Norm Adaptation of Autonomic Electronic Institutions with Multiple Goals

Eva Bou[1] Maite López-Sánchez[2] J. A. Rodríguez-Aguilar[1]

[1] IIIA - CSIC Artificial Intelligence Research Institute, Campus UAB
08193 Bellaterra, Spain
Email: {ebm, jar}@iiia.csic.es
http://www.iiia.csic.es

[2] WAI, Volume Visualization and Artificial Intelligence, MAiA Dept.,
Universitat de Barcelona, Gran Via de les Corts Catalanes 585, 08007 Barcelona, Spain
Email: maite@maia.ub.es
http://www.maia.ub.es/ ∼maite

Abstract: Electronic institutions (EIs) have been proposed as a means of regulating open agent societies. EIs define the rules of the game in agent societies by fixing what agents are permitted and forbidden to do and under what circumstances. And yet, there is the need for EIs to adapt their regulations to comply with their goals despite coping with varying populations of self-interested agents. This paper focuses on the extension of EIs with autonomic capabilities to alllow them to yield a dynamical answer to changing circumstances through the adaptation of their norms and their performative structure.

Keywords: autonomic electronic institutions, multiagent systems, adaptation.

1. Introduction

The growing complexity of advanced information systems in the recent years, characterized by being distributed, open and dynamical, has given rise to interest in the development of systems capable of self-management. Such systems are known as self-* systems [21], where the * sign indicates a variety of properties: self-organization, self-configuration, self-diagnosis, self-repair, etc. A particular approximation to the construction of self-* systems is represented by the vision of autonomic computing [18], which constitutes an approximation to computing systems with a minimal human interference. Some of the many characteristics of autonomic systems are: it must configure and reconfigure itself automatically under changing (and unpredictable) conditions; it must aim at optimizing its inner workings, monitoring its components and adjusting its processings in order to achieve its goals; it must be able to diagnose the causes of its eventual malfunctions and reparate itself; it must act in accordance to and operate into a heterogeneous and open environment.

In what follows it is argued that are EIs [7] a particular type of self-* system. When looking at computer-mediated interactions Electronic Institutions (EI) can be regarded as regulated virtual environments wherein the relevant interactions among participating agents take place. EIs have proved to be valuable to develop open agent systems [17]. However, the challenges of building open systems are still considerable, not only because of the inherent complexity involved in having adequate interoperation of heterogeneous agents, but also because the need for adapting regulations to comply with institutional goals despite varying agents' behaviors. Particularly, when these are self-interested.

The main goal of this work consists in studying how to endow an EI with autonomic capabilities that alllow it to yield a dynamical answer to changing circumstances through the adaptation of its regulations. Among all the characteristics that define an autonomic system the paper will focus on the study of self-configuration as pointed out in [18] as a second characteristic: "An autonomic computing system must configure and reconfigure itself under varying (and in the future, even unpredictable) conditions. System configuration or "setup" must occur automatically, as well as dynamic adjustments to that configuration to best handle changing environments".

The paper is organized as follows. Section 2 introduces the notion of autonomic electronic institution as an extension of the classic notion of electronic institution along with a general model for adaptation based on transition functions. Section 3 details how these functions are automatically learned. Section 4 details a case study to be employed as a scenario wherein to test the model presented in section 2. Section 5 provides some preliminary, empirical results. Finally, section 6 summarizes some conclusions and outlines paths to future research.

2. Autonomic Electronic Institutions

The idea behind EIs [24] is to mirror the role traditional institutions play in the establishment of "the rules of the game"–a set of conventions that articulate agents' interactions– but in our case applied to agents (be them human or software) that interact through messages whose (socially relevant) effects are known to interacting parties. The essential roles EIs play are both descriptive and prescriptive: the institution makes the conventions explicit to participants, and it warrants their compliance. EIs involve a conceptual framework to describe agent interactions as well as an engineering framework [1] to specify and deploy actual interaction environments.

In general, an EI regulates multiple, distinct, concurrent, interrelated, dialogic activities, each one involving different groups of agents playing different roles. For each activity, interactions between agents are articulated through agent group meetings, the so-called *scenes*, that follow well-defined interaction protocols whose participating agents may change over time (agents may enter or leave). More complex activities can be specified by establishing networks of scenes (activities), the so-called *performative structures*. These define how agents can legally move among different scenes (from activity to activity) depending on their role.

Although EIs can be regarded as the computational counterpart of human institutions for open agent systems, there are several aspects in which they are nowadays lacking. According to North [25] human institutions are not static; they may evolve over time by altering, eliminating or incorporating norms. In this way, institutions can adapt to societal changes. Nonetheless, neither the current notion of EI in [7] nor the engineering framework in [1] support norm adaptation so that an EI can self-configure. Thus, in what follows, it is studied how to extend the current notion of EI in [7] to support self-configuration.

First of all, notice that in order for norms to adapt, we believe that a "rational" view of EIs must be adopted (likewise the rational view of organizations in [8]) and thus consider that *EIs seek specific goals*. Hence, EIs continuously adapt themselves to fulfill their goals. Furthermore, we assume that an EI is *situated* in some environment that may be either totally or partially observable by the EI and its participating agents.

With this in mind, it can be observed that according to [7] an EI is solely composed of: a dialogic framework establishing the common language and ontology to be employed by participating agents; a performative structure defining its activities along with their relationships; and a set of norms defining the consequences of agents' actions. From this follows that further elements are required in order to incorporate the fundamental notions of goal, norm configuration, and performative structure configuration as captured by the following definition of *autonomic electronic institution*.

Definition 1 Given a finite set of agents A, we define an Autonomic Electronic Institution (AEI) as a tuple $\langle PS, N, DF, G, P_i, P_e, P_a, V, \delta\rangle$ where:

- PS stands for a performative structure;
- N stands for a finite set of norms;
- DF stands for a dialogic framework;
- G stands for a finite set of institutional goals;
- $P_i = \langle i_1, \ldots, i_s\rangle$ stands for the values of a finite set of institutional properties, where $i_j \in \mathbb{R},\ 1 \leq j \leq s$ contains the value of the j-th property;
- $P_e = \langle e_1, \ldots, e_r\rangle$ stands for the values of the environment properties, where each e_j is a vector, $e_j \in \mathbb{R}^{n_j}\ 1 \leq j \leq r$ contains the value of the j-th property;
- $P_a = \langle a_1, \ldots, a_n\rangle$ stands for the values that characterize the institutional state of the agents in A, where $a_j = \langle a_{j_1}, \ldots, a_{j_m}\rangle\ \ 1 \leq j \leq n$ stands for the institutional state of agent A_j;
- V stands for a finite set of reference values; and
- $\delta : N \times G \times V \rightarrow N$ stands for a normative transition function that maps a set of norms into a new set of norms given a set of goals and a set of values for the reference values; and
- $\gamma : PS \times G \times V \rightarrow PS$ stands for a performative structure transition function (henceforth referred to as PS transition function) that maps a performative structure into a new performative structure given a set of goals and a set of values for the reference values.

