

Self-Organizing Multi-Agent Systems for the Control of Complex Systems

Jérémy Boes*, Frédéric Migeon

*IRIT, University of Toulouse
118, route de Narbonne
F-31062 Toulouse Cedex 9, France*

Abstract

Because of the law of requisite variety, designing a controller for complex systems implies designing a complex system. In software engineering, usual top-down approaches become inadequate to design such systems. The Adaptive Multi-Agent Systems (AMAS) approach relies on the cooperative self-organization of autonomous micro-level agents to tackle macro-level complexity. This bottom-up approach provides adaptive, scalable, and robust systems. This paper presents a complex system controller that has been designed following this approach, and shows results obtained with the automatic tuning of a real internal combustion engine.

Keywords: Multi-Agent Systems, Control, Self-Organization, Complex Systems, Internal Combustion Engines

1. Introduction

Controlling a system means being able to perform the adequate modifications on its inputs in order to set the outputs on a desired state. Over the course of History, humans made tremendous efforts to control systems that are more
5 and more complex: nonlinear, dynamic, noisy, with a large number of inputs and outputs, and so on. Yet, the law of requisite variety [1] implies that the

*Corresponding author

Email addresses: boes@irit.fr (Jérémy Boes), migeon@irit.fr (Frédéric Migeon)

complexity of a controller has to be greater than or equal to the complexity of the target system. Thus, the design of a controller involves the design of a complex system. This is a challenge for engineering.

10 Complexity is often tackled a posteriori, to study existing systems. On the contrary, methods enabling the design of complex systems that meet strict requirements are quite rare. The main feature of a complex system is that its behavior can not be easily predicted [2]. Usual design methods, for instance in software engineering, seek to a priori eliminate any unexpected event. The
15 design process must ensure that everything will be smooth at runtime. But, as any other complex system, complex programs sometimes have unexpected, unpredictable behaviors, and these classical methods fail.

For instance, in the field of system control, the usual methods in the industry rely on the construction of a fine mathematical model of the target system, that
20 is later used to compute the commands to perform, given some setpoints. The cost and difficulty of the construction (and the tuning) of a mathematical model is high. An often used alternative is machine learning. Giving the ability to learn to a controller enables it to learn the behavior of the target system and build a model from data. However, this method shows its limits when used
25 with complex systems. Nonlinearities in the learnt model lead to overcostly or impossible computations in the control system. Another possibility exists: directly learning the adequate commands, instead of a model that will later lead to the said commands. We then focus only on the inputs and outputs of the controlled system, without trying to decipher its internal mechanisms.

30 Another difficulty is scalability. While various control methods exist, they (almost) all fail to scale when a large number of inputs and outputs are involved. Most advanced solutions rely on the distribution of the control. Instead of letting a central controller handle all the inputs, each input is controlled by one local controller, and all controllers try to cooperate to control the whole system.

35 Multi-Agent Systems (MASs), composed of autonomous entities, are naturally distributed. They can be very useful to the problem of the control of complex systems, for instance with multi-objective optimization[3]. Moreover, they

bring innovative design methods. In particular, Adaptive Multi-Agent Systems (AMASs) are designed to be able to self-adapt at runtime to any unexpected
40 event. Instead of wasting time trying to cope with any possible event during the design phase, we let the system deal with the unexpected at runtime. Driven by cooperation principles, agents self-organize locally to produce and maintain the desired global function.

This paper presents experimental results obtained with an AMAS designed
45 to control complex systems, and applied to the calibration of real heat engines. This system is fully described in English for the first time in this paper. Able to learn and control simultaneously, it provides a generic and robust solution to the problem of control. It is a good example of the ability of AMASs to be efficient in real life conditions.

50 Section 2 gives a quick background on control. Section 3 introduces our approach and section 4 presents our system. Results, obtained in simulated as well as in real conditions, are shown in section 5. Section 6 concludes with our perspectives.

2. Related Works

55 Our work is at the crossroad of the fields of complex systems, control, and machine learning. It is inspired by the ideas of Edgar Morin on complexity [4], which we apply here to the design of self-adaptive control systems.

2.1. Complex Systems

The notion of complexity reflects the difficulty to analyze a system and
60 to forecast its behavior. Nonlinearities, inner feedback loops, large number of inputs/outputs/inner parts, uncertainty on the measures, and unpredictable behaviors are some of the recurring features of complex systems. However, there is no common agreement on a definition. For instance, Kolmogorov defines the complexity of a string as the length of the shortest description of said string [5].
65 While it is largely accepted, this measure implies that a purely random string

is of maximal complexity, as it can only be described by its full enumeration. However this contradicts one of the key features of complexity: it is situated somewhere between total order and total chaos [2]. Moreover, a complex system is dynamic, it is able to spontaneously change its state. It is important not to neglect this aspect during the analysis or the design of a system. Measures such as Kolmogorov complexity give too much attention to static, structural features of systems, and not enough to their dynamics. To this end, *dynamical depth* is based on the idea that the degree of complexity of a system is not given by its part and their causal relations, but by the imbrication of the different dynamics that drive its behavior [6].

Furthermore, the general system theory states that the classical analytical approach can only be applied on systems whose parts are linear and share negligible interactions [7]. This lets a lot of systems out of its scope, in particular complex systems. We need to follow a different approach than the reductionist top-down analysis for complex systems control as well as for complex systems design. The Adaptive Multi-Agent Systems theory is being developed in this regard.

2.2. Control

Control approaches also find their limits when faced with complexity. Artificial Intelligence (AI), and in particular machine learning, are used to overcome these limits.

The objective with AI in control is to automatically learn either the model of the target system, the tuning of the model, the calibration of the controller, or directly control laws from observations. For instance, [8] uses a genetic algorithm to learn the optimal tuning of PIDs. This approach gives excellent results but needs a large number of iterations. Moreover, if the behavior of the controlled system changes over time (for instance, because of mechanical wear), the tuning must be entirely redone, it is not adaptive.

The biggest difficulty of dual control is to find the correct balance between probe actions and control actions. A way to do this is to use neural networks

to learn this balance from data [9]. This approach is limited to control affine systems, i.e. systems that reacts linearly to modifications on their inputs.

The most promising approach for scaling up, i.e. for controlling a large number of inputs with many criteria on many outputs, is to distribute the control. For instance, [10] and [11] control road traffic junction signals on several crossroads. In these approaches, there is no central controller that handles all the traffic junctions, each crossroad is controlled by a local controller. [10] uses learning classifier systems, while [11] uses a combination of neural networks, genetic algorithms and fuzzy logic. They obtained very interesting results, but the difficulty to instanciate their approaches to real life problems is a severe drawback.

Our approach uses feedback loops to learn not the model of the controlled system but the control laws themselves, and distributes a controller on each controlled input. Inner feedback loops ensure an adequate balance between exploration and exploitation of the model.

2.3. Machine Learning

A program learns when it is able to improve its fonctionnality using its experience, i.e. data acquired during its execution [12]. Machine learning has been heavily influence by the way we think the human mind works. The two well-known methods for machine learning are supervised learning and unsupervised learning, whether examples of the expected results are presented to the learning program or not. However, this distinction is merely technical and does not allow to highlight the fundamental differences between machine learning algorithms. We prefer the following five categories: Behaviourism, Cognitivism, Connectivism, Evolutionism, and Constructivism.

Behaviourism considers the learner as a black-box. Learning occurs when the observed behaviour changes in response to the dynamics of the environment. In machine learning, the behaviour is then a product of the initial state of the program and its progressive conditioning by its environment through a feedback loop. Reinforcement learning can be considered as a behaviourist machine

learning approach. It is notably popular in robotics [13]. Its most well-known algorithm is Q-Learning [14].

On the contrary, cognitivists consider that what is important is not what the learner does but what he knows. Cognitivist machine learning algorithms classically relies on symbol manipulation, and thus on a predefined set of symbols, which is not adequate when dealing with complexity [15].

Connectionism considers learning at a lower level in the brain: the dynamic interconnection of neurons. In machine learning, it regroups all the artificial neural network algorithms, from back-propagation perceptrons to the more recent Kohonen maps [16] and deep learning algorithms [17]. They show impressive results but need a huge amount of data and computing power.

