extended abstract

THE WINK PROJECT FOR VIRTUAL ENTERPRISE NETWORKING AND INTEGRATION*

D. BENEVENTANO[1] AND S. BERGAMASCHI[1,3] AND D. GAZZOTTI[2] AND G. GELATI[1] AND F. GUERRA[1] AND M. VINCINI[1]

1 - Dipartimento di Ingegneria dell'Informazione - Università di Modena e Reggio Emilia
VIA VIGNOLESE 905 – MODENA

2- Gruppo Formula S.p.A.
VIA MATTEOTTI, 5
40055 VILLANOVA DI CASTENASO (BO)

{beneventano.domenico,bergamaschi.sonia, gelati.gionata, guerra.francesco, vincini.maurizio}@unimo.it, Davide.Gazzotti@formula.it

Abstract

To stay competitive (or sometimes simply to stay) on the market companies and manufacturers more and more often have to join their forces to survive and possibly flourish. Among other solutions, the last decade has experienced the growth and spreading of an original business model called Virtual Enterprise. To manage a Virtual Enterprise modern information systems have to tackle technological issues as networking, integration and cooperation. The WINK project, born form the partnership between University of Modena and Reggio Emilia and Gruppo Formula, addresses these problems. The ultimate goal is to design, implement and finally test on a pilot case (provided by Alenia), the WINK system, as combination of two existing and promising software systems (the WHALES and MIKS systems), to provide the Virtual Enterprise requirement for data integration and cooperation amd management planning.

1. INTRODUCTION

In modern economies companies and manufacturers have to struggle against market complexity. Complexity can from time to time refer to large-scale markets involving lots of competitors or to market characterised from high (and expensive) technological hurdles. To stay competitive (or sometimes simply to stay) on the market companies and manufacturers more and more often have to join their forces to survive and possibly flourish. It is the case, for instance, of large-scale engineering projects, where companies with adequate know-how, references and financial resources to sustain the project, are the prime contractors and outsource specific components and services to other firms through several forms of sub-contracting. Among other solutions, the last decade has experienced the growth and spreading of an original business model called Virtual Enterprise (VE). A VE can be generally seen as an alliance of enterprises that team up for a limited period of time to take advantage or share competencies, resources and skills in order to better respond or eventually gain business opportunities, and whose *integration* and *cooperation* are supported by computer *networks*. This general definition highlights three relevant aspects of a VE, namely integration, cooperation and networks. Integration refers to the original heterogeneity in organisations, structures, know-how and application domains that each enterprise teaming up brings within the new partnership. Cooperation identifies the fundamental spirit and attitude of the new enterprise constituted by the partners. Networks are meant to unify and connect the set of enterprises now recognised as a unique VE, enabling wide communication and information interchange. As a VE is based upon a technological connection among the partners, all these aspects arise technological issues. Integration addresses information accessing, interchanging and processing according to a common model within the VE. This entails the information system of a VE should be capable of solving conflicts and discrepancies between heterogeneous and distributed data. Cooperation implies a common understanding of business planning and management procedures to be respected through the VE activities and projects. This involves the adoption of common procedures and mechanism to deal with cooperative situations. Networks naturally address the way the partners are organised (hierarchically, peer-to-peer and so on) and which role they play within the VE. It follows that information systems have to be considered the very heart of a VE. There is therefore a need for finding suitable solutions to the challenges in terms of integration, cooperation and networks proposed in the field of information systems for VEs.

* WINK is EU Project, IST-2000-28221.