Notice that a major challenge in the design of an AEI is to learn a *normative transition function*, δ, along with a *PS transition function*, γ, that ensure the achievement of its institutional goals under changing conditions. Next, the new elements composing an AEI are dissected.

Agents participating in an AEI have their social interactions mediated by the institution according to its conventions. As a consequence of his interactions, only the *institutional (social) state* of an agent can change since an AEI has no access whatsoever to the inner state of any participating agent. Therefore, given a finite set of participating agents $A = \{A_1, \ldots, A_n\}$ where $n \in \mathbb{N}$, each agent $A_i \in A$ can be fully characterized by his institutional state, represented as a tuple of observable values $\langle a_{i_1}, \ldots, a_{i_m}\rangle$ where $a_{i_j} \in \mathbb{R},\ 1 \leq j \leq m$. Thus, the actions of an agent within an AEI may change his institutional state according to the institutional conventions.

The main objective of an AEI is to accomplish its goals. For this purpose, an AEI will adapt. We assume that the institution can observe the environment, the institutional state of the agents participating in the institution, and its own state to assess whether its goals are accomplished or not. Thus, from the observation of environment properties(P_e), institutional properties (P_i), and agents' institutional properties (P_a), an AEI obtains the reference values required to determine the fulfillment of goals. Formally, the reference values are defined as a vector $V = \langle v_1, \ldots, v_q\rangle$ where each v_j results from applying a function h_j upon the agents' properties, the environmental properties and/or the institutional properties; $v_j = h_j(P_a, P_e, P_i),\ \ 1 \leq j \leq q$.

Finally, regarding institutional goals, an example of institutional goal for the Traffic Regulation Authority could be to keep the number of accidents below a given threshold. In other words, to ensure that a reference value satisfies some constraint.

Formally, lets define the goals of an AEI as a finite set of constraints $G = \{c_1, ..., c_p\}$ where each c_i is defined as an expression $g_i(V) \lhd [m_i, M_i]$ where $m_i, M_i \in \mathbb{R}$, $\lhd$ stands for either $\in$ or $\notin$, and g_i is a function over the reference values. In this manner, each goal is a constraint upon the reference values where each pair m_i and M_i defines an interval associated to the constraint. Thus, the institution achieves its goals if all $g_i(V)$ values satisfy their corresponding constraints of being within (or not) their associated intervals.

2.1 Norm Transition

An AEI employs norms to constrain agents' behaviors and to assess the consequences of their actions within the scope of the institution. Although there is a plethora of formalizations of the notion of norm in the literature, this paper adheres to a simple definition of norms as effect propositions as defined in [12]:

Definition 2 An effect proposition is an expression of the form

$$A \textbf{ causes } F \textbf{ if } P_1, \ldots, P_n$$

where A is an action name, and each of $F, P_1, \ldots, P_n (n \geq 0)$ is a fluent expression. About this proposition we say that it describes the effect of A on F, and that $P_1, \ldots, P_n$ are its preconditions. If n = 0, 'if' disappears and the effect proposition simply becomes A causes F. From this definition, changing a norm amounts to changing either its pre-conditions, or its effect(s), or both. Norms can be parameterized, and therefore we propose that each norm $N_i \in N$, $i = 1, \ldots, n$, has a set of parameters $\langle p^N_{i,1}, \ldots, p^N_{i,m_i}\rangle \in \mathbb{R}^{m_i}$. Hence, changing the values of these parameters means changing the norm. In fact this parameters correspond to the variables in the *norm transition function* that will allow the institution to adapt under changing situations.

2.2 PS Transition

As mentioned above, an EI involves different groups of agents playing different roles within scenes in a performative structure. Each scene is composed of a coordination protocol along with the specification of the roles that can take part in the scene. Notice that we differentiate between institutional roles (played by staff agents acting as the employees of the institution) and external roles (played by external agents participating in the institution as users). Furthermore, it is possible to specify the number of agents than can play each role within a scene.

Given a performative structure, it is necessary to choose the values that are aimed at changing in order to adapt it. This involves the choice for a set of parameters whose values will be changed by the PS transition function. In our case, we choose as parameters the number of agents playing each role within each scene. This choice is motivated by our intention to determine the most convenient number of institutional agents to regulate a given population of external agents.

Scenes can be parameterized, and therefore, we propose that each scene in the performative structure, $S_i \in PS, i = 1, \ldots, t$, has a set of parameters $\langle p^R_{i,1}, ..., p^R_{i,q_i} \rangle \in \mathbb{N}^{q_i}$ where $p^R_{i,j}$ stands for the number of agents playing role r_j in scene S_i.

3. Learning Model

Adapting EIs amounts to changing the values of their parameters. This paper proposes to learn the *norm transition function* (δ) and the *PS transition function* (γ) by exploring the space of parameter values in search for the ones that best accomplish goals for a given population of agents. In this manner, if it is possible to automatically adapt an EI to the global behavior of an agent population, then, it becomes also possible to repeat it for a number of different agent populations and thus characterize both δ and γ.

Fig. 1 describes how this learning process is performed for a given population of agents (A) using an evolutionary approach. We have an initial set of individuals $\langle I_1, \ldots, I_k \rangle$, where each individual represents the set of norm and role parameters defined above $\{\langle p^N_{1,1}, \ldots, p^N_{1,m_1} \rangle, \ldots, \langle p^N_{n,1}, \ldots, p^N_{n,m_n} \rangle, \langle p^R_{1,1}, \ldots, p^R_{1,q_1} \rangle, \ldots, \langle p^R_{t,1}, \ldots, p^R_{t,q_t} \rangle\}$. Each individual represents a specific AEI configuration, and therefore, the institution uses each configuration to perform a simulation with the population of agents A. The corresponding configuration can then be evaluated according to a fitness function that measures the satisfaction degree of institutional goals (*configuration evaluation*). Finally, the AEI compiles the evaluations of all individuals in order to breed a new generation from the best ones *configuration adaptation*. This process results with a new set of individuals (*New configurations*) to be used as next generation in the learning process.

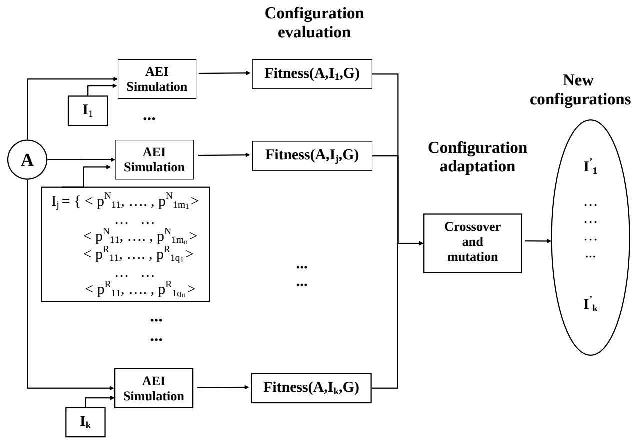

Fig. 1. Example of a step in EI adaptation using an evolutionary approach.