Evolutionism considers learning at the scale of a species rather than an individual. Evolutionary algorithms evolve a population of solutions towards better solutions by evaluating them, mutating them, and crossing the best individuals. These algorithms are interesting because they can tackle problems for which there is no known solution, but they are time-consuming and the fitness function can be difficult to obtain [18].

Constructivism is the idea that humans have the ability to construct knowledge in their own mind through interactions with the environment. Constructivist artificial intelligence aims at designing self-constructive systems [19]. In such systems, not only the knowledge but the means to acquire it are learned. The focus is made on self-organization and bottom design methods.

Note that there are no hard boundaries between these categories. Most advanced machine learning algorithms actually take simultaneously from several of them. For instance Learning Classifier Systems stems from Behaviourism since they are reinforcement learning algorithms, but they also incorporate an evolutionary component [20].

The Adaptive Multi-Agent Systems approach is constructivist: it focuses on self-organization and shares the same long term goal of designing a fully self-constructed artificial intelligence. It also has a link with connectionism with the idea that a complex task can be achieved by a set of several simple entities.

3. Approach

Top-down classical methods have severe shortcomings when it comes to complexity: scale, integration, and flexibility [19]. This section presents the
160 Adaptive Multi-Agent Systems (AMASs) theory, that aims at overcoming these limitations thanks to the natural modularity of MASs and cooperative self-organization of agents.

3.1. Adaptive Multi-Agent Systems

Wooldridge defines an agent as follows: *An agent is a computer system that*
165 *is situated in some environment, and that is capable of autonomous action in this environment in order to meet its delegated objectives.* [21] "Autonomous action" means an agent takes its own decision on what to do and when to do it. It indefinitely follows a lifecycle of perception, decision and action without any external control.

170 A system composed of several agents in interaction in the same environment is called a Multi-Agent System (MAS) [22]. Knowledge, computation, and control are distributed among the agents of a MAS. Such systems are based on collective problem solving, the idea that local behaviors within a group can ensure the achievement of a given global task. Multi-agent systems provides in-
175 teresting features when dealing with complexity, such as scalability, robustness and adaptivity [23].

The function of a MAS is dependent on its organization (the agents, their relations, their behavior). A change in the organization of the MAS is a change of its global function. When agents decide themselves to dynamically change their
180 behavior or their relations, the system is self-organizing. Di Marzo Serugendo et al. define self-organization as the process with which a system changes its structure without any external control to respond to changes in its operating conditions and its environment [24]. It is very natural and powerful for a MAS to perform learning and self-adaptation this way.

185 The Adaptive Multi-Agent Systems approach aims at facilitating the design of multi-agent systems for solving complex problems by designing simple agents

that self-organize to generate a complex global function [25]. In this approach, the process of self-organization is driven by cooperation principles. Local decisions from each agent may provoke local changes that in turn lead to changes
190 in the global function of the system.

This approach is based on the theorem of functional adequacy [25]. Applied to MASs, one of the consequences of this theorem is the assurance that the global function of a system is adequate if all agents maintain interactions with their environment that are favorable to themselves and to their environment
195 (they are said to be in a *cooperative state*). Then, the challenge is to find the behavior for each agent that enables each of them to remain in a cooperative state despite changes in their environment.

To this end, each agent has two sets of rules. *Nominal* rules enable an agent to achieve its function when it is already in a cooperative state. However, it
200 is highly probable that the agent eventually finds itself unable to achieve its function, due to changes in its environment, or to a simple lack of knowledge. Such cases are called *Non-Cooperative Situations* (NCSs) and are probable cause of failure for the global system to achieve its task. There are seven types of NCSs:

- 205 • Incomprehension: the agent is not able to extract information from the perceived signal.
- Ambiguity: the agent can interpret the perceived signal in several different manners.
- Incompetence: the agent is not able to decide anything based on its current
210 knowledge and skills.
- Unproductiveness: the decision of an agent is to do nothing.
- Concurrence: the agent thinks its action will have the same effects as the action of another agent.
- Conflict: the agent thinks its action is discordant regarding the action of
215 another agent.

- Uselessness: the agent thinks that its action will have no consequences on its environment.

When a NCS occurs, the involved agents switch from their nominal behavior rules to their *cooperative* rules, which seek to solve the NCS by provoking changes in the MAS (in other words, by triggering self-organization). An agent has several means to solve a NCS: tuning internal parameters, reorganizing its relations with other agents, creating of a new agent, or self-destructing.

In the current state of the approach, the AMAS designer has to design the cooperative behavior for each NCS. A methodology named ADELFE (French acronym for Toolkit for Developing Software with Emergent Functionalities) guides the design of AMASs [26]. It is a bottom-up and iterative design process that encourages the designer to focus on the local function of each agent, and to forget the global function of the system. A strong focus is put on decomposing the problem instead of the solution. The resulting agents will often be following simple (yet intricated) reactive behavioral rules, and thus will seem too simple to solve anything. It is the point of our approach: dodging complexity by thinking exclusively within a local scope. If agents behave accordingly to the AMAS principles of cooperation, the emerging global function shall be adequate. Originally based on the Rational Unified Process [27], ADELFE incorporates specific steps and guidelines to help identify the entities of the problem and which ones should become agents, and find their NCSs and their cooperative behaviors. It has been used for the development of the system presented in 4.

3.2. Objectives in Terms of Control

Other than pushing forward the experimental verification of the AMAS approach, the main objective of this work is to design a system able to learn in real time how to put a complex system in a desired state. In our case, the controlled system may have multiple inputs and outputs (MIMO), and the desired state is described as a combination of criteria. A criterion may affect one or several inputs or outputs. There are three types of criteria:

- 245 • Constraint: a threshold to meet
- Setpoint: a target value
- Optimization: a value to minimize or to maximize

An additional requirement is that the controller must be easy to implement for real-life complex systems. In particular, this means the controller should not
250 need a heavy tuning and should not require any predefined model. In other words, prerequisite knowledge on the controlled system has to be minimal.

Moreover, the learning process has to be perpetual and in real time. It has to occur simultaneously to the control, so the controller adapts itself to changes in the controlled system (such as failures, wear, etc). Our controller sees the
255 controlled system as a black box: it only has access to the inputs and the outputs of the black box, not to the internal processes that drives its behavior.

4. **ESCHER, an Adaptive MAS to Learn the Control of Complex Systems**

In this section we present a multi-agent system called **ESCHER**, for Emergent Self-adaptive Controller for Heat Engine calibration. Thanks to cooperative self-organization, it is able to learn in real-time the control of a system. It has been designed and tested during a project revolving around automotive thermal engines, but has been design under the assumption that nothing is known about the controlled system, except its number of inputs and outputs.

265 The goal is to make the controller generic enough to be used on any other systems without any modifications other than the interface. Following a "black box" approach of the control, **ESCHER** plays with the inputs of the controlled system, observes the effects on the outputs and infers the actions that will lead to compliance with the user-defined criteria.

270 *4.1. System Overview*

The environment of **ESCHER** is composed of the controlled system and of the user defined criteria. This means that **ESCHER** observes the inputs and

outputs of the controlled system, and also the control criteria defined by the users. Among the inputs of the controlled system, there may be some that are
275 not controlled by ESCHER but have an impact on the controlled system. For instance the atmospheric pressure cannot be controlled but can significantly alter the output of a thermal engine. If such a sensor is available, it can be taken into account by ESCHER.

ESCHER itself is composed of four types of agents:

- 280 • Variable Agents are the eyes of the system, there is one Variable Agent for each input and output of the controlled system.
- Criterion Agents represent user-defined criteria, the desired state of the controlled system.
- 285 • Context Agents can be seen as the memory of the system, they represent a part of the state space of the environment for which the consequences of a given action are known.
- Controller Agents are the hands of the system, they interact with a set of Context Agents to find the most adequate action to perform in the environment.

290 Figure 1 shows an overview of the system, with the links between the four types of agents. Note that this view is intended for the reader, agents do not have a global view of the system.

4.1.1. Context Agents and Controller Agents

Each Controller Agent is coupled with a set of Context Agents whose mem-
295 orized action is related to the effector associated to this same Controller Agent. The Controller Agent selects the next action to perform among the received suggestions and notifies the Context Agents which has sent a suggestion. There is no direct interaction between Context Agents, neither between Controller Agents. The only link between them is through the environment: the action of
300 a Controller Agent will have an impact on the controlled system which will be




Figure 1: A View of all the Agents of ESCHER

perceived from other Controller Agents through Variable Agents and Criterion Agents.

