Implementation is done on top of the JADE system, a modular and scalable platform that satisfies FIPA standards.

**Keywords:** multi-agent systems, data integration, ontologies

## 1. INTRODUCTION

The SEWASIE (SEmantic Webs and AgentS in Integrated Economies) project (IST-2001-34825) is an European research project that aims at designing and implementing an advanced search engine enabling intelligent access to heterogeneous data sources on the web, in a rich semantic (ontological) framework. Thus, it provides the basis for precise and effective access to information, and efficient web based communications. The goal of the architecture is to support a flexible set of actors enabling data providers, intermediaries, and data seekers to meet and exchange both information and meta-information, the ultimate goal being the ability to support large scale communities and economies. The project started on May 2002, and took place for three years. SEWASIE is a collaboration between the Università degli Studi di Modena e Reggio Emilia, the Università degli Studi di Roma "La Sapienza", the Rheinisch Westfaelische Technische Hochschule Aachen, and the Free University of Bozen/Bolzano, with CNA SERVIZI Modena s.c.a.r.l, Thinking Networks AG, IBM Italia SPA, and Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung eingetragener Verein as industrial partners.

In this paper we present the design and development of the system architecture, describing the agent technology that supports the main com-

ponents and their functionality. In section 2 we introduce the architecture of the system, and in section 3 we summarize the process of query answering in terms of agent behaviors. In section 4 we introduce some characteristics of current system deployment, and in section 5 we compare SEWASIE with some related systems.

## 2. SEWASIE ARCHITECTURE

We can think to the SEWASIE architecture [1] as composed by a logical topology and a semantic network. The logical topology is realized through the adoption of an agent framework. The semantic network is realized by agents offering services and interacting in semantically rich manners. It is shown in figure 1.

*Brokering agents* (BAs) are the peers responsible for maintaining a view of the knowledge handled by the network. This view is maintained in *ontology mappings*, that are composed by the information on the specific content of the *SEWASIE Information Nodes* (SINodes) which are agents under the direct control of the BA, and also by the information on the content of other BAs. Thus, BAs must provide means to publish the locally held information within the network. There are different roles for BA which depend on the business model of the company which deploys the BA. A company may establish a BA to manage access to its sources which it makes available via SEWASIE. A company specialized on information brokering may establish a BA that combines ontologies provided by several other BA.

*Query agents* (QAs) are the carriers of the user query from the user interface to the SINodes, and have the task of solving a query by interacting with the BAs network. Once a BA is contacted, it informs the QA a) which SINodes under its control contain relevant information for the query, and b) which other BAs may be further contacted. Therefore, the QA translates the query according to the ontology mappings of the BA, and directly ask the SINodes for collection of partial results. Also, it decides whether to continue the search with the other BAs. Once this process is finished, all partial results are in-

tegrated into a final answer for the user.

The life cycle of a QA is initiated by an invocation to its service for solving a query, and it is finalized when delivering the results. Therefore, a single QA is attached to only one query processing. In principle, it seems preferable that QAs reside in a single server node, sending remote messages to BA and SINodes on other server nodes, so mobility is not an issue for QA. Anyway, it might be possible for a query agent to decide that the SEWASIE Server node in which it is residing is overloaded, and therefore prefers to move to other Server node in order to improve query answering performance (load balance). These mobility is supported and/or implemented by the underlying agent platform.

The SINodes are mediator-based systems[2], each including a global view of the overall information managed within. The managed information sources are heterogeneous collections of structured, semi-structured or unstructured data, e.g., relational databases, XML or HTML documents. SINodes are accessed by QAs in order to obtain data, and also by the managing BAs for building the ontology mappings. In order to create and maintain a global view of its information sources, SINodes require an ontology builder. This component performs in a semi-automatic way the enrichment process to create the SINode *ontology*. In turn, this SINode ontology is also integrated with other similar components into the BAs ontology mappings.

The user interface is a group of modules which work together to offer an integrated, easy to the user interaction with the SEWASIE system. It includes a query tool that guides the user in composing queries. In doing so, it requires the ontology of a *starting brokering agent* which helps in the interface presentation and behavior. Each instance of the query tool includes a *Query Tool Agent* (QTA) that is responsible to carry out communications with other agents in the SEWASIE system. In general, QTAs are needed to obtain the initial ontology from a BA, and also to create QAs that will solve the queries generated by the user.

Two extra elements of the user interface are the visualization and the communication tools. The visualization tool is responsible for monitoring information sources according to user interests which are defined in *monitoring profiles*. To perform this task, the visualization tool generates one *Monitoring Agent* (MA) for each topic of interest. Each MA contains a fixed internal ontology (so-called domain model) which is linked to higher level SEWASIE ontologies. Agents of this type regularly set up QAs to query the SEWASIE network, filter the results, and fill *mon-*

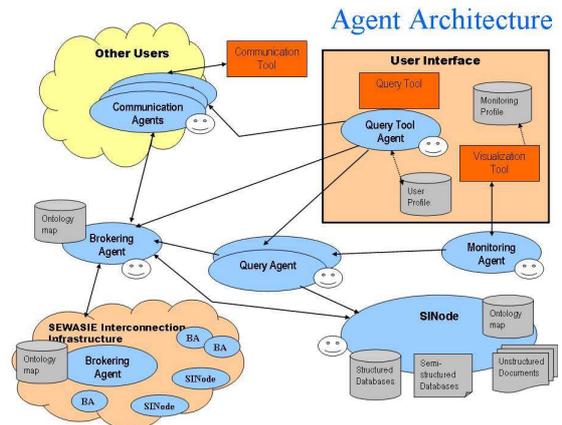


Figure 1: The SEWASIE system agent architecture.

*itoring repositories* with observed documents.

The communication tool supports negotiation between the user and other parties. Any query including contact information sets the context to launch the communication tool. This tool create several types of Communication Agents (CAs) that help in finding and contacting potential business partner, asking for initial offers, and ranking them. The human negotiator can then decide and choose the best offer to begin negotiating, with support from the communication tool. In order to fulfill their tasks, some of these agents periodically create QAs to search for information in the SEWASIE network.

CAs belong to one of the following four types:

- *Initiation agents*: that maintain the pre-negotiation phase for each search.
- *Filtering and Ranking agents*: that evaluate offers and rank them according to given user preferences.
- *Resource Management agents*: that check the capacity of the negotiator and notify him by over-commitment.
- *Negotiation agents*: that conducts the negotiation when a well defined state is achieved, and the user releases the control.

We presented in this section the different types of agents that conform the SEWASIE architecture. In next section we will describe in more detail how they interact in order to solve a query.

### 3. QUERY ANSWERING PROCESS

We will start by describing our assumptions on the query language, followed by the query building process. The whole system is supported by formally defined reasoning services, as described

below. Our aim is to be as less restrictive as possible on the requirements for the ontology language. In this way, the same technology can be adopted for different frameworks, while the user is never exposed to the complexity (and peculiarities) of a particular ontology language. For a more detailed description the reader should refer to [3, 4].

### Query Language

In the SEWASIE context, an ontology is composed by a *set of predicates* (unary, binary), together with a *set of constraints* restricting the set of valid interpretations (i.e., databases) for the predicates. The kind of constraints which can be expressed defines the expressiveness of the ontology language. Note that these assumptions are general enough to take account of widely used modelling formalisms, like UML for example.

The query tool is build around the concept of classes and their properties, so we consider conjunctive queries composed by unary (classes) and binary (attribute and associations) terms. Query expressions are compositional, and their logical structure is not flat but tree shaped; i.e., a node with an arbitrary number of branches connecting to other nodes. This structure corresponds to the natural linguistic concepts of noun phrases with one or more propositional phrases. The latter can contain nested noun phrases themselves.

A query is composed by a list of terms coming from the ontology (classes); e.g., “Supplier” and “Multinational”. Branches are constituted by a property (attributes or associations) with its value restriction, which is a query expression itself; e.g., “selling on Italian market”, where “selling on” is an association, and “Italian market” is an ontology term.

### Query Building

In the query tool, the body of a query can be considered as a graph in which variables (and constants) are nodes, and binary terms are edges. A query is connected (or acyclic) when for the corresponding graph the same property holds. Given the form of query expressions composed by the interface above introduced, we restrict ourselves to acyclic connect queries. This restriction is dictated by the requirement that non expert user must be comfortable with the language itself.<sup>1</sup> Note that the query language restrictions do not affect the ontology language, where the terms are defined by a different (in our case more expressive) language. The complexity of the ontology language is left completely hidden to the

user, who doesn't need to know its details. The process of transforming a query expression in a conjunctive query is described in [7].

