An Information-Rich, Virtual Trading Environment

Les Green
Faculty of Information Technology
University of Technology, Sydney
PO Box 123, NSW 2007,
Australia
lesgreen@it.uts.edu.au

KEYWORDS
E-Market, Synthetic Environment, Business Process

ABSTRACT

The E-Markets group at the University of Technology, Sydney has been undertaking research on the evolution of trading environments. The main focus is on mechanisms for innovation. Work to date includes: simulation experiments to investigate the relative impact and cost of the fundamental evolutionary mechanisms, the construction of an e-Market kernel embedded in a virtual world, the construction of smart mining bots to extract and condense information from the market context, and a process management system. In this project every transaction is managed as a business process. Future work focuses on understanding the evolution of business networks in the virtual marketplace. This paper presents the current state-of-art project in information rich, virtual trading environments.

INTRODUCTION

The overall goal of this work is to derive fundamental insight into how the next generation of e-market will evolve. The work is mapping out and exploring future generation trading environments. e-Marketplaces will be complete and immersive. **Complete** in that all transactions, including requests for information, will be handled within the e-marketplace and so it contains all information sources. **Immersive** in that e-marketplaces are 'virtual realities' in which players interact as lifelike avatars.

Another key aim is to understand the evolution of business networks that are established in a virtual marketplace. Three fundamental evolutionary mechanisms are innovation, imitation and improvement of existing procedures. The main focus is on mechanisms for innovation. The perturbation of market equilibrium through entrepreneurial action is the essence of market evolution by innovation. Entrepreneurship relies both on intuition and information discovery. The term 'entrepreneur' is used here in its technical sense (Watkins 2002). Smart systems that deliver timely information to support the market evolutionary process in a virtual marketplace are described here. The deep research goals that motivate this work are to understand the processes that will drive both market evolution and the evolution of business relationships in a future generation trading environment.

Work to date includes: simulation experiments to investigate the relative impact and cost of these three fundamental evolutionary mechanisms, the construction of an e-Market kernel, the construction of smart mining bots to extract and condense information from the market context, and a process management system that handles the transactions in such a place. The goal of the bots constructed to date is to identify timely information for traders in an e-market. The traders are the buyers and sellers.

The e-market was designed by Professor John Debenham and Professor Simeon Simoff. It is a virtual marketplace that is embedded in virtual worlds technology. The actors' assistant is being constructed in the Faculty of Information Technology at the University of Technology, Sydney. It is funded by an Australian Research Council three-year Discovery Grant, and various other grants. Our e-marketplace consists of a market kernel embedded in an information-rich environment. It is implemented within virtual reality space.

The smart information system addresses the problem of identifying timely information for e-markets with their rapid, pervasive and massive flows of data. This information is distilled from individual signals in the markets themselves and from signals observed on the unreliable, information-overloaded Internet. Distributed, concurrent, time-constrained data mining methods are managed using business process management technology to extract timely, reliable information from this unreliable environment.

One feature of the whole project is that every transaction is treated as a business process and is managed by a process management system. In other words, the process management system makes the whole thing work. These transactions include simple market transactions such as "buy" and "sell" as well as transactions that assist the actors in buying and selling. The process management system is based on a robust multiagent architecture and has been constructed in Jack. The use of multiagent systems is justified first by the distributed nature of e-business, and second by the critical nature of the transactions involved.

The overall goal of this work is to investigate the evolution of e-markets. We look at some factors that cause changes in electronic markets. A market is said to be in *equilibrium* if there is no opportunity for arbitrage—ie. no opportunity for no-risk or low-risk profit. Real markets are seldom in equilibrium. Market *evolution* occurs when a player identifies an innovative opportunity for arbitrage and takes advantage of it possibly by introducing a novel form of transaction. At a deeper level, business networks that support market activity also evolve and are the key focus of this investigation.

A trading agent strives to make informed decisions in an information-rich environment. Its design was provoked by the observation that agents are not always utility optimizers, and by the quotation: "Good negotiators, therefore, undertake integrated processes of knowledge acquisition combining sources of knowledge obtained at and away from the negotiation table. They learn in order to plan and plan in order to learn" (Watkins 2002). The agent attempts to fuse the negotiation with the information generated both by and because of it. It reacts to information derived from its opponent and from the environment, and proactively seeks missing information that may be of value.

