

Adaptation of Autonomic Electronic Institutions through norms and institutional agents.

Eva Bou¹ and Maite López-Sánchez² and J. A. Rodríguez-Aguilar¹

¹ IIIA - CSIC Artificial Intelligence Research Institute, Campus UAB 08193
Bellaterra, Spain, email: {ebm, jar}@iia.csic.es

² WAI, Volume Visualization and Artificial Intelligence, MAiA Dept., Universitat de
Barcelona, email: maite@maia.ub.es

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 external agents. In this paper we focus on the extension of EIs with autonomic capabilities to allow them to yield a dynamical answer to changing circumstances through norm adaptation and changes in institutional agents.

Key words: 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 [1], 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 [2], 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 processing in order to achieve its goals; it must be able to diagnose the causes of its eventual malfunctions and repair itself; and it must act in accordance to and operate into a heterogeneous and open environment.

In what follows we argue that EIs [3] are a particular type of self-* system. When looking at computer-mediated interactions we regard Electronic Institutions (EI) as regulated virtual environments wherein the relevant interactions among participating agents take place. EIs have proved to be valuable to develop open agent systems [4]. 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 dealing with self-interested agents.

The main goal of this work consists in studying how to endow an EI with autonomic capabilities that allow it to yield a dynamical answer to changing circumstances through the adaptation of its regulations. Among all the characteristics that define an autonomic system we will focus on the study of self-configuration as pointed out in [2] 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. In section 2 we introduce 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 empirical results. Finally, section 6 summarizes some conclusions and related work and outlines paths to future research.

2 Autonomic Electronic Institutions

The idea behind EIs [3] is to mirror the role traditional institutions play in the establishment of "the rules of the game" –a set of conventions that articulate participants' interactions. The main goal of EIs is the enactment of a constrained environment that shapes open agent societies. EIs structure agent interactions, establishing what agents are permitted and forbidden to do as well as the consequences of their actions.

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 [5] 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 nor the engineering framework in [6] support their adaptation so that an EI can self-configure. Thus, in what follows we study how to extend the current notion of EI to support self-configuration.

First of all, notice that in order for EIs to adapt, we believe that a "rational" view must be adopted (likewise the rational view of organizations in [7]) 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, we observe that according to [3] 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 $\langle PS, N, DF, G, P_i, P_e, P_a, V, \delta, \gamma \rangle$ 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;
- $P_i = \langle i_1, \dots, i_s \rangle$ 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;
- $P_e = \langle e_1, \dots, e_r \rangle$ 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;
- $P_a = \langle a_1, \dots, a_n \rangle$ stands for the values that characterize the institutional state of the agents in A , where $a_j = \langle a_{j_1}, \dots, a_{j_m} \rangle$, $1 \leq j \leq n$ stands for the institutional state of agent A_j ;
- V stands for a finite set of reference values;
- $\delta : N \times G \times V \rightarrow N$ 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, we dissect the new elements composing an AEI.

2.1 Goals

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 $A = \{A_1, \dots, A_n\}$ 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 $\langle a_{i_1}, \dots, a_{i_m} \rangle$ 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 $V = \langle v_1, \dots, v_q \rangle$ 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, we can turn our attention to 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 we define the goals of an AEI as a finite set of constraints $G = \{c_1, \dots, c_p\}$ where each c_i is defined as an expression $g_i(V) \triangleleft [m_i, M_i]$ where $m_i, M_i \in \mathbb{R}$, \triangleleft 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.

2.2 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, in this paper we adhere to a simple definition of norms as effect propositions as defined in [8]:

Definition 2. *An effect proposition is an expression of the form*

$$A \text{ causes } F \text{ if } P_1, \dots, P_n$$

where A is an action name, and each of $F, P_1, \dots, P_n (n \geq 0)$ is a fluent expression. About this proposition we say that it describes the effect of A on F , and that P_1, \dots, P_n are its preconditions. If $n = 0$, we will drop it and write simply 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 $N_i \in N$, $i = 1, \dots, n$, has a set of parameters $\langle p_{i,1}^N, \dots, p_{i,m_i}^N \rangle \in \mathbb{R}^{m_i}$. Hence, changing the values of these parameters means changing the norm. In fact these parameters correspond to the variables in the *norm transition function* that will allow the institution to adapt under changing situations.

