Electronic Institutions as a framework for Agents' Negotiation and mutual Commitment

Ana Paula Rocha, Eugénio Oliveira

LIACC, Faculty of Engineering, University of Porto,
R. Dr. Roberto Frias, 4200-465 Porto, Portugal
{arocha,eco}@fe.up.pt

Abstract. Electronic transactions are of increasing use due to its openness and continuous availability. The rapid growth of information and communication technologies has helped the expansion of these electronic transactions, however, issues related to security and trust are yet limiting its action space, mainly in what concerns business to business activity. This paper introduces an Electronic Institution framework to help in electronic transactions management making available norms and rules as well as monitoring business participants' behaviour in specific electronic business transactions. Virtual Organisation (VO) life cycle has been used as a complex scenario encompassing electronic transactions and where Electronic Institution helps in both formation and operation phase. A flexible negotiation process that includes multi-attribute and learning capabilities as well as distributed dependencies resolution is here proposed for VO formation. "Phased commitment" is another concept here introduced for VO operation monitoring through the Electronic Institution.

1 Introduction

An Electronic Institution is a framework for enabling, through a communication network, automatic transactions between parties according to sets of explicit institutional norms and rules.

An Electronic Institution helps on both providing tools and services for and on supervising the intended relationships between the parties.

Usually, parties engaged in electronic transactions and joint actions are software agents mediated. We therefore believe that an appropriate Electronic Institution can be implemented as an agent-based framework where external agents can meet together according to a set of established and fully agreed mutual constraints.

A good example illustrating the need of an Electronic Institution can be found in a software framework providing the automatic services needed for helping on the Virtual Organisations' life cycle.

A Virtual Organisation (VO) is an aggregation of autonomous and independent organisations connected through a network (possibly a public network like Internet) and brought together to deliver a product or service in response to a customer need. Virtual Organisation management should be supported by efficient information and communication technology through its entire life cycle.

Tools for Virtual Organisations formation process, through the use of an Electronic Market providing enhanced protocols for appropriate negotiation can easily be accommodate into the Electronic Institution available services. In our approach, these services include automatic capabilities for adaptive bid formulation, accepting a qualitative feedback for the sake of keeping the information as much as possible private to each one of the negotiating agents, as well as multi-issue bid evaluation.

Contrary to other approaches [1, 2] that use a-priori fixed values for weighting attributes' values in the bid evaluation function, we here advocate weighting those values, reflecting the deviation from the preferred values, according to the relative importance of the attributes. Further more, the negotiation protocol we are introducing, here called Q-Negotiation, includes the capability for Agents that are trying to supply mutually constrained items, to deal with the respective inter-dependencies while keeping their self-interestedness.

During the self-interested agents Q-Negotiation process, they may have to agree on supplying sets of mutually constrained items whose values, although not being optimum for each one of the individual agents, correspond to a minimal joint decrement of the agents' maximum utility. However, it would not be fair that this agreement, in the name of the best possible joint utility, benefits one agent more than the other. We therefore propose that, in the case of agreements based on the agents' joint utility, the agents concerned should equally distribute the joint decrement of their maximum utility. Through the calculation of appropriate compensations, agents that have the most beneficial, know they have to transfer some utility to those who have the least.

The outcome of the Virtual Organisation formation stage is, for our proposed Electronic Institution, a "phased commitment" through which different parties commit themselves to specific future actions. We intend to link commitments, which are the result of previous negotiation during the Virtual Organisation formation stage, to the next stage, the VO operation stage through a monitoring process.

Commitments have to be agreed and represented in such a way that at several different future points in time, they can be verified. What follows after that verification should also be considered in the agreed contract. This paper introduces an appropriate framework we find suitable for dealing, in a safe and structured way, with Virtual Organisations life cycle. We also introduce the concepts of both Meta Institution and Electronic Institution which is the responsible for making available functionalities including an adaptive multi-criteria based negotiation protocol together with features for solving agents mutual dependencies. We also introduce the notion of "phased commitment" and its importance for subsequent stages of VO life cycle. An Electronic Institution also provides is the Knowledge Representation capabilities suitable for the Virtual Organisation activities.

This paper includes, besides the introduction, four more sections. The next section introduces the concepts of Electronic Institutions and Meta Institutions in the context of Virtual Organisations. A third section details our proposed Q-Negotiation algorithm used in the VO formation stage. Section 4 briefly describes the knowledge representation needed for agents (as enterprise delegates) interoperability. Section 5 describes phased commitments in the VO operation stage and, finally, section 6 gives some conclusions and directions for future work.

2 Electronic Institutions and Meta Institutions

2.1 Meta-Institutions

It seems intuitive that when automated agents' interactions become more sophisticated and agents' autonomy more evident, a problem related with confidence, trust and honesty may arise. Moreover, agreements like deals made by different companies or their delegates (automated or not) always claim for a common and non-ambiguous ground of understanding. We need to design a framework, accepted by all the parties, to encompass all the automatic activities that are going to take place between agents representing different individual or collective entities.

An Electronic Institution is a framework for enabling, through a communication network, automatic transactions between parties according to sets of explicit institutional norms and rules.

An Electronic Institution helps on both providing tools and services for and on supervising the intended relationships between the parties.

However, each Electronic Institution will be, at least partially, dependent on the specific application domain of business it has been designed for. Our proposal starts with the definition of a Meta-Institution, which has to be independent of the application domain.

We see a Meta-Institution as a shell for generating specific Electronic Institutions. A Meta-Institution is a set of Electronic facilities to be requested and used in order to help on the creation of suitable Electronic Institutions according to a set of established social rules and norms of behaviour. Those rules of behaviour should apply to many different kinds of (automatic) entities interacting through the web. Besides enforcing those general rules into the social interactions, a Meta-Institution also provides tools for some important stages of the interaction process.

The main goal of a Meta-Institution is, however, to be able to make available suitable Electronic Institutions that will leave all along a particular business process in a specific application domain.

Electronic transactions between distributed and autonomous entities are becoming more and more software agents mediated. We therefore believe that an appropriate Electronic Institution can be implemented as an agent-based framework where external agents can meet together according to a set of established and fully agreed norms, rules, and mutual constraints.

2.2 Virtual Organisations and Electronic Institutions

A good example illustrating the need of an Electronic Institution (EI) can be found in providing the automatic services needed for helping on the Virtual Organisations' life cycle.

This aggregation of autonomous organisations is advantageous in the sense that it will reduce complexity – today's products and services are increasingly complex and require close coordination across many different disciplines – and most important, will enable the response to rapidly changing requirements. Virtual Organisation will only exist for a temporary time duration, that is the time needed to satisfy its purpose.

The VO life cycle is decomposed in four phases [3, 4] that will also be reflected in the EI framework:

1. *Identification of Needs*: Appropriate description of the product or service to be delivered by the VO, which guides the conceptual design of the VO.
2. *Partners Selection*: Automatic selection of the individual organisations (partners), which based in its specific knowledge, skills, resources, costs and availability, will integrate the VO.
3. *Operation*: Control and monitoring of the partners' activities, including resolution of potential conflicts, and possible VO reconfiguration due to partial failures.
4. *Dissolution*: Breaking up the VO, distribution of the obtained profits and storage of relevant information for future use of the Electronic Institution.

We expect from a Meta-Institution, the application independent framework we have defined in section 2.1, the capability to directly help on the "Identification of needs" stage of the VO life cycle, as follows:

– First helping agents in describing their needs in such a way that can be understood by other potential members that may join later a specific and appropriate Electronic Institution and
– Second, providing the searching tools to look for potential partners who know how to achieve those described needs.

The Meta-Institution will then generate that specific Electronic Institution for the particular application domain through the instantiation of some modules according to the explicit VO goals.