Business networks are the result of the *interaction* between the players (traders) in the electronic market environment. These interactions are facilitated by synthetic characters presented in more detail in the next section.

SYNTHETIC CHARACTERS

To represent autonomous agents in the information rich environment, anonymous social agents have been embedded in the UTS eMarket environment. Avatars represent these agents visually in the eMarket space. They provide casual, but context-related conversation with anybody who chooses to engage with them. This feature markedly distinguishes this work from typical role-based chat agents in virtual worlds. The latter usually are used as "helpers" in navigation, or for making simple announcements. They are well informed about recent events, and so "a quick chat" can prove valuable. A demonstration that contains one of these synthetic characters is available at: http://research.it.uts.edu.au/emarkets/ — first click on "virtual worlds" under "themes and technologies" and then click on "here". To run the demonstration you will need to install the Adobe Atmosphere plugin.

The following issues have to be addressed in the design of any synthetic character:

- its appearance
- its mannerisms
- its sphere of knowledge
- its interaction modes
- its interaction sequences

and how all of these fit together.

Appearance is a key issue in initiating contact but is not addressed here. Mannerisms are equally important—a researcher associated with the UTS group is investigating the impact of facial expressions on interaction. This is a major issue that is beyond the scope of this discussion.

An agent's sphere of knowledge is crucial to the value that may be derived by interacting with it (Barthelemy et al. 2004). We have developed extensive machinery, based on unstructured data mining techniques, that provides our agent with a dynamic information base of current and breaking news across a wide range of financial issues. For example, background news is extracted from on-line editions of the Australian Financial Review.

Our agent's interaction modes are presently restricted to passive, one-to-one interaction. They are passive in that these agents do not initiate interaction. We have yet to deal effectively with the problem that occurs when a third party attempts to "barge in" on an interaction.

The interaction sequences are triggered by another agent's utterance. The machinery that manages this is designed to work convincingly for interactions of up to around five exchanges only. Our agents are designed to be "strangers" and no more—their role is simply to "toss in potentially valuable gossip".

The aspects of design discussed above cannot be treated in isolation. If the avatars are to be "believable" then their design must have a unifying conceptual basis. This is provided here by the agent's *character*. The first decision that is made when an agent is created is to select its character using a semi-random process that ensures that multiple instances of the agent in close virtual proximity have identifiably different characters.

The dimensions of character that we have selected are intended specifically for a finance-based environment. They are:

- Politesse
- Dynamism
- Optimism
- Self-confidence

The meaning of each of these dimensions is (moderately) independent of the others:

- Politesse means the use of polite words, phrases and forms
- Dynamism is the tendency to react rapidly, succinctly and vigorously
- Optimism here means a tendency to use up-beat phrases and the tendency to not use negations
- Self-confidence here means the tendency to respond with declarative statements rather than tentative propositions or questions

The selection of the character of an agent determines: its appearance (ie: which avatar is chosen for it) and the style of its dialogue. Future plans will address the avatars mannerisms. The selection of agent's character provides an underlying unifying framework for how the agent appears and behaves, and ensures that multiple instances of the agent appear to be different.

The selection on an agent's character does not alone determine its behavior. Each agent's behavior is further determined by its moods that vary slowly but constantly. This is intended to ensure that repeated interactions with the same agent have some degree of novelty. The dimensions of moods that we have identified are:

- Happiness
- Sympathy
- Angry

Synthetic characters provide only one of the components of the immersive trading environment. In the next section we present the negotiation support.

AUTOMATED NEGOTIATION

We illustrate the automated negotiation approach on the example of negotiation between two agents in an information-rich environment. That is, the agents have access to general information that may be of use to them. They also exchange information as part of the negotiation process. The two agents are called "me" and my opponent ω. The environment here is the Internet, in particular the World Wide Web, from which information is extracted on demand using special purpose 'bots'. So my agent, "me", may receive information either from ω or from one of these bots. In a 'real life' negotiation, the sort of information that is tabled during a negotiation includes statements such as "this is the last bottle available", "you won't get a better price than this", and so on. To avoid the issue of natural language understanding, and other more general semantic issues, the interface between each of the negotiating agents and these information sources is represented using the language of first-order, typed predicate logic, and a set of pre-agreed, pre-specified predicates. All information passed to "me" is expressed in this way.

As well as being unambiguous, the use of first-order typed predicate logic has the advantage of admitting metrics that describe, for example, how "close" two statements are. These metrics are useful in enabling the agent to manage the information extracted from the environment in a strategic way.

Negotiation between two trading agents is a two-stage process. First, the agents exchange offers whilst acquiring and exchanging information. Second, they attempt to reach a mutually satisfactory agreement in the light of that information. This process aims to reach informed decisions in eMarket bargaining by integrating the exchange of offers and the acquisition and exchange of information drawn from the environment.

Negotiation then proceeds by a loose alternating offers protocol that is intended to converge when the agents believe that they are fully informed.

Each agent exchange offers at each discrete time. The agents enter into a commitment if one of them accepts a standing offer. The protocol has three stages:

1. Initial offers from both agents;
2. A sequence of alternating offers, and
3. An agent quits and walks away from the negotiation.

The negotiation ceases either in the second or final round if one of the agents accepts a standing offer or in the final round when one agent quits and the negotiation breaks down.

PROCESS MANAGEMENT

To successfully manage a large number of negotiation threads, every market transaction is managed as a business process. To achieve this, suitable process management machinery is required. To investigate what is 'suitable' the essential features of these transactions are related to two classes of process that are at the 'high end' of process management feasibility. The two classes are goal-driven processes and knowledge-driven processes. The term "business process management" is generally used to refer to the simpler class of workflow processes (Fischer 2000), although there notable exceptions using multiagent systems (Jennings et al. 2000) (Huhns and Singh 1998).

Goal-Driven Processes

A *goal-driven process* has a process goal, and achievement of that goal signals the termination of the process. The process goal may have various decompositions into possibly conditional sequences of sub-goals such that these sub-goals are associated with (atomic) activities and so with atomic tasks. Some of these sequences of tasks may work better than others, and there may be no way of knowing which is which (van der Aalst and van Hee 2001). A task for an activity may fail outright, may fail to achieve its goal, or may violate process constraints (Debenham 2004). In other words, a central issue in managing goal-driven processes is the management of task failure. Hybrid multiagent architectures whose deliberative reasoning mechanism is based on "succeed/fail/abort plans" (Rao and Georgeff 1995) are well suited to the management of goal-driven processes.

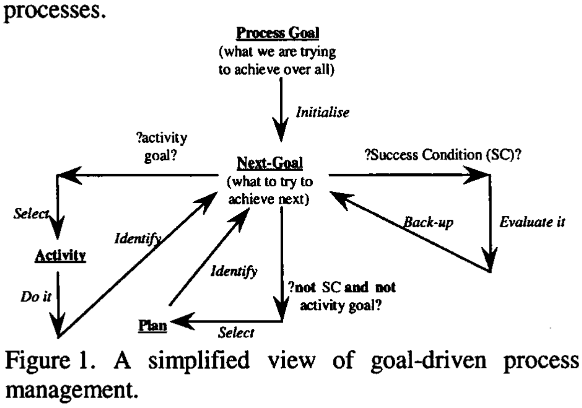

Figure 1. A simplified view of goal-driven process management.

Knowledge-Driven Processes

A second class of process, whose management has received little attention, is called knowledge-driven processes (Debenham 2000). *Process knowledge* is all the knowledge that is relevant to a process instance. It includes common-sense knowledge, knowledge that was available when an instance is created, and knowledge acquired during the time that the instance exists. A *knowledge-driven process* may have a process goal, but the goal may be vague and may mutate. In so far as the process goal gives direction to goal-driven—and activity-driven—processes, the process knowledge gives direction to knowledge-driven processes. The body of process knowledge is typically large and continually growing and so knowledge driven processes are seldom considered as candidates for process management. They are typically supported, rather than managed, by CSCW systems (van der Aalst and van Hee 2001). But, even complex knowledge-driven processes are "not all bad"—they typically have goal-driven sub-processes which may be handled as described above. *Knowledge-base processes* are a special type of knowledge-driven process for which the process knowledge *can* be represented and accessed by a process management system. This proves to be a useful concept for managing e-marketplace transactions.

