Chapter XI

Contributions to an Electronic Institution Supporting Virtual Enterprises' Life Cycle

Ana Paula Rocha, University of Porto, Portugal

Henrique Lopes Cardoso, University of Porto, Portugal

Eugénio Oliveira, University of Porto, Portugal

ABSTRACT

Electronic commerce competitiveness, due to market openness and dynamics, enabled the arising of new organizational structures, as it is the case with virtual enterprises. The virtual enterprise (VE) concept can effectively answer to new demanding market requirements, as it combines the core competencies of independent and heterogeneous enterprises that collaborate in a temporary and loosely linked network, thereby presenting high flexibility and agility. However, institutional and social laws must be introduced here to enforce and regulate individual enterprises' behavior. An electronic institution is a framework that enables through a communication network automatic transactions between electronic business parties, according to sets of explicit institutional norms and rules. This chapter presents and discusses tools for automatic negotiation and operation monitoring that make an electronic institution a suitable framework for helping in the two most important stages of a VE's life cycle: formation and operation. Moreover, the electronic contract concept is defined and discussed.

INTRODUCTION

The growth of information and communication technology has changed the way traditional commerce has been done by eliminating time and space restrictions. A new way of commerce, based on network communications, encompasses two fields: the business-to-consumer (B2C) and the business-to-business (B2B) electronic commerce. In the B2C electronic commerce, business participants are individual buyers and sellers that announce and negotiate over a final product or service. In the B2B electronic commerce, contrary to what happens in B2C, the goal of the business transaction is not a final product, and generally, business participants are enterprises that need to include in their own processes products that are outside of their expertise domain or resources they do not own. The work reported here is related to the last-mentioned type of electronic commerce, that is, the B2B electronic commerce.

The electronic commerce has increased the business competitiveness, due both to the market openness and dynamics. Enterprises try to answer these new market requirements by engaging themselves in temporary corporations, thereby presenting a flexible structure that changes dynamically according to current market situations. This new agile organizational structure is called VE. All those enterprises collaborate for a global goal with their competencies, knowledge, and resources. Agility is possible, because individual enterprises that belong to the VE are loosely coupled in this networked structure, and, although working for the VE global goal, enterprises maintain their autonomy.

A computing platform named *ForEV* (acronym for *Virtual Enterprises Formation* equivalent in Portuguese) was developed for supporting the VE formation stage. The VE formation stage has as its primary objective the creation of an organization able to compete as well as respond to the demanding requirements coming from an open market, by including in that organization those enterprises that have either the higher competence or present the best transaction conditions for that business opportunity. Our approach includes an iterative, adaptive, multiattribute negotiation protocol using qualitative argumentation (the "Q-negotiation" algorithm).

The negotiation that takes place during the VE formation stage leads to the agreement of an electronic contract that should be signed by all individual enterprises selected as partners in the VE. The VE operation stage uses this electronic contract to monitor the VE activity. The electronic contract describes the rights and duties of all VE partners, as well as penalties to apply to those that do not satisfy the agreement.

The rationale of this chapter includes the understanding of the VE concept, a definition of a generic model of an electronic institution, our proposal of tools enabling the electronic institution's role in helping in the VE formation stage, the exploitation of electronic-contracting services within an electronic institution that helps in the VE operation stage, and conclusions and directions for future work.

VIRTUAL ENTERPRISE

The VE is generally associated with the concept of a network of enterprises. However, a network of enterprises is not, necessarily, a VE. Figure 1 summarizes and clarifies several networked organizations categories according to two dimensions: uncertainty and mutual dependency (Camarinha & Afsarmanesh, 1999; Jagers, Jasen, &

Steenbakkers, 1998). The uncertainty level measures the uncertainty found by one enterprise when initiating a business relationship with other enterprises in the network. The mutual dependency level measures the enterprise's autonomy.

Figure 1. Networked organizational structures.

An *extended enterprise* can be seen as a network of enterprises where one is dominant, and thereby subcontracts other (dominated) enterprises by outsourcing the products it needs.

In a *strategic alliance*, all enterprises have interest in each others' success, because their activities are mutually dependent.

The organizational structure named *VE* presents a more democratic structure than the extended enterprise, where all its members are equally important. The main difference between a *stable VE* and a *dynamic VE* (or simply VE) is that in the first case, members are chosen from a closed set of already known enterprises, while in the second case, enterprises are in an open network and are not known in advance.

A *virtual organization* differs from the VE, because its members can be any kind of organization (with or without profit means), and not necessarily enterprises.

The *networked organization* encompasses all organizational structures where participants are entities linked with a computational network.

The subject of our study is the VE structure, that is, a set of independent networked enterprises that cooperate to a global goal. Following this general VE definition, different visions can even be formulated according to different authors. These are summarized in the following three topics:

- A temporary enterprise network (Fischer, Muller, Heimig, & Scheer, 1996; Peterson & Gruninger, 2000)
- A permanent network of enterprises (Camarinha & Lima, 1998)

Figure 1. Networked organisational structures

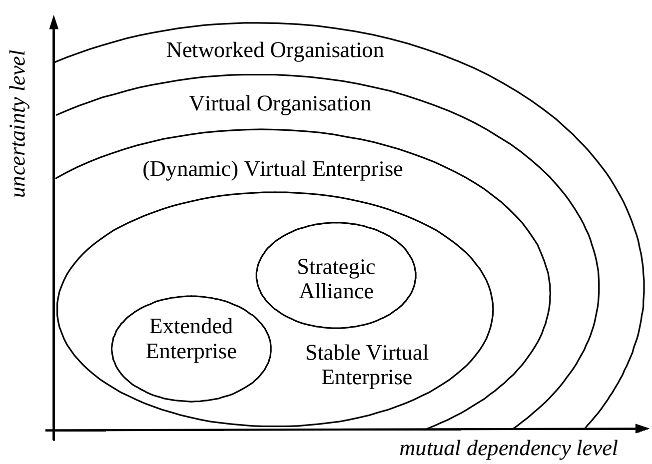

Virtual images of an enterprise structure and available data (Shmeil & Oliveira, 1997)

Our perspective on a VE is related to the first of these views, and a complete definition can be formulated as follows (Rocha & Oliveira, 2002): "A Virtual Enterprise is a temporary aggregation of autonomous and independent enterprises connected through a network and brought together to deliver a product or service in response to a customer need (Rocha & Oliviera, 2002, p. 232)."

In an electronic market, because of its openness, transactions complexity is increased due both to the huge amount of available information and the environment dynamics. The presence of a large number of business participants also originates higher market competition and increases the customers' demands. The response to these market requirements implies a new organization's concept that needs to have a (virtual) very large size needed for satisfying all the required skills. However, this type of organization, contrary to what happens with large traditional organizations, has to be flexible enough in order to deal with the dynamics of the market. VEs can satisfy these new challenges, as they combine the core competencies of several autonomous and heterogeneous enterprises aggregated in a temporary network, thereby presenting high flexibility and agility.

The VE life cycle is decomposed in four stages (Fischer et al., 1996), as follows:

- *Identification of needs*: Appropriate description of the product or service to be delivered by the VE, which guides the conceptual design of the VE
- *Formation (Partners Selection)*: Automatic selection of the individual organizations (partners), which based in its specific knowledge, skills, resources, costs, and availability, will integrate the VE
- *Operation*: Controlling and monitoring the partners' activities, including resolution of potential conflicts, and possible VE reconfiguration due to partial failures
- *Dissolution*: Breaking up the VE, distributing the obtained profits, and storing relevant information for future use for the electronic institution

Electronic tools for helping on the automatic VE life cycle imply the need of a framework for secure and reliable agents' encounters. The next section describes the electronic institution, which provides the means for helping on several stages of the VE life cycle.

ELECTRONIC INSTITUTION

An electronic institution (EI) is a framework that enables, through a communication network, automatic transactions between parties, according to sets of explicit institutional norms and rules. Thereby, the EI ensures the trust and confidence needed in any electronic transaction. However, each EI will be dependent on the specific application domain for which it has been designed. Here, we need to introduce the notion of a meta-institution, which is a shell for generating specific electronic institutions for particular application domains. The meta-institution includes general modules related to social and institutional behavior norms and rules, ontology services, as well as links to other institutions (financial, legal, etc.). The main goal of a meta-institution is to generate specific electronic institutions through the instantiation of some of these modules that are domain dependent according to the current application domain.

Figure 2. General architecture of an electronic institution

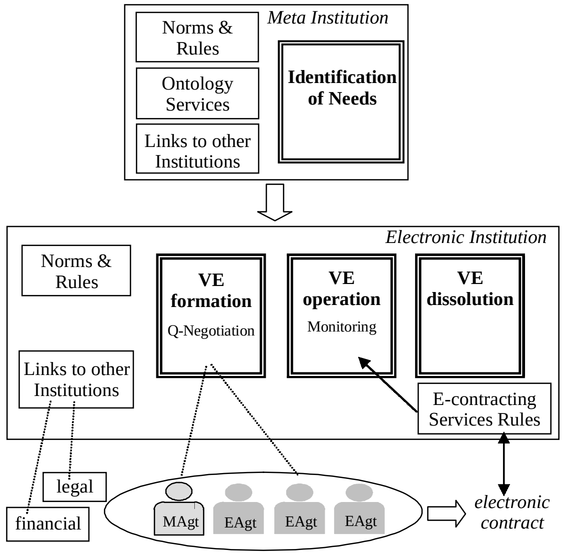

In this chapter, the electronic institution framework is analyzed in the VE scenario, and it can effectively help in making automatic several aspects of the VE's life cycle. It helps in both providing tools and services for supervising the intended relationships between parties. Figure 2 presents the general architecture of an electronic institution used in the VE scenario.

The meta-institution, as said previously, generates specific electronic institutions through the instantiation of some of these modules that are domain dependent. So, it can help in the first VE life cycle stage (the Identification of Needs), where a particular customer/market need is identified, that will be the goal of the future VE. The VE Identification of Needs stage will not be detailed here, we only note that the result of this stage is the instantiation of an electronic institution for a particular application domain.

In the next section, we discuss the electronic institution services to help in the second VE life cycle stage, that is, the VE formation process. The VE formation stage has as its primary objective the creation of an organization composed of several independent and possibly heterogeneous enterprises, which have higher competence for that business opportunity. Our approach includes an iterative, adaptive, multiattribute negotiation protocol using qualitative argumentation (the "Q-negotiation" algorithm).

In electronic transactions, in general, and in the VE formation, in particular, an important issue is to preserve the enterprise's private information during the negotiation process. An entity participating in a business transaction, and an enterprise in particular, tries to hide from the market its own private evaluation of the goods under negotiation. Adaptation is another important characteristic to be included in any entity present in an electronic market. Moreover, simultaneous partial interdependent negotiations may arise during the VE formation stage. The following sections detail the Q-negotiation algorithm, highlighting these advanced negotiation issues.

THE NEGOTIATION PROCESS

In the VE formation process, participants in the negotiation can be either market or enterprise agents. The market agent plays the role of organizer, meaning that it is the agent that starts and guides all the negotiation process. The enterprise agents play the role of respondents, meaning that they are those who are willing to belong to the future VE, and therefore, they have to submit proposals during the negotiation phase.
We consider that the VE goal is decomposed in a set of components, and for each of these components, an independent negotiation process takes place. Multiple negotiations are done simultaneously during the VE formation process.

The Negotiation Protocol

In order to agree in a VE structure, agents (market and several enterprises) naturally engage themselves in a sequential negotiation process composed of multiple rounds of proposals (sent by enterprises to market) and counterproposals that are actually comments to past proposals (sent by market to enterprises). This is what really happens in traditional commerce, where humans exchange proposals and counterproposals trying to convince each other to modify the issues' values that they evaluate the most. A negotiation protocol should then be defined in order to select the participants that, based on capabilities and availability, will be able to make the optimal deal according to its own goals.

Our proposed Q-negotiation algorithm introduces new and important advanced features in electronic markets' negotiation: *multiple-attribute negotiation*, *learning in negotiation*, *distributed dependencies resolution*. These features are detailed in the next sections.

Figure 3. Negotiation protocol

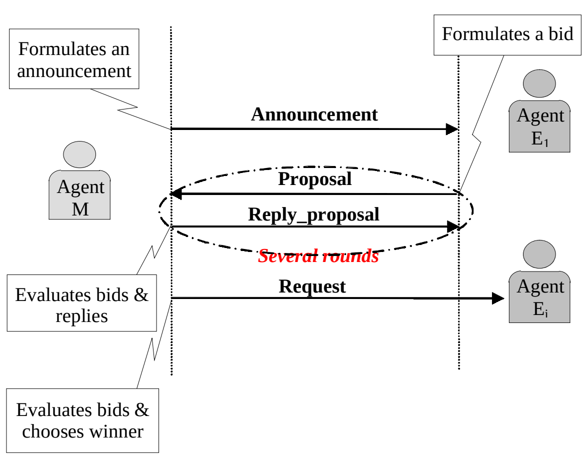

Multiattribute Bid Evaluation

Negotiation implies, for most of the economic transactions, that not only one, but multiple, attributes for defining the goods under discussion be taken into consideration. For instance, although the price of any good is an important (perhaps the most important) attribute, delivery time and quality can also be, and generally are, complementary issues to include in the decision about to buy or sell or not a specific good.

Attaching utility values to different attributes under negotiation solves the problem of multiattribute evaluation. Generally, an evaluation formula is a linear combination of the attributes' values weighted by their corresponding utility values. In this way, a multiattribute negotiation is simply converted in a single attribute negotiation, where the result of the evaluation function can be seen as this single issue (Vulkan & Jennings, 1998).

However, in some cases, it could be difficult to specify absolute numeric values to quantify the attributes' utility. A more natural and realistic way is to simply impose a preference order over attributes. The multiattribute function presented in formula (1) encodes the attributes' and attributes values' preferences in a qualitative way and, at the same time, accommodates intradependencies of the attributes.

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

where n = number of attributes that defines a specific component,

$$V_x = f(V_1, \ldots, V_n), \quad x \notin \{1, \ldots, n\}, \text{ and}$$

$$dif(\mathrm{PrefV_i}, V_i) = \begin{cases} \dfrac{\mathrm{V_i} - \mathrm{PrefV_i}}{\mathrm{max_i} - \mathrm{min_i}}, & if\ \ continuous\ domain \\ \dfrac{Pos(V_i) - Pos(\mathrm{PrefV}_i)}{nvalues}, & if\ \ discrete\ domain \end{cases}$$

A proposal's evaluation value is calculated by the market agent, as the inverse of the weighted sum of the differences between the optimal ($PrefV_i$) and the real (V_i) value of each of the attributes. In the formula, each parcel should be presented in increasing order of preference, that is, attributes identified by lower indexes are least important than attributes identified with higher indexes. The proposal with the highest evaluation value so far is the winner, because it is the one that contains the attributes' values more closely related to the optimal ones from the market agent point of view.

The negotiation process is realized as a set of rounds (see Figure 3) where enterprise agents concede, from round to round, a little bit more, trying to approach the market agent preferences, in order to be selected as partners of the VE. The market agent helps enterprise agents in their task of formulating new proposals by giving them some hints about the directions they should follow in their negotiation space. These hints are given, by the market agent, as comments about attributes' values included in current proposals.

Qualitative Feedback Formulation

The response to proposed bids is formulated by the market agent as qualitative feedback, which reflects the distance between the values indicated in a specific proposal and the optimal one received so far. The reason why the market agent compares a particular proposal with, not its optimal, but the best one received so far, can be explained by the fact that it is more convincing to say to an enterprise agent that there is a better proposal on the market than saying that its proposal is not the optimal one.

A qualitative feedback is then formulated by the market agent as a qualitative comment on each of the proposal's attributes values, which can be classified in one of three categories: sufficient, bad, or very bad.

Enterprise agents will use this feedback information to its past proposals, in order to formulate, in the next negotiation rounds, new proposals trying to follow the hints included in the feedback comments.

Learning in Bid Formulation

The Q-negotiation algorithm uses a reinforcement learning strategy based on Q-learning for the formulation of new proposals. The Q-learning algorithm (Watkins & Dayan, 1992) is a well-known reinforcement learning algorithm that maps evaluation values (Q-values) to paired state/action.

The selection of a reinforcement learning algorithm seems to be appropriate in the negotiation process that acts as a conduit to VE formation, because organization agents evolve in an, at least, partially unknown environment. And, in particular, Q-learning enables online learning that is an important capability in our specific scenario where agents will learn in a continuous way during the negotiation process, with information extracted from each one of the negotiation rounds, and not only in the end with the negotiation result.

Q-learning is based on the idea of rewarding actions that produce good results, and punishing those that produce bad results, as indicated by parameter *r* in the correspondent formula [see Equation (2)].

$$Q(s,a) \quad = \quad Q(s,a) + \alpha \left(r + \gamma \max_b Q(s',b) - Q(s,a) \right) \tag{2}$$

In the Q-negotiation process, we assume that:

- A state is defined by a set of attributes' values, thus representing a proposal.

$$s = \langle v_1, v_2, \ldots, v_n \rangle \quad \begin{array}{l} , n = number\ of\ attributes \\ , v_x\ :value\ of\ attribute\ x \end{array}$$

- An action is a relationship that is a modification of the attributes' values through the application of one of the functions: increase, decrease, or maintain.

$$a = \langle f_1, f_2, \ldots, f_n \rangle \quad \begin{array}{l} , n = number\ of\ attributes \\ , f_x \in \{increase, decrease, maintain\} \end{array}$$

The adaptation of the Q-learning algorithm to our specific scenario, the formulation of new proposals in the negotiation to become VE partners, leads to the inclusion of two important features: how to calculate the reward value and what part of the exploration space to consider.

The reward value for a particular state is calculated according to the qualitative feedback received from the market agent, in response to the proposal derived from this state [see formula (3)].

$$r = \begin{cases} n & , \ if\ winner \\ \frac{n}{2} - \sum_i penalty_i & , \ if\ not winner \quad (0 \le penalty_i \le 1) \end{cases} \qquad (3)$$

The exploration space, which can became very large and thus implies a long time to learn, is reduced in order to include only those actions that can be considered as promising actions. A promising action is an action that can be applied to a previous state proposed to the market agent hints included in the feedback formulated by this agent.

Distributed Dependencies Resolution

One of the requirements for the negotiation protocol we are here proposing, besides dealing with intradependencies of attributes, is the capability to deal with attributes' interdependencies. This is an important requirement to be considered in our scenario, because in the VE formation process, interdependent negotiations take place simultaneously, and proposals received from different organization agents may have incompatible dependent attributes' values. Therefore, agents should negotiate in order to agree between them on mutual admissible values, what can be seen as a distributed dependencies satisfaction problem.

The distributed dependencies satisfaction problem has been the subject of attention of other researchers, addressing the study of both single (Yokoo, Durfee, Ishida, & Kuwabara, 1992) and multiple dependent variables (Armstrong & Durfee, 1997; Parunak, Ward, & Sauter, 1999). In the VE formation process, dependencies may occur between multiple variables, making the latter approaches more relevant to our research. The first mentioned paper (Armstrong & Durfee, 1997) describes algorithms to reach one possible solution, not the optimal one. The second paper (Parunak et al., 1999) introduces an algorithm that, although reaching the optimal solution, imposes that all agents involved in the mutual dependencies resolution process have to know all agents' private utility functions.

Different from all these proposals, our distributed dependencies satisfaction algorithm, besides reaching the optimal solution, keeps agents' information as much as possible private.

Each agent involved in the distributed dependent problem resolution should know its space of states, that is, all possible values for its own dependent attributes. Agents will then exchange between them alternative values for the dependent attributes, in order to approach an agreement. As in any iterative negotiation process, agents start the negotiation by proposing its optimal (from a local point of view) solution and, in the next rounds, start conceding trying to reach a consensus.

In order to properly understand the way the algorithm works, first we should introduce the concept of "decrement of the maximum utility" of an alternative state. State transitions are due to relaxation of one or more state variables. The decrement of the maximum utility of a particular alternative proposal can be calculated as the difference between the evaluation values of this alternative proposal and the optimal one. We will abbreviate "decrement of the maximum utility" to "decrement of the utility," meaning that the successive amount of utility agents has to concede compared to the (local) optimal bid. Formula (4) represents the decrement of utility for agent *i*, corresponding to the particular state s^k, where $s*$ is the agent's optimal state (proposal).

$$du_i^k = Ev(s^*) - Ev(s^k) \quad (4)$$

At each negotiation step, the agent selects as a new proposal the one that has the lowest decrement of the utility of those not yet proposed. During the negotiation process, agents do not reveal their own state's utility, but only the state's decrement utility, which enables keeping important information private.

This process ends when all agents cannot select a next state better than one already proposed. In this way, agents, although remaining self-interested, will converge for a solution that is the best possible for all of them together, because it represents the minimum joint of decrement of the utility. The joint decrement of the utility is calculated according to formula (5):

$$jdu^k = \sum_{dag} du_{dag}^k \quad , dag = \{1 \ldots n\} \text{ set of mutual dependent agents} \quad (5)$$

After agreement in a global solution, agents involved in the dependencies resolution process generally get different local decrement of utility values, and, therefore, some agents become more penalized than others. In order to guarantee that all agents involved in the distributed dependencies resolution get the same real decrement of utility (*rdu*), the joint decrement of the utility will be distributed between them according to formula (6):

$$rdu = \frac{jdu^m}{n}, \quad n = \text{number of agents} \quad (6)$$

As a consequence, some agents have to pay or get a compensation value to others. Once agent *i* has previously calculated dui^m as its local decrement of utility, the compensation value is calculated according to formula (7):

$$cValue_i = rdu - du_i^m \quad (7)$$

If the agent's real decrement of the utility is greater than its local decrement of the utility, it will pay a compensation value to others, that is calculated as the difference of these two values. If not, the agent will get a compensation value.

Through all the steps mentioned before (multiattribute bid evaluation, learning in negotiation, and distributed dependencies resolution), the VE formation stage is accomplished. In the next sections, we will discuss how to formalize through an e-contract all the commitments that have been made.

E-Contracting

The result of the negotiation process leading to the VE formation should be "compiled" in an electronic contract that establishes rights and duties, as well as associated penalties, for all the individual enterprises included in the final agreement. This electronic contract can be used in the VE operation stage, for the sake of monitoring all the VE activities.

In B2B electronic commerce, more attention has been given recently to contract formation and fulfillment. In fact, this issue is part of the so-called B2B life cycle model, as presented in He, Jennings, and Leung (2003). Approaches to B2B contract handling (e.g., Goodchild, Herring, & Milosevic, 2000) identify the need to specify and represent contracts, and further to monitor and enforce them.

Figure 4 shows some of the services that may be available within an electronic institution. When considering contracts as the result of a business negotiation process, we can identify certain typified relations that can be assisted through the use of contract templates. After the negotiation phase, the obtained contract must be checked for compliance to existing business norms; it is then registered with a notary. The business relation is then carried out, and services like contract monitoring and enforcement may be provided, ensuring coherent behavior between the parties and registering the fulfillment of transactions. In the following subsections, we develop these issues.

Figure 4. E-contracting services

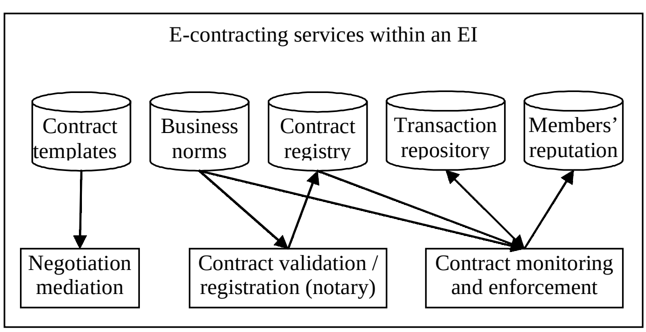

E-Contracts and Norms

Contracts are formalizations of the behavior of a group of agents that jointly agree on a specific business activity. Contracts are used as a means of securing transactions between the involved parties, forming a normative structure that explicitly expresses their behavior interdependencies. *Electronic contracts* are virtual representations of such contracts. The aim of e-contracting is to improve the efficiency of contracting processes, supporting an increasing automation of both e-contract construction (using automated tools) and execution (integrating with business processes).

The components of a contract include the identification of the participants, the specification of the products and/or services included, and a discrimination of the actions to be performed by each of the participants. These actions are normally accompanied with time and precedence constraints. Typified business relations can recurrently use preformatted contracts. In this case, contracts usually have a set of identified *roles* to be fulfilled by the parties involved in the relation.

The core of a contract is composed of contract clauses. These clauses can specify different types of behavior *norms* that will guide the interaction between the parties. This normative conception of contracts is generally adopted (e.g., Dignum & Dignum, 2001; Kollingbaum & Norman, 2002; Sallé, 2002). Broadly speaking, three types of norms can exist within a contract structure:

- *Obligation*: an agent has an obligation to another agent to bring about a certain state of affairs (by executing some action), before a certain deadline
- *Permission*: an agent is allowed to execute some action, within a given window of opportunity (specified either by a deadline or more generally by a state of affairs)
- *Prohibition*: an agent is forbidden to bring about a certain state of affairs (some action is interdicted)

A formal approach to model such norms is *deontic logic* (von Wright, 1950; Meyer, 1988), which is also known as the logic of normative concepts, a branch of modal logic. The normative concepts obligation, permission, and prohibition are analogous to the modal concepts of necessity, possibility, and impossibility, respectively.

When representing contracts, another fundamental concept is typically added to the norms above: the *sanction*. Any obligation must be accompanied by at least one sanction, as obligations without sanctions are ineffective (Kollingbaum & Norman, 2002). Thus, obligations are not absolute but are relative to their associated sanctions in case of nonperformance (Sallé, 2002). Prohibitions can be addressed in an analogous way. A prohibition is sometimes handled as a negated obligation, that is, a duty for not performing some action (see, for instance, Kollingbaum & Norman, 2002).

Approaches to the automation of contractual relationships necessarily include this sanction component. Particularly when that automation is based on the autonomous agent paradigm, norms cannot be taken as constraints on the behavior of each contractual party. Each agent is able to deliberatively reason about its goals and the norms to which it has committed; hence, the notion, in Castelfranchi, Dignum, Jonker, and Treur (2000), of *deliberative normative agents*. An agent can violate a norm in order to accomplish a private goal that it considers to be more important. When doing so, the agent is aware of the sanction to which it will be subject. Some researchers address the advantages of anticipating sanctions (also called decommitment penalties) in multiagent contracting, introducing the concept of a *levelled commitment contract* (Sandholm & Lesser, 2001), and study reasoning decision processes that consider strategic breaches (Sandholm & Lesser, 2001; Excelente-Toledo, Bourne, & Jennings, 2001).

Norms and Electronic Institutions

Contracting is normally subject to contract law. This law is enforced by the court and can be seen as a normative system that contracts must respect. Furthermore, we can conceptualize norms at different levels of abstraction. For instance, we can consider those that are applied to contracts in general (thus being inherited in all established contracts), those that refer to particular contractual domains, and those that are created when a specific business relation is formalized.

Electronic institutions, while regulating the interactions that can take place between agents, can represent normative systems that limit the behavior of participants and describe the penalties incurred when norms are violated (Dignum, 2001). Contractual relations created inside the institution must abide to the imposed norms, specifying the details of a particular business relation. This two-level conception of normative agent interactions is proposed by some researchers. In Dignum and Dignum (2001), the authors model a society of agents distinguishing between an institutional level (where social norms and rules are specified) and an operational level (dependent on the goals of each agent).

E-Contracting Life Cycle

Any contractual relationship can be said to evolve through a number of steps. These can be resumed to the following three stages:

- *Information discovery*: Clients find potential suppliers.
- *Contract negotiation*: The parties negotiate the contract terms—the result of this stage is a legally binding contract, reflecting the agreement made.
- *Execution*: The contract terms are fulfilled by the parties, namely involving product delivery or service rendering, and the corresponding payments.

The first stage thus comprises the brokering phase of B2B electronic commerce. One can also conceptualize it as a *precontractual phase* involving a definition of the products or services sought/sold by clients/suppliers, and the utilization of yellow-pages services allowing potential partners to contact each other.

The second stage is devoted to the negotiation of the terms of an agreement—it is the *contractual phase*, because a contract is being constructed. That agreement will express a number of steps to be performed by the contractual parties. Hence, the parties negotiate not only attributes of products/services (as explained in the ForEV platform) but also details of how those products/services will be delivered/rendered and paid. The document that represents the agreement reached is a legally binding contract, signed by those involved. Typically, it will also specify how to handle exception conditions, such as those related to nonfulfillment of duties (e.g., late delivery or nonpayment).

The third stage is the *postcontractual phase*, that is, after the contract is established, it is time to proceed as agreed. It is also referred to as the fulfillment phase. In more complex and integrated interactions, the parties involved will eventually engage their business processes, forming an interenterprise workflow.

Representing E-Contracts

A normative conception of contracts is normally used for contract representation. Hence, languages for representing norms in contracts have been proposed.

Normative statements, based on the operators of deontic logic, can be formally represented as follows (Sallé, 2002):

$$\text{ns: } \varphi \rightarrow \theta_{s,b} (\alpha < \psi)$$

where *ns* is a label; φ is an activation condition; θ is a deontic operator (obligation, permission, or prohibition); s is the subject of q; b is the beneficiary of q; α is the action to perform or the state of affairs to bring about; and ψ is a deadline.

In this approach, obligations are not absolute but are relative to their associated sanctions. That is, deviation from prescribed behavior is admitted and properly addressed through sanctions. These are defined just like the other normative statements, but by specifying as the activation condition the nonfulfillment of a given obligation. Sanctions may give rise to other obligations or prohibitions: either the beneficiary of the violated norm is granted a right (the subject has a new obligation toward the beneficiary) or the subject of the violated norm is refused a right (he is forbidden to do something).

A number of standards for contract representation are also emerging, mostly founded on rule-based markup languages (see Angelov & Grefen, 2001, for a survey). Developing on these standards, Grosof and Poon (2002) described the SweetDeal system as a rule-based approach to the representation of business contracts. Emerging Semantic Web standards for knowledge representation of rules (RuleML) are combined with ontology representation languages (DAML+OIL).

E-Contract Negotiation

When the outcome of the negotiation phase is to be formally represented and (eventually) automatically executed, it is beneficial to consider this formal representation in the negotiation phase.

The need for a starting ground in contracting is acknowledged by several researchers (see, for instance, Kollingbaum & Norman, 2002; Sallé, 2002). In fact, starting a negotiation where nothing is fixed represents a problem that is too ill-structured to consider automating. The importance of a *contract template* resides on its ability to provide a structure on which negotiation can be based.