4. Case Study: Traffic Control

Traffic control is a well-known problem that has been approached from different perspectives, which range from macro simulation for road net design [27] to traffic flow improvement by means of multi-agent systems [22]. This paper tackles this problem from the Electronic Institutions point of view, and therefore, this section is devoted to specify how traffic control can be mapped into Autonomic Electronic Institutions.

In this manner, the Traffic Regulation Authority is considered to be an Autonomic Electronic Institution, and cars moving along the road network are regarded as agents interacting inside a traffic scene. Considering this set-up, traffic norms regulated by Traffic Authorities can therefore be translated in a straight forward manner into norms belonging to the Electronic Institution. Norms within this normative environment are thus related to actions performed by cars (in fact, in our case, they are always restricted to that). Additionally, norms do have associated penalties that are imposed to those cars refusing or failing to follow them. In our case study, we assume that the Traffic Authority is always aware of norm violations: cars may or may not respect rules, but they are not able to avoid the consequences of their application. Furthermore, our Electronic Institution is able to change norms based on its goals – just as traffic authorities do modify their traffic rules– and, therefore, it is considered to be autonomic.

Our AEI sets up a normative environment where cars do have a limited amount of credit (just as some real world driving license credit systems) so that norm offenses cause credit reductions. The number of points subtracted for each traffic norm violation is specified by the sanction associated to each norm, and this sanction can be changed by the regulation authority (that is, our AEI) if its change leads –or contributes to– the accomplishment of goals. Eventually, those cars without any remaining points are forbidden to circulate. On the other hand, we assume a non-closed world, so expelled cars are replaced by new ones having the total amount of points.

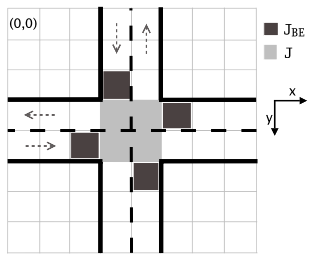

Fig. 2. Grid environment representation of a 2-lane road junction.

Getting into more detail, this paper focuses on a two-road junction. It is a very restrictive problem setting, but it is complex enough to allow us to tackle the problem without losing control of all the factors that may influence the results. In particular, no traffic signals (neither yield or stop signals nor traffic lights) are considered, therefore, cars must only coordinate by following the traffic norms imposed by the AEI. Our institution is required to define these traffic norms based on general goals such as minimization of the number of accidents or deadlock avoidance.

We model the environment as a grid composed by road and field cells. Road cells define 2 orthogonal roads that intersect in the center (see Fig. 2). Discretization granularity is such that cars have the size of a cell. As section 4.2 details, our model has been developed with the Simma tool [20]. Although the number of road lanes can be changed parametrically, henceforth we assume the 2-lane case. Next subsections are devoted to define this “toy problem” and present our solution proposal in terms of it. But before that, some nomenclature definitions are introduced :

- A_i: an agent i, agents correspond to cars.
- t: time step. Our model considers discrete time steps, also known as ticks. Time is necessary for sequentiality considerations (i.e., time ordering).

- (x_0^J, y_0^J): top left cell inside the road junction area.
- (J_x, J_y): size in x, y of our road junction area. Thus, $J_x \times J_y$ is the total amount cells inside the junction. If we consider 2-lane roads, then $(J_x, J_y) = (2, 2)$.
- J: inner road junction area $J = \{(x, y) \mid x \in [x_0^J, x_0^J + J_x - 1],\ y \in [y_0^J, y_0^J + J_y - 1]\}$. Considering the 4 J cells in the junction area of Fig. 2: $J = \{(x_0^J, y_0^J), (x_0^J + 1, y_0^J), (x_0^J, y_0^J + 1), (x_0^J + 1, y_0^J + 1)\}$.
- J_{BE}: Junction Boundary Entrance, set of cells surrounding the junction that can be used by cars to access it. They correspond to cells near by the junction that belong to incoming lanes. Fig. 2 depicts $J_{BE} = \{(x_0^J, y_0^J - 1), (x_0^J - 1, y_0^J + J_y - 1), (x_0^J + J_x - 1, y_0^J + J_y, (x_0^J + J_x, y_0^J))\}$. Nevertheless, the concept of boundary is not restricted to adjacent cells: a car can be also considered to be coming into the junction if it is located one –or even a few– cells away from the junction.
- (x_i^t, y_i^t): position of car A_i at time t, where $(x, y) \in \mathbb{N} \times \mathbb{N}$ stands for a cell in the grid.
- (h_{ix}^t, h_{iy}^t): heading of car A_i, which is located in (x, y) at time t. Heading directions run along x, y axes and are considered to be positive when the car moves right or down respectively. In our orthogonal environment, heading values are: 1 if moving right or down; -1 if left or up; and 0 otherwise (i.e., the car is not driving in the axis direction). In this manner, car4's heading on the right road of Fig. 3 is (-1,0).

4.1 AEI specification

4.1.1 Environment

As mentioned above, we consider the environment to be a grid. This grid is composed of cells, which can represent roads or fields. The main difference among these two types is that road cells can contain cars. Indeed, cars move among road cells along time.

Fig. 2 depicts a 8×8 grid example. The top left corner of the grid represents the origin in the x, y axes. Thus, in the example, cell positions range from (0,0) in the origin up to (7,7) at the bottom-right corner. Additionally, a cell is a road if one of its x, y coordinates belong to J inner junction area (see previous definition).

This grid environment is defined as:

$$P_e = \langle (x, y, \alpha, r, d_x, d_y) \mid 0 \le x \le max_x,\ 0 \le y \le max_y,\ \alpha \subseteq P(A),\ r \in [0, 1],\ d_x \in [-1, 0, 1],\ d_y \in [-1, 0, 1]\ \rangle$$

being x and y the cell position, α defines the set of agents inside the grid cell (x, y), r indicates whether this cell represents a road or not, and, in case it is a road, d_x and d_y stand for the lane direction, whose values are the same as the ones for car headings. Notice that the institution can observe the environment properties along time, we use P_e^t to refer the values of the grid environment at a specific time t. This discretized environment can be observed both by the institution and cars. The institution observes and keeps track of its evolution along time, whilst cars do have locality restrictions on their observations.

4.1.2 Agents

Lets consider $A = \langle A_1, ..., A_n \rangle$ to be a finite set of n agents in the institution. As mentioned before, agents correspond to cars that move inside the grid environment, with the restriction that they can only move within road cells. Additionally, agents are given an account of points which decreases with traffic offenses. The institution forbids agents to drive without points in their accounts.