2.3 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, we must choose the values that we aim 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, $S_i \in PS$, $i = 1, \dots, t$, has a set of parameters $\langle p_{i,1}^R, \dots, p_{i,q_i}^R \rangle \in \mathbb{N}^{q_i}$ where $p_{i,j}^R$ 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. We propose 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 we can automatically adapt an EI to the global behavior of an agent population, then, we can repeat it for a number of different agent populations and thus characterize both δ and γ .

Figure 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 $\langle I_1, \dots, I_k \rangle$, where each individual represents the set of norm and role parameters defined above $\{ \langle p_{1,1}^N, \dots, p_{1,m_1}^N \rangle, \dots, \langle p_{n,1}^N, \dots, p_{n,m_n}^N \rangle, \langle p_{1,1}^R, \dots, p_{1,q_1}^R \rangle, \dots, \langle p_{t,1}^R, \dots, p_{t,q_t}^R \rangle \}$. 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. Since we are working with a complex system, we propose use an evolutionary approach for learning due to the fact that the institutional objective function can be naturally mapped to the fitness function and an evolutionary approach provides a solution good enough. Notice that the AEI does not learn any agent parameter, it learns the best parameters by simulation for a certain population of agents, that is whose values will be changed by the normative transition function and by the PS transition function. It is a first step learning where the AEI learns

by simulation the best parameters for a list of populations, thus, in a next step the AEI could use this learning in a real environment to adapt itself to any population of agents (e.g., using Case-Based Reasoning (CBR) problem solving technique).

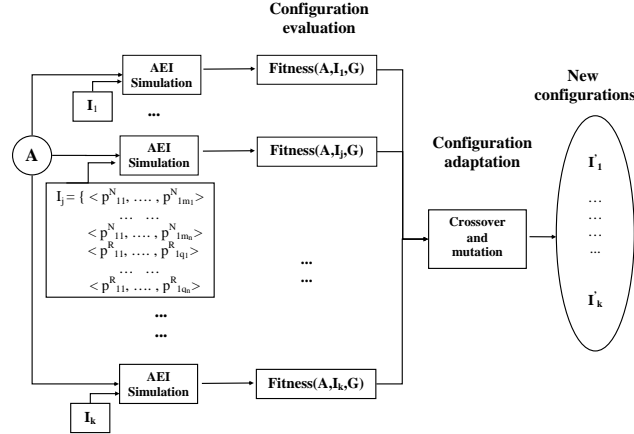


Fig. 1. Example of a step in EI adaptation using an evolutionary approach.

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 [9] to traffic flow improvement by means of multi-agent systems [10]. We tackle this problem from the Electronic Institutions point of view, and therefore, this section is devoted to specify how traffic control can be mapped into Autonomous Electronic Institutions.

In this manner, we consider the Traffic Regulation Authority as an Autonomous Electronic Institution, and cars moving along the road network as external agents interacting inside a traffic scene through driving actions. Additionally, indirect communication is established by means of stop, rear and turn signal indicators. 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. On the other hand, institutional agents in the traffic scene represent Traffic Authority employees. In our case study, we assume institutional agents to be in charge of detecting norm violations so that we will refer to them as police agents. Each police agent is 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 directly affects the number of detected norm violations, and thus, the overall quantity of penalties imposed to car agents. Furthermore, our Electronic Institution is able to adapt both norms and the number of deployed police agents based on its goals

– just as traffic authorities do modify them– 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 if its change leads 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.

Getting into more detail, we focus 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 dead-lock 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. In previous work [11] we introduce some nomenclature definitions about the environment.

4.1 AEI specification

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.