Certain kinds of business relations are formally typified (for instance, sales and purchases). In this sense, instead of beginning from scratch a new contractual relation, two (or more) agents can use an electronic contract template, which is a contract outline containing 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). If all goes well, the result of the negotiation will be an actual contract, instantiated from the template, that will be signed by the parties. Templates thus provide a structure that allows negotiation, as a process of cooperative construction of a business relation, to be focused on those elements that, when instantiated, will distinguish the agreement obtained from other contractual relationships. Meanwhile, the common elements in relations of the same type will be preserved. These common elements might include, for example, outlined commitments of the involved parties, which when instantiated through negotiation will detail their concrete objects (eventually including technical properties) and temporal references.

As already described within the *ForEV* framework, negotiation mediation services are important mechanisms that allow business agreements to be obtained in a regulated fashion. Preestablished protocols can be used, taking advantage of template structures and ensuring that resulting contracts are in accordance with business norms.

Electronic institutions can provide means of validating contracts, by checking them against existing institutional norms. Valid (and signed) contracts are then stored using notary services, in order to ensure their legal existence.

E-Contract Execution

The execution of an e-contract consists of the parties following the norms they committed to when signing the contract. If any deviations from the prescribed behavior should occur, sanctions can be applied as specified in the contract or in its normative system of reference. However, the parties involved will typically not voluntarily submit themselves to such penalties. Therefore, appropriate mechanisms are needed to monitor and enforce norm execution. Within the framework of electronic institutions, monitoring and enforcement services can be rendered by the institution. Only a trusted third party can enable the necessary level of confidence between the parties involved in a business relation.

In Kollingbaum and Norman (2002), a supervised interaction framework is proposed, where a trusted third party is included as part of any automated business transaction. Agents are organized in three-party relationships between two contracting individuals (a client and a supplier) and an authority that monitors the execution of contracts, verifying that errant behavior is either prevented or sanctioned. This authority enables the marketplace to evaluate participants, keeping reputation records on the basis of past business transactions.

In Sallé (2002), a *contract fulfillment protocol (CFP)* is proposed, a collaborative protocol based on the normative statements' life cycle. The idea is that as contractual relationships are distributed, there is a need to synchronize the different views each agent has about the fulfillment of each contractual commitment. Each norm has a set of states it might go through. For instance, an obligation is first *agreed* (when the contract is signed), it then becomes *pending*, and later on might be *refused* (triggering appropriate sanctions) or *accepted*. In the latter case, it will become *in progress* and afterwards *executed*. When executed, it might be *rejected* (again requiring correction measures) or considered as *fulfilled*. Agents use this life cycle to communicate their intentions on fulfilling contractual norms, allowing their contractual partners to know what to expect from them. This ability is referred to as *dynamic forecasting of partners' behavior*, and it permits a fluent and prompt execution of contracts, as agents do not have to wait for the fulfillment of their partners' obligations to start executing their own (hence the collaborative nature).

The real-world application of agents in automated contract fulfillment is challenged by the presence of complex legal issues and subjective judgments of agent compliance (He et al., 2003). Some work on these matters has been made, for instance, in Daskalopulu, Dimitrakos, and Maibaum (2001), where an e-market controller agent (a third party) is suggested to resolve disputes arising from subjective views on contract compliance, thereby playing the role of a judge. This agent holds a representation of the contract, and when a conflict occurs, it collects evidence from the involved parties and obtains information from independent advisors, such as certification authorities, regulators, or controllers of other associated markets.

FUTURE TRENDS AND CONCLUSIONS

ForEV is an agent-based tool we have developed aiming at facilitating automatic partners selection in the context of VEs.

Appropriate negotiation protocols, for multiattribute evaluation, keeping information private as much as possible, and solving mutual constraints between attributes, are efficient tools to be used by agents as delegates of enterprises in finding the best temporary consortium to respond to an opportunity of business.

The VE formation stage has been tested in a simplified textile example using the *ForEV* framework, and a coherent consortium was established.

However, we soon realize that these facilities need to be made available in the context of a larger framework representing some secure, trustful institution, responsible for supervising the entire VE life cycle.

An electronic institution, as we are proposing here, encompasses all the facilities needed to help in the VE formation and operation processes, making it possible to follow these steps and enforcing agents to comply with norms and rules according to their specific roles in the VE.

Our next move will be the inclusion of a real, flexible although compliant electronic contract, as a result of the negotiation process between all the parties, so as to be explored during the VE operation phase. The electronic contract will be the guarantee that what was previewed and agreed upon is being accomplished and that the right measures will be taken whenever agent misbehavior occurs. Also, learning facilities to derive new rules of behavior from past events and situations have to be included in the electronic institution, enabling both an evolution along time and some specialization of the general rules for specific scenarios.

We see, thus, *ForEV* as a seed for a more complex tool helping VEs along their life cycles.

REFERENCES

Angelov, S., & Grefen, P. (2001). *B2B eContract Handling—A Survey of Projects, Papers and Standards*. University of Twente: CTIT Technical Reports.

Armstrong, A., & Durfee, E. (1997). Dynamic Prioritization of Complex Agents in Distributed Constraint Satisfaction Problems. In M. Pollack (Ed.), *Proceedings of the 15th International Joint Conference on Artificial Intelligence* (620-625), Nagoya, Japan. Morgan Kaufmann Publishers.

Camarinha-Matos, L., & Lima, C. (1998). A Framework for Cooperation in Virtual Enterprises. In J. Mills & F. Kimura (Eds.), *Information Infrastructure Systems for Manufacturing* (pp. 305–321). Dordrecht: Kluwer.

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). Heidelberg: Springer.

Daskalopulu, A., Dimitrakos, T., & Maibaum, T. (2001). E-Contract Fulfilment and Agents' Attitudes. Presented at *Proceedings of the ERCIM WG E-Commerce Workshop on the Role of Trust in e-Business*. Zurich, Switzerland.

Dignum, F. (2001). Agents, Markets, Institutions and Protocols. In F. Dignum & C. Sierra (Eds.), *Agent Mediated Electronic Commerce: The European Agentlink perspective* (pp. 98–114). Heidelberg: Springer.

Dignum, V., & Dignum, F. (2001). Modelling Agent Societies: Co-ordination 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). Heidelberg: Springer.

Excelente-Toledo, C. B., Bourne, R. A., & Jennings, N. R. (2001). Reasoning about Commitments and Penalties for Coordination Between Autonomous Agents. In E. André, S. Sen, C. Frasson, & J. P. Müller (Eds.), *Proceedings of the Fifth International Conference on Autonomous Agents* (pp. 131–138). New York: ACM Press.

Fischer, K., Muller, J., Heimig, I., & Scheer, A. (1996). Intelligent Agents in Virtual Enterprises. In B. Crabtree and N. Jennings (Eds.), *Proceedings of the 1th International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology* (pp. 205-223), London, UK. The Practical Application Company Ltd.

Goodchild, A., Herring, C., & Milosevic, Z. (2000). Business Contracts for B2B. In H. Ludwig, Y. Hoffner, C. Bussler, & M. Bichler (Eds.), *Proceedings of the CAISE'00 Workshop on Infrastructure for Dynamic Business-to-Business Service Outsourcing (ISDO'00)* (pp. 63–74). CEUR Workshop Proceedings.

Grosof, B., & Poon, T. C. (2002). Representing Agent Contracts with Exceptions using XML Rules, Ontologies, and Process Descriptions. In *Proceedings of the International Workshop on Rule Markup Languages for Business Rules on the Semantic Web, First International Semantic Web Conference*. Sardinia, Italy.

He, M., Jennings, N. R., & Leung, H. (2003). On agent-mediated electronic commerce. *IEEE Transactions on Knowledge and Data Engineering, 15*(4), 985–1003.

Jagers, H., Jansen, W., & Steenbakkers, G. (1998). Characteristics of Virtual Organisations. *Working Paper 98-02*. Department of Information Management, University of Amsterdam, The Netherlands.

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 First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 1* (pp. 272–279). New York: ACM Press.

Meyer, J. -J. C. (1988). A Different Approach to Deontic Logic: Deontic Logic Viewed as a Variant of Dynamic Logic. *Notre Dame Journal of Formal Logic, 29*(1), 109–136.

Parunak, H., Ward, A., & Sauter, J. (1999). The MarCon Algorithm: A Systematic Market Approach to Distributed Constraint Problems. *Artificial Intelligence for Engineering Design Analysis and Manufacturing Journal, 13*, 217–234.

Peterson, S., & Gruninger, M. (2000). An Agent-Based Model to Support the Formation of Virtual Enterprises. In *Proceedings of the International ICSC Symposium on Mobile Agents and Multi-Agents in Virtual Organisations and E-Commerce*. Woolongong, Australia. ISC Academic Press.

Rocha, A., & Oliveira, E. (2001). Electronic Institutions as a Framework for Agents' Negotiation and Mutual Commitment. In P. Brazdil & A. Jorge (Eds.), *Proceedings of the Second Workshop on Multi-Agent Systems, Logic Programming, and Constraint Solving* (pp. 232-245). Springer.

Sallé, M. (2002). Electronic Contract Framework: An Overview. In *Proceedings of the AAAI-2002 Workshop on Agent-Based Technologies for B2B Electronic Commerce*. Edmonton, Alberta, Canada.

Sandholm, T. W., & Lesser, V. R. (2001). Leveled Commitment Contracts and Strategic Breach. *Games and Economic Behavior, 35,* 212–270.

Shmeil, M., & Oliveira, E. (1997). The Establishment of Partnerships to Create Virtual Organisations: A Multi-Agent Approach. In L. Camarinha-Matos (Eds.), *Re-engineering for Sustainable Industrial Production* (pp. 284–294). London; New York: Chapman & Hall.

Vulkan, N., & Jennings, N. (1998). Efficient Mechanisms for the Supply of Services in Multi-Agent Environments. In *Proceedings of the First International Conference on Information and Computing Economics* (pp. 1-10). Charleston, SC. ACM Press.

von Wright, G. (1950). Deontic logic. *Mind, 60*, 1–15.

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

Yokoo, M., Durfee, E., Ishida, T., & Kuwabara, K. (1992). Distributed Constraint Satisfaction for Formalizing Distributed Problem Solving. In *Proceedings of the 12th International Conference on Distributed Computing Systems.* Yokohama, Japan. IEEE Computer Society.