

Curious Negotiator

Simeon Simoff and John Debenham

Faculty of IT, University of Technology, Sydney

{simeon, debenham}@it.uts.edu.au

ABSTRACT

In negotiation the exchange of information is as important as the exchange of offers. The curious negotiator is a multiagent system with three types of agent. Two negotiation agents, each representing an individual, develop consecutive offers, supported by information, whilst requesting information from its opponent. A mediator agent, with experience of prior negotiations, suggests how the negotiation may develop. A failed negotiation is a missed opportunity. An observer agent analyses failures looking for new opportunities. The integration of negotiation theory and data mining enables the curious negotiator to discover and exploit negotiation opportunities. Trials will be conducted in electronic business.

1. INTRODUCTION

The *curious negotiator* is a multiagent system of competitive agents [Saunders 2001] supporting multi-attribute negotiation where the set of issues is not fixed [Gerding et al, 2000]. In 2001 initial work designed and built an e-exchange [Maes et al., 1999] in a joint collaboration with Bullant Australasia Pty Ltd – an Australian software house with a strong interest in business-to-business (B2B) e-business. The e-exchange was designed by the authors. Present work in the existing project is using the e-exchange and is focusing on the initial goal of building a suite of tools to identify arbitrage opportunities [Bjork, 1999]. These tools support the e-Speculator actor class. The working hypothesis is that these tools will work as a “pump primer” in practical experiments that will lead us to the deeper question of market evolution.

The curious negotiator is designed to complement and complete the work in an existing project in e-Markets. The existing project, funded for three years from 2002 by the Australian Research Council, is investigating the evolutionary mechanisms in e-Markets focusing on business conducted in electronic exchanges. The authors are chief investigators in that project. The set of actor classes that unify the existing project and the curious negotiator is based on [Wise & Morrison, 2000]. The classes are illustrated in Fig. 1. We have extended this model, adding the sell-side “Content aggregator” component. Central logical components are “e-exchanges” in which one-off deals are done, and “solution providers” through whom contracts are negotiated and business relationships developed. E-speculators take short-term positions in an e-exchange and look for medium term arbitrage opportunities. Sell-side “Asset Exchange” components exchange or share assets between sellers. Content Aggregators—acting as forward aggregators—coordinate and package goods and services from various sellers. Specialist Originators—acting as reverse aggregators—coordinate and package orders for goods and services from various buyers. In simple terms the existing project is considering all except the Solution Provider class. The machinery that the Solution Provider class will need to negotiate is not yet clearly understood [Wise & Morrison, 2000]; the curious negotiator presented in this paper aims to build this machinery and so to complete the whole e-market picture [Tennenholtz, 1999] [Bichler, 2001].

Both projects are concerned with the relationship between negotiation (including complex bargaining [Muthoo, 1999] and simple dealing [Osborne and Rubinstein, 1990]) and knowledge. Before presenting the design of the curious negotiator we discuss negotiation process in terms of mechanisms, strategies and information discovery about the context in which the negotiation takes place.




Fig. 1. Major actor classes in the e-market

1.1. Negotiation process and negotiation mechanisms

Negotiation is the process whereby two (or more) individual agents with conflicting interests reach a mutually beneficial agreement on a set of issues. The term *bargaining* is often used in the literature to refer to negotiation between just two individuals. A negotiation may be assisted by a third-party *mediator* who facilitates and accelerates the process. The individuals involved in a negotiation operate in accordance with a set of rules called the *negotiation mechanism*. Negotiation mechanism specifies how the negotiation will take place, i.e. what each of the negotiating individuals may or should do as the negotiation proceeds, and when they should do it; when offers made are binding and what happens if such commitments are broken. The key components of negotiation include the players, issues, information, time, mechanism and links between the current and other negotiations. More recently the work on the theory and practice of negotiation has been extended to support interaction between rational autonomous software agents [Rosenschein & Zlotkin, 1998; Krauss, 2001]. [Faratin, 2002] describes tactics for the manipulation of the utility of deals, trade-off mechanisms that manipulate the value rather than the overall utility of an offer, and manipulation mechanisms that add and remove issues from the negotiation set. Further research in electronic markets has produced negotiation mechanisms for those markets [Bichler, 2001]. The open research issues identified by these researchers include multi-issue negotiation mechanisms, combinatorial negotiation mechanisms (that support simultaneous negotiation on a number of items) and mechanisms for the sequential sale of related items. In the past two years taxonomies have been developed to classify the mechanisms (e.g., the “London Classification¹”). Curious negotiator is focused on utilization of contextual information in the negotiation process: how the process creates demands for information; how those demands may be met; and how the information influences the offers and issues and so influences the whole negotiation process.

1.2. Negotiation strategies and contextual information

Given a negotiation mechanism, an individual will develop a *negotiation strategy* that aims to ensure the negotiation proceeds in the individual’s interests—whether an agreement is reached or not. Early negotiating agents in agent-mediated electronic commerce used several simplistic fixed strategies [Klusch, 2001], e.g., MIT Kasbah agent uses three predefined strategies. A negotiation strategy should generally rely on information drawn from the context of the negotiation. The significance of information to the negotiation process was analysed formally in the seminal paper by Milgrom and Weber [1982] in which the Linkage Principle, relating the revelation of contextual information to the price that a purchaser is prepared to pay, was introduced. Their analysis is limited to single-issue negotiation using conventional auction mechanisms.

¹ <http://enegotiations.wu-wien.ac.at>

The information generated during the negotiation process should assist each player to gauge and accommodate the interests of their opponent [Benn et al., 1999]. “Good negotiators, therefore, undertake integrated processes of knowledge acquisition that combine 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 curious negotiator encapsulates this observation. The flow of information between negotiating agents is seen as an integral part of a negotiation. During a negotiation, an agent may actively acquire *contextual information* that it may, or may not, choose to place on the negotiation table. It includes information that can be extracted from: (i) *results* of previous negotiations such as market information on deals struck or being struck; (ii) the *behaviour* of the negotiating parties like a view on the opponent’s current negotiation strategy; (iii) general *sources*, such as general information on the opponent (eg: reputation) and negotiation issues.

1.3. Negotiator ‘curiosity’, data mining and information discovery

Contextual information can be characterised as unexpected, interesting, and even novel. The main streams involved in the research in computational novelty (i.e. ‘curiosity’ and ‘creativity’) include knowledge-based systems, (“constrained exploration” within the parameters of the given knowledge), grammars and heuristic search (in which a solution is constructed gradually, with heuristics employed to decide how to choose each successive part) [Boden, 1998]. Data mining and evolutionary computation take the research in computational creativity to a new level. Data mining focuses on algorithms for extracting valid, potentially useful and understandable structure from data [Hand et al., 2001]. Curious negotiator is designed to incorporate data mining and information discovery methods that operate under time constraints, including methods from the area of topic detection and event tracking research [Franz, et al., 2001]. Curious negotiator agent system utilises the research in integration of data mining technology in e-commerce applications [Ansari et al., 2001]. We consider novelty and unexpectedness of discovered patterns with respect to a belief system (the current belief system of the negotiator), similar to [Padmanabhan & Tuzhilin, 1998]. Evolutionary systems approached in this project will be used to identify novel strategies that emerge during the negotiation process (e.g., by analogy to the approach presented in [Smith et al. 2002]).

As people and agents are the participants in negotiation, trust and reputation are other key elements of contextual information in negotiation, together with information about people types. To enable agents to evaluate the reputation of other agents (parties) involved, the curious negotiator incorporates a computational representation of trust. Closer to the curious negotiator is the computational approach of reputation management developed in [Yu & Singh, 2000]. Agents assist and represent people and businesses involved in electronic activities (commerce, government, politics, education), forming an interacting *electronic community*. Our approach includes evaluation metrics and mechanisms for: reputation rating and propagation of that rating; incorporating estimates from different sources and; selection rewards and penalties that affect individual trust estimates. These mechanisms employ data mining techniques for network evaluations of negotiating parties, for identifying possible influences and interactions. [Domingos & Richardson, 2001] offer a good starting point for the curious negotiator. To be able to evaluate types of people in terms of character features, curious negotiator employs techniques like enneagram²

1.4. Negotiation process, process management and multi-agent systems

Negotiation is goal-directed in the sense that individual agents involved in a negotiation may—probably will—have agendas of their own. But the agendas of the negotiating agents may be

² <http://www.enneagraminstitute.com/>

incompatible—there may be no solution that satisfies them all. Further the existence of a solution is unlikely to be known when the negotiation commences. So it may not be useful to consider negotiation as a search problem because the solution space may be empty whilst the negotiating agents may believe that it is not so. If the negotiation is a multi-issue negotiation for which the issue set is open [i.e. it can change at any stage in the negotiation] then the agendas of the individual negotiating agents must necessarily be at a higher level than the issues because the issues are unknown, and may even be issues that ‘had never occurred’ to one of the agents. So for multi-issue negotiation the agendas of the agents can not in general be an even high level goal such as ‘to maximise profit on the deal’ as the deal space is unknown.

In summary, there appears to be little virtue in attempting to manage negotiation as a goal-directed process. What then drives negotiation? Using an alternating offers mechanism, multi-issue negotiation consists of a sequence of offers where each offer is derived from the history of offers considered, including the current offer, and from information that is either part of the offers or has been derived from the negotiation context because of the offers. So the direction that a negotiation takes is determined by the agents’ responses to this accumulated wisdom. All of this has much in common with knowledge-driven processes that are the hallmark of emergent processes. A multiagent system for emergent process management is described in [Debenham, 2000].

The direction of a knowledge-driven process is determined by the accumulated wisdom of all of the players who interact with the process. For example, “Fly to Brussels and see what you can find out about XYZ” is an example of a knowledge-driven process. Although there is an implied goal in the statement quoted, that goal does not determine what should actually happen. What a person actually does in satisfaction of this process will be determined by what they already know, and by what they find out along the way. In general then knowledge driven processes can not be managed as the accumulated wisdom, which includes all relevant prior wisdom, can not be represented in non-trivial examples. The negotiation process is knowledge-driven in the sense described above. Further, the contextual investigations that are an integral part of the negotiation process can typically be conducted by managing a suite of data and text mining bots as time—and maybe cost—constrained goal-driven processes. From a process management point of view, negotiation processes are interesting in that they are knowledge-driven emergent processes that can be fully managed provided that, first, full authority to negotiate is delegated to the agent and, second, sufficient contextual information can be derived from the market data, from the sources, available on the Internet (news feeds, company white papers, specialised articles, research papers) and other sources by the data mining bots.

2. OUTLINE DESIGN OF CURIOUS NEGOTIATOR

The overall goal of this design to exploit the interplay between contextual information [Gomes and Jehiel, 2001] and the development of offers in negotiation conducted in an electronic environment. The curious negotiator is a multiagent system containing three types of agents: negotiation, mediation and observer agents. *Negotiation agents* apply the negotiation strategies in the negotiation process [Krauss, 2001]. The term *negotiation strategies* here is used in a rich sense; it includes strategies for developing the set of issues in an offer as well as identifying, requesting and evaluating contextual information including determining what information to table as the negotiation proceeds. The impartial *mediation agents* assist two or more negotiation agents. The role of *observer agents* is to observe and analyse what is happening on the ‘negotiation table’ and to look for opportunities particularly from failed negotiations. The initial design of a ‘curious negotiator’ includes two negotiation agents, one mediation agent and one observer agent, as illustrated in Fig. 2.

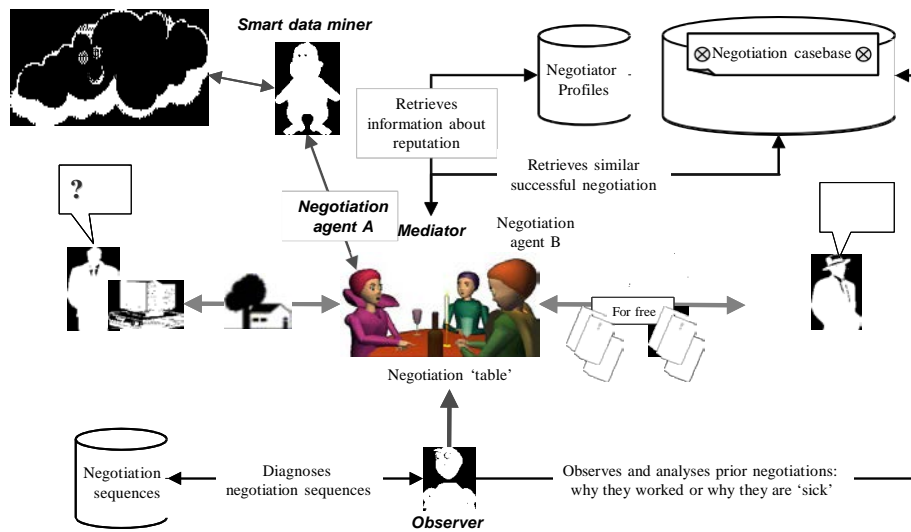


Fig.2. The initial design of the curious negotiator (includes Negotiation agent, Mediator, Observer and the Smart data miner)³

The curious negotiator is a general-purpose negotiation architecture designed for multi-issue negotiations; it proactively acquires and exploits contextual information in the negotiation process. This architecture is designed to either assist or enable two or more individuals to reach an agreement on a set of issues on which those individuals have conflicting interests. Successful negotiation relies on an understanding of how to ‘play’ the negotiation mechanism [Ströbel, 2001] and on contextual information. This contextual information is derived from what happens at the bargaining table and away from it. To do this the curious negotiator includes a *smart data mining system* [Hand et al., 2001] — the “Smart data miner” in Fig.2, which operates in tandem with the negotiation agent. The smart data miner extracts contextual information from relevant markets and from the Internet generally. The main components of the ‘curious negotiator’ are described in more detail below.

2.1 Negotiation agents

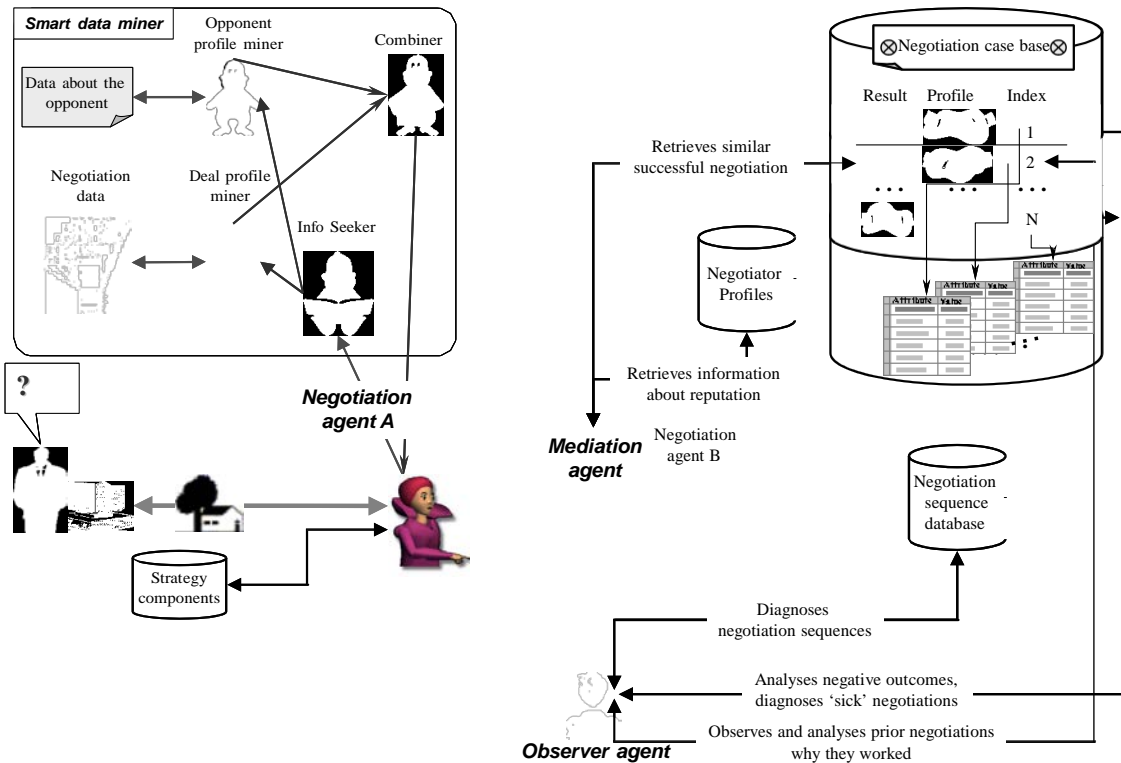
The negotiation agents are the core agents in the ‘curious negotiator’. The negotiation agents operate in two modes: as an assistant that works with an individual who makes some of the decisions, and as an autonomous agent that has delegated authority to negotiate [Wong et al., 2000]. To operate as an assistant, the negotiator will require an interface that enables high-level interaction with its user. Hence, virtual worlds are considered as a possible implementation technology⁴. If given the authority, the negotiator agent will negotiate autonomously on behalf of a player, and so will manage the gathering and verification of information using process management techniques. Each negotiation, including the information gathering, verification, combination [Parsons 2001] and distillation, is managed as a business process. This is a novel approach, especially in managing data mining bots [Lin and Pantel, 2001] under tight time and cost constraints by ensuring that the best available advice is provided when required.

The negotiation agent architecture is illustrated in Fig. 3a. The negotiation agent strategies [Fatima et al , 2001] are governed by the rules of the negotiation mechanism [Faratin, 2002]. In addition to governing the development of the offers, the negotiation mechanism governs, in part, the information exchange [Milgrom, 2002], hence negotiation strategies deal with information gathering and verification as well as with issue modification. A strategy determines: (i) the

³ Avatars, used to denote software agents in the diagram, are adapted from MIT project “BodyChat” and Bonzi Buddy (www.bonzi.com).

⁴ Current experiments are conducted in an environment based on Java3D.

modification of the existing issues, and modification of the issue set; (ii) requests for information from the opponent agent and from information gathering bots, as illustrated in Fig. 3a; (iii) verification of, or establishment of a measure of belief in, information provided. Obtaining and verifying information takes time and resources—the negotiation strategy will accommodate those delays and will manage those costs as part of the overall cost of the negotiation. To reduce some of the delays, the Smart data miner in Fig 3a can also operate in ‘pre-emptive’ mode, in which it ‘pre-fetches’ some of the information that is expected to be necessary for a scheduled negotiation.



a. Negotiation agent
b. Mediation and observer agents
Fig.3. The three types of agents (including the smart data mining system).

In real world negotiations, the corpus of contextual information is dynamically changing. There are a number of challenges that the smart data mining system needs to address, including critical pieces of information being held in different repositories; non-standard nomenclatures; radically different data types and models; possible duplicative, inconsistent and erroneous data; and possible high rate of change of the models representing data content. The mining and discovery procedures include: (i) mining the opponent’s profile information (this is a broad group of methodologies in which adapt and further develop: user-centric and site-centric data mining methods, methods for mining social networks in electronic communities for information about opponents reputation [Yu & Singh, 2000] [Marsh, 1994], text data mining methods, including discovering unexpected information about the opponent from competitors sources, methods for topic detection in communication transcripts); (ii) mining deal profiles information — these methods analyse the preconditions of negotiations, and the dynamics of change in negotiation issues; (iii) event sequence mining — will extract behaviour patterns of negotiating parties from the ‘utterances’, sequences of key events that can change negotiation (based on past experiences and current situation on the ‘negotiation table’). These information seeking procedures can reveal *unexpected information*, that is, a piece of (contextual) information that is relevant to the

negotiation process, but is unknown to the agents, or contradicts their existing beliefs (expectations). The complete integration of contextual support into the negotiation process is a novel addition to negotiation systems. To find unexpected information about specific issue and from specific source the data mining system will take in account the existing knowledge (beliefs) that the curious negotiator possesses on that issue and will try to challenge that knowledge.

2.2 Mediation agents

The role of the *mediation agents* is to assist negotiations to converge. They are independent and impartial agents. They observe all negotiations and gather information on the development of offers, particularly those offers that result in a satisfactory conclusion. That information is related to the players' profiles and is retained. Then, by establishing the type of the players and the type of the negotiation, a mediation agent suggests issues that have proven to be constructive in the past. For example, if a student is negotiating (using the Curious Negotiator) the purchase of her first computer system from a second-hand computer store, the mediator agent may be able to suggest issues that add value to the negotiation for both the student and the store. The choice of such an issue is made on the basis of the type of the two negotiation agents (for example, one could be representing a student), the subject of the negotiation (for example, a personal computer system), the negotiation history and the current negotiation state. Achievement of this functionality requires broadly-based shallow reasoning [Leake and Kolodner, 2001]. The inclusion of mediation agents in the curious negotiator resolves an intractable problem by partly removing the need for each negotiation agent to model its opponent—that could be a very costly exercise.

The mediation agent is a non-trivial intelligent system that views both sides as part of an electronic community and maintains the reputation network [Lee et al, 2001]. The mediation agent maintains a “profile book” of the negotiations that it has observed and mediated. The ‘input’ to the mediation agent is a long history of negotiations each of which has an identifiable type. Each player has attributes that enable them to be typed as well. The mediation agent then accelerates the negotiation process by suggesting new issues or combinations of issues triggered by the state of the negotiation, the type of the negotiation and the types of the players. The initial design of the curious negotiator employs a combination of case-based reasoning and ‘collaborative filtering’ for its mediation agent. The idea is presented in Fig.3b. The mediation agent operates only over the positive examples from the negotiation case base. When confronted with a situation where mediation is necessary, it retrieves a successful negotiation that is the “closest” (according to the implemented distance metric [Hilderman & Hamilton, 1999]) to the current situation on the negotiation table. The mediation agent also provides the negotiation agent with information about opponent's reputation that is based on previous experience with that negotiator. In a sense, the mediation agents use methods developed in the area of ‘collaborative filtering’, which have proved to be effective in supporting on-line shopping. The mediation agent is also a repository for anecdotal appraisals of negotiation agents by other agents. This information is freely available and is unedited. The presence of the mediation agent should at least ensure that a negotiator confronted with a particular negotiation (such as the purchase of a second-hand computer system) for the first-time does not fall into any ‘first time traps’.

2.3. Observer Agents

The *observer agents* do not contribute directly to an existing negotiation process. They observe negotiations looking for failed, or otherwise unsatisfactory, negotiations. Failed negotiations are lost business opportunities. The observer agent analyses these failed negotiations to determine why they failed and then synthesises these reasons for failure into new forms of transactions designed to prevent similar lost opportunities in the future. The introduction of an observer agent enables the curious negotiator to discover innovative, new forms of transaction. The data mining algorithms that support the observer agent are oriented towards discovery of unexpected information about the lost opportunities — information that is relevant to the negotiation process,

but is unknown to the negotiation and mediation agents, and may even contradict their beliefs (expectations). Some of the analysis techniques that are included in the initial design are: (i) knowledge-based diagnostic methods for diagnosing 'sick' negotiations; (ii) mining trust chains [Khare and Rifkin, 1997] [Castelfranchi & Tan, 2000] and relating them to the negotiating agents and the results of negotiation (we will investigate what will the length and branching of such chain tell us about the failure); and (iii) mining outliers in the behaviour of negotiating agents and mediation agent (outliers usually are thrown out, however, they can be containing the answer to what went wrong.) The observer agents address the deep issue in the curious negotiator of devising creative, innovative forms of transactions and ways of negotiating that should prevent lost opportunities. To do this they will challenge beliefs about "the way business should be done" by proposing leaps of faith [Pease et al., 2001] that may at first appear to be absolutely crazy [Bentley & Come, 2002] [Boden, 1998]!

3. CURIOUS NEGOTIATOR IN ACTION

Earlier, in section 1.4, we mentioned that negotiation is a knowledge-driven process. Debenham [2000] argues that a significant amount of support can be provided to knowledge-driven processes. Further, within any knowledge driven process there are typically goal-driven processes that can still be managed in the traditional sense even if the way in which the goals may be achieved is not known. Debenham [2000] uses a multiagent process management system based on a generic three-layer BDI agent architecture to do this. However, despite the emergent nature of negotiation processes (see section 1.4), the approach to managing emergent process described in [Debenham, 2000] does not apply directly to negotiation processes partly because negotiating agents are competitive, whereas the process agents in [Debenham, 2000] are cooperative. This is no surprise as the competitive vs. cooperative nature of agents is reflected in a clear division of work on the practice and theory of agent interaction mechanisms.

3.1. Managing negotiations Processes

Curious Negotiator deals with multi-issue negotiation with an open negotiation set using an alternative offers mechanism. After the process commences, each agent receives an offer, checks the set of issues in the offer for consistency, evaluates the offer and determines a response. All of this is done using the process knowledge and information that can be gleaned from the context and from the opponent. The construction of the response can reasonably be expected to be achieved within a certain time, and so the whole business of gathering information to determine the response is in general time constrained. This is illustrated in Fig. 4.

An offer may contain apparent inconsistencies such as "the item has a full and unconditional warranty for twelve months" and "once the item has been supplied and delivered to the purchaser the vendor is not responsible for any subsequent transportation charges". The determination of a consistent offer is concerned with the removal of such apparent inconsistencies. This is a complex problem even if the terms of the offer can be represented in Horn clause logic due to the amount of common and background knowledge required.

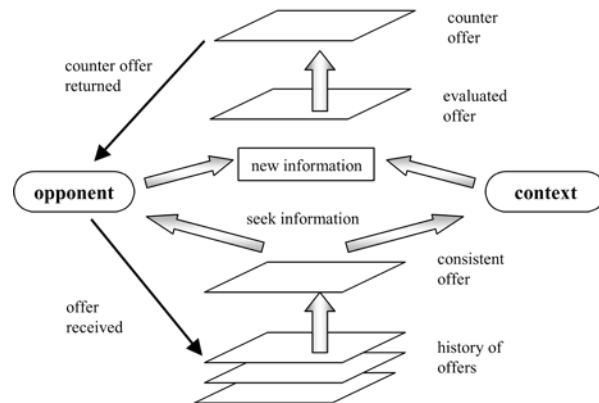


Fig.4. High-level view of the e-market and player

3.2. Curious Negotiator at the e-Market

The virtual space in which Curious Negotiator operates includes the e-market place, the World Wide Web and other Internet sources. The e-market place is designed so that it accommodates the actors described in Fig.1 (see Section 1). The e-exchange in Fig.1 is designed as a virtual space in which a variety of markets can take place. Each market is advertised on a notice board which shows the start and stop time as well as what the market is and the regulations that apply to players who wish to participate in it. The market regulations include the specification of the market mechanism [Weiss, 1998]. Designing market mechanisms is an active area of research. One important feature of a mechanism is the 'optimal' strategy that a player should use, and whether that strategy is "truth revealing" [Jennings/Wooldridge, 1998]. A human player works through a PC from which the player communicates with the negotiation agent of the Curious Negotiator, situated in the e-market.

Negotiation agent's world beliefs [Jennings/etal, 2000] are derived either from reading messages received from the player, the e-market place, or from the Smart data miner. Beliefs play two roles. First, they can be partly or wholly responsible for the agent committing to a goal, and may thus initiate an intention (eg: a plan to achieve what a message asks, such as "please buy xyz"). 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 plan. This is *reactive reasoning*. Deliberative reasoning is responsible for selecting a plan for a goal and for determining realistic constraints for each sub-goal in that plan. A plan for a goal-directed process can not necessarily be relied upon to achieve its goal even if all of the sub-goals on the chosen path through that plan have been achieved. The success condition (SC) is a procedure whose goal is to determine whether the plan's goal has been achieved. The final sub-goal on every path through a plan is a success condition. The success condition is a procedure; the execution of that procedure may succeed (\checkmark), fail (\hat{u}) or abort (A). If the execution of the success condition fails then the overall success of the plan is unknown (?). So the four possible plan exits resulting from an attempt to execute a plan are as shown in Fig. 5a. A plan body is represented as a directed AND/OR graph, or state-transition diagram, in which some of the nodes are labelled with sub-goals. The generic plan is shown in Fig. 5b.

Reactive reasoning play two roles: first, a plan is aborted if its specified abort condition is satisfied, and second, data is passed to partly executed plans for goals an agent is committed to achieve. Of these two roles the first takes precedence over the second. Reactive reasoning is achieved by rules of the form:

if <trigger state> **and** <belief state> **then** <action> **and** <trigger state>

where the <trigger state> is a device to determine whether the trigger is active or not, and <belief state> is something that may be in the agent’s world beliefs. Data is passed to partly executed plans using *procedure triggers*. For example, the sub-goal $[SG_2, c_2]$ of the plan illustrated in Fig.6a is [chair purchased, cost < 30]. This sub-goal may be achieved if “chair has been purchased for \$28” is present in the agent’s world beliefs. So until such a belief materialises this sub-goal may “hang”. This situation is managed by linking the sub-goal to a procedure trigger “**if** waiting for chair purchase **and** chair has been purchased for \$x **then** goal [chair purchased, cost < 30] is satisfied and not waiting for chair purchase”. This procedure trigger “watches” the agent’s world beliefs.

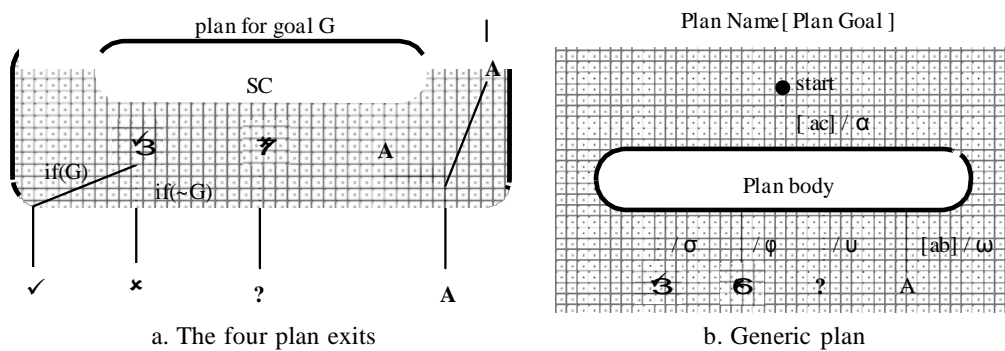




Fig.5. Plans and goals

Curious negotiator is engaged in variety of e-market transactions, spanning from simple trading orders to buy and sell in an e-exchange to multi-issue negotiations and requests for market data as well as requests for information extracted from newsfeeds and other Internet data. All e-market transactions are managed as constrained goal-driven processes. The “plan-goal” mechanism is illustrated on a simple example of a ‘buy’ trading order - “buy and chair and a desk for less than \$100”. This is represented (see Fig.6a) as a plan with goal [desk and chair have been purchased, cost < 100]. This plan has sub-goal $SG_1 =$ ‘prices of chairs, desks and delivery are xyz’, $[SG_2, c_2] =$ [chair purchased, cost < 30], $[SG_3, c_3] =$ [desk purchased, cost < 50], and $[SG_4, c_4] =$ [desk and chair delivered, cost < 20]. This purchase order is represented as a plan whose goals have constraints.

An example of a request for information is “find out all you can about ABC Corp within five minutes”. This triggers a process to locate, extract, validate, condense and combine information from the Internet. All, except combining information, is achieved by data/text mining bots that are described in [Debenham & Simoff, 2001]—they are not considered here. Combining information is achieved by a Bayesian network. The data/text mining bots produce output in the form [data, belief]—i.e. some data and a measure of the belief held in the validity of that data. This request for information is first represented as a goal/constraint pair: [find_info_about (‘ABC Corp’): time_upper_limit = now + 5mins]. Given a goal/constraint pair a plan (see Fig.6a) is selected for it—the mechanism for selecting a plan is described in [Debenham, 2002]. A *plan* for a goal/constraint pair is a possibly-conditional state-chart of other goal/constraint pairs. For a ‘find_info_about’ process, the plan uses a Prospector-style Bayesian network (see Fig.6b) to combine results $[D_i, b_i]$, in the form [data, belief], extracted from the Internet by a suite of data/text mining bots. The Bayesian network actually does more than combine information. If the level of belief, b_r , in a result, R, derived by the network is below a set threshold then a ‘reverse calculation’ identifies ‘inputs’ whose belief levels are responsible for the low level of belief in R. Then further data/text mining is initiated in an attempt to raise this level of belief.



a. A plan for goal [G, c]



b. Combining information with Bayesian network
Fig.6.