The institution can observe the $P_a = \langle a_1, \ldots, a_n \rangle$ agents' institutional properties, where

$$a_i = \langle x_i,\ y_i,\ h_{ix},\ h_{iy},\ speed_i,\ indicator_i,\ offenses_i,\ accidents_i,\ distance_i,\ points_i \rangle$$

These properties stand for: car A_i's position within the grid, its heading, its speed, whether the car is indicating a trajectory change for the next time step (that is, if it has the intention to turn, to stop or to move backwards), the norms being currently violated by A_i, wether the car is involved in an accident, the distance between the car and the car ahead of it; and, finally, agent A_i's point account. Notice that the institution can observe the agent properties along time, we use a_i^t to refer the agent A_i's properties at a specific time t.

4.1.3 Reference values

In addition to car properties, the institution is able to extract reference values from the observable properties of the environment, the participating agents and the institution. Thus, these reference values are computed as a compound of other observed values. Considering our road junction case study, it is possible to identity different reference values:

$$V = \langle num\_collisions,\ num\_crashed,\ num\_offenses,\ num\_blocked \rangle$$

where **num_collisions** indicates total number of collisions for the last t_w ticks $(0 \le t_w \le t_{now})$:

$$num\_collisions = \sum_{t=t_{now}-t_w}^{t_{now}} \sum_{e \in P_e^t} f(e_{\alpha^t}) \tag{1}$$

being P_e^t the values of the grid environment at time t, e_{α^t} the α^t component of element $e \in P_e^t$ and

$$f(e_{\alpha^t}) = \{^{1\ if\ |e_{\alpha^t}|>1}_{0\ otherwise} \tag{2}$$

Similarly, **num_offenses** indicates the total number of offenses accumulated by all agents during last t_w ticks $(0 \le t_w \le t_{now})$:

$$num\_offenses = \sum_{t=t_{now}-t_w}^{t_{now}} \sum_{i=0}^{|A|} offenses_i^t \tag{3}$$

Furthermore, *num_crashed* counts the number of cars involved in accidents for the last t_w ticks:

$$num\_crashed = \sum_{t=t_{now}-t_w}^{t_{now}} \sum_{i=0}^{|A|} accidents_i^t \tag{4}$$

And finally, *num_blocked* describes how many cars have been blocked by other cars for last t_w ticks:

$$num\_blocked = \sum_{t=t_{now}-t_w}^{t_{now}} \sum_{i=0}^{|A|} blocked(a_i, t) \tag{5}$$

where $blocked(a_i, t)$ is a function that indicates if the agent a_i is blocked by another agent a_j in time t.

$$blocked(a_i, t) = \begin{cases} 1 & if\ \exists\ e \in P_e^t \mid (e_{x^t} = x_i^t + h_{ix}^t\ \&\ \\ & e_{y^t} = y_i^t + h_{iy}^t\ \&\ |e_{\alpha^t}| \ge 1\ \&\ \\ & \exists\ a_j \in e_{\alpha^t}\ so\ that\ speed_j^t = 0) \\ \\ 0 & otherwise \end{cases} \tag{6}$$

being $e_{x^t},\ e_{y^t},\ e_{\alpha^t}$ the $x^t,\ y^t,\ \alpha^t$ components of element $e \in P_e^t$.

4.1.4 Goals

Goals are in fact institutional goals. The aim of the traffic authority institution is to accomplish as many goals as possible.The institution tries to accomplish these goals by defining a set of norms (see subsection 4.1.5).

Table 1. Right priority norm.

	Action	Pre-conditions	Consequence
Right priority	$in(a_i, J_{BE}, t-1)\wedge$ $in(a_i, (x_i^{t-1} + h_{ix}^{t-1}, y_i^{t-1} + h_{iy}^{t-1}), t)\wedge$ $\neg indicator(a_i, right, t-1)$	$right(a_i, a_j, t-1)$	$points_i^t = points_i^t - fine_{right}$

Institutional goals are defined as constraints upon a combination of reference values. Considering our scenario, it is possible to define restrictions as intervals of acceptable values for the previous defined reference values (V) so that the institution considers it has accomplished its goals if V values are within their corresponding intervals. In fact, the aim is to minimize the number of accidents, the number of traffic offenses, as well as the number of blocked cars by establishing the list of institutional goals G as:

$$G = \langle\ g(num\_collisions) \in [0, MaxCollisions],\\ g(num\_offenses) \in [0, MaxOffenses],\\ g(num\_crashed) \in [0, MaxCrashed],\\ g(num\_blocked) \in [0, MaxBlocked]\ \rangle$$

Having more than one institutional goal requires to combine them. We propose an **objective function** [26] that favors high goal satisfaction while penalizing big differences among them:

$$O(V) = \sum_{i=1}^{|G|} w_i \sqrt{f(g_i(V), [m_i, M_i], \mu_i)} \qquad (7)$$

where $1 \le i \le |G|$, $w_i \ge 0$ are weighting factors such that $\sum w_i = 1$, g_i is a function over the reference values, $\mu_i \in [0,1]$ and f is a function that returns a value $f(x, [m, M], \mu) \in [0,1]$ representing the degree of satisfaction of a goal:

$$f(x, [m, M], \mu) = \begin{cases} \dfrac{\mu}{e^{k\frac{m-x}{M-m}}} & x < m \\ 1 - (1-\mu)\dfrac{x-m}{(M-m)} & x \in [m, M] \\ \dfrac{\mu}{e^{k\frac{x-M}{M-m}}} & x > M \end{cases} \qquad (8)$$

4.1.5 Norms

Autonomic Electronic Institutions use norms to try to accomplish goals. Norms have associated penalties that are imposed to those cars refusing or failing to follow them. These penalties can be parameterized to increase its persuasiveness depending on the agent population behavior.

Considering a road junction without traffic signals, priorities become basic to avoid collisions. We consider, as in most continental Europe, that the default priority is to give way to the right. This norm prevents a car A_i located on the Junction Boundary Entrance (J_{BE}) to move forward or to turn left whenever there is another car A_j on its right. For example, car 1 in Fig. 3 must wait for car 2 on its right, which must also wait for car 3 at the bottom J_{BE}. The formalization in table 1 can be read as follows: "if car A_i moves from a position in J_{BE} at time $t-1$ to its next heading position at time t without indicating a right turn, and if it performs this action when having a car A_j at the J_{BE} on its right, then the institution will fine A_i by decreasing its points by a certain amount" (see Fig. 4).

Where the predicate $\mathbf{in}(a_i, Region, t)$ in table 1 is equivalent to $\exists(x, y, \alpha^t, r, d_x, d_y) \in P_e^t$ so that $(x, y) \in Region\ and\ a_i \in \alpha^t$ and $\mathbf{right}(a_i, a_j, t)$ is a boolean function that returns true if car A_j is located at J_{BE} area on the right side of car A_i. For the 2-lane J_{BE} case in Fig. 2, it corresponds to the formula: $(x_i^t - h_{iy}^t + h_{ix}^t J_x,\ y_i^t + h_{ix}^t + h_{iy}^t Jy) = (x_j^t, y_j^t)$.

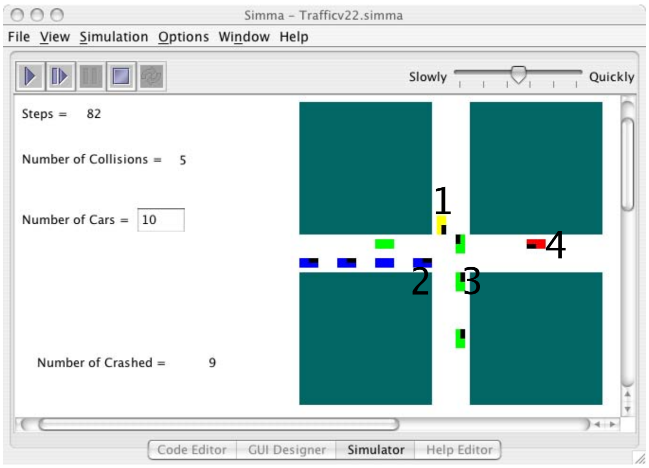

Fig. 3. Priority to give way to the right (Simma tool screenshot).

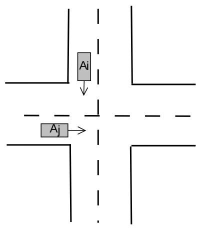

Fig. 4. Priority to give way to the right.

Similarly, it is possible to define an additional norm that is somehow related to the previous 'right priority norm'. Lets name it 'front priority norm'. It applies when two cars A_i, A_j reach Junction Boundary Entrance areas (J_{BE}) located at opposite lines, and one of them (A_i in Fig. 5) wants to turn left. Car A_i turning left may interfere A_j's trajectory, and therefore, this norm assigns priority to A_j so that A_i must stop until its front J_{BE} area is clear. Otherwise A_i will be punished with the corresponding $fine_{front}$ fee.

Table 2 shows the formalization of this norm where $\mathbf{front}(a_i, a_j, t)$ is a boolean function that returns true if car A_j is located in front of car A_i at time t. In an orthogonal environment, this function can be easily computed by comparing car headings $((h_{ix}^t, h_{iy}^t), (h_{jx}^t, h_{jy}^t))$ by means of the boolean formula $(h_{ix}^t h_{jx}^t + h_{iy}^t h_{jy}^t) == -1$.

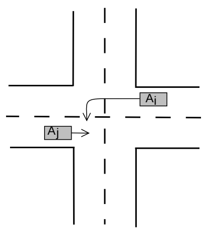

Fig. 5. Priority to give way to the front.

Table 2. Front priority norm.

	Action	Pre-conditions	Consequence
Front priority	$in(a_i, J_{BE}, t-1) \wedge$ $in(a_i, (x_i^{t-1} + h_{ix}^{t-1}, y_i^{t-1} + h_{iy}^{t-1}), t) \wedge$ $indicator(a_i, left, t-1)$	$in(a_j, J_{BE}, t-1) \wedge$ $front(a_i, a_j, t-1)$	$points_i^t = points_i^t - fine_{front}$

4.1.6 Performative Structure

As introduced in section 2, an AEI involves different groups of agents playing different roles within scenes in a performative structure. Each scene is composed of a coordination protocol along with the specification of the roles that can take part in the scene. Our case study particularizes the Performative Structure component so that we define it as being formed by a single traffic scene with a single agent role, i.e. agent cars. Regarding performative structure transition, γ, we do not exploit it in our current implementation of the traffic case study so that we just fix the number of car agents that interact within the scene. Nevertheless it could be done by defining new institutional agent roles contributing to the EAI goal fullfilment.

4.2 Experimental Settings and Design

As a proof of concept of our proposal in section 2, we have designed an experimental setting that implements the traffic case study. In this preliminary experiment we consider two institutional goals related to both $num\_collisions$ and $num\_offenses$ reference values; and both right and front priority norms in tables 1 and 2. Institutional goals are combined with the objective function introduced in section 4.1.4, assuming both weights are equal to 0.5 so that both goals are considered to be equally important. On the other hand, norms are parameterized through its fines (i.e., points to subtract to the car falling to follow the corresponding norm).

The 2-road junction traffic model has been developed with Simma [20], a graphical MAS simulation tool shown in Fig. 3, in such way that both environment and agents can be easily changed. In our experimental settings, we have modeled the environment as a 16×16 grid where both crossing roads have 2 lanes with opposite directions. Additionally, the environment is populated with 10 cars, having 40 points each.

Our institution can observe the agents properties for each tick and can keep a record of them in order to refer to past ticks. In fact, the institution usually determines traffic offenses by analyzing agent actions along time. Agent actions are observed through consecutive car positions and indicators (notice that the usage of indicators is compulsory for cars in this problem set up). During our discrete event simulation, the institution replaces those cars running out of points by new cars, so that the cars' population is kept constant. Cars follow random trajectories at a constant 1-cell/tick speed and they collision if two or more cars run into the same cell. In that case, the involved cars do remain for two ticks in that cell before they can start following a new trajectory.

Cars correspond to agents without learning skills. They just move, within the traffic scene, based on their trajectories and institutional norms. Agents have local information about their environment (i.e., grid surrounding cells) and know whether their next movements will violate a norm and what fine will be thus applied. Agents decide whether to comply with a norm based on three parameters: $\langle fulfill\_prob$, $high\_punishment$, $inc\_prob\rangle$. Being $fulfill\_prob \in [0,1]$ the probability of complying with norms that is initially assigned to each agent, $high\_punishment \in \mathbb{N}$ the fine threshold that causes an agent to consider a fine to be high enough to reconsider the norm compliance, and $inc\_prob \in [0,1]$ the probability increment that is added to $fulfill\_prob$ when the fine threshold is surpassed by the norm being violated. In summary, agents decide whether they keep moving regardless of violated norms or they stop in order to comply with norms based on a probability that is computed as:

$$final\_prob = \begin{cases} fulfill\_prob & fine \leq hp \\ fulfill\_prob + inc\_prob & fine > hp \end{cases} \quad (9)$$

where hp is $high\_punishment$.

Our goal is to adapt norms to agent behaviors by applying Genetic Algorithms (GA)[1] to accomplish institutional goals, that is, to maximize the objective function, which comprises both number of collisions and number of offenses. We propose to learn norms for different agent population behaviors by simulation as mentioned in section 3. Once specified different agent populations, it is possible to run a genetic algorithm per population. Therefore, norm adaptation is implemented as a learning process of the "best" norm parameters. In our experiments, Genetic Algorithms run 15 generations of 10 individuals. An individual corresponds to a list of a binary codifications of specific values for the right-norm-penalty and front-norm-penalty institution parameters. Crossover among individuals is chosen to be singlepoint and a mutation rate of 5% is applied. The fitness function for individual evaluation corresponds to the objective function described above, which is computed as an average of 5 different 2000-tick-long simulations for each model setting (that is, for each set of parameters):

$$O(V) = \frac{1}{2}\sqrt{f(g(num\_collisions), [0, MaxCollisions], \frac{1}{2})} +$$

$$\frac{1}{2}\sqrt{f(g(num\_offenses), [0, MaxOffenses], \frac{1}{2})} \quad (10)$$

where $g(num\_collisions)$ and $g(num\_offenses)$ correspond to average values of both reference values averaged from 5 different simulations; and $f(x, [m, M]) \in [0,1]$ represents the goal satisfaction.

