

Investigating the Evolution of Electronic Markets

John Debenham and Simeon Simoff

Faculty of Information Technology, University of Technology, Sydney
{debenham, simeon}@it.uts.edu.au

Abstract. Markets evolve through ‘entrepreneurial’ intervention which is based on intuition and on timely information. An electronic market has been constructed in the laboratory as a collaborative virtual environment to identify timely entrepreneurial information for e-markets. This information is distilled from individual signals in the markets themselves and from signals observed on the Internet. Distributed, concurrent, time-constrained data mining methods are managed using business process management technology to extract timely, reliable information from this inherently unreliable environment.

1. Introduction

E-business is growing fast and with its growth there is a need for a better understanding of the mechanisms that govern e-markets and the way that they evolve. An electronic market has been constructed in the laboratory as a collaborative virtual environment or virtual ‘place’. This market is designed to support investigations into how electronic markets will evolve. It supports a comprehensive range of basic market transactions. These market transactions are managed using business process management technology. Market evolution takes place as the result of entrepreneurial intervention which may either be through the creation of new goods, services, needs or requirements, or through the novel combination of existing ones. Entrepreneurial intervention relies on intuition and on information. In this market, a suite of data mining tools are available to the would-be entrepreneur to ‘discover’ this information. To be of use to an entrepreneur these tools need to produce information and to estimate the value (including the validity) of that information within a set time. This information is distilled from individual signals in the markets themselves and from signals observed on the unreliable, information-overloaded Internet. This electronic market is designed to support experiments in business-to-business (B2B) markets and in business-to-consumer (B2C) markets.

The term “entrepreneur” is used here in the sense of [1]: “The aspect of knowledge which is crucially relevant to entrepreneurship is not so much the substantive knowledge of market data as alertness, the ‘knowledge’ of where to find market data.”. In general usage, the term entrepreneur refers to various aspects of business including the characteristics of successful and innovative business people and firms. The *entrepreneurial opportunity discovery process* lies at the heart of market innovation and evolution, in which new knowledge is created by bringing together partial knowledge of different elements of a market.

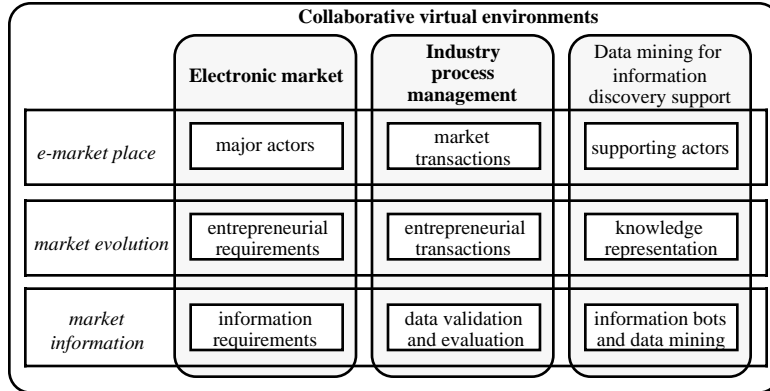


Fig. 1. Conceptual framework integrating the domains involved

The overall conceptual framework for this market place is shown in Fig. 1. It has been constructed as a collaborative virtual environment. At the *e-market place* layer, basic market transactions are managed as industry processes using a multiagent business process management system. These transactions include the creation of particular markets, the buying and selling of goods, and the packaging of goods, services, needs and requirements. The basic market transactions take place between a set of ‘major actor classes’ in the market—the decision-makers, which are described in Sec. 2. These transactions are monitored by ‘supporting actor classes’ that include data mining actors—these are described in Sec. 3. At the *market evolution* layer the entrepreneurial transactions are also managed by business process management technology. This includes the timely extraction of reliable information and knowledge—as described in Sec. 4—as well as entrepreneurial intervention. At the *market information* layer, atomic signals from information bots and other sources feed into the data mining actors, the reliability of these atomic signals is evaluated before the signals are combined by the process management system at the market evolution layer. The implementation of the market as a virtual ‘place’ is described in Sec. 5.

2. Electronic Markets

E-markets enable the study of the market evolution process using intelligent software and the vast resources of the Internet. First, e-markets are in themselves a fruitful arena for hunting entrepreneurial opportunities—witness the proliferation of new market forms and players [2], [3] and the characterisation of the new network economy as an “opportunity factory” [4]. Second smart computer systems may be applied to leverage the rich flows of information in e-markets leading to the discovery of potentially profitable opportunities and thereby invigorate the market evolution process. “[T]he value of the Internet and IT lies in their capacity to store, analyse and communicate information instantly, anywhere at negligible cost” [Economist, 27 Sep. 2000].