[3] CSITE-CNR viale Risorgimento 2, 40136 Bologna, Italy

Among other research activities, the database group at the University of Modena and Reggio Emilia together with Gruppo Formula S.p.a (Bologna) have joined their efforts and competencies, based on previous projects, to design, develop and test a system intended to answer the VE technological needs as mentioned above. The rationale behind this partnership lies in the experiences matured from the database group in the field of information integration and software agents and the know-how hold by Gruppo Formula in the field of network enterprises. In 2001, this partnership called forth the WINK system proposal, which is a project funded by the European Commission, EU Project IST-2000-28221.
The WINK system is intended to provide the required services to answer the VEs needs for data collection and management planning by integrating two existing advanced software systems matured in the field of information integration and management, namely the *WHALES* (*Web-linking Heterogeneous Applications for Large-scale Engineering and Services*) and the *MIKS* (*Mediator agent for Integration of Knowledge Sources*) systems ([Gelati01, Bergamaschi02]. WHALES project (IST-99-12538) run at Gruppo Formula (www.gformula.com, Bologna, Italy) and the MIKS project is under development at the University of Modena e Reggio Emilia (www.dbgroup.unimo.it, Modena, Italy). The WHALES system is a web-based planning and management infrastructure for complex distributed organisations working on large scale engineering projects involving non-repetitive production. The MIKS provides a framework for managing software agents performing information extraction and harmonization/integration from structured and semi-structured data sources. As the two systems tender matching functionalities, they are of a kind to face the challenge set out by advanced industrial scenarios.
The ultimate goal is to design and implement a new system that fulfil the requirements in terms of both information collection, access and integration and management planning. Further, the new WINK system will then be tested in an engineering pilot case gently provided by Alenia Space (www.aleniaspace.com, Turin, Italy), partner of the space consortium building the International Space Station. This will take the combined software module one step from becoming a market product suitable to extensive commercial use.

The outline of this paper is the following: Section 2 explains the motivations leading to the WINK project. Section 3 introduces to the WINK platform together with the two systems it combines (the MIKS and the WHALES systems), Section 4 presents the related works, finally Section 5 provides some conclusions.

2. PROJECT MOTIVATION

These issues related to information integration and networking in VEs are far from being exhaustively tackled and generally solutions far from being effectively plunged into real world scenarios.
Data handling entails operating/coping with diverse network models, diverse network processes, diverse data models and structures and finally concurrent data updates and changes. Management planning entails flexibly arrange resource availability, resource coordination, partnership configurations and carrying out monitoring phases and consequent re-planning.
The WINK project will in fact support large-scale multi-supplier, and thus multi-site, projects by implementing integration network process model. Enterprises will therefore spend less human and IT resources coordinating and maximising possible project partnership configurations. They will also save resources during the monitoring phases required for the complex projects they are carrying out. The WINK tool will result in improved planning and budgeting, leading to better forecasting which will reduce re-planning overheads during the course of the projects themselves. It will improve monitoring, making use of the customizable alert systems, while greatly reducing resources tied to cost and risk assessment. WINK will grant effective contingency management thanks to improved management of contingent factors applying pro-active analysis of deterministic and casual risk factors, re-planning to face deviations, impact verification and solution comparisons, re-alignment of plans and budgets for the involved project units, change-tracking mechanisms for revision evaluations. WINK will allow higher flexibility and efficiency by allowing a quicker reaction to customer requests and to contingencies, taking into consideration the appropriate resources and skills at all project sites. Furthermore, WINK will ensure better exploitation of the network potential, by describing its resources, competencies and operating processes, and speeding up and automating information processes during the entire project lifecycle, spanning companies and organisational unit boundaries. The integrated WINK environment will facilitate the updating of network databases and it will lower costs tied to system integration or consortium reassessment (new partners joining the project network).

3. THE WINK PLATFORM

The overall technical objective of the WINK project is to validate a combination of two promising software platforms, which, when joint will perform automated, efficient data collection and management planning in complex “Virtual Enterprises”. WINK will support these “Virtual enterprises” by integrating, adapting and implementing the very promising technologies developed during the distinct WHALES and MIKS projects. We will shortly present these two systems in the following two sub-sections.

3.1 WHALES: a system for management planning

The WHALES system aims at providing a planning and management infrastructure for distributed organisations. These are alleged as extended enterprises, i.e. enterprises working on large scale projects which are characterised by large materials and human resources investments and by concurrent independent activities such as manufacturing, design and services. Figure 1 depicts a feasible common scenario:

Figure 1. A feasible scenario of extended enterprise Figure 1. A feasible scenario of extended enterprise Figure 1. A

Figure 1. A feasible scenario of extended enterprise

Carrying out complex projects involving distributed resources and partners, extended smart enterprises typically have to deal with high risk due to contingent factors and uncertainty concerning budgets and profit margins. The WHALES system therefore endows with services expressively meant for more reliable project plans and budgets, monitoring and costs/risks assessment, improved management of contingent factors and higher flexibility and efficiency.
The WHALES project focus on leveraging the virtual enterprise potential in order to support visibility and collaboration inter-company processes, which take place among the cooperating companies. Thus WHALES provides more reliable project plans and budgets to these organizations; in addition to that, it also provides a detailed model of project activities and requirements (materials, resources) and an updated and consistent view on available resources, by aggregating and harmonising data from heterogeneous applications and different domains and by identifying and comparing alternative solutions. All this helps speed up information flows across companies or units boundaries. The mentioned functionalities are available through a common Web-based working environment.
The design philosophy behind the WHALES system is a typical three-tier architecture where we find the client side, a middle tier, which comprises the business logic, and finally the database tier where data are physically stored and maintained. Figure 2 shows the tiers.

