

Fencing the Open Fields: Empirical Concerns on Electronic Institutions (Invited Paper)

Pablo Noriega

Institut d'Investigació en Intel·ligència Artificial
IIIA-CSIC. Bellaterra, Barcelona, Spain

Abstract. The regulation of multiagent systems may be approached from different stand-points. In this paper I will take the perspective of using a certain type of devices, electronic institutions, to regulate agent interactions. Furthermore, in this paper I am concerned with the tasks of design and construction of actual electronic institutions and I will explore some of the empirical aspects that one may encounter in such activities. More specifically, I will focus on those empirical aspects that are characteristic of electronic institutions rather than those that may be typical of multi-agent systems development in general or other types of software engineering. I use three examples of actual electronic institutions that show different and complementary features in order to motivate a number of distinctions that may be used to treat empirical features in a systematic way.

1 Introduction

Social interaction in everyday life is structured in many ways. When I buy fruit in the local market I exchange information with the fruit lady to find out what is today's price of the best fruit available and, if I like the options she offers me, I pay her and I get my apples. That simple interaction entailed conventions for simple bargaining and payment which in turn involved a common understanding of fruit features, money and the delivery of goods. Other interactions may involve conventions that are far more complex than the ones my fruit lady and I need to share. Traveling by bus, getting medical attention or passing a law through Parliament would be impossible if we didn't share with those with whom we interact some conventions that guide our individual behavior and facilitate us to achieve our intended goals; Society has developed organizations, contracts, standard procedures, markets, laws and many other such devices to *regulate* human interactions and make them effective.

In multiagent systems interactions may need to be structured as well. Generally the structuring is part of the design of the multiagent system because the participating agents are built by the same designer that builds the environment where the agents interact or because the rational components of agents are in some way accessible to the system. However, when the multiagent system is open to the participation of unknown agents or when the autonomy of participating

agents may have undesirable social consequences, the structuring of agent interactions may need to resort to regulation devices analogous to those that are practical in human affairs.

MAS literature has addressed the problem of structuring agent interactions from different approaches: coordination, agent communication languages, communication and interaction protocols, teams and coalitions, negotiation, institutions, organizations and norms. Sometimes these efforts have taken an agent-centric perspective in which the prevalent issues are how the agent receives, adopts or contravenes the conventions, while other works have taken a social perspective where the objective is the design of conventions that provoke the intended aggregate behavior of agents and the prevalent issues are those that affect the system components that are shared or used by participants, like the expressiveness of languages, effectiveness of interaction protocols, enforcement of commitments. The motivation of much of the MAS community's work along these lines has been theoretical, inspired by Logic, Game-Theory, (Economics) Mechanism Design or Sociology and Social Psychology. Nevertheless there has also been considerable work derived from applications of multiagent systems to domains like supply networks, auctions, virtual organizations or conflict resolution that address different forms of structuring or regulating agent interactions.

In this paper I will address the subject of regulating interactions in multiagent systems. I will look into the problem from a social perspective and will focus my discussion around one particular approach: electronic institutions. I will concern myself only with the empirical aspects involved in the development and use of electronic institutions, motivated by the real-world application of these devices. Although in this paper I will merely test the waters, my purpose is to throw some light on a significant but elusive subject that, I believe, deserves a systematic treatment.

In the next section I outline the generic notion of an electronic institution and then make explicit some compromises adopted to make such generic view operational. Next, I discuss three examples of actual electronic institutions whose features will allow me to outline, in Sec. 4, some empirical aspects involved in the development of electronic institutions.

2 Institutional Intuitions

The easiest way to describe electronic institutions is as the computational counterpart of traditional institutions. Traditional institutions are conventions that a group of agents follows in order to accomplish some socially agreed upon objective. Although we take institutions to be distinct from the agents that interact within them, it is not unusual to abuse language and identify an institution — the set of conventions — with an entity — a firm, company, organization — which is the warrant of those conventions ([5], p.5).¹ We can picture an institution as a nicely fenced plot in an open field. Things in the open field may be confuse and unpredictable, but inside the institution agents are able to play on a safe level ground.

¹ In keeping with this abuse of language we refer to the implementation of an electronic institution as an electronic institution.

Traditional —and electronic institutions— are used to regulate interactions where participants establish commitments and to facilitate that these commitments are upheld. It makes sense to institute some conventions if the establishment of commitments between participants is a process that is repeated with the same or different participants, but always under those same conventions. In that way participants are liberated of devising a process for establishing the commitments and concentrate on the decision-making tasks. The institutional conventions are devised so that those commitments can be established and fulfilled in an effective fashion and therefore participants be willing to submit to those conventions.