We define this grid environment as:

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

being x and y the cell position, α defines the set of external agents inside the grid cell (x, y) (notice that $\alpha \subseteq A$), 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.

Agents We consider $A = \langle A_1, \dots, A_n \rangle$ to be a finite set of n external agents in the institution. As mentioned before, external agents correspond to cars that move inside the grid environment, with the restriction that they can only move within road cells. Additionally, external agents are given an account of points which decreases with traffic offenses. The institution forbids external agents to drive without points in their accounts. The institution can observe the $P_a = \langle a_1, \dots, a_n \rangle$ 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 , whether the car is involved in an accident, the distance between the car and the car ahead of it; and, finally, external agent A_i 's point account. Notice that the institution can observe the external agent properties along time, we use a_i^t to refer the external agent A_i 's properties at a specific time t .

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 as well as the institution. Thus, these reference values are computed as a compound of other observed values. Considering our road junction case study, we identify different reference values:

$$V = \langle col, crash, off, block, expel, police \rangle$$

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

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

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

$$f(e_{\alpha^t}) = \begin{cases} 1 & \text{if } |e_{\alpha^t}| > 1 \\ 0 & \text{otherwise} \end{cases}$$

Similarly, considering last t_w ticks ($0 \leq t_w \leq t_{now}$): off indicates the total number of offenses accumulated by all external agents; $crash$ counts the number of cars involved in accidents; $block$ describes how many cars have been blocked by other cars; $expel$ indicates the number of cars that have been expelled out of the environment due to running out of points; and finally, $police$ indicates the percentage of police agents that the institution deploys in order to control the traffic environment.

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 and by specifying how many police agents should be deployed on traffic scene.

Institutional goals are defined as constraints upon a combination of reference values. Considering our scenario, we define restrictions as intervals of acceptable values for the previous defined reference values (V) so that we consider the institution accomplishes 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, the number of blocked cars, the number of cars that are expelled from the traffic scene, as well as the percentage of deployed police agents. In order to do it, we establish the list of institutional goals G as:

$$G = \langle g(col) \in [0, maxCol], g(off) \in [0, maxOff], g(crash) \in [0, maxCrash], \\ g(block) \in [0, maxBlock], g(expel) \in [0, maxExpel], g(police) \in [0, maxPolice] \rangle$$

Having more than one institutional goal requires to combine them. We propose an *objective function* [12] that favors high goal satisfaction while penalizing big differences among them:

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

where $1 \leq i \leq |G|$, $w_i \geq 0$ 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:

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

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 external agent population behavior.

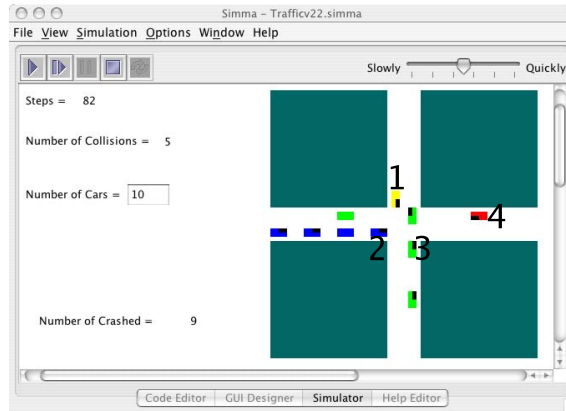


Fig. 2. Priority to give way to the right (Simma tool screenshot).

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 figure 2 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”.

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, it corresponds to the formula: $(x_i^t - h_{iy}^t + h_{ix}^t J_x, y_i^t + h_{ix}^t + h_{iy}^t J_y) == (x_j^t, y_j^t)$.

Similarly, we define an additional norm that is somehow related to the previous ‘right priority norm’. We 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 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 1 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 $(h_{ix}^t h_{jx}^t + h_{iy}^t h_{jy}^t) == -1$.

Table 1. Right and Front priority norms.