There is also a wealth of established work in economic theory, that describes the behaviour of rational agents performing basic transactions in traditional markets. For

example, this work includes the theory of auctions, the theory of bargaining, the theory of contracts and applied negotiation theory [5]. This work may not necessarily describe the behaviour of rational agents operating in electronic markets [6]. Little work has been done on how electronic market places will evolve, although the capacity of the vast amount of information that will reside in those market places, and on the Internet generally, to assist the evolution process is self-evident.

The environment of electronic markets is quite different to the traditional market environment, and so the scope for entrepreneurial activity in electronic markets can not be assumed to be the same. Substantial amounts of data may be acquired quickly and cheaply from the markets themselves. Further, valuable signals may be observed by searching the Internet, by reading news feeds, by watching corporate pages, and by reading background market data such as stock exchange data. If all of these individual signals can be distilled into meaningful information then it may be of use to an entrepreneur. The scope of possible entrepreneurial intervention will determine, and is determined by, the range of the support provided.

The basic market transactions take place in the electronic market place between a set of *major actor classes*; these are based on an extension of the model given in [3]. There are eight major actor classes—see Fig. 2. A central logical component is an “e-exchange” in which individual markets are created within which deals are done. This is actor class is essentially an “empty virtual box” for holding markets and for advertising their existence. The *buyer* and *seller* classes include those wishing to buy or sell any form of good or service, or any combination of goods or services, in a market. The remaining actor classes are all entrepreneurial and so are directly instrumental in market evolution. Members of the *content aggregator* actor class act as forward aggregators, they package goods and services, possibly from different sellers, and place these packages in markets in the e-exchange. Members of the *solution providers* class act as intermediaries in the negotiation of contracts and the development of business relationships. *E-speculators* take short term positions in markets in the e-exchange. Sell-side *asset exchanges* exchange or share assets between sellers. *Specialist originators* act as reverse aggregators, they coordinate and package orders for goods and services on behalf of various buyers. In our e-market model, the major actors, shown in Fig.2, are provided with distilled information by the supporting actors, which utilise a variety of data mining and information discovery methods.

3. Data mining and knowledge discovery

E-market is a killer domain for the new generation of data analysis and supporting techniques, collectively labeled as data mining methods. The methods have been successfully applied to stock market analysis, predictions and other financial and market analysis applications [7]. The application of data mining methods in e-business to date has predominantly been done within the framework of the B2C model. This research combines both e-commerce and “classical” data mining methods [8]. Rather than following the “classical” sequence, the data mining process in the e-market place is organized as a concurrent process supported by a number of specialised agents. There are seven supporting actor classes, illustrated in Fig. 3. The central actor, the




Fig 2. The major actor classes.




Fig. 3. Supporting actor classes.

information kiosk handles all the internal data and information extracted from this data, notifying the other actors of what is available in the format that they can understand. *Scanners* are the data collectors, responsible for monitoring company web sites, bringing news text and image data, retrieving data from stock markets and other signals. *Preprocessors* are basically data filters and formatters. They convert signal data to particular format and packed it in a “ready-to-use” format, adding information about the format, size, time/date stamp and other service information. *Miners* are the classes that perform the actual data mining exercise. Each miner implements particular machine learning method (the idea is an analogy of the approach taken by [9] in the development of their class library). Such schema permits to realise combined and/or hierarchical learning schemata, when the output produced by one miner becomes an input to another mining agent. *Composer* plays additional coordinating role between the mining agents. The *info-provider* and *info-visualiser* are supporting actors, which actually provide the distilled information to the major actors and human entrepreneurs.

An example of supporting actors in action is illustrated in Fig. 4. The initial data set, a Web page, supplied by the scanner, is preprocessed by the preprocessor in some structured form. The new data set is clustered by a SOM¹ miner, which produces a table of weights for each node. Such table is stored back in the information kiosk. It can be passed to a major actor in this form or the composer can invoke another mining actor, in this case and Apriori miner (an association rule miner which implements a version of the Apriori algorithm (for more details about the different versions of the Apriori algorithm see [8])). On the request from a major actor, a rule provider (a type of info-provider actor) supplies distilled rules. A map visualiser (a type of info-visualiser actor) provides an output from the table of weights in a form comprehensible for

¹ Self-Organising Map – a type of artificial neural network

formulate a hypothesis that there is a change in the Board of Directors in a large multinational corporation. This means supplying to supporting actors the name of the corporation and at least one person, who is expected to be changed. Such hypotheses can be either refined during the exploration, partially or completely reformulated or finally rejected.