Institutions, in general terms, are established to facilitate effective interactions, and in order to do so they are devised to deal with a few complementary concerns, the most salient are:

- Establish the institutional conventions. So that these conventions have an objective reference that participants may invoke to understand the conventions, follow them, be accountable for their satisfaction and contend the wrongdoing of other participants.
- Assure permanence and stability of the conventions. In order that participant may hold sufficient certainty of the requirements and outcomes of their interactions and that they may expect and choose to participate on different opportunities without undue adjustment of their participation requirements.
- Enforce satisfaction of institutional commitments. So that all participants may rely and be held responsible for their institutional actions as far as the institutional conventions state.
- Guarantee accountability of institutional interactions. Be able to allocate risk and blame in an objective and effective manner. In most institutions, participants may be liable when they establish a commitment and if these participants are unreliable or even malevolent, there is risk involved that the institution is intended to allocate properly and limit damage effectively.
- Manage access and identity of participants. Validate that they satisfy the requirements of capability, resources of entitlement as long as they act within the institution, in order to be held accountable for their institutional actions.

From this intuitive description, it is not difficult to conceive electronic institutions as devices that facilitate on-line interactions: *coordination artifacts* that constitute —in Herbert Simon’s engineering design image [8]— an interface between the internal rational decision-making capabilities of agents and the social effect of their interactions.

Two features of that description are readily apprehended and I will take for granted from now on: the fact that participants are willing and able to *interact* and that these participants may be *human or software agents*. However, two more aspects of that description need further discussion because they may be operationalized in different ways and give ground to electronic institutions of different flavors. The first subtle issue is what constitutes an interaction, the other is the way interactions are structured to achieve the shared or common social purpose.

Institutional Interactions and Constitutive Rules. What would a bare-bones interaction within an institution may be? I take it that if we think of humans interacting with software agents, or software agents interacting amongst themselves, in any significant way, the least one requires of any two participants that interact is that they exchange some sort of a *message* that is mutually *intelligible* to both of them.² Intelligible messages is all we need, provided intelligibility involves some communication conventions that entail syntactic compatibility and some ontology alignment so that, in particular, the message could satisfy presumable conditions and have foreseeable effects that are acknowledged by the speaker and the receiver of the message. To achieve intelligibility, electronic institutions ought to be virtual entities that *establish* –define and uphold– the shared communication conventions. In this sense the electronic institution will then be not only the set of communication conventions that regulate agent interactions, but also the warrant of the conventions that make messages intelligible.

In order to fulfill these functions electronic institutions need to institute a connection with the real world through some *constitutive rules*. Constitutive rules fix the socially shared meaning of messages by linking the utterance of illocutions within the institution with conditions and effects those messages have in the real world. Hence, institutional interactions are messages that comply with the conventions for interactions of the institution, but institutional interactions *count as* real-world interactions when participants are bound to their meaning and effects by the constitutive rules of the electronic institution where they participate.

Notice that in the previous paragraphs I have taken a strong dialogical stance by assuming that all institutional interactions are messages and only messages, however these messages do have a connection with the real world through the constitutive rules that make them count as true actions³. So, from now on, I will assume electronic institutions to contain a set of constitutive rules on one hand and, on the other, a set of interaction conventions that regulate institutional actions properly.

We may think of interaction conventions as a way of establishing the pragmatics of institutional illocutions, that is, what are the admissible messages, what their proper sequencing and, in general, what their pre and post-conditions are. We can also look at interaction conventions as sets of norms that institutional illocutions are bound to satisfy. In the first approach we may think of the institution as commitment-based interaction protocols, as structured dialogues or as some sort of workflow. In this case, the interaction conventions constrain the class of potential interactions to an acceptable subclass and determine how a

² Here I take “message” to be an ostensible manifestation (a string of characters, a coded signal, or movements) with an ostensible effect (a change of state in the systems or their environment).

³ Nothing extraordinary here. Recall, for example, that a trial involves merely an exchange of statements that are linked to the purported criminal action, evidence, and derived actions. A death sentence, for example, is based on an institutionally valid sequence of illocutions and the prisoner’s life ended by force of a constitutive rule

given interaction forces or prevents certain future interactions. From the normative perspective, an institution may be formally construed as some logical system whose components involve all the meta-normal resources involved in the issuance, adoption, compliance and enforcement of norms. In both cases we need to express a mixture of declarative and procedural requirements that may be expressed in different ways. The choice is made on pragmatic and formal factors like the ease of specification, the ease of communication and adoption of the conventions by participating agents, the completeness of the specification, its enforceability or the computational complexity of the formalism, or, of particular relevance for this paper, the implementation constraints.

3 Three Examples

The following three examples illustrate various empirical features that are pertinent for regulated MAS development.

3.1 Example 1: Compranet, a Public Procurement Institution

Compranet is the on-line public procurement system developed and managed by the Comptrollers Office –now part of the Ministry for Public Administration– of the Mexican Federal Government.⁴ All Federal Government ministries, agencies, departments and offices that are entitled to perform any contracting or acquisition on their own and the publicly owned companies –like the very large national oil (PEMEX), electricity (CFE, CLyFC)— are, by law, required to use this system for all of their purchases, service contracting as well as all building and construction of public facilities contracting.⁵

The more significant motivations for instituting Compranet were transparency of government transactions, making information on demand and contract settlements readily available to all potential and actual participants, enabling convenient access of SME and international suppliers to the large federal market, lowering transaction costs and motivating the adoption of IT technology by SMEs. The system was conceived and a prototype designed in 1995, it was gradually deployed, starting in 1996 until it reached its full functionalities around 2002.

Description of the Compranet Institution. The system runs in a centralized location managed by the Comptroller’s Office. Each time a purchasing agency starts a procurement process, a new procurement thread is opened in the central site. All processes follow essentially the same interaction protocol with minor variations in requirements and sometimes also in time. The protocol is comprehensive of the whole process. The CFP is posted in a public database and RFQs made available to paying participants who then may ask for clarifications of the RFQ conditions. These requests and the procurer’s responses are

⁴ <http://www.compranet.gob.mx/>

⁵ Ley de Adquisiciones, Arrendamientos y Servicios del Sector Público. Out of convenience, many State and Municipal governments also use this system although they are not bound to this law.

made available to all participating suppliers. Bids may then be submitted electronically. Bids are kept secret until the contract is awarded, then the winning bid (only) is posted in a public database. The awarded contract is registered in the system and landmarks are registered and audited by the Comptroller's Office. The process lasts from one to three months depending on the purchasing modality. All interactions are asynchronous but subject to deadlines enforced by the system. Appeals follow also a due process that may involve the Comptroller's Office and may even force a new enactment of the procurement process. Each buying party may activate multiple processes and suppliers may simultaneously participate in as many processes as they wish. In every step of the process, electronic documents are issued by the corresponding parties and copies of those documents are kept in the system for appeals and auditing purposes. Compranet's main functions are outlined in Figure 1. The diagram on the left shows the preparatory process of procurement, from issuance of CFPs to the purchasing of the RFQ by supplier agents. The diagram on the right shows the ensuing, electronic bidding and contract awarding phases of the process.

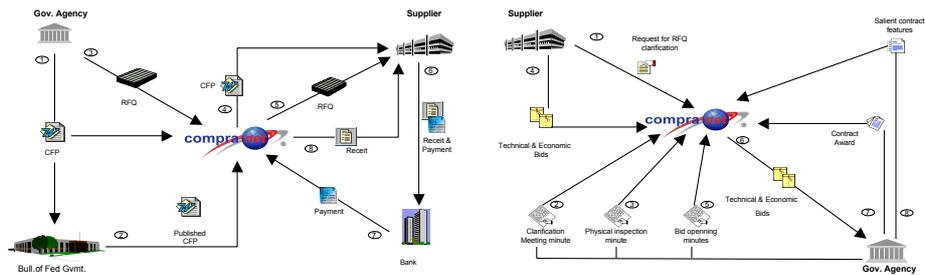


Fig. 1. The *Compranet* public procurement processes: (a) CFP and RFP cycle; (b) Electronic bid submission cycle

3.2 Example 2: MASFIT, On-line Fish Market Auctions

MASFIT (Multi Agent System for FISH Trading) is a MAS-enabled electronic marketplace that allows buyers to bid on-line in different fish auctions simultaneously. The main motivation for the system is to expand the daily market of fresh coastal fish catches by increasing the number of potential buyers and by aggregating the offer of several local fish markets (cf. [1]).

MASFIT was a joint proposal of, on one side, a firm that provides the back-office systems to local and electromechanical technology for handling fish (weighting, labeling) and controlling the bidding clock and the the electromechanical devices for face-to-face bidding (buttons on a desk or infrared remote-control devices), AUTECA, and, on the other side, a consortium formed by the Office for Livestock and Fisheries of the Catalonian regional government and the *lonjas* (fish markets) of three different ports. The MAS technology was originally developed through an EU Take-up Action (IST-2000-28221) designed to profit from the

IIIA's Fishmarket developments ([4, 7, 2]) and continues through a Spanish government grant involving AUTEK and the IIIA. There is a working system undergoing user adoption tests and commercial deployment is pending on AUTEK business model satisfaction.

Description of the MASFIT Institution. The MAS is subject to three design requirements:

1. That face-to-face bidding conditions of each local marketplace continue exactly as before, except for the possibility of the participation of remote buyers.
2. That on-line bidding follows exactly the same conventions —information flows, auctioneer, timing and interaction protocol— as face-to-face bidding.
3. That a remote buyer may participate simultaneously in all the fishmarkets that subscribe to MASFIT.

Because of condition 1, the auctioning protocols were already defined. All involve the same —dutch auction—conventions for bidding but have slight differences on admission and accounting procedures. The consortium creates a federation of markets but it still allows each market to have direct relations with its customers. Buyers will need to sign a contract and establish some guarantees to participate in any or all local markets. For human buyers in the actual *lonja* sites, the only difference from the current situation is that they may lose a round against a buyer that is not physically present in the auction house. Remote buyers bid through a remote device (a PDA or some other web client), or software agents acting on their behalf. Each remote buyer may have as many buyer agents as he or she wants and these may participate in one or many *lonjas* simultaneously. Each buyer agent is activated in a virtual *lonja* where a governor (owned and controlled by the virtual *lonja*) is attached to it. That governor controls all information flows between the agent and the specific *lonjas* where the owner of that buying agent wants to bid. Buyer agents have access in real time to all the information that is institutionally becoming available and to historical market information. Figure 2 depicts the trading architecture.

