

Chapter LI

Virtual Organization Support through Electronic Institutions and Normative Multi-Agent Systems

Henrique Lopes Cardoso

University of Porto, Portugal

Ana Paula Rocha

University of Porto, Portugal

Eugénio Oliveira

University of Porto, Portugal

ABSTRACT

The multi-agent system (MAS) paradigm has become a prominent approach in distributed artificial intelligence. Many real-world applications of MAS require ensuring cooperative outcomes in scenarios populated with self-interested agents. Following this concern, a strong research emphasis has been given recently to normative MAS. A major application area of MAS technology is e-business automation, including the establishment and operation of business relationships and the formation of virtual organizations (VOs). One of the key factors influencing the adoption of agent-based approaches in real-world business scenarios is trust. The concept of an electronic institution (EI) has been proposed as a means to provide a regulated and trustable environment, by enforcing norms of behavior and by providing specific services for smooth inter-operability. This chapter exposes our work towards the development of an agent-based EI providing a virtual normative environment that assists and regulates the creation and operation of VOs through contract-related services. It includes a presentation of the EI framework, knowledge representation structures for norms in contracts, and a description of two main institutional services, namely negotiation mediation and contract monitoring.

INTRODUCTION

The multi-agent system (MAS) paradigm has become a prominent approach in distributed artificial intelligence. These systems start with the individual—the agent—and evolve to populated environments where the most important feature is the interaction between agents.

The concept of intelligent agents is today dominant in artificial intelligence (Russell & Norvig, 2003). Agents are described as entities having a set of intrinsic capabilities, such as *autonomy* (control of own decision making), *reactivity* (response to changes in the environment), and *pro-activeness* (goal-directed behavior) (Wooldridge & Jennings, 1995). Moreover, the so-called strong notion of agency considers agents as entities having mental states, including *beliefs*, *desires*, and *intentions* (the BDI architecture; Rao & Georgeff, 1995). Another important capability—*social ability*—emphasizes the fact that agents do not act in an isolated environment, and will inevitably have to interact with other agents.

Thus, while agent theory has its inspirations in psychology and cognitive science, MAS research is influenced by organizational and social sciences, distributed computing, and economics. Furthermore, many MAS researchers look for inspiration in nature, both from collective cooperative behavior in groups of animals (e.g., ant colonies) and from our social interactions as (not necessarily cooperative) human beings.

Typical MAS applications include inherently distributed or complex domains. While some problems require system architectures including cooperative agents developed so as to accomplish an overall goal, in other cases agents may represent independent self-interested entities, with no presupposed cooperation besides mere interaction efforts. The former types of problems may be addressed through a central-

ized design, producing a top-down specification of a MAS environment with an overall purpose. The latter types are usually conceived as open environments, where heterogeneous agents arising from different sources interact either cooperatively or competitively. In this setting, agents may form organizations that dynamically emerge in a bottom-up fashion from the individuals, which together agree, usually through a negotiation process, to cooperatively perform some task not doable individually.

As one might expect, although decentralized and dynamic systems are much more appealing, they must be handled with hybrid approaches, since a minimum set of requirements is necessary to allow for heterogeneous and independently developed agents to successfully interact. One way of achieving such a common milieu is by defining communication standards, such as those proposed by FIPA (2002). However, an important issue arises when attempting to apply agents in real-world settings: how to ensure cooperative outcomes in scenarios populated with self-interested agents. A possible answer to this problem is to regulate the environment, enforcing appropriate types of agent behavior. This should provide a level of trust necessary for the development of real-world applications of open MAS.

Following these concerns, a strong research emphasis has been given recently to normative MAS (Boella, van der Torre, & Verhagen, 2005). A normative system is a set of interacting agents whose behavior can usefully be regarded as governed by norms (Jones & Sergot, 1993). Agents are subject to these norms, which influence their decision making. Therefore, besides their goals, agents must take into account the norms that apply to them. However, considering autonomy as a central property of agents, norms are often used as a means to regulate the environment by providing incentives for cooperative behavior through norma-

tive constraints, while allowing agents to choose whether to obey or to violate them (Castelfranchi, Dignum, Jonker, & Treur, 2000).

A major application area of MAS technology is e-business automation, comprising not only the use of information gathering and filtering agents, but also the establishment and operation of business relationships. The business-to-business case is particularly important, and includes not only sell and purchase operations, but also the formation of enterprise consortiums, which have been addressed through the notion of virtual enterprises or virtual organizations. This application domain clearly matches the aforementioned MAS research concepts (Fischer, Muller, Heimig, & Scheer, 1996): agents can represent the individual interests of each business entity, and negotiate in order to establish cooperative business agreements. Agent organizations can thus emerge establishing a set of commitments among them.

One of the key factors towards the adoption of agent-based approaches in real-world business scenarios is trust. When attempting to automate the creation and operation of business relationships, the behavior of agents must be made predictable, by creating a regulated environment that enforces agents' commitments. The concept of an *electronic institution* (EI) (Dignum & Dignum, 2001; Lopes Cardoso, Malucelli, Rocha, & Oliveira, 2005) has been proposed as a means to provide such a regulated and trustable environment, by enforcing norms and providing specific services.

In this chapter we expose our work towards the development of an agent-based EI providing a virtual normative environment that assists and regulates the creation and operation of virtual organizations (VOs). These are composed of agents representing different real-world entities, which adhere to a set of norms regulating the established cooperation commitments. We will start by introducing the frame-

work of our EI and by exposing its normative environment. We will then address the issue of organizational emergence among agents within the institutional environment, and the formal means to explicitly state agents' commitments. A presentation of the main institutional services will follow—namely negotiation mediation and contract monitoring, focusing on the enforcement of norms. The chapter will conclude discussing the real-world application of the described approach and its relation to relevant contract law concepts.

ELECTRONIC INSTITUTION FRAMEWORK

The application of open MAS to real-world problems raises trust concerns. It is essential to provide guarantees that (software) agents will comply with coordination efforts. Moreover, when considering scenarios where agents agree on cooperative joint activities, it is necessary to provide mechanisms that force them to comply. The intrinsic nature of agents within MAS research does not provide cooperative assumptions, as agents are defined as autonomous and self-interested entities. This case is exacerbated when dealing with heterogeneous, independently developed, and privately owned agents.