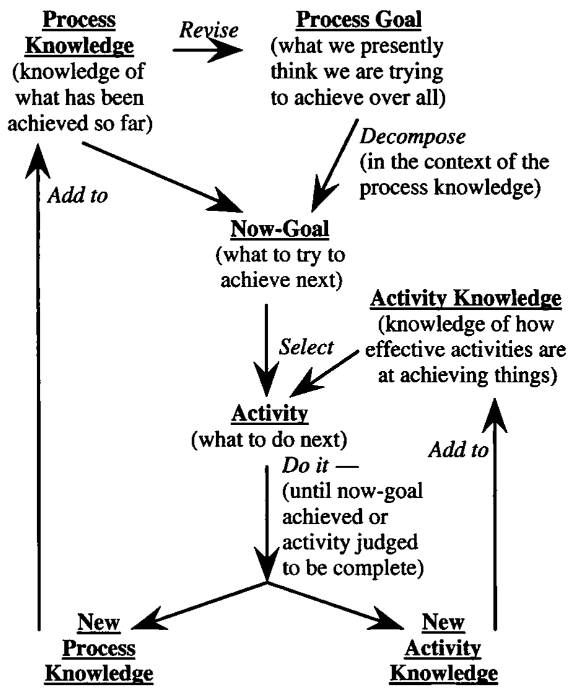

Figure 2. A simplified view of knowledge-driven process management.

The management of goal-driven and knowledge driven processes are radically different. Goal-driven processes may be managed by a goal/plan decomposition process (see Fig. 1), and knowledge-driven processes are managed by continually reviewing the growing corpus of process knowledge—this is illustrated in Fig. 2. That Figure is deceptively simple in that the business of managing the process knowledge and of revising the process-goal and next-goal in the light of that growing body of knowledge is far from trivial in even simple examples.

This approach allows us to incorporate and synchronise the work of the main players – the trading agents in the virtual trading environment with the requests and provision of the necessary market information, performed by data and text mining agents.

DATA AND TEXT MINING

Data and text mining agents are responsible for extracting, verifying, combining, filtering and condensing information drawn from the trading environment and from the Internet. This information is then fed to the players in the e-marketplace. A data stream can include both structured (numerical and symbolic data) and unstructured (text, images and video data). Specialised data mining 'bots' (software agents), each of which applies a specific data mining method, are developed for extracting patterns from these data streams.

The agents in this project support data mining as a concurrent process, which allows us to identify potentially useful patterns within a large variety of rapidly-changing signals. These technologies include mining text data from news feeds as well as text and numerical data from financial reports and web pages. Text mining bots identify particular topics and trace those topics in the news stream. An initial prototype of a multi-agent system for identifying relevant events in news feeds and tracking those events from news sources on the Internet is fully operational. Potential future work includes the development of image and video mining technologies which complement the text mining tools, allowing us to extract context information from additional (alternative) media sources.

This work does not, for example, attempt to use gathered data to predict whether the US dollar will rise or fall against the UK pound. What it does attempt to do is predict the value that an actor in the system will place on the data (Han et al. 1999). So the feedback here is provided by the 'bots' in the form of a rating of the material found. A five point scale runs from 'totally useless' to 'very useful' to classify this data.

As mentioned in section 1, the environment is complete, i.e. all transactions are handled within. Information about these transactions is collected in the application level logs. This information is the source for a further type of data mining – Network Mining, which aims at discovery and analysis of business networks and their evolution.

NETWORK MINING

Uncovering the patterning of people's interaction has been the subject of social network analysis

(Wasswerman and Faust 1994). Relational data mining (Dzeroski and Lavrac 2001) extends these techniques beyond the comparative static methods typical of existing analysis systems (e.g. UCINET, Inflow) to the examination of large continuous flow data sets representing links between a variety of entities (e.g. friendship ties, links between places, things, ideas, data, web pages, events). Existing data mining techniques that are currently considered include a number of probabilistic relational techniques (Kersting and De Raedt 2001).

The work involves the analysis of activities and the diffusion of information in networks of virtual communities leading to the emergence of new business models and networks (Hagel and Armstrong 1997). Relevant background to this project are recent works on identifying the "network value" of a person (node) within a social network structure (e.g. (Domingos and Richardson 2001)), based on person interactions and information diffusion, and research on the role and importance of social capital and network competence in business performance and innovation (Ritter et al. 2002). Our work uses the continual, extensive and complex data flows from the virtual market to detect, measure and track static and dynamic patterns of change in social, business and knowledge networks and relate this to innovation events and the evolution of the players' strategies.