Initially the user is presented with a choice of different query scenarios which provide a meaningful starting point for the query construction. The interface guides the user in the construction of a query by means of a diagrammatic interface, which enables the generation of precise and unambiguous query expressions.

The focus paradigm is central to the interface user experience: manipulation of the query is always restricted to a well defined, and visually delimited, subpart of the whole query (the *focus*). The compositional nature of the query language induces a natural navigation mechanism for moving the focus across the query expression (nodes of the corresponding tree). A constant feedback of the focus is provided on the interface by means of the kind of operations which are allowed. The system suggests only the operations which are “compatible” with the current query expression; in the sense that do not cause the query to be unsatisfiable. This is verified against the formal model (the ontology) describing the data sources.

One of the main requirements for the interface is that it must be accessed by any HTML browser, even in presence of restrictive firewalls. This constraints its design. The interface is composed by three functional elements. The first one (top part) shows the tree structure of the query being composed, and the current focus. The second one is the query manipulation pane (bottom part) providing tools to specialize the query. Finally, a query result pane containing a table representing the result structure. The first two components are used to compose the query, while the third one is used to specify the data which should be retrieved from the data sources.

Reasoning services w.r.t. the ontology are used by the system to drive the query interface. In particular, they are used to discover the terms and properties (with their restrictions) which are proposed to the user to manipulate the query. In fact, the terms and the properties proposed by the system depend on the overall query expression, not only on the focus. This means that subparts of the query expression, taken in isolation, would generate different suggestions w.r.t. those in their actual context in the query.

We do not impose general restrictions on the expressiveness of the ontology language; however, we require the availability of two *decidable* reasoning services: satisfiability of a *conjunctive query*, and containment test of two conjunctive queries, both w.r.t. the constraints. Because of space restrictions, the details of reasoning services are not included, and more information can

<sup>1</sup>Our technique can deal with disjunction of conjunctive queries, even with a limited form of negation applied to single terms. See [5, 6] for the technical details.

be found in [8].

### Query Management

Query management in the SEWASIE framework involves different tasks corresponding to the two-level integration scheme between agents. These can be summarized as follows:

**global level** Given a query expressed in terms of an ontology provided by a brokering agent,

1. Single out the SINodes managed by the BA that are relevant for computing the answer to the query, and reformulate in terms of queries to be posed to the SINodes.
2. Individuate alternative BAs (different from the original one used for query formulation) that may have links to SINodes containing relevant information for computing the answer. The query is then reformulated in order to obtain single queries which should be forwarded to the appropriate BAs.
3. Reconstruct the answer to the QTA on the basis of the received partial answers.

Note that the second step above is recursive, in the sense that answering a query posed to other BAs is done through the very same process we are describing. This means that the overall strategy for query management must deal with the problem of how to stop recursion.

**local level** Given a query posed in terms of the virtual global view associated to an SINode, retrieve the answer to the query. This task is carried out by the query manager of the SINode of interest, and is characterized as follows. The first goal of this task is to derive a query plan that is able to correctly access the data sources under the control of that SINode. The second goal of this task is to execute the query, thus computing the corresponding answer.

Let us first consider the case in which there is a single BA in the system (or no other BAs are relevant for a given query). In this scenario the key components involved in the query answering process are the BA Ontology (used to compose the query) and the mappings between this ontology and the Global Virtual Views exported by the known SINodes.

Query processing is based on two query reformulation steps. The first step reformulates the query in terms of the SINodes known by the Brokering Agent, and the second step reformulates each of the SINode query obtained in the first step in terms of the data sources known by the SINodes. The actual query processing phases are the following:

1. **Query expansion:** the query posed in terms of the brokering agent ontology is expanded to take into account the explicit and implicit constraints in the brokering agent ontology.
2. **BA class materialization:** the atoms in the expanded query are materialized by taking into account the mapping from the classes in the global schemas of the SINodes to the classes of the Brokering agent ontology.
3. **Evaluation of the expanded query:** when all the BA ontology classes that are relevant for the query are materialized, the expanded query is submitted to the Query Engine to obtain the answer of the original query.

Contrary of the structure among BA and SINodes, there is no hierarchical relation among different BAs. In fact, every brokering agent acts (and cooperates) at the same level. To model the data integration problem underlying the interaction between BAs, one should come up with a formal framework which is of different nature with respect to the one adopted in the characterization of SINodes. Whereas in a SINode, data integration is based on the existence of the global virtual view, such a notion does not show up when trying to formalize the interconnection between brokering agents. The formal framework adopted within SEWASIE follows the Peer to Peer (P2P) paradigm [9, 10, 11]. Roughly speaking, each BA can be considered a peer, and the interconnection between BAs can be seen as mappings between peers.