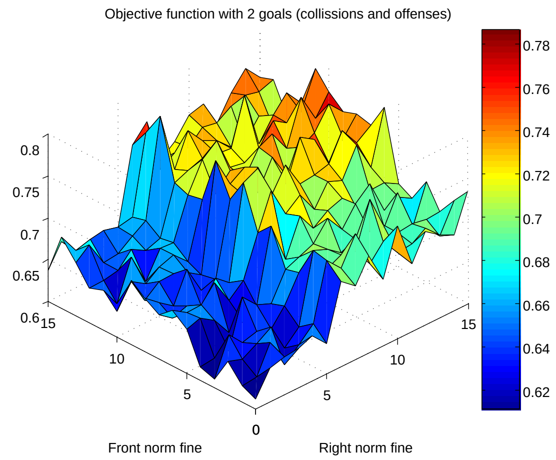

Fig. 6. Objective function.

[1] We use a genetic algorithm Toolbox [3].

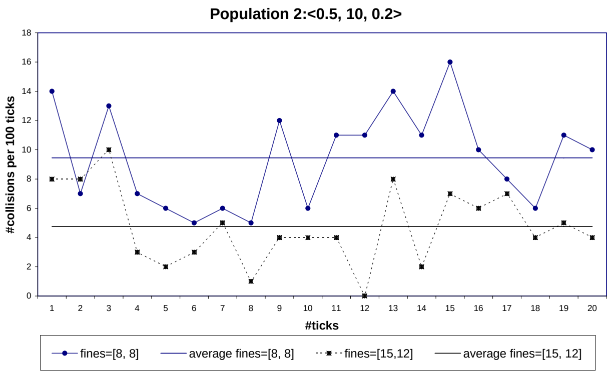

a)

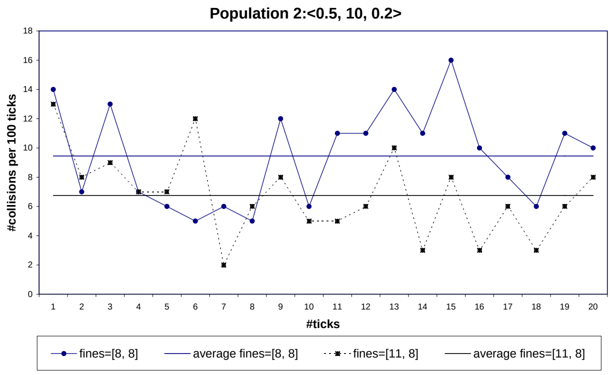

b)

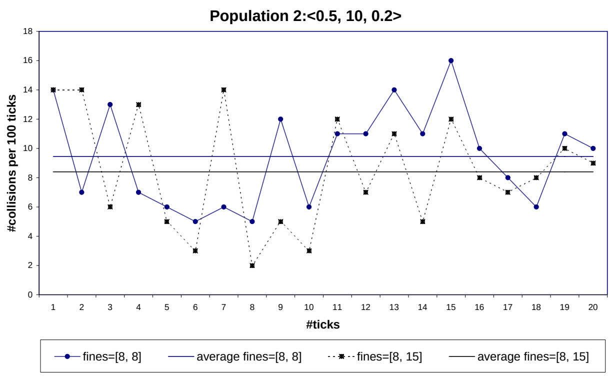

c)

Fig. 7. Result comparison for different fine values in terms of number of collisions per 100 ticks along a 2000-tick simulation. (8,8) fines vs. a) (15,12), b) (11,8), and c) (8,15).

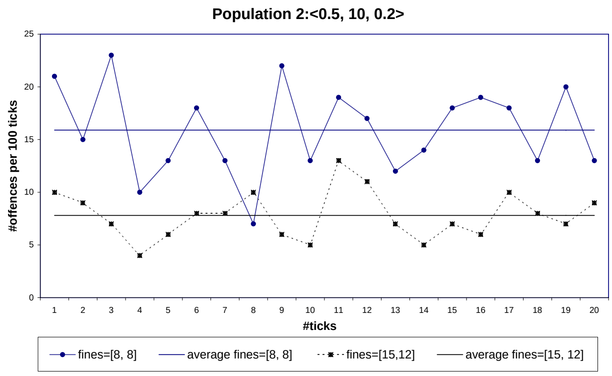

a)

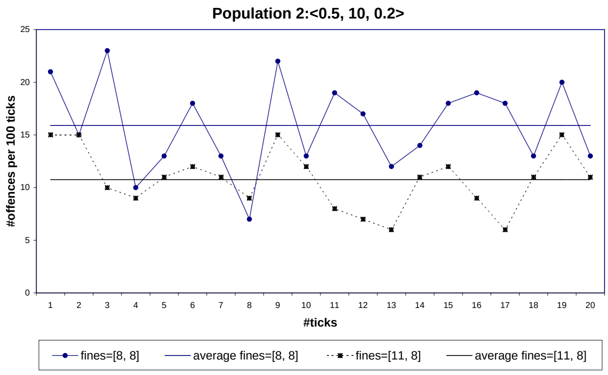

b)

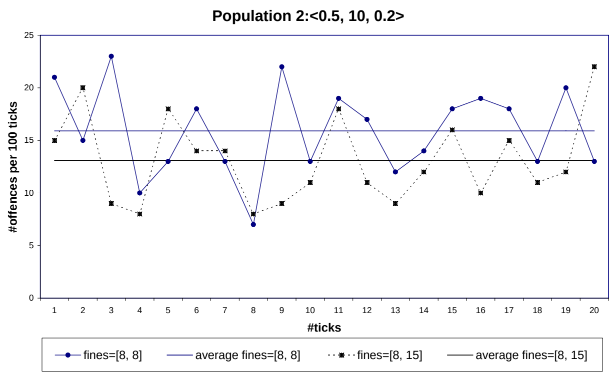

c)

Fig. 8. Result comparison for different fine values in terms of number of offences per 100 ticks along a 2000-tick simulation. (8,8) fines vs. a) (15,12), b) (11,8), and c) (8,15).

5. Results

From the experimental settings specified above, we have run experiments for three different agent populations. These populations are characterized by their norm compliance parameters, being $fulfill\_prob = 0.5$ and $inc\_prob = 0.2$ for the three of them whereas $high\_punishment$ varies from 5 for the first, to 10 for the second, up to 14 for the third (see table 3).

Since both right and front priority norms contribute to reduce accidents, our AEI must learn how to vary its fine parameters to increase its persuasiveness for agents, and eventually, to accomplish the normative goals of minimizing the total number of collisions and the total number of offenses. Table 3 shows the fine vectors that our AEI has learned for population1, population2, and population3 respectively: $(14, 11)$, $(15, 12)$, and $(15, 15)$, where the first parameter stands for $fine_{right}$ and the second one for $fine_{front}$. Agent's behavior is so that its probability of complying with norms increases –by inc_prob– only when the fine is larger than $high\_punishment$ (see equation 9). Therefore, any fine value higher than the population's $high\_punishment$ value will have the same effect, and thus, will generate equivalent goal satisfaction degrees. As a result, the AEI must learn the best combination of parameters ($fine_{right}$ and $fine_{front}$) according to the 2-goal objective function and to the agents' behavior.

Henceforth, the obtained results allow us to state that the

AEI succeeds in learning the norms that best accomplish its overall goal because both learned fines are larger than the population's *high_punishment* value in all three cases.

In order to further analyse the obtained results, next figures 7 and 8 give some more detail about the performance of agent populations for different norm fine values. In particular, these figures illustrate the performance of agent population2 having *high_punishment* threshold equal to 10. First chart in Fig. 7 compares the number of collisions per 100 ticks when fines are $(8, 8)$ with the resulting number of collisions when they are $(15, 12)$, which correspond to the previously learnt values for this population. First chart in Fig. 8 provides an analogous comparison, but for number of offenses instead. Similarly, second and third charts in figures 7 and 8 compare results for the same fine values of $(8, 8)$ with fine values $(11, 8)$ in second charts and with $(8, 15)$ values in third charts. For all three cases in both figures, it can be observed that the number of collisions and the number of offenses for fines $(8, 8)$ keep above the other ones both in average and along the curve that results from a simulation of 2000 ticks. Differences are greater for first charts, which compare with learnt fine values. As expected, the reason is that both values $(8, 8)$ are smaller than the *high_punishment* value for the second agent population (which is 10), whereas both learnt fine values $(15, 12)$ are greater than it. This implies agents to increment its probability of complying with norms. On the contrary, only one fine value is greater than the *high_punishment* value for second and third charts and, therefore, the decrease in the number of collisions and offenses is moderated.

The effect of *high_punishment* can also be appreciated in Fig. 6, wich shows the overall goal function for population1. This 3D chart depicts all the values of the goal function using $fine_{right}$ and $fine_{front}$ parameters, in the 16 domain, that correspond to x and y axis. As it can be appreciated, a discontinuity appears for x and y values higher than 5, which corresponds to the *high_punishment* value for population1. In fact, the region defined for $fine_{right},\ fine_{front} > 5$ includes maximal satisfaction degrees, and thus the genetic algorithm should provide solutions belonging to this area, which has been the case in our experiments.

Table 3. Learning results for three different agent populations.

Parameters	population1	population2	population3
$fulfill\_prob$	0.5	0.5	0.5
$high\_punishment$	5	10	14
inc_prob	0.2	0.2	0.2
Learned right fine	**14**	**15**	**15**
Learned front fine	**11**	**12**	**15**

6. Discussion and Future work

Within the area of Multi-Agent Systems, adaptation has been usually envisioned as an agent capability. In this manner, works such as the one by Excelente-Toledo and Jennings [9] propose a decision making framework that enables agents to dynamically select the coordination mechanism that is most appropriate to their circumstances. Hübner et al. [15] propose a model for controlling adaptation by using the $\mathcal{M}$OISE+ organization model. Agents in this model adapt their MAS organization to both environmental changes and their own goals. In [11] Gasser and Ishida presented a general distributed problem-solving model which can reorganize its architecture, in [16] Ishida and Yokoo introduce two new reorganization primitives that change the population of agents and the distribution of knowledge in an organization, whereas Horling et al. [14] propose an approach where the members adapt their own organizational structures at runtime. On the other hand, it has been long stated [4] that agents working in a common society need norms to avoid and solve conflicts, make agreements, reduce complexity, or to achieve a social order. Both approaches –i.e. adaptation and norms– have been considered together by Lopez-y-Lopez et al. [19], where agents can adapt to norm-based systems and they can even autonomously decide its commitment to obey norms in order to achieve associated institutional goals. This adaptation from the point of view of agents in these related works is the most remarkable difference with the approach presented in this paper, which focuses on adapting the institution –that is, the authority issuing norms– rather than adapting the agents. Institution adaptation is accomplished by changing norms autonomously (as opposite to the work by Hoogendoorn et al. [13], which is based on design considerations). Therefore, we do not select norms at design stages as it is done by Fitoussi and Tennenholtz [10], who do it so by proposing the notions of minimality and simplicity as selecting criteria. They study two basic settings, which include Automated-Guided-Vehicles (AGV) with traffic laws, by assuming an environment that consists of (two) agents and a set of strategies available to (each of) them. From this set, agents devise the appropriate ones in order to reach their assigned goals without violating social laws, which must be respected.

Regarding the traffic domain, MAS has been previously applied to it [22] [6] [5]. For example, Camurri et al. [2] propose two field-based mechanisms to control cars and traffic-lights. Its proposed driving policy guides cars towards their (forward) destinations avoiding the most crowded areas. On the other hand, traffic light control is based on a linear combination between a distance field and the locally perceived traffic field. Additionally, authors combine this driving policy and traffic light control in order to manage to avoid deadlocks and congestion.

Traffic has also been widely studied outside the scope of MAS, for example, the preliminary work by [23] used Strongly Typed Genetic Programming (STGP) to control the timings of traffic signals within a network of orthogonal intersections. Their evaluation function computed the overall delay.

This paper presents AEI as an extension of EIs with autonomic capabilities. In order to test our model, we have implemented a traffic AEI case study, where the AEI learns two traffic norms in order to fulfill its goals while adapting to different agent populations. We are currently working on extending the performative structure adaptation to include institutional agents with a new role representing the Traffic Authority employees. They would be in charge of detecting norm violations so that we could refer to them as police agents. Each police agent would be able to detect only a portion of the total number of norm violations that car agents actually do. Therefore, the number of police agents in the traffic scene would directly affect the number of detected norm violations, and thus, the overall quantity of penalties imposed to car agents.

As future work, and since this basically represents a centralized scenario, we plan to develop a more complex traffic network, adopting a decentralized approach such that the rules of a regulatory institution are distributed across different areas (e.g., junctions). In this manner, rules may have different local scopes. Furthermore, different areas containing different rules are expected to coordinate and cooperate in order to ensure that the system's performance fulfils the required reference values. On the other hand, we plan to extend the autonomic capabilities to other than self-configuration. In fact, additional research is required in order to extend both our traffic model and the institutional adaptation capabilities so that an autonomic electronic institution will not only learn the most appropriate norms for a given agent population, but it will be able to adapt to any change in the population.