The MASFIT system includes two important additional services for buyers: a training environment and an agent-builder toolbox. In the training environment a user may test and tailor his or her buyer agents using data from past auctions or the information that is being generated in current auctions. The agent-builder tool box facilitates the assembly of an agent shell that is capable of following the *lonja* conventions with a decision-making model developed by the owner of an agent. AUTEK also provides complete agents with different parametric decision-making models that human buyers may instantiate.

3.3 Example 3: Framework for EI-Enabled Information Systems for Organizations

This example refers to a work in progress ([6]). The general idea is to have a framework to develop and deploy corporate information systems (CIS) whose

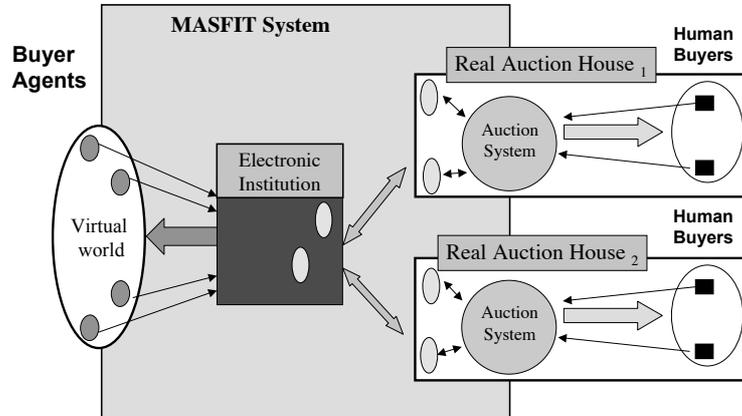


Fig. 2. MASFIT virtual auction federation. Traditional fish markets (right) allow human and software buyers to participate concurrently in simultaneous auctions.

operation is regulated by a prescriptive description of the way the organization is intended to function. The intuitions are rather simple: We take organizations to be groups of individuals that work together to achieve their shared goals the best way they can. Furthermore, we assume these individuals should follow some *institutional* conventions that make their interactions structured and predictable. Our framework, then should allow us to connect the institutional conventions that prescribe the procedures as well as the guidelines staff members should follow in their everyday activities, with the way those activities actually happen as reflected in the organization’s CIS.

The framework we are developing is outlined in Figure 3. Staff and client interactions are coordinated by an institutional convention which, in this example, is specified, enforced and enacted through the type of electronic institutions proposed by the IIIA (top layers of the diagram). Agent interactions, are mapped onto the CIS through a grounding language that establishes a correspondence between the linguistic interactions that take place in the institutional layer and actions that take place in the business domain (bottom layers).

The framework is designed for the development of actual corporate systems of significant complexity. These CIS involve the usual CIS components: data repositories, human users as well as business forms and procedures that are *agentified*.⁶ Institutional conventions are captured as interaction protocols that take care of procedural conventions, and as in-house software agents whose behavior — specified and implemented by the organization— is subject to the organizational guidelines, policies or norms.

Figure 4 is an illustration of a typical organization, a hotel, whose activities are organized as a network of interrelated “business contexts”. The illustration

⁶ Simplifying things, we have conventional CIS components handled through front-end devices —that we build— that are reified as *server* agents.

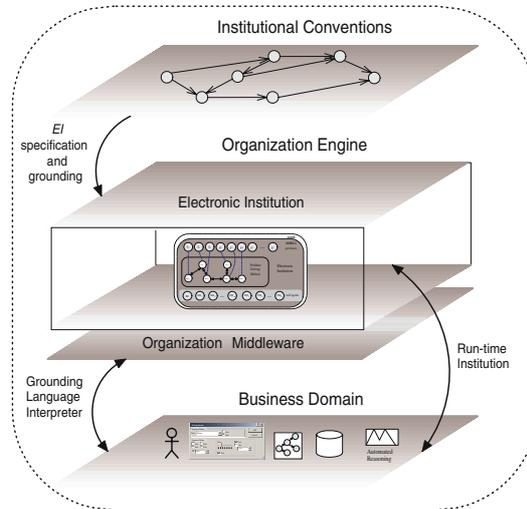


Fig. 3. EIO: An institutional description of an organization (top level) is implemented as an electronic institution that controls the operation of the organization’s conventional information system (bottom level)

indicates how each business context is implemented as a standard institution that involves client and organizational agents that act on behalf of those users of the CIS, plus server agents that translate institutional illocutions in terms of the CIS components (users, forms, databases, business rules, ...) and actions (database updates, PDA messages, procedure executions, ...).

4 A Timid Proposal

I propose to look into three “dimensions” that involve the design decisions I have found most significant in the development of electronic institutions that are intended for use. The rationale for choice is that they are closely linked to the concerns of institutions and to the type of conventions electronic institutions implement that I mentioned in Section 2. Moreover, I believe these dimensions apply also to regulated multiagent systems in general and are peculiar to them in the sense that they are not equally significant for other types of MAS applications or conventional IS development.

It is worth mentioning here that the design of electronic institutions, and regulated MAS, is in practice a matter of organizational design [3], and as such involves engineering and design technologies, methodological approaches and validation assessments that are part of that discipline. My remarks will take them for granted.