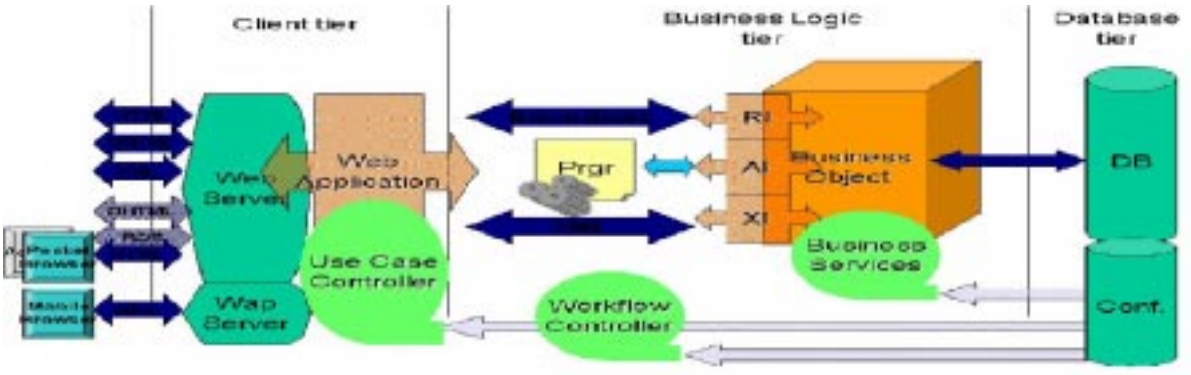

Figure 2. The three tiers of the WHALES architecture

The intelligence core of the entire system resides in the business logic tier. Here business objects retrieved from (and posted to) the data base tier are processed taking into consideration the specific business services required. A few interfaces are available for handling data as recordsets and XML documents and exchanging

data with applications. The client side runs web applications and servers, while the database tier holds some DBMSs.
As for the methodologies, WHALES works thanks to a real-time notification of events and "alert" conditions and an impact evaluation for deviations, i.e. delays and changes on downstream activities. The system performs also a pro-active analysis of risk factors, both deterministic and casual and a re-planning activity to face deviations, with impact verification and solutions comparison and a re-alignment activity of plans and budgets. Different levels of visibility and update rights can be assigned to different units together with the roles they play in the project. Finally, it is important to highlight the integration of WHALES with local applications running at every ERP node (accounting, job order management, purchase) or with applications such as with planning and analysis tools.

3.2 MIKS: a framework for accessing and integrating heterogeneous information

The **MIKS** system [Bergamaschi01, Bergamaschi02, Gelati01] aims to integrate data from structured and semistructured data sources. **MIKS** is an infrastructure for semiautomatic integration of heterogeneous sources schema (relational, object, XML and semistructured sources); it carries out integration following a semantic approach, which uses Description logics-based techniques, clustering techniques and an ODM-ODMG extended model to represent extracted and integrated information, $ODMi_{I3}$. Using the ODL_{I3} language, referred to the $ODMi_{I3}$ model, it is possible to describe the sources (local schema) and the system supports the designer in the generation of an integrated view of all the sources (Global Virtual View), by building a "common thesaurus" that expresses inter and intra-schema knowledge in the form of terminological knowledge. This shared ontology provides a reference on which to base either the identification of classes' candidate to integration and the subsequent derivation of their global representation. Moreover, the MIKS system performs revision and validation of the various kinds of knowledge used for the integration, using Description Logics, affinity-based clustering techniques and lexical knowledge derived from WordNet lexical system (www.cogsci.princeton.edu/wn).
The common thesaurus and the global virtual view are expressed using XML standard. The use of XML in the definition of the Global Virtual View lets to use MIKS infrastructure with other open integration information systems by the interchange of XML data files. In addition, the Common Thesaurus translated into XML files may provide a shared ontology that can be used by different semantic ontology languages. Figure 3 shows the architecture of the system. In particular, it is possible define three levels in which the system may be divided: the user level, that provides the interaction between the user (or the user application) with the system, the mediator level, that provides the Global Virtual View and executes the query phase, and the data level, that manages all the interactions with the sources.
The MIKS system extends the MOMIS system [Benetti02, Bergamaschi99] introducing intelligent and mobile agents to improve the system capabilities. Basically, the exploitation of agents could help improving some of the existing MOMIS features and show up possible new functionalities we could add to the system. One feasible improvement is to introduce intelligent and mobile agents in order to carry out the integration process in a more efficient and flexible way. The exploitation of agents could be particularly useful for broadening the volume of available data and for query processing.