4. CONCLUSION

The curious negotiator, of which the negotiation agent is a key component, is our “grand vision” for automated negotiation systems. It will blend ‘strategic negotiation sense’ with ‘strategic information sense’ as the negotiation unfolds. This will require a system that can readily adapt to changes in its environment and so agent architecture is indicated. This multi-agent architecture will continually revise its actions in the light of possibly conflicting signals from the negotiation and its context. This will involve the construction of subtle plans and reactive apparatus within the agent as illustrated in Section 3. An intractable part of the negotiation agent’s architecture is the interface between the contextual information and the negotiation strategy. To resolve this intractability the curious negotiator will be trialled in an area in electronic business.

The smart data mining systems that, as part of the Curious Negotiator, support the negotiation agent, are expected to operate under time-constraints and over dynamically changing corpus of information. They will need to determine the sources of information, the confidence and validity of these sources and a way of combining extracted information (models). Some of the information can be pre-fetched, before actual negotiation. All transactions, including complex requests for information and combination of results, are managed as business processes. Overall Curious Negotiator is a novel negotiation system, which will assist in understanding the interlay between negotiation strategies and the context in which negotiation takes place.

REFERENCES

- [Ansari et al., 2001] Ansari, S., Kohavi, R., Mason, L. and Zheng, Z. Integrating E-Commerce and Data Mining: Architecture and Challenges. Proceedings IEEE International Conference on Data Mining, ICDM'01, 2001.
- [Benn, et al., 1999] Wolfgang Benn, Otmar Görlitz and Ralf Neubert. An Adaptive Software Agent for Automated Integrative Negotiations, Int. J. of e-Business Strategy Management, 1 (2), 1999
- [Bentley & Corne, 2002] Peter J. Bentley and David W. Corne. An Introduction to Creative Evolutionary Systems, In: Peter J. Bentley and David W. Corne (eds). Creative Evolutionary Systems. Academic Press, London, UK, 2002, pp. 1-78.
- [Bichler, 2001] Martin Bichler. The Future of E-Markets: Multi Dimensional Market Mechanisms. Cambridge University Press, 2001.
- [Bjork, 1999] Bjork, Tomas. Arbitrage Theory in Continuous Time. Oxford University Press, 1999.
- [Boden, 1998] Margaret A. Boden. Creativity and Artificial Intelligence. Artificial Intelligence 103 (1-2): 347-356, 1998.
- [Castelfranchi & Tan, 2000] Castelfranchi, C. and Tan, Y. H. (eds). Trust and Deception in Virtual Societies. Kluwer Academic Publishers, 2000

- [Debenham, 2000] Debenham, J. K. Supporting knowledge-driven processes in a multiagent process management system. In proceedings 20th 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, 2002] Debenham, J. K. (2002). E-Market negotiation managed in the context with a multi-agent system. Proceedings 22th International Conference on Knowledge Based Systems and Applied Artificial Intelligence, ES'2002.
- [Debenham & Simoff, 2001] Debenham, JK and Simoff, S (2001). Investigating the Evolution of Electronic Markets. In proceedings Sixth International Conference on Cooperative Information Systems, CoopIS 2001, Trento, Italy, September 5-7, 2001, pp344-355.
- [Faratin, 2002] Peyman Faratin. Multi-Agent Contract Negotiation. In: Socially Intelligent Agents - creating relationships with computers and robots (eds. K. Dautenhahn, A. Bond, L. Canamero, B. Edmonds), Kluwer 2002.
- [Fatima et al , 2001] S.S. Fatima, M. Wooldridge and N. R. Jennings. Optimal negotiation strategies for agents with incomplete information. In: Proc. 8th Int. Workshop on Agent Theories, Architectures and Languages (ATAL), Seattle, USA, 2001.
- [Gerding et al, 2000] Gerding, E.H., van Bragt, D.D.B. and La Poutre, J.A. Multi-issue negotiation processes by evolutionary simulation: validation and social extensions. In proceedings Workshop on Complex Behavior in Economics. Aix-en-Provence, France, May 4-6, 2000.
- [Gomes and Jehiel, 2001] Gomes, A. and Jehiel, P. Dynamic process of social and economic interactions: on the persistence of inefficiencies. Centre for Economic Policy Research, CEPR, London, 2001.
- [Hand et al., 2001] David Hand, Heikki Mannila and Padhraic Smyth. Principles of Data Mining. The MIT Press, Cambridge, MA, 2001.
- [Hilderman & Hamilton, 1999] R.J. Hilderman and H.J. Hamilton. Knowledge discovery and interestingness measures: A survey. Technical Report CS 99-04, Department of Computer Science, University of Regina, October 1999.
- [Khare and Rifkin, 1997] Rohit Khare and Adam Rifkin. Weaving a Web of Trust. World Wide Web Journal, 2 (3), 77-112.
- [Klusch, 2001] Klusch, M. Agent-Mediated Trading: Intelligent Agents and e-Business. In: Hayzelden, A. L. G. and Bourne, R. A. Agent Technology for Communication Infrastructures, John Wiley and Sons, Chichester, 2001, pp. 59-76.
- [Krauss, 2001] Sarit Kraus. Strategic Negotiation in Multiagent Environments. MIT Press, 2001.
- [Leake and Kolodner, 2001] Leake, D. and Kolodner, J. Learning through case analysis. In: Encyclopedia of Cognitive Science. Macmillan, London 2001
- [Lee et al, 2001] Lee, M., Alem, L. and Kowalczyk, R. Supporting Electronic Negotiation for Intelligent Trading. In A. Gangopadhyay (Ed.) Managing Business with Electronic Commerce: Issues and Trends. Idea Group Publishing, 2001.
- [Lin and Pantel, 2001] Dekang Lin and Patrick Pantel. DIRT – Discovery of Inference Rules from Text. In Proc. 7th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining KDD 2001, pp. 323-328.
- [Maes et al., 1999] Maes, P., Guttman, R.H. and Moukas, A. G. Agents That Buy and Sell. Communications of ACM, 42 (3), 1999, 81-91.
- [Marsh, 1994] P. Steven Marsh. Formalising Trust as a Computational Concept. Ph.D. Dissertation. University of Stirling, Department of Computing Science and Mathematics, 1994.
- [Milgrom, 2002] Milgrom, P. Auction Theory for Privatization. Cambridge UP, 2002.
- [Milgrom, P. and Weber, R. 1982] Milgrom, P. and Robert Weber, R. A Theory of Auctions with Competitive Bidding. Econometrica 50(5) 1982.
- [Muthoo, 1999] Muthoo, A. Bargaining Theory with Applications. Cambridge UP, 1999.