## 4. SYSTEM IMPLEMENTATION AND DEPLOYMENT

So far we have presented the functional architecture of the system. We now want to shortly describe its deploying architecture. The SEWASIE system is intended to operate in networked environments where heterogeneity and distribution of information arise. Peers expose their ontologies on the network and software agents act as a glue among the different peers. Peers are recognized as being part of the SEWASIE system as long as they register their ontologies by a brokering agent. From a deployment view point, what is distributed is the multi-agent system. As the scope of the SEWASIE project is to focus on the application of software agents and not in providing a general toolkit for building multi-agent systems, the choice was to use existing tools and practices. The key features we were looking for were:

- a high-level language in order to focus on application programming;

- portability in order to allow for multiple platforms to become part of the SEWASIE system in a transparent way;
- FIPA compliant in order to be aligned with the current standards in the agent technology;
- support and maintenance in order to meet deployment needs.

Currently, there exist plenty of alternative toolkits for development multi-agent systems, see [12] for a comparison. Our choice fell on the Java Agent DEvelopment (JADE) [13] developed by TILab. JADE is currently one of the most evolving toolkits and is an open source projects where both professionals and researchers take part. JADE is written in Java and exploits Java RMI for managing software distribution in the environment.

A JADE multi-agent system (or *platform*) is a logical space that can be distributed over diverse physical hosts. Each host participating to the platform has its own Java Virtual Machine (JVM) running. Each JVM is an agent container, i.e., a runtime environment that allows agents to concurrently execute. In order to boot the platform, a main container has to be created. The main container hosts the services necessary to support agent life cycle, migration and communication. Technically, the main container activates the RMI registry that JADE uses to allow containers and agents to reside on multiple hosts. Containers eventually residing on remote hosts can be added to the platform at runtime. No matter where containers are located, the agent platform is seen as a uniform logical space, where all containers can be reached simply knowing their name. Recently, JADE introduced the support of security for multi-agent systems. Security for agents is seen as an extension of the Java security model and in particular of the JAAS interface. Besides the JADE security extension we have exploited tunneling techniques in order to address security issues related to network configurations. This has been necessary to deploy the system in firewalled environments.

JADE agents are implemented as a set of “behaviors”. A behavior represents a task that an agent can carry out, and the platform provides a set of general behaviors already implemented. In order to make an agent execute the task implemented by a behavior, it is sufficient to add the behavior object to the agent’s behaviors list. Behaviors can be added anytime, when the agent starts, or within other behaviors. An agent can execute several behaviors concurrently. However, scheduling of concurrent behaviors is not preemptive in Jade, but cooperative. This means

that when a behavior is scheduled for execution, its action method is called and runs until it returns. Therefore, it is the programmer who defines when an agent switches from the execution of one behavior to the execution of the next one.

In order to implement QAs and BAs based on JADE behaviors, we must take into account that these behaviors need in general to be dynamically created and deleted. Therefore, we need a “deliberative” behavior that does the control reasoning, deciding when and how the creation of other behaviors is done. These other behaviors will exhibit a typically “reactive” implementation: they are activated when messages are received (the only sensor information SEWASIE agents have about their environment), and do some processing actions. This control architecture for agents is similar to some proposed hybrid architectures for robots [14]. The deliberative behavior aims at adding and removing reactive behaviors in the agent, being activated at agent creation and it never completes.

Our deployment architecture foresees therefore that each host activates a web server. The web server acts as a gate to the network environment. Messages and objects to and from an agent container belonging to the platform are HTTP requests going through the web server. This is made possible because Jade manages remote objects and remote calls using Java RMI. When an RMI server is activated, a registry to keep track of all (possibly remote) objects registered is initiated which listen to incoming requests on a given port number. An RMI client can call this service in order to remotely connect and use objects. The web server can make accessible the RMI server in two ways either through a CGI script or by means of a servlet activated in an application server. While the CGI script requires less infrastructural component, the servlet offers higher performance. This represents a tradeoff.

## 5. RELATED WORK

Several agent-based information retrieval systems are known. In order to compare to similar systems, we now emphasize SEWASIE main characteristics:

- two-level data integration scheme: strongly tied local nodes are integrated into SINodes; BAs provide globally integrated ontologies by means of weaker mappings.
- query management: query building assisted by a query tool, query rewriting in the two levels of data integration following local ontologies using sound and complete algorithms.
- additional tools: negotiation and monitoring

tools integrated in the same agent architecture.