3.2.1 AGENT-BASED MODELS

As the Internet popularity and use spread all round the world, distributed applications become key factors into the development of any information system. A number of solutions to distributed computing have grown and mature in the last decade. One of the most promising approaches for developing such applications is the multi-agent system paradigm. Multi-agent systems have recently moved from pure research studies to more practical and sometimes commercial applications. Examples are the fields of web databases, electronic commerce systems and information gathering systems.

According to the software agent paradigm, an application can be thought as composed by a dynamic group of mobile computational entities, each characterised with specific capabilities and trying to coordinate with other entities in order to accomplish some given tasks. Compared to the object-oriented paradigm, the agent approach extends communication as intended in RPC distributed communication mechanism. According to the latter, a message is just a procedure call whereas in the agent world a message is an object with state and functionality [Samaras99]. This approach is very powerful when we have to design applications intended to run in distributed and heterogeneous environments. The purpose is to leverage the flexibility of the services our applications provide. Within the WINK project, the assistance of a multi-agent system could enhance diverse features such as distributed query execution and interoperability between systems. Distributed query execution could exploit autonomous intelligent agents that move to local sources and execute local queries. Rather than submitting to each source only the local query concerning its data, *query agents* can coordinate themselves in order to carry out the overall querying process, sharing the partial query results.