CONCLUSIONS AND FUTURE RESEARCH

E-markets in their current form must evolve into Virtual Trading Environments where well informed decisions are made from an information rich, immersive environment. In order for these decisions to be made, information must be gathered from the environment as well as other actors and extracted from the interactions which happen between them. Data and text mining as well as network mining techniques have been applied in this area.

In the discussed work, all transactions are handled by a process management system where both knowledge based and goal based processes are discussed. The environment has been expressed by a virtual world where avatars represent the actors involved and other contributing characters such as informative 'bots'.

The E-Markets group at the University of Technology, Sydney has developed an appropriate framework to handle the evolution of a Virtual Trading Environment and facilitate effective operation.

The project recognises that social relationships and knowledge networks (networks of ideas) established in the virtual marketplace are central to the recognition of opportunities, productive links among existing ideas, and to their exploitation and implementation. The future research will investigate the nature of business networks that can be established in an electronic marketplace, their evolution and possible ways in assisting and guiding their development.

REFERENCES

Barthelemy F, Dosquet B, Gries S & Magnant X. 2004, Believable Synthetic Characters In A Virtual Emarket. in *proceedings International Conference on Artificial Intelligence And Applications AIA 2004*, February 16-18, 2004, Innsbruck, Austria

Debenham, JK. 2000, Supporting knowledge-driven processes in a multiagent process management system. In proceedings Twentieth International Conference on Knowledge Based Systems and Applied Artificial Intelligence, ES'2000: Research and Development in Intelligent Systems XV, Cambridge UK, December 2000, pp273—286.

Debenham J K. 2004, Electronic Market Transactions Managed as Knowledge-Driven Processes. In proceedings The 9th IEEE International Conference on Engineering of Complex Computer Systems. Florence, Italy, 14-16 April, 2004

Domingos, P. and Richardson, M. 2001, Mining the Network Value of Customers. Proc. 7th Int. Conf. on Knowledge Discovery and Data Mining. San Francisco, CA 57-66

Dzeroski, S. and Lavrac, N. (eds), 2001, Relational Data Mining. Springer-Verlag, Berlin

Fischer, L. (Ed). 2000, Workflow Handbook 2001. Future Strategies

Hagel III, J. and Armstrong, A.G. 1997, Net Gain: Expanding Markets Through Virtual Communities. Harvard Business School Press, Boston, Massachusetts

Han, J. Lakshmanan, L.V.S. and Ng, R.T. 1999, Constraint-based multidimensional data mining. IEEE Computer, 8, 46-50.

Huhns, M.N. and Singh, M.P. 1998, *Managing heterogeneous transaction workflows with cooperating agents.* In N.R. Jennings and M. Wooldridge, (eds). Agent Technology: Foundations, Applications and Markets. Springer-Verlag: Berlin, Germany, pp. 219—239.

Jennings, N.R., Faratin, P., Norman, T.J., O'Brien, P. and Odgers, B. 2000, Autonomous Agents for Business Process Management. Int. Journal of Applied Artificial Intelligence 14 (2) 145-189

Kersting, K. and De Raedt, L. 2001, "Adaptive Bayesian Logic Programs. In: Rouveirol, C. and Sebag, M. (eds): Proceedings of the 11th OInternational Conference on Inductive Logic Programming. Springer-Verlag, Heidelberg (2001), 104-117

Rao, A.S. and Georgeff, M.P. 1995, "BDI Agents: From Theory to Practice", in proceedings First International Conference on Multi-Agent Systems (ICMAS-95), San Francisco, USA, pp 312—319.

Ritter, T., Wilkinson, I.F. and Johnston, W.J. 2002, Measuring network competence: some international evidence. Journal of Business and Industrial Marketing. 17 (2/3) (2002) 119-138.

Van der Aalst, W.M.P & Van Hee, K.M. 2001, *Workflow Management: Models, Methods, and Systems.* MIT Press

Wasserman, S. and Faust, K. 1994, Social Network Analysis: Methods and Applications. Cambridge University Press, Cambridge

Watkins, M. 2002, *Breakthrough Business Negotiation - A Toolbox for Managers.* Jossey-Bass.