A Controller Agent and its set of Context Agents can be seen as an autonomous MAS. Its environment would be made of Variable Agents and Criterion
 305 Agents. A Context-Controller "sub-MAS" is able to synchronize its actions with the other sub-MASs by observing the controlled system's inputs and outputs variations. A Controller Agent does its best to decrease the critical levels by performing actions on only one input, locally, without caring about how the other inputs are handled. There is no global decision process to find the ad-
 310 equate actions on each input at once. This feature is the key to scalability. Moreover, the distribution of control makes ESCHER modular. The addition or the removal of a new Controller Agent does not impact the others.

4.1.2. Variable Agents and Criteria Agents

To fulfill its function, each agent besides Variable Agents, needs to know the
 315 current state of the controlled system. This is why Variable Agents send value

update to every other types of agents (the relevant Criterion Agents, every Context Agent, every Controller Agent). This broadcast may seem harmful for scalability, but it is not. Indeed, agents of ESCHER are not physically distributed, the cost of message sending is very low. On the contrary, the
320 cost of reading the value of a physical sensor is high, since it involves external systems, and probably networking. Hence, it is way more efficient to have an agent per sensor, broadcasting its value to others than to give access to a sensor to every agent needing this particular value.

Criterion Agents transform the variable values into critical levels, representing the satisfaction of the criteria (i.e. a idea of how far from the desired state
325 is the current state of the controlled system). Variable Agents and Criterion Agents give ESCHER a complete representation of its environment.

At a given moment, if every agent in the system is able to properly perform its function, then ESCHER is in a cooperative state and its global function is
330 adequate. However, numerous cases exist where at least one of the agents is unable to execute its function. These cases are the Non-Cooperative Situations, that are presented in section 4.3.

4.2. Function and Nominal Behavior of ESCHER Agents

This section presents a decomposition of the activity of control in elementary
335 tasks. Agents in charge of these tasks are detailed.

4.2.1. Observing the Controlled System

The first thing we need when it comes to controlling a system with a "black box" standpoint is to be able to observe it. A specific type of agents is in charge of the perception of the controlled system: *Variable Agents*. To each input and
340 output of the system is associated a Variable Agent. During its lifecycle, a Variable Agents perceives the value of its designated variable on the controlled system and forwards it to the other agents which may need this information. If necessary, a Variable Agents may embed a noise filtering algorithm.




Figure 2: Examples of Criticality Functions

4.2.2. Representing Control Criteria

345 The controller needs to have an internal representation of the objectives of the user, of the desired state for the controlled system. Giving such a representation is the function of *Criterion Agents*. There are three types of Criterion Agents:

- Threshold: the agent expresses the will to maintain a variable either above
350 or below a user-defined threshold.
- Setpoint: the agent expresses the will to set a variable to a user-defined value.
- Optimization: the agent expresses the will to minimize or to maximize the value of a variable.

355 Each Criterion Agent receives updates from the relevant Variable Agents, computes a *critical level*, and sends it to other agents which may need this information. This critical level reflects the satisfaction of the criterion represented by the agent. The critical level ranges from zero (the criterion is fully satisfied) to 100 (the criterion is far from being satisfied).

360 Figure 2 shows examples of criticality functions used by Criterion Agents to compute their critical level. For instance the threshold criticality function returns zero if the threshold is met, otherwise a value up to 100. The criticality function for a setpoint returns zero only when the target value has been reached. The criticality function of an optimization criterion is asymptotic to zero. The

365 curves of these functions can be adjusted by the user to define the relative
significance of its needs.

Criterion Agents apply a transformation from the space of the controlled
system variables to the space of the criteria. The critical levels decrease when
their criterion is being satisfied. Hence, agents perceiving critical levels seek
370 to decrease them. The only way to do so is to perform adequate actions on
the input of the controlled system. Finding these adequate actions requires the
analysis of the current state of variables and criteria to try to understand the
dynamics of the system.

4.2.3. Analyzing the State of the Environment

375 With the Variable Agents and the Criteria Agents, ESCHER has an internal
distributed representation of its environment. To be able to decide which actions
to perform, an analysis of this environment is needed. This is the function of
Context Agents.

A Context Agent memorizes the effect, on each critical level, of a particular
380 action performed on a particular effector. The agent also memorizes the state of
the environment when the action is performed. This provides information about
the expected consequences of a particular action if the action is performed while
the environment is in a particular state.

Concretely, a Context Agent is composed of:

- 385 • an action, i.e. an offset to be performed on an input of the controlled
system,
- a set of forecasts, which contains a value for each Criteria Agent, repre-
senting the expected variations of critical level,
- a set of *validity ranges*, which contains a value range for each Variable
390 Agent, representing the state of the controlled system.

A Context Agent receives value updates from Variable Agents, and critical level
updates from Criterion Agents. When the current value of each Variable Agent

is inside their corresponding validity range, the Context Agent is said *valid*. This means the controlled system is in a state in which the forecasts of the agent are relevant. When a Context Agent becomes valid, it sends a notification which contains its action and its forecasts. This notification is actually an action suggestion. Let p a suggestion (1):

$$p := (a, F) \tag{1}$$

where a is an action and F is a set of critical levels forecasting functions. Thus, a function $f^i \in F$ returns the critical level of Criterion Agent i forecasted by the Context Agent if a is performed. Such a function can be expressed as (2):

$$f^i(a) = c^i + \delta^i(a) \tag{2}$$

where c_i is the current critical level of Criterion Agent i , and δ^i is a function resulting from the learning of the Context Agent. In practice, a Context Agent sends an action suggestion together with a set of values $f^i(a)$, not a set of computable functions f^i . We only show expression (2) to explicit a part of the
395 learning of Context Agents, which will be discussed later.

A notification is also sent when the Context Agent becomes non-valid so its suggestion is withdrawn. These suggestions and notifications are received by the *Controller Agent* in charge of the affected effector. This new type of agent is presented in the next paragraphs.

400 4.2.4. *Performing the Most Adequate Action*

A Controller Agent is associated to each input controlled by ESCHER. The function of a Controller Agent is to perform the most adequate action on this input, i.e. the action which will provoke the greatest decrease of critical level. An action may be increasing, decreasing, or maintaining the value of the input.

Let u_t the current value of the input controlled by the Controller Agent, and a_t the action performed by the Controller Agent at its lifecycle t . The next value of the input is given by equation 3.

$$u_{t+1} = u_t + a_t \tag{3}$$

405 At each lifecycle t , the Controller Agent chooses a_t according to its internal representations, which are composed of:

- \mathcal{C}_t , the set of critical levels, updated at lifecycle t .
- \mathcal{P}_t , the set of action suggestions from valid Context Agents at lifecycle t .

Among \mathcal{P}_t (the received suggestions), the Controller Agent chooses the action
 410 associated with the greatest decrease of the highest critical level. If the highest critical level is not expected to vary, according to the forecasts, then the Controller Agent seeks to decrease the second highest critical level, and so on.

Hence, for each suggestion $p_t^k \in \mathcal{P}_t$, the Controller Agent looks at $f_{max}^k \in F_t^k$, the function which returns the highest critical level (in other words, the function corresponding to the most critical Criterion Agent). This function is defined by equation 4.

$$f_{max}^k := f_t^k \in F_t^k, f_t^k(a_t^k) = \max_{f \in F_t^k} (f(a_t^k)) \quad (4)$$

The chosen a_t is the action from the suggestion with the lowest $f_{max}(a)$, while being lower to the current highest critical level (equation 5).

$$a_t := a^i \in \mathcal{A}_t, f_{max}^i(a^i) = \min_k (f_{max}^k(a^k)) \wedge f_{max}^i(a^i) \leq \max \mathcal{C}_t \quad (5)$$

where \mathcal{A}_t is the set of actions contained in the suggestions from \mathcal{P}_t .

The Controller Agent then performs the action a_t and sends:

- 415 • an acceptance notification to the currently valid Context Agents whose action has been selected and performed,
- a rejection notification to the currently valid Context Agents whose action has not been selected,
- in case of the current action is different from the action of the previous
 420 step, a waiver notification to the Context Agents which suggested the previous action.