Acknowledgements

This work was partially funded by the Spanish Education and Science Ministry as part of the Web-i-2 project (TIC-2003-08763-C02-01) and by the Spanish Council for Scientific Research as part of the 2006 5 OI 099 project. The first author enjoys an FPI grant (BES-2004-4335) from the Spanish Education and Science Ministry.

References

[1] J L Arcos, M Esteva, P Noriega, J A. Rodríguez-Aguilar, and C. Sierra, Engineering open environments with electronic institutions. *Engineering Applications of Artificial Intelligence*, (18):191–204, 2005.

[2] M Camurri, M Mamei, and F Zambonelli, Urban traffic control with co-fields. In *Proc. of E4MAS Workshop at AAMAS 2006*, pages 11–25, 2006.

[3] A Chipperfield and P Fleming, The matlab genetic algorithm toolbox. *IEE Colloquium Applied Control Techniques Using MATLAB*, pages 10–14, 1995.

[4] R Conte, R Falcone, and G Sartor, Agents and norms: How to fill the gap? *Artificial Intelligence and Law*, (7):1–15, 1999.

[5] A Doniec, S Espié, R Mandiau, and S Piechowiak, Dealing with multi-agent coordination by anticipation: Application to the traffic simulation at junctions. In *EUMAS*, pages 478–479, 2005.

[6] K Dresner and P Stone, Multiagent traffic management: An improved intersection control mechanism. In *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 471–477, New York, NY, July 2005. ACM Press.

[7] M Esteva, *Electronic Institutions: from specification to development.* IIIA PhD Monography. Vol. 19, 2003.

[8] A Etzioni, *Modern Organizations.* Prentice-Hall, Englewood Cliffs, NJ, 1964.

[9] C B Excelente-Toledo and N. R. Jennings, The dynamic selection of coordination mechanisms. *Autonomous Agents and Multi-Agent Systems*, 9(1-2):55–85, 2004.

[10] D Fitoussi and M Tennenholtz, Choosing social laws for multi-agent systems: Minimality and simplicity. *Artificial Intelligence*, 119(1-2):61–101, 2000.

[11] L Gasser and T Ishida, A dynamic organizational architecture for adaptive problem solving. In *Proc. of AAAI-91*, pages 185–190, Anaheim, CA, 1991.

[12] M Gelfond and V Lifschitz, Representing action and change by logic programs. *Journal of Logic Programming*, 17:301–321, 1993.

[13] M Hoogendoorn, C Jonker, and J Treur, Redesign of organizations as a basis for organizational change. In *Proc. of the Workshop on Coordination, Organization, Institutions, and Norms in Agent Systems (COIN @ ECAI 2006)*, page to appear, 2006.

[14] B Horling, B Benyo, and V Lesser, Using Self-Diagnosis to Adapt Organizational Structures. *Proceedings of the 5th International Conference on Autonomous Agents*, pages 529–536, June 2001.

[15] J F Hübner, J S Sichman, and O Boissier, Using the $\mathcal{M}$oise+ for a cooperative framework of mas reorganisation. In *LNAI - Proc. of the 17th Brazilian Symposium on Artificial Intelligence (SBIA'04)*, volume 3171, pages 506–515. Springer, 2004.

[16] T Ishida and M Yokoo, Organization self-design of distributed production systems. *IEEE Trans. Knowl. Data Eng.*, 4(2):123–134, 1992.

[17] N R Jennings, K Sycara, and M Wooldridge, A roadmap of agent research and development. *Autonomous Agents and Multi-agent Systems*, 1:275–306, 1998.

[18] J O Kephart and D M Chess, The vision of autonomic computing. *IEEE Computer*, 36(1):41–50, 2003.

[19] F López-López, M Luck, and M d'Inverno, Constraining autonomy through norms. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 674–681, New York, NY, USA, 2002. ACM Press.

[20] M López-Sánchez, X Noria, J A Rodríguez-Aguilar, and N Gilbert, Multi-agent based simulation of news digital markets. *International Journal of Computer Science and Applications*, 2(1):7–14, 2005.

[21] M Luck, P McBurney, O Shehory, and S Willmott, *Agentlink Roadmap.* Agenlink.org, 2005.

[22] S Luke, C Cioffi-Revilla, L Panait, and K Sullivan, Mason: A new multi-agent simulation toolkit. In *Proceedings of the 2004 SwarmFest Workshop*, page 8, 2004.

[23] D J Montana and S Czerwinski, Evolving control laws for a network of traffic signals. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 333–338, Stanford University, CA, USA, 28–31 July 1996. MIT Press.

[24] P Noriega, *Agent-Mediated Auctions: The Fishmarket Metaphor.* IIIA Phd Monography. Vol. 8, 1997.

[25] D North, *Institutions, Institutional Change and Economics Perfomance.* Cambridge U. P., 1990.

[26] J Rodríguez-Aguilar, A Reyes-Moro, M López-Sánchez, and J Cerquides, Enabling assisted strategic negotiations in actual-world procurement scenarios. *Electronic Commerce Research*, to appear 2006.

[27] Q Yang, *A Simulation Laboratory for Evaluation of Dynamic Traffic Management Systems.* PhD thesis, MIT, 1997.

Author Bios

Eva Bou is a Ph.D. student at Artificial Intelligence Research Institute (IIIA) of the Spanish Research Council (CSIC). In 2003 she got a Computer Science Engineer degree (B.Sc.) from the Universitat Autonoma de Barcelona (UAB). Her current focus is upon electronic instutions, learning and multiagent systems. She is co-author of some publications in multiagent areas.

Maite López-Sánchez (PhD) is a researcher at the WAI: Volume Visualization and Artificial Intelligence Research Group at the University of Barcelona (UB). Previously, she was a research manager at iSOCO's innovation department and visiting researcher at the Robotics Laboratory from the University of Southern California USC (USA). She is a Computer Science Engineer (a B.Sc.) and holds a Master's degree in Computer Vision & Artificial Intelligence (AI). She doctorated in AI at the Artificial Intelligence Research Institute (IIIA-CSIC). During her career, she has been doing research in areas such as multi-agent systems, classification, robotics, or e-commerce. She has published papers at international journals and conferences as well as contributions to books. Additionally, she has participated in several research projects (Spanish, DARPA and European IST).

Juan A. Rodríguez-Aguilar (PhD) is a permanent researcher at the Artificial Intelligence Research Institute (IIIA) of the Spanish Research Council (CSIC). Previously, he was contracted as a researcher at the Massachusetts Institute of Technology (MIT) as well as at iSOCO's innovation department. He doctorated in AI at the Artificial Intelligence Research Institute (IIIA-CSIC). During his career, he has been doing research in the areas of multi-agent systems and electronic commerce. He has participated in many international research projects and has published his works at books, international conferences as well as relevant journals.