The concept of social structure is a central issue in the social sciences: it is assumed to exist in order to impose a sense of order and predictability within a group of individuals. Social structure typically involves a framework of norms attached to roles played by members of a society. A role represents the way someone is expected to behave in particular situations, by having an associated set of normative expectations.

From a social perspective, different types of norms can be identified, with different abstrac-

tion levels. On one hand we have values, conventions, and abstract norms that are implicitly adhered to and may not have an explicitly defined social response in case of deviation. On the other side of the spectrum, we may have formal (legal) norms and prescriptions that include explicitly defined punishments to deal with rule violations.

Norms can also play an important role in open artificial agent systems, where they improve coordination and cooperation (Conte, Falcone, & Sartor, 1999) and allow for the development of trust and reputation mechanisms. As in real-world societies, norms provide us a way to achieve social order (Castelfranchi, 2000) by controlling the environment and making it more stable and predictable.

Having norms is not sufficient by itself, since agents will not voluntarily submit themselves to associated penalties in case of deviation. Therefore, appropriate mechanisms are needed to enforce norm compliance. This is where an electronic institution comes about: it provides the necessary level of trust by ensuring an enforceable normative environment. Other approaches towards the concept of EI (Esteva, 2003) have considered restrictive settings in which agents are not allowed to violate norms. In our perspective, to enforce norms is not the same as preventing their violation. This approach allows us to maintain the autonomous nature of agents, while influencing their decision making by ensuring that certain consequences will hold in case of non-compliance.

Institutional Reality

In order to provide a trustable social (agent) environment, the EI must have means to register what is going on. In general, two main types of events compose this “institutional reality” (Lopes Cardoso, & Oliveira, 2005, inspired by

Searle, 1995), which is based on the creation of “institutional facts.”

Most events will depend on agents’ activity. In the case of software agents, we must take into account their observable actions: illocutions (messages exchanged between agents which must be monitored by appropriate institutional services). On the other hand, there are events that do not depend on the performance of any action. The most obvious example is the passage of time, which is important to verify the fulfillment or violation of norms (more on this later on). Therefore, institutional facts may come about from these two sources. Taking into account the moment when an institutional fact comes about, we represent it using the following structure:

ifact(InstitutionalFact, Timestamp)

When applying the EI framework to real-world scenarios like e-contracting, it is essential that external events (those that take place in the real world) are reflected inside the EI’s virtual environment. This issue is addressed with institutional roles, as explained below.

Norm Specification

Formal models of norms typically rely on deontic logic (von Wright, 1951), a branch of modal logic. Also known as the logic of normative concepts, it embraces the notions of obligation, permission, and prohibition. Although traditionally used to analyze normative reasoning in law, applications of deontic logic in computer science exist (Wieringa & Meyer, 1993), not limited to the domain of legal knowledge representation and legal expert systems; other applications include, for example, authorization mechanisms and electronic contracting.

Extensions to the original work on deontic logic have been made so as to allow its practical

use. These include approaches to handling norm violations and associated sanctions, which may consist of contrary-to-duty obligations (Jones & Carmo, 2001): obligations that come into force in sub-ideal situations (i.e., when other obligations are violated). Other important variations consider the use of conditional and temporal aspects: obligations are often made conditional on the occurrence of another event and have an associated deadline (Dignum, Meyer, Dignum, & Weigand, 2002; Dignum, Broersen, Dignum, & Meyer, 2005).

Taking a simplified approach, we can consider norms as describing the expected behavior of agents when certain circumstances arise. In order to achieve a computational representation of norms, we can thus distinguish two main parts—the situation in which the norm applies, and its prescription:

Situation → *Prescription*

The *Situation* describes when the norm is in place. The *Prescription* specifies what should be accomplished in order for the norm to be fulfilled. With this representation, norms lend themselves to a rule-based implementation. Rule-based systems (Giarratano & Riley, 1998) allow us to encode simple reasoning procedures through rules. Each rule has two main components: the antecedent part (where the rule's conditions are specified, also known as the *left-hand side* of the rule) and the consequent part (where the rule's conclusions are stated, also known as the *right-hand side* of the rule). The conditional element of norms is satisfied according to institutional facts. Moreover, the prescriptive part indicates what else should be accomplished.

The most obvious way of prescribing certain behavior through a norm is by indicating obligations. Thus, an agent may be obliged to bring about a certain state of affairs. For practical

purposes, obligations typically include a deadline by which they are to be satisfied:

obligation(Bearer, InstitutionalFact, Deadline).

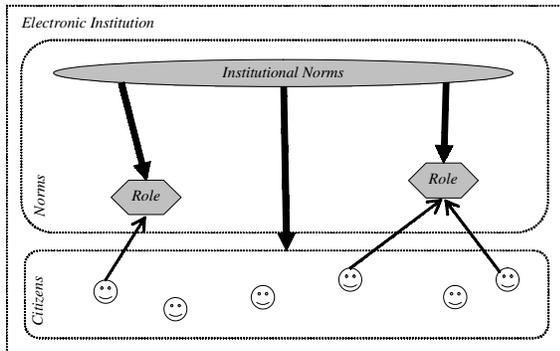
The *InstitutionalFact* that the *Bearer* must bring about by a certain *Deadline* will be part of institutional reality when the obligation is fulfilled. Notice that, instead of dictating the exact action an agent must perform, we prescribe the institutional fact that he must bring about. This enables an agent to delegate or outsource tasks conducting to the accomplishment of such state of affairs, while still being responsible before the institution for the (un)fulfillment of the obligation.

Other approaches towards norm implementation (e.g., Vázquez-Salceda, Aldewereld, & Dignum, 2004; García-Camino & Rodríguez-Aguilar, 2005) have also considered conditional permissions and prohibitions, subject to temporal references for both their activation and ceasing. With such an approach, it is straightforward to represent permissions and prohibitions with the same structure as the obligation above. Yet, we may adopt the deontic logic formalizations of permission and prohibition with respect to that of obligation: an agent is permitted to do something if it is not obliged to not do it; an agent is forbidden to do something if it is obliged not to do it.

Norms and Roles

The prescription of institutional norms comprises the top-down specification of the MAS environment (see Figure 1). Agents become subject to a set of norms when they assume a social dimension to which those norms refer. Within the EI framework, norms may apply to all citizens, to citizens enacting specific roles, or to specific groups of agents. The next section

Figure 1. Institutional norms as a top-down specification of the MAS environment



addresses the case where a group of agents may adopt a set of norms to regulate their joint activity.

