

Supervised Interaction – A Form of Contract Management to create Trust between Agents

Martin J. Kollingbaum, Timothy J. Norman

Department of Computing Science, University of Aberdeen, Aberdeen AB24 3UE,
Scotland, UK
{mkolling,tnorman}@csd.abdn.ac.uk

Abstract. Supervised interaction is concerned with the problem of establishing trust between contracting agents in electronic markets. It is designed to put safeguards in place that ensure that errant behaviour in business transactions is either prevented or sanctioned. Supervised Interaction consists of three elements: an organisational framework, a contract specification language and a contract management protocol. The organisational framework emphasises the importance of introducing a trusted third party into any automated business transaction. The normative positions of the agents involved in an automated business transaction are explicitly expressed within the contracts that govern agents' behaviour during supervised interaction. This interaction model is designed to provide the web of trust necessary for successful deployment of agent-mediated electronic markets.

1 Introduction

Supervised Interaction is designed to support the automation of business transactions between software agents in electronic commerce environments. Its specific concern is to establish trust relationships between agents acting in open electronic markets. Software agents, under the command of their human organizations, are sent into such virtual environments to offer services or goods, negotiate deals and initiate the exchange of commodity for money. Such environments or electronic markets provide means for interaction the performance of such business transactions [6,22], and in such markets, many agents are required in different roles such as buyers, sellers, auditors, information vendors, financial institutions and other intermediaries [1]. In these transactions, the “exchange of money for commodity” is a delicate issue, and event more so when automated because of the loss of direct human control over the process.

Contracts are a traditional means to regulate and secure business transactions. They make explicit the dependencies between the contract participants and contain all the norms that govern their interaction. Contracts also embed such a transaction into the social context of the contracting partners. The relationship to such a context is important, as it provides the means of control and law enforcement to sanction defective (or non-ideal) behaviour [10]. Contracts, therefore, are just as important for electronic business scenarios as they are for traditional business transactions. The

reason is simple: whatever deals agents negotiate in an automated fashion in electronic environments, these deals are, in effect, contracts between human organizations and, therefore, these human organizations will be held responsible for their correct execution. Current contract management models are unsuited to dealing with defective behaviour of business agents. Agent interaction is often based on the assumption that agents involved will not display unexpected behaviour. It is, therefore, essential that either potential trading partners are recognizable as trustworthy, or mechanisms are put in place that establish trust.

Supervised Interaction introduces the necessary tools for automated contracting with a specific emphasis on trust: an organizational framework, a contract specification language and a specific contract management process. The organizational framework proposed for Supervised Interaction is presented in section 2. Section 3 introduces the contract specification language, which is used to describe the normative positions of the signatories to a contract: their rights in respect to the activities involved in the execution of the contract. The contract management procedure is presented in section 4. It takes place in three main phases: the registration phase, where a coalition is formed between the agents involved in the contract, a negotiation phase where the contract is instantiated and, finally, the contract execution phase.

2 Organisation

Castelfranchi et al. [2,4] emphasize the importance of a witness or “trusted third party” in the contracting process as a means to enforce social commitments. As a third force, it enables the creation of relationships between two contracting agents under a situation of trust. Three-party relationships are a powerful means to create trust in agent societies and it is used as a basic organizational element in Supervised Interaction.

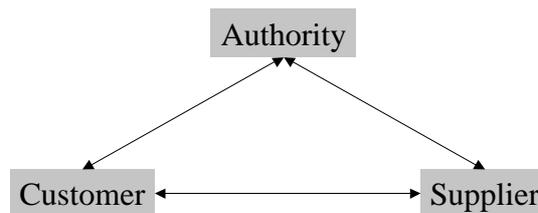


Fig. 1. Three-Party Relationship between Authority and contracting agents.

Agents are organized in a three-party relationship between two contracting individuals (or organizations), the customer and the supplier, and the “authority as the trusted third party (figure 1). The interaction between customer and supplier is determined by commitments.