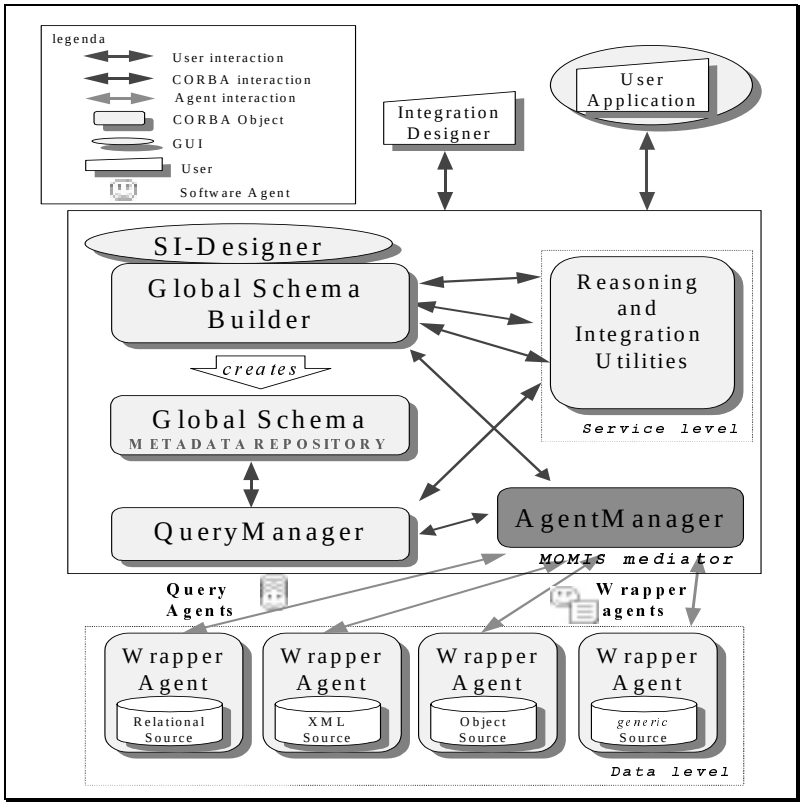

Figure 3. The Miks Architecture

Computation can be thus distributed among all hosts composing the WINK system network. Further, communication costs could possibly decrease due to the lower amount of data to be transmitted. The key issue related to the deployment of multi-agent systems for distributed query is coordination. Coordination among query agents means creating an execution plan to answer the global query and forming a coalition of agents with the precise goal of carrying out a distributed computation. This entails agents should share both the execution plan and the partial answers by means of a proper communication mechanism supported by agent mobility in order to migrate code and data among diverse machines. Interoperability is referred to communication and interaction between entities running on some hosts of the network. Entities can be then systems, agents, applications and data sources. As we are dealing with integration systems the approach to interoperability is based not only on structural terms but also on semantics. This is particularly difficult to achieve seen the heterogeneity of entities (and subsequently data structures) connected through the network. The help of intelligent and mobile agents is critical at this stage. Where new services can be directly programmed as agent capabilities, which are from the beginning explicitly created to support interaction and cooperation, legacy and already existing applications can be shielded by agents from the new network they become parts of, thus opening the function of their services. This process is technically known as wrapping. We want to define a class of agents, namely *wrapper agents*, charged to wrap the information entities of our system using semantic techniques [Bergamaschi01, Bergamaschi02]. For instance, semantic wrapping can be applied to data sources to translate the data they store into a universal format according to some domain ontology, or to legacy applications to make their functionalities openly accessible.

To realise a multi-agent system with the above features a number of *application frameworks* could be taken into consideration. Our framework should certainly offer a set of basic services such as data base query execution and handling, source and application wrapping and information retrieval. Implementing such a framework raises some operational issues. Particularly significant to us are security and performance analysis. Security impacts both the way agents have access information on local hosts and the way the retrieved data are carried by agents or sent through the network. No general security policy can be outlined as it adjusts to the special business case considered according to the cooperation ruled enforced within the VE. Performance analysis has a twofold interest: besides showing the behaviour of the multi-agent system, it is helpful to make some comparison with more traditional way of executing the same tasks. Agent technology has not widely spread yet and understanding its benefits and drawback helps enriching the picture of this fascinating and promising approach. As far as distributed systems are concerned, we experience a lack of methods and tools to fairly evaluate performance due to both the density of a distributed environment and the dynamic nature of multi-agent systems and applications. In literature some studies have been conducted though. In [Papastavrou99] authors have already shown how we can benefit from software agents when accessing remote sources on the Internet. In [Samaras99] authors propose a method for benchmarking the performance of agent platforms such as IBM's Aglet and Mitsubishi's

Concordia. Thus, one of the goals of the WINK project is to define metrics to highlight specific advantages and possible drawbacks when using multi-agent systems in distributed applications.

3.3 THE WINK SYSTEM

The main goal in designing the WINK system is to mould the MIKS and WHALES architectures and functionalities in order to get hold of a new system with the desired features. The design process should go over some major steps: requirements specification, merging of existing functionalities and introduction of essential extensions.

3.3.1 Reference specifications

First, there is a need for collecting reference requirements about the software integration process. This entails analysing thoroughly how the two systems work, understanding at which level they operate and where they can be interfaced to each other. Results should suggest the guidelines and the constraints under which the WINK architecture together with the services to be provided can be specified.

The insight has put forward some of the issues to solved in the software integration process. The interoperability appears to be the most relevant obstacle to the integration of the two applications since the two systems refer to two different technologies (Visual Basic and DCOM vs. Java and CORBA) as far as their design and implementation are concerned. A suitable solution to interoperability should take into consideration the following points:

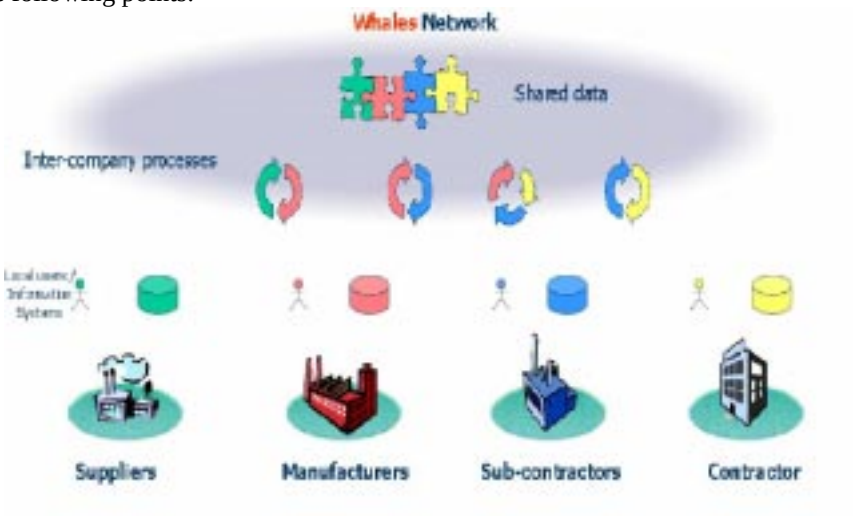

Figure 4. The Whales Network model

- WHALES Network Model system (see Figure 4) should propose a standard language (or formalism) and support for data interchange by identifying the relevant information to be integrated by the MIKS system;

- MIKS system should instead make provision of services relevant for Virtual Enterprise applications/scenarios and the components to be used to collect network data for WHALES.