Of course, at any given time, a Controller Agent may not be able to make a good decision (i.e. a decision that will lead to the decrease of critical levels),

because of false or incomplete information. These cases are Non-Cooperative Situations (NCSs). They occur when ESCHER has not sufficiently learned and is not fully adapted to its environment. For instance, if the condition 6 is not met, then equation 5 cannot be applied.

$$\exists p_t^i \in \mathcal{P}_t, \exists f_{max}^i \in F_t^i, f_{max}^i(a^i) \leq \max \mathcal{C}_t \quad (6)$$

The occurrence of a NCS triggers a specific behavior (the cooperative behavior) of the involved agents to solve it and set the agents in a cooperative state. Solving NCSs drives the whole system towards a state of functional adequacy.
 425 NCSs and their resolution are presented in section 4.3.

4.3. Non-Cooperative Situations

This section explains how agents detect and solve NCSs. Since they provoke changes in the organization of the system, NCSs and their resolution are the key to the self-adaptativeness of AMASs. Each agent locally solves the NCSs it
 430 detects, thanks to specific actions. In ESCHER, NCSs mainly occur for Context Agents and Controller Agents. They motivate the system to self-organize, in particular by creating, modifying, or deleting Context Agents.

4.3.1. NCS 1 : Controller Agent Incompetence

Detection: When a Controller Agent doesn't receive any action suggestion,
 435 $\mathcal{P}_t = \emptyset$, hence $\mathcal{A}_t = \emptyset$. In this situation, the agent is not able to choose an adequate action using equation 5: it finds itself in a NCS of incompetence.

Resolution: The resolution of this NCS has two steps. First, the Controller Agent has to choose an action on its own. Its choice is based on the effects of its previous action. If the critical levels are increasing, the new action is chosen as the opposite of the previous action, otherwise the previous action is repeated (equation 7).

$$a_t := \begin{cases} a_{t-1} & \text{if } \max \mathcal{C}_t < \max \mathcal{C}_{t-1} \\ -a_{t-1} & \text{otherwise} \end{cases} \quad (7)$$

If $t = 0$, then the new action is randomly chosen.

If the previous action had been selected from \mathcal{P}_{t-1} and is continued, the Controller Agent does not send a waiver notification to the Context Agents that had suggested it at $t-1$, even if they are now non-valid. They may need this information to learn (see NCS 6).

Otherwise, after having determined its new action, but before performing it, the Controller Agent creates a new Context Agent. This new Context Agent is initialized with the new action, and memorizes the current value of all variables. While the highest critical level decreases, the Controller Agent continues the same action. During this time, the new Context Agent observes the variations of all critical levels to set its forecasts. Finally, when the action is abandoned, the Context Agent sets its validity ranges with the minimum and maximum observed on each variable.

4.3.2. NCS 2: Controller Agent Unproductiveness

Detection: When none of the received action suggestions contains forecasts of a decrease of the highest critical level (condition 6 is not met), the Controller Agent is in a NCS of unproductiveness. Its nominal decision process (select the action associated to the biggest decrease of the highest critical level) does not produce any action. There are two ways of solving this NCS, depending on the received suggestions. Let \mathcal{A} the set of all possible actions for the Controller Agent, at each time step t we have $\mathcal{A}_t \subseteq \mathcal{A}$.

Resolution 1: If $\mathcal{A}_t = \mathcal{A}$, in other words if every type of actions (increment, decrement, stay) has been suggested, the Controller Agent thinks that the highest critical level can not be decreased, whatever the agent may do. Then, the agent attempts to decrease the second highest critical level (without increasing the highest critical level). If it is not possible, it will look at the third highest critical level, and so on. If there is no forecasted decrease at all, the agent chooses the least harm: the action associated with the smallest increase of the highest critical level is chosen (equation 8).

$$a_t := a^i \in \mathcal{A}_t, f_{max}^i(a^i) = \min_k(f_{max}^k(a^k)) \quad (8)$$

Resolution 2: The second case is when $\mathcal{A}_t \neq \emptyset \wedge \mathcal{A}_t \neq \mathcal{A}$. It means that some actions have not been suggested, they have not been tested in the current state of the environment. Since none of the received action suggestions contains forecasts of decrease of the highest critical level, they actually contain actions to avoid. Let $\mathcal{A}_c = \mathcal{A} - \mathcal{A}_t$ the set of candidate actions, i.e. actions that are not currently suggested. The Controller Agent then decides to select an action among the ones which are not suggested (which we call candidate actions). The selection of the new action is similar to the resolution of the NCS 1 but is, this time, conditioned by the presence of this action in \mathcal{A}_c (9).

$$\begin{cases} a_t = a_{t-1} & \text{if } a_{t-1} \in \mathcal{A}_c \wedge \max \mathcal{C}_t < \max \mathcal{C}_{t-1} \\ a_t = -a_{t-1} & \text{if } -a_{t-1} \in \mathcal{A}_c \wedge \max \mathcal{C}_t \geq \max \mathcal{C}_{t-1} \\ a_t = \text{rand}(\mathcal{A}_c) & \text{otherwise} \end{cases} \quad (9)$$

With the same conditions than in the NCS 1, the Controller Agent may create a new Context Agent, initialized in the same manner.

4.3.3. NCS 3: Controller Agent Conflict

Detection: When a Controller Agent applies an action suggested by a Context Agent, it expects that the critical levels will vary in the way indicated by the forecasts. If the Controller Agent notices that it is not the case, it thinks that the action that has just been performed may be harmful, it is a conflict NCS.

Resolution: The action must be stopped. The Controller Agent abandons the action and notifies the Context Agents which had suggested it when it was selected. Moreover, if the Context Agents which were wrong are still valid, they will be temporarily ignored in future step.

4.3.4. NCS 4: Context Agent Conflict (false forecasts)

Detection: When the action of a valid Context Agent is being performed, said agent observes the variations of critical levels. When the action is terminated, the agent compares the observed variations with its forecasts. There is a conflict

NCS if at least one of the observed variation contradicts the forecast (their direction of variation is different).

Resolution: An error in the direction of variation of a forecast is probably more
485 than a simple mistake in the initial observation, it is not a problem of forecast
adjustment. This rather indicates that the Context Agent should not have
sent its suggestion, it should not have been valid. To correct this situation, the
Context Agent will reduce its validity ranges, bringing closer the nearest bound
to the current value of the corresponding variable.

490 *4.3.5. NCS 5: Context Agent conflict (inaccurate forecasts)*

Detection: This NCS is similar to NCS 4. But this time, the observed vari-
ations are in the same direction as the forecasts, but not of the same amount.
This kind of observation is sensitive to noise on the perception of variable values,
hence small differences (under 5% of criticality) are ignored.

495 **Resolution:** An error in the amplitude of variation is less serious than an
error in the direction of variation. The Context Agent only needs to adjust its
forecast. Thus, in this case, the agent does not change its validity ranges, but
rather increase or decrease the erroneous forecasts so they fit its observations.

4.3.6. NCS 6: Context Agent Incompetence

500 **Detection:** It happens that a Context Agent whom action is being performed
becomes non-valid, but does not received any reject nor waiver notification from
the Controller Agent (it is a possible outcome of NCS 1). The Context Agent
is then in an incompetence NCS, this situation is not covered by its nominal
behavior.

505 **Resolution:** From its standpoint, this situation means that the Controller
Agent considered that its action can be kept a little longer. Hence, to keep
sending what could be a good suggestion, the Context Agent extends the validity
ranges that make him non-valid.

4.3.7. NCS 7: Context Agent Uselessness

510 **Detection:** Sometimes, after several NCS 4, some validity ranges of a Context Agent have been so shrunked that their amplitude is near zero. If the amplitude of at least one validity range falls under the threshold of minimal size, the Context Agent is in a uselessness NCS, the chances of being valid are too low. By default, the threshold is equal to one hundredth of the domain of the variable.
515 This NCS is ignored for unbounded variables.