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

Performative Structure As introduced in 1, 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 two possible agent roles. On one hand, there is an institutional role played by police agents, whereas, on the other hand, the external role is played by car agents. Notice also that it is possible to specify the number of agents than can play each role within a scene.

4.2 Experimental Settings and Design

As a proof of concept of our proposal in section 3, we have designed an experimental setting that implements the traffic case study. In this preliminary experiment we consider four institutional goals related to *col*, *off*, *expel*, and *police* reference values; and both right and front priority norms in table 1. Institutional goals are combined with the objective function introduced in section 4.1, assuming corresponding weights are 4/10, 4/10, 1/10, 1/10 so that the first two goals are considered to be most important. On the other hand, norms are parameterized through its fines (i.e., points to subtract to the car failing to follow the corresponding norm).

The 2-road junction traffic model has been developed with Simma [13], a graphical MAS simulation tool shown in Figure 2, 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 external agent properties for each tick and can keep a record of them in order to refer to past ticks. Institutional police agents determine traffic offenses by analyzing a portion of car actions along time. External 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). Furthermore, 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 external agents without learning skills. They just move based on their trajectories, institutional norms and the percentage of deployed agents on the traffic scene. Car 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 applied if a police agent sees the infraction. Car agents decide whether to comply with a norm based on four parameters: $\langle \textit{fulfill_prob}$, $\textit{high_punishment}$, $\textit{inc_prob}$, $\textit{police} \rangle$. Being $\textit{fulfill_prob} \in [0, 1]$ the probability of complying with norms that is initially assigned to each agent, $\textit{high_punishment} \in \mathbb{N}$ the fine threshold that causes an agent to consider a fine to be high enough to reconsider the norm compliance, $\textit{inc_prob} \in [0, 1]$ the probability increment that is added to $\textit{fulfill_prob}$ when the fine threshold is surpassed by the norm being violated, and $\textit{police} \in [0, 1]$ the percentage (between 0 and 1) of police agents that the traffic authority has deployed to the traffic environment. 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:

$$\textit{final_prob} = \begin{cases} \textit{police} \cdot \textit{fulfill_prob} & \textit{fine} \leq \textit{high_punishment} \\ \textit{police} \cdot (\textit{fulfill_prob} + \textit{inc_prob}) & \textit{fine} > \textit{high_punishment} \end{cases}$$

Our goal is to adapt the institution to agent behaviors by applying Genetic Algorithms (GA)³ to accomplish institutional goals, that is, to maximize the objective function, which comprises the number of collisions, the number of offenses, the number of expelled cars and the percentage of police agents that should be deployed to control the traffic environment. We shall notice, though, that these offences do not refer to offences detected by police agents but to the real offences that have been actually done by car agents.

As section 3 describes, we propose to adapt the institution to different external agent population behaviors by running a genetic algorithm for each population. Therefore, institution adaptation is implemented as a learning process of the “best” institution parameters. In our experiments, Genetic Algorithms run 50 generations of 20 individuals. An individual corresponds to a list of a binary codifications of specific values for the following institution parameters: right norm penalty, front norm penalty, and percentage of police agents. Crossover among individuals is chosen to be singlepoint and a mutation rate of 10% 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):

$$O(V) = \frac{4}{10} \cdot \sqrt{f(g(col), [0, maxCol], \frac{1}{2})} + \frac{4}{10} \cdot \sqrt{f(g(off), [0, maxOff], \frac{1}{2})} + \\ \frac{1}{10} \cdot \sqrt{f(g(expel), [0, maxExpel], \frac{3}{4})} + \frac{1}{10} \cdot \sqrt{f(g(police), [0, maxPolice], 0)}$$

where $g(col)$, $g(off)$, $g(expel)$ and $g(police)$ correspond to average values of each reference value averaged for 5 different simulations; and $f(x, [m, M], \mu) \in [0, 1]$ represents the goal satisfaction.