4.1 Grounding

This dimension is concerned with the relation between the actions that take place within the institution and the relation they have with the real world. How

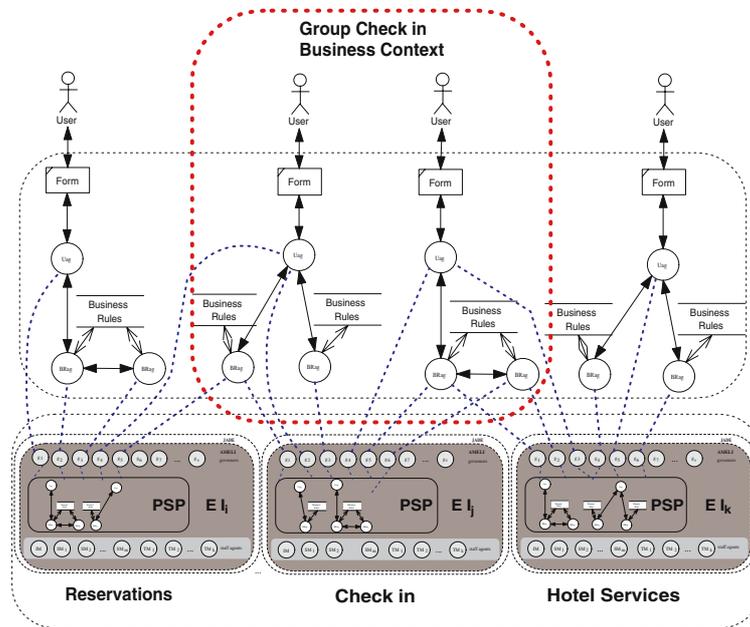


Fig. 4. EI-enabled corporate information system as a federation of electronic institutions each corresponding to the business units of the corporation

institutional actions represent, correspond and get to count as legitimate actions in the world, i.e. how to *establish* the constitutive conventions of an electronic institution

It is worth distinguishing two types of grounding, the first is legitimating the institution, i.e. making it legitimate or, more plainly, making it exist in the world. The other type of grounding is achieved by establishing a working correspondence between entities that are involved in institutional actions, within the institution, and the entities of the real world that should affect and be affected through institutional actions. The first involves, usually some constitutional act like a contract between participants, a public charter for the institution or a legal regulation that declares the achievement of a legitimate status—in the actual social world—of the institution and the commitments established therein. The second type involves the establishment of a sort of isomorphism between the language of the institution and the application domain where the institution applies.

Our three examples provide good illustration of different grounding mechanisms.

Compranet was created as an electronic institution to support the actual compliance with a law. It came to exist as an act of authority from the office in charge of enforcing and interpreting the law that regulates public procurement in the Federal Government. In fact, once Compranet became operational, the regulations that determined the procedures involved in public procurement were

rewritten to be a textual description of a functional specification of the Compranet sytem. Those regulations made precise the correspondence between institutional actions and the real world, by stating, for instance, the conditions that a company needed to fulfill in order to participate in any procurement contract, how RFP should be paid, what were the requirements for a valid “electronic bid” or how to appeal an award resolution.

In MASFIT constitution comes about through a contract that binds the technology supplier and the *lonjas*, on one side and another contract between the *lonjas* and the participants on the other. The first contract makes the virtual institution become the actual regulator and enforcer of the conventions for on-site trading as well as remote on-line trading. The other contracts make the trading regulations applicable to participants. The first contract establishes the ontological and procedural grounding that translate virtual exchanges into real exchanges. The second agreement guarantees that both parties will be liable.

Notice that the MASFIT contracts need to be concerned with very concrete matters like setting up guarantees to cover misbehaviors of different sorts in order to make the grounding work. For instance, the *lonja* establishes the obligation to pay the seller every item sold in an auction and to deal with the eventuality of a defaulting buyer, the *lonja* falls back on a credit line—or some escrow mechanism—that buyers need to establish when signing their corresponding contract with the *lonja*. Likewise the technology supplier commits to a certain level of service and some penalties in case of system malfunction, that may for instance be underwritten in an insurance policy.

The case of EIO is of a different nature. The company owns the institution and “owns” the staff that is supposed to work under the institutional conventions. Grounding in this case involves the obvious constitutional act of making the system operational but the isomorphism between institutional actions and activities in the world takes a very characteristic form. First note that grounding is in fact made not in the *physical* world properly but in the *virtual* world of the corporate information system on which the company operates. The links between the CIS and the real world are the ones we are familiar with and do not require further comment, but the link between the institutional conventions and the CIS involves an ontology alignment between illocutionary language terms and CIS entities and the instrumentalization of the institutional actions catalogue into functionalities of the CIS (e.g. Database diagrams correspond to constants and relations in the EI; functions—like making a payment—correspond to table look-ups, execution of business rules and updating database registers).

An important concern in the design of an electronic institution is—as with traditional ones—to determine the interplay of interests involved in the agents’ interactions and the proper allocation of responsibility and control that such interplay requires. It is essentially a matter of choosing appropriate checks and balances to make the interplay conducive to the stated objectives of the institution. It is also a matter of deciding what to make an institutional convention and how.