The combination of data-driven and hypothesis driven process aims at providing a mechanism for meeting the tight time constraints. The data driven mining does not require specific request and can be handled in a “prefetched” mode, when the support agents actually are bringing data and extracting the necessary information, storing the results in the information kiosk. Consequently, some of the information is anticipated to be distilled before it is even required. On the other side, the major actors can access only what is currently distilled in the information kiosk. Thus, to deal with potentially incomplete information, it will be necessary to introduce several measures for an information discovery – importance, relevance to the transaction, potential impact on future transactions.

Managing and synchronizing the activities of the major and supporting actors is not a trivial task. In this e-market place this task is handled by methods and techniques developed in industry process management.

4. Industry process management

The term *industry processes* is used here to refer to processes that are triggered by both e-market transactions and entrepreneurial activities. This is in-line with the term *industry process re-engineering* which concerns the re-engineering of trans-corporate processes as electronically managed processes. Industry process re-engineering addresses the four issues of complexity, interoperability, communication and management. Of these four issues, complexity and management are particularly relevant here. Communication and management are dealt with by an agent architecture based on ‘near-failure-proof’ plans—as described below.

All e-market transactions are managed as industry processes. The problem of delivering information to an entrepreneur at the right granularity and at the right time is also a process management problem. As is the problem of effecting any intervention that may be triggered indirectly by that information. A single business process system manages both the e-market transactions and the entrepreneurial activities. The complexity in managing e-market transactions stems from the distributed nature of e-markets and from the complex time constraints that govern the operation of e-markets. The complexity in providing entrepreneurial support here stems from the unreliable and unpredictable nature of the Internet, and the problem of delivering something useful and reliable in the required time. This means that a powerful process management system is required, particularly in its capacity to manage heavily constrained and possibly interdependent processes. The system used is a multiagent system that is based on a three-layer, BDI (Belief-Desire-Intention) hybrid agent architecture. A high-level description of this architecture is given below.

One fundamental e-market transaction is the creation by a seller (or buyer) of a particular *market* within the e-exchange. To do this a real (or virtual) seller specifies the

goods or services, the market mechanism to be used, and any temporal constraints on the market including when the market will clear. This specification triggers the creation of an instance of the generic *market process* that first requests the creation of certain publicly accessible spaces in the e-exchange in which the market is to be represented. Once this has been achieved this process then advertises the existence of the market and the terms that apply to it, eventually it clears and closes the market. Following the advertisement of the existence of a market, bids (or asks) from a virtual buyer (or seller) trigger an instance of the generic *trading process* which manages the bids, settles the market entry fee (if any), meets the commitments of the bid if it is successful, and manages the retrieval (or delivery) of the goods or services traded. Trading processes may contain time constraints such as “if an attempt to reach an agreement with X on purchasing Y has not succeeded before 2.30pm then abandon that attempt and place a bid for Y in the market Z”. These two types of generic process are managed using plans.

The entrepreneurial activities are typically more complex than the e-market transactions. A fundamental entrepreneurial activity is a request for information. Such a request triggers an instance of an *information process* that may invoke any number of a set of basic information discovery and data mining methods (supporting actors) as well as hand-crafted analytical tools. These methods and tools may be applied to either the market data or to the Internet. Information processes may contain critical time constraints such as “get me the best information that you can before 3.00pm”. This information is derived from observing a large amount of market traffic in real-time, and from the inherently-unreliable, information-overloaded Internet. So there are four basic problems here. First, determining where to access the individual signals. Second, assessing some measure of confidence in the validity of those signals. Third combining those signals—which may be invalid and which may have been observed at different times and different places—into reliable advice. Fourth to do all this within tight time and cost constraints. Another entrepreneurial activity is *entrepreneurial intervention*. These transactions are innovative in nature. A level of generic support is provided for entrepreneurial activities that fall into the major actor classes of specialist originator, e-speculator, content aggregator and asset exchange. The solution provider class is very broad and hand crafting is usually required. These classes are all illustrated in Fig. 2.