5 Results

From the experimental settings specified above, we have run experiments for five different agent populations. These populations are characterized by their norm compliance parameters, being *fulfill_prob* = 0.5 and *inc_prob* = 0.2 for the five of them whereas *high_punishment* varies from 5 for the first, to 8 for the second, to 10 for the third, to 12 for the fourth, up to 14 for the fifth (see table 2).

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 goal of minimizing the total number of collisions. Nevertheless, it is also important for the AEI to reduce the total number of offenses, as well as, to a lesser extent, the number of expelled cars and the police deployment percentage. Each institutional agent has an associated cost, so that the AEI pursues the success of the traffic environment (i.e., a few collisions, agents respecting traffic norms and not having many expelled agents) at minimum cost. Thus, the AEI must learn what is the minimum percentage of police agents that should be deployed to control the traffic environment.

³ We use a genetic algorithm Toolbox [14].

Table 2. Agent populations.

Parameters	population1	population2	population3	population4	population5
<i>fulfill_prob</i>	0.5	0.5	0.5	0.5	0.5
<i>high_punishment</i>	5	8	10	12	14
<i>inc_prob</i>	0.2	0.2	0.2	0.2	0.2

Learning AEI parameters is a rather complex task because individual AEI’s goals are interrelated and can generate conflicts (for example, increasing police helps with collisions but raises costs). Furthermore, goals are related to agents’ behaviors. As explained before, agent’s behavior is so that its probability of complying with norms is proportional to the percentage of police, and therefore, since norms contribute to reduce accidents, collisions increase when police decrease. Moreover, the more percentage of police is deployed, the more number of fines are applied, and thus, the higher number of cars are expelled. Nevertheless, agents generate less offences when the police percentage increases. Additionally, the number of expelled cars decreases proportionally, not only to the police percentage, but also to the amount of applied fines. On the other hand, it may be worth recalling that agent’s behavior also increases its probability of complying with norms when the fine is larger than *high_punishment*. Therefore, any fine value higher than the population’s *high_punishment* value will have the same effect, and thus, will generate equivalent individual goal satisfaction degrees. As a result, the AEI must learn the best combination of parameters (*fine_{right}*, *fine_{front}* and *police*) according to the 4-goal objective function and to the agents’ behavior.

When learning, we have repeated tests for each setting three times –i.e., three separated learning runs for each agent population and setting–. Table 3 shows the learned parameters, where columns Learned *fine_{right}*, Learned *fine_{front}*, and Learned *police* include the learned values for each corresponding parameter and agent population. Each cell in the table contains three values: one per repeated experiment. Thus, for example, considering population1, learned values for *fine_{right}* are 15 for the first test, 12 for the second one and 7 for the third test. Notice that, due to the agent’s behavior, any fine value higher than 5 (*high_punishment* value) will have the same effect. Table 3 also shows the goal satisfaction value obtained for each test (this value corresponds to the objective function value explained above ($O(V)$) using $maxCol = 150$, $maxOff = 200$, $maxExpel = 200$ and $maxPolice = 1$).

As it can be seen, learned fines are larger than the population’s *high_punishment* value except for the third test in population5 (where the GA fails to find the maximum). Therefore, the institutional goals are successfully reached in fourteen of the fifteen tests. In this manner, we can rather state the AEI succeeds in learning the norms that better accomplish its goals. Relating to the police percentage, learned values are close to the 90%. This is due to its low associated weight in the objective function. Notice that the objective function is weighted in such a way that goals aiming to decrease the number of collision and offenses are considered to be significantly more important than those that pursue to decrease the number of expelled cars and the police percentage.

Table 3. Learning results for five different agent populations.