Finally, high-level requirements to be specified for the WINK system concern cost/performance evaluation, methods for accessing the shared data at local sources and infrastructure and protocols to maximise security of source servers.

3.3.2 Functionality integration

Right from the start, some of the features provided by WHALES and MIKS can be inherited by the WINK system. First, to establish a VE we usually need to arrange some facility acting as communication channel among the partners. This facility should further consider that “Virtual Enterprises” are not generally based on peer relationships, forming a more or less complex business hierarchy. The web interface of WHALES could be used as centralising facility to allow data sharing and exchanging. Thus, the WHALES web interface plays the role of web portal. The business logic behind the interface allows project managers to set up a flexible collaboration system based on visibility right modelling. The available data, either stored or virtually collected by the WHALES system, are partitioned according to any desired access policy. During the WINK project we should identify which services of the portal better suit the Alenia business case.

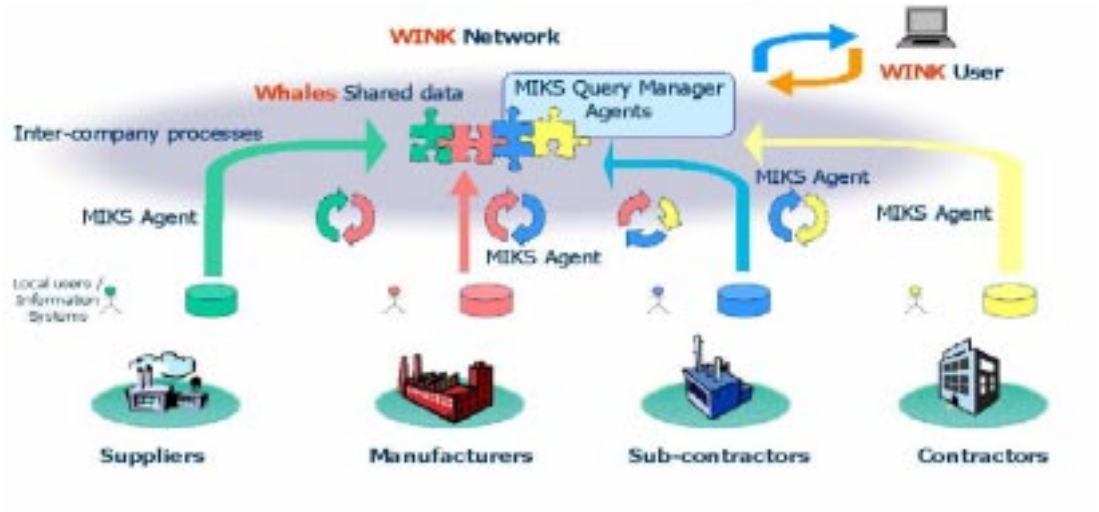

Figure 5. The WINK platform architecture

Secondly, sharing information through a centralised web portal entails dealing with heterogeneous data coming from different partners. These data are unlikely to be stored in the original format, as they have to be proposed in a common format to all partners. At this stage the features of the MIKS system come into play.
The MIKS framework could be exploited inside the WINK system in order to help the designer integrate ALENIA Business data within the WHALES shared data model. For doing this, we can use the MIKS system to define the mapping rules between revised WHALES shared data and ALENIA local data and finally export the Integrated Schema (Global Virtual View) in XML format. Finally, a flexible information management service should always allow dynamic updates to its contents and thus users should always be allowed querying it. While there is a centralised point of contact (the web portal) there are diverse local sources storing relevant data for the project (such as technical specifications or project documents). On one hand, these data are often protected and only some of their content can be made accessible. On the other, they contain essential pieces of information for the project. The query manager of the MIKS system allows global query handling, decomposing the user query into queries that can be executed at each local source. Adding this feature, the WINK system will be able to support remote query management.