Institutional norms represent the commitments of agents towards the EI. By assuming roles, agents become committed to their associated norms. A model of normative specifications based on roles can be found in Pacheco and Carmo (2003). The idea of ascribing norms to roles allows us to specify and anticipate the expected performance of the parts of a complex system (e.g., an agent-based organization) before specific agents start taking part in the system. When agents take on roles, they inherit their associated norms.

Within an EI, we can define a set of general roles (e.g., supplier or customer) that agents may adhere to, and prescribe the expected behavior of agents enacting those roles. When an agent enters the institutional environment (by registering in the EI) and announces itself as performing some institutionally defined role(s), it implicitly becomes committed to the institutional norms attached to the role(s) it plays.

Another important part that roles play in our EI setting concerns the creation of institutional reality. As mentioned before, agents' actions concern the exchange of illocutions. Two is-

ssues are worth mentioning here. First, since agents are autonomous: they are free to utter any illocutions they like. Second, part of the institutional reality concerns what happens in the "real world." This poses a problem of action certification. If, for instance, an institutional fact reports a payment of a certain amount to have taken place, we cannot rely on the bearer of such a payment stating that it is fulfilled. Instead, we would trust an independent financial third party, providing a certified institutional service.

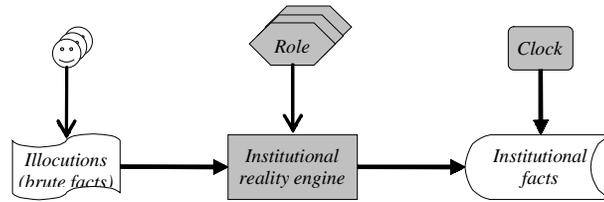
Therefore, we must define a set of institutional roles with which we assign powers concerning the creation of institutional reality. Following Searle (1995), inside the EI we distinguish between brute facts, which concern agents' illocutions, and institutional facts (see Figure 2). Rules implementing authoritative relations between roles and institutional reality are defined, making a connection between what is said and what is taken for granted. According to Searle (1995), these are "constitutive rules." An agent performing a given institutional role is said to be empowered to achieve the effects expressed in its role-related constitutive rules.

Until now, we have described an institutional environment within which agents interact by taking roles. We have also presented a means to specify norms that regulate agent behavior and how institutional reality comes about by connecting agents' illocutions to the roles they play. In the next section we describe how to specify, with contracts, commitments resulting from agents' interactions.

VIRTUAL ORGANIZATION EMERGENCE

Forming temporary coalitions of agents' efforts is an important strategy for live entities to deal with hard, complex tasks. The creation of vir-

Figure 2. Creation of institutional reality



tual organizations is, likewise, a major trend in cooperative business, where different companies try to align their businesses to perform integrated and more complex activities. B2B players are becoming more focused on their core businesses, and rely on outsourcing and dynamic consortiums. The autonomous and self-interested nature of agents fits this scenario. Agents may represent different business units or enterprises, which come together to address new market opportunities by combining skills, resources, risks, and finances no single partner can alone fulfill (Dignum & Dignum, 2002).

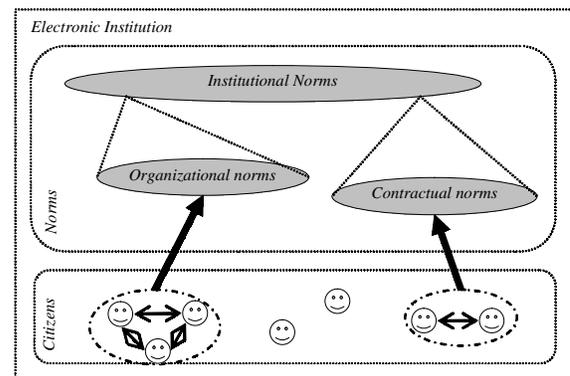
Agents may thus form organizations that dynamically emerge from the individuals, in a bottom-up fashion, establishing a set of commitments among partners. Agents agree to align their activities in a cooperative setting, which is regulated by specific norms, usually agreed upon as a result of a negotiation process. Differently from institutional norms, these are voluntarily written and adhered to by agents, instead of being pre-established in the institutional environment. Moreover, these norms apply only to the subgroup of agents participating in the agreement. Agents commit to cooperative agreements because it is in their interest to do so. Figure 3 illustrates bottom-up norm emergence.

Any cooperation activity requires trust between the involved partners. In open MAS, agents' performance records may not be as-

sessable. If such information is available, reputation ranks can be used by agents when choosing appropriate partners. Nevertheless, a VO may potentially include agents that have never worked together in the past. In this case, trust must be based on third parties, which verify whether agents comply with their promises. This makes it necessary to formulate agents' commitments explicitly through e-contracts, which are then monitored and enforced by the EI.

In the following specifications we use the Prolog (Clocksin & Mellish, 1981) notation conventions for variables and relations, although our rule-based approach is founded on general first-order logic (and thus does not assume the Prolog clause syntax).

Figure 3. Bottom-up norm emergence from the individual agents



Contract Specification

In Dignum et al. (2002), a logical formalism for describing interaction in an agent society, including social norms and contracts, is presented. Focusing on contract representation, it gives emphasis to conditional obligations with deadlines. Other approaches to contract representation through norms include Sallé (2002), considering normative statements that can comprise obligations, permissions, and prohibitions. Sanctions are seen as obligations or prohibitions activated by the violation of another obligation. Also, Kollingbaum and Norman (2002) propose the inclusion of obligations, permissions, and sanctions in a contract specification language. Each of these constructs has a similar structure, including activation and expiration conditions (which together define a “window of opportunity” within which an agent must or can act).

The use of sanctions as contrary-to-duty obligations is what makes normative multi-agent systems effective. Sanctions and the desire to avoid them are the motive for agents to comply with norms. Agents reason about the effects of their actions both in terms of positive and negative consequences: they consider both the norm and the sanction to decide whether to fulfill the norm or not (Sallé, 2002).

An approach to contract representation based on event calculus can be found in Knottenbelt and Clark (2005). This representation incorporates how a contract is to be fulfilled (that is, which events initiate or terminate obligations), making it a heavier structure.

Following the use of norms as explained before, we concentrate on specifying contracts using norms that prescribe obligations, where these concern the achievement of institutional facts by certain deadlines. However, since these norms are not institutional, but relative to a certain context—that of the agreement es-

tablished between agents—we extend the norm representation to include its scope:

[Context] Situation → Prescription

This allows us to have, within the same normative environment, norms applicable to different contexts.

We also introduce a distinction between one-shot contracts and longer-termed cooperation agreements. The aim here is to embrace not only contracts formalizing sell and purchase operations, but also more complex settings, such as a VO that involves specific interactions during a certain timeframe (Lopes Cardoso & Oliveira, 2004). The VO is regulated by a set of norms that regulate its continuous nature. Inside this ongoing relationship, contracts implementing the agreed-upon cooperation can be established, comprising a set of instantiated norms.

Simple Purchase Contracts

General contracts for purchase operations between a customer and a vendor can be specified by an institutional fact stating that such a contract (identified by *IdPC*) is in place:

`ifact(purchase_contract(IdPC, Customer, Vendor, Item, Quantity, Price), PCTime)`

Additionally, a set of behavior norms representing contract clauses prescribe the expected behavior of contracting partners. The following templates represent this information:

`[purchase_contract:IdPC] obligation(Vendor, delivery(Vendor, Item, Quantity, Customer), PCTime+10)`

`[purchase_contract:IdPC] fulfilled(Vendor, delivery(Vendor, Item, Quantity, Customer), TD)`

Virtual Organization Support

→ obligation(Customer, payment(Customer, Price, Vendor), TD+30)

According to the clauses above, a vendor's obligation exists to deliver the contract's item to the customer within 10 time units; when this is done, the customer is obliged to pay to the vendor the agreed amount. Within a norm, literals either in the situation or in the prescription part are dependent on the context of the norm. This means that prescribed obligations and their fulfillment occur within a certain context. The same applies to institutional facts, which are often related to a given context (as for instance a payment in the context of a certain purchase).

The contract is typically complemented with sanctioning norms applicable to sub-ideal (contrary-to-duty) situations. These are activated when other norms are violated. Their structure looks like the following:

[Context] violated(Bearer, IFact, Deadline) → obligation(Bearer, NewIFact, NewDeadline)

where *NewIFact* consists of a new state of affairs supposedly harder than the previously prescribed and violated one. Exemplifying, two additional clauses could be added to the contract above:

[purchase_contract:IdPC] violated(Vendor, delivery(Item, Quantity, Customer), Deadline)

→ obligation(Vendor, delivery(Item, Quantity, Customer), Deadline+5) ∪

*obligation(Vendor, payment(10%*Price, Customer), Deadline+5)*

[purchase_contract:IdPC] violated(Customer, payment(Price, Vendor), Deadline)

→ *obligation(Customer, payment(Price*110%, Vendor), Deadline+15)*

The first of these clauses indicates that if the vendor violates its obligation to deliver the item, it becomes obliged to pay the customer 10% of the price, while keeping the original obligation with a new deadline. The second clause states that, should the customer fail to pay the price to the vendor, it becomes obliged to pay a 10% increased price by a new deadline.

Eventually, problems may arise when an agent chooses to violate a sanction. A discussion of possible approaches to such situations is offered later.

Virtual Organization Cooperation Agreements

The creation of contracts regulating VO agreements is becoming common in the B2B world, where parties create consortiums that exist during a period of time. A distinguishing feature between these and the former purchase contracts is that they do not terminate after a predetermined normative path (that is, a sequence of fixed obligations). Instead, VO cooperation agreements have an ongoing nature: they may include repetitive yet unscheduled interactions.

Let us illustrate the normative specification of a VO cooperation agreement, considering a case where the intended cooperation consists of the exchange of resources between partners. The agreement will aggregate the organization's constitutional information, including the cooperation effort parties commit to and their general business process flow. Institutional facts register the standing cooperation agreement, which resulted from a successful negotiation process. We consider the following templates:

[] ifact(cooperation_agreement(IdCA, Participants, Resources), CATime)

[cooperation_agreement:IdCA]

coop_effort(Participant, Resource, MinQuantity, MaxQuantity, Frequency, UnitPrice)

[cooperation_agreement:IdCA] business_process(From, Resource, To)

Besides the initial fact stating who the participants are and what resources are considered within this cooperation agreement, cooperation efforts indicate quantity ranges for the supply of resources within a given frequency, together with agreed prices. Business process entries indicate the resources that are supposed to flow between participants. Their effective transfer, however, is dependent on appropriate requests, as we shall see. A minimum of one *coop_effort* fact per participant and resource should be present. Likewise, there should be at least one *business_process* fact for each *coop_effort*, and it must conform to the latter. Both these kinds of facts refer to a certain cooperation agreement (*IdCA*) and are thus contextualized by it.

Although the commitments of each party are implicitly defined by such institutional facts, the possibility for their verification requires that they are made explicit. The following norm addresses this need:

[cooperation_agreement:IdCA]

ifact(request(Requester, Resource, Quantity, Answerer), TR) Û

business_process(Answerer, Resource, Requester) Û coop_effort(Answerer, Resource, MinQt, MaxQt, Freq, _) Û

calculate_performed_effort(Answerer, Resource, Freq, TR, PE) Û PE+Quantity<=MaxQt

→obligation(Answerer, acknowledge(Answerer, Resource, Quantity, Requester), TR+10)

This norm states that if a predicted request (considering the stated business process and cooperation effort) is made in the context of a cooperation agreement, then the envisaged agent is obliged to accept it. An institutional procedure (*calculate_performed_effort*) is invoked for calculating the effort already performed by the agent within the timeframe indicated in the cooperation effort frequency, taking into account the request time. If the agent does not exceed its promised efforts, the obligation comes into effect.