In EIO the company that owns the CIS that is regulated by the EI is the same company that defines, deploys and runs the EI. The choice of conventions and their enforcement mechanisms is made by the same company that will abide by them. The choice of conventions and their enforcement mechanisms is precisely a matter of institutional design, in classical terms, that the designer of the company undertakes in order to shape everyday activities to better serve the company needs and in the case of EIO the chosen checks and balances, departmental structure, lines of authority, decision guidelines and standard procedures *are* the conventions that define the EI. The EI is just making the normative specification becoming operational.

The case of Compranet is quite different. The EI facilitates the interactions of buyers and sellers and is in fact an independent third party that guarantees fair play and ideally contributes to make the market more effective. Consequently, fairness and effectiveness are the leading design features. In Compranet, for example, the choice of having a unique centralized procurement clearinghouse over a distributed procurement mechanism –a little Compranet in every government agency– was to better serve the objective of guaranteeing that the public procurement market was fair and reliable. In that light, the choice of a single trading room housed at the Comptroller’s Office signaled its ostensible unquestionable authority for that market. In addition, although centralization concentrated risk of technical and political failure, it greatly facilitated the adoption of the institutional practices and, in fact, their gradual deployment and uniform applicability. Both aspects make the third party strong *vis a vis* the sometimes conflictive objectives of buyers and sellers, and specially *vis a vis* a corruption-prone market culture.

MASFIT is an interesting case where the technology supplier creates a virtual institution that is a market-maker, an independent third party between buyers and sellers whose ownership is shared by the technology supplier, and a consortium of the regional government with the *lonjas* –which in turn are owned by the fishermen guilds and hence twice-removed from the sellers. However, the same technology supplier may enable buyers with *buyer agents* that perform aptly in the virtual institution thus blurring the border between the independent market-maker and the buyer. Notice then that the stakeholders in the fish trade have different relationships with the technology supplier who is on one hand in charge of enforcing the conventions, and on the other facilitating the participation of buyers through a technology that would be hard to develop by them but that unless the buyers have it they could hardly profit from the advantages of the virtual institution over the traditional *lonjas*.

4.2 Degree of Agentification

The fact that agents are present in one form or another in an electronic institution is again a matter of institutional design in the sense that agents are a component that is brought into the design in order to achieve certain functionalities or realize certain advantages, but that as other components the use of agents needs to be managed in order to achieve the intended benefits without

undesirable side effects. The characteristic features that agents may bring into the institutional design are persistence, automated rationality and ubiquity. They need to be assessed against their effect on reliability of the interactions, identity or entitlement of the participants, and their competency as suggested before, and upon this assessment decide where, to what extent and in what capacity agents are conducive to a better articulation of interactions in the application. By degree of agentification I want to refer to those choices, the type and level of functionalities that is delegated to software agents in the system and to the way such delegation is managed in the electronic institution. I hope the examples clarify what I mean.

I have used the notion of electronic institution as a rather generic coordination artifact without committing to many specific features. In particular I have not required, nor assumed, the need of software agents anywhere. Not as a constitutive component of the institution, not as participants. Thus I am confident to talk about Compranet as an electronic institution even if it didn't involve any agents whatsoever. In fact it doesn't assume the need of software agents in its design. It was designed to support them and some care was taken in order to guarantee that all institutional interactions could be performed by software agents but that is as far as the "agentification" of Compranet went. It was only *agent-compatible*. In Compranet no agent is involved in the operation of the institution, although agents may be —are being— used by suppliers to find interesting CFPs, to analyze competitors or market behavior, and by authorities to audit contracts, keep track of incidents and to gather evidence of punishable misbehavior.

The case of MASFIT is another extreme. The federation of *lonjas* is accessible only through software agents that represent the human buyers. In MASFIT human buyers may still bid in a physical *lonja* exactly as they used to, and they may get to participate in the different physical *lonjas* in the same circumvented and limited way they used to: having a partner present in another *lonja* to bid on their behalf and using a telephone to coordinate with that partner. Notice that, from the buyers perspective, the real benefits of the virtual market are realized to their full potential when software agents are making bidding decisions autonomously for two main reasons. First, because a human user may deploy software agents that can participate simultaneously in all the *lonja* according to whatever buying strategy the buyer delegates on them —notably strategies that involve real-time information from other *lonja* and coordination of the buyer agents; second because those software agents may profit from all the information that is available in the market whose volume and speed is excessive for human users and are able to exploit it in whatever automated reasoning technology their owner puts in their decision-making strategies. On the market-maker side agents are also essential as internal "staff" to govern buyer agents, manage bidder's interactions, access to the virtual auction floors and on-line clearing of bids and keeping track of the commitments incurred by the totality of agents belonging to each buyer. MASFIT is *agent-based* in a very strong sense: it is feasible only if software agents are involved.

While in Compranet agentification is dispensable, in MASFIT it is essential. The case of EIO is still a different type of MAS agentification, it is *agent-pervasive*. You start with a human organization and a traditional CIS that supports its operation, and you end-up with agents all over: agents that mediate interactions with external users, agents that encapsulate the decision-making functions of the organization's staff, agents that mediate the interactions human staff users of the CIS still need to perform and, finally, agents that manage the resources of the CIS, that is, server agents that interact with CIS components in order to get or pass information to a data base, activate a business rule or a standard procedure.

4.3 Autonomy

In the previous subsection I purposely left autonomy out of the characteristic features of agents that need to be assessed for the design of an electronic institutions. I left autonomy out because, I believe, it deserves a more systematic treatment along two aspects that are fundamental in electronic institution design: the openness of the electronic institution and the way institutional conventions are enforced.

Openness can be understood in two ways. First, as the extent of requirements imposed on participants to join an institutions, second as an indication of how structured or rigid are the conventions the institution upholds. Evidently, both have to do with the flexibility agents may or need to have in order to participate in a given institution and to a certain extent to the type of autonomy they are entitled or forced to exercise.

MASFIT is a good example of a rather generous openness of access and highly inflexible interaction conventions. In MASFIT any potential buyer is admitted as long as he or she provides good enough guarantees to cover his or her purchases and accept other conventional contractual obligations. Once the human buyer is accepted he or she may deploy in the federated market any (external) software agent of whatever structure or composition he or she wishes. No requirements are imposed on the agent's capabilities, truthfulness, livelihood, benevolence and none is validated in the institution, thus in the first sense of the term, that electronic institution is completely open. However, the external buyer agent has no choice over the way it would go about buying fish, it has to abide strictly by the MASFIT rules. The institutional interaction conventions in MASFIT are explicit and comprehensive and the buyer agent is only allowed to utter admissible utterances at admissible moments. The contents of the buyer agent utterances are up to the agent's internals and the institution has no business in determining how or why a given utterance has such or such content. MASFIT either admits it or refuses it, but if MASFIT admits it the buyer agent is held to the commitments entailed by that utterance in a strict unavoidable manner. The autonomy of the buyer agent is limited only to its choice of parameters for the admissible illocutions it decides to make. And that is quite enough for an auction market.

In the Compranet case the situation is similar but less definite. Access to any particular procurement process is restricted to suppliers who have first documented their personality, entitlements, pertinence and competitiveness in the thorough manner dictated by the Law of Acquisitions and its associated regulations, once that process is completed the admissible suppliers need to buy their right to submit a bid by paying for the RFQ of the specific procurement event. Once these —grounding conditions— are met, the supplier may participate by following the legally established protocol. The protocol conventions are explicit but are open to adjustment and interpretation to a certain extent. Institutional actions are electronic documents that register bids, clarifications, protocol adjustments (time-schedules, updated conditions), award resolutions, signed contracts, certifications of termination and acceptance, etc. All may be performed electronically and the procedure doesn't impose any condition on the agents' rationality in performing those actions, only in their preconditions and effects. Compranet as a third party in the process keeps a register of all institutional actions and as I will comment below, sees to it that commitments are satisfied. Hence, the institution is rather open to access and the statement and satisfaction of interaction conventions is not as explicit and inflexible as MASFIT.

The EIO model leans towards the opposite balance. It is mostly closed in the sense that many interactions involve agents that are owned by the organization, and it is somewhat open in the sense that the organizational staff does have contact with external agents whose motivation and worthiness are opaque to the organization. For the organization, those external agents are black boxes that are dealt with like buyer agents are dealt with in MASFIT: each external agent is governed by an internal agent that controls all information flows in and out of that agent and imposes on it the interaction conventions of the organization in a strict manner. However, having control over internal agents allows the designer to use autonomy in a rather fruitful way. The point of having a prescriptive description of the organization in the top-level of the EIO model is to govern interactions of the members or the organization in such a way that all participants are aware of what is expected of them and do what they are expected to do in *foreseeable situations*. In practice, this means that stable institutional conventions govern explicitly procedures and also some of the decision-making processes some staff agents are endowed with. Evidently there is a problem of granularity, the detail with which procedural conventions need to be expressed is related to the complexity of the process, its variability and the amount of local decision-making needed to make the conventions work. But notice that having control of the deliberative components of staff agents allows the designer to rely on their autonomy by specifying simpler standard procedures and program, in some autonomous staff agents, the decision-making capabilities to contend with non-standard situations *institutionally*, that is, according to the prescriptive definition of the organization. These individual decision-making capabilities can in fact be considerable since nothing prevents the designer from making the statement of the convention for a case —standard or not— in a way that prescribes —declaratively— the intended final outcome and leaves the —procedural— implementation up to the intervening

staff agents that have been endowed with sufficient knowledge and dexterity to carry out the task in question.