The complete set of these commitments comprises the contract. Each commitment is “directed” from one agent giving this commitment (the “addressee”) towards its contracting partner (the “counter-party”), who receives this commitment. For example, a customer will commit to pay money to the supplier, whereas the supplier commits to deliver the goods to the customer. With that, a flow of goods/services can be observed from supplier to customer and a flow of money from customer to supplier. In this way, the bilateral business relationship between customer and supplier is effectively based on two unilateral sub-relationships on the level of their commitments. In terms of these commitments, both customer and supplier find themselves as the the “addressee” or “counter-party” depending on the focus of the commitment. On the level of commitments, the three roles – authority, addressee and counter-party – represent the basic building blocks of any unilateral interaction based on Supervised Interaction between two contracting agents and an authority:

- The *authority* acts as a witness to the contract that is established and executed between two agents willing to collaborate and is in a position to give an unbiased judgement on the outcomes of the contract.
- The *addressee*, one of the contracting agents under observation by the authority, takes on a commitment. In terms of a flow of goods, a supplier will commit to the supply of goods, becoming the addressee in this transfer. In terms of a flow of money, a customer will commit to pay money, becoming the addressee in this transfer.
- The *counter-party*, again under observation by the authority, is the recipient of the goods/services (flow of goods) or money (flow of money) and, therefore, gains rights over the addressee under the contract. A customer gains rights to receive payment and a customer gains a right to receive the goods.

It is important to note that a contract establishes right/duty relationships between agents, but a separate mechanism is necessary so that these rights are actually enforced and duties are correctly executed. A web of trust is needed that secures such a transaction. It is the authority that has to establish this trust relationship. The authority observes the correct execution of the contract.

In the work presented here a contract contains normative characterizations for the three agents participating in an interaction. These norms are expressed in terms of obligations, permissions and prohibitions for the contracting partners. The required and allowed behaviour is specified for each agent. The authority has an exceptional role within such an interaction, as it must have ascribed certain powers to enforce correct contract execution. This kind of empowerment is established by a separate set of behavioural definitions in the contract, called “sanctions”. These are actions an authority commits to deploy in case an agent acts in such a way that its obligations are not fulfilled or that it performs acts that are forbidden. The duty for sanctioning eventually backtracks to the human organisation represented by the authority agent in the ongoing interaction. It also means that Supervised Interaction depends on the embedding in a legal and social environment and that legal institutions must extend into electronic environments to provide services of trust and contract enforcement.

3 Contract Specification

Contracts are the central means of Supervised Interaction to create norm-governed behaviour of agents within a three-party relationship. A contract establishes this three-party relationship by specifying the participating agents, their role within the contract and their obligations, permissions and prohibitions. Supervised Interaction is designed to facilitate the creation and management of binding contracts between agents and hence between the human organizations represented by these agents. These contracts must, therefore, capture the essence of real contracts between human organizations in a form that may be interpreted and executed by agents. Real contracts describe interactions between business partners in such detail that the creation and automated negotiation of such contracts from scratch is a highly complex problem, and, arguably, an inappropriate problem for agents to solve. Legal experts are far better suited to using their expertise in the nuances of contract law in the generation of sound contracts. Existing automated negotiation mechanisms concentrate on the establishment of agreements to singular issues such as price, delivery date, quality etc. Whereas in real contracts, complex interaction schemes or business protocols provide a specification of the procedure for the enactment of the business transaction.

To allow, on the one hand, the capturing of a complex business protocol with all obligations, permissions and clauses of exceptions and sanctions as in real contracts and, on the other hand, to limit the effort for the actual negotiation task between agents, so-called “Contract Templates” are introduced.

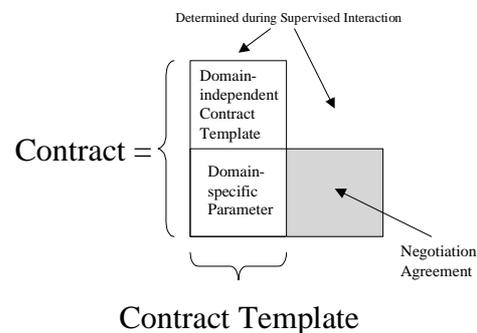


Fig. 2. Contract Template and Contract Instantiation

Contract Templates are pre-fabricated contract outlines that encode domain-independent schemata or “business protocols” such as, for example, the widely used “Letter of Credit”. Here, “domain-independent” means that such a Contract Template describes in detail the protocol for the business partners to follow, but does not specify the actual commodity or service or the current business domain. Contract Templates are formulated in such a way that they can be mapped onto any business case. The actual contract is instantiated from this template by the contracting agents in negotiating the required “domain-specific” parameters (figure 2). These are for example the price for a specific commodity or service, required quality criteria,

delivery dates etc. To allow this instantiation, Contract Templates contain “place holders” or variables for this domain-specific information.

To enable the creation of such contract templates that capture the essence of real contracts, a contract specification language has been proposed [25]. This contract specification language draws its influences from theories from the Philosophy of Law where the legal positions of individuals and groups are of interest [13,19], and the use of such languages in the specification of computer systems [10,11,17].

```

<contract> ::= "contract" <name> "(" <content> ")"
<content> ::= ( <agent> | <role> | <norm> )+
<agent>   ::= "agent" "(" <agent_name> "," <role_name> ")"
<role>    ::= "role" "(" <role_name> ")"
<norm>    ::= ( "obligation" |
                "permission" |
                "sanction" ) <norm_body>
<norm_body> ::= "(" <norm_id>      ","
                <activity>      ","
                <activation>    ","
                <expiration>    ")"
<activity> ::= <perform> | <achieve> | "not" <activity>
<perform>  ::= "do" "(" <role_name> "," <action> ")"
<achieve>  ::= "achieve" "(" <role_name> "," <condition> ")"
<action>   ::= <action_name> "(" <parameterlist> ")"
<activation> ::= <condition>
<expiration> ::= <condition>
<condition> ::= <predicate> (( "and" | "or" ) <predicate>)* |
                "not" <predicate>
<predicate> ::= <pred_name> "(" <parameterlist> ")" |
                "(" <condition> ")" |
                "TRUE" | "FALSE"

```

Fig. 3. BNF Syntax of the Contract Specification Language

Figure 3 shows the syntax specification of this language. It consists of following principle elements: role specification, agent-role assignment and normative statements. Contract templates contain role specifications only. The actual agent-role assignment takes place during contract instantiation.

Normative statements are role-specific. They express obligations, permissions and sanctions relevant to agents taking on such roles in the instantiated contract. Beside obligations and permissions, sanctions are explicitly introduced as the third form of normative statements. Sanctions are specified in correspondence to obligations and have to be enacted, if an agent does not meet its specified obligations. The role specified in sanctions is assumed to be taken on by the authority. The authority is therefore responsible to impose sanctions specified in a contract.

Obligations, permissions and sanctions are represented by the same language construct in this specification language. A normative statement consists of a <norm_id>, the required <activity>, and <activation> and <expiration> conditions. The purpose of the <norm_id> is to uniquely identify

a norm within a contract. The <activity> specification is the main focus of a norm. Following [16], a clear distinction is made between the achievement of a goal and the performance of an action. This is reflected in the language, as its norm constructs allow the formulation of statements such as “agent x sees to it that the state of affairs p holds” and “agent x sees to it that the action a is performed”. Furthermore, the language allows the expression of, for example, obligations of particular roles (and hence agents) to *not* achieve states of affairs and to *not* perform acts. Therefore, the <activity> construct in a norm statement may express four different forms of agent activity:

- An agent (taking on some role) sees to it that a state of affairs is achieved
- An agent sees to it that a state of affairs is not achieved
- An agent sees to it that an action is performed
- An agent sees to it that an action is not performed

It is worth noting here that, following theoretical models of these act expressions, stating that an agent sees to it that a state of affairs is achieved does not force the agent concerned to actually achieve this state of affairs. It may do so, but it may also delegate this activity to some other agent. It does, however, remain responsible for its achievement [16].

The conditions <activation> and <expiration> determine those states of affairs, under which obligations, permissions and sanctions become operative. Naturally, the activation condition indicates for a norm to become operative, whereas the expiration condition puts a norm out of operation. In case of a permission, both activation and expiration condition describe a *window of opportunity*, within which the agent may act. For obligations, they delimit the period in which the agent must act (possibly giving deadlines for the fulfillment of an obligation).

One of the key elements to the contract specification language is the explicit specification of sanctions within a contract. Sanctions describe actions that may be taken (or not taken) or states of affair that may be achieved or not achieved. Any obligation should be accompanied by at least one sanction, as obligations without sanctions would not be effective in determining the actions of an agent. In relationship to other norm specifications such as obligations and permissions there are important issues of consistency:

- A permission must be consistent with other permissions and obligations specified in the contract
- A sanction can override any obligation or permission specified in the contract. This means that, for example, a sanction could be defined as the removal of a permission.
- Obligations themselves have to be consistent, contradicting obligations such as “agent x has the obligation to achieve a certain state of affairs” and “agent x has the obligation to *not* achieve a certain state of affairs” are not allowed to be specified in a contract at the same time.

In Supervised Interaction, sanctions have to be assigned to the authority role. The imposition of a sanction may become an option for the authority (the activation condition indicates that it becomes operative) if one of the contracting agents does not act according to its obligations.

