Levels of Emergent Behaviour in Agent Societies

Adina Magda Florea, Eugenia Kalisz, Andrei Olaru

Department of Computer Science
"Politehnica" University of Bucharest
Spl. Independentei 313, Bucharest, Romania
E-mail: adina.florea@cs.pub.ro, eugenia.kalisz@cs.pub.ro, cs@andreiolaru.ro

Abstract
Artificial agents societies are well suited to design and implement open distributed systems. As the complexity of such systems grows, the design of agent societies with a complete pre-defined behaviour is a significant challenge due to the dynamic interactions among agents and between agents and the environment. To overcome existing difficulties, agent systems with emergent behaviour are a fertile area of research and span a large range of applications. The paper presents an analysis of the anticipatory capabilities of multi-agent systems endowed with emergent behaviour, by considering both reactive multi-agent systems and cognitive multi-agent systems. Our approach in analysis is driven by the different levels on which predicted behaviour can be achieved and is illustrated by two scenarios, one for each case of agent system, grounding the view for the cognitive case in our previous related work on self-organizing systems with cognitive agents.

Keywords: multi-agent systems, emergent behaviour, anticipatory behaviour

1. Introduction

In open systems, situated agents are evolving in a continuously changing environment that requires adequate behaviour and responses. In this context, the design of agents and multi-agent systems with a complete pre-defined behaviour is a significant challenge due to the complex possible interactions among agents and between agents and the environment. To overcome existing difficulties, agent systems with emergent behaviour are a fertile area of research. Several definitions of emergence exist but none is yet generally accepted. Most of the time, emergence is defined by its effect as the resulting properties of a system that are not captured by properties of the parts, such as the formation of patterns in the structure or behaviour of the overall system.

Current work on emergent behaviour in multi-agent systems is mainly focused on properties and behaviour in systems formed of a large number of reactive agents. Reactive agents have simple behaviour and need very limited computational capacity, while emergent functions may be more complex and the structure is robust. However, limitations regarding deliberative capabilities and adaptation do not always make reactive multi-agent systems a suitable choice for modelling complex problems. Cognitive agents explicitly represent their knowledge and beliefs, have goals they desire to fulfil, make plans and take action in order to fulfil them, and may be endowed with elaborated reasoning mechanisms. Recently, there is a growing interest in the design of cognitive multi-agent systems with emergent behaviour, in which pre-defined properties and functionality is gradually enriched with emergent behaviour. Both at the level of reactive and cognitive multi-agent systems, anticipated behaviour of the system and of the individual parts is a cornerstone in the design of such systems, as emergent properties may be either desired or undesired ones.

There is a large range of applications for such systems, especially in distributed solutions built up with nodes having a significant degree of autonomy. Among these, we can mention grid computing systems, loosely-coupled distributed data storage featuring caching and replication, P2P systems, web applications, including information searching and filtering. Using agents to model and implement such systems is a good choice considering the flexibility and dynamic character of the paradigm. Cognitive agents with emergent behaviour, for example, need not be very heavy and such systems may be implemented even on small devices, like the ones used in pervasive computing and sensor and actuator networks. Even if the agents are light, like the reactive ones, the emergent behaviour may still be sufficiently complex.

The work presented in this paper provides a brief overview of artificial multi-agent societies with emergent behaviour and an insight into the anticipatory capabilities of multi-agent systems endowed with such behaviour, by considering reactive multi-agent systems, cognitive multi-agent systems, and different levels of analysis of the predictive ability and anticipated behaviour of such systems. The analysis takes into account both the "self" view of one agent as regarding properties and behaviour, and the observer's view, focused on whether the system satisfies or not the desired behaviour. For both types of agent societies, reactive and cognitive, the paper proposes a view of a system endowed with emergent behaviour and anticipatory abilities.

The paper is structured as follows. Section 2 presents an overview of the concept of emergence, systems with emergent behaviour, and why prediction is important in such systems. Section 3 deals with the emergence in reactive agent systems and implications of emergent properties on the predictive ability of the system, if any, while Section 4 describes emergence in cognitive multi-agent systems and analyses different levels of anticipatory capabilities of the system. Both Section 3 and 4 propose a model in which the prediction of emergents can be identified, and present associated case studies. Section 5 is devoted to conclusions and further work.

2. Emergence, self-organizing systems and anticipation

The notion of an emergent effect was first proposed in 1875 by the philosopher George Lewes [10] to describe non-additive effects of causal interactions, to be contrasted with resultants. Several other theories of emergence appeared since then, in which the concept of emergence as a relation between simultaneous causes and their joint effect was translated to consider the upward causation of composition. Nowadays, emergence theories are mainly focused on properties rather than dynamical interactions, by considering the relationship between components and the whole they compose. This view on emergence was held in contrast to reductionism mechanism, the ideal that all apparently different kinds of matter are the same stuff, differing only in the number, arrangement and movement of their constituent components [19].

Throughout recent literature, there have been many attempts to define emergence; however, to the moment, none of the definitions is generally accepted. A simple definition is that "emergence is the concept of some new phenomenon arising in a system that wasn't in the system's specification to start with" [21]. A more elaborated definition is based on the fact that emergence needs *two levels of perspective*: the inferior, or micro-level of the individual entities, and the superior, or macro-level of the whole system. Considering this view, "a system exhibits emergence when there are coherent emergents at the macro-level that dynamically arise from the interactions between the parts at the micro-level. Such emergents are novel with respect to the individual parts of the system" [4]. The "emergent" in this definition is a general concept that denotes the result of the emergence, which can be a property, a structure, behaviour, or some other phenomenon.

The emergence of order and structures in nature can be explained by the dynamics and attractors of complex systems - systems composed of a large number of interacting individual entities [1, 11]. They result from collective patterns of interacting elements in the sense of many-bodies problems that cannot be reduced to the features of single elements in a complex system. Interactions among many components in complex systems often have synergetic effects, which can neither be traced back to single causes nor be forecasted in the long run or controlled in all details. Obviously, self-organization leads to the emergence of new phenomena on sequential levels of evolution. Nature has demonstrated that self-organization is necessary, in order to manage the increasing complexity on these evolutionary levels.

Based on the definitions above but also on some other definitions [4, 8, 9], there are some important features of emergence that can be identified:

- emergence occurs out of the interactions between the parts of a complex system;
- emergence is defined in relation with two levels – it is manifested at the higher level, arising from interactions at the lower level;
- the representation of the emergent cannot be reduced to the specifications of the individuals;
- emergents only arise after a certain time in which the system has evolved;
- once they occurred, emergents will maintain identity over time;
- emergents arise without any centralised or exterior control;
- the emergent phenomena are robust and flexible, i.e., they are not influenced by damage on the system.

Considering the large scale distribution and complexity of current interconnected computing systems, it became apparent that pre-programmed models of behaviour and communication will be inadequate. In this context, the modelling and design of artificial systems with emergent behaviour became a prevalent interest in both scientific and industrial communities. For the multi-agent systems community, emergence is interesting because it can produce, at a higher level (the level of the system), properties that are not explicitly needed or implemented at the lower level (the level of the individual agents). In the study of agent systems, emergence is almost always related to the notion of self-organisation [9] – the organisation of the system emerges dynamically from interactions among agents, with no external or centralised control, resulting in an organisation that is robust and redundant.

Since its application to designing artificial systems, emergence has been mostly explained in terms of levels of observation, by identifying two levels of description: the micro-description, namely the description of the system entities, and the macro-description, namely the description of the system as a whole. If the micro-description is clear as being the design of the individual parts of the system, the macro-description, however, