AISB'02
Convention

Proceedings of the AISB'02
Symposium on

Intelligent Agents in Virtual Markets

3rd - 5th April 2002

Imperial College of Science,
Technology and Medicine
University of London

Published by

The Society for the Study of Artificial Intelligence and Simulation of Behaviour

http://www.aisb.org.uk/

ISBN 1 902956 29 9

A generic model for contract negotiation

Philippe Mathieu , Marie-Hélène Verrons
Équipe SMAC
Laboratoire d'Informatique Fondamentale Lille – CNRS upresa 8022
Université des Sciences et Technologies de Lille
59 650 VILLENEUVE D'ASCQ, FRANCE
mathieu@lifl.fr , verrons@lifl.fr

Abstract

During an interaction between several software entities trying to solve a problem, two kinds of attitudes can occur : *cooperation* if all entities have the same goal, *negotiation* if self-interests are in conflict. Cooperation has largely been studied, whereas the later has not really been computationally studied. In this paper, we define a generic agent state model and a generic protocol for contract negotiations, called ANTS, and we propose a three-level methodology to build generic negotiation applications based on a communication protocol, a generic internal structuration for problem representation and a strategy level to obtain better rationality. Before comparing our approach with other researches on the same topic, we give an example of application of this work to solve a concrete virtual market problem and another exemple of application for appointment taking. The aim of this paper is to show that it is possible to describe precisely a generic model that we can use in several negotiation and marketplace problems. These researches on negotiation take place in software engineering works for artificial intelligence and multi-agent systems.

1 Introduction

For a couple of years, multi-agent systems (MAS) have been studied from different points of view. Whether for Distributed Artificial Intelligence (DAI) (CL90; MCL88), for distributed algorithmic or even for robotics (Bro86), communication and cooperation between agents have been the main points studied. If many multi-agent environments aim to make agents cooperate to realise a common goal Visiter Hoster (Syk89), other systems use agents which have different interests (Kle91; MCL88; RZ96) as in Baazar (Syk89) and Kasbah (CM96) systems, which needs negotiation between them. After the fashion of negotiation between customers and providers who use negotiation rules to go towards a median position, competitive software agents have to negotiate with their environment to lead to a position agreed by everyone. The realisation of a society of agents working on this objective leads to many problems. These problems occur when an agent must choose between several competitive negotiations and/or when these negotiations lead to peals of renegotiations. As agents have a partial knowledge of the environment and as this environment continuously evolves, these problems are even more difficult. Moreover, the utility of using an agent during negotiations is perfectly justified by the explosion of the number of messages exchanged between agents due to the cascaded process. In certain cases, the number of messages can be in $O(m^n)$ if n is the depth of the cascaded process and m the number of agents involved in one negotiation.

As part of works on negotiation, we will base our arguments on the following consensual definition, which can be applied to many fields such as auctions, appointment taking systems, and even in certain multi-agent games.

Definition 1 *Negotiation is carried out on a contract to obtain common resources and on the request of an agent called the initiator. It brings together an initiator and a set of agents to whom the contract is proposed by the initiator agent, these agents are called the participants, and runs until an agreement satisfying a percentage of participants is reached. Participants equally try to obtain the best possible solution for themselves while giving a minimum set of information to the others.*

To realise a negotiation application, three well identified phases need to be completed. **The internal representation** which contains the management of data structures and speech acts necessary for agents to evolve their knowledge; **the communication layer** allowing agents to send messages in a centralised way if agents are on the same machine, or in a distributed way if they are on different machines; **the strategic level** allowing agents to reason on the problem and infer on the knowledge obtained from the others.

In this paper, we will realise agent communication via the Magique API (BM97), which allows the system to run on an heterogeneous and distributed network. The success of a negotiation depends of course on adapted strategies. We will not discuss here about strategies, which are different according to the kind of negotiation done. This

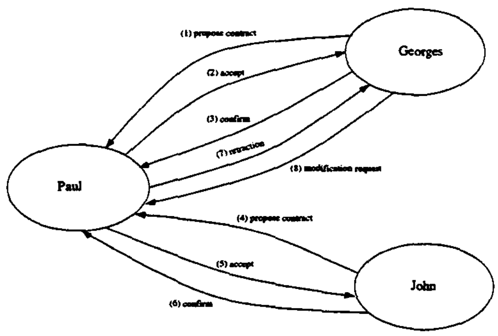

Figure 2: Communication between agents

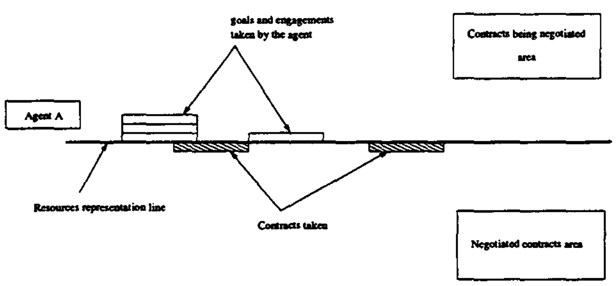

Figure 3: Resources representation

the contract (*confirm* message). Then Paul takes John's contract and cancels Georges's one (*retraction* message). When Georges receives this message, he decides to ask Paul a modification for the contract (*modification request* message).

2.3 Internal objects

Any realisation in computer science needs an adaptive data structure. Here, there is a need to code simultaneous negotiations, negotiation progressing, initiators, participants and of course resources and contracts.

To implement our protocol, we define several objects common to any form of negotiation : resources and their representation, contracts and their properties, goals and engagements.

Resources Resources are abstract and can represent as well articles as hours or fruit boxes. Resources are the objects that each agent is trying to obtain. Any object can be a resource, including a contract as a datatype, but not the contract being negotiated itself.

Contracts and their properties

Definition 2 *A contract is an object created at the initiation of a negotiation. It contains resources we want to obtain, the initiator of the negotiation and a limit date to answer to avoid deadlocks.*

A contract over several resources means that you want to obtain the whole resources as a lot, ie. the resources are all needed for the contract to be confirmed, if you can't have them all together the negotiation will fail, whereas several contracts on one resource means that you want the resources separately, ie. if you don't succeed in having some of them it doesn't mind, all negotiations are independent, the failure of one has no importance on the success or failure of the others.

The contract does not contain the set of participants in order to protect their "private life", so a participant doesn't know if there are other participants nor their answer to the proposition or the modification they propose.

Definition 3 *The properties of the contract contain the contract itself, the set of participants, the answer delay, the default answer, the minimal number of agreements to confirm the contract and the maximum number of authorised renegotiations. It also keeps in memory the set of participants who have accepted the contract, the number of received agreements, and the number of renegotiations done.*

They are basic "memory" objects which contain properties common to any type of negotiation.

Goals A *goal* is an object created at the initiation of a contract. In fact, it is described by the initiator agent automata on Figure 1. The initiator agent creates the contract and then a goal to which he gives the contract and relevant information. Once the goal created, the whole managing process of the contract is done by it.

Engagements When an agent receives a contract proposition, he creates an *engagement* and gives it the contract. It is then the engagement which manages the contract negotiation. It is described by the participant agent automata on Figure 1.

Graphical representation of negotiation To represent resources, we define a graphic model as shown in Figure 3. A graduated line is used to represent available resources, each resource corresponds to a segment. Below the line, we find contracts already taken by the agent, the corresponding resources being no longer free. Above the line, we find goals and engagements which are currently in negotiation on corresponding resources. This upper part of the representation describes a matrix where columns corresponds to resources and where goals and engagements fall in a first-in first-out list. In fact, this matrix behaves like Tetris game (Figure 4) : arriving contracts are put one onto another in conflicting cases. The active goal/engagement is the one placed at the bottom line of the matrix, the other ones are waiting. When a negotiation finishes, the corresponding goal or engagement is removed from the matrix, and the one following it is waked up, that is to say the corresponding negotiation can begin.

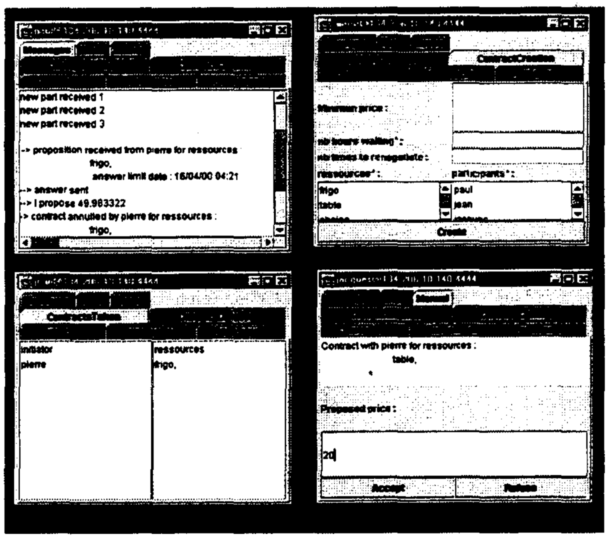

Figure 6: Four agents in an Auction application

3.1.2 From the participant point of view : engagement and associated strategy

When a participant receives an auction *proposition*, he first checks if the proposed article interests him. If yes, he *agrees* with the contract and proposes a price. Otherwise, he *refuses* it.

When a confirmation of auction is received, the participant adds the article in his bag and virtually pays the seller.

When an *ask for modification* is received, the participant checks his/her amount of money and sends a new price greater than the previous one to the initiator if he has enough money or a price equals to 0 if he doesn't want to carry on the auction.

To compute the price to propose in the case of a modification request, the engagement follows the strategy specified by the user, namely a linear, a quadratic , ... method.

Figure 6 shows interface of 4 agents negotiating auctions with our API.

The top left screen is an agent showing his window for visualising messages sent and received by him. It permits to see different propositions received and advancement of negotiation (answer sent, confirm, cancel, ask for modification, ...). The top right screen is an agent showing the new contract input interface, the bottom left one displays contracts taken with the name of the initiator and the negotiated resources. The last one shows the display of a proposition of contract for manual mode.

The advantages of this application are numerous, so we will cite here the most important ones. First, this application helps the user in auctionning and auctions for him when he is not there, according to the strategy he defined. Secondly, this application can easily be extended to another type of auctions, such as English, Dutch, Vickrey ... auctions. And thirdly, this application is portable, as a matter of fact, the agents can be placed on an heterogeneous network, on PDAs, ...

3.2 An appointment taking application

This is a negotiation application for appointment taking where each agent must be able to negotiate rendezvous for the user. Each user defines a diary with time slots free or not. In addition, he gives preferences on slots and on persons with who he prefers to take appointments. Each user can initiate an ask for a rendezvous with one or more participants on one or more time slots.

This problem is a full featured one because it needs preferences over persons, for example, my boss has a greater priority than my colleague, but also priorities over resources (here time slots), eg. I don't want to have appointments at lunch time or before 8 am, the corresponding time slots will then have a lower priority. Moreover, appointment taking is an application where there are typically many renegotiations and retractions, because it is difficult to find time slots that fit everyone.

3.2.1 From the initiator point of view : goal and associated strategy

The first thing a goal do is sending the proposition to the set of participants. Simultaneously, he starts a timer (associated to a default answer, here a rejection) in order not to wait infinitely for each participant's answer.

When an agreement is received, the number of received agreements is incremented. Once all answers have been received, the goal decides whether to *confirm* the appointment if the minimum number of agreements is reached, to *request a modification* if it is not the case and if the maximum number of renegotiations isn't reached too, in other case, it *cancels* the contract.

Here, a modification is a proposition of new time slots which better fits the participant.

When a *contract modification* is received, the initiator computes a note to the proposed time slots according to the priority of the participant who proposed them and the priority of these time slots for the participant.

Each time slot is noted in the following way :

$$note(h_i) = priority(h_i, init) * coef_{init} + \sum_{j=1}^{n} priority(h_i, part_j) * priority(part_j) * coef_{part}$$

Where $priority(h_i, init/part_j)$ stands for the priority of time slot i for the initiator/participant (if it is not in the received list, it is given a discriminatory note), $coef_{init/part}$ is the multiplicative coefficient of the initiator/participant, fine his own purchase strategy. The architecture of AuctionBot is asynchronous, it stores auctions in a database and can manage several auctions simultaneously. It is AuctionBot which manages auctions according to a set of auction rules, not the initiator of the contract.

Adept (AGJ+96) realised at the Queen Mary and Westfield College in collaboration with British Telecom in order to develop a service-oriented negotiation model (SFJ97) to evaluate the cost of their settings. In Adept, agents use a bilateral model of negotiation which allow them to formulate offers, to evaluate them and to formulate counter-propositions.

ZEUS (NDNC) is a generic Java API realised by British Telecom in order to easily conceive cost-based negotiation applications between autonomous agents. Zeus proposes a negotiation protocol between two agents (an initiator and a participant) and on a single resource per contract. The protocol consists of a call-for-bids, and no mecanism of counter-proposition is provided. Moreover, it is possible to negotiate simultaneously different contracts on the same resource, that we don't allow. Another difference with our protocol is that retraction is not possible with Zeus. Once a contract is taken you can't retract yourself. Moreover, Zeus provides only cost-based strategies, and so is less generic than our protocol which is not dedicated to cost-based contracts.

These previous works, like our, are based on the general **Contract Net Protocol** model (Smi80) which works on bids invitation between a Manager agent and Contractor agents. From all these works, Magnet is probably the one which is closest to what we present. Nevertheless, none of them takes into account at the same time generic aspects, automatic renegotiations and a mechanism to manage conflicts between simultaneous negotiations, that we propose in ANTS.

5 Conclusion

In this paper, we have defined a generic agent state model and a generic protocol for contract negotiations, called ANTS. This is a potentially useful aspect of what happens in a component of a virtual market. We have also presented a Java API suitable for different kinds of negotiation such as auctions or appointment taking, independent of systems thanks to Java. It allows negotiation from 1 to many agents, simultaneous negotiation of several contracts, and the management of deadlocks in conversation. Our API is based on a protocol and a specific internal representation. Three distinct levels were defined : the knowledge representation level allowing the agent viewing the advancement of his/her negotiations, the communication level which we realised with a multi-agent platform allowing physical distribution, and the strategic level for which we propose generic strategies adaptable to any kind of problem. Each level can be easily extended by the developer as he wants to map with his application. These works are a part of software engineering and distributed artificial intelligence works with the same objective than British Telecom's Zeus platform. Many implementation perspectives of these works on different software supports are possible (distributed, centralised, WEB) and strategic level enhancement for different specific problems is considered. This API will now be applied to different problems like distance teaching, network games, workflow systems.

References

[AGJ+96] J.L. Alty, D. Griffiths, N.R. Jennings, E.H. Mamdani, A. Struthers, and M.E. Wiegand. ADEPT - Advanced Decision Environment for Process Tasks : Overview and Architecture, 1996.

[BM97] Nourredine Bensaid and Philippe Mathieu. A Hybrid and Hierarchical Multi-Agent Architecture Model. In *Proceedings of the Second International Conference and Exhibition on the Practical Application of Intelligent Agents and Multi-Agent Technology*, pages 145–155, London-United-Kingdom, 21st-23rd April 1997. PAAM'97. The Practical Application Company Ltd.

[Bro86] R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, 1986.

[CL90] P.R. Cohen and H.J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[CM96] Anthony Chavez and Pattie Maes. Kasbah : An Agent Marketplace for Buying and Selling Goods, 1996.

[CYJ+98] J. Collins, B. Youngdahl, S. Jamison, B. Mobasher, and M. Gini. A Market Architecture for Multi-Agent Contracting. In *2nd Int'l Conf on Autonomous Agents*, pages 285–292, Minneapolis, May 1998.

[Kle91] M. Klein. Supporting conflict resolution in cooperative design systems. *IEEE Transactions on System, Man, Cybernetics*, 21(5):1379–1390, 1991.

[MCL88] R.A. Meyer, S.E. Conry, and V.R. Lesser. Multistage negotiation in distributed planning. In A. Bond and L. Gasser, editors, *Reading In Distributed Artificial Intelligence*, pages 367–386, Los Altos, CA, 1988. Morgan Kaufmann Publishers.

Safety and Liveness of Negotiation Protocols

Shamimabi Paurobally and Jim Cunningham
Department of Computing, Imperial College, London
{sp1, rjc}@doc.ic.ac.uk

Abstract

Automated negotiation is facilitated by sharable, verified and unambiguous protocols with desirable properties. A negotiation protocol may be expressed as a theory in our multi-modal meta-language, ANML, thereby enabling the proof of its properties and its correctness. This paper focuses on specifying and proving safety and liveness properties of negotiation protocols. A safe and sound interaction allows no unpredictable states and transitions, and negotiation states that become valid are those that are defined by the protocol and none other. We specify three properties - liveness, termination and soundness - for a protocol in ANML. As an example, we examine a negotiation between two agents following a bilateral protocol. We prove that the latter protocol is sound, terminates and exhibits liveness.
Keywords: negotiation, protocol, ANML, soundness, termination, liveness.

1 Introduction

Negotiation is an important aspect of commerce, (Cunningham et al. 1998). It is a form of decision-making where two or more parties jointly search a space of possible solutions with the goal of reaching a consensus, (Rosenschein and Zlotkin, 1998). Negotiation is part of those wider dynamic processes whereby commercial goals are achieved by the parties to a contract. Usually in a real world negotiation, all the parties have to be physically present during the process, which may extend over a long period of time, leading to increased transaction costs. In online negotiation, participants bargain for self and mutual gain through electronic commerce frameworks. Online auctions such as Ebay and Amazon remove the need for all participants to be co-located but still require users to manage their auctions. Automated negotiation has the potential to save both time and computational costs where agents may be able to find better deals in strategically complex environments without the drawbacks of human ego and prejudices. Software agents can participate in parallel and complex transactions and monitor these over long time periods, which would otherwise be difficult for a single human.

A key aspect in both real and automated negotiation is the rich, sometimes implicit, interaction between participants, using languages and protocols. Setting up a negotiation involves deciding on a common language and ontology and choosing a common protocol. There have been a number of specifications for a standardised agent communication language e.g. KQML, (Finin et al. 1993), and FIPA ACL, (FIPA, 2001). Having a standardised electronic language removes the need for agents to find a common language for interaction but there have been some criticisms about the lack of formalisation and ambiguous nature of existing agent communication languages, (Pitt and Mamdani, 1999).

Agents often do more than send single messages during an interaction. Rather they engage in task-oriented and shared sequences of messages i.e. a protocol to which they jointly comply for accomplishing tasks. There is currently a need for tools for conversation specification, verification and sharing. Our research aims to fulfil this need for automating negotiation using AI techniques. In a previous paper, (Paurobally and Cunningham, 2000), we presented a first version of our meta-language, ANML. The latter can be used to build and verify protocols and to provide an abstract theory of a conversation to coordinate agents in their goal seeking activities. ANML is based on propositional dynamic and multi-modal logic so as to provide elegant theories for negotiation. An ANML protocol specifies the structure of and sequence in a conversation in a given language, while exchanged messages can be in agent communication languages like FIPA ACL or KQML. As an analogy to real life conversations, ANML may be considered as a tool to construct and check the grammar in a conversation where the words being communicated are in English (or ACL).

A protocol may be expressed as a logical theory in ANML thereby enabling the proof of its properties and its correctness. In this paper, we briefly describe our meta-language as a precursor for showing how protocols represented in it can be proved to exhibit desirable properties. We focus on safety and liveness properties by defining axioms for the liveness, termination and soundness of negotiation protocols. A liveness property asserts that a process eventually enters a desirable state and that there are no deadlocks. There is a relation between liveness, termination, soundness and safety properties. A safety property asserts that nothing bad happens and termination asserts that something happens so both can be used to given by its protocol. Given the theory of a protocol, we can logically derive the possible states and paths.

Let ε denote the set of all states S given by the function *States*(ϕ, λ) and let ψ represent the set of all processes α given by *Proc*(ϕ, λ, ω). Let μ represent the ANML theory of a protocol.

$$\text{state-set} = \{s, s_1 : \varepsilon \mid \mu \models ((s \vee (s_1 \rightarrow s) \vee (s \rightarrow s_1)) \bullet s\}$$
$$\text{transition-set} = \{s, s_1 : \textit{state-set},\ t : \omega \mid \mu \models (s \rightarrow [t]\ s_1) \bullet t\}$$
$$\text{path-set} = \{s, s_1 : \textit{state-set},\ p : \psi \mid \mu \models (s \rightarrow [p]\ s_1) \bullet p\}$$
$$\text{sub-path}(p_1, p_2) \leftrightarrow (\exists\ p_3 \in \text{path-set}\ (p_1 ; p_3 = p_2))$$

state-set is the set of states which are given in protocol μ, containing parent or sub-states. *transition-set* is the set of all atomic transitions in the theory μ of a protocol. A process is a path containing allowed transitions. *path-set* denotes the set of possible paths in μ. It is useful to define the function *sub-path* which returns *true* if p_1 is a sub-path of p_2.

2.2 Safety and Liveness Properties

Safety and liveness properties were first introduced by Lamport (1980) and have since then been extensively studied in the verification of concurrent programs. A safety property asserts that something bad never happens i.e. the program never enters an unacceptable state. A liveness property expresses that something good will eventually happen i.e. the program eventually enters a desirable state. Other properties may be classified into safety and liveness properties, which allows for choosing a suitable proof method to show that a program exhibit that property. Different techniques are used to prove safety and liveness. Owicki and Lamport, (Owicki and Lamport, 1982), use proof lattices or well-founded induction to prove liveness properties. Methods based on global invariants have been used to prove safety properties, (Manna and Pnueli, 1995). Temporal logic has been proposed by Pnueli, (Manna and Pnueli, 1995; Pnueli, 1977) for the specification and verification of concurrent programs since it provides a convenient language for expressing properties of reactive systems. This has been followed by several versions of temporal logics, (Burstal, 1974; Lichtenstein et al. 1985; Manna and Sipma, 1998; Sistla and Clarke, 1985) in the field.

Safety properties include partial correctness, mutual exclusion and absence of deadlock. Partial correctness states that from a true precondition a process never terminates with a false postcondition, (Owicki and Lamport, 1982). Partial correctness may be the only safety property that is required in sequential programs. Mutual exclusion means that two different processes are never both in their critical section where only one process can be. A liveness property includes program termination which is usually proved through induction. Liveness may also mean absence of livelock or the existence of cyclically recurring states. There are a number of verification techniques for proving properties e.g. deductive verification, (Gordon and Melham, 1993), verification diagrams, (Bjorner et al. 1999), model checking and abstraction, (Bjorner et al. 1999; Shostak, 1984).

2.3 Liveness Property

A liveness property asserts that something good eventually happens, that is there is no deadlock. Safety properties describe allowed behaviour while liveness properties describe required behaviour. A liveness property and termination ensure the absence of deadlock and livelock. We define the behaviour of an execution and analyse the possible sequence of actions to prove liveness and the absence of deadlock. In a protocol, an agent either starts the negotiation process or responds to a previous action, both of which may possibly change the current state. In case of deadlocks, ANML easily allows addition of the occurrence of internal events such as *timeout* or *overdue*.

Liveness axiom: $\mu \models \forall\ s_i : \textit{state-set},\ a_i : \textit{path-set}\ (s_0 \leftrightarrow [X_1.a_1]s_1 \vee \ldots \vee [X_n.a_n]s_n) \rightarrow \neg\exists\ b : \textit{path-set}\ (((Y.b?; X_i.a_i) \wedge (X_i.a_i?; Y.b)) \vee [\textit{timeout}]Z)$
where $0 \leq i \leq n$, Y and $X_i \in$ *A_group*

A liveness property asserts that it is not the case that an agent (or a group of agents) X waits for another agent (or group of agents) Y to proceed while Y is waiting for X. Otherwise there is deadlock and none of the agents are able to proceed in their execution since they are waiting for each other. To do a_i, X_i must wait for Y to do b and for Y to do so, it must in turn wait for X_i to perform a_i. The only option would be termination of the process through a timeout event. A protocol μ satisfies the liveness property if the *liveness axiom* applies for all possible states, *St.*, and paths *Pt.* in that protocol where *St.* $\in$ *state-set*, *Pt.* $\in$ *path-set*.

2.4 Termination Property

Termination analysis, (Codish and Taboc, 1999), attempts to determine whether evaluation of a given expression will definitely terminate. The evaluations of a constant or of a non-recursive function terminate, in the case where their sub-procedures terminate. A recursive function can be shown to terminate if it can be proved that there is a base case that will inevitably be executed. Termination analysis gives either a program "definitely terminates" or "don't know". There has been increasing work in the automatic termination analysis of logic programs, such as transforming a consistent logic program into a term rewriting system and proving termination for the latter, (Ganzinger and Waldmann, 1993; Ohlebusch 1999). In our case, a protocol shows termination if it can be proved that a negotiation process complying to that protocol eventually reaches a terminal state. To do so, we prove that not only from all states in a protocol there is a path leading to a terminal state without getting stuck in an infinite communication loop, but an offer only when Y has enabled the *requested(Y)* state through a *request* or an *initial_request.* Let us only consider *Y.request* for simplicity and since both cases are relatively similar. This yields in s_0 the condition

(offered(X)?Y.agree)∧(requested(Y)?X.offer) i.e. *(X.offer?;Y.agree)∧(Y.request?;X.offer).* This condition does not unify with the condition for deadlock, *(X.action2?; Y.action1)* ∧ *(Y.action1?; X.action2).* We do not have a deadlock since although *X* has to make an offer before *Y* can agree, *X* is not waiting for *Y* to agree before *X* can make its offer. In fact *X* is waiting for *Y* to do another action i.e. a *request.* In order to make a request *Y* does not depend on *X* making an offer or itself to agree - *Y.request* is not possible from an *offered* or *agreed* state. In this sequence of atomic actions, two transitions do not depend mutually on each other but rather occur sequentially. As we have shown above there is no deadlock in *requested(X),* thus *Y.request?;X.offer* eventually happens enabling *X.offer?;Y.agree.* The testing of *offered(X)* and the atomic action *agree* terminates rendering *agreed(Y)* true. We have proved that *X.offer?;Y.agree* is not a deadlocked action and it causes no deadlock in the terminal *agreed* state. Therefore there are no deadlocks in the *open* state by the induction hypothesis and neither in its successor states, being terminal states.

Case5: When s_0 is *closed, rejected, timedout or agreed(X).* This is a trivial case. These four states being terminal states, there are no transitions possible from them. Deadlock is absent.

Case6: When s_0 is *¬negotiating.* The three possible entry actions are *initial_request, initial_offer, initial_propose.* The induction hypothesis (and base case) asserts that there is no deadlock at *¬negotiation* and *X* can perform one of these actions without waiting for *Y* as these are entry actions. All the transitions are atomic and terminates making states *requested(X), (offered(X)∧¬proposed(X))* or *proposed(X)* eventually true. None of these states contain a deadlock as there are no actions possible before *¬negotiating. Y* does not need to perform any action before *X* can make these entry actions.

Case7: When s_0 is *negotiating*
This state is dealt by analysing all its sub-states in all the above cases.

We have proved all transitions from an arbitrary state s_0 maintain liveness. We have done so for all possible states s_0. **QED**

3.2 Termination

We prove that the termination property holds for the bilateral protocol by showing that all paths from any state lead to the *closed* state. We first derive the sub-paths from a given state to *closed.*

Termination Axiom: $\mu \models \forall s\text{: state-set}\ \forall p_s\text{:path-set}\ \exists p_t\text{:path-set}\ ((\neg\exists p\text{:path-set}, s_i\text{:state-set}\ ((s_t \rightarrow [p]s_j) \wedge \neg(s_t \leftrightarrow s_i)) \rightarrow (s \leftrightarrow [p_s;p_t]\ s_t))$
where $s_i, s_t \in$ state-set

The simplest sub-paths to a successful or unsuccessful *closed* state consists of an atomic transition and are given by *{X ∪ Y}.exit_paths = {X ∪ Y}.(agree ∪ reject ∪ timeout). {X ∪ Y}.exit-paths* represents the possible actions leading to *closed* in a negotiation between agents *X* and *Y.* The sub-paths to a *closed* and *failed* state is *{X ∪ Y}. unsuccesful_exit = {X ∪ Y}.reject ∪ {X ∪ Y}.timeout.* According to the protocol, *unsuccessful_exit* transitions may be chosen at any point in a negotiation. By set theory, the relation between *exit_paths* and *unsucessful_exit* is given by *{X∪Y}.exit_paths ={X∪Y}.(agree ∪ unsucessful_exit).*

We analyse the logical theory of the bilateral protocol to derive the possible paths to *closed* from the *open* states *offered(X), proposed(X) and requested(X)* given by the paths *from_offered(X), from_proposed(X), from_requested(X)* respectively.

from_offered(X) = offered(X)? (Y.agree ∪
 {X ∪ Y}.reject ∪ {X ∪ Y}.timeout)
= offered(X)? exit_paths

from_proposed(X)
= proposed(X)? ((Y.request; from_requested(Y)) ∪
 Y.agree ∪ {X∪Y}.reject ∪ {X ∪ Y}.timeout)
= proposed(X)? ((Y.request; from_requested(Y)) ∪
 {X ∪ Y}.exit_paths)

from_requested (X)
= requested(X)? ((Y.offer; from_offered(Y)) ∪
 (Y.propose; from_proposed(Y)) ∪
 (Y.suggest;from_requested(Y)) ∪ {X∪Y}.reject ∪
 {X ∪ Y}.timeout)

The above 3 path decompositions are the possible paths from an *open* sub-state to *closed.* We can now deduce an abstract compositional path for a negotiation process by sequencing an entry to an *open* negotiation with all possible paths to a *closed* state i.e. *negotiation-paths = entry-actions;paths-to-closed.* There are 3 entry actions as given in the protocol and the *paths-to-closed* are given above. We are left to prove that *negotiation-paths* terminates, thereby proving that all paths reach the *closed* state and that the bilateral protocol shows the termination property.

negotiation_paths = (X.initial_offer; from_offered(X))
 ∪ (X.initial_propose(X); from_proposed(X)) ∪
 (X.initial_request;from_requested(X))

We first expand and simplify each sub-paths occurring after the entry actions in *negotiation_paths* for the purpose of expressing it in terms of entry and exit paths. In *from_requested,* we replace the path *from_offered* and *from_proposed* with their sub-paths ({XUY}.*reject* ∪ {X ∪ Y}.*timeout*) is equivalent to unsuccessful exiting i.e. *unsucessful_exit.*

from_requested(X) = *requested(X)*? (
(*Y.offer*; *offered(Y)*?; {XUY}.*exit_paths*) ∪
(*Y.propose*; *proposed(Y)*? ((*X.request* ;
from_requested(X)) ∪ {XUY}.*exit_paths*)) ∪
(*Y.suggest*;*from_requested(Y)*)∪ *unsucessful_exit*)

It is obvious that the corresponding state holds after a successful action i.e. $X.do_{action}$;$done_{action}$? can be abbreviated to $X.do_{action}$ e.g. *X.propose*;*proposed(X)*?; *Y.do_action* to *X.propose*;*Y.do_action.* We can omit the testing of *proposed(X)*? after *X.propose* because it always succeeds by virtue of reaching this test because *X.propose* must have succeeded previously leading to *proposed(X)*. We remove from *from_requested(X)* tests that are successful by default of reaching them.

from_requested(X) = *requested(X)*? (
(*Y.offer*; {XUY}.*exit_paths*) ∪ (*Y.propose*;
((*X.request* ; *from_requested(X)*) ∪
{XUY}.*exit_paths*)) ∪ (*Y.suggest*;
from_requested(Y)) ∪ *unsucessful_exit*)

Distributing *Y.propose* in ((*X.request*;
from_requested(X)) ∪ {XUY}.*exit_paths*), we have

from_requested(X) = *requested(X)*? (
(*Y.offer*; {XUY}.*exit_paths*) ∪
(*Y.propose*; (*X.request* ; *from_requested(X)*)) ∪
(*Y.propose* ;{XUY}.*exit_paths*) ∪
(*Y.suggest*;*from_requested(Y)*)∪ *unsucessful_exit*)

We further simplify *from_requested* by distributing and grouping the terms with *exit_paths* along the set operators by e.g.
(*Y.offer*;{XUY}.*exit_paths*) ∪
(*Y.propose*;{XUY}.*exit_paths*) is equivalent to
(*Y.offer* ∪ *Y.propose*); {XUY}.*exit_paths*

from_requested(X)
= *requested(X)*? (((*Y.offer* ∪ *Y.propose*);
{XUY}.*exit_paths*) ∪ (((*Y.propose*; *X.request*) ∪
Y.suggest); *from_requested(Z)*) ∪
unsucessful_exit) (*f_r1*)

from_requested(X) has a nested *from_requested* and is therefore an iterative path set. Currently we have *from_requested(X)* in the form *m* = *r?* ((*a* ∪ *b;m*) ∪ *c*) i.e. *m* = *r?*(*d* ∪ *b;m*) = *r*? (*d;b**). Using this equation, we can explain the termination of *from_requested(X)*.

r? is analogous to *requested(X)*? and terminates since it is just a logical test. *a is* ((*Y.offer* ∪ *Y.propose*); {XUY}.*exit_paths*) and terminates since both transitions terminates and are followed by exit_paths. *b* is (*Y.propose*; *X.request*) ∪ *Y.suggest* and terminates as these transitions are finite. *c is unsucessful_exit* terminates being exit points. Thus in the expression *m* = *r?* ((*a* ∪ *bm*) ∪ *c*), the paths *a*, *b*, *c* and *r* terminates.

m as *from_requested* is a recursive path equation. Following our earlier definition of termination for a recursive function, the base case is *d* - ((*Y.offer* ∪ *Y.propose*); {XUY}.*exit_paths*) ∪ *unsucessful_exit.* This base case terminates as all its sub-paths terminates. *m*=*r*? (*d;b**) can be solved using algebra and thus all paths *from_requested* lead to a terminal state.

We similarly simplify *negotiation_paths* by substituting for *from_proposed* and *from_offered* with their sub-paths as follows:

negotiation_paths = (*X.initial_offer*; *from_offered(X)*)
∪ (*X.initial_propose*; *from_proposed(X)*) ∪
(*X.initial_request*;*from_requested(X)*)
= (*X.initial_offer*; *offered(X)*?; *exit_paths*) ∪
(*X.initial_propose*; *proposed(X)*?; ((*Y.request*;
from_requested(Y))∪{XU Y}.*exit_paths*)) ∪
(*X.initial_request*;*from_requested(X)*)

Grouping terms over *from_requested* and *exit_paths* and removing redundant state testing, we obtain

negotiation_ paths = ((*X.initial_offer* ∪
X.initial_propose) ; {XU Y}.*exit_paths*) ∪
((*X.initial_request* ∪ (*X.initial_propose*;
Y.request)); *from_requested(Z)*)

We have shown that the termination axiom holds for *from_requested(X)*. All the other sub-paths in *negotiation_paths* terminates being atomic transitions or exit paths. Therefore all paths in *negotiation_paths* lead to *closed* (a terminal state). **QED.**

We further manipulate *negotiation_paths* to show an interesting representation in terms of entry and exit paths. Replacing *from_requested* with its sub-paths from (*f_r1*)

negotiation_ paths = ((*X.initial_offer* ∪
X.initial_propose) ; {XU Y}.*exit_paths*) ∪
((*X.initial_request* ∪ (*X.initial_propose*;
Y.request)); *requested(X)*? (
((*Y.offer* ∪ *Y.propose*); {XUY}.*exit_paths*) ∪
(((*Y.propose*; *X.request*) ∪ *Y.suggest*);
from_requested(Z)) ∪{XU Y}. *unsucessful_exit*))

Let one entry-point in a negotiation be called *entry-point1* where the action *entry_point1* = *X.initial_offer* ∪ *X.initial_propose.* Let the second entry-point to a negotiation into a *requested* state be *entry_to_requested* where the path *entry_to_requested* = *X.initial_request* ∪ (*X.initial_propose*; *Y.request*). For clarity let {XUY}.*exit_paths* be abbreviated to the path *exit_paths* and the path {XUY}.*unsucessful_exit* be *unsucessful_exit.* Using these two entry points in a negotiation, removing any redundant term, testing in the expression *negotiation_paths* and distributing *entry_to_requested* over the sub-paths of *from-requested* we obtain

negotiation_ paths =
entry_point1; *exit_paths* ∪
(*entry_to_requested*; (*Y.offer* ∪ *Y.propose*) ;
exit_paths) ∪(*entry_to_requested* ; ((*Y.propose*;
X.request) ∪ *Y.suggest*); *from_requested*(Z)) ∪
(*entry_to_requested* ; *unsucessful_exit*)

Let paths to an *offered* state be denoted by *to-offered* where the *offered* state succeeds a *requested* state.
to-offered = *entry_to_requested*; (*Y.offer* ∪ *Y.propose*)
Let *to_requested* denote the sequence ((*Y.propose*; *X.request*) ∪ *Y.suggest*).

negotiation_ paths =
((*entry_point1* ∪ *to-offered*); *exit_paths*) ∪
(*entry_to_requested* ; ((*to_requested*;
from_requested(Z)) ∪ *unsucessful_exit*))

In *negotiation_ paths*, ((*entry_point1* ∪ *to-offered*); *exit_paths*) terminates since it is ultimately followed by the exit-paths and both *entry_point1* and *to-offered* terminates containing atomic transitions. In the path (*entry_to_requested* ;((*to-requested*; *from_requested*(Z)) ∪ *unsucessful_exit*)), the sub-paths *entry_to_requested*, *to-requested*, *unsucessful_exit* all terminate. We have previously shown the paths in *from_requested*, which is the recursive part of a negotiation, lead to a terminal state. Hence all the sub-paths in *negotiation_ paths* lead to a terminal (*closed*) state, as shown earlier.

Examining the latest expression for *negotiation_paths*, we can see that it consists of direct termination via the *offered* state and a recursive path through *requested*. The base case is ((*entry_point1* ∪ *to-offered*); *exit_paths*) showing that a negotiation ultimately passes through the *offered* state for successful exit paths. The recursive case concerns entering and iterating in the *requested* state via the sequential path (*entry_to_requested*;((*to_requested*;*from_requested*(Z)) ∪ *unsucessful_exit*)).

3.3 Soundness

All states, processes and transitions occurring in a negotiation instance, δ, are only those defined by the protocol, μ, if μ is sound.

Soundness axiom for process p: $\forall s, s_1$: *States S* ,$\forall p$:*Processes* α $(\delta \models (s \vee (s \rightarrow [p]\, s_1))$ iff $(\mu \models (s \vee (s \rightarrow [p]\, s_1))$

If μ is sound, δ does not yield any unexpected states. This property ensures the validity of negotiation paths and states, e.g. an *agree* action cannot directly follow a *request* action nor can a state be both *agreed* and *rejected*. There are no erroneous states or transitions. We prove by induction that the soundness property holds for the bilateral protocol by showing that all possible transitions from any state preserve soundness.

The base case proves that the protocol is sound after one atomic transition. This case is covered by the logical theory of the protocol. As induction hypothesis, we assume that the soundness property holds at an arbitrary state *s1* after performing the path p_d. $p_d \in$ *path-set* and may be a *null* process. We then prove that the negotiation stays sound when P_d is either followed or preceded by any valid action (i.e. that is an element of *transition-set*). This means a negotiation instance is sound after the execution of any paths.

Base Case. The base case for the proof of the soundness of the protocol is given by the action-condition rules in theory 1. Atomic transitions in *transition-set* lead from a source state to a target state, both states being well-defined and in *state-set*. The protocol dictates the valid states and atomic transitions in *state-set* and *transition-set* respectively.

From ¬*negotiating*, entry transitions lead to *open* and are given by rule (8) in theory 1. Likewise atomic transitions from *requested, offered, proposed* and *open* are given by rules (9) to (12). The rules (1) to (7) on the relation between states and sub-states ensure that the parent states of the current sub-state are *true* and all other states are *false* as required.

For all atomic transition *p*: *transition-set*, there is a source state s_1: *state-set* and a well-defined target state s_2:*state-set*. If *p* is equivalent to *null* then s_1 and s_2 are the same state e.g. *closed* states are terminal states with no transitions. Atomic transitions in the form $s_1 \leftrightarrow [p]\, s_2$ are given by theory 1. *Induction hypothesis*. For an arbitrary state s_n: *state-set* there is a path p_d: *path-set* that yields a well-defined state s_d: *state-set* given by both the protocol μ and the negotiation instance δ. $(\mu \models s_n \leftrightarrow [p_d] s_d) \wedge (\delta \models s_n \leftrightarrow [p_d] s_d)$.

Induction Proof. We prove that soundness is maintained by the path p_d followed or preceded by an atomic transition, *a*:*transition-set* i.e. the path (p_d;*a*) or (*a*; p_d).

Case1. We first prove soundness for the case (p_d;*a*) by showing that the transition $s_n \leftrightarrow [p_d\, ; a]\, s_{d+1}$ yields a well-defined state s_{d+1}. From the semantic definition of ANML, a path consisting of sub-processes as here in [α; β] can be decomposed in the execution of each of its sub-process with an intermediate state after each sub-process here, α.

$[\alpha; \beta]\, state_\beta \leftrightarrow ([\alpha] state_\alpha \wedge (state_\alpha \rightarrow [\beta]\, state_\beta)) \leftrightarrow ([\alpha]\, state_\alpha?; [\beta]\, state_\beta) \leftrightarrow [\alpha][\beta]\, state_\beta$ where $state_\beta$, $state_\alpha \in$ *state-set*

$s_n \leftrightarrow [p_d ; a]\ s_{d+1}$ is equivalent to $s_n \leftrightarrow [p_d][a]\ s_{d+1}$ which means that there is an intermediary state, s_i:*state-set*, between executing the path p_d and the atomic transition a.

$$s_n \leftrightarrow [p_d]\ s_i?;\ [a]\ s_{d+1}$$
$$s_n \leftrightarrow [p_d]\ s_i \wedge s_i \leftrightarrow [a]\ s_{d+1}$$

From the induction hypothesis we know that $s_n \leftrightarrow [p_d]\ s_d$ gives a sound negotiation and a well-defined state s_d. From the logical theory and the base case, the atomic transition $s_i \leftrightarrow [a]\ s_{d+1}$ preserves soundness. Thus we make s_d equivalent to s_i . We know that s_d is well-defined and therefore so is s_i, which is thus in *state-set.* The resulting two transitions are $s_n \leftrightarrow [p_d]\ s_d$ and $s_d \leftrightarrow [a]\ s_{d+1}$. Since s_d and s_i are in *state-set* and action a is in *transition-set*, then s_{d+1} is in *state-set* given by corresponding action-condition rules in μ. Both transitions yield a sound system (induction hypothesis and logical theory).

The path $(p_d ; a)$ has s_d as the intermediary state. Using theory 1, s_d may be instantiated to an appropriate intermediate state by analysing which state after path p_d is the pre-condition to the execution of a. The transition $s_n \leftrightarrow [p_d ; a]\ s_{d+1}$ preserves the soundness of the negotiation since its sub-paths p_d and a do so. A negotiation in the state s_n is sound and remains sound at s_{d+1} after an atomic transition. Having proved that an atomic transition in *state-set* preserves the soundness of a negotiation, and assuming that a path p_d preserves soundness, we have proved that the path consisting of $p_d ;a$ also yields a sound negotiation.

Case2. The proof for soundness when preceding the path p_d with an action is similar. We prove that $s_{n-1} \leftrightarrow [a;\ p_d\ ]\ s_d$ maintains a sound negotiation for a:*transition-set.* From the base case, any atomic transition in *transition-set* provides a sound negotiation. The induction hypothesis asserts that a path p_d preserves the soundness property where $s_n \leftrightarrow [p_d]\ s_d$ gives a well-defined state s_d. We need to prove that the path $a;\ p_d$ also yields a sound negotiation at all states.

By the definition of process composition in ANML, $s_{n-1} \leftrightarrow [a;\ p_d\ ]\ s_d$ is equivalent to $s_{n-1} \leftrightarrow [a]\ [p_d]\ s_d$ where there is an intermediate state, s_{i1}, after executing atomic transition a and before executing the process p_d.

$$s_{n-1} \leftrightarrow [a]\ s_i?;\ [p_d]\ s_d$$
$$s_{n-1} \leftrightarrow [a]\ s_i \wedge s_i \leftrightarrow [p_d]\ s_d$$

We equate state s_n to the intermediary state s_i making thus s_i to be in *state-set*. The action, a, is in *transition-set*, therefore from the logical theory and the base case, s_{n-1} is in *state-set* and is well-defined. Since the path p_d and the action a preserves soundness giving s_{n-1} and s_d to be in *state-set*, the path $a;\ p_d$ preserves the soundness property as its sub-paths do so. **QED**

We have proved that an atomic transition gives a sound negotiation, and assuming that a path p_d preserves soundness, we have shown that the path consisting of $a;\ p_d$ also maintains soundness. By induction, we have thus shown that all paths in a bilateral negotiation will preserve the soundness property. The theory of the bilateral protocol is sound.

4 Conclusion

We can use our meta-language for defining and proving the properties of negotiation protocols represented in it. Compared to some temporal logic systems, it is relatively simple to express intuitive axioms for a property in ANML. Our meta-language is more suitable for providing a computational representation of a negotiation process in terms of its states and processes. After defining axioms in ANML for liveness, termination and soundness properties for a protocol, we have shown how to prove them for a bilateral protocol in ANML. Other protocols may be similarly represented and reasoned about using our framework. Our proof methods can be further generalised for analysing process executions in state-transition systems-like.

There are a number of other properties that we have examined in our work such as fairness, performance and complexity, consistency, reliability, serial. We could also analyse game theoretic properties of a protocol. These include Pareto efficiency, stability, possible equilibrium, deception-free and conflict resolution, (Rosenschein and Zlotkin, 1998).

However proving that a protocol satisfies a property may not always be feasible. This problem may be remedied by monitoring the executions of a negotiation for violations of the property. While it is possible to monitor an interaction for violations of safety properties, this is not so easy for liveness properties. Hence, one may prove safety properties for a protocol but be unable to establish the liveness property.

References

M. Barbuceanu and M. S. Fox. "COOL: A Language for Describing Coordination in Multi-Agent Systems". *In Proc. of the International Conference on Multi-Agent Systems*, San Francisco, CA. 1995.

N. Bjorner, A. Browne, M. Colon, B. Finkbeiner, Z. Manna, H. Sipma and T. Uribe. "Verifying temporal properties of reactive systems: A step tutorial". *Formal Methods in System Design.* 1999.

R. M. Burstall. "Program proving as hand simulation with a little induction." *In Information Processing 74.* Stockholm pp 308-312. 1974.

P. R. Cohen and H. J. Levesque. "Teamwork". *Nous*, 35, 25(4):487—512. 1991.

J. Cunningham, S. Paurobally, A. Diacakis, S. McConnell, G. Gross, L. Lorenzen. "Requirements for Electronic Commerce". *In proc. Trends in Distributed Systems for Electronic Commerce*, Hamburg, Springer LNCS 1402. 1998.

T. Finin, J. Weber, G. Wiederhold, M. Genesereth, R. Fritzson, D. McKay, J. McGuire, R. Pelavin, S. Shapiro, S. Buffalo, C. Beck. "Specification of KQML, Agent Communication Language" *The DARPA Knowledge Sharing Initiative External Interfaces Working Group, February, 1993.*

FIPA. Agent Communication LanguageTechnical report. Foundation for Intelligent Physical Agents. http://www.fipa.org. 2001.

R. Goldblatt. Logics of Time and Computation, *Lecture Notes, Center for the Study of Language and Information.* 1987.

M. Gordon and T. Melham. "Introduction to HOL: A Theorem Proving Environment for Higher-Oder Logic". *University Press*, Cambridge, 1993.

D. Harel, M. Politi. "Modeling reactive systems with statecharts, The Statemate approach". *McGraw-Hill.* 1998.

R. Keller. "Formal Verification of Parallel Programs". *Communications of the ACM*, vol. 7. 1976.

Y. Labrou and T. Finin. "Semantics for an agent communication language". *The fourth Int. workshop on Agent Theories, Architectures and Languages.* Rhode Island, USA. 1997.

Y. Labrou. "Standardising Agent Communication". *In Proceedings of 9th ECCAI Advanced Course and EASSS 2001.* Springer, LNAI 2086. 2001.

L. Lamport. "Sometimes is sometimes "not never" - on the temporal logic of programs". *In Proceedings of the 7th ACM Symposium on Principles of Programming Languages*, pages 174--185, January 1980.

O. Lichtenstein, A. Pnueli, L. Zuck. "The glory of the past". *Proc. Workshop on Logics of Programs*, Brooklyn, 1985, Springer-Verlag, LNCS 193, pp. 196--218. 1985.

Z. Manna and A. Pnueli. "Temporal verification of reactive systems – safety". *Springer-Verlag ISBN 0387944591.* 1995.

Z. Manna and H. Sipma. "Deductive Verification of Hybrid Systems using SteP". *In Hybrid systems: computation and control. International Workshop* LNCS 1386, Springer-Verlag. 1998.

M. Nodine and A. Unruh. "Facilitating open communication in agent systems: the InfoSleuth infrastructure". *In Proc. of the Fourth Int. Workshop on Agent Theories, Architectures, and Languages*, 1997. MCC-INSL-113-98.

S. Owicki and L. Lamport. "Proving Liveness Properties of Concurrent Programs." *ACM Transactions on Programming Languages and Systems.* Vol. 4, No. 3, pages 455 – 495, July 1982.

D. H. Van Parunak. "Visualizing agent conversations: Using enhanced dooley graphs for agent design and analysis". *In Proc. of the 2nd International Conference on Multi-Agent Systems* (ICMAS '96). 1996.

S. Paurobally and J. Cunningham. "Specifying the Processes and States of Negotiation". *In Agent Mediated Electronic Commerce. The European AgentLink Perspective.* C. Sierra and F. Dignum, editors. ISBN: 3-540-41671-4. Springer Verlag. pages 61-77. 2000.

J. Pitt and A. Mamdani. "Communication protocols in multi-agent systems". *In Working Notes of the Workshop on Specifying and Implementing Conversation Policies, Third International Conference on Autonomous Agents* pages 39-48. 1999.

A. Pnueli. "The temporal semantics of concurrent programs." *Theoretical Computer Science* 13:1-20. 1977.

J. S. Rosenschein and G. Zlotkin. "Rules of Encounter. Designing Conventions for Automated Negotiation among Computers". MIT Press. ISBN 0262181592. 1998.

R. E. Shostak. "Deciding combinations of theories". *Journal of the ACM*, 31(1):1--12, 1984.

A. P. Sistla and E. M.Clarke. "The complexity of propositional linear temporal logics". *Journal of the ACM*, 32(3):733--749, 1985.

A. P. Sistla. "Safety, liveness and Fairness in Temporal Logic". *Formal Aspects in Computing*, 1994, vol. 6, pp 495-511. 1994.

R. G. Smith. "The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver". *IEEE Transactions in Computers. Volume 29, Number 12*, (1104-1113). 1980.

Towards a Trust Management Model for Electronic Marketplaces

Theo Dimitrakos
Business and Information Technology Department, Central Laboratory of the Research Councils, Rutherford Appleton Laboratory, Oxfordshire, OX11 0QX

Abstract

Motivated by an industrial application, we discuss how role-based modelling can be used to support the analysis of the conditions under which trust may be transferred between agents in an e-market. Following an analysis of our working model of trust, we survey a number of formalisations of complementary aspects of trust. We explain why each of them alone fails to interpret our model, and we suggest an integration, which appears to provide an adequate formalisation. We conclude by sketching how the trust modelling approach presented in this paper can be put together with a suitable risk management scheme, so as to provide a basis for a trust management framework.

1 Introduction

Differing trust relationships can be found among the parties involved in a contract, and the emerging virtual communities require richer models of trust, in order to distinguish between them, and accommodate them in the context of a specific service. To build consumer confidence, e-commerce platform providers need to improve the existing technology in order to capture, measure and manage the trusting relationships that underlie such services. If e-commerce is to achieve the same levels of acceptance as traditional commerce, trust management has to become an intrinsic part of it.

Current solutions fail to incorporate in their decision making evidence or opinions collected by an agent through the agent's own experience, or via communication with other agents who cohabit the system. This makes the evolution of e-commerce systems harder and impedes their ability to adapt to changes in trust and to set up new relationships. In order to be able to handle trust dynamics, future solutions will have to incorporate methods to simulate learning, reasoning and analysing transactional and environmental risks with respect to the agents' view of the system they inhabit.

We suggest that the needs for flexibility and scalability are better addressed by separating the trust management framework from the purpose of the application, and we emphasise that risk management and role-based modelling can be combined to support trust management solutions. More specifically, we find roles to be well suited for modelling service-specific aspects of trust and particularly helpful for identifying and analysing cases where trust may be transferable. We also anticipate that risk management can guide an agent through transforming a mere *inclination* to trust into a carefully considered *intention* to trust, and through endorsing dependable *behaviour* as a realisation of the agent's dependable *intentions* to trust.

Our motivation stems from an industrial application outlined in section 2. In section 4 we present a role-based model of trust in e-markets which supports the analysis of trust relationships between agents on the basis of the definition of trust in e-services (section 3). In section 5 we investigate the formalisation of this model and in section 6 we sketch a trust management scheme which can be superimposed on our trust model in order to provide a basis for trust management solutions.

2 Motivating example

As a motivating example, we summarise the Home Shopping Tool (HST) component[1], which will be used as one of the test-beds in the CORAS project[2]. HST delivers a personalised, targeted marketing experience through the realisation of a variety of services including personalised shopping, catalogue information, product search and recommendations, sales negotiation, e-payment, and user management facilities. Specific consumer information is gathered for the purpose of behaviour analysis and can be made available to the platform operator.

Notably, consumers and suppliers are provided with an agent-based automated bargaining mechanism that al-

[1] Developed in the ESPRIT EP 27046 project ACTIVE.

[2] CORAS (IST-2000-25031) is building a tool supported framework for model-based risk analysis

lows customers to find and negotiate products with various suppliers, and suppliers to promote their products and attract consumers. The agents get involved in a negotiation process and try to reach a mutual agreement according to the mandates given by their creators.

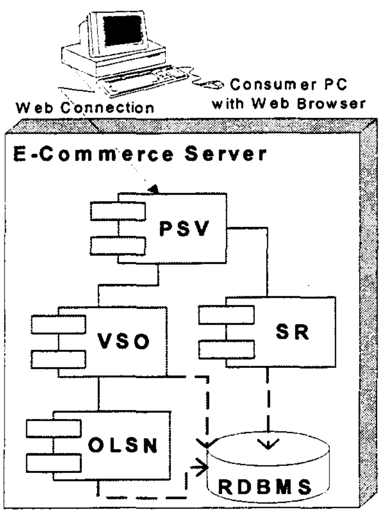

Figure 1. Home Shopping Tool Components (HST)

These services are offered with the help of the following software modules depicted in *Figure 1*:

- *Virtual Shopping Operator (VSO),* which is responsible for the provision of basic purchasing services, access to product information through electronic catalogues and Point of Sales (POS) type of services (i.e. secure order processing and payment).
- *Shopping Recommender (SR),* which provides recommendations to the users about products purchasing i.e. discounts, special offers, new products, etc. taking into account the consumer's profile, which is registered in the "Consumer Information Model" (CIM) database. In addition, recommendations are provided on the basis of other consumers' comments.
- *On-line Sales Negotiator (OLSN),* which uses software agents to discover, negotiate and purchase products on behalf of the shoppers and retailers, following a set of user-specified constraints (e.g., a highest or lowest acceptable price, technical specification, a date by which to sell or buy). Launching of agents is optional depending on consumer or retailer requests.
- *Personalised Store Visualiser (PSV),* which provides an interface based on consumer preferences. (E.g. the consumer may wish to see only products in pre-selected categories or to use a Virtual Reality interface instead of catalogue navigation.

The e-commerce platform presented in this section will be modelled in order to perform security risk analysis in the course of the CORAS project. We will further analyse these models in relation to e-commerce trust (in parallel to, and independently of, CORAS) and use them as indicative examples to relate basic security and fairness properties to trust. During this analysis we will assess the effectiveness of, and further develop, the model of trust outlined in this paper.

3 Defining trust in e-services

In the physical world, we derive much of our notions of trust from the tangible nature of the entities in our environment. Our trust relies on personal contact, the tangibility of the objects and services, the difficulty of fraudulence and a clearly defined legal framework. Personal contact in virtual communities is limited, the legal framework is vague and the objects and services under negotiation are less tangible. The traditional notions of trust need to be rethought, and suitable models of trust in e-commerce have to be developed. In (Dimitrakos 2001) we provide a rigorous definition of trust in e-commerce, following a survey (Dimitrakos & Bicarregui 2001, Grandison & Sloman 2000) of recent attempts to formalise this concept.

Definition 1.
Trust of a party A in a party B for a service X is the measurable belief of A in B behaving dependably for a specified period within a specified context in relation to the service X.

Remarks 2.

- A party can be an individual entity, a collective of humans or processes, or a system; (obviously, the trustor must be an entity that can form a belief).
- The term service is used in a deliberately broad sense to include transactions, recommendations, issuing certificates, underwriting, etc.
- The above mentioned period may be in the past, the duration of the service, future (a scheduled or forecasted critical time slot), or always.
- Dependability is used broadly to include security, safety, reliability, timeliness, and maintainability.
- The term context refers to the relevant service agreements, service history, technology infrastructure, legislative and regulatory frameworks that may apply.
- Trust may combine objective information with subjective opinion formed on the basis of factual evidence and recommendation by a mediating authority.
- Trust allows one agent to rely for a critical period on behaviour or on information communicated by another agent. Its value relates to the subjective

probability that an agent will perform a particular action (which the trustor may not be able to monitor) within a context that affects the trustor's own actions.

Notably, our definition differs from (Jøsang 1996) and (Kini et al. 1998) with respect to the trusting subjects. Intelligent agents who negotiate can be either humans or programs and in both cases they need to manifest trust intentions and establish trusting relationships. Intelligent software agents are adaptive autonomous programs featuring the ability to acquire knowledge and to alter their behaviour through learning and exercise. Their decision making can be enhanced so that they form trust intentions and make decisions relying on trust. Our definition differs from (Grandison & Sloman 2000, S. Jones 1999), and (Yahalom 1993) with respect to the inherent measurability and the subjective nature of trust. It also differs from (Gambetta 1990, Castelfranchi 1998, Grandison & Sloman 2000, Jøsang 1996) in that trust differentiates between services and it is active for critical periods of time.

We also note that distrust, accounting to what extent we can ignore one's claims about her own or a third party's trustworthiness and their proclaimed actions or commitments, is modelled as a measurable belief in that a party will behave non-dependably for a critical period within a specified context.

Definition 3.
Distrust of a party A in a party B for a service X is defined as A's measurable belief in that B behaves non-dependably for a specified period within a specified context in relation to service X.

Distrust is useful in order to revoke previously agreed trust, obstruct the propagation of trust, ignore recommendations, and communicate that a party is "blacklisted" for a class of potential business transactions.

4 A role-based model of trust

In this section we characterise the notion of trust captured in Definition 1 (section 3) by discussing some general properties of trust relationships. These properties provide the foundation of a role-based model that is able to support an elaborate analysis of some basic trust relationships and the structural properties that underpin the propagation of trust in e-services.

4.1 General properties of trust

The particular characteristics of trust may differ from business to business. Nevertheless, there are some common delimiters that indicate the existence of general principles governing trust in e-commerce.

Proposition 4.
The following are general properties of trust and distrust.

P1. *Trust is relativised to some business transaction.* A may trust B to drive her car but not to baby-sit.

P2. *Trust is a measurable belief.* A may trust B more than A trusts C for the same business.

P3. *Trust is directed.* A may trust B to be a profitable customer but B may distrust A to be a retailer worth buying from.

P4. *Trust exists in time.* The fact that A trusted B in the past does not in itself guarantee that A will trust B in the future. B's performance and other relevant information may lead A to re-evaluate her trust in B.

P5. *Trust evolves in time, even within the same transaction.* During a business transaction, the more *A* realises she can depend on *B* for a service *X* the more *A* trusts *B*. On the other hand, *A*'s trust in *B* may decrease if *B* proves to be less dependable than *A* anticipated.

P6. *Trust between collectives does not necessarily distribute to trust between their members.* On the assumption that A trusts a group of contractors to deliver (as a group) in a collaborative project, one cannot conclude that A trusts each member of the team to deliver independently.

P7. *An agent is aware of its degree of trust in itself.* Mary may trust her lawyer to win a legal case more than she trusts herself to do it. Self-assessment underlies an agent's ability to delegate or offer a task to another agent, so as to improve efficiency or reduce risk.

Remarks 5

- P1 states that trust depends on the *tasks* that each entity is expected to perform in the context of some particular business.
- P2 states that there are different *degrees of trust*: an agent *A* may trust agent *B* more than *A* trusts agent *C* for the same task and in the same business context.
- P3 states that different parties with different roles in a transaction may have different views on trust in each other or in third parties; to an extend trust is subjective.
- P4 states that trust depends on a particular business transaction and momentum. Trust may be relativised to a service that is relevant for a critical interval, thus trust statements may be time-stamped and become irrelevant outside this interval.
- P5 emphasises dependence of trust on a sequence of events.
- P6 distinguishes trust in a collective from trust in its members. On the assumption that *A* trusts a group of

contractors to deliver (as a group) in a collaborative project, one cannot conclude that *A* trusts each member of the team to deliver in the project. A potentially bad performance of a member of the group can be balanced by an excelling performance of another.

- P7 supports the ability of an agent to delegate or offer a task to another agent in order to improve efficiency or reduce risk.

4.2 Trust propagation and transferability

As we elaborate in the sequel, at least unintentional transferability of trust within a locus may be acceptable in specific contexts. *Note that "transferability" in our case corresponds to influencing the level of trust rather than relational transitivity.* We distinguish three special *roles* that entities mediating in a trust relationship can play. These roles are *guarantors, intermediaries*, and *advisors*. Note that an entity may play more than one mediating role in a business relationship.

Guarantor is a party taking the responsibility that the obligations of the parties she acts as a guarantor for are fulfilled at an agreed standard. Guarantors assist the establishment or facilitate the increase of trust for a specific transaction by underwriting (a part of) the risk associated with the transaction. A typical example is a credit card company.

Intermediary is a party that intervenes between other parties in a business transaction and mediates so that they establish a business relationship with or without their knowledge. We distinguish the following types of intermediary:

- *Transparent:* an intermediary that identifies the parties she is mediating between to each other. An example is Lloydstsb.com, a bank, who offer to their on-line customers a comprehensive car rental and flight booking service powered by Expedia.co.uk.
- *Translucent*: an intermediary that identifies the existence of the parties she mediates between but not their identity. An example is a retailer advertising product delivery by courier without identifying which courier.
- *Overcast:* an intermediary that hides the existence of the parties she is mediating between from each other. Examples include enterprises that are selectively outsourcing tasks to unidentified strategic allies.
- *Proxy:* an intermediary who is authorised to act as a substitute of another entity.

Advisor is a party that offers recommendations about the dependability of another party. Examples include credit scoring agencies and reputation systems. Advisors include the agents maintaining blacklists for a community.

Proposition 6.
Trust and distrust propagate according to the following rules:

P8. *(Dis)trust is not transferred along an overcast intermediary.* Assume that A (dis)trusts an overcast intermediary T for a service X provided by B. Since A is not aware that B provides the service, her (dis)trust is placed in T.

P9. *Trust is transferred along transparent intermediaries – distrust is not.* Assume that, for a service X, A trusts a transparent intermediary T mediating for B. By agreeing to the service, A expresses trust in B for X instigated by T's mediation.

P10. *(Dis)trust in a subcontractor of a transparent intermediary is transferred to (dis)trust in the intermediary.* If a party A (dis)trusts a subcontractor of a transparent intermediary T for a service X, then A is inclined to (dis)trust T for this particular service.

P11. *Trust is transferred anonymously along translucent intermediaries – distrust is not.* Assume that A trusts a translucent intermediary T for X and T trusts B to subserve for X. By agreeing to the service, A effectively expresses trust in a third party to subserve for X without necessarily knowing identity of that party.

P12. *Trust in an advisor is transferred to the recommended party - distrust is not.* The more A trusts T the more she relies on her recommendation.

P13. *Distrust in a recommended party is transferred to the advisor – trust is not.* A's distrust in a party B recommended by T for a service X prompts A to question T's competence as an advisor for X.

P14. Advisors distinguish between recommendations based on "first hand" and "second hand" evidence. In the latter case they ought to identify their sources. If T_1 and T_2 both pass to A advise by T as their own observations then T gains an unfair advantage in influencing A. See (Jøsang 1999) for further analysis.

P15. Distrust propagates through trust and it obstructs the propagation of trust. If A distrusts an intermediary T for a service X then A will ignore T's mediation to the extent of the distrust.

Note that P9, P10 and P12, P13 allow for trust and distrust to be transferred in opposite directions. This does not necessarily result in a conflict. The opposite initial values will affect each other and the final decision will depend on the resulting balance between trust and distrust in each party, and the tendencies of the trustor.

5 Towards a formal model of trust

There have been some attempts to formalise aspects of trust relationships within a logical framework (e.g. Firozabadi & Sergot 1999, Jones & Sergot 1996, Firozabadi & Jones 2000, Jøsang 1999, Jøsang 2001, Marsh 1994

Yahalom 1993). The majority of these attempts involves variants of first order logic or tailored modal logics with distinctive deontic elements.

In general, one can expect the logic used for modelling trust to be rich enough to represent actions and interactions between distributed agents, temporal constraints and deontic statements relating to duty and obligation as ethical concepts, while providing an effective mechanism for assigning and inferring subjective valuations of trust. Modal logics can be used to express possibility, necessity, belief, knowledge, temporal progression and other modalities. (We consider deontic, temporal and dynamic logic to be special subclasses of modal logic.) The underlying concepts of possible worlds and accessibility can be useful when modelling trust. Possible worlds can be used to describe future states or alternative realities, for a statement may take a different meaning depending on the world it is in. Evidence based and subjective reasoning (e.g. Jøsang 2001) is also useful for assigning trust values in the presence of uncertainty. It addresses the problem of forming a measurable belief about the truth or falsity of an atomic proposition denoting a state, event or identifying an agent, in the presence of uncertainty.

The rest of this section is structured as follows. First, we elaborate on a number of pre-selected formalisms, which are focusing on complementary aspects of trust. Then, we investigate a potential integration of these formalisms that appears to be rich enough to provide the basis for formalising the proposed trust-modelling framework.

5.1 Modelling patterns of trust behaviour

Jones and Sergot (1996) provide a formal theory to reason about situations where an agent is given institutionalised power to ensure a certain state of affairs. They focus on a non-monotonic[3] conditional relation $\Rightarrow_X$ called *"counts as"* which is used for encoding the idea that, within a given institution, the performance of an act in a given context by a designated agent *"counts as"* a way of establishing a particular institutional fact. Conditional sentences of the form $\Rightarrow_X$ describe that according to an institution *X*, establishing that the state of affairs described by *A* is, counts as a means of establishing that the state of affairs described by *B* is.

Jones and Firozabadi (2000) extended this to a formal theory about trust and deception. They integrate a modal action logic (developed by Kanger, Pörn (1989 & 1997) and Lindahl) to specify the actions of an agent, a belief logic and a deontic logic to describe commitment. The action component logic is axiomatically similar to a relativised classical modal system of type ECT, with one further axiom schema asserting that logical truths fall outside anyone's agency. This action logic uses a family of modalities E_i where $E_i\varphi$ denotes that *"agent i brings it about that φ"*. In contrast to the modal action logics commonly used in computer science this logic focuses on who the agent is and the results of her agency, abstracting away both temporal aspects and considerations by means of which a particular state is brought about. The belief component logic is axiomatically similar to a relativised KD45, reading expressions of the form $B_i\varphi$ as *"agent i believes that φ"*. The deontic component logic is axiomatically similar to a relativised normal modal system of type KD. Sentences of the form $O_X\varphi$ are read as *"the optimal functioning of system X requires the establishment of φ"*. As it is also indicated in the following example, the resultant formalism is rich enough to support expressing deception and reasoning about patterns of (dogmatic) trust including an agent's belief that another agent makes something happen.

Example 7.
The following statements are expressible in Jones & Firozabadi (2000) framework.

- A statement of the form *"agent 1 believes that agent 2 sees to it that φ"* corresponds to the sentence $B_1E_2\varphi$.
- A statement "agent 1 believes that within context X agent 2 makes φ happen" corresponds to the sentence $((E_2\psi \Rightarrow_X O_X\varphi) \wedge B_1E_2\psi) \rightarrow B_1\varphi$.
- Finally deception is captured by sentences of the form $\neg B_2\varphi \wedge E_2B_1\varphi$ denoting that "agent 2 does not believe φ but gets agent 1 to believe it".

This formalism is well suited for describing patterns of trusting behaviour but it relies on the existence of trust relations and cannot accommodate reasoning about procedures that establish trust between agents. Furthermore the use of Standard Deontic Logic as a component formalism makes it inappropriate for security: as it is noted in (Damianou et al. 2001), Standard Deontic Logic and its variants support a view of permission of β as the opposite of the obligation of $\neg\beta$ which is counterintuitive for access control in open systems. Finally, there is also little or no support for capturing distrust or uncertainty in an agent's trust.

Nevertheless Jones & Firozabadi (2000) capture the right kinds of reasoning one has to combine in order to build agents able to reason about trust and deception:

- A non-monotonic consequence relation $\Rightarrow_X$ between states of affairs in a given context *X*.
- A system to reason about actions in different modalities (including time).

[3] Non-monotonic in the sense that sentences of the form $(\alpha \Rightarrow_X \varphi) \rightarrow ((\alpha \wedge \beta) \Rightarrow_X \varphi)$ are not tautologies.

- A system to reason about obligations.
- A belief system.

In the following subsection we summarise an alternative belief system, which could be adapted to replace the relativised variant of KD45 in order to introduce measurability in the above-mentioned integrated reasoning system and enhance it to support subjective reasoning in the presence of uncertainty.

5.2 Modelling subjective beliefs

The problem of assigning trust values in the presence of uncertainty is addressed though, in Jøsang's trust model (1996) which incorporates the concept of an *"opinion"* based on subjective logic (Jøsang 2001). Subjective logic addresses the problem of forming a measurable belief about the truth or falsity of an atomic proposition denoting a state, event or identifying an agent, in the presence of uncertainty. It is defined by integrating classical logic and a theory of subjective probabilities, based on an extension of the Dempster-Shafer theory of evidence (Shafer 1976).

Definition 8.
An agent's A opinion about a proposition φ is a representation of a belief that is modelled as a triple $\omega_A(\varphi)=\langle b_A(\varphi), d_A(\varphi), u_A(\varphi)\rangle$ where:

- $b_A(\varphi)$ measures A's belief, represented as the subjective probability that a proposition φ is true;
- $d_A(\varphi)$ measures A's disbelief, represented as the subjective probability that a proposition φ is false;
- $u_A(\varphi)$ measures A's uncertainty, represented as the subjective probability that a proposition φ is either true or false;
- $b_A(\varphi), d_A(\varphi), u_A(\varphi) \in [0\ldots1]$ and $b_A(\varphi)+d_A(\varphi)+u_A(\varphi)=1$, for any proposition φ and agent A.

For brevity, in the following we may omit the subscript denoting the agent and/or the proposition about which the opinion is formed if this is clear from the context.

Remarks 9.
Jøsang (2001) defines an opinion as a quadruple $\omega=\langle b,d,u,a\rangle$ where a denotes relative atomicity. Roughly, assuming a frame of discernment Θ, which delimits a set of possible states of affairs exactly one of which is assumed to be true at any one time, relative atomicity normalises uncertainty by incorporating the percentage of the elementary states of affairs that are covered by the state of affairs about which the opinion is formed. In short, relative atomicity indicates the likelihood that uncertainty should have been belief. For the purpose of this paper, we have taken the simplifying assumption that $a=0.5$ and have therefore omitted relative atomicity in order to simplify the example. In real applications, though, relative atomicity is important as it relates uncertainty to probability expectation $E(\omega)$, where $E(\omega)=b+a\cdot u$.

A strong correlation between this opinion model and the probability density functions associated with the beta distribution ensures that opinions can be deterministically established if all available evidence can be analysed statistically (Jøsang 2001). Subjective logic has been used in (Jøsang 1999) in order to reason about decisions involved in authenticating public keys based on recommendations and certificates and to illustrate that, in this case, trust in remote agents can be determined by embedding trust recommendations inside public key certificates.

5.2.1 An overview of subjective logic operators

The following operations of subjective logic provide the essential means for combining trust beliefs from different agents and incorporating recommendations from advisors:

Conjunction: The conjunction $\omega(\varphi)\wedge\omega(\psi)$ is an opinion $\langle b,d,u\rangle$ such that:

$b=b(\varphi)\cdot b(\psi)$, $d=d(\varphi)+d(\psi)-d(\varphi)\cdot d(\psi)$,
$u=b(\varphi)\cdot u(\psi)+u(\varphi)\cdot b(\psi)+u(\varphi)\cdot u(\psi)$

Conjunction forms an opinion about the *conjunction of two propositions that refer to distinct judgements by the same agent.* Belief in the conjunction is viewed as belief in both propositions being true simultaneously, which is measured by multiplying the probabilities of each of the two propositions being true individually. Disbelief in the conjunction is viewed as belief in either of the two propositions being false, which amounts to the union of the probabilities of each proposition being false individually. Uncertainty in the conjunction is viewed as the belief that either one proposition is true and the truth of the other one is uncertain or that the truth of both propositions is uncertain.

Disjunction: The disjunction: $\omega(\varphi)\vee\omega(\psi)$ is an opinion $\langle b,d,u\rangle$ such that

$b=b(\varphi)+b(\psi)-b(\varphi)\cdot b(\psi)$,
$d=d(\varphi)\cdot d(\psi)$,
$u=d(\varphi)\cdot u(\psi)+u(\varphi)\cdot d(\psi)+u(\varphi)\cdot u(\psi)$

Disjunction forms an opinion about the *disjunction of two propositions that refer to distinct judgements by the same agent.* Belief in the disjunction is viewed as belief in either of the two propositions being true, which amounts to the union of the probabilities of each proposition being true individually. Disbelief in the disjunction is viewed as belief in both propositions being false simultaneously, which is measured by multiplying the probabilities of each of the two propositions being false individually. Uncertainty in the conjunction is viewed as the belief that either one proposition is false and the truth of the other one is uncertain or that the truth of both propositions is uncertain.

Negation: The negation: $\neg\,\omega(\varphi)$ is an opinion $\langle b,d,u\rangle$ such that: $b=d(\varphi)$, $d=b(\varphi)$, $u=u(\varphi)$

The negation of an opinion about a proposition being true amounts to the opinion about the same proposition being false. This is in accordance with the fact that subjective logic is based on a foundational integration of classical logic with a theory of posterior probabilities. Had another component calculus being used as the basis, one would of course expect the subjective logic operator to reflect the truth-function transformation associated with negation in the component calculus.

Recommendation: $\omega_A^B(\varphi)=\omega_A(i_B)\otimes\omega_B(\varphi)$ is an opinion $\langle b,d,u\rangle$ such that:

$$b=b_A(i_B)\cdot b_B(\varphi),$$
$$d=b_A(i_B)\cdot d_B(\varphi),$$
$$u=d_A(i_B)+u_A(i_B)+b_A(i_B)\cdot u_B(\varphi)$$

The recommendation operator[4] $\otimes$ combines agent A's opinion about agent B's advice with agent B's opinion about a proposition φ expressed as an advice from agent B to agent A. Note that B's opinion about φ is weighted against A's trust in B's advice, while A's distrust and uncertainty in B's advice increases A's uncertainty about φ. This is consistent with the trust model introduced in (Dimitrakos 2001) and elaborated in sections 3 and 4 of this paper. Here, agent B plays the role of the intermediary (viz. an advisor) for assessing the truth of φ. As it is hard to comprehend physical belief discounting, the definition of an operator for discounting opinions lends itself to different interpretations. In particular, it is not obvious how to describe the effect of A disbelieving that B will give good advice. In this case, the latter contributes to agent A's uncertainty about the truth-value of φ regardless of what B's advice is.

The operator $\otimes$ is associative, but not commutative. This means that, in the case of a chain of recommendations, the same discounting of opinions will be calculated by starting at either end of the chain. However, the order of the recommendations is important. Finally, independence of opinion dictates that the same advisor is not allowed to appear more that once in a recommendation chain. This, in effect, requires that advisors refrain from mixing "first-hand" and "second-hand" evidence in their recommendations: an advisor should always include the source of the evidence in her recommendation (e.g. first hand evidence, or agent X's recommendation).

Consensus: The consensus operator $\oplus$ combines the opinions of two agents A and B about the same proposition φ in a fair and equal way, resulting in the opinion of an imaginary agent $A+B$ about φ. The consensus $\omega_{\{A,B\}}(\varphi)=\omega_A(\varphi)\oplus\omega_B(\varphi)$ is an opinion $\langle b,d,u\rangle$ such that:

$$b=\frac{b_A(\varphi)\cdot u_B(\varphi)+u_A(\varphi)\cdot b_B(\varphi)}{k},$$
$$d=\frac{d_A(\varphi)\cdot u_B(\varphi)+u_A(\varphi)\cdot d_B(\varphi)}{k},$$
$$u=\frac{u_B(\varphi)\cdot u_A(\varphi)}{k}$$

where $k=u_A(\varphi)+u_B(\varphi)-u_A(\varphi)\cdot u_B(\varphi)$ and $\oplus$ is undefined when $u_A(\varphi)=0=u_B(\varphi)$.

Consider that agents A and B observe a machine over two different time intervals and have potentially different opinions about it depending on the behaviour of the machine in the respective periods. The consensus corresponds to the opinion that a single agent $A+B$ would have after observing the machine over both periods.

This operator $\oplus$ assumes a set of mutually consistent opinions and reduces uncertainty. Uncertainty is therefore interpreted as "room for influence", meaning that it is possible to reach a consensus only with somebody who maintains some uncertainty. The effect of this operator is analogous to the situation in court where several witnesses give consistent testimony, hence, amplifying each other's views in the judge's opinion. This operator is both associative and commutative; prescribing the order in which opinions are combined to form a consensus is irrelevant. However, opinion independence has to be assumed in order to avoid allowing for an agent to unfairly influence consensus. This basically means that the opinion of an agent should not be counted more than once. Finally, note that dogmatic opinions cannot be combined into a consensus. This is because uncertainty is interpreted as "room for influence". Having conflicting dogmatic opinions is philosophically counter intuitive (unless if deception is admitted as a

[4] This operator is also called "discounting" (Jøsang 2001) and is similar to the "discounting" function of (Shafer 1976). In fact, if c denotes the "discounting" rate of (Shafer 1976), then, by setting $c=d_A(i_B)+u_A(i_B)$, the two definitions become equivalent

possibility). In this model, opinions about the same "real" observation either incorporate some uncertainty or are necessarily equal.

Example 10. Consider the authentication graph depicted in
Figure 2. Agent H is aware of authentication authorities 1 and 2, which have information about P's authenticity. Authority 1 has direct evidence about P's authenticity whereas authority 2 uses indirect information from authorities 3 and 4, both of which are based on direct evidence from authority 5. Agent P states her confidence in the authorisation means k she uses. The trust values of each principle are encoded in the opinions provided in
Figure 3. The trust level of H in P's authenticity is modelled by the formula: $\omega_{H.D}(k) = \omega_{H.D}(i_p) \otimes \omega_p(k)$ where:

$$\omega_{H.D}(i_p) = \left(\omega_H(i_2) \otimes \left(\left(\omega_2(i_3) \otimes \omega_3(i_5)\right) \oplus \left(\omega_2(i_4) \otimes \omega_4(i_5)\right)\right) \otimes \omega_5(i_p)\right) \oplus \left(\omega_H(i_1) \otimes \omega_1(i_p)\right)$$

The confidence of H in the reliability of P is calculated by forming a consensus of the two recommendation branches headed by authorities 1 and 2. Notably, a consensus of the opinions of authorities 3 and 4 about authority 5 has to be calculated before the recommendation of the latter is taken under consideration. Otherwise, 2's opinion about P would be biased as 5's recommendation about P would be counted twice.

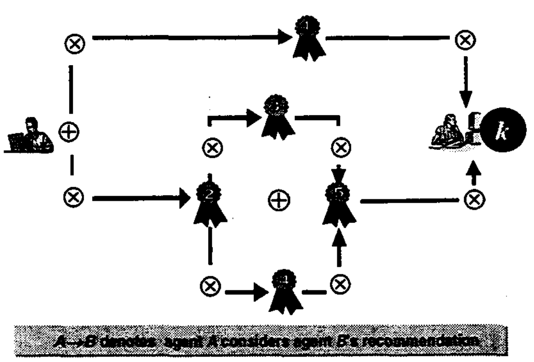

Figure 2: Authentication graph

H to 1	$\omega_H(i_1)$	H to 2	$\omega_H(i_2)$
1 to P	$\omega_1(i_p)$	2 to 3	$\omega_2(i_3)$
2 to 4	$\omega_2(i_4)$	3 to 5	$\omega_3(i_5)$
4 to 5	$\omega_4(i_5)$	5 to P	$\omega_5(i_p)$
P about k	$\omega_p(k)$		

Figure 3: Opinions

5.3 Integration scenario

None of the above mentioned formalisations is rich enough to capture the mixture of measuring trust in the presence of uncertainty, on the one hand, and relating trust to commitment or expectation to exhibit a prescribed behaviour, on the other, that is prominent in our analysis.

However, it seems plausible that an integration of suitable formalisations, such as subjective logic (Jøsang 2001) as a belief system, the consequence relation of (Jones & Sergot 1996) as a means of imparting contextual information, and a state machine interpretation of deontic reasoning similar to (Meyer) in the spirit of (Jones & Firozabadi 2000) may provide the basis for building machine-assisted reasoning support for our model. As we see no fundamental incompatibility between these two reasoning systems we will further investigate integration scenarios in view of our analysis of trust relationships. In doing so, we have to extend subjective logic with operators to support forming and combining opinions about temporal, deontic and dynamic behaviour, which is not supported in its current form.

A defect that will survive such integration is that one cannot guarantee that users will assign values appropriately and therefore cannot ensure determining trust consistently. However, analogous problems of determining measures of probability are well recognised in risk analysis in the process industry and finance. In principle, this is a matter of achieving the right means of abstracting information from reality into a mathematical model. We claim that by using logical reasoning along side risk analysis, one can produce the right guidelines and metrics to make the logical models work. In the following section we provide evidence to support this claim while we expand our trust model by introducing a suitable trust management scheme.

6 A trust management scheme

Trust management aims to provide a coherent framework for determining the conditions under which a party A takes the risk to depend on a party B with respect to a service X for a specific period within a specific context, and even though negative consequences are possible. Increasing the levels of trust facilitates processes to become more efficient but also increases the risk of allowing for the exploitation of vulnerabilities. One would consequently aim, in principle, to maximise trust while minimising risk. Hence, trust management subsumes and relies on risk management:

1. One may employ tailored risk analysis in order to analyse environmental risks and assess the most

tangible aspects of trust (e.g. the dependability of the information technology infrastructure).
2. Risk management allows us to weight transaction risk against trust, evaluate the impact of a failure in trust and help device countermeasures.

Note that the above two analyse different types of risk. Trust management becomes more tractable in the presence of a conceptual classification of the different aspects of trust and the corresponding ways they influence behaviour. For this purpose, we have adapted the conceptual framework proposed in (McKnight 2000). Our adaptation extends the approach proposed in (Povey 1999) and includes the concepts summarised in Figure 4. In (Dimitrakos 2001) we elaborate on this trust management scheme emphasising the role of risk management as a means of controlling the transition from one layer of this classification to another.

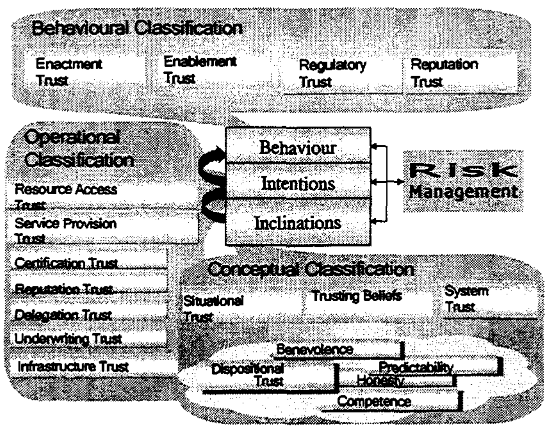

Figure 4. Overview of the trust-management scheme

7 Conclusion and further work

The pliability of the emerging communication media, the complexity of plausible interactions in virtual communities and the frequency of critical interactions among people who are relative strangers lead to problems that may not arise in traditional social settings. The same pliability allows an unprecedented degree of engineering and supports solutions to many of these problems. Effective solutions require interdisciplinary approaches requiring the integration of tools from cognitive sciences and economics in addition to telecommunications and computing.

In the this paper we introduced a rigorous model of trust in e-commerce and presented general properties of trust that underlie e-services, highlighting a role-based approach to the analysis of (unintentional) transfer of trust. We provided evidence of emerging methods, formalisms and conceptual frameworks, which can provide the basis for trust management solutions in e-markets. However, there is still a long way to go. Further work in the area of agent e-trading service provision and foreseen research challenges include:

- *To formalise and evaluate the proposed role-based model of trust in e-commerce.* Preliminary results are reported in (Dimitrakos & Bicarregui 2001).
- *To develop risk management techniques supporting the transition between trust inclinations, intentions and dependable behaviour.* An output is to produce practical guidelines for the attention of regulators and technology providers on how to maximise trust and minimise risk in different e-service scenarios.
- *To embody trust elements in contract negotiation, execution monitoring and arbitration.* This involves modelling legal issues concerning the status of electronic agents as participants in the process of contract formation. Preliminary results in this direction are reported in (Daskalopulu, Dimitrakos, Maibaum 2002).
- To experiment with developing a virtual marketplace from scratch, taking trust issues into account throughout the development lifecycle.

References

Blaze M., *et al. Proc. IEEE Conference on Security and Privacy*, Oakland, CA. May 1996.

Castelfranchi C., and R. Falcone. Principles of Trust for MAS: Cognitive Anatomy, Social Importance, and Quantification. In Proc. of the 3rd International Conference on Multi-Agent Systems, IEEE C.S., 1998.

CORAS Web http://www.nr.no/coras CLRC web- page: http://www.itd.clrc.ac.uk/Activity/CORAS

Damianou N., Dulay N., Lupu E., Sloman M. "The Ponder Policy Specification Language" Proc. *Policy 2001: Workshop on Policies for Distributed Systems and Networks*, Springer-Verlag LNCS 1995

Daskalopulu A., Dimitrakos T., Maibaum T. Evidence-based Electronic Contract Performance Monitoring. *INFORMS Journal of Group Decision and Negotiation, Special Issue: Formal Modeling of Electronic Commerce*, 2002.

Dimitrakos, T. System Models, e-Risk and e-Trust. In *Towards the E-Society: E-Business, E-Commerce, and E-Government*, (1st IFIP Conf. on e-business, e-commerce, e-government) Kluwer Ac. Pub, 2001.

Dimitrakos T., Bicarregui J. Towards A Framework for Managing Trust in e-Services. *4th International Conf. on Electronic Commerce Research*, ATSMA, IFIP, INFORMS vol. 2 p.360-381 ISBN 0-9716253-0-1

Dimitrakos T., Bicarregui J. Towards modelling e-trust. *Proc. 3rd Panhellenic Symposium on Logic*, 2001.

Firozabadi B. S., Sergot M. "Power and Permission in Security Systems". In *7th International Workshop In Security Protocols*, LNCS, Springer-Verlag, 1999.

Gambetta D., Can We Trust Trust? In *Trust: Making and Breaking of Cooperative Relations*. Basil Blackwell, Oxford, 1990.

Grandison T., Sloman M. "A Survey of Trust in Internet Applications" In *IEEE Communications Surveys and Tutorials*, Fourth Quarter 2000.

Jones A.J.I, Sergot M.J. A Formal Characterisation of Institutionalised Power. In *J. of the IGPL, 4(3)* 1996

Jones A.J.I., Firozabadi B.S. On the characterisation of a Trusting agent. In *Workshop on Deception, Trust and Fraud in Agent Societies*, 2000.

Jones S., TRUST-EC: requirements for Trust and Confidence in E-Commerce, European Commission, JRC, 1999.

Jøsang A. "A logic for uncertain probabilities". In *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, Vol 9. No. 3*, June 2001

Jøsang A. "An Algebra for Assessing Trust in Certification Chains". In *Proc. Network and Distributed Systems Security Symposium*. The Internet Society, 1999.

Jøsang A. The right type of trust for distributed systems. *New Security Paradigms Workshop*, ACM, 1996.

Kanger, S. "Law and Logic". In *Theoria, 38(105-132).*

Kini A., Choobineh J., "Trust in Electronic Commerce: Definition and Theoretical Consideration". Proc. *31st International Conf. on System Sciences*, IEEE, 1998.

Laprie J.C., *Dependability: Basic Concepts and Terminology*, Springer-Verlag, 1992.

Lindahl, L. Position and Change - A Study in Law and Logic, *Synthese Library112,* D.Reidel, Dordrecht, 1977

Luhmann, N. *Trust and Power*, New York: Wiley, 1979.

Marsh, S. Formalising Trust as a Computational Concept, PhD Thesis, University of Stirling, April 1994

Meyer Ch. A Different Approach to Deontic Logic: Deontic Logic Viewed as a Variant of Dynamic Logic. *Notre Dame J. of Formal Logic*, 29(1).

McKnight D.H., Chervany N.L. "What is Trust? A Conceptual Analysis and an Interdisciplinary Model". Proc. *the 2000 Americas Conference on Information Systems (AMCIS2000)*, August 2000.

Pörn, I. "Action Theory and Social Science". *Synthese Library 120*, D.Reidel Pub. Comp., Holland, 1997.

Pörn, I. "On the nature of a social order". In *Logic, Methodology and Philosophy of Science VIII*, J.E. eds. Fenstad et al., Elsevier Science Pub., 1989.

Povey D. "Developing Electronic Trust Policies Using a Risk Management Model". In *LNCS 1740*, 1999.

Shafer, G. *A Mathematical Theory of Evidence.* Princeton University Press, 1976.

Yahalom R., et al. Trust relationships in secure systems -A distributed authentication perspective. *IEEE Symposium on Research in Security and Privacy*, 1993.

A Negotiation Policy for Interacting Agents in e-Markets

Larbi Esmahi ; Petre Dini
Center for Computing & Information Systems
Athabasca University, 1 University Drive,
Athabasca, Alberta, T9S 3A3, Canada
Tel: +1 780 6756657 Fax: +1 780 6756186
Cisco Systems Inc./Concordia University
Internet OSS Group
725 Alder Drive, Milpitas, CA 95035, USA
Tel: +1 408-853-4050, Fax: 1-408-527-6351
E-mail: larbie@athabascau.ca; E-mail: pdini@cisco.com

Abstract

E-transactions via shopping agents constitute a promising opportunity in the e-markets. It becomes relevant what kind of information and what kind of exchange policies between competing sites can better model the e-market world. There are several steps for building e-business: (i) attracting the customer, (2) knowing how they buy, (iii) making transactions, (iv) perfecting orders, (v) giving effective customer service, (vi) offering customers recourse for problems such as breakage or returns, (vii) providing a rapid conclusion such as electronic payment. In the distributed e-market paradigm, these functions are abstracted via agents representing both contractual parts. The paper presents a policy-based negotiation considering the nature of dependency relationships between cooperating agents, focusing on contract negotiations between agents. The general procedure for mediating between interacting agents may be used by the agent managers to validate agent agreements based on agent dependencies.
Keywords: negotiation policy, agent dependency, e-market.

1 Introduction

The Internet introduced a set of facilities covering easy access to news, video-on-demand, news-on-demand, virtual shopping, electronic transactions, etc. Many shopping Internet centers are now currently in USA, with various specialties such books, flowers, wines, pizzas, etc. Other centers proposed to the users virtual negotiation on things to be sold, auction sites, Internet electronic withdraws, etc. Customers are facing multiple possibilities in performing auctions, due to the auction mechanism itself; in order to win a bid, a customer must offer the highest price. In the e-market era, the roles of smart agents represent the interests of different players willing to conclude or negotiate a contract.

The proposed protocols in game theory consider rational agents with definite targets for maximizing their profits without any consideration the history of the negotiation process. Moreover, it is assumed that agents have complete information on their partners (matrix of profits), while the agent dependency is totally ignored. For example, two on-line selling agents may exchange information on the selling side, as buyer agents can do the same. Equally, two selling agents may cooperate to sell items having shared profits as a common target [14]. The auction models can be applied only for marketing domain and they are vulnerable to collusion. Other issues may relate to the control on whether cooperating agents may form an alliance on the selling side to unfairly persuade a buyer agent, while a coalition of buyer agent may affect the offer strategies of selling agents. With disregard to the final issues, we consider that, in both cases, considering agent dependencies may optimize a common target. On the other hand, two cooperating agents (belonging to the opposite camps, i.e., sellers and buyers) involved into a negotiation process to achieve a deal may adopt a trade-off procedure to achieve their apparently opposed individual goals, but with a common target, i.e., making a profitable deal.

While e-market business follows a similar pattern of the human-based trading [15], special agent dependencies for e-transactions may influence the procedure to achieve a deal, and consequently, a transaction manager (trading supervisor) should consider these dependencies.

The paper presents a policy-based negotiation considering the nature of dependency relationships between cooperating agents, focusing on contract negotiations between agents. The general procedure for mediating between interacting agents may be used by the agent managers to validate agent agreements based on agent dependencies. Section 2 succinctly presents related topics commonly used in game theory, auction models, and contract protocols (how to get a contract or an item). Section 3 continues with interaction analysis by presenting the e-market bargaining process and identifying relevant types of dependency relationships between market agents. It evaluates goal satisfaction under weak or strong dependency as well as the associated costs. Section 4 defines a negotiation policy based on the agent dependencies refining the procedures for achieving a contract under different types of dependencies. The conclusion presented in Section 5 summarizes the proposal and pave the further way concerning the truth in the negotiation strategy and the use of temporal aspects on commitments and executions of contracts.

2. Current achievements

Previous work and significant achievements are reported on various related fields of research and concrete solutions. Hereafter, we present and evaluate some concepts developed in the game theory, auction models, and contract protocols.

2.1. Game theory models

Game theory address many aspects of the agents' interaction: contract elaboration, profit repartition and conflict resolution. Many negotiation models have been proposed in this topic [1, 2, 3, 4, 5, 6]. These models have some desirable properties, such as insuring the negotiation convergence, the Nash-equilibrium and the pareto-optimality. The main representative works in this domain are those presented by Zlotkin and Rosenschein [1, 3, 7]. The authors propose a formal model that allows agents to select the pareto-optimal solution that maximize their utilities. The agents communicate their desires explicitly by exchanging messages and may accept concessions that allow them to elaborate contracts that satisfy their goals. A contract may concern task repartition (Task Oriented Domains), utility value repartition (Worth Oriented Domains) or decision making on the next state of the environment (State Oriented Domains). Different types of contracts have been studied: pure contracts where the agent's role in the joint plan is fix and mixed contracts where the agent's role depends on a probability.

If we consider a negotiation between two agents A1 and A2, the authors propose a protocol that can be summarized as follow:

1. At each step t >= 0, both agents propose their deals $\delta_1(t)$ and $\delta_2(t)$ such that those deals satisfy two conditions: (1) the deals must be individually rational to their respective agents ($\forall$ A_i δ_i, the utility $U_i(\delta_i) >= 0$), and (2) for each $A_i \in \{A_1, A_2\}$, $\forall$ t>0 we have $U_i(\delta_i(t)) <= U_i(\delta_i(t-1))$.
2. The negotiation finishes at a step t when one of the two situations happens:
   * the agents agree on a deal. $\exists$ $i \neq j \in \{A_1, A_2\}$, such that $U_j(\delta_i(t)) >= U_j(\delta_j(t))$.
   * the agents run on a conflict. $\forall$ $A_i \in \{A_1, A_2\}$, $U_i(\delta_i(t)) = U_i(\delta_i(t-1))$ (i.e. no more concession is possible for both agents).

The advantage of the proposed protocols in game theory consists of their suitability for rational cooperating agents that work for maximizing their profits. However, the main drawbacks of those models consist of (i) their inability to take into consideration the history of the negotiation process, and (ii) in the fact that each step is processed as a stand-alone step. Furthermore, the agents are supposed to have complete information on their partners, especially by knowing all their matrix of profits. The agents are also supposed of being self-sufficient, while the complementarity and dependency between agents is ignored.

2.2. Auction models

Auction theory analyzes the protocols and strategies used by agents during an auction sale. Many protocols have been proposed in auction theory [8]:

English auction: in the English auction, the bidding process is public, so each bidder has complete information about the auction. Any time, each agent is free to raise his bid. When no bidder is willing to raise anymore, the auction ends, and the highest bidder wins the item at the price of his bid. The agent's strategy consists of a series of bids, where the bidding value is function of his private value, his prior estimates of other bidder's valuations, and the past bids of others. An agent's dominant strategy is to always bid a small amount greater than the current highest bid, and stop when his private value price is reached.

Sealed Bid Auction: in the sealed-bid auction, each bidder submits one bid without knowing the others' bids. The highest bidder wins the item and pays the amount of his bid. The agent's strategy consists of determining his bid as a function of his private value and prior beliefs of others' valuations. In general there is no dominant strategy for acting in this auction.

Dutch Auction: in the Dutch auction, the seller or the auction manager continuously lowers the price until one of the bidders takes the item at the current price. The Dutch auction is strategically equivalent to the sealed-bid auction, because in both games, an agent's bid matters only if it is the highest and no relevant information is revealed during the auction process.

Vickrey Auction: the Vickrey auction is similar to the sealed-bid auction at some detail exception. In fact, each bidder submits one bid without knowing the others' bids and the highest bidder wins, but it pays only the price of the second highest bid. The agent's strategy consists of determining his bid as a function of his private value and prior beliefs of others' valuations. The dominant strategy in Vickrey auctions is to bid one's true valuation. If an agent bids more than that, and the increment hits the difference between winning or not, he will end up with a loss if he wins. If he bids less, there is a smaller chance of winning, but the winning price is unaffected. The dominant strategy result of Vickrey auctions means that an agent is best off bidding truthfully no matter what the other bidders are like: what are their capabilities, operating environments, bidding plans, etc. This has two desirable sides namely, (i) the agents reveal their preferences truthfully, which allows globally efficient decisions to be made, and (ii) the agents need not waste effort in counter-speculating other agents, because they do not matter in making the bidding decision.

The auction models have two main drawbacks: they can be applied only for marketing domain and they are vulnerable to collusion. In the context of their vulnerability to collusion, the English auction presents the worst example. In fact, in English auction the biding process is public and the agents can verify if the collusion is respected or not and in case of the collusion broke the agent can use all his purchase power without loosing anything. In the other protocols (Sealed-bid auction, Dutch auction and Vickrey auction), the agents don't receive any information on the auction process and then the collusion formation is not so evident.

2.3. Contract-Net protocol

In the contract-net protocol (CNP) [9, 10], a contract is an explicit agreement between an agent who announces its tasks (the manager) and an agent who proposes to achieve those tasks (the contractor). The manager is responsible for the tasks management and results processing, while the contractor is responsible for the tasks processing and results transmission. The protocol proposes a set of performatives that the agents can use either for the contract negotiation and the contract execution. The following diagram summarizes the process for elaborating contracts.

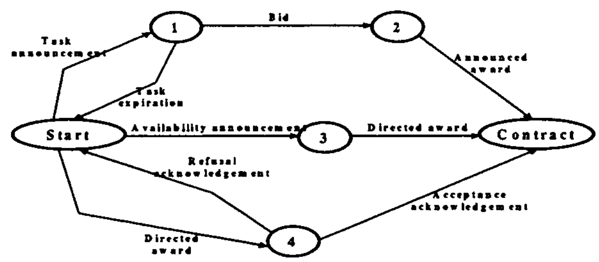

Figure1: Contract-Net Protocol

In this first version of the CNP as presented here above, there's no mechanisms for evaluating offers. This means that the protocol can be used only in contexts with cooperative agents and modular tasks.

Many extensions have been proposed for the CNP in the last years. One of the most important works done in this topic is the extension proposed by T. Sandholm & V. Lesser [11, 12], which is based on marginal costs. In fact, agents process locally their marginal costs for achieving tasks and use those costs either for announcing a task, or bidding, or awarding a contract. Partners' selection is based on profit evaluation using the marginal costs. Hence, extending the protocol's domain of application to competitive and self-interest agents is required.

The CNP is one of the most complete models for cooperating agents. The notion of profit and utility recently added to it, resolves the problem of motivation for competitive agents. However, the protocol still has some drawbacks, such as considering the agents as passive actors because no persuasion mechanisms can be used to convince partners for accepting an agreement. Also, dependency relations between agents are not considered.

3. Interaction process analysis in e-markets

The interaction in e-markets may be analysed according to different views. Since our concern is about contract negotiation, the main interaction to analyse in this context is the bargaining process inside the market place. In this context, we will introduce the dependency relations that exist between agents with respect to contract achievements.

in compensation to the first agent. The value of exchange must be the same, when measured in terms of lost or gain.

The exchange contract becomes more interesting with agents that have mutual dependencies. A general form of the exchange contract is where more than two items are exchanged between agents.

Grouping contract: This type of contract is built by grouping a set of items in the same contract, in such a manner that the global cost of the contract becomes profitable for the two agents, even if the individual cost of some items is greater than the accepted price. The grouping contract becomes more interesting with agents that have multi-dependency relation. The interest for this type of contract grows if we remember that in real trading, the price of a group of items is always lower than the sum of their individual prices.

Multi-agent contract: In this type of contract, a group of agents take a commitment to buy/sell a set of items. The contract is built in such a manner that is profitable for all participating agents. A multi-agent contract may include exchanging and grouping of items.

4.2. Generating contracts based on agent dependencies

In order to achieve the mediation for conflicts and generate agreements, the manager uses the classes of dependencies and contracts defined above. When mediation is requested, the manager first searches for mutual dependencies between the two agents, in order to construct a solution using an exchange contract. If no mutual dependencies are found, the manager then searches for multi-dependencies in order to construct a grouping contract. In the case where neither mutual dependencies nor multi-dependencies can be found, the manager uses the graph of dependencies to construct multi-agent contract.

A general algorithm for mediation can be summarized as following:

1- construct and propose a exchange contract using mutual dependencies.

1.2- search the graph for all mutual dependencies of the buyer towards the seller

1.3- sort the dependencies based on their strength

1.4- construct an exchange contract using the mutual dependency

1.5- propose the exchange contract

1.6- if the contract is accepted; then success, then stop

1.7- else select the next mutual dependency in the list and go to step 1.4

1.8- if the set of mutual dependencies is exhausted without success then go to step 2

2- construct a grouping contract using multi-dependencies.

2.1- search the graph for all multi-dependencies of the seller towards the buyer

2.2- sort the dependencies based on their strength

2.3- select the first dependency in the list

2.4- construct a grouping contract based on the multi-dependency relation

2.5- propose the grouping contract

2.6- if the contract is accepted; then success, then stop

2.7- else select the next multi-dependency in the list and go to step 2.4

2.8- if the set of multi-dependencies is exhausted without success, then go to step 3

3- construct a Multi-agent contract using other agents dependencies

3.1- search the graph for a balanced cycle of dependencies that closes the path between the buyer and the seller

3.2- construct a Multi-agent contract using the cycle of dependencies

3.3- propose the Multi-agent contract to the group of agents

3.4- if all agents agree on the contract then success, stop

3.5- else go to step 3.1

3.6- if there is no cycle of dependencies between the buyer and the seller then go to step 4

4- use other creative process to generate a solution (e.g., Case-Based reasoning, etc.).

5. Conclusion and open problems

The problem of negotiation policy in E-markets is not so simple! The existing protocols and strategies don't accommodate the real world. The virtual community, E-commerce and E-market are the real world. Many other factors that influence the agents' negotiation must be taken into consideration: the agent's mental state, the dependency relations between negotiators and the degree of truth that the agent grants to his partners.

The main idea we are pursuing in this research consists of integrating to the gents' mental state a model of their world especially a model of their partners and to use this knowledge in the negotiation strategy.

Many research issues are open for exploration in this project:

1. Exploring the ways of integrating the truth component in the negotiation strategy. Two ways may be

considered: representing the truth as knowledge in mental state of the agent and using this knowledge for negotiation or representing truth as a value (may be fuzzy value) and using this value in the agents utility function. The utility function is the base for evaluating offers, constructing counter-offers and determining concession to do in each negotiation step.

2. Extending negotiation model based on utility function by integrating the notion of time, temporal utility functions and temporal penalty functions in order to accommodate delays in executing contacts and commitments and also in order to accommodate the real behaviour in market place for prices and values.
3. Exploring the use of social norms as arguments for convincing a partner to accept an offer or to change his position in a conflict. Social norms may be defined as some standard community rules or the best is that social norms may emerge from the community interaction. In the last case, a hierarchical structure will be established for the emerging norms and a norm can move from one level to another according to the score granted by the community to this norm. The power of the norm as argument depends also on the level to witch this norm belongs. Agents must not only respect norms in their interaction or propositions but also make sure that norms are respected inside the community. (Those are socially committed agents!)

References

[1] J. S. Rosenschein, M. R. Genesereth. Deals among rational agents. In Proceedings of the Ninth International Joint Conference on Artificial Intelligence, pp: 91-99, Los Angeles, California, August 1985.

[2] M. R. Genesereth, M. L. Ginsberg, J. S. Rosenschein. Cooperation without communication. In Proceedings of the Fifth National Conference on Artificial Intelligence, pp 51-57, Philadelphia, PA, August 1986.

[3] G. Zlotkin, J. S. Rosenschein. Cooperation and conflict resolution via negotiation among autonomous agents in noncooperative domains. IEEE Transactions on Systems, Man and Cybernetics, 21(6), pp:1317-1324, November/December 1991.

[4] T. Khedro, M. R. Genesereth. Progressive negotiation: A strategy for resolving conflicts in cooperative distributed multidisciplinary design. In Proceedings of the Conflict Resolution Workshop, IJCAI-93, Chambery, France, September 1993.

[5] E. Ephrati, J.S. Rosenschein. Constrained intelligent action: Planning under the influence of a master agent. In Proceedings of the Tenth National Conference on Artificial Intelligence, pp:263-268, San Jose, California, July 1992.

[6] S. Kraus, J. Wilkenfeld and G. Zlotkin Multiagent Negotiation Under Time Constraints, Artificial Intelligence journal, Vol:75, No 2, pp : 297-345, 1995.

[7] J. S. Rosenschein, G. Zlotkin. Rules of Encounter. MIT Press, 1994.

[8] E. Rasmusen. Games and Information. Basil Blackwell, 1989.

[9] R. G. Smith, R. Davis. Frameworks for cooperation in distributed problem solving. In IEEE transactions on systems, man and cybernetics, pp : 61-70, N°1, Vol 11, 1981.

[10] R. G. Smith. The Contract Net Protocol : high-level communication and control in a distributed problem solver. In Readings in distributed artificial intelligence (A.H. Bond & L. Gasser), pp : 357-366, Morgan kauffmann publishers, 1988.

[11] T. W. Sandholm. An implementation of the Contract Net protocol based on marginal cost calculations. In proceedings 11th national conference on artificial intelligence, pp :256-262, 1993.

[12] T. W. Sandholm. Negociation among self-interested computationally limited agents. PhD Thesis, University of Massachusetts, 1996.

[13] L. Esmahi, P. Dini, J.C. Bernard. MIAMAP : A virtual market place for intelligent agents. The 33rd Hawaii International Conference on System Sciences (HICSS-33), Maui, Hawaii, January 4-7, 2000.

[14] O. A. Bodale, P. Dini. Using Nomadic Code in Negotiating e-Transactions, The 8th IEEE International Conference on Telecommunications, (ICT 2001), Bucharest, Romania, June 3-7, 2001

[15] J. Kong. http://techdivas.com/outside_is_in.htm, E-business: The Outside is now In, Diva Networks

COULD USERS BE ASSOCIATED WITH SOFTWARE AGENTS IN E-COMMERCE?

Zakaria Maamar; Fahim Akhter
Software Agents Research Group@ZU
College of Information Systems
Zayed University
PO Box 19282, Dubai
United Arab Emirates
zakaria.maamar@zu.ac.ae; fahim.akhter@zu.ac.ae

Abstract

This paper aims at answering three questions related to associating users with software agents. This association applies to e-commerce field in which users access online several products and services. Such users delegate their transactions to agents that act on their behalf. Why, How, and When this delegation occurs is discussed in this paper.

1. Overview

In one of their papers (Froehlich et al. 1999), the authors asked the following question: what makes e-commerce different from normal commerce? In fact, their question was motivated by the claim that since computers mediate currently almost all business transactions, all commerce is electronic. Garry et al. answered that e-commerce is not just the presence of computers or absence of papers, but it implies more such as:

- Using a non-proprietary open network along with issues of security and reliability.
- Servicing 24 hours a day, 7 days a week along with requirement of reliability.
- Customers and suppliers can be geographically distributed worldwide.
- A change in the relationships between trading parties. The role of 3rd party brokers changes. They must provide value-added services if they do not want to be bypassed.
- Establishing the identities of parties without requiring physical contact.

In our research work, we aim at understanding the value-added of Software Agents (SAs) (Jennings and Sycara and Wooldridge 1998) to e-commerce systems. We study why, how, and when users can entrust a part of the e-commerce operations they undertake to SAs. Usually, most of these operations are complex and though repetitive with a large segment suitable for computer aids and automation. In addition, users are already overwhelmed with information that needs to be filtered and sorted out before these information can be used efficiently and effectively. To assist such users in e-commerce operations, we suggest associating users with SAs and decomposing an e-commerce scenario into three phases: investigation, negotiation, and settlement. This is illustrated in Figure 1 where texts in italic summarize the operations and their outcomes.

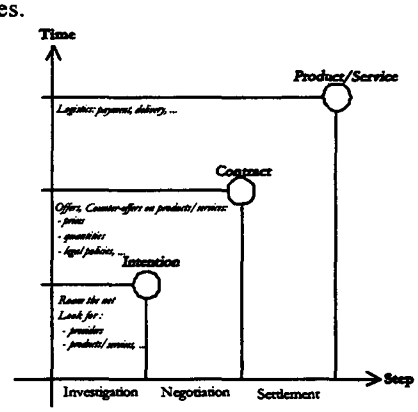

Figure 1 Proposed chronology of operations in an e-commerce scenario

2. Steps in e-commerce

Investigation, negotiation, and settlement phases are explained as follows (cf. Figure 1).

- Investigation phase: the purpose is to look for the providers that have the products/services that satisfy users' needs. To this end, agents assist users.

These agents roam networks on behalf of users, identify the providers that interest their users, and exchange messages with these providers about different matters, such as prices, availability, and return policies. Finally, the agents report their findings to their users. If users approve their agents' suggestions, agents are mandated to inform the potential providers about their intentions to request their products/services. The investigation phase outcome is an intention to request products/services.

- Negotiation phase: the purpose is to trigger an offer and counter-offer process, if needed. This process concerns the requested products/services and applies to their prices, quantities, payment conditions, return conditions, usage restrictions, delivery, and last but not least legal policies. The outcome of the negotiation phase is a contract; it documents the agreed upon obligations and commitments of both users and providers. SAs and users have to be tightly involved in negotiations. In fact, agents need to get their users' agreements before making decisions.
- Settlement phase: the purpose is to execute the contract and the clauses this contract contains. The operations that implement the clauses are initiated, for instance shipping and delivering products, notifying users, paying for the received products/services, etc. In the settlement phase, agents follow-up the progress of these operations, according to the agreed upon schedule. In case of delays, users have to be notified and corrective actions need to be taken. The settlement phase outcome is the product/service delivery.

Table 1 illustrates when users/SAs combination is desired. We advocate that the negotiation phase is the one in which users and agents have to be "tied" together. Although a SA has to be autonomous, it has to inform its user about the progress of its negotiation, for follow-up and quality control purposes. If a negotiation strategy is not followed appropriately, it can "harm" the user. Supporting electronic negotiation is one of the MEMO (for MEdiating & MOnitoring electronic commerce) project's recommendations to the expansion of e-commerce (MEMO 2000).

Table 1 Desired combination of users and SAs

Step	Combination	Outcome
Investigation	Agents	Intention
Negotiation	Users & Agents	Contract
Settlement	Agents	Product/Service

In an e-commerce scenario, SAs carry out multiple of operations that are listed as follows:

- Investigation phase:
  - Build users' profiles based on their needs and interests;
  - Map users' needs into requests;
  - Suggest modifications to users' requests;
  - Recommend products/services based on users' requests and agents' experiences;
  - Be aware of the market trends for follow-up and notification purposes;
  - Recommend alternate products/services in case of the investigation phase fails;
  - And, compare products/services.
- Negotiation phase:
  - Keep track of changing negotiation conditions;
  - Recommend appropriate negotiation strategies based on users' requirements;
  - Suggest to users to relax/not-relax certain constraints in case of the negotiation phase fails.
  - Switch from one provider to another during negotiations, while retaining negotiation contexts;
  - And, compare negotiations' results.
- Settlement phase:
  - Pay for the agreed upon products/services;
  - Enforce the clauses of the signed contracts;
  - Ensure that the agreed upon products/services are delivered;
  - Notify users in case of delays and assess the consequences of these delays;
  - And, suggest corrective actions.

A SA is usually defined as an autonomous entity that acts on behalf of user, satisfies his needs, and interacts with other agents when deemed appropriate. Several attributes make the difference between an agent and other components, such as objects. These attributes are: autonomy, goal-orientation, collaboration, flexibility, self-starting, temporal continuity, character, communication, adaptability, and mobility. Another argumentation about the important role of SAs in e-commerce scenarios can be suggested using these attributes. Table 2 depicts this argumentation where rows correspond to the agent's attributes and columns correspond to the three phases of an e-commerce scenario. The following notation is used: "2" stands for strongly appreciated, "1" stands for weakly appreciated, and finally "0" stands for not needed.

Table 2 Attributes of a SA vs. phases of an e-commerce scenario

	Phases

SA's attributes	Investigation	Negotiation	Settlement
Autonomy	2	2	2
Goal-oriented	2	2	0
Collaborative	1	2	0
Flexible	1	2	0
Self-starting	0	2	2
Temporal continuity	0	0	0
Character	0	0	0
Communicative	2	2	1
Adaptive	2	2	0
Mobile	2	0	0

In Table 2, the autonomy attribute is strongly appreciated in both investigation and negotiation phases and weakly appreciated in the settlement phase. For instance, a SA can make decisions when it searches for products and negotiates with these providers' services. However, the agent autonomy is less needed in the settlement phase. At this phase, the agent has a monitoring role. Regarding the mobility attribute, it is strongly appreciated that SA has to be embodied with this attribute at the investigation phase. In fact, the SA may have to roam networks of providers, visiting several sites. However, this attribute is not needed whether at the negotiation phase or at the settlement phase.

3. Software agents and types of e-commerce exchange

In (Guttman and Moukas and Pattie 1998), the authors did a survey on software-agent systems in the context of consumer buying behavior. Different systems were cited on the basis of the stages that can constitute a typical e-commerce scenario. These stages are needs identification, product brokering, merchant brokering, negotiation, purchase and delivery, and finally product service and evaluation. These stages are similar to the ones we presented in the beginning of this paper (cf. Figure 1). However, the authors did not discuss the role of agents according to the type of exchanges in which consumers can be involved. This is what we aim at doing below. Five types of exchanges exist (Liand and Huang 2000): bartering, bargaining, bidding, auctioning, and clearing.

According to (Liand and Huang 2000), customers (i.e. users), providers, and brokers can be involved in different types of e-commerce exchange. Each type has its set of operations to be carried out. A well-known type of exchange is the bilateral, i.e. users and providers interact directly. Trilateral exchanges are possible and involve brokers that support in diverse ways both users and providers. An overview of bilateral and trilateral exchanges is available in (Liand and Huang 2000). Bartering, bargaining, and bidding are bilateral exchanges. Auctioning and clearing are trilateral exchanges. Below, we discuss in which type of exchange agents can be considered relevant for e-commerce. Bartering is not discussed.

- Bargaining: involves a user that negotiates with a provider until an acceptable agreement for both is reached. First, the user looks for a provider, consults his products/services, and negotiates with him for an agreement. If negotiations fail, the user searches repeatedly for other providers till an agreement is reached with one of them. Agents would be suitable for looking for providers, negotiating on behalf of users, and saving the interactions their users have had with previous providers for evaluation purposes.
- Bidding: implies that one user and several potential providers participate. First, the user calls for bids. Next, he compares the offers he receives from providers. Finally, the user selects the best offer, i.e. the lowest offer. Agents would be suitable for looking for providers, initiating bids, accepting bids from providers, comparing bids, and notifying the winner provider.
- Auctioning: involves one provider, several potential users, and a broker. Successively, users make an offer on the product/service to be provided. First, a provider fixes the lowest price of the product/service to be auctioned. Through the Broker, the provider advertises the product/service and calls for auctions. Afterwards, the users make offers to the Broker. Finally, the broker selects the user who made the highest offer regarding the initial-provider's offer[1]. Agents would be suitable for finding the broker, monitoring users' offers, sending offers to the broker, and following-up the auctioning progress on behalf of users.
- Clearing: requires that several users, several providers, and one broker take part. Users and providers submit their requests to the Broker, regarding their needs and offers respectively. Next, the Broker matches needs with offers. If there is a success match, the Broker informs both users and providers about this match. Agents would be suitable for finding the broker, monitoring users' and providers' offers, and sending offers to the broker

[1] The described auctioning is called English. Other types of auctions exist and are Dutch and Vickrey.

auctioning is provider-driven, and cleaning is broker-driven. By user-driven, we mean that users have the initiative to lead the entire e-commerce scenario. However, users have less "opportunity" to regulate the progress of the e-commerce scenario in provider/broker-driven situations. Thus, SAs would be more suitable for user-driven situations, i.e. in term of roles to play, rather than for provider/broker-driven situations. In user-driven situations, user-agents would be actors. However, in provider/broker-driven situations user-agents would mainly be spectators. Our actor/spectator comparison is supported by the work done in (Arpinar and Dogac and Talbul 2000) in term of active and passive role.

In order to understand how SAs fit in the user/provider/broker-driven classification, we suggest representing the agent involvement (cf. Table 3). The scale is: 0 for not involved, 1 for weakly involved, and 2 for strongly involved. As stated above, user-agents are strongly involved in bargaining and bidding situations. However, they are weakly involved in auctioning and cleaning situations. In these situations, the functionalities that user-agents used to carry out are now entrusted to broker-agents. Among these functionalities, we cite looking for providers and matching users' needs with providers' offers.

Table 3 Agent involvement in e-commerce exchanges

	User-driven	Provider-driven	Broker-driven
Bargaining	2	1	0
Bidding	2	1	0
Auctioning	1	2	1
Clearing	1	1	2

4. Summary

Understanding the role of SAs in e-commerce is an important step towards the current trend of agent-oriented e-commerce development that sees a wider deployment of SAs for the following uses (Chong and Heng and Liu 2001):

- SAs will help customers to identify and locate the products or services that they require.
- SAs will help to keep track and inform customers of new offers that match their preferences.
- SAs will help customers to negotiate electronically with the SA sellers in order to buy and sell goods or services that is in the best

or services on their behalf.

References

S. Arpinar, A. Dogac, and N. Talbul. An open electronic marketplace through agent-based workflows: MOPPET. *International Journal on Digital Libraries*, 3(1) : 36-59, 2000.

S. Chong, M. S. H. Heng, and K. Liu. The social consequences of using agent-based systems. in *Proceedings of the Pacific Asia Conference on Information Systems (PACIS'2001)*, Seoul, Korea, 2001.

G. Froehlich et al. Application framework issues when evolving business applications for electronic commerce. in *Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS'32), minitrack on electronic commerce technologies*, A. Tsalgatidou and M. Papazoglou editors, Hawaii, 1999.

R. H. Guttman, A. G. Moukas, and M. Pattie. Agent-mediated electronic commerce: a survey. *Knowledge Engineering Review*, June 1998.

N. Jennings, K. Sycara, and M. Wooldridge. A roadmap of agent research and development. *Autonomous Agents and Multi-Agent Systems*, 1(1): 7-38, 1998.

MEMO-Consortium. MEMO home page. http://abnamro.com/memp. Visited September 2000.

T. P. Liand, and J. S. Huang. *A framework for applying intelligent agents to support electronic trading*. Decision Support Systems. 28(4), June 2000.

Towards an Active Network Architecture

D.L. Tennenhouse and D.J. Wetherall

Presented by
Pakorn Waewsawangwong

Objectives

- Envision an active network architecture.
- Outline an approach to its design.
- Survey the current related technologies.

What's the active network?

- A network that allows users to inject customized programs into network nodes.
- It is *active* in 2 ways:
  - Switches perform computation on the user data flowing through them.
  - Users can inject programs into the network.

Approaches to active network

- Two approaches: discrete and integrated (capsule).
- Capsule approach is preferred.
  - Every message is a capsule.
  - Capsule = program fragment + (embedded data)

Capsule Processing

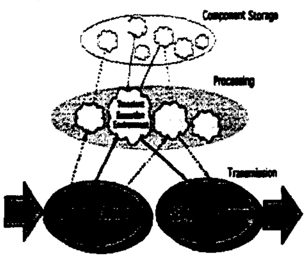

Scenario:

1. Framing identifies capsule boundaries.
2. Capsule contents are dispatched to TEE.
3. TEE may access external resources.
4. Execution results (0..* capsules) are scheduled and transmitted.
5. TEE is destroyed.

Towards ... an Interoperable Programming Model

- To be interoperable, computational model needs to be standardized.
- It's standardized in 2 aspects:
  - Capsule Programs (instruction set)
  - Node Resources available to capsule programs.

Node Resources (1)

Goal 1: *Interoperability*

- Solution: Resource is specified with a small set of platform independent abstractions:
  - Transmission Bandwidth
  - Instruction Execution (CPU)
  - Transient Storage
  - Logical Resources
  - Active Storage

Node Resources (2)

Goal 2: *Resource Safety*

- Solution: use 3 types of safety activity:
  - Dynamic Assignment of resources to specific capsules.
  - Validation of user requests for resources assignment, through authentication and verification of their authorizations.
  - Automated Delegation of resource authorizations.

Architectural Considerations

- How is interoperability achieved?
- Isn't the trend towards less functionality in the network?
- What's the impact on the layered reference model?
- What about the end-to-end arguments?

Mapping Service Level Agreements to the Adaptive Management of DiffServ Networks
Version: 0.1

1. Background

A SLA is a contract between a customer and a service provider that specifies the forwarding service. SLAs can be either static or dynamic. An important distinction between static and dynamic is that static SLAs are based on a negotiation between human agents-that is, between a customer and service provider; dynamic SLAs, on the other hand, change without human intervention and therefore require an automated agent [RFC 2475]. A second distinction can be made among quantitative, qualitative and relative service models. For example, the service definition related to packet-loss ratio could be as follows:
Qualitative service: Low loss ratio.
Quantitative service: Less than 2% packet-loss ratio.
Relative service: Packet-loss ratio on service level A is smaller than on service level B.

Part of the [RFC 2475] assumes that service differentiation definitely means different levels of quantitative service.

In his book on Differentiated Services [Killki], Killki presents four primary categories for a SLA between a customer and his/hers service provider:

- ✓ Guaranteed connections (dynamic bandwidth)
- ✓ Leased-line connections (permanent bandwidth)
- ✓ Dynamic importance (dynamic precedence)
- ✓ Resource sharing (permanent share)

These four approaches for SLA are shown in Figure1.

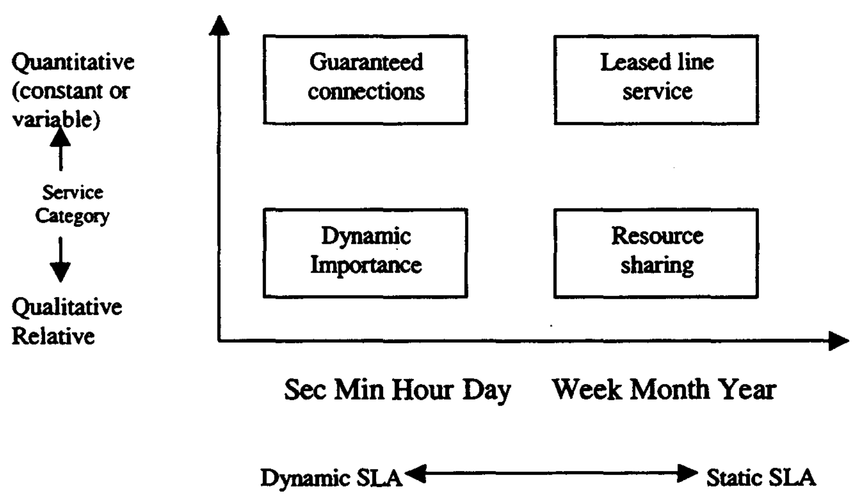

Figure1 : Four main SLA categories

2. Related Work

Current research in the Policy-based management area is concentrated on SLA specification, assuming as in RFC 2475 that the type of SLA is static.

2.1 Work within the TEQUILLA project (George Pavlou)

2.2 Dinesh Verma's work

3. Ponder SLAs

A SLA is as an aggregation of SLOs (or SLSs), where each SLO is a performance objective for a particular service or a group of services. This objective can be seen as a rule that governs the behavior of the managed system, i.e. it specifies how the managed system must treat the packets that belong to each service. Therefore, we can represent an SLO as an obligation rule, which configures the network in order the later to provide the requested QoS characteristics to each relevant service.

The general representation of a SLO within the Ponder framework can be the following:

```
inst oblig /Policies/SLAPolicies/SLAPolicy1/ServiceLevelObjective1
{
subject      /SLA_Agents/SLA_Agent1;
on     RequestForSLOEnforcement (SLO_parameters[]);
do     DiffServPolicy = select (SLO_parameters[]);
       DiffServPolicy's_parameters[] = calculate ( SLO_parameters[]);
       DiffServPolicy.enable();
       EventService.GenerateEvent         (DiffServPolicy'sObligationEvent,
       DiffServPolicy's_parameters[]);
}
```

Figure X: Specification of a generic SLO policy

Explanation:

As in the Policy 2002 paper, we propose two levels of policies. SLA management then maps to the adaptive management of DiffServ by selecting and enabling lower-level network QoS policies upon a request for enforcement of a specific SLO. The higher level is constituted by SLO policies, which are interpreted by higher level Policy Management Agents (/SLA_Agents). In the lower level, there are DiffServ QoS policies (e.g. a policy that configures the network to provide Expedited Forwarding service or alternatively a PDB policy for Premium Service).

The interaction between these two levels of policies is presented in the following figure.

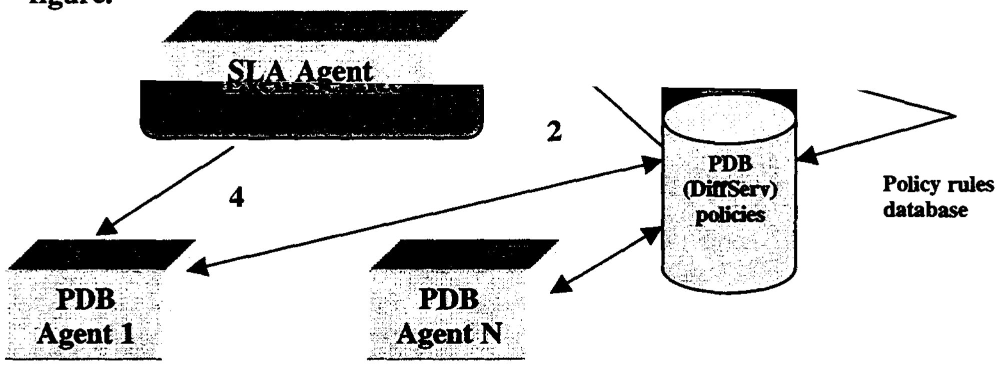

Figure X: Enforcement Architecture

1. **The SLA Agent receives an event that requires the enforcement of the SLO policy, which is loaded and enabled to this agent.**
2. **The SLA Agent decides at run-time which DiffServ policy must be chosen from the set of DiffServ policies. The selected DiffServ policy will fulfil the service's need, i.e. the DiffServ policy will actually impose the new network configuration. The selected DiffServ policy is then enabled to the relevant PDB Agent (s).**
3. **The DiffServ policy's obligation event with the calculated parameters is sent to the Event Service.**
4. **The latter obligation event is forwarded to the corresponding to the DiffServ policy PDB (DiffServ) Agent. Upon receiving this event, the PDB Agent will impose the new network configuration to the target devices (DiffServ policy's target domain, e.g. /Routers/CoreRouters/Domain1).**

In the following, we will present some examples that demonstrate how this framework can cater for different SLA types, as presented in Section 1.

3.1 Static SLAs

An example of a static SLA is one that guarantees to the customer a Leased Line Service (Figure 1). A leased line service is a permanent guaranteed connection, which is independent of the actual use of the service. This type of service between a customer and an ISP provider can be implemented with the DiffServ technology, as the capacity reservations are permanent enough (change every month or more), so there is no need for signaling.

Leased line service is implemented in DiffServ networks with the Expedited Forwarding (EF) PHB. The SLO Policy that provides the customer this type of network behavior can be the following:

allocation (existing EF allocation plus the application's requested bandwidth), etc.

Good Faith and Fair Dealing in the Context of Electronic Agents

Emily M. Weitzenböck
Norwegian Research Centre for Computers & Law, Faculty of Law, University of Oslo, P.O. Box 6706, St. Olavs plass, 0130 Oslo, Norway
emily.weitzenboeck@jus.uio.no

Abstract

The principle of good faith is an important guideline for contractual behaviour that permeates civil law systems. This paper examines how this principle is applied during the negotiation stage of a contract. Select examples from civil law literature of precontractual duties of good faith, and of precontractual behaviour that is deemed to be contrary to good faith, are discussed. This is followed by a discussion of the extent to which such duties are recognised, or such behaviour proscribed, in common law jurisdictions. Some common standards for precontractual behaviour in civil and common law systems are identified. These standards, in the situation where contracts are being negotiated by or through electronic agents, would need to be reflected in the way such agents behave.

1 Introduction

Electronic agents are gradually developing from mediators in the electronic marketplace into more active players that may initiate, negotiate and conclude contracts and implement the contractual obligations undertaken. Some of the issues being discussed in disciplines such as Computer Science, Cognitive Science and Logic are how such agents should behave in order to fulfil their contractual obligations. Central notions of trust, risk and reliability are vital for there to be confidence in such technology. Contractual behaviour is regulated, to a greater or lesser extent, by existing legal norms which impose duties, conditions or requirements on a contracting party. One important guideline for contractual behaviour (Lando, 2000) that permeates civil law systems, particularly in contract law, is the requirement that parties negotiate, conclude and carry out contracts in good faith.

1.1 Good Faith as a behavioural criterion

Broadly speaking, in continental civil law systems based traditionally on the Roman legal tradition, the juridical notion of good faith ('bona fides') is an important behavioural criterion to assess and regulate the conduct of the contracting parties. In common law literature, the term "fair dealing" is often used, rather than the term good faith.[1] Together with the norms which regulate the formation[2] and execution of contracts, one finds the good faith norms which lay down a general criterion of behaviour for the contracting parties, that is, that parties are expected to act in good faith towards each other in the negotiation, formation and execution of contracts. However, the extent and scope of application of the good faith principle varies from one civil law jurisdiction to another, and this may give rise to confusion as to what its essence is and what duties stem from it.

Moreover, in other legal systems such as English law, there is no general rule requiring the parties to conform to good faith. However, many of the results which in civil law systems are achieved by requiring good faith, have been reached in English law through a different, piecemeal approach.

1.2 Research Approach

The aim of this paper is to examine the principle of good faith and fair dealing as a guideline for contractual behaviour. It will try to map the constitutive elements of good faith and fair dealing by focusing on the stage of contract negotiation and formation, in order to determine what standards and norms of behaviour are expected from parties entering into a contract. It is presumed that where contracts are negotiated by or

[1] See further on this point, section 2.2 infra.

[2] For example, civil law systems generally recognise four essential elements of a contract: (i) capacity and (ii) consent of the contracting parties, as well as the requirement that a contract should have a definite object (iii) and a lawful (iv) "causa". See for example, article 1108 of the French Civil Code, and articles 1325 and 1425 of the Italian Civil Code of 1942.

The recognition of "obligations d'information" (obligations to inform or disclose) is another development related to good faith in French law. Similar duties to provide information also emanate in Italian contract law.

2.2 Is there a duty of good faith in common law countries?

An important principle that permeates both civil and common law contract law is the freedom of the parties to enter into contractual relations or to choose not to enter into contractual relations. Cohen (1995) describes the former as "the positive freedom of contract" in that the parties are free to create a binding contract reflecting their will, and the latter as "the negative freedom of contract" which means that the parties are free from obligations so long as a binding contract has not been concluded. However, in civil law countries, the negative freedom of contract is subject to the principle of good faith and other doctrines based on good faith such as the doctrine of abuse of right and unjustified enrichment.[5]

In English law, according to many writers, there is no general rule requiring the parties to negotiate in good faith (O'Connor, 1990; Whittaker, 2000b). This does not mean that there is a free-for-all, with no controls on contracting parties. The traditional rules proscribing duress, undue influence and fraud, still apply. Other than that, in English law, either party is entitled to break off negotiations at any stage before the final conclusion of the contract. Liability for pre-contractual behaviour is only imposed under limited circumstances such as fraudulent representation or negligent misstatement, as discussed more fully below.[6]

Although English common law recognised a principle of good faith in contractual dealings in the mid-eighteenth century,[7] the modern view is that in English law, good faith is, in principle, irrelevant. Bingham L.J. stated in Interfoto Picture Library Ltd v. Stilletto Visual Programmes Ltd (1989) that, "[i]n many civil law systems, and perhaps in most legal systems outside the common law world, the law of obligations recognises and enforces an overriding principle that in making and carrying out contracts parties should act in good faith. This does not simply mean that they should not deceive each other, a principle which any legal system must recognise; its effect is perhaps most aptly conveyed by such metaphorical colloquialisms as 'playing fair', 'coming clean' or 'putting one's cards face upwards on the table'. It is in essence a principle of fair open dealing ... English law has, characteristically, committed itself to no such overriding principle but has developed piecemeal solutions in response to demonstrated problems of unfairness."

Many English writers often use the term "fair dealing" rather than "good faith" on the basis that the latter term "nowadays has a fuzzy sound to the ears of English lawyers" (Harrison, 1997) and may appear as a vague literary concept. The term "fair dealing" connotes observance of fairness which appears as a more objective test to common law lawyers. However, it should be emphasised that this objective connotation is also included in the civil law notion of "good faith".

As regards the United States, the Uniform Commercial Code (UCC)[8] provides in section 1-203 that "[e]very contract ... imposes an obligation of good faith in its performance or enforcement." This is mirrored in §205 of the Restatement of Contracts Second[9] which states that "[e]very contract imposes upon each party a duty of good faith and fair dealing in its performance and its enforcement."

Good faith is defined in the UCC as "honesty in fact in the conduct or transaction concerned".[10] In the case of a merchant, the UCC provides that good faith means "honesty in fact and the observance of reasonable commercial standards of fair dealing in the trade."[11]

According to American jurists, similar to English law, the requirement of good faith in American law does not apply to contract negotiations.

3 Precontractual behaviour

The focus of this paper is the behaviour of the parties at the pre-contractual stage, that is, just before the parties conclude a contract or else decide not to enter into a contract.

As abovementioned, both civil and common law systems regard the freedom of the contracting parties as sacrosanct. Parties should be free to decide whether to enter into contractual relations or not. However, the question that arises is what happens where, because of certain blameworthy conduct of a contracting party at

[5] See further on this, O'Connor, 1990.

[6] See section 3.1.2 infra.

[7] Several legal writers make reference to the famous dictum of Lord Mansfield C.J. in 1766, that "[t]he governing principle is applicable to all contracts and dealings. Good faith forbids either party by concealing what he privately knows, to draw the other into a bargain, from his ignorance of that fact, and his believing the contrary. But either party may be innocently silent, as to grounds open to both, to exercise their judgment upon." See, for example, Beale, 1999.

[8] In 1960, the UCC was being introduced in, and adopted by the American state legislatures.

[9] The Restatement Second was introduced in 1979 with official promulgation in 1981.

[10] See §1-201(19) of the UCC.

[11] See §2-103(1)(b) of the UCC.

France has a highly developed notion of duties of disclosure which has evolved both through direct legislation[13] and through case law. A party should disclose only such information as is relevant, having regard to the subject matter of the contract and to the obligations undertaken by the parties. For instance, in a contact of sale of a machine, the seller should inform the buyer of the conditions of use of the machine; however, the same obligation does not arise for the person who repairs it. Duties of disclosure arise due to the parties' unequal information. However, a party's lack of knowledge should be legitimate. The Cour de cassation states that the contracting party who made a mistake by being too gullible or careless in checking some information has only himself to blame (Fabre-Magnan, 1995). A party is allowed to be unaware of information if it was impossible for him to know it or if he could legitimately rely on the information given by the other party.

Since in most cases where the duty to disclose was violated the other party, if correctly informed, would have abstained from entering into the contract, such party will not be able to recover the value of the promised performance (the expectation interest, i.e. the benefit anticipated) but the reliance interest (for relying on the validity of the contract).

3.1.2 The situation in common law systems

In common law, the situation appears to be different to that prevalent in civil law countries. As abovementioned, there is no duty to negotiate in good faith. However, Harrison (1997) explains that "the duty 'bites' when the contract is made. Something rather like a contractual Judgment Day occurs at this point, exposing the seller's evasions, pregnant half-truths and the like, and penalising him for failing in his duty of good faith by deeming the contractual provisions to have included the duty of good faith."

Although the main rule is "caveat emptor" (let the buyer beware) and thus "there is no legal obligation on the vendor to inform the purchaser that he is under a mistake, not induced by the act of the vendor" (Smith v. Hughes, 1871), certain implied warranties developed in English sale of goods law which, to a certain extent, have reduced the effect of "caveat emptor" in this area. As Whittaker et al (2000a) explain, "where English law sees the circumstances as ones in which one of the parties to the contract is in a position to know of the characteristics of the subject matter of the contract and the other is not, it may simply place the responsibility for those characteristics being present on the shoulders of the typically knowledgeable party by means of an implied term." Examples are the implied terms about title (that the seller has a right to sell the goods), and the implied terms about quality or fitness.[14]

[13] For example, to protect consumers of goods and services.