Enforcement of Conventions. While the EIOs model can take advantage of the autonomy of its staff agents, inside the institution autonomy is nil for all external agents. Likewise in MASFIT. In both cases, that is a significant design decision. In both cases, the explicit definition of the interaction conventions and their strict enforcement is necessary in order to assure reliability, fairness, trust and accountability *vis a vis* their clients. Although in Compranet trust, reliability, fairness, accountability and transparency are also relevant design features, the enforcement of conventions is more discretionary. The reason is the way constitutive conventions are established. MASFIT and the EIO are autarchic: Their constitutive conventions legitimize a social space where they define the rules, they control them and they let agents participate in that space if and only if those agents are willing to abide by the rules and, notably, MASFIT and the EIO are capable of strictly enforcing those rules. Strict enforcement is possible because the rules are explicit, and because of the way interaction conventions are implemented in that private space, their observance is mandatory and infractions are impossible. Through the constitutive conventions, MASFIT and the EIO have the authority, capability and power to enforce interaction conventions universally and strictly.

By contrast, Compranet is legitimized through an implementation that institutes only very basic procedural interactions, leaving participants enough leeway in their compliance of the conventions that are declaratively instituted by law and practice. Compranet as an electronic institution leaves considerable autonomy to participants at many stages of the process and its interaction conventions are not explicit enough to prevent discretionary interpretation. Compranet solves this problem through two classical mechanisms self-enforcement and authority. Self-enforcement is achieved by the fact that pertinent information is made public (transparency), that supplier and buyer interests are opposed, by having checks and balances among the roles that different buyer staff agents perform and by establishing significant penalties for misbehavior. Authority is centralized and final, has access to every institutional commitment, resolves any interpretation disagreements and has power to impose sanctions swiftly.

5 Closing Remarks

The three examples I discussed are representative of a variety of applications that regulated MAS may have. In particular, the highly structured model behind MASFIT is adequate for applications that are heavy in individual agent's decision-making but light in interaction, however high the liability of participants may be. Applications of this sort are typical in electronic commerce, customer complaints management and TRAMITES, the model is also convenient for classical mechanism design. Compranet is a token of due processes, whose purpose is to channel a complex agreement or coordination process into a manageable

—potentially intricate— sequence of standardized interaction stages that facilitate fair and objective conditions for the parties involved. Common examples of these processes arise for instance in conflict resolution, judiciary and legislative practice or in the execution of publicly sponsored programs. The EIO model is being developed to apply to large corporate information systems and should be applicable to recursively decomposable complex systems (i.e. that are decomposable into a few similar subsystems of less complexity, and these decomposable in turn). Hence, hospitals, retail chains and franchises, emergency response organizations are natural applications of it.

The comments I made around the three examples show how tentative and unsystematic my understanding of the subject still is. Nevertheless I hope that they serve to facilitate the establishment of a rudimentary set of distinctions and considerations that may be beneficial for the development of applied regulated agent systems.

Although I made my remarks around the notion of electronic institution I believe that most of what I said applies to other ways of regulating MAS because the main concerns of an institution that I postulated are also in the domain of most regulation of agent interactions. Notwithstanding this intended generality, I did persist in holding to a single major bias along my discussion: I have assumed that interactions among agents are repeated and structured around rather stable processes. It is a justified bias in the sense that (as stated above) it is appropriate for a large number of applications and also in the sense that it has proven feasible to implement a set of tools that have been adequate for a variety of applications. Nevertheless, the assumption imposes regulatory features that are unnecessary in some cases and unacceptable in others. I trust it can be weakened by focusing on the essential features required for peer to peer interaction I mentioned in Section 2.

When addressing the more general problem of regulating interactions in open systems we may profit by holding another metaphor that frees us from a building a neat enclosure and allows us to regulate directly the essential communication and commitment making assumptions that permit agents to interact “in the wild”. So far we have learned how to “fence the open fields” we may now consider “hanging the bell on the cat”.

Acknowledgements

The research reported in this paper is partially supported by the Spanish MEC project Web-i (2) (TIC-2003-08763-C02-01).

References

1. Guifré Cuní, Marc Esteva, Pere Garcia, Eloi Puertas, Carles Sierra, and Teresa Solchaga. Masfit: Multi-agent systems for fish trading. In *16th European Conference on Artificial Intelligence (ECAI 2004)*, Valencia, Spain, August 2004.
2. Marc Esteva. *Electronic Institutions: from specification to development*. Number 19 in IIIA Monograph Series. PhD Thesis, 2003.

3. James G. March and Herbert A. Simon. *Organizations*. John Wiley and sons, New York, USA., 1958.
4. Pablo Noriega. *Agent-Mediated Auctions: The Fishmarket Metaphor*. Number 8 in IIIA Monograph Series. PhD Thesis, 1998.
5. Douglass C. North. *Institutions, Institutional change and economic performance*. Cambridge University press, 40 west 20th Street, New York, NY 10011-4211, USA, 1990.
6. Armando Robles and Pablo Noriega. A framework for building EI-enabled intelligent organizations using MAS technology. In *Proceedings of the Third European Conference on Multiagent Systems. EUMAS 2005*, page 344354, 2005.
7. Juan Antonio Rodriguez-Aguilar. *On the Design and Construction of Agent-mediated Electronic Institutions*. Number 14 in IIIA Monograph Series. PhD Thesis, 2003.
8. Herbert A. Simon. *The Sciences of the Artificial*. MIT Press, Cambridge, MA, 1969.