Then, we state that an operational contract (containing operations that implement the cooperation agreement) comes into existence when an agent fulfils its obligation to accept a request:

[cooperation_agreement:IdCA]

fulfilled(Answerer, acknowledge(Answerer, Resource, Quantity, Requester), TA)

→ ifact(operational_contract(IdOC, Requester, Answerer, Resource, Quantity), TA)

Finally, we define how operational contracts are to be handled within a cooperation agreement. For instance, if delivery and payment are to take place:

[cooperation_agreement:IdCA, operational_contract:IdOC]

obligation(Answerer, delivery(_, Resource, Quantity, Requester), OCTime+10)

[cooperation_agreement:IdCA,
operational_contract:IdOC]

fulfilled(Answerer, delivery(_, Resource,
Quantity, Requester), TD) Û

coop_effort(Answerer, Resource, _, _, _,
UnitPrice)

→ obligation(Requester, payment(_,
UnitPrice*Quantity, Answerer), TD+30)

in which *IdOC* remains unbound, as these norms apply to all operational contracts which will be created in the future within agreement *IdCA*.

Facilitating Contract Formation

The creation of contracts may become cumbersome if we assume that they must predict all possibilities concerning their performance. Moreover, when considering the possibility to automatically negotiate, using software agents, a contractual relationship, it is necessary to provide a structure on which negotiation can be based. This is the role of a contract template: it contains domain-independent interaction schemata and variable elements (such as price, quantity, deadlines, and so on) to be filled in with domain-specific data resulting from a negotiation (Kollingbaum & Norman, 2002).

In general, three possibilities exist concerning the creation of a contract (Field & Hoffner, 2005):

- using monolithic contract templates, which are rigid and require no assembling from smaller granularity clauses, only filling in the relevant information from a business agreement;
- building compositional contracts, starting from contract clause templates (the build-

ing blocks) that have to be assembled, resulting in a complete contract conforming to certain composition rules; and

- taking a hybrid approach, combining monolithic templates with a number of contract clause templates that can be assembled in different ways.

In any case, complete contracts that spell out in complete detail the exact duties of each party under every possible situation are imaginary; in practice, no real contract ever achieves this level of completeness. As such, incomplete contracts must rely on a legislative system that resolves any issues not explicitly addressed by the parties. In fact, an important concept in contract law theory is the use of “default rules” (Craswell, 2000): rules that define the parties’ obligations in the absence of any explicit agreement to the contrary. These exist with the intent of facilitating the formation of contracts, allowing them to be underspecified by defining default clauses or default values.

It is therefore imaginable that institutional norms could embrace default rules. These are particularly useful for handling abnormal contract performance (that is, dealing with violations through default sanctions), since these situations may be (at least in principle) not likely to occur in practice.

INSTITUTIONAL SERVICES

According to our EI rationale, the main goals of an EI include (Lopes Cardoso et al., 2005):

- to support agent interaction as a coordination framework, making the establishment of business agreements more efficient; and
- to provide a level of trust by offering an enforceable normative environment.

The way to achieve these goals is by providing a set of institutional services, covering the phases of the lifecycle of contractual relationships: information discovery, contract negotiation, and execution. These phases have also been identified as *pre-contractual*, *contractual*, and *post-contractual*, respectively. This nomenclature emphasizes the core of any business relationship: the contract. The EI's normative environment is devoted to monitoring and enforcing contracts that are expressed through sets of norms regulating a cooperative business activity. We consider not only simple contractual relationships between two parties, but also the formation of VOs composed of multiple partners.

Most currently available support to e-business is devoted to the first phase: we can identify typical e-market functions such as yellow-page support, customer aggregation mechanisms, and recommender systems. We concentrate on the subsequent two stages, where we think of services assisting the negotiation of contracts (negotiation mediation, ontology-mapping, contract-building tools through templates) and their execution (contract monitoring and enforcement).

In this section we will focus on negotiation mediation, and on contract monitoring and enforcement services. These are aligned as shown in Figure 4 (adapted from Lopes Cardoso et al., 2005).

Negotiation Mediation

Different negotiation protocols may be devised that enable the establishment of B2B contracts. Most work in this area has focused on bilateral interactions, developing on multi-attribute utility theory for handling multiple issue negotiations. Extensive work has also been done with auctions, addressing in most cases single-issue interactions (based on the price of goods). There is a comparatively smaller amount of

research concerning the formation of a VO, where the outcome of a negotiation is the establishment of an agreement between a set of partners. Here we present our approach to a negotiation process devoted to the formation of a VO (Rocha, Lopes Cardoso, & Oliveira, 2005).

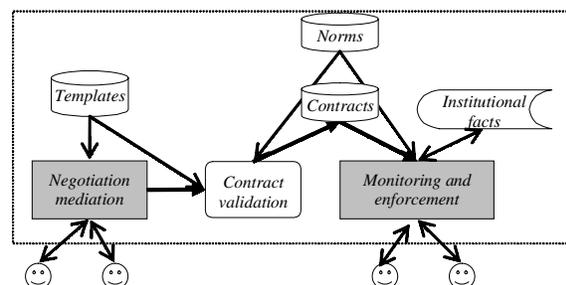
Our proposed negotiation protocol has in mind to make electronic commerce activity closer to what happens in traditional markets. Intelligent trading agents, representing the individual organizations (*Organization Agents*) and the VO organizer (*VO Organizer Agent*) engage themselves in a negotiation process by exchanging proposals and counter-proposals, trying to convince opponents to modify their bidding values. Two important features of agent-mediated electronic commerce are dealt with in this negotiation process that leads to VO formation: multi-issue and adaptation.

The Multi-Issue Negotiation Protocol

In commerce transactions, goods under negotiation should be described along several issues. Each one of these issues has a relative importance to its owner, which should be taken into account during negotiation.

The *VO Organizer Agent* starts the process with an announcement and evaluates received bids. Bid evaluation is done through a

Figure 4. Negotiation and contract execution



multi-issue function that encodes the preferences for both the issues and issues' values, which are private to the agent that plays the role of *VO Organizer*. This multi-issue evaluation function is defined by the following formula:

$$Ev = \frac{1}{Deviation}$$

$$Deviation = \frac{1}{n} * \sum_{i=1}^n dif(PrefV_i, V_i)^i$$

where n is number of issues that defines a specific good. In the *Deviation* formula, parcels are presented in increasing order of preference. The issue identified by number 1 ($i=1$) is the one that the *VO Organizer Agent* classifies as the least important, and issue n ($i=n$) is the most important. It seems intuitive that the *VO Organizer Agent* chooses a proposal that contains the values nearest to the preferable ones, for those issues the agent evaluates the most. The function $dif(PrefV_i, V_i)$ computes, for an issue i , the degree of acceptability of the current value (V_i) proposed by a specific *Organization Agent* when compared to the *VO Organizer Agent's* preferred value ($PrefV_i$).