Altogether these points make the SEWASIE system unique among the agent-based information retrieval systems.

Some systems are strong on data integration. CARROT II [15] is an agent-based architecture for distributed information retrieval and document collection management. It consists of an arbitrary number of agents providing search services over local document collections or information sources. They contain metadata describing their local document store which are sent to other agents that act as brokers. Like in SEWASIE, these metadata have an unstructured form, without a central control. But there are anyway several differences with the SEWASIE architecture. First, data integration is done in only one level. In this sense, CARROT II agents play the role of a brokering agent and an SINode at the same time. Second, there is no support for the user in creating the query. Metadata information is not reflected in the process of query building. Finally, the most important difference is that agents in this system only produce a routing of the query to relevant information sources, no query rewriting is done in this step. In SEWASIE the query is reformulated following brokering agent's ontology before asking SINodes, which contain the information sources. Several other information retrieval systems are known with routing agents, like HARVEST [16], CORI [17] and InfoSluth [18].

Other systems, like TSIMMIS [19], include some rewriting rules against predefined query patterns. There are several steps of query processing also in the MISSION project [20]. In these cases, data integration technology is not present, or in TSIMMIS limited to automatic generation of wrappers [21] and mediators [22] from web pages. In SEWASIE, the data integration techniques [23] adopted by SINodes apply not only to unstructured, or semi-structured data sources, but also to relational databases.

## 6. CONCLUSIONS

This paper presented the work done, within the perspective of the agent technology, in the SEWASIE project. We showed how the agent-based architecture has been tied up with an ontology based approach to provide the users a transparent access to heterogeneous data sources.

The SEWASIE system provides an ease of access to the data without requiring an in depth knowledge of the sources. This is achieved by leveraging both the agent based collaboration between the different components of the system,

and the heavy use of ontologies as a means of abstraction.

With respect to the query tool, SEWASIE has presented the first well-founded intelligent user interface for query formulation support in the context of ontology-based query processing. Our work has been done in a rigorous way both at the level of interface design and at the level of ontology-based support with latest generation logic-based ontology languages such as description logics, DAML+OIL and OWL. However, there are open problems and refinements which have still to be considered in our future work.

Another important aspect to be worked out is the understanding of the effective methodologies for query formulation in the framework of this tool, a task that needs a strong cooperation of the users in its validation. This is going in parallel with the interface user evaluation.

The other crucial aspect is the efficiency and the scalability of the ontology reasoning for queries. We are currently experimenting the tool with various ontologies in order to identify possible bottlenecks.

## ACKNOWLEDGEMENTS

This work is supported in part by the 5th Framework IST programme of the European Community through project SEWASIE within the Semantic Web Action Line. The SEWASIE consortium comprises the Universities of Modena-Reggio Emilia (Sonia Bergamaschi, the coordinator of the project), Aachen RWTH (M. Jarke), Roma La Sapienza (M. Lenzerini, T. Catarci), Bolzano (E. Franconi), as well as IBM Italia (G. Vetere), Thinking Networks AG (C. Engels) and CNA (A. Tavernari) as user organization.

## REFERENCES

- [1] Bergamaschi, S., Fillottrani, P., Gelati, G.: The SEWASIE multi-agent system. In: Proc. of the 3rd. Intl. Workshop on Agents and Peer-to-Peer Computing (AP2PC 2004). (2004)
- [2] Wiederhold, G.: Mediators in the architecture of future information systems. *IEEE Computer* **25** (1992) 38–49
- [3] Lenzerini, M., Majkić, Z., Beneventano, D., Mandreoli, F.: Techniques for query reformulation, query merging, and information reconciliation part a. Technical report, SEWASIE consortium (2003)
- [4] Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Query reformulation over ontology-based peers. In: Proc. of the 12th Italian Conf. on Database Systems (SEBD 2004). (2004)