Resolution: A useless Context Agent can do nothing else than delete itself to solve this situation. Indeed, a Context Agent can only learn if its action is selected while valid. If the agent is never valid, it never brings information to the system and never learns. By deleting itself, the agent frees computation
520 ressources. This NCS is not pivotal for ESCHER. The presence of useless agents does not prevent the adaptation and functional adequacy of the whole system. But this NCS avoids overages of Context Agents. To avoid that too many deletions and a loss of memory, we advise to set

4.3.8. NCS 8: Context Agent Unproductiveness (validity ranges)

525 **Detection:** This NCS concerns a Context Agent which has been valid, selected, then became non-valid, and observed a decrease of critical levels. This is an ideal cases, everything went fine. This is why a Context Agent in this situation considers that its action may still be relevant, even if the agent itself is now non-valid. This is an unproductiveness NCS: the nominal decision process results
530 in doing nothing (since the agent is not valid), while there is good chances that sending an action suggestion should be a good thing to do.

Resolution: The Context Agent expands the validity ranges that make it non-valid, so it is now valid. The agent also sends an action suggestion. If the agent was wrong to send a suggestion, a NCS 4 will occur and the ranges will
535 be shrunked. Likewise NCS 7, this situation is not crucial for the system, but enables a finer adaptation for a limited risk.

4.3.9. NCS 9: Context Agent Unproductiveness (suggested action)

Detection: A Context Agent whose action has been selected several times in a row considers itself in unproductiveness NCS. Indeed, the agent thinks that the
540 ideal case would be that its action should provoke a greater decrease of critical level so it only has to be performed once. The Context Agent hence seek to adjust the amplitude of the suggested action, in a way to maximize the decrease (or minimize the increase) of critical levels.

Resolution: The adjustment of the amplitude of the action is based on the
545 estimation of the effects of the variation of the amplitude on the variation of critical levels. The idea is to increase or decrease the amplitude of the action in a way to accelerate the decrease (or slow down the increase) of critical levels. To this end, a Context Agent which has been selected several times in a row slightly and randomly changes the amplitude of the suggested action and correlates this
550 variation with the speed variation of critical levels. Hence, if the highest critical level is decreasing:

- quicker while the amplitude has been increased: the Context Agent keeps increasing the amplitude ;
- quicker while the amplitude has been decreased: the Context Agent keeps
555 decreasing the amplitude ;
- slower while the amplitude has been increased: the Context Agent decreases the amplitude ;
- slower while the amplitude has been decreased: the Context Agent increases the amplitude ;

560 The Context Agent does the exact opposite if the highest critical level is increasing, although this rarely happen since it is not frequent that an action is continued if it has provoked a rise of the highest critical level. Note that a maximal amplitude can be set in order to avoid too brutal actions.

4.3.10. Conclusion on Non-Cooperative Situations

565 This section has presented the NCSs encountered by the agents of ESCHER. In particular, the resolution of these situations provokes the creation, the deletion, and the modification of Context Agents, which are the memory of the system. In other words, NCSs provoke the memorizing, the forgetting, and the correction of knowledge based on observations of the real system: their resolution
570 enables ESCHER to learn and self-adapt.

NCSs 1 and 2 correspond to the acquisition of new informations. They occur when ESCHER is discovering a new part of the state space of its environment. They open the system as they add new Context Agents.

NCS 3 enables ESCHER to not persist in error. It is partially solved thanks
575 to the reorganization of the relations between a Controller Agent and some of its Context Agents. Indeed, the Controller Agent ignores some of the Context Agents if they have been wrong.

Context Agents always self-evaluate. Hence, NCSs 4 to 9 are detected if one of the parts is no longer adapted to the environment. They are solved by the
580 adjustment of the agents (except for NCS 7 which is solved thanks to openness). Hence, ESCHER is always self-evaluating and self-adapting.

4.4. Learning and Adjustment

A large part of the learning of the system relies on the tuning of Context Agents' internal parameters during the resolution of a NCS. All these parameters
585 are tuned thanks to Adaptive Value Trackers (AVT, [28]). These parameters are: the boundaries of the validity ranges, the amplitude of the suggested action, and the values of the forecasts.

An AVT converges towards a value thanks to binary feedbacks: *lower* if the real value is lower, or *greater* if the real value is greater. Both the value and
590 the variation step of the tracker are dynamically tuned. The variation step is increased when two consecutive feedbacks are equal, and decreased otherwise. These variations follow user-defined coefficients. Figure 3 shows an example of the variation of the value of an AVT with standard settings (two equal con-




Figure 3: Typical Convergence of an Adaptive Value Tracker

secutive feedbacks double the variation step, two different consecutive feedback
 595 divide the variation step by three). A plus sign means the AVT received a
greater feedback, a minus sign means it received a *lower* feedback.

A Context Agent transforms its observations and received notifications into
 feedbacks for its numerous AVTs. For instance, a Context Agent
 in NCS 5 observing a greater variation of critical levels than what its forecast
 600 indicates will send a *greater* feedback to the corresponding AVT. The tracker
 then increases its value. Of course, the new value of the forecast may not
 be equal to the observation. But given the dynamics of the environment and
 the inevitable noise on real sensors, perfectly fitting to the observations is not
 desirable.

605 AVTs quickly converge toward a value, are able to stabilise, and to move
 again quickly toward a new further value. They match our needs, as the pa-
 rameters of agents often have to change, often drastically.

4.5. Comparison with Existing Approaches

ESCHER has been presented as a control system because it has been de-
 610 signed to control. Nevertheless, learning plays a crucial role in this system.
 This section explores this two complementary sides of our system and their
 links through comparisons with the Dual Control Theory and with Learning
 Classifier Systems.

4.5.1. Comparison with Dual Control.

615 In the Dual Control Theory, the controlled system is partially known. The controller applies either probe actions to learn and refine its model of the controlled system, or control actions to put the controlled system in the desired state [29]. Too many probes hampers the control, but too many control actions makes a small gain. Finding the balance between probe actions and control ac-
620 tions requires to solve the difficult Bellman equation, which is not easily feasible in real cases.

Like dual controllers, ESCHER faces unknown systems and learns from its actions. However, it learns from all of its actions and all of its actions seek to put the controlled system in the desired state. All of its actions are probes and
625 control actions at the same time. Moreover, unlike dual controllers, ESCHER does not need a predefined model that is later adjusted by learning.

The need to lower the critical levels (even when no agent indicates how to do it), combined to the fact that ESCHER learns from each of its actions, can be seen as an approach to solve the problem of balance between probes and control
630 actions. The control process drives the learning process towards interesting states of the environment, while getting closer to the desired state (and thus preventing to stray away and visit uninteresting distant states).

4.5.2. Comparison With Learning Classifier Systems.

Learning Classifier Systems (LCSs) are reinforcement learning systems [20]. They
635 are composed of a set of behavior rules, a pairing system which matches states of the environment with rules conditions, a selection mechanism between simultaneously triggered rules, and a genetic algorithm to tune the set of rules.

There are several similarities between a LCS and a Controller Agent coupled with its set of Context Agents. Context Agents play the same role than the pair-
640 ing system (with their validity ranges) and the set of rules (each Context Agent can be seen as a behavior rule since it suggests an action under certain conditions). The Controller Agent plays a similar role than the selection mechanism, choosing an action among several suggestions from valid Context Agents.

The main difference comes from the fact that Context Agents are autonomous,
645 they learn by themselves. On the contrary, the rules of a LCS are processed by
a genetic algorithm, to withdraw the weakest and generate new and presumably
more adapted rules. The fitness function of this algorithm is usually a reward
signal, perceived from the environment. A great difficulty in the instantiation
650 of a LCS is to adequately split the reward between the different rules. This dif-
ficulty does not exist in ESCHER, because of the autonomy of Context Agents.
They evaluate their adequacy themselves, and adjust themselves if needed. On
certain aspects, the notion of critical levels may be assimilated to the reward
signal, as it enables to evaluate the adequacy of the rules.

By self-adjusting, Context Agents suggest actions that are more and more
655 adequate, with a more and more adequate timing, along with more and more
reliable forecasts. Thus, the learning process feeds the control process.

4.6. Settings

For ESCHER to be easy to instantiate to a particular system, the number
of parameters has to be as low as possible, and setting them should not require
660 the use of elaborate calibration methods.

The only knowledge about the controlled system that ESCHER needs is
quite simple:

- the number of controlled variables, and their references;
- the number of observed variables, and their references.

665 It is possible to give the lower and higher bound for each variable. ESCHER
works without this information, but it can be of use for the criticality func-
tions. Anyway, this is basic knowledge about the controlled system, it is not an
obstacle.