Industry processes are modelled as state and activity charts [12] and are managed by plans that can accommodate failure [13]. The planning system provides the deliberative reasoning mechanism in the agent architecture. The successful execution of a plan in an industry process is not necessarily related to the achievement of that plan’s goal. For example, the plan may simply be bad, alternatively the goal may have been achieved by some other plan or externally to the management system. So if a plan executes successfully then its goal may not have been achieved; also, if a plan fails then its goal may have been achieved. This means that each plan should terminate with a *success condition* (SC) that is a procedure whose goal is to determine whether the plan’s goal has been achieved. The success condition is the final sub-goal on *every* path through a plan. The success condition is a procedure; the execution of that procedure may succeed (✓), fail (✗) or abort (A). If the execution of the success condition does not succeed then the overall success of the plan is unknown (?).

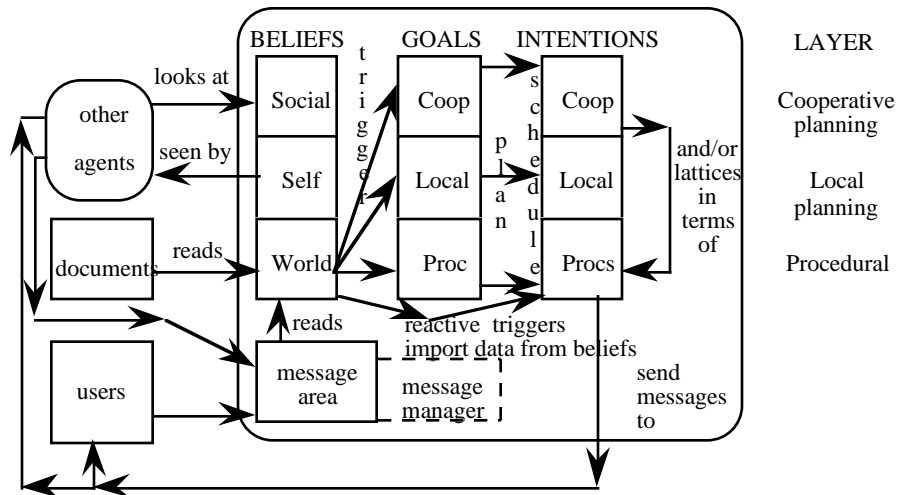


Fig. 5. Conceptual agent architecture

One well reported class of agent architecture is the three-layer, BDI hybrid architecture. One member of this class is the INTERRAP architecture [14], which has its origins in the work of [13]. The process agent architecture described here extends the INTERRAP architecture; it is based on a set of basic concepts or “mental categories”. The *conceptual architecture* describes how the agents are specified. The *control architecture* describes how the agents operate.

The conceptual architecture is shown in Fig. 5. In that architecture world beliefs are derived *either* from reading messages received from a user, *or* from reading the documents involved in the process instance, *or* from reading messages received from other agents. These activities are fundamentally different. Documents are “passive” in that they are read only when information is required. Users and other agents are “active” in that they send messages when they feel like it. The social beliefs of an agent are the observed self-beliefs of other agents in the system. Beliefs play two roles. First, they can be partly or wholly responsible for the activation of a local or cooperative trigger leading to the agent committing to a goal, and may thus initiate an intention (eg. the execution of a plan to achieve what a message asks, such as “find the value of BHP Corp’s ordinary stock”). A cooperative intention will involve other agents such as in a process generated by a content aggregator (eg. will you join with X by providing a service to maintain the computers that X sells thus creating a package deal). This is *deliberative reasoning*. Second, they can be partly or wholly responsible for the activation of a ‘procedure trigger’ that will pass data to a partly executed procedure and so may advance the agent’s progress in achieving some goal to which it is committed (eg. the message “the value of BHP Corp’s ordinary stock at 12.20pm was \$12.20” could lead to a belief that activates a procedure trigger). This is *reactive reasoning*. Reactive reasoning also involves “hard-wired” procedures such as a hard-wired procedure trigger that watches for the arrival of messages such as “market Z has closed”. The generic agent is implemented in Java; it is implemented as an interpreter of high-level

agent specifications. This interpreter enables agents to be built quickly and enables the specification of agent plans to be modified during execution. Implementation as an interpreter also simplifies general maintenance, which only has to deal with high level specifications.