The bid that presents the highest evaluation value is selected as the winner bid in the current round, since it is the solution that contains issue values closer to the preferred values. The winning bid in this current round is compared with the winning bid in all past rounds, and the best one is selected.

All other bids are compared with the winning bid in order to inform the respective agents about the reason why their bids were not selected. A comparison is made for the values proposed for the most important issues. These values are then classified in a qualitative way: low/high, very-low/very-high, or extremely-low/extremely-high. This classification is used as feedback to the agents that made the proposals.

Improving Offers by Learning

Another important characteristic that should be available in automated EC negotiation is the capability of learning, since it is naturally present in any commerce activity. Similar to what happens with human behavior, agents use past experience to influence future negotiations. Past proposals in a negotiation process can, and should, constrain the value of the next participant's proposal. The learning methodology we have chosen is a non-supervised learning called *reinforcement learning*. Reinforcement learning (RL) systems learn how to behave through trial-and-error interactions with its environment. RL is based on the idea of rewarding actions that produce good results and punishing those that produce bad results. Agents will use their experience to improve performance over time. "The idea that we learn by interacting with our environment is probably the first to occur to us when we think about the nature of learning" (Sutton & Barto, 1998).

We use the Q-learning algorithm (Watkins & Dayan, 1992). This is a reinforcement learning algorithm that maps state-action pairs to values, called Q-values. The *Organization Agent* receives a scalar evaluation from the environment for each action selected in a specific situation. When an agent in current state s performs action a , receives reward r , and arrives at next state s' , the correspondent Q-value is updated as follows (Sutton & Barto, 1998):

$$Q(s, a) = Q(s, a) + \alpha \left(r + \gamma \max_b Q(s', b) - Q(s, a) \right)$$

where $0 \leq \alpha \leq 1$ (learning rate) and $0 \leq \gamma \leq 1$ (discount factor). The discount factor is used to decrease the weight of reinforcements received in the future. It causes immediate reinforce-

ments to have more importance than future reinforcements.

An agent is faced with a dilemma in order to choose what action to perform in each specific situation. Should it choose the most promising one? Or should it try to explore new states? If no exploration is made, the agent, although using a greedy strategy, may miss the dynamics of the market. If too much exploration is made, the agent can make a lot of poor choices. It is necessary to ensure that sufficient exploration is made while choosing actions with high known Q-values; the so-called *Boltzmann exploration* (Sutton & Barto, 1998) equation allows us to deal with this issue.

We have adapted the Boltzmann exploration to our application domain, by not exploring all the possible actions, but reducing the exploration to the set of actions that are constrained by means of the feedback to past bids.

In our application domain, a state is defined by the n-uple:

$$s = (V_1, V_2, \dots, V_n)$$

V_i is the value for issue i . The state will also represent the bid sent to the *VO Organizer Agent*.

We define an action by the n-uple:

$$a = (A_1, A_2, \dots, A_n), A_i \in \{increase, decrease, maintain\}$$

When an *Organization Agent* receives a feedback message from the *VO Organizer Agent* to its previous bid (this bid represents the state s that results from performing action a_p in state s_p), it tries to formulate a new bid by executing the following steps:

- Calculate a reward value for the previous bid (that also defines the current state s),

using the qualitative evaluation included in the received feedback message.

- Update the state/action pair Q-value ($Q(s_p, a_p)$) using the Q-learning formula presented above.
- Derive all new promising actions a' taking into account the specific feedback message.
- Move to the state s^* with greatest reward value, that is, choose the action a^* with greatest probability value, according to the Boltzmann exploration formula (note that actions never tried before are stored in state s^* with Q-value=*default*).
- Perform action a^* by sending the new bid (state) out to the *Organizer Agent*.

This introduced negotiation protocol and adaptive negotiation strategy is detailed in Rocha et al. (2005). It has been tested with simulation experiments in simple scenarios of the textile industry domain.

Contract Monitoring and Norm Enforcement

Having a contract specified through a set of norms, we need appropriate mechanisms to monitor and enforce norm execution. Within the framework of an EI, monitoring and enforcement services can be rendered by the institution itself. Only a trusted third party can enable the necessary level of confidence between the parties involved in a business relation.

According to Vázquez-Salceda et al. (2004), an operational semantics for norms comes down to either: (1) defining constraints on unwanted behavior, or (2) detecting violations and reacting accordingly. In MAS, the absence of control over autonomous agent behavior leads us to the latter practice.

Monitoring the compliance of parties to norms touches some problems of deontic logic, namely

the need to consider violations and their handling, as well as the inclusion of deadlines. These issues must be resolved in order to enable the implementation of norms in practical applications. As explained before, our approach consists of prescribing conditional norms with deadlines and defining sanctioning norms, also known as contrary-to-duty obligations (Jones & Carmo, 2001).

Since we essentially rely on obligations, verifying norm compliance consists of detecting the fulfillment and violation of obligations. For this, we define rules that are triggered by corresponding events: the achievement of obligations' institutional facts (for fulfillment) and the reach of deadlines (for violation).

Contextualized institutional facts are used to verify the fulfillment of obligations. For this, we define an obligation *fulfillment rule* applicable to all contexts (that is, to any contractual relationship monitored by the EI):

$$\begin{aligned} & [\text{Context}] \quad \text{ifact}(\text{IFact}, \quad T) \quad \wedge \\ & \text{obligation}(\text{Bearer}, \text{IFact}, \text{Deadline}) \wedge \\ & T < \text{Deadline} \end{aligned}$$
$$\rightarrow \text{fulfilled}(\text{Bearer}, \text{IFact}, T)$$

This rule indicates that if an institutional fact prescribed by an obligation is achieved before its deadline, then that obligation is fulfilled. As with behavior norms, literals within the rule are dependent on its context. That is, if an obligation within a certain contract is accomplished, the fulfillment of such obligation occurs, obviously, inside that same contract.

This rule is fundamental for enabling the chaining of obligations within a contractual relationship. It establishes a connection between the institutional facts that are added and the pending obligations.