The only difficulty in the instantiation ESCHER is the definition of the
670 criticality functions. Controller Agents focus on the most critical Criterion
Agent. This means that the compromise between several criteria is expressed
through the definition of the criticality functions. For instance, in an absurd

Parameters	Significance
Number of controlled variable	Important
Number of observed variables	Important
Variables references	Important
Criticality functions	Important
Variation ranges	Optional
Maximal size of an action	Incidental
Minimal size of a validity range	Incidental
Minimal step of an AVT	Incidental
Coefficients of an AVT	Incidental

Table 1: Parameters of ESCHER and their significance

case, if we want to maximize and minimize the same variable, ESCHER will stabilize on the value where the two criticality functions meet. This knowledge
675 concerns not only the controlled system, but also the objectives of the user.

Finally, some other parameters are secondary. They have a very limited impact on the overall performance of the system, they do not require to be specifically set each time, their default values work fine. It is, for instance, the minimal size of validity ranges (that triggers NCS 7), the maximal size of an
680 action (to prevent ESCHER to perform brutal actions, for safety reasons), or the internal parameters of AVTs. The strong and quick adaptiveness of the agents reduces the impact of these parameters. Table 1 shows all the parameters of ESCHER and their significance.

5. Experiments: Real-Time Control of Combustion Engines

685 The first experiments presented in this section have been conducted on automatically generated synthetic black boxes. Then, experiments on a real combustion engine are shown. The implementation of ESCHER used for these experiments is a prototype written in Java 1.7 using Eclipse and a component-based multi-agent architecture generator called Make Agent Yourself [30]. It

690 runs on a laptop with an Intel i7 2.67GHz CPU and 4GB of RAM. The du-
 ration of a lifecycle of ESCHER (i.e. a lifecycle of each of its agents) depends
 mainly on the number of agents. It is approximately 20 ms with 10 agents, and
 500 ms with 800 agents. This is something that should be improved by code
 optimization, but this is not the immediate concern for ESCHER. Here the goal
 695 is to show that the agents are indeed able to learn how to control several inputs
 of an unknown system, regarding several criteria.

5.1. Criticality Functions

The function¹ used in our experiments to compute critical levels is defined
 over \mathbb{R} as follows (equation 10) :

$$f(x) = \begin{cases} 100 & \text{if } x \leq 0 \\ \gamma \frac{(x-\eta)^2}{2\eta} + \gamma(x-\eta) + \delta & \text{if } 0 < x \leq \eta \\ -\gamma \frac{(x-\eta)^2}{2(\epsilon-\eta)} + \gamma(x-\eta) + \delta & \text{if } \eta < x \leq \epsilon \\ 0 & \text{if } \epsilon < x \leq sup - \epsilon \\ -\gamma \frac{(sup-x-\eta)^2}{2(\epsilon-\eta)} + \gamma(sup-x-\eta) + \delta & \text{if } sup - \epsilon < x \leq sup - \eta \\ \gamma \frac{(sup-x-\eta)^2}{2\eta} + \gamma(sup-x-\eta) + \delta & \text{if } sup - \eta < x \leq sup \\ 100 & \text{if } sup < x \end{cases} \quad (10)$$

with

$$\gamma = -2 \frac{100}{\epsilon}$$

and

$$\delta = -\gamma \frac{(\epsilon - \eta)}{2}$$

Parameters sup , ϵ et η are defined by the user. The curve of this function is
 symmetrical with respect to the center of $[0; sup]$, it decreases on $[0; \epsilon]$, and
 700 increases on $[sup - \epsilon; sup]$. Parameter η defines the inflection point, sup acts
 as the upper bound of the function, above this value the critical level is always
 100, and ϵ defines the interval $[\epsilon; sup - \epsilon]$ where the critical level is always zero.

¹Function whose formula was proposed by our colleague Sophie Jan, at the Toulouse In-
 stitute of Mathematics

In our implementation, it is possible to shift the function so the slopes happen in an arbitrary interval instead of $[0; sup]$. It is also possible to make the function
705 asymmetrical by defining different ϵ and η for each half of the interval. For instance, by setting $\epsilon = 0$ for the left half only, we obtain a curve similar to the threshold one from figure 2.

It is worth remembering that each Criterion Agent has its own function set differently. It is up to an expert of the controlled system domain to set the pa-
710 rameter of each criticality function. This is how the balance between all criteria is expressed to ESCHER, as it always try to lower the most critical criterion before the others. However, our prototype has a simplified procedure regarding the experiments. The user does not have to directly manipulate equation 10, she or he only needs to select each critical variable and indicate whether the
715 function he or she wants is a threshold, a setpoint, a minimization or a maximization. After specifying the threshold value or the setpoint value, ϵ and η are generated automatically.

5.2. Experiments on Synthetic Black Boxes

The use of a black box generation tool [31] enabled us to test ESCHER over
720 50 cases of various complexity, with up to dozens of inputs and outputs. We present here two very simple cases to provide a better understanding on how ESCHER reaches a compromise between several criteria, and how it is robust to perturbations. In these experiments, a cycle corresponds to a lifecycle of each agent followed by a simulation step of the black box.

725 5.2.1. Optimizing Two Criteria

In this experiment, the black box has one input (I1) and two outputs (O1 and O2) varying from zero to 100. The setpoint on both outputs is 50. There are two Criterion Agents, one for each output, each with the same criticality function. Hence, both criteria have the same weight. However, this setpoint
730 is not reachable on both output at the same time, there is no value for the

input that put both output at 50. ESCHER has to find a compromise, i.e. to minimize the highest critical level.




Figure 4: Optimization of two Criteria

Figure 4 shows the variations of the input and outputs of the controlled black box, of the number of Context Agents in the system, and of the critical levels. The input is initialized to 1.1, which sets O1 to 21.8 and O2 to 1.8. O2 is further from the setpoint than O1, its critical level is therefore higher. ESCHER has no preliminary knowledge on the black box. Its action at the first step is a mistake, ESCHER slightly increase the input which provokes a small increase of both critical levels. A Context Agent for this action is created. The following step, ESCHER corrects this mistake, and find the action which push the outputs towards the setpoint. A second Context Agent is created, which action is kept until the highest critical level stops decreasing.

The critical level of O1 reaches 0 at lifecycle 76. However, the critical level of

O2 is then at 26.1, and still decreasing. The action is continued, since the highest
745 critical level is decreasing, even though the other critical level is increasing.

At lifecycle 96, critical level of O1 becomes higher than critical level of O2. In
consequence, ESCHER modifies its action, and critical levels cross again. A serie
of oscillations follows, during which 3 new Context Agents are created. Finally,
the value of the input is stabilized, slightly oscillating around 3. O1 oscillates
750 around 60 and O2 around 40. Both critical levels oscillate around 5. ESCHER
has reached the best compromise (according to the criticality functions), since
the highest critical level is the lowest possible.

This experiment shows how a Controller Agent is able to deal with an input
that control several outputs with antinomic criteria. Different criticality func-
755 tions would have lead to a different compromise. For instance, one can prioritize
one output over the other by making a criticality function always greater than
the other.

5.2.2. Robustness

This experiment shows how ESCHER reacts to perturbations in its environ-
760 ment. Here, ESCHER controls two of the three inputs (I1 and I2) of a black
box. The third input (I3) is manually controlled. These three inputs have an
influence on the same output (O1), on which a setpoint criterion is applied.
First, we let ESCHER make O1 meet the setpoint by modifying I1 and I2.
Then, we manually change the value of I3, provoking a perturbation on O1,
765 which abruptly go away from the setpoint. ESCHER must adapt itself to this
modification by finding new values for I1 and I2.

Figure 5 shows the variations of the input and outputs of the black box,
along with the number of Context Agents and the critical level of the setpoint
criterion. Inputs are initialized to 1, which sets the output to 68. The setpoint
770 is 50. ESCHER reaches the setpoint in less than 100 lifecycles by increasing I2
only.

At lifecycle 160, I3 is manually set to 50. This makes O1 decrease, jumping
out of the setpoint, resulting in a peak of crtitical level, which rises from 0 to 12.




Figure 5: Robustness to Perturbations at Runtime

This is resorbed by ESCHER, which decreases I2 until the setpoint is reached
 775 again.