A distinction has been made between *silence* as such, i.e. mere nondisclosure, and *active suppression* which constitutes fraud. And, as Kessler (1964) notes, it has become increasingly difficult to draw a clear line between them. Kessler quotes Turner v. Harvey (1821) where it was held that "a very little is sufficient to affect the application of [caveat emptor] ... If a word, if a single word be dropped which tends to mislead the vendor, that principle will not be allowed to operate." As the New York case of Donovan v. Aeolian Co. (1936) further explains, "it depends upon the circumstances of each case whether failure to disclose is consistent with honest dealing. Where failure to disclose a material fact is calculated to induce a false impression, the distinction between concealment and affirmative misrepresentation is tenuous. Both are fraudulent." Moreover, facts which are true when said must be corrected if they have become untrue by the time the contract is entered into.

An action will lie in tort[15] for negligent misrepresentation causing loss to the representee where the relationship of the parties is such as to give rise to a duty of care (Hedley Byrne & Co Ltd v. Heller and Partners Ltd (1964)). A negligent misrepresentation is one which is made carelessly, or without reasonable grounds for believing it to be true.

Therefore, while a party who positively misleads the other party (even if innocently) will in principle be faced with rescission of the contract, a person who says nothing will be secure.[16] While parties to a contract should not mislead each other as to the subject matter of the contract, whether innocently or fraudulently, they should not in general have to act so as to protect the other's interests, but may act in their own interest.

[14] The common law position was enshrined in the Sales of Goods Act 1979 (amended by the Sale and Supply of Goods Act 1994) which provides in section 14(1) that, except as provided in that section, there is no implied term about the quality or fitness for any particular purpose of goods supplied under a contract of sale. However, section 14(2) then goes on to provide that "[w]here the seller sells goods in the course of a business, there is an implied term that the goods supplied under the contract are of satisfactory quality." The quality of the goods "includes their state and condition and ... fitness for all the purposes for which goods of the kind in question are commonly supplied." Moreover, according to Section 14(3) "where the seller sells goods in the course of business and the buyer, expressly or by implication, makes known ... to the seller ... any particular purpose for which the goods are being bought, there is an implied term that the goods supplied under the contract are reasonably fit for that purpose, whether or not that is a purpose for which such goods are commonly supplied ..." See further on this Harrison, 1997 and Zimmermann et al, 2000.

[15] I.e. outside contract, and irrespective of whether a contract will be concluded or not.

[16] See further on this discussion, Whittaker et al, 2000a, and Kessler, 1964.

However, the situation is rather different in the case of fiduciary contracts such as, for example, insurance contracts and agency.[17] In contracts of insurance, both parties have a duty to observe "the utmost good faith" towards the other party, which imposes a duty of disclosure. A special characteristic of insurance contracts is that they are based on facts and knowledge of facts which are almost invariably under and within the exclusive control of one of the parties (usually the insured), and a proper assessment of the risk could not be made unless there were full disclosure of all material facts. The insurer, very often, does not have independent means, or has impaired means, of finding out risk-relevant information, and would have to rely on disclosures made by the person applying for an insurance policy.

Similarly, in agency, an agent is expected to act with perfect good faith, requiring full disclosure of all the material circumstances, when he deals with his principle. He is required to make full disclosure of all material circumstances and the exact nature and extent of his personal interest in a transaction where this might conflict with his duty to his principal. Furthermore, he must not use his position or his principal's property to acquire benefits for himself.[18]

Harrison (1997) puts forward the thesis that the duty of good faith or fair dealing as it applies in the formation of contracts of sale, is normally a twin duty of *candour and accuracy*. This is the duty to give proper information or none at all about what is being sold in contracts outside the area of fiduciary contracts. Harrison states that this is a presumption of law and operates both as an obligation in interpreting the contract and as an additional implied term where there are no relevant express terms to be interpreted. She holds that it does not operate as regards matters which it would be normal and possible for the buyer to investigate himself. Most importantly, Harrison states that a pre-contractual breach of this duty ***has no effect unless a contract is made.*** Thus, the effect on the parties only occurs when a contract is made, but not if negotiations break down.

As regards remedies in English law, *where a contract was concluded*, before the Misrepresentation Act 1967, where a party was induced to enter into a contract as a result of a misrepresentation by the other party and the misrepresentation never became incorporated as a contractual term, the representee was entitled to rescind the contract whether the misrepresentation was fraudulent, negligent or wholly innocent. Beale et al (1999) explain further that at common law, the right to rescind was confined to cases in which the misrepresentation was fraudulent or in which there was a total failure of consideration, but in equity there was a right to rescind even for innocent misrepresentation. Since the Act was passed, the right of rescission is qualified (except in cases of fraud) by the court's power to refuse rescission and award damages in lieu, and there remain certain bars to recession, but this falls outside the scope of this paper.[19]

Where negotiations broke down and a contract was not concluded, a remedy in tort for fraudulent misrepresentation or negligent misrepresentation, may be available in the circumstances mentioned above.

3.1.3 Concluding remarks on the precontractual duty of disclosure

To sum up, in most civil law systems, there is a duty of disclosure where each party is bound to disclose such matters as are clearly of importance for the other party's decision, provided the latter is unable to procure the information himself and the nondisclosing party is aware of the fact.

As regards English law, legal writers hold that there is no such duty of disclosure at the stage of contract negotiation, save for fiduciary contracts as abovementioned. Where there has been fraudulent representation or negligent misstatement, a remedy would be available in tort. However, when negotiations have led to the conclusion of a contract, the silence of one party could be problematic for such party (who could be liable for damages and/or find the contract rescinded) where the information suppressed relates to a fact that is deemed to be an implied term. Hence the wisdom for negotiating parties, to act with "candour and accurately" (Harrison, 1997).

One may thus note the important role played by the applicable law that is used to determine a particular dispute because the rules and remedies vary as explained above, according to whether the applicable law chosen is that of a country which has a common law system or a civil law (and even in the latter, there are some differences among civil law countries, as seen above). However, some common threads may be identified and certain behaviour should be refrained from at the stage of negotiation, irrespective of whether a contract is later on concluded or not. Thus:

[17] Harrison (1997) also holds that there is a duty of disclosure by the debtor to his/her surety, i.e. where property is put up as security for another person's debt, in the case where the surety could not be expected to find out the information in question for himself or herself.

[18] For a further discussion of these fiduciary duties in agency law, see O'Connor, 1990.

[19] For a more detailed discussion of these bars – i.e. affirmation of the contract, lapse of time or the acquisition by a third party of the rights in the sub-matter of the contract - see Beale et al, 1999.

(1) one should refrain from fraudulent misrepresentation, that is where the party has actual knowledge of the incorrectness of a certain (material) fact or facts presented as true: This appears to be closely linked to the attribute of *veracity* discussed in (technical) agent literature (Wooldridge et al, 1995) which states the assumption that an agent will not knowingly communicate false information.
(2) one should refrain from negligent misstatement: each party should be careful not to make remarks on facts material to the goods which fact(s) turn out to be incorrect, where such remarks are made carelessly or without reasonable grounds for believing them to be true.

Put positively, the parties should behave, as Harrison (1997) opines, with *candour and accuracy* and give proper information about what is being sold in contracts outside the area of fiduciary contracts.

If a contract is concluded, then, if a certain feature of the subject matter of the contract is of paramount or fundamental importance to one of the parties, such party should make its existence an essential term or condition of the contract.

3.2 Sudden and unjustified rupture of negotiations

3.2.1 To what extent is an offer binding?

Before the conclusion of a final contract, each party is free to withdraw from the negotiations, each party bears his own expenses and acts at its own risk. However, as Cohen (1995) explains, strict adherence to freedom from contract might transform it into a freedom to manipulate the rules of the game. Freedom of action which is the underlying idea of freedom of contract, may be abused.

English law tends to enable the contracting party to act according to the rules of the game and to benefit from the freedom granted by these rules. Therefore, an offer promised to be irrevocable, is still revocable under English law because it is considered not to be supported by consideration and therefore not binding (Cohen, 1995).

Nor could one try to argue that though a contract was not concluded, there was *an agreement to negotiate* since a contact to negotiate is not recognised in English law. In Courtner & Fairbairn Ltd v. Tolaini Brothers (Hotels) Ltd (1975), such a contract was held to be unenforceable because "[i]f the law does not recognise a contract to enter into a contract (when there is a fundamental term yet to be agreed) ... it cannot recognise a contract to negotiate. The reason is because it is too uncertain to have any binding force. No court could estimate the damages because no one can tell whether the negotiations would be successful or would fall through: or if successful, what the result will be."

By contrast to English law, continental systems empower the offeror to bind himself by an irrevocable offer. As Cohen explains, civil law systems empower the offeror to bind himself/herself by an irrevocable offer, if he/she so wishes. "Negative freedom in the negotiations is not sacred: even a non-contractual promise or a mere expectancy may have a certain binding force under the doctrine of good faith in negotiations". Thus, in German law, unless the offeror by appropriate language has given fair warning to the offeree that the offer is not binding, any offer once communicated is binding.[20] In France, an ordinary offer is revocable until accepted, but a firm offer is binding.

The following two sections (3.2.1 and 3.2.2) shall focus on the case of a sudden and unjustified rupture of negotiations.

3.2.2 Sudden and unjustified rupture of negotiations: Perspectives under the civil law

A sudden and unjustified rupture of precontractual negotiations by one party may also be deemed to be in breach of good faith, when the other party had good reason to rely on the future conclusion of the contract and had, for example, incurred some expenses in preparation of the fulfilment of its obligations. This is the opinion of Galgano (2000), writing on Italian law. This is also the case in a number of other civil law countries. In fact, Lando et al (2000) mention that the German Supreme Court has held a person liable if, without good reason, he refuses to continue negotiations after having conducted himself in such a way that the other party had reason to expect a contract to come into existence with the content which had been negotiated. They explain further that the same also applies in Austria, Belgium, Denmark, France, The Netherlands and Portugal.

A similar notion is also found in the Principles of European Contract Law drawn up by the Commission on European Contract Law under the chairmanship of Prof. Ole Lando (Lando et al, 2000), and in the UNIDROIT Principles of International Commercial contracts (UNIDROIT, 1994). The Principles of European Contract Law and the UNIDROIT Principles are neither treaties nor international conventions.[21]

[20] See German BGB, §145.
[21] However, among the stated aims of the Principles of European Contract Law, is to suggest a modern European 'lex mercatoria', and to help bring about harmonisation of general contract law within the European Union. Of course, these Principles are also available for immediate use by individual contracting parties who may choose to

Article 2.301 of the Principles of European Contract Law starts with the general rule in sub-article (1) that "a party is free to negotiate and is not liable for failure to reach agreement." This is identical to Article 2.15 (1) of the UNIDROIT Principles.

However, the Principles of European Contract Law then provide that a party who breaks off negotiations contrary to good faith and fair dealing is liable for the losses caused to the other party (Article 2.301 (2)). The UNIDROIT Principles contain a mirror provision in Article 2.15 (b) which however provides that a party who breaks off negotiations *in bad faith* is liable for the losses caused to the other party.

3.2.3 Sudden and unjustified rupture of negotiations: The common law position

In England, the basic principle of freedom of contract, and the absence of any legally relevant intermediate stage between contract and no-contract, often makes it difficult to identify a possible cause of action for breaches of good faith in the negotiation stage (O'Connor, 1990). In general, a party will not be held liable for breaking off negotiations. However, a remedy in tort for negligent misrepresentation may be available to the innocent party who relied on a negligent misstatement by the other party who led him to believe that a contract would be concluded, whereby the innocent party suffered loss and on the facts there was a special relationship between the parties (Lando et al, 2000).

Common law authors mention two other bases on which recovery could be made (i) by holding that a collateral contract had come into existence between the parties, (ii) by way of a claim for restitution.

The former was the basis for the decision in Brewer Street Investments Ltd v. Barclays Woollen Co. Ltd. (1954) and in Harvela Investments Ltd. v. Royal Trust Co. of Canada (CI) Ltd. (1986). In Brewer Street, defendants were negotiating the lease of plaintiffs' premises and, in the expectation shared by both parties that a lease would be agreed, defendants had requested that the plaintiffs have certain work done on the premises which was otherwise of no benefit to them. The defendants had, however, expressly undertaken that they would be responsible for the cost of this work. However, before the work was completed, it became clear that the lease would not be concluded. Plaintiffs stopped work and sued for the amounts which they had paid to the contractors in respect of the work carried out. The Court held there was a contract between the parties for the carrying out of the work on the premises, despite the fact that there was no contract of lease concluded. Recovery was granted on the basis of a contractual 'quantum meruit', i.e. a reasonable sum for the work done (which was set at the amount which the plaintiffs had paid their contractors) as the defendants had agreed to pay for the cost of the work.

The other possible basis for recovery is by a claim for restitution. This is the case where one party, upon request by the other party, carries out certain work which was not intended to be gratuitous but which was intended to be compensated for out of the profit which the former would have made out of the future contract, and this work *benefited* the other party (William Lacey (Hounslow) v. Davis, 1957). One here notes the element of unjustified enrichment.[22]

Another solution offered by some common law jurisdictions is the rule of promissory estoppel which is a hybrid notion, comprising elements of contract (promise) and tort (reliance). This institution has been highly developed in the United States and Australia, but is of somewhat limited application in English law.[23]

One thus notes that common law judges ingeniously provided a basis for recovery, without entering into the notion of good faith, by using the notion of collateral contact, restitution and the law of torts. Cohen (1995) postulates that the collateral contract and the tort of negligence currently serve as the main tools for imposing pre-contractual liability. As Furmston (1998) explains, whether or not common law courts ultimately embrace good faith, there is an inherent strength in the common law to police bad faith.[24]

3.3 No real intention to contract

As a specific example of breach of good faith, the Principles of European Contract Law mention the case where a party enters into or continues negotiations with no real intention of reaching an agreement with the other party (Article 2:301(3)). The UNIDROIT Principles of International Commercial Contracts go one step further and clearly state that such behaviour constitutes bad faith.

subject their contract to these Principles. Similarly, parties may wish to subject their contract to the UNIDROIT Principles. See further on this, Lando et al, 2000.

[22] In civil law systems, the notion of unjustified enrichment is founded on good faith. For further discussion on restitution see Furmston et al, 1998 and Zimmermann et al, 2000.

[23] According to this doctrine, sometimes known as 'forbearance in equity', where a person promises not to enforce a legal right against another person and the latter relies on this promise to his detriment, then the court may, if it is equitable to do so, prevent the promisor from going back on the promise.

[24] Furmston et al (1998) were here referring to English and Australian common law.

Although legal systems may vary as to the scope of precontractual duties, generally, they do not permit one party to break off negotiations with impunity in pursuance of a scheme never to come to agreement. In civil law, a party who has used negotiations solely to induce the other party to take a desired course of action and then terminates them after his goal has been accomplished, will have to answer in damages to the party (Kessler et al, 1964).

This sort of behaviour is also likely to be proscribed in common law as fraudulent representation. As Furmston (1998) explains, "[s]uppose for instance it can be shown that the defendant never intended to enter into contract with the plaintiff and simply entered negotiations in order to deflect the plaintiff from negotiating with somebody else. Granted that one's state of mind is a question of fact, this is capable of being a fraudulent representation."

However, although such behaviour could also be proscribed in common law, this is not because of any duty to negotiate in good faith, but because it would be tortuous behaviour, i.e. fraudulent representation. Therefore, a party should be cautioned not to enter into or continue negotiations with no intention of contracting.

4 Concluding Remarks

This paper is the result of ongoing research on the role of good faith and fair dealing in contract. While the focus in this paper has been on the negotiation and contract formation stage, future work will focus on good faith in the execution of contracts.

Throughout this paper, one could note the marked difference of approach between civil law and common law jurisdictions. While the notion of good faith has effected contract law and is applied from the stage of contractual negotiation in most civil law countries, English and American jurists start from the premise that there is no general rule to negotiate in good faith. In these common law countries, piecemeal solutions have been developed to problems of unfairness. However, it has been shown that there are marked similarities in the behaviour expected of negotiating parties.

It is also important to bear in mind that electronic commerce transactions facilitate borderless trade and that negotiating parties may be located (i.e. resident, domiciled or established) in different jurisdictions. An important role is played by the applicable law that will determine a particular dispute because the rules and remedies may vary from one country to another. Nevertheless, efforts to identify standards for behaviour that are common in different legal systems, in order to then ensure that such behaviour is followed by one's electronic agent, will help minimise risk of falling foul of national law. Or at least it is so hoped ...

References

H.G. Beale, gen. ed. *Chitty on Contracts - Vol. 1: General Principles*, 28th ed., Sweet & Maxwell, London, 1999

N. Cohen. Pre-contractual duties: Two freedoms and the contract to negotiate. In J. Beatson and D. Friedmann, ed. *Good Faith and Fault in Contract Law*. Clarendon Press, Oxford, 1995

M. Fabre-Magnan. Duties of disclosure and French contract law: Contribution to an economic analysis. In J. Beatson & D. Friedmann, *Good Faith and Fault in Contract Law*. OUP, Oxford: 1995

I.S. Forrester, S.L. Coren and H.M. Ilgen. *The German Civil Code*. 1975, Excerpts available at http://www.hull.ac.uk/php/lastcb/bgbengl.htm#§305

M. Furmston, T. Norisada, J. Poole. *Contract Formation and Letters of Intent*. John Wiley & Sons, Chichester, 1998

F. Galgano. *Diritto Privato*. 3rd ed., CEDAM – Casa Editrice Dott. Antonio Milani, Padova: 1985

R. Harrison. *Good Faith in Sales*. Sweet & Maxwell, London, 1997

F. Kessler and E. Fine. 'Culpa in contrahendo', Bargaining in Good Faith and Freedom of Contract: A Comparative Study. *Harvard Law Review* 77(3):401-449, 1964

O. Lando and H. Beale. *Principles of European Contract Law – Parts I and II Combined and Revised*. Kluwer Law International, The Hague, London, Boston, 2000

UNIDROIT Principles of International Commercial contracts, 1994. International Institute for the Unification of Private Law (UNIDROIT), Rome, available at http://www.jus.uio.no/lm/unidroit.contract.principles.1994/doc.html (last visited 11.3.2002)

S. Whittaker and R. Zimmermann (a). Coming to terms with good faith. In R. Zimmermann and S. Whittaker, eds. *Good Faith in European Contract Law*. Cambridge University Press, Cambridge, 2000

S. Whittaker and R. Zimmermann (b). Good faith in European contract law: surveying the legal landscape. In R. Zimmermann and S. Whittaker, eds. *Good Faith in European Contract Law.* Cambridge University Press , Cambridge, 2000

M. Wooldridge and N. Jennings. Intelligent Agents: Theory and Practice. Knowledge Engineering Review 10(2), 1995

R. Zimmermann and S. Whittaker, eds. *Good Faith in European Contract Law.* Cambridge University Press , Cambridge, 2000

Court judgments

Brewer Street Investments Ltd v. Barclays Woollen Co. Ltd. (1954) 1 QB 428

Courtner & Fairbairn Ltd v. Tolaini Brothers (Hotels) Ltd. (1975) 1 WLR 297

Donovan v. Aeolian Co. (1936) 270 NY

Harvela Investments Ltd v. Royal Trust Co. of Canada (CI) Ltd. (1986) 1 AC 207

Hedley Byrne & Co Ltd v. Heller and Partners Ltd. (1964) AC 465

Interfoto Picture Library Ltd v. Stilletto Visual Programmes Ltd. (1989) 1 QB 433, 439

Smith v. Hughes. (1871) LR 6 QB 597

Turner v. Harvey. (1821) Ch. 814, 818

William Lacey (Hounslow) v. Davis. (1957) 1 WLR 932