- [5] Calvanese, D., De Giacomo, G., Lenzerini, M.: On the decidability of query containment under constraints. In: Proc. of the 17th ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS'98). (1998) 149–158
- [6] Horrocks, I., Sattler, U., Tessaris, S., Tobies, S.: How to decide query containment under constraints using a description logic. In: Logic for Programming and Automated Reasoning (LPAR 2000). Volume 1955 of Lecture Notes in Computer Science., Springer (2000) 326–343
- [7] Dongilli, P., Fillottrani, P.R., Franconi, E., Tessaris, S.: A multi-agent system for querying heterogeneous data sources with ontologies. In: Proc. of the 13th Italian Sym. on Advanced Database Systems (SEBD 2005). (2005) 75–86
- [8] Dongilli, P., Franconi, E., Tessaris, S.: Semantics driven support for query formulation. In: Proceedings of the 2004 International Workshop on Description Logics (DL-04). Volume 104 of CEUR Workshop Proceedings. (2004)
- [9] Bernstein, P., Giunchiglia, F., Kementsietsidis, A., Mylopoulos, J., Serafini, L., Zahrayeu, I.: Data management for peer-to-peer computing: A vision. Workshop on the Web and Databases, WebDB (2002)
- [10] Halevy, A., Ives, Z., Suci, D., Tatarinov, I.: Schema mediation in peer data management systems. In: Proceedings of the 19th International Conference on Data Engineering (ICDE'03). (2003)
- [11] Franconi, E., Kuper, G., Lopatenko, A., Serafini, L.: A robust logical and computational characterisation of peer-to-peer database systems. In: International VLDB Workshop On Databases, Information Systems and Peer-to-Peer Computing (DBISP2P'03). (2003)
- [12] Fillottrani, P.: SEWASIE architecture agentization, SEWASIE deliverable 1.6. SEWASIE Deliverable D1.6 (2003)
- [13] Bellifemine, F., Caire, G., Poggi, A., Rimassa, G.: JADE a white paper. [jade.csel.tu.it](http://jade.csel.tu.it) (2003)
- [14] Murphy, R.: An Introduction to AI Robotics (Intelligent Robotics and Autonomous Agents). MIT Press (2000)
- [15] Klusch, M., Ossowski, S., Shehory, O., eds.: Integrating Distributed Information Sources with CARROT II. In Klusch, M., Ossowski, S., Shehory, O., eds.: Cooperative Information Agents VI, 6th International Workshop, CIA 2002, Madrid, Spain, September 18–20, 2002, Proceedings. Volume 2446 of Lecture Notes in Computer Science., Springer (2002)
- [16] Bowman, C.M., Danzig, P., Hardy, D.R., Manber, U., Schwartz, M.F.: The Harvest information discovery and access system. Computer Networks and ISDN Systems **28** (1995) 119–125
- [17] Callan, J.P., Lu, Z., Croft, W.B.: Searching distributed collections with inference networks. In Fox, E.A., Ingwersen, P., Fidel, R., eds.: SIGIR'95, Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. Seattle, Washington, USA, July 9–13, 1995 (Special Issue of the SIGIR Forum), ACM Press (1995) 21–28
- [18] Woelk, D., Tomlinson, C.: Infosleuth: Networked exploitation of information using semantic agents. In: COMPCON Conference. (1995)
- [19] Garcia-Molina, H., Papakonstantinou, Y., Quass, D., Rajaraman, A., Sagiv, Y., Ullman, J.D., Vassalos, V., Widom, J.: The TSIMMIS approach to mediation: Data models and languages. Journal of Intelligent Information Systems **8** (1997) 117–132
- [20] McClean, S.I., Karali, I., Scotney, B.W., Greer, K., Kapos, G.D., Hong, J., Bell, D.A., Hatzopoulos, M.: Agents for querying distributed statistical databases over the internet. International Journal on Artificial Intelligence Tools **11** (2002) 63–94
- [21] Hammer, J., McHugh, J., Garcia-Molina, H.: Semistructured data: The tsimmis experience. In: Proceedings of the First East-European Symposium on Advances in Databases and Information Systems (ADBIS'97), St.-Petersburg, September 2–5, 1997. Volume 1: Regular Papers, Nevsky Dialect (1997) 1–8
- [22] Papakonstantinou, Y., Garcia-Molina, H., Widom, J.: Object exchange across heterogeneous information sources. In Yu, P.S., Chen, A.L.P., eds.: Proceedings of the Eleventh International Conference on Data Engineering, March 6–10, 1995, Taipei, Taiwan, IEEE Computer Society (1995) 251–260
- [23] Bergamaschi, S., Castano, S., Beneventano, D., Vincini, M.: Retrieving and integrating data from multiple sources: the MOMIS approach. Data and Knowledge Engineering **36** (2001) 215–249