Population	Learned $fine_{right}$	Learned $fine_{front}$	Learned $police$	Goal satisfaction
population1	15, 12, 7	8, 14, 13	0.93, 0.93, 0.93	0.699, 0.7, 0.691
population2	13, 13, 14	10, 11, 9	0.93, 0.93, 0.93	0.689, 0.694, 0.691
population3	15, 12, 15	14, 11, 15	0.93, 0.87, 0.93	0.685, 0.681, 0.685
population4	15, 13, 15	14, 13, 13	0.93, 0.93, 0.87	0.676, 0.686, 0.68
population5	15, 15, 15	15, 15, 8	0.93, 0.93, 0.93	0.668, 0.674, 0.677

For the seek of clarity, figure 3 shows the overall goal functions for population1, population3 and population5 respectively. These 3D charts depict all the values of the goal function using only two parameters ($fine_{right}$ and $police$ ⁴ but not $fine_{front}$), so that search space for the learning algorithm is kept 3 dimensional. The domain for each is $16 \times 16 \times 1$. The figure shows the dependency between both parameters: when the police percentage is 100% the effect of the norm fine ($fine_{right} > high_punishment$) is greater than for smaller values of $police$, and becomes null when the police percentage goes down to 0%.

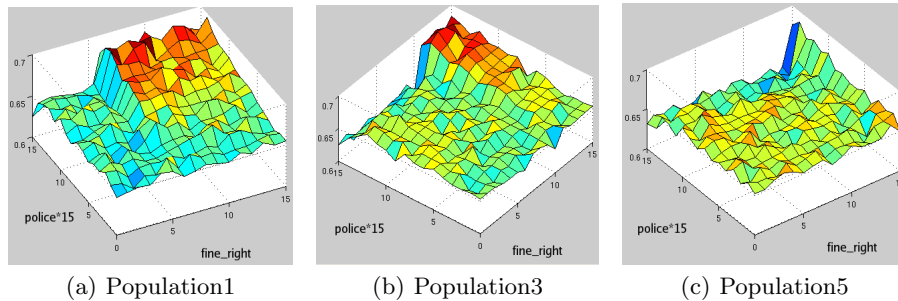


Fig. 3. Objective functions with 4 goals: col,off,expel,police. (a)Population1 ($high_punishment = 5$), (b)Population3 ($high_punishment = 10$), (c)Population5 ($high_punishment = 14$).

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 [15] propose a decision making framework that enables agents to dynamically select the coordination mechanism that is most appropriate to their circumstances. In [16] Gasser and Ishida presented a general distributed problem-solving model which can reorganize its architecture, in [17] Ishida and Yokoo introduce two new reorganization primitives that change the population of agents and the distribution of knowledge in an organization; and Horling et al. [18] propose an approach where the members adapt their own organizational structures at runtime. On the other hand, it has been long stated [19] 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. [20], where

⁴ Notice that the parameter $police$ is scaling to 15.

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. Therefore, we do not select norms at design stages as it is done by Fitoussi and Tennenholtz [21], 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. Our approach does not use regulative norms to represent agent goals, in contrast to Boella et al. [22], since we do not focus on modeling agents but the institution itself. For this same reason, agents are not allowed to change norms, which it is done through the institutional adaptation. Although the idea of autonomous normative system and goals of a normative system has been promoted before, [22], our proposal differs from it since we focus on the extension of the current notion of Electronic Institutions as an autonomous system (AEI) that could adapt to societal changes by changing the norms and the performative structure of the institution.

Regarding the traffic domain, MAS has been previously applied to it [10] [23], [24]. But traffic has been also widely studied outside the scope of MAS, for example, the preliminary work by [25] 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.

Preliminary results in this paper provide soundness to our AEI approach. We plan to perform the same experiments with other norms and with more goals. Nevertheless, we plan to extend both our traffic model and the institutional adaptation capabilities so that the AEI 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.

References

1. Luck, M., McBurney, P., Shehory, O., Willmott, S.: Agentlink Roadmap. Agentlink.org (2005)
2. Kephart, J.O., Chess, D.M.: The vision of autonomic computing. *IEEE Computer* **36** (2003) 41–50
3. Esteva, M.: Electronic Institutions: from specification to development. IIIA PhD Monography. Vol. 19 (2003)
4. Jennings, N.R., Sycara, K., Wooldridge, M.: A roadmap of agent research and development. *Autonomous Agents and Multi-agent Systems* **1** (1998) 275–306
5. North, D.: *Institutions, Institutional Change and Economics Performance*. Cambridge U. P. (1990)

6. Arcos, J.L., Esteva, M., Noriega, P., Rodríguez-Aguilar, J.A., Sierra, C.: Engineering open environments with electronic institutions. *Engineering Applications of Artificial Intelligence* (2005) 191–204
7. Etzioni, A.: *Modern Organizations*. Prentice-Hall, Englewood Cliffs, NJ (1964)
8. Gelfond, M., Lifschitz, V.: Representing action and change by logic programs. *Journal of Logic Programming* **17** (1993) 301–321
9. Yang, Q.: *A Simulation Laboratory for Evaluation of Dynamic Traffic Management Systems*. PhD thesis, MIT (1997)
10. Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K.: Mason: A new multi-agent simulation toolkit. In: *Proceedings of the 2004 SwarmFest Workshop*. (2004) 8
11. Bou, E., López-Sánchez, M., Rodríguez-Aguilar, J.: Self-configuration in autonomic electronic institutions. COIN (Coordination, Organization, Institutions and Norms in Agent Systems) Workshop at ECAI Conference (to appear 2006)
12. Rodríguez-Aguilar, J., Reyes-Moro, A., López-Sánchez, M., Cerquides, J.: Enabling assisted strategic negotiations in actual-world procurement scenarios. *Electronic Commerce Research* (to appear 2006)
13. López-Sánchez, M., Noria, X., Rodríguez-Aguilar, J.A., Gilbert, N.: Multi-agent based simulation of news digital markets. *International Journal of Computer Science and Applications* **2** (2005) 7–14
14. Chipperfield, A., Fleming, P.: *The matlab genetic algorithm toolbox*. IEE Colloquium Applied Control Techniques Using MATLAB (1995) 10–14
15. Excelente-Toledo, C.B., Jennings, N.R.: The dynamic selection of coordination mechanisms. *Autonomous Agents and Multi-Agent Systems* **9** (2004) 55–85
16. Gasser, L., Ishida, T.: A dynamic organizational architecture for adaptive problem solving. In: *Proc. of AAAI-91, Anaheim, CA* (1991) 185–190
17. Ishida, T., Yokoo, M.: Organization self-design of distributed production systems. *IEEE Trans. Knowl. Data Eng.* **4** (1992) 123–134
18. Horling, B., Benyo, B., Lesser, V.: Using Self-Diagnosis to Adapt Organizational Structures. *Proceedings of the 5th International Conference on Autonomous Agents* (2001) 529–536
19. Conte, R., Falcone, R., Sartor, G.: Agents and norms: How to fill the gap? *Artificial Intelligence and Law* (1999) 1–15
20. López-López, F., Luck, M., d’Inverno, M.: Constraining autonomy through norms. In: *AAMAS ’02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems, New York, NY, USA, ACM Press* (2002) 674–681
21. Fitoussi, D., Tennenholtz, M.: Choosing social laws for multi-agent systems: Minimality and simplicity. *Artificial Intelligence* **119** (2000) 61–101
22. Boella, G., van der Torre, L.: Regulative and constitutive norms in normative multiagent systems. In: *Proc. of Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR’04), Whistler (CA)*. (2004)
23. Dresner, K., Stone, P.: Multiagent traffic management: An improved intersection control mechanism. In: *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, New York, NY, ACM Press* (2005) 471–477
24. Doniec, A., Espié, S., Mandiau, R., Piechowiak, S.: Dealing with multi-agent coordination by anticipation: Application to the traffic simulation at junctions. In: *EUMAS*. (2005) 478–479
25. Montana, D.J., Czerwinski, S.: Evolving control laws for a network of traffic signals. In: *Genetic Programming 1996: Proceedings of the First Annual Conference, Stanford University, CA, USA, MIT Press* (1996) 333–338