We can see both the general architecture and the role of a Meta-Institution in the light of the emergence of Virtual Organisations as it is depicted in the figure below:

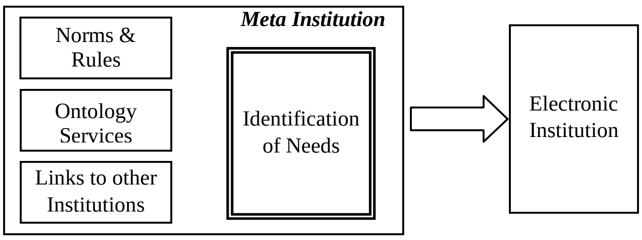

Fig. 1. General architecture and role of a Meta-Institution

The Electronic Institution, inherits appropriate rules as well as important links to other institutions that may play a crucial role for all the VO process, that are inherited from the Meta-Institution, will provide the framework for dealing with, at least, the two main stages of the VO life cycle: VO formation and VO operation monitoring.

We will describe in detail, in the next two sections, how an Electronic Institution provides the means for helping on those two stages by making available an electronic market agent-based tool together with appropriate knowledge representation mechanisms.

Figure 2, below, shows the main modules of an Electronic Institution putting in evidence that, at the end of the VO formation process, a special commitment will be reached. It is this result of the first stage that will make it possible for the Electronic Institution to be useful in the next step of the VO life cycle – VO operation.

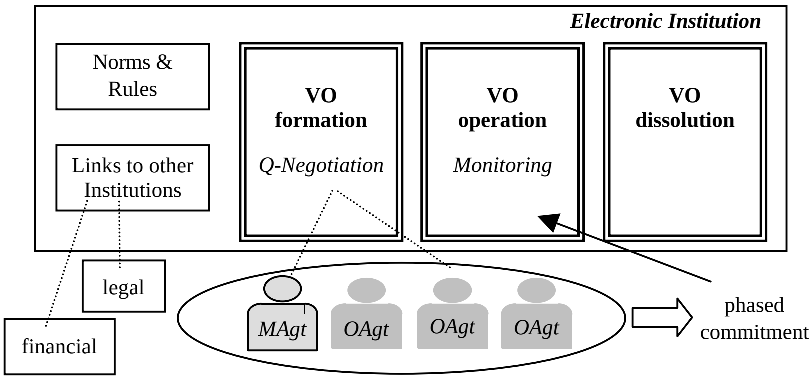

Fig. 2. General architecture and role of an Electronic Institution

3 Advanced features for Negotiation in VO formation

In the scenario of an agent-based Virtual Organisation formation process, the negotiation mechanism should enable the selection of the individual organisations that, based on their own competencies and availability, will constitute the optimal group to satisfy the previously described VO needs. In such a scenario, the adopted automatic negotiation mechanism has to be powerful enough to satisfy three important requirements:

- Capabilities for negotiating about what are the most promising organisations that should belong to VO. This means that agents, as representative of organisations, have to negotiate over goods or services those organisations are able to provide. In realistic scenarios, goods/services are described through multiple attributes, which imply that the negotiation process must be enhanced with the capability to both evaluate and formulate multi-attribute based proposals.
- Agents that are willing to belong to the Virtual Organisation may compete between them. An agent does not know other ones' negotiation strategies or even current proposals. Learning techniques can help agents to better negotiate in these partially unknown environments, by reasoning about past negotiation episodes as well as about the current one improving its own behaviour.
- In the VO formation process, each one of the individual organisations will contribute with at least one of its own capabilities (good or service) to the VO. All these contributions may be, and they usually are, mutually dependent. The negotiation process has to be able to deal with those inter-dependencies, reaching a coherent solution as the final one to be accepted by all the agents.

An important characteristic that must be considered in VO scenario is the fact that any organisation has as its main objective to maximize its own profit. In order to do that, negotiation process has to take into account agents individual rationality and self-interestedness. Keeping information private prevents loosing negotiation power to competitors, since others will never know or be able to deduce how close they are to another agent's preferences.

It is our claim that our proposed negotiation algorithm can effectively deal with these three important requirements in the VO scenario while keeping agents' information private as much as possible. In the next sections, we further detail the proposed model for agents' negotiation. First, we present a formal description of our negotiation model. Then we describe how to specifically deal with each one of the three requirements mentioned above (multi-attribute bidding, adaptive bidding and mutual dependencies resolution).

3.1 The Negotiation Model

In the VO formation process, participants in the negotiation can be either market or organisation agents. The Market Agent plays the role of organizer, meaning that it is the agent that starts and guides all the negotiation process. The Organisation Agents play the role of respondents, meaning that they are those who are willing to belong to the future VO and, therefore, they have to submit proposals during the negotiation phase.

In order to agree in a VO structure, agents (Market and several Organisations) naturally engage themselves in a sequential negotiation process composed of multiple rounds of proposals (sent by Organisations to Market) and counter-proposals which are actually comments to past proposals (sent by Market to Organisations).

The Market Agent playing a central role as organizer, models the negotiation process through the Neg^{MA} triplet as follows:

$$\mathrm{Neg}^{MA} = <\mathrm{Cmpt, LAgts, H}> \qquad (\mathbf{1})$$

where:

- *Cmpt* identifies the component under negotiation
- *LAgts* is the list of respondent (organisation) agents that can provide component *Cmpt*.
- *H* is the negotiation history. Each element of H contains information related to a single negotiation round. Each negotiation round includes all proposals received during that round.

$$\mathrm{H} = \{\mathrm{H}_t\}, \quad \mathrm{Ht} = \{<\mathrm{Prop}_{ti}, \mathrm{Eval}_{ti}>\} \qquad (\mathbf{2})$$

$$\mathrm{Prop}_{ti} = \{\mathrm{V}_1, \ldots, \mathrm{V}_n\}$$

where:

- $Prop_{ti}$ is the proposal sent by organisation agent i, in the negotiation round t.
- Val_x is the proposal's value of attribute *x*.
- $Eval_{ti}$ is the evaluation value of proposal $Prop_{ti}$, from the Market Agent point of view.

Each one of the Organisation Agents model the negotiation process through the Neg^{OA} n-uple as follows:

$$Neg^{OA} = \langle Cmpt, MAgt, H, Q \rangle \quad (\mathbf{3})$$

- *Cmpt* identifies the component under negotiation
- *MAgt* identifies Market Agent that is the organizer of this negotiation process.
- *H* is the negotiation history. Each element of H contains information related to a single negotiation round. Each negotiation round includes the proposal sent by each specific Organisation Agent to the Market Agent plus the feedback comment received from the Market Agent.

  $$H = \{\langle Prop_t, Comment_t \rangle\}, \quad Comment_t \in \{winner, \langle Eval_{t1}, \ldots, Eval_{tn} \rangle\} \quad (\mathbf{4})$$

  where:
  - $Prop_t$ is the proposal sent during the negotiation round *t*.
  - $Comment_t$ is the comment received from Market Agent to proposal $Prop_t$. This comment indicates if the proposal is either the winner in the current round (*winner*) or includes a qualitative appreciation for each one of the attribute values under negotiation.
- The *Q* parameter includes relevant information to be used by the learning algorithm used for next bid formulation. This particular topic is discussed in a later section. At this point, it is only important to say that *Q* is described as follows:

$$Q = \{Q_t\}, \quad Qt = \{\langle State_t, Action_t, QValue_t \rangle\}$$

3.2 Multi-Attribute Bid Evaluation

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

Attaching utility values to different attributes under negotiation solves the problem of multi-attribute evaluation. Generally, an evaluation formula is a linear combination of the attributes' values weighted by their corresponding utility values. In this way, a multi-attribute negotiation is simply converted in a single attribute negotiation, where the result of the evaluation function can be seen as this single issue. Examples of this method are presented in [1, 2].

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 multi-attribute function presented in formula (5) encodes the attributes' and attributes values' preferences in a qualitative way and, at the same time, accommodates attributes intra-dependencies.

$$Ev = \frac{1}{Deviation}, \quad Deviation = \frac{1}{n} * \sum_{i=1}^{n} \frac{i}{n} * dif(PrefV_i, V_i), \qquad (5)$$

where:

n = number of attributes that defines a specific component,

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

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

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 one 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, since 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 where Organisation 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 VO. The Market Agent helps Organisation Agents in their task of formulating new proposals by giving them some hints about the direction 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 a 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 Organisation Agent that there is a better proposal in 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' attributes values, which can be classified in one of three categories: *sufficient*, *bad* or *very_bad*.

Organisation 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.

3.3 Learning in Bid Formulation

The Q-Negotiation algorithm uses a reinforcement learning strategy based in Q-learning for the formulation of new proposals. The Q-learning algorithm [6] is a well known reinforcement learning algorithm that maps evaluation values (Q-values) to pairs state/action.

The selection of a reinforcement learning algorithm seems to be appropriate in the negotiation process that conduits to the VO formation, since organization agents evolve in an, at least, partially unknown environment. And in particular, Q-learning enables on-line learning, which is an important capability in our specific scenario where agents will learn in a continuous way during all 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 in the idea of rewarding actions that produces good results, and punishing those that produce bad results, as indicated by parameter *r* in the correspondent formula (see equation (6)).

$$Q(s,a) \quad = \quad Q(s,a) + \alpha \left(r + \gamma \max_{b} Q(s',b) - Q(s,a) \right) \tag{6}$$

In the Q-Negotiation process, we assume that:

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

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

- 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.

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

The adaptation of the Q-learning algorithm to our specific scenario, the VO formation negotiation, leads to the inclusion of two important features we will briefly enumerate in next paragraphs, and are detailed elsewhere [6].

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 7).

$$r \quad = \quad \begin{cases} n & , \quad \textit{if winner} \\ \dfrac{n}{2} - \sum\limits_{i} \textit{penalty}_i & , \quad \textit{if not winner} \quad (0 \le \textit{penalty}_i \le 1) \end{cases} \tag{7}$$

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. As an example, if the Market Agent, as a proposal's feedback, classifies the value of attribute *x* as *bad*, one promising action should be increase a little bit this attribute and maintain all the others.

3.4 Distributed Dependencies Resolution

One of the requirements for the negotiation protocol we are here proposing, besides dealing with attributes intra-dependencies, is the capability to deal with attributes' inter-dependencies. This is an important requirement to be considered in our scenario, because in the VO formation process interdependent negotiations take place simultaneously, and proposals received from different organisation 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 [7] and multiple dependent variables [8, 9, 10]. In the VO formation process, dependencies may occur between multiple variables, making the latter approaches more relevant to our research. The first two mentioned papers, [8, 9] describe algorithms to reach one possible solution, not the optimal one. The third paper [10] 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.

Differently 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 the successive amount of utility agents has to concede compared to the (local) optimal bid. Formula (8) represents the decrement of utility for agent i, corresponding to the particular state s^k, where s^* is the agent's optimal state (proposal).

$$du_i^k = Ev\left(s^*\right) - Ev\left(s^k\right) \tag{8}$$

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, what enables keeping important information private.

This process ends when all agents cannot select a next state better than one already proposed in the past. 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 proposed distributed dependencies satisfaction algorithm can be described as follows:

1. Each agent *i* select its next preferable alternative state, from those not yet proposed before. Let us suppose this is state *a.*

   $$du_i^a = \min_s \left(du_i^s \right)$$

2. Each agent *i* sends out to others:
   - its own preferable state as a new proposal
   - its own local decrement of the utility for that state
3. When agent *j* receives the proposal (state a) from agent *i*, it calculates:
   - its own local decrement of utility (du_j^a)
   - the joint decrement of the utility:

     $$jdu^a = \sum_{dag} du_{dag}^a \quad , dag = \{1 \ldots n\} \; set\ of\ mutual\ dependent\ agents$$

   - the minimum joint of the decrement of the utility already known (jdu^m)
4. and selects:
   - its next preferable state. Suppose it is state *b*.
   - if $du_j^b < jdu^m$, agent *j* proposes state b to other agents
   - else agent *j* accepts state *m* as the final proposal and negotiation ends.

Transfer of Compensations

After agree 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 (9):

$$rdu = \frac{jdu^m}{n}, \quad n = number of\ agents \qquad \textbf{(9)}$$

As a consequence, some agents have to pay or get a compensation value to others, according to formula (10).

$$cValue_i = rdu - du_i^m \qquad \textbf{(10)}$$

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

4 Knowledge Representation

An Electronic Institution provides the means for helping in VO formation stage, not only providing an appropriate negotiation algorithm, but also supporting other generic services, including a knowledge representation scheme.

In order to enable different organisations to understand each other, ontology [11] is needed. It is mandatory that agents agree on a world model, in order to reach mutual understanding of the information they exchange, that is they have to agree on a common ontology. Ontology is a description of domain objects, concepts and relationships that is designed to enable knowledge sharing among heterogeneous agents that coexist in the same community. In the VO context, different organisations are, in principle, heterogeneous and not known from each other in advance. Therefore they will need to share the semantics of the context they are interacting about. Adopting a common ontology guarantees information consistency and compatibility for a community of agents. The information consistency is satisfied when each specific expression has the same meaning for any agent in the Electronic Institution. The information compatibility is verified when any concept is described by the same expression, for all the agents.

The ontology should be defined in such a way that encompasses all VO life cycle scenarios. However, for now, we are only applying ontology concepts needed for the VO formation stage. Items are described by means of a set of components, and components are described by a set of attributes, which are instantiated with values.

The ontology also includes multi-attribute definition as well as attributes' intra and inter dependencies. In that context, ontology is then the knowledge representation scheme suitable for properly identify items, components, attributes and values, together with relations that map items to components, components to attributes and attributes to values. Such ontology can be represented by the following structure:

$$Ont = \langle Item, Comp, Att, Val, IC_r, CA_r, AV_r, Deps \rangle \text{, where:}$$

- *Item* is the set of items' identifiers.
- *Cmpt* is the set of components' identifiers.
- *Att* is the set of attributes' identifiers.
- *Val* is the set of attribute values' identifiers. Each one of these values is represented by the tuple:

  $Val_i = \langle Type, Domain \rangle$, where:

  $$Type = \{integer, real, string\} \quad \text{and} \quad Domain = \{continuous, discrete\}$$
- $IC_r : Item_i \rightarrow \{Comp\}, \quad \forall_{Item_i} \in Item$

  A relationship that assigns to each item in *Item*, a set of components in *Comp*.
- $CA_r : Comp_i \rightarrow \{Att\}, \quad \forall_{Comp_i} \in Comp$

  A relationship that assigns to each component in *Comp*, a set of attributes in *Att*.
- $AV_r : Att_i \rightarrow Val_k, \quad \forall Att_i \in Att, \exists^1 Val_k \in Val$

  A relationship that assigns to each attribute in *Att*, a specific value in *Val*.
- $Deps : \{Dep_{ij}\}, \quad Dep_{ij} = f(Val_{ki}, Val_{mj}), \quad \forall Att_i, Att_j \in Att$

  A set of relationships defining the dependencies between attributes' values.

The Electronic Institution will provide this knowledge representation scheme in order to facilitate both definition of needs and interactive message contents.

5 Phased Commitments

We have seen how agents, representing individual autonomous enterprises, may reach an agreement through appropriate negotiation procedures. The contract that formalises that agreement should explicitly state all the commitments that those agents are due to satisfy all along the VO life cycle.