I3 is once again modified at lifecycle 220, from 50 to 100. This provokes a
 huge increase of the output, therefore a rise of critical level (from 0 to 72). Once
 again, ESCHER self-adapts to this perturbation. First, I1 is increased, then I2.
 The critical level is brought back to 0 at lifecycle 350, while new Context Agents
 780 have been created. Two other perturbations are later performed. Each time,
 ESCHER is able to bring back the output on the setpoint.

This experiment shows that ESCHER is able to react to perturbations on
 the controlled system. It self-adapts to changes to maintain an adequate control.
 Here, each perturbation is big enough to provoke the creation of new Context
 785 Agents.

5.3. Experiments in Real Conditions

The results presented in this section have been obtained during tests driven on a 125cc monocylinder fuel engine. The engine was instrumented so ESCHER has access to temperatures, pressures, and others, via the Engine Control Unit (ECU) and a gas analyzer.

The link between the engine and the ECU is assured thanks to various specific instruments. A Controller Area Network (CAN) bus enables the communication of external systems with the ECU. CAN buses are widely used in the automotive industry. A computer software called ControlDesk enables the reading on the ECU (in particular of the variables measured by the sensors), the computation of values from read variables, and the modifications of parameters (such as the ignition advance). ESCHER is connected to ControlDesk via a specific communication protocol, MCD-3 (stands for Measurement, Calibration, Diagnostics) over Ethernet, enabling our system to read and write values on the ECU. Finally, a gaz analyzer is plugged onto the engine exhaust. It measures gas concentration of various pollutants (carbon monoxide, for instance), and sends data via a serial output (RS232/DB25) interfaced with the USB port of the computer on which ESCHER runs. Figure 6 shows this set-up. For these experiments, ESCHER had to be slowed down and wait at least 10 seconds between each lifecycle in order to let the engine stabilize after changing its parameters. For the second experiment, ESCHER had to wait 10 more seconds between each lifecycle for the gas analyser to provide data.

5.3.1. Torque Optimization

In this experiment, the engine is put at 5000 rpm, with a load of 870 mbar in the intake manifold. ESCHER controls the total injected fuel mass and the ignition advance. The only control criterion is to maximize the indicated mean effective pressure (IMEP), which reflects the torque.

The injected fuel mass is measured in milligrams per shot (mg/shot), and the ignition advance in crankshaft degrees ($^{\circ}$ c), i.e. the position of the piston in the cylinder when the combustion is triggered. IMEP is measured in bars. IMEP




Figure 6: Experimental Set-Up for the Tests on a Real Engine

is a very unstable variable, in particular with monocylinder engines. Working at high rpm and high load, as it is the case in this experiment, reduces the instability.

The criticality function is strictly decreasing (since we want to maximize
 820 IMEP). We do not know a priori what is the maximal reachable IMEP, therefore we can't set the criticality function in a way that it returns 0 when the maximal PMI is reached. Thus, we do not expect the critical level to be zero at the end of the test, but we do expect it to be lower at the end than at the start. This is true for every criticality function used with the real engine.

825 Figure 7 shows the variations of the controlled inputs, the optimized output, the number of Context Agents and the critical level. At the start, the injected fuel mass is low (7 mg/shot) regarding the current operating point. The engine is on the verge of stalling. Of course, ESCHER which doesn't have any knowledge about the engine, is not aware of this fact. Its first action is a mistake: ESCHER
 830 decreases both parameters, which leads to a drop of IMEP (and a rise of critical level).

ESCHER quickly finds a way to make the critical level decrease, by increasing first the injected fuel mass, then the ignition advance. IMEP finally reaches its maximum (about 9 bars), the critical level stops decreasing. ESCHER stabilizes
 835 itself at 11.50 mg/shot of injected fuel, with a 2424°c ignition advance. The decrease of these inputs at lifecycle 24 is explained by nose on the IMEP. But




Figure 7: IMEP Optimization while Controlling Two Parameters

the system quickly corrects itself.

ESCHER managed to improve the IMEP by 3 bar in 9 lifecycles (about 90 seconds), reaching the maximal IMEP possible for the considered operating point. Obtaining the same result takes a skilled engineer, used to this particular engine, around 20 minutes with usual methods.

5.3.2. Multi-Objective Optimization

For this test, the engine is put in another operating point (2500 rpm, 750 mbar). ESCHER controls the injected fuel mass the ignition advance, but also the start of injection (SOI). This new parameter is the timing of the injection relatively to the position of the piston, it is measured in crankshaft degrees. There are criteria on four outputs:

- IMEP must be maximized ;




Figure 8: Inputs and Critical Levels during a Multi-Objective Optimization

- fuel consumption, measured in g/kWh, must be minimized ;
- 850 • hydrocarbons (HC) emission must be under 500 ppm (parts par million) ;
- carbon monoxide concentration (CO) must be under 3%.

The last three criteria are contradictory with the first one. Indeed, the most efficient way to improve IMEP is to inject more fuel. However, this also increase fuel consumption and pollutants emissions. We need to adjust ignition advance and SOI to extract more power from the combustion. This is what ESCHER has to learn.

855 Figure 8 shows the variations of the controlled parameters and the critical levels, while figure 9 shows the variations of the outputs. At the beginning, the highest critical levels is that of fuel consumption. Thus, ESCHER seeks




Figure 9: Engine Outputs during a Multi-Objective Optimization

860 to decrease the fuel consumption critical level in priority. The system manages to do so during the first 20 lifecycles, in particular by increasing the ignition advance from 10 to 26°c and by decreasing the SOI from -150 à -400°c, while the fuel injection oscillates between 6 and 7 mg/shot.

At lifecycle 10, IMEP maximization becomes the most critical criterion, 865 however, its critical level is decreasing, so the same actions are continued. At lifecycle 20, the CO threshold is crossed, its critical level rises. ESCHER explores new actions to solve this problem. It continues to decrease SOI but start to decrease ignition advance. This lead to a peak of consumption and a drop of IMEP between lifecycles 45 and 50, along with small excesses of hydrocarbons. 870 Finally, after some oscillations, ESCHER manages to put the pollutants under their respective thresholds, while maintaining a high IMEP and a low

consumption.

At the end of the test, IMEP is around 8 bar (2 bar higher than the begin-
ing), while fuel consumption is around 275 g/kWh (165 g/kWh less than the
875 initial value). Pollutants emissions are higher than their initial values, but they
meet their threshold. ESCHER has successfully completed a standard engine
optimization (i.e. optimizing torque and consumption while respecting pollu-
tion thresholds) without having any prior knowledge about engines. This test
lasted 123 lifecycles, around 41 minutes (ESCHER has to wait for the gaz ana-
880 lyzer). This is about twice as fast than a human expert with usual methods for
a similar end result.

6. Conclusion and Perspectives

This article presented ESCHER, an system that illustrates the contributions
of the AMAS approach to the field of control systems and calibration. This ar-
885 ticle focused on the full presentation of the system, and showed results obtained
both with unrelated black-box simulations and real engines. The goal with
the experiments on black-boxes was to illustrate how ESCHER works on basic
cases. Experiments on the real engine show its applicability in real conditions
and its robustness to noisy data. Overall, the automatic calibration performed
890 by ESCHER is faster than methods used in the industry for a similar result.
However these experiments highlight a limitation of ESCHER. We had to make
it wait between its lifecycles for the engine to stabilize and for the gaz analyzer
to provide data. This is due to its inability to correlate actions and effects
if the effects become sensible too long after the action. Further papers will
895 present comparisons with other learning methods, detailing the advantages and
limitations of each.

The AMAS approach breaks with the traditional top-down design of artificial
systems. It focuses on the local behavior of agents, leaving them the task of
controlling their own organization. An adequate global function emerges from
900 this local self-organization process. We hope this is the first step towards a fully

self-reconfigurable ECU.

Other AMASs have tackled the problem of learning and control with similar Context Agents, for instance with model generation [32] and ambient robotics [33]. Context Agents are being generalized and standardized to become a pattern for context learning in a multi-agent system [34].

AMASs are a young technology compared to the majority of AI methods used in intelligent control, such as artificial neural networks or genetic algorithms. Our future work must focus on the formalization of the approach to enable a priori proofs of AMAS properties. This is a work in progress, which first steps have been made with Event-B [35] and continuous approximation [36].

Acknowledgements

This work is part of the ORIANNE project, in partnership with Aboard Engineering (leader), FH Electronics, IRIT, IRSEEM, CERTAM, CEVAA, and Renault.

References

- [1] W. R. Ashby, An Introduction to Cybernetics, Chapman & Hall, London, UK, 1956.
- [2] F. Heylighen, Complexity and self-organization, in: M. N. Bates, Marcia J. & Maack (Ed.), Encyclopedia of Library and Information Sciences, 3rd Edition, Taylor and Francis, 2008, pp. 1215–1224.
- [3] M. A. Khamis, W. Goma, Adaptive multi-objective reinforcement learning with hybrid exploration for traffic signal control based on cooperative multi-agent framework, Engineering Applications of Artificial Intelligence 29 (2014) 134 – 151.
- [4] E. Morin, On complexity, Hampton Press, 2008.
- [5] A. N. Kolmogorov, On tables of random numbers, Theoretical Computer Science 207 (2) (1998) 387–395.

- [6] T. Deacon, S. Koutroufinis, Complexity and dynamical depth, *Information* 5 (3) (2014) 404–423.
- 930 [7] L. Von Bertalanffy, *General System Theory: Foundations, Development, Applications*, George Braziller, New York, 1968.
- [8] I. S. Jesus, R. S. Barbosa, Tuning of fuzzy fractional $pd^\beta+i$ controllers by genetic algorithm, in: *Proceedings of the 10th International Conference on Informatics in Control, Automation and Robotics (ICINCO 2013)*, INSTICC Press, Reykjavik, Iceland, 2013, pp. 282–287.
- 935 [9] S. G. Fabri, M. K. Bugeja, Kalman filter-based estimators for dual adaptive neural control: A comparative analysis of execution time and performance issues, in: *Proceedings of the 10th International Conference on Informatics in Control, Automation and Robotics*, INSTICC Press, Reykjavik, Iceland, 2013, pp. 169–176.
- 940 [10] L. Bull, J. Sha’Aban, A. Tomlinson, J. D. Addison, B. G. Heydecker, Towards distributed adaptive control for road traffic junction signals using learning classifier systems, in: *Applications of Learning Classifier Systems*, Springer, 2004, pp. 276–299.
- 945 [11] M. C. Choy, D. Srinivasan, R. L. Cheu, Neural networks for continuous online learning and control, *Neural Networks, IEEE Transactions on* 17 (6) (2006) 1511–1531.
- [12] T. M. Mitchell, *The discipline of machine learning*, Carnegie Mellon University, School of Computer Science, Machine Learning Department, 2006.
- 950 [13] J. Kober, J. A. Bagnell, J. Peters, Reinforcement learning in robotics: A survey, *The International Journal of Robotics Research* 32 (11) (2013) 1238–1274.
- [14] C. Watkins, P. Dayan, Q-learning, *Machine learning* 8 (3-4) (1992) 279–292.

- [15] V. V. Raghavan, V. N. Gudivada, V. Govindaraju, C. R. Rao, Cognitive
955 Computing: Theory and Applications, Elsevier, 2016.
- [16] C. A. Astudillo, B. J. Oommen, Topology-oriented self-organizing maps: a
survey, *Pattern analysis and applications* 17 (2) (2014) 223–248.
- [17] L. Deng, D. Yu, Deep learning: methods and applications, *Foundations
and Trends® in Signal Processing* 7 (3–4) (2014) 197–387.
- 960 [18] J. C. Bongard, Evolutionary robotics, *Communications of the ACM* 56 (8)
(2013) 74–83.
- [19] K. R. Thórisson, A new constructivist ai: From manual methods to self-
constructive systems, in: *Theoretical Foundations of Artificial General In-
telligence*, Springer, 2012, pp. 145–171.
- 965 [20] R. J. Urbanowicz, J. H. Moore, Learning classifier systems: a complete
introduction, review, and roadmap, *Journal of Artificial Evolution and Ap-
plications* 2009 (2009) 1.
- [21] M. Wooldridge, *An introduction to multiagent systems - second edition*,
John Wiley & Sons, 2009.
- 970 [22] J. Ferber, *Multi-agent systems: an introduction to distributed artificial
intelligence*, Addison-Wesley Reading, 1999.
- [23] W. Ren, Y. Cao, *Distributed Coordination of Multi-agent Networks: Emer-
gent Problems, Models, and Issues*, Springer Publishing Company, Incor-
porated, 2013.
- 975 [24] G. Di Marzo Serugendo, M.-P. Gleizes, A. Karageorgos, Self-organising
systems, in: G. Di Marzo Serugendo, M.-P. Gleizes, A. Karageogos (Eds.),
Self-organising Software: From Natural to Artificial Adaptation, Natural
Computing Series, Springer Berlin Heidelberg, 2011, pp. 7–32.
- [25] J.-P. Georgé, M.-P. Gleizes, V. Camps, Cooperation, in: G. Di Marzo Seru-
980 gendo, M.-P. Gleizes, A. Karageogos (Eds.), *Self-organising Software: :*

From Natural to Artificial Adaptation, Natural Computing Series, Springer Berlin Heidelberg, 2011, pp. 193–226.

- 985 [26] N. Bonjean, W. Mefteh, M.-P. Gleizes, C. Maurel, F. Migeon, Adelfe 2.0, in: M. Cossentino, V. Hilaire, A. Molesini, V. Seidita (Eds.), Handbook on Agent-Oriented Design Processes, Springer Berlin Heidelberg, Berlin, Heidelberg, 2014, pp. 19–63.
- [27] P. Kruchten, The rational unified process: an introduction, Addison-Wesley Professional, 2004.
- 990 [28] S. Lemouzy, V. Camps, P. Glize, Principles and properties of a MAS learning algorithm: A comparison with standard learning algorithms applied to implicit feedback assessment, in: 2011 International Conference on Web Intelligence and Intelligent Agent Technology, Springer, Lyon, France, 2011, pp. 228–235.
- 995 [29] A. A. Feldbaum, Dual control theory, I-IV, Automation Remote Control 21-22.
- [30] V. Noël, Component-based Software Architectures and Multi-Agent Systems: Mutual and Complementary Contributions for Supporting Software Development, Phd thesis, Université de Toulouse, Toulouse, France (2012).
- 1000 [31] J. Boes, P. Glize, F. Migeon, Mimicking Complexity: Automatic Generation of Models for the Development of Self-Adaptive Systems, in: International Conference on Simulation and Modeling Methodologies, Technologies and Applications, INSTICC Press, Reykjavik, Iceland, 2013, pp. 243–250.
- 1005 [32] J. Nigon, M.-P. Gleizes, F. Migeon, Self-adaptive model generation for ambient systems, in: Proceedigns of the 7th International Conference on Ambient Systems, Networks and Technologies (ANT 2016), Vol. 83, Elsevier, 2016, pp. 675–679.
- [33] N. Verstaevel, C. Régis, M.-P. Gleizes, F. Robert, Principles and Experimentations of Self-Organizing Embedded Agents Allowing Learning From

- Demonstration in Ambient Robotics, *Future Generation Computer Systems*, Emerging, Ambient and Ubiquitous Systems 64 (2016) 78–87.
- 1010
- [34] J. Boes, J. Nigon, N. Verstaevel, M.-P. Gleizes, F. Migeon, The Self-Adaptive Context Learning Pattern: Overview and Proposal, in: *International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT)*, no. 9405 in *LNAI*, Springer, Larnaca, Cyprus, 2015, pp. 91–104.
- 1015
- [35] Z. Graja, F. Migeon, C. Maurel, M.-P. Gleizes, L. Laibinis, A. Regayeg, A. H. Kacem, A pattern based modelling for self-organizing multi-agent systems with event-b, in: *International Conference on Agents and Artificial Intelligence*, INSTICC Press, Angers, France, 2014, pp. 223–236.
- 1020
- [36] S. Stuker, F. Adreit, J.-M. Couveignes, M.-P. Gleizes, Continuous approximation of a discrete situated and reactive multi-agent system: Contribution to agent parameterization, in: H. K. Dam, J. Pitt, Y. Xu, G. Governatori, T. Ito (Eds.), *PRIMA 2014: Principles and Practice of Multi-Agent Systems: 17th International Conference, December 1-5, 2014*, Springer International Publishing, Gold Coast, Australia, 2014, pp. 365–380.
- 1025