Multiagent technology is an attractive basis for industry process re-engineering [15] [16]. A multiagent system consists of autonomous components that interact with messages. The scalability issue is “solved”—in theory—by establishing a common understanding for inter-agent communication and interaction. KQML (Knowledge Query and Manipulation Language) is used for inter-agent communication [17]. Specifying an inter-agent communication protocol may be tedious but is not technically complex. Standard XML-based ontologies will enable data to be communicated freely [18] but much work has yet to be done on standards for communicating expertise. Any process management system should take account of the “process knowledge” and the “performance knowledge. *Process knowledge* is the wisdom that has been accumulated, particularly that which is relevant to the process instance at hand. *Performance knowledge* is knowledge of how effective people, methods and plans are at achieving various things.

5. E-Market Place as a Collaborative Virtual Environment

Market activities are intrinsically social. E-market is a computer-mediated market. There is a number of ways to implement the amalgamation of the actors, representations, analysis methods and techniques presented in the previous sections. We seek a convenient and understandable integration of the underlying technology, which also supports human interactions within the market. Although, by its nature market activities are distributed across borders and places, handling activities in one “logical place” provides convenient mechanism for observing and collecting data about the behaviours, strategies and activities in the market. A significant effort in the research and development in collaborative virtual environments (CVEs) [19] has gone in the implementation of virtual ‘place’ models, which offer a convenient way of structuring and handling information, communication, exchange and other functionality in such environments [20]. Virtual worlds are a class of CVE environments that implement one possible elaboration of a virtual place and have the potential to provide scalable professional working environments that integrate communication, interaction, negotiation and other exchange between actors. In our approach the virtual world provides a place where the complete range of market activities is fully supported.

The implementation of an e-market lab for investigating the evolution of an e-market should be able to accommodate new emerging markets, perhaps as new e-market places. Fig. 6 illustrates the idea. The distributed e-market, populated with different e-market places is organised around the concept of an “universe”, populated with numerous “worlds”, respectively. Active Worlds virtual world technology implements this metaphor, providing means not only for developing universes and worlds in 3D representations, but also for populating these spaces with programmable objects.

One aspect of the virtual place model is the establishment of the identity of the agents (whether human or computer) in that place. In the context of our e-market place,




Fig. 6. The idea of Universe/World implementation of a distributed e-market as a collection of e-market places

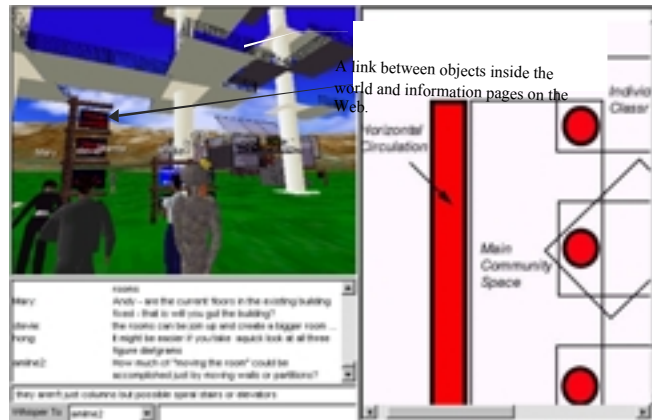


Fig. 7. An avatar in an e-market place represents an actor

the representation of the agent should be understandable to both sides, so to a human in the world an agent may appear as an avatar, as shown in Fig. 7. Thus in our e-market an avatar is a 3D model of an actor and shows where they are, where they are looking, and what gestures they want to communicate.

Technically, the ontology of a virtual world defines the way the agents interact in it. Fig. 8 illustrates the idea. The interaction happens via the Behavior and Transformation Group Nodes, when the representation is placed in a Shape3D node.

6. Conclusion

One of the innovations in this research is the development of a coherent environment for e-market places and a set of virtual agents for representing the actors. Virtual worlds are a smart way of doing this because they provide a single environment, whose consistent organisation can significantly reduce the cognitive and information overload, and they provide a coherent environment for interactive data mining and decision making. Data (market, communication, and strategic) is warehoused in one place, and is processed by the agents that populate the e-market place. The use of a powerful business process

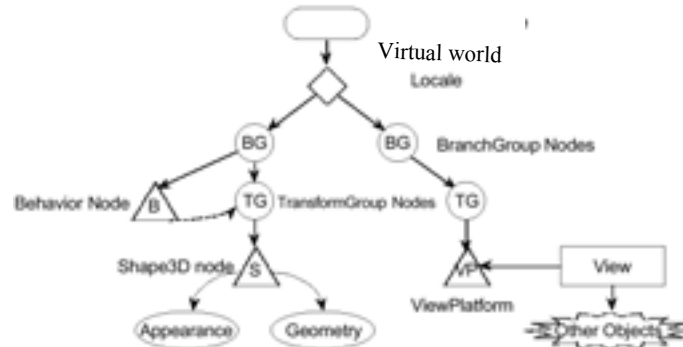


Fig. 8. An ontology of a virtual universe (adapted from Java3D ontology)

management system to drive all the electronic market transactions unifies the whole market operation. By merging the functionality of e-business and the functionality of the place provided by a virtual world, a new, more effective and coherent environments for e-business will be developed. The development of computational models of market actors, deploying those models in the e-market place, and including them as part of the building blocks for creating e-market places is significant in providing a practical instrument for further research and development in electronic markets.

References

1. Israel M. Kirzner. *Competition & Entrepreneurship*. University of Chicago Press, 1973.
2. Kaplan Steven and Mohanbir Sawhney. *E-Hubs: The New B2B Marketplace*. Harvard Business Review 78 May-June 2000 97-103.
3. R. Wise & D. Morrison. *Beyond the Exchange; The Future of B2B*. Harvard Business review Nov-Dec 2000, pp86-96.
4. Kelley, Kevin. *New Rules for the New Economy*. New York, Penguin Books, 2000.
5. Peyman Faratin. *Automated Service Negotiation Between Autonomous Computational Agents*. PhD dissertation, University of London (Dec 2000).
6. Moshe Tennenholtz. *Electronic Commerce: From Economic and Game-Theoretic Models to Working Protocols*. Invited paper. Proceedings Sixteenth International Joint Conference on Artificial Intelligence, IJCAI'99, Stockholm, Sweden.
7. B. Kovalerchuk & E. Vityaev. *Data Mining in Finance: Advances in Relational and Hybrid Methods*. Kluwer, 2000.
8. Han, J. & Kamber, M (2001). *Data Mining: Concepts and Techniques*, Morgan Kaufmann, San Francisco, CA, 2001.
9. Witten, I. & Frank, E. (2000). *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*, Morgan Kaufmann Publishers, San Francisco, CA.
10. Uschold, M. and Gruninger, M.: 1996, *Ontologies: principles, methods and applications*. Knowledge Engineering Review, 11 (2), 1996.
11. Guarino N., Masolo C., and Vetere G., *OntoSeek: Content-Based Access to the Web*, IEEE Intelligent Systems 14(3), May/June 1999, pp. 70-80

12. Muth, P., Wodtke, D., Weissenfels, J., Kotz D.A. and Weikum, G. "From Centralized Workflow Specification to Distributed Workflow Execution." In *Journal of Intelligent Information Systems (JIIS)*, Kluwer Academic Publishers, Vol. 10, No. 2, 1998.
13. Rao, A.S. and Georgeff, M.P. "BDI Agents: From Theory to Practice", in proceedings First International Conference on Multi-Agent Systems (ICMAS-95), San Francisco, USA, pp 312—319.
14. Müller, J.P. "The Design of Intelligent Agents" Springer-Verlag, 1996.
15. Jain, A.K., Aparicio, M. and Singh, M.P. "Agents for Process Coherence in Virtual Enterprises" in *Communications of the ACM*, Volume 42, No 3, March 1999, pp62—69.
16. Jennings, N.R., Faratin, P., Norman, T.J., O'Brien, P. & Odgers, B. Autonomous Agents for Business Process Management. *Int. Journal of Applied Artificial Intelligence* 14 (2) 145—189, 2000.
17. Finin, F. Labrou, Y., and Mayfield, J. "KQML as an agent communication language." In Jeff Bradshaw (Ed.) *Software Agents*. MIT Press (1997).
18. Robert Skinstad, R. "Business process integration through XML". In proceedings XML Europe 2000, Paris, 12-16 June 2000.
19. Capin, T.K., Pandzic, I. S., Magnenat-Thalmann, N. and Thalmann, D. *Avatars in Networked Virtual Environments*. John Wiley and Sons, Chichester, 1999.
20. J.A. Waterworth. Spaces, places, landscapes and views: experiential design of shared information spaces. In A. J. Munro, K. Hook and D. Benyon (eds), *Social Navigation of Information Space*, Springer, London, pp. 132-154, 1999.