The connection between behavior norms and sanctioning norms is achieved through violation

detection mechanisms. These are based on *violation detection rules*, which fire when deadlines have elapsed. For this we consider time events, which are generated as institutional facts by a clock triggering mechanism, corresponding to the time when obligations are due.

$$\begin{aligned} & [\text{Context}] \quad \text{ifact}(\text{time}, \text{Deadline}) \wedge \\ & \text{obligation}(\text{Bearer}, \text{IFact}, \text{Deadline}) \wedge \\ & \text{not}(\text{fulfilled}(\text{Bearer}, \text{IFact}, _)) \end{aligned}$$
$$\rightarrow \text{violated}(\text{Bearer}, \text{IFact}, \text{Deadline})$$

This violation detection rule states that in any context, if a deadline referring to an obligation was reached, and such obligation was not fulfilled, then a violation occurred (in the same context).

These two rules comprise the core of our normative environment, enabling the chaining of norms. The implementation of such a framework is achievable using a rule-based engine, like Jess (Friedman-Hill, 2003), which enables a forward-chaining approach based on the occurrence of events.

DISCUSSION

The concept of virtual enterprise or virtual organization has been applied to many forms of cooperative business, such as outsourcing, supply chains, or temporary consortiums. We are mostly concerned with the latter case, considering flexible and dynamic relationships. Their temporary nature requires quick set-up phases, allowing the VO to start operating as soon as possible.

Taking this into account, tools are being developed that allow for automating part of the process of setting up VOs. These include negotiation mediation and contract formation services, as we have described in this chapter. We

have also developed a model that addresses the operation phase of a VO, allowing the monitoring of parties' compliance to contractual commitments. Other developments would include regulating the structure of VOs by defining norms that specify when and how partners can leave or enter a given consortium, and when and how the consortium should terminate.

The approach presented in this chapter is inspired by our social interactions as humans. In order to regulate the real-world environment, we have created institutions that deal with specific aspects of our daily lives, giving us some sense of order. The Electronic Institution concept is proposed to regulate the interactions of computational agents, providing trust by enforcing appropriate types of behavior. Furthermore, we address the issue of commitment creation between individual agents, which form a VO that is regulated by norms enforced by the EI.

In the real world, when contractual obligations are assumed by parties, they are not automatically enforced. For this reason, in non-electronic contracts, contrary-to-duty structures are not common. The violation of an obligation entitles the offended party to invoke legal power on a court of law, which may prescribe a secondary obligation to be imposed on the prevaricator (Daskalopulu & Maibaum, 2001). Besides, parties are not willing to stipulate handling procedures for all possible circumstances, deferring them to when and if situations arise. In our framework, the use of default rules (Craswell, 2000) as institutional norms addresses this problem, while enabling the automation of contract monitoring and enforcement.

Another significant intricacy concerns how to approach situations in which sanctions are violated. In such cases, more restrictive measures are needed. Instead of prescribing further sanctions, we can constrain the behavior of

agents, for example by preventing their access to institutional services. Other approaches consist of indirect punishment measures through reputation mechanisms.

Relational contract theory (Hviid, 2000) studies continuing relations that are naturally self-enforceable. Instead of a detailed enforceable contract based on a third party, a relational contract is based on repeated interactions and social norms, representing an informal agreement sustained by the value of future relationships. Formal contracts are preferred when establishing relationships between unknown parties. On the other hand, regular partners generally rely on implicit relationships, supported by trust and by the threat of withholding business from anyone who has broken a promise in the past. The information on agents' reputation may also be used, if not as a ruling out factor, at least when deciding the level of detail a contract should have.

The adoption of MAS in domains such as the B2B world is challenged by matters of user trust. It is arguable what kinds of tasks human decision makers will be likely to delegate to software agents. The delegation of business-critical tasks is not likely to be a reality in the near future. Therefore, we envision the use of agent-based infrastructures—such as our Electronic Institution environment—as supporting tools that assist contractual relationships. Besides facilitating contract formation, our framework also enables monitoring contract fulfillment, and may be used to alert real-world entities concerning their upcoming obligations, the consequences of their actions, and their partners' compliance.

REFERENCES

Boella, G., van der Torre, L., & Verhagen, H. (2005). Introduction to normative multiagent

systems. *Proceedings of the 1st International Symposium on Normative Multiagent Systems*, Hatfield, England.

Castelfranchi, C. (2000). Engineering social order. In A. Omicini, R. Tolksdorf, & F. Zambonelli (Eds.), *Engineering societies in the agents' world* (pp. 1-18). Springer.

Castelfranchi, C., Dignum, F., Jonker, C., & Treur, J. (2000). Deliberative normative agents: Principles and architectures. In N. Jennings & Y. Lesperance (Eds.), *Intelligent agents VI: Agent theories, architectures, and languages* (pp. 364-378). Springer.

Clocksin, W. F., & Mellish, C. S. (1981). *Programming in Prolog*. Springer-Verlag.

Conte, R., Falcone, R., & Sartor, G. (1999). Introduction: Agents and norms: How to fill the gap? *Artificial Intelligence and Law*, 7(1), 1-15.

Craswell, R. (2000). Contract law: General theories. In B. Bouckaert & G. De Geest (Eds.), *Encyclopedia of law and economics, volume III: The regulation of contracts* (pp. 1-24). Cheltenham: Edward Elgar.

Daskalopulu, A., & Maibaum, T. (2001). Towards electronic contract performance. *Proceedings of Legal Information Systems Applications, the 12th International Conference and Workshop on Database and Expert Systems Applications* (pp. 771-777).
Dignum, F., Broersen, J., Dignum, V., & Meyer, J.-J. (2005). Meeting the deadline: Why, when and how. In M. G. Hinchey, J. L. Rash, W. F. Truszkowski, C. Rouff, & D. Gordon-Spears (Eds.), *Formal approaches to agent-based systems* (pp. 30-40). Berlin: Springer-Verlag (LNAI 3228).

Dignum, V., & Dignum, F. (2001). Modeling agent societies: Coordination frameworks and institutions. In P. Brazdil & A. Jorge (Eds.),

Progress in artificial intelligence: Knowledge extraction, multi-agent systems, logic programming, and constraint solving (pp. 191-204). Berlin: Springer-Verlag (LNAI 2258).

Dignum, V., & Dignum, F. (2002). Towards an agent-based infrastructure to support virtual organizations. In L.M. Camarinha-Matos (Ed.), *Collaborative business ecosystems and virtual enterprises* (pp. 363-370). Kluwer.