3.3.3 Extensions

Aiming at actively combining them, the software integration process also impacts on the two systems in terms of extensions related to the particular domain of usage the WINK platform has been studied for.
A feasible extension of the WHALES system functionalities to be used inside the WINK system should enlarge (or modify if necessary) the WHALES shared data model to manage the ALENIA business data and define web-based user interfaces to interact with MIKS Query Manager agents. Further, we could investigate the SOAP protocol to realize interoperability between WHALES and MIKS.
Software agents can extend the MIKS system functionalities to be used inside the WINK system. In particular, three types of agents may be helpful: search agents, integration agents and query manager agents. Search agents are charged to pump the ALENIA local data to the WHALES shared data, using XML-Schema standard for data interchange. Integration agents are charged to translate XML data using the XSLT standard. Query Manager agents allow submitting non-predefined queries over WINK data, to query on-demand the ALENIA local sources, to give a unified answer by XML-Schema data format and to use search agents to reach the local sources. Web-based user interfaces should be added in order to help users interact with the Query Manager Agent module.

3.3.4 The Wink Architecture

Figure 5 shows the WINK architecture. In general, the WINK network comprises of all the main actors participating actively or passively to the project. Access points, i.e. users, can always be added to the network at any time, the only constraint being the institutional rules enforced within the Virtual Enterprise. Only authorised users can access the web portal. The centralised web portal makes provision of all services necessary to access shared data and information concerning the project management planning. Different users have in general different types of visibility over the shared data. The shared-data warehouse has to be built by integrating data coming from heterogeneous and dispersed sources. The underlying technology to this process is agent technology. Agents operate according to the MIKS framework to build a virtual view that synthesises the original data.
Security is an ever-open issue of information technology. Within the WINK project it impacts on several aspects. The shared-data model is one of them. The shared-data have to be accessible to different categories of users, each with its own view and restriction over the status of the project. The WHALES system puts

forward a method to face this kind of data protection. It is based upon the definition of the overall organisation of the Virtual Enterprise (relationships and hierarchy), the morphology of the established network and the desired security level each partner is subject to. This allows outline some precise visibility rights that each single user holds on the shared-data. As we have already seen agent technology embodies another security concern.

4. RELATED WORK

The WINK project addresses manifold research and practical themes.
As for the sphere of information systems focused on enterprise management and networking several types of applications have been developed especially at commercial level (as the WHALES system is). Rather than literature we are here concerned with products and their application domains. A possible classification groups applications into the following classes:

- **Project planning**: these applications support the step of decision making, identifying the business structure and the capacity profile of the enterprise. As for the management, they help sort the activities that have to be carried out;
- **Human resources**: these applications support knowledge management and working force budgeting. At execution level they propose time management and scheduling according to the available resources and the chosen organisational model;
- **ERP logistics**: these applications support demand and supply planning together with order management. This implies handling inventory, orders and the required supply relations;
- **ERP production**: these applications are particularly voted to production planning. This entails dealing with bills of material and order status.
- **ERP project management**: these application support management staffs accounting and evaluating project costs and business structure.

In the last few years, trends have shown this traditional classification to be dogmatic: nowadays customers require applications with features that cross over more classes. Therefore, there is a need for answers that tackle both project and human resources and ERP related to the different project areas. The Whales system reflects this attitude.
As for information agent systems focused on heterogeneous information integration literature is rich. In this area, many projects based on mediator architecture have been developed A significant work is the MCC InfoSleuth [Perry98, Fowler99]. It is an agent-based system for information gathering and analysis tasks performed over networks of autonomous information sources. A key motivation of the InfoSleuth system is that real information gathering applications require long-running monitoring and integration of information at various levels of abstraction. To this end, InfoSleuth agents enable a loose integration of technologies allowing: (1) extraction of semantic concepts from autonomous information sources; (2) registration and integration of semantically annotated information from different sources; and (3) temporal monitoring, information routing, and identification of trends appearing across sources in the information network. Another experience is the RETSINA multi-agent infrastructure for in-context information retrieval [Sykara99]. In particular, the LARKS description language [Sykara99b] is defined to realize the agent matchmaking process (at both syntactic and semantic level) by using several different filters: Context, Profile, Similarity, Signature and Constraint matching. Differently from our approach, both InfoSleuth and RETSINA does not take advantage from the mobility feature of the agents, which we consider fundamental in a dynamic and uncertain environment such as the Internet.
Finally, as for performance analysis, numerous studies ([Dongarra93, Dikaiakos98, Culler98]) have tackled the field of benchmarking to evaluate performance properties of computer systems and applications, and to predict possible developments to improve features. Previous research has rationalized the advantages of the software agent approach over other state-of-the-art techniques [Papastavrou99] and provided models for benchmarking multi-agent system platform performance [Samaras99].