- [Osborne and Rubinstein, 1990] Martin J. Osborne and Ariel Rubinstein. *Bargaining and Markets*. Academic Press, 1990.
- [Parsons 2001] Simon Parsons, *Qualitative Methods for reasoning under uncertainty*, MIT Press, Cambridge, MA.
- [Pease et al., 2001] Alison Pease, Daniel Winterstein and Simon Colton. Evaluating Machine Creativity. In: Proc. Workshop on Creative Systems: Approaches to Creativity in AI and Cognitive Science – Int. Conf. on Case-Based Reasoning ICCBR-2001, Vancouver, Canada, 2001, Washington, DC: Naval Research Laboratory.
- [Saunders, 2001] Rob Saunders. *Curious Design Agents and Artificial Creativity: A Synthetic Approach to the Study of Creative Behaviour*. Ph.D. Dissertation. University of Sydney, Department of Architectural and Design Science, 2001.
- [Scott and Markovitch, 1989], Paul D. Scott and Shaul Markovitch (1989). Learning Novel Domains Through Curiosity and Conjecture. International Joint Conference on Artificial Intelligence, Vol. 1, pp. 669 - 674, Palo Alto, CA: Morgan Kaufmann.]
- [Ströbel, 2001] Ströbel, M. Design of Roles and Protocols for Electronic Negotiations. *Electronic Commerce Research Journal*, Special Issue on Market Design, 2001.
- [Tennenholtz, 1999] 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.
- [Watkins, 2002] Watkins, M. *Breakthrough Business Negotiation—A Toolbox for Managers*. Jossey-Bass, 2002.
- [Wise & Morrison, 2000] R. Wise & D. Morrison. Beyond the Exchange; The Future of B2B. *Harvard Business review* Nov-Dec 2000, pp86-96.
- [Wong et al., 2000] Wai Yat Wong, Dong Mei Zhang and Mustapha Kara-Ali. Towards an Experience Based Negotiation Agent. In: *Cooperative Information Agents IV: The Future of Information Agents in Cyberspace* (eds M. Klusch and L. Kerschberg), Springer, 2000, pp. 131-142
- [Yu & Singh, 2000] Bin Yu and Munindar P. Singh, A Social Mechanism of Reputation Management in Electronic Communities. In: *Cooperative Information Agents IV: The Future of Information Agents in Cyberspace* (eds M. Klusch and L. Kerschberg), Springer, 2000, pp. 154-165

- [Kenney & Raiffa, 1993] R.L. Kenney and H. Raiffa. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. Cambridge University Press. 1993.
- [Faratin, 2000] Peyman Faratin. Automated Service Negotiation Between Autonomous Computational Agents. Ph.D. Dissertation. University of London, Queen Mary College, Department of Electronic Engineering, 2000.
- [Guttman et al, 1998] R. Guttman, A. Moukas, and P. Maes. Agent-mediated Electronic Commerce: A Survey. Knowledge Engineering Review, June 1998.
- [Padmanabhan & Tuzhilin, 1998] Padmanabhan, B. and Tuzhilin, A. A Belief-Driven Method for Discovering Unexpected Patterns. In: Proc. 4th ACM SIGKDD Conf. On Knowledge Discovery and Data Mining, KDD-98, 1998, pp. 27-31.
- [Rea & Skevington, 1998] Tim Rea and Peter Skevington. Engendering trust in electronic commerce. British Telecommunications Engineering, 17 (3), 150-157, 1998.
- [Bentley & Corne, 2002a] Peter J. Bentley and David W. Corne (eds). Creative Evolutionary Systems. Academic Press, London, UK, 2002.
- [Bentley, 1999] Peter J. Bentley (ed.) Evolutionary design by computers. Morgan Kaufmann, 1999.
- [Foner, 1999] Lenny Foner. Political Artifacts and Personal Privacy: The Yenta Multi-Agent Distributed Matchmaking System. Ph.D. Dissertation. Massachusetts Institute of Technology, School of Architecture and Planning, 1999
- [Franz, et al., 2001] Martin Franz, Abraham Ittycheriah, J. Scott McCarley, Todd Ward. First Story Detection: Combining Similarity and Novelty Based Approaches (available from http://www.nist.gov/speech/tests/tdt/tdt2001/PaperPres/ibm-pres/tdt2001_nn.ppt)
- [Plott, 2002] Plott, C. Handbook of Results in Experimental Economics. Elsevier, 2002.
- [Raiffa, 1982] Howard Raiffa. The Art and Science of Negotiation. Harvard, 1982.
- [Rosenschein & Zlotkin, 1998] Jeffrey S. Rosenschein and Gilad Zlotkin. Rules of Encounter. MIT Press, 1998.
- [TDT, 2001] Topic Detection and Tracking (<http://www.nist.gov/speech/tests/tdt/index.htm>)