A fragment of a contract template is shown in figure 4. It shows part of an encoding of the “Letter of Credit” business protocol in the proposed contract specification language.

```

Contract LetterOfCredit
(
  role ( ?customer )
  role ( ?supplier )
  role ( ?bank      )

  obligation ( customer_account_request,
              do ( ?customer, open_LoC_account ( ?bank, ?customer ) ),
              TRUE,
              granted ( ?account_no ) )
  permission ( customer_account_reply,
              do ( ?bank, provide_account ( ?account_no ) ),
              account_request_received ( ?customer ),
              granted ( ?account_no ) )
  obligation ( customer_deposit,
              do ( ?customer, transfer_deposit ( ?bank, ?account_no, ?deposit ) ),
              granted ( ?bank, ?account_no ),
              deposit_transferred ( ?customer, ?deposit ) )

  . . .
  . . .

  sanction ( withhold_deposit,
            do ( ?bank, withhold_deposit ( ?supplier ) ),
            not_received_before ( ?delivery, ?deadline ),
            received_before ( ?delivery, ?deadline )
  )
)

```

Fig. 4. Fragment of a Contract Template

The template contains role specifications and declarations of obligations, permissions and sanctions. Variable elements such as `?delivery` or `?deadline` are subject to negotiation. An instantiated contract will contain explicit information instead of these placeholders. The first obligations for the customer role (as part of a “Letter of Credit” specification) are outlined. The obligation `customer_account_request` determines that the customer must deposit the money for the purchase with the bank (the authority). This amount is subject to negotiation and therefore one of the variable elements of the Contract Template. This obligation specification therefore contains an action specification for the customer to

open a “Letter of Credit” account with the bank. The activation condition for this obligation is specified here as `TRUE`, which means that this would be the first norm activated as soon as the contract execution is started by the agents. This obligation is fulfilled, or expires, as soon as the account is granted. As the following permission specification (indicating that the bank “is allowed” to grant such a request) has the same expiration condition; both norms expire at the same time. One sanction is illustrated, describing a `withhold_deposit` action in case of defection. Its activation condition specifies when this sanction is activated – in case that the goods are not delivered in time. The expiration condition is very similar to the activation condition; this is intentional, it specifies that the sanction immediately expires (and is, therefore, never active) if the goods arrive on time.

The contract management procedure of Supervised Interaction, presented in the next section, builds upon both three-party relationships as the chosen organizational structure and the contract specification language presented above. This contract management procedure provides a framework, within which agents may instantiate contracts and enact these contracts under the supervision of an authority.

4 Contract Management

The previous section shows, how a contract can be constructed using Contract Templates and a contract specification language. For contract instantiation and execution, Supervised Interaction provides a detailed contract management procedure. This procedure is the third key element of Supervised Interaction. This management activity takes place in three main phases: registration, contract negotiation and contract execution.

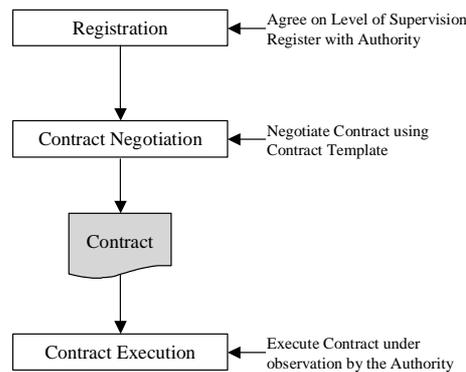


Fig. 5. Contract Management Process

- *Registration.* The purpose of the registration phase is to set up subsequent phases of the contract management procedure. Most importantly, a customer, a supplier

and authority have to create a three-party relationship. This requires (i) a match-making or yellow pages facility, and (ii) a means of deciding how to proceed with the following stages (contract negotiation phase). First of all, a set of potential business partners (customers and suppliers) have to decide that they want to engage in an interaction under the supervision of an authority. Then, they are required to agree in principle on issues open to negotiation. These issues are the type of supervisory service required from the authority and the purpose of the contract that will be negotiated in the following phase. The type of service requested from the authority is the business protocol under which the business partners intend to pursue their transaction. This service is expressed as a contract template put forward by a chosen authority

- *Contract negotiation.* In this phase, the focus is on the domain-specific content of the contract following the template agreed in the registration phase. Issues determined important in the registration phase must be negotiated, for example price, quality or delivery dates
- *Contract execution.* The fully negotiated contract is executed by the three agents under the supervision of the authority

Figure 6 shows the three phases of Supervised Interaction. The result of the registration phase is an agreement between a number of agents on how to proceed during the negotiation phase. During negotiation, the instantiation of the contract is finalized and the agents involved may proceed to execute the contract. The whole process may fail, if the agents find no agreement in the negotiation phase. The process must then be re-initiated with a new registration attempt. In the execution phase, defective behaviour of one of the agents could result in the imposition of the sanctions declared in the contract, and if so, this may disrupt the contract management process as well. It should be emphasized, that Supervised Interaction does not depend on a specific negotiation mechanism, the agents can agree on any form of negotiation in the registration phase. This can range from simply accepting a price from a catalogue for “off-the-shelf” purchases to specific negotiation about a specialized product from a limited set of suppliers. If there is a commodity/service provided by many suppliers then it would be reasonable for customers to simply advertise their needs against an anonymous crowd of suppliers. In such business transactions, the trusted third party may provide, for example, an auction service, such as the market model proposed by Dellarocas [6] or the Fishmarket [22,15,21].