Dignum, V., Meyer, J.-J., Dignum, F., & Weigand, H. (2003). Formal specification of interaction in agent societies. In M. Hinchey, J. Rash, W. Truszkowski, C. Rouff, & D. Gordon-Spears (Eds.), *Formal approaches to agent-based systems* (pp. 37-52). Berlin: Springer-Verlag (LNAI 3228).

Esteva, M. (2003). *Electronic institutions: From specification to development*. PhD thesis, Technical University of Catalonia.

Field, S., & Hoffner, Y. (2005). Dynamic contract generation for dynamic business relationships. In G. D. Putnik & M. M. Cunha (Eds.), *Virtual enterprise integration: Technological and organizational perspectives* (pp. 207-228). Hershey, PA: Idea Group Publishing.

FIPA (Foundation for Intelligent Physical Agents). (2002). *FIPA-ACL message structure specification*. Retrieved from <http://www.fipa.org/specs/fipa00061/>

Fischer, K., Muller, J., Heimig, I., & Scheer, A. (1996). Intelligent agents in virtual enterprises. *Proceedings of the 1st International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, London.

Friedman-Hill, E. (2003). *Jess in action*. Manning Publications Co.

García-Camino, A., & Rodríguez-Aguilar, J.A. (2005). Implementing norms in electronic insti-

- tutions. *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS2005)*, Utrecht, The Netherlands.
- Giarratano, J., & Riley, G. (1998). *Expert systems: Principles and programming* (3rd ed.). PWS.
- Hviid, M. (2000). Long-term contracts and relational contracts. In B. Bouckaert & G. De Geest (Eds.), *Encyclopedia of law and economics, volume III: The regulation of contracts* (pp. 46-72). Cheltenham: Edward Elgar.
- Jones, A., & Carmo, J. (2001). Deontic logic and contrary-to-duties. In D. Gabbay (Ed.), *Handbook of philosophical logic* (pp. 203-279). Kluwer.
- Jones, A., & Sergot, M.J. (1993). On the characterization of law and computer systems: The normative systems perspective. In J.-J. Meyer & R. J. Wieringa (Eds.), *Deontic logic in computer science: Normative system specification* (pp. 275-307). Chichester, UK: John Wiley & Sons.
- Knottenbelt, J., & Clark, K. (2005). Contract-related agents. *Proceedings of the 6th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA VI)*, London.
- Kollingbaum, M. J., & Norman, T. J. (2002). Supervised interaction—Creating a web of trust for contracting agents in electronic environments. In C. Castelfranchi & W. Johnson (Eds.), *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems: Part 1* (pp. 272-279). New York: ACM Press.
- Lopes Cardoso, H., Malucelli, A., Rocha, A. P., & Oliveira, E. (2005). Institutional services for dynamic virtual organizations. *Proceedings of the 6th IFIP Working Conference on Virtual Enterprises (PRO-VE'05)*, Valencia, Spain.
- Lopes Cardoso, H., & Oliveira, E. (2004). Virtual enterprise normative framework within electronic institutions. In M.-P. Gleizes, A. Omicini, & F. Zambonelli (Eds.), *Engineering societies in the agents world V* (pp. 14-32). Springer.
- Lopes Cardoso, H., & Oliveira, E. (2005). Towards an institutional environment using norms for contract performance. *Proceedings of the 4th International Central and Eastern European Conference on Multi-Agent Systems*, Budapest, Hungary.
- Pacheco, O., & Carmo, J. (2003). A role-based model for the normative specification of organized collective agency and agents interaction. *Journal of Autonomous Agents and Multi-Agent Systems*, 6(2), 145-184.
- Rao, A. S., & Georgeff, M. P. (1995). BDI agents: From theory to practice. In V. R. Lesser & L. Gasser (Eds.), *Proceedings of the 1st International Conference on Multiagent Systems* (pp. 312-319). Cambridge, MA: MIT Press.
- Rocha, A. P., Lopes Cardoso, H., & Oliveira, E. (2005). Contributions to an electronic institution supporting virtual enterprises' lifecycle. In G. Putnik & M. M. Cunha (Eds.), *Virtual enterprise integration: Technological and organizational perspectives* (pp. 229-246). Hershey, PA: Idea Group Publishing.
- Russell, S., & Norvig, P. (2003). *Artificial intelligence: A modern approach* (2nd ed). Englewood Cliffs, NJ: Prentice-Hall.
- Sallé, M. (2002). Electronic contract framework for contractual agents. In R. Cohen & B. Spencer (Eds.), *Proceedings of Advances in Artificial Intelligence, the 15th Conference of the Canadian Society for Computational Studies of Intelligence* (pp. 349-353). Springer.

Searle, J. R. (1995). *The construction of social reality*. New York: The Free Press.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Vázquez-Salceda, J., Aldewereld, H., & Dignum, F. (2004). Implementing norms in multiagent systems. In G. Lindemann, J. Denzinger, I.J. Timm, & R. Unland (Eds.), *Multiagent system technologies* (pp. 313-327). Springer.

von Wright, G. (1951). Deontic logic. *Mind*, 60, 1-15.

Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3/4), 279-292.

Wieringa, R.J., & Meyer, J.-J. (1993). Applications of deontic logic in computer science: A concise overview. In J.-J. Meyer & R. J. Wieringa (Eds.), *Deontic logic in computer science: Normative system specification* (pp. 17-40). Chichester, UK: John Wiley & Sons.

Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2), 115-152.

KEY TERMS

Agent: A goal-directed entity capable of interacting with the surrounding environment in an autonomous fashion; a software program representing an individual business entity.

Contract: An explicit formulation of agents' commitments.

Cooperation Agreement: A contract with an ongoing nature, regulating a virtual organization relationship that exists for a period of time.

Electronic Institution: A software platform providing a set of contract-related services for agent interaction and a normative environment for enforcing contracts.

Multi-Agent System: A system where multiple agents interact.

Norm: A prescription of behavior.

Normative Multi-Agent System: A set of interacting agents whose behavior can be regarded as governed by norms.

Role: A set of normative expectations and empowerments.

Virtual Organization: A consortium of organizations (enterprises) that align their businesses to perform integrated and more complex activities.