5. CONCLUSIONS

We have presented the business scenario in which the WINK system should operate and we have shown the originality of the WINK approach to networking, integration and cooperation. The approach is based on the integration of two promising system, namely the WHALES and MIKS systems. The two systems have matching functionalities that can be tailored to our particular domain of interest. Further, the WINK project requirements give the opportunity to enhance and extend those functionalities, resulting in an advanced application in the field of information systems. State-of-the-art open standards such as those proposed by W3C or FIPA will be taken into account throughout all phases of the project. Finally, the prototype will be tested on a pilot case proposed by Alenia Space. This will take the combined software module one step from becoming a market product suitable to extensive commercial use.

REFERENCES

[Benetti02] I. Benetti, D. Beneventano, S. Bergamaschi, F. Guerra, M. Vincini, "An Information Integration Framework for E-Commerce", IEEE Intelligent Systems Magazine, January/February 2002.

[Bergamaschi99] S. Bergamaschi, S. Castano, D. Beneventano, and M. Vincini, 2001. 'Semantic integration of heterogenous information sources'. Journal of Data and Knowledge Engineering, 36(3):215–249, 2001.

[Bergamaschi01] S. Bergamaschi, G. Cabri, F. Guerra, L. Leonardi, M. Vincini, F. Zambonelli : "Supporting Information Integration With Autonomous Agents", Fifth International Workshop CIA-2001 on COOPERATIVE INFORMATION AGENTS September 6 - 8, 2001 Modena, Italy.

[Bergamaschi02] S. Bergamaschi, G. Cabri, F. Guerra, L. Leonardi, M. Vincini, F. Zambonelli: "Exploiting Agents to Support Information Integration", Special Issue of the International Journal on Cooperative Information Systems vol. 11(3), 2002, to be published.

[Borgida89] A. Borgida, R.J. Brachman, D.L. McGuinness and L.A. Resnik, ``CLASSIC: A Structural Data Model for Objects," in SIGMOD, pages 5867, Portland, Oregon, 1989.

[Culler98] D. Culler, J. Singh and A. Gupta, Parallel Computer Architecture. A Hardware/software approach. Morgan Kaufmann, 1999.

[Dikaiakos98] M. D. Dikaiakos, A. Rogers, K. Steiglitz, Performance Modeling through Functional Algorithm Simulation. In Zobrist, G., Bagchi, K., and Trivedi, K., editors, Advanced Computer System design, chapter 3, pages 43-62. Gordon and Breach, 1998.

[Dongarra93] J. Dongarra, W. Gentzsch, eds, Computer Benchmarks, North Holland, 1993.

[Fowler99] J. Fowler, B.Perry, N.H. Nodine and B. Bargmeyer. Agent-based semantic interoperability in Infosleuth™, SIGMOD Record, 28(1):60-67, 1999.

[Gelati01] G. Gelati, F. Guerra, M. Vincini: "Agents Supporting Information Integration: the MIKS Framework", AI*IA Notizie, Periodico dell'Associazione Italiana per l'Intelligenza Artificiale, AnnoXIV, N.4, Dicembre 2001.

[Perry98] B .Perry, M. Taylor, and A. Unruh, "Information Aggregation and Agent Interaction patterns in Infosleuth". In Proceedings of the Fourth IECIS International Conference on Cooperative Information Systems, 1998.

[Samaras99] G. Samaras, M. D. Dikaiakos, C. Spyrou, A. Liverdos, Mobile Agent Platfomrs for Web Databases: A Qualitative and Quantitative Assessmen, Proceedings of ASAMA, pages 50-64, 1999.

[Papastavrou99] S. Papastavrou, G. Samaras, E. Pitoura, Mobile Agents for WWW Distributed Database Access, Proceedings of the Fifteenth International Conference on Data Engineering, pages 228-237, IEEE, March 1999.

[Sykara99] K. Sykara, In-Context Information Management truough Adaptative Collaboration of Intelligent Agents, Intelligent Information Agents, M. Klusch Ed. – Springer, 1999.

[Sykara99b] K. Sykara, M. Klusch, F. Widoff and J. Lu, Dynamic Service Matchmaking among agents in open information environments, SIGMOD-RECORDS, 28(1), March 1999.