The three phases are outlined in more detail in the following sections.

4.1 Registration Phase

The registration phase has to produce a result that enables the subsequent negotiation of the contract and its execution. This requires that the participants of the negotiation phase be identified along with the roles that they may play following successful negotiation of a contract, the template for the contract itself along with the domain-specific parameters that are open to negotiation, and, finally, the negotiation mechanism that is to be used. This partly depends of the relationship between the

contracting agents. It can be characterized as three types of a customer-supplier relationship:

1. One customer, many suppliers (1:N). In such a configuration, typically negotiation mechanisms such as the contract net protocol are used.
2. Many customers, one supplier (N:1). This is typical for classical auctions such as the Dutch Auction used in the FishMarket system [15,22]
3. One customer, one supplier (1:1). This contracting situation is more likely in a situation where the service or commodity required is specialised, and more sophisticated argumentation-based negotiation mechanisms may need to be employed. This is also typical where there are few (if any) issues open to negotiation; for example, the identified supplier is willing only to give a “take-it or leave-it” quotation and a set of possible delivery times

A couple of decisions are made in the registration phase that are essential to the complete contract management process. During an initial match-making step, the potential candidates for a business transaction must be identified. Customers have to retrieve information about suppliers, for example from a brokering source or an auction agent. An upfront decision about the willingness to interact has to be made by the single agents.

The outcome of the registration phase is an “Agreement in Principle” between a customer and supplier to pursue a business transaction. This agreement is signed with an authority and determines the so-called “Level of Supervision” represented by a chosen contract template and the negotiation mechanism that should be used in the subsequent negotiation phase. The registration phase can yield agreements between a large set of business agents, but it is assumed that all agree on a single authority.

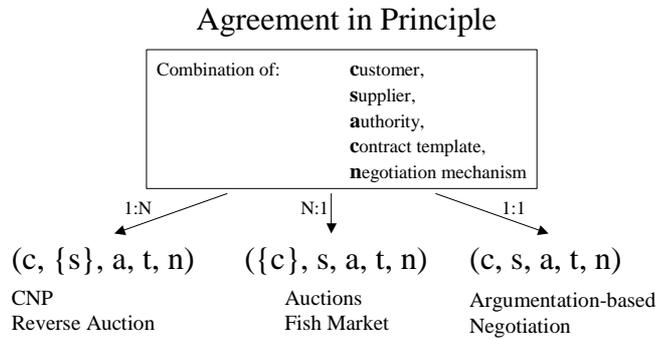


Fig. 6. Agreement in Principle

The Agreement in Principle bundles a set of information important for the complete contract management procedure. It exists between a single authority, a single contract template (partially instantiated with domain-specific information about the commodity under negotiation and the issues open to negotiation), a negotiation

mechanism and customer(s) and supplier(s) as shown in figure 6. To create an “Agreement in Principle”, the business agents have to clarify a variety of sub-issues:

- Find an agreement on the “Level of Supervision”. The “Level of Supervision” is the business protocol the business partners choose to use for their interaction. This business protocol is encoded as a Contract Template, An authority must be found that is willing to support such a protocol and can put forward the appropriate Contract Template.
- Find an agreement on an authority. According to Supervised Interaction, a customer will interact with a supplier under the supervision of an authority. Both business partners have to establish an agreement about which authority to approach and negotiate with the authority itself to gain this required support (the authority itself is a supplier of an authority service).
- Instantiate the domain independent Contract Template with the details of the commodity or service that is the object of the contract and the business transaction and identify issues that are open to negotiation.
- Find an agreement on the negotiation mechanism to be used.

The registration phase will yield an “Agreement in Principle”, in case of a 1:N relationship between one customer and many suppliers. The agreement includes one authority and one Contract Template and negotiation.

The “Agreement in Principle” can come into existence in a variety of ways. In case of a 1:N relationship between customer and supplier, the customer can (a) simply propose a completely pre-arranged “Agreement in Principle” or (b) leave all sub-issues subject to negotiation with its suppliers. In case of a pre-arranged “Agreement in Principle”, the customer has to make pre-negotiations with potential authorities, before it can propose one in the agreement. The choice of an authority is similar to that of a supplier. The authority itself is a supplier of services. Service fees and reputation are typical criteria. The authority, when approached, must decide, if it is capable and willing to provide its services, depending on, for example, its current volume of business. The authority may even offer different services, it may support different Contract Templates. After the Agreement in Principle is proposed, the potential suppliers may then indicate their willingness or unwillingness to be involved in the contracting process. An indication of willingness is considered to be its assent to this Agreement in Principle to be involved in the proposed contracting process.

In case of an N:1 relationship between a set of customers and one supplier, the supplier can act in a similar way to the customer in the previous case and eventually propose the completely pre-arranged “Agreement in Principle” to any potential customer. This scenario applies for example to auctions, where the auction house takes on the role of the authority. A customer assents to the rules of the auction and the contract template supported by the authority offering this auction service. By going through the “signing on” process involved in entering the auction house, the agent agrees to the rules of the auction: the negotiation mechanism and the Contract Template. By engaging in a specific auction within that auction house, the customer instantiates the Contract Template with the details of the commodity being auctioned. In such a situation, it is typical for the only issue open to negotiation to be the price of the commodity being auctioned.

4.2 Negotiation Phase

After an “Agreement in Principle is established, the negotiation phase can proceed. The contract template is partially instantiated, reflecting the purpose of the transaction, but will leave the issues open to negotiation unspecified. These issues must be resolved during the negotiation phase.

The registration phase may yield an agreement between a large set of agents, as described previously. The negotiation phase has to narrow down this set to one single contract establishing a three-party relationship between a customer, supplier and authority. As already mentioned, the process of Supervised Interaction is not dependent on any particular model of a negotiation dialogue. Anything from the insertion of the cost of commodity or service from a catalogue and the selection between available delivery times through an auction protocol to argument-based mechanisms could equally be used. The participants must, of course, support these mechanisms.

The completely instantiated contract is known to all three participating agents and is lodged with the authority. It contains declarations of obligations, permissions and sanctions for each contracting party. These declarations will guide the execution of the contract.

4.3 Execution Phase

In the execution phase, a signed contract will guide the behaviour of the agents. The execution phase can be demonstrated with the “Letter of Credit” protocol that is well established in the business world. It is employed in situations where there is no trust between business partners, but they (individually) trust the third party. The Letter of Credit proceeds as follows:

1. Customer deposits money with authority (the bank in this example illustrated in the Contract Template fragment in figure 4).
2. Customer receives a Letter of Credit
3. Authority informs supplier about Letter of Credit
4. Supplier transfers commodity to customer
5. Customer gives LoC to supplier
6. Supplier sends LoC to authority
7. Authority hands over money to supplier

The fragment of a contract template in figure 4 shows how this protocol could be encoded. With its definitions of obligations and permissions it describes the necessary and allowed actions and moves for the participating agents. It also shows how normative statements interact within a contract. The first obligation creates a situation that activates the second normative statement, which is a permission for the bank to provide the customer with an account. This protocol introduces a strict regime regarding the flow of money between customer and supplier. The bank acts as an intermediary and provides a deposit service. The money will be handed out under fixed circumstances.

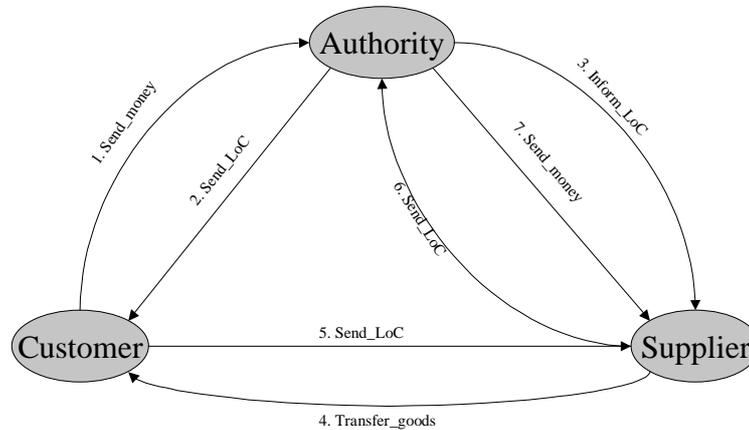


Fig. 7. Example business protocol

Figure 7 gives the graphical representation of this protocol. The seven steps of this protocol described above show how the complete transaction takes place, when the participating agents follow their duties. This protocol is organized in such a way that the money transfer takes place via a trusted third party. Of specific interest are situations where agents fail to comply with an obligation. Deviations from activities specified in the contract result in activation of sanctions. One sanction is presented contract template in figure 4, describing an activity for the bank to withhold the money under certain activation conditions.

Apart from simply not complying with obligations, agents could engage into a re-negotiation process to change their existing contract. With that, deadlines could be extended or financial penalties specified in sanctions could be adapted.

5 Related Work

This paper is based on three bodies of related work: (a) normative system specification, (b) agent-mediated electronic commerce, electronic institutions and virtual organisations and (c) models of trust and reputation.

Jones and Sergot [10, 11] and Pacheco and Carmo [17] (influenced, among others, by the seminal works of Lindahl [13], Pörn [19]) investigate the modelling of complex organisations and organisational behaviour using normative models. Pacheco and Carmo emphasise the importance of contracts as the central element to bind agents into societies. They analyse human institutions to derive properties relevant for contract specification. They describe the concept of a “role” taken on by agents as essential for modelling such an agent society. Contracts bind agents to specific roles within an institution. Roles correspond to qualities of agents and are associated with the deontic notions of obligation, permission and prohibition (cf. the specification of

roles and relationships discussed in [5, 22]). The contract specification language put forward by Pacheco and Carmo includes role specifications, deontic characterisations for these roles, statements of representation, the attribution of roles to agents and the relations between roles. In this paper, we explore extensions to this contract specification language. We consider the explicit specification of sanctions that may be imposed by the specific “authority” role introduced in Supervised Interaction. We also consider explicit activation and expiration conditions for normative statements in a contract, to clearly specify the time window during which a normative activity is operative.

Dellarocas [16] proposes the “Contractual Agent Societies” as a model for building open multi-agent market systems. In such societies, agents representing different interests join a virtual institution, where they may organise themselves through a set of dynamically negotiated contracts. An agent that joins a society undergoes a process of “socialisation”, and a set of contracts defines the shared context for interaction with members of the society. Contractual agent societies, therefore, are an abstraction of systems such as the fish market [22, 15, 21]. Agents entering the fish market undergo a registration process (socialisation) and interact (albeit through a specific market protocol) in the establishment of contracts for the sale of boxes of fish. The contractual agent society model, in common with the fish market, provides a means of social control that discourages agents from violating their commitments by detecting defective behaviour by reporting to reputation agents. Dellarocas specifies contracts in terms of beliefs, preferences and objectives.

As already mentioned, agents engage in business transactions only, if there is a certain level of trust between the business partners. Even if such transactions become more and more automated, the agents still act on behalf of human organisations and, eventually, these organisations will be held responsible for the activities of their agents. Castelfranchi and Falcone [3] state that “[t]rust is as important in multi-agent systems as it is in human societies. The notion of an agent implies the concept of delegation and delegation is based on trust.” According to Castelfranchi (and others, e.g. Marsh [14]), trust is a mental attitude, where delegation is an action that results in a specific “trusted” relationship between agents. Models of trust are established as a means for estimating the “trustworthiness” of agents. A number of attempts to quantify trust for this purpose are presented in literature ([3, 14, 24]). Yu and Singh [24] describe a model of reputation or trust management that is influenced by techniques used in recommender systems [20]. They model an electronic community, where agents assist users by maintaining contact to other agents and recommending potential and trustworthy partners. The reputation of a participant in such a community depends on capability (or expertise) and helpfulness. Agents will recommend the most helpful and reliable parties. To build and manage representations of trust, agents accumulate their own experience with a specific participant and combine it with reputation transmitted from other agents.

A model of quantifying, or identifying in some way, the level of trust of one agent in another agent with respect to specific activities is not presented in this paper. Supervised Interaction in contrast is interested to put forward a machinery that will produce upfront safeguards that generate a web of trust around a transaction between two interacting business agents.

6 Conclusion

In this paper, Supervised Interaction is presented as an interaction schema to create trust between contracting agents. It contains three key elements: an organizational structure based on a three-party relationship between two contracting agents and a trusted third party, a contract specification language and a contract management procedure. Contract Templates are used as pre-fabricated contract outlines to simplify negotiation efforts between agents. During the contract management procedure, agents register with an authority, negotiate the details required for the instantiation of full contracts from Contract Templates, and execute this contract under the supervision of an authority.

References

1. Feldman, S.: Electronic Marketplaces, *IEEE Internet Computing*, July/August (2000)
2. Castelfranchi, C.: Commitments: From Individual Intentions to Groups and Organizations, *Proceedings of the First International Conference on Multi-Agent Systems ICMAS'95*, San Francisco, (1995)
3. Castelfranchi, C., Falcone, R.: Principles of Trust for MAS: Cognitive Anatomy, Social Importance, and Quantification, *ICMAS'98* (1998)
4. Castelfranchi, C., Tan, Y.-H.: The Role of Trust and Deception in Virtual Societies, *Proc. 34th Hawaii Intl. Conf. On System Sciences* (2001)
5. Cavedon, L., Sonenberg, L.: On social commitment, roles and preferred goals. In *Proceedings of the Third International Conference on Multi-Agent Systems*, pages 80-86 (1998)
6. Dellarocas, D.: Contractual Agent Societies: Negotiated Shared Context and Social Control in open Multi-agent Systems, *2000 Workshop on Institutions and Norms in MAS, Autonomous Agents 2000*, Barcelona, Spain (2000)
7. Dignum, F., Morley, D., Sonenberg, E.A., Cavedon, L.: Towards socially sophisticated BDI Agents, *ICMAS 2000*, pp.111-118 (2000)
8. Jennings, N.R., Faratin, P., Norman, T.J., O'Brian, P., Odgers, B.: Autonomous Agents for Business Process Management. *International Journal of Applied Artificial Intelligence*, 14(2):145-189 (2000)
9. Jennings, N.R., Parsons, S., Sierra, C., Faratin, P.: Automated Negotiation, *Proc. 5th Int. Conf. On the practical Application of Intelligent Agents and Multi-Agent Systems (PAAM-2000)*, Manchester, UK (2000)
10. Jones, A.J.I., Sergot, M.: On the Characterisation of Law and Computer Systems: The Normative Systems Perspective, In J.-J.Ch. Meyer, R.J. Wieringa (editors), *Deontic Logic in Computer Science: Normative System Specification* (1992)
11. Jones, A.J.I., Sergot, M.: A Formal Characterisation of Institutionalised Power, *Journal of the IGPL*, 4(3), pp.429-445 (1996)
12. Kraus, S., Sycara, K., Evenchil, A.: Reaching agreements through argumentation: A logical model and implementation. In *Artificial Intelligence*, 104, pages 1-69 (1998)
13. Lindahl, L.: *Position and Change*, D. Reidel Publishing Company, Dordrecht-Holland/Boston-U.S.A. (1977)
14. Marsh, S.P.: *Formalising Trust as a Computational Concept*, PhD Thesis, University of Stirling (1994)
15. Noriega, P.: *Agent Mediated Auctions: The Fishmarket Metaphor*, PhD Thesis, Universitat Autònoma De Barcelona (1997)

16 **Martin J. Kollingbaum, Timothy J. Norman**

16. Norman, T.J., Reed, C.: Delegation and Responsibility, UKMAS 2000, Oxford (2000)
17. Pacheco, O., Carmo, J.: A Role Based Model for the Normative Specification of Organized Collective Agency and Agents Interaction, Journal of Autonomous Agents and Multi-Agent Systems, in press (2001)
18. Parsons, S., Sierra, C., Jennings, N.R.: Agents that reason and negotiate by arguing, In Journal of Logic and Computation, 8 (3) 261-292 (1998)
19. Pörn, I.: The Logic of Power, D. Reidel Publishing Company, Dordrecht – Holland (1970)
20. Resnick, P., Varian, H.R.: Recommender Systems, *Commun. ACM*, 40(3):56-58 (1997)
21. Rodriguez, J.A., Noriega, P., Sierra, C., Padget, J.: FM96.5 A Java-based Electronic Auction House, Proceedings of the Second International Conference on the Practical Applications of Intelligent Agents and Multi-Agent Technology PAAM-97 (1997)
22. Sierra, C., Dignum, F.: Agent-Mediated Electronic Commerce: Scientific and Technological Roadmap, In (F. Dignum and C. Sierra eds.) Agent-mediated Electronic commerce (The European AgentLink Perspective), LNAI 1991, pp. 1-18 (2000)
23. Wooldridge, M., Jennings, N.R., Kinny, D.: The Gaia Methodology for Agent-Oriented Analysis and Design. In Int Journal of Autonomous Agents and Multi-Agent Systems, 3 (3) (2000)
24. Yu, B., Singh, M.P.: A social Mechanism of Reputation Management in Electronic Communities, Cooperative Information Agents, pp.154-165 (2000)
25. Kollingbaum, M.J., Norman, T.J.: Supervised Interaction – Creating a Web of Trust for Contracting Agents in Electronic Environments, AAMAS'02, Bologna (2002)