A full commitment contract may be unable to deal with possible future, partially unexpected, events. This fact has already been recognized since the definition of the old contract net protocol [12] where the possibility of a contract cancellation was envisaged. More recently, other authors like [13] have approached this subject in the context of decommiting in the meeting scheduling application. However, it was Sandholm [14, 15] who gave a more systematic and relevant contribution for this issue through the introduction of the concept of "leveled commitment" and associated penalties. Contrary to the game theoretic approach where contingency contracts are established according to the existence or not of future events, Sandholm [15] allows unilateral decommitments through the payment of calculated penalties. Resulting contracts are then called "leveled commitment contracts".

Three main aspects are related to this issue:

- First, the problem of how to represent, in an unambiguous form, such a commitment including all relevant information about future agents' attitudes.
- Second, how to explore this knowledge in order to correctly monitoring the next stages of the VO life cycle?
- Third, what to do in case of failure of what was previously negotiated and agreed and finally stated in the accepted contract. Can parties back out? In what circumstances can that happen? And, if the answer is yes, what procedure should follow? Should a new negotiation process start or (and) should appropriate penalties be enforced on the agents?

We envisage representing the negotiation outcome as a kind of frame where each slot represents pre-conditions plus a set of rules to be selected for possible application. Those rules whose conditional part has been verified indicate the appropriate action to be taken in those specific circumstances.

A contract, including a set of phased commitments, can be represented as:

$$Contract = \langle\ LAgts,\ Agts - P/S, Verification\_procedure \rangle$$, where:

- *LAgts* represents the list of agents that accept that contract.
- *Agt-P/S* ties each agent together with the contribution (product or service) is commited to give to the VO.

$$Agt - P/S = \{Agt - P/S_i\}, \quad Agt - P/S_i = \{\langle Agt_i, Prod/Serv_i \rangle\}$$

- *Verification_procedure* indicates how and when to monitoring the operation procedures agreed through the contract. *Verification_procedure* is represented as:

Verification_procedure=<Pre-cond, Rule-set>, where:

- *Pre-cond* ∈ {*event*, *time_point*}
  - *event* is a specific type of arrived messages.
  - *time _point* is a pre-specified point in time for checking current conditions.
- *Rule-set* = $\{<cond_i, action_i>\}$
  - $cond_i$ is a set of conditions to be checked after *Pre-cond* is true.
  - $action_i \in$ *{dpenalty, dpenalty + re_negt, noact}*, and:
    - *dpenalty* represents a decommitment penalty value.
    - *re_negt* represents the re-negotiation action.
    - *noact* represents the case where no action is to be done.

Decommitment penalties have to be calculated according to the other VO partners' respective losses. It is our intention to enhance the negotiation protocol in such a way that the agents already know these penalties at the end of the negotiation phase.

6 Conclusions and Discussion

Electronic Institutions are general frameworks for helping in collaborative work in electronic environments. Electronic Institutions provide, sometimes enforce, rules and norms of behaviour and make available service facilities supporting both interaction and operation monitoring of computational entities.

A Virtual Organization is a powerful example of the need of such collaborative work, once different enterprises have to join together, temporarily, to achieve a common business oriented goal. This paper elaborates on how Electronic Institutions can effectively help during the VO life cycle.

For the VO formation stage we have introduced a new negotiation algorithm, called Q-Negotiation, which includes appropriate features for dealing with the specific requirements of the VO scenario. An important requirement in the VO scenario, is that information must be kept private to individual enterprises, since they are competitive by nature and do not want to reveal their market strategy to others. The Q-Negotiation algorithm has the ability to maintain information private to individual enterprises, and at the same time, includes the capability to evaluate multi-attribute proposals, to learn during the negotiation process, and to resolve attributes' inter dependencies. Let us discuss each one of these features separately. First, multi-attribute evaluation is done assigning relative preferences to attributes. Other studies in multi-attribute evaluation [1, 2] generally impose the use of a real concrete value that captures each attribute's importance, which sometimes can be difficult to quantify. Second, learning is performed by using an on-line reinforcement learning algorithm during the negotiation process, through a qualitative feedback that is the opponent's comment to each proposal. Third, the inter attributes' dependencies resolution process proposed in Q-Negotiation reach the optimal solution keeping information private as much as possible. Other known approaches related with distributed dependencies resolution, either reach non-optimal solutions [8, 9], or impose the knowledge of other agents' private information to be made public [10].

For the VO operation stage we here propose the exploration of a phased commitment that is established in the end of the negotiation process during VO life cycle previous stage. This commitment is then specified in a contract. Through phased commitments, the Electronic Institution has the capability to monitor the behaviour of the participant entities at pre-specified moments previewed in the contract according to time points or future events.

References

1. Vulkan, N., Jennings, N.R.: Efficient Mechanisms for the Supply of Services in Multi-Agent Environments. 1st Int Conf. on Information and Computation Economics, Charleston (1998)
2. Matos, S., Sierra, C.: Evolutionary Computing and Negotiation Agents. Workshop on Agent Mediated Electronic Trading, Minneapolis, USA (1998)
3. Faisst, W.: Information Technology as an Enabler of Virtual Enterprises: A Life-cycle-oriented Description. European Conference on Virtual Enterprises and Network Solutions – New Perspectives on Management, Communication and Information Technology, Paderborn, Germany (1997)
4. Fischer, K., Muller, J.P., Heimig, I., Scheer, A.: Intelligent Agents in Virtual Enterprises. 1st International Conference on the Pratical Application of Intelligent Agents and Multi-Agent Technology, London, UK (1996)
5. Watkins, C.J.C.H., Dayan, P.: Q-Learning. Machine Learning, Vol. 8, N. 3/4. Kluwer Academic Publishers, Boston (1992) 279-292
6. Oliveira, E., Rocha, A.P.: Agents advanced features for negotiation in Electronic Commerce and Virtual Organisations formation process. In: Dignum, F., Sierra, C. (eds.): Agent Mediated Electronic Commerce, the European AgentLink Perspective. Lectures Notes in Artificial Intelligence, Vol. 1991. Springer-Verlag (2000) 77-96
7. Yokoo, M., Durfee, E., Ishida, T., Kuwabara, K.: Distributed Constraint Satisfaction for Formalizing Distributed Problem Solving. 12th International Conference on Distributed Computing Systems, Yokohama, Japan (1992)
8. Armstrong, A., Durfee, E.: Dynamic Prioritization of Complex Agents in Distributed Constraint Satisfaction Problems. 15th International Joint Conference on Artificial Intelligence, Nagoya, Japan (1997)
9. Yokoo, M., Hirayama, K.: Distributed Constraint Satisfaction Algorithm for Complex Local Problems. 3rd International Conference on Multi-Agent Systems, Paris, France (1998)
10. Parunak, H.V.D., Ward, A., Sauter, J.: The MarCon Algorithm: A Systematic Market Approach to Distributed Constraint Problems. Artificial Intelligence for Engineering Design, Analysis and Manufacturing J., Vol. 13. Cambridge Univ. Press (1999) 217-234
11. Gruber, T.R.: A Translation Approach to Portable Ontology Specifications. Knowledge Acquisition, Vol.5, N.2, (1993) 199-220
12. Smith, R.G.: The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver. IEEE Trans. on Computers, Vol.29, N.12 (1980) 1104-1113
13. Sen, S., Durfee, E.: The role of commitment in cooperative negotiation. International Journal on Intelligent Cooperative Information Systems, Vol.3, N.1 (1994) 67-81
14. Sandholm, T.W., Lesser, V.R.: Advantages of a Leveled Commitment Contracting Protocol. In Proceedings of the Nat. Conf. on Artificial Intelligence (AAAI), Portland (1996) 126-133
15. Sandholm, T.: Agents in Electronic Commerce: Component Technologies for Automated Negotiation and Coalition Formation. Autonomous Agents and Multi-Agent Systems, Vol,3, N.1. Kluwer Academic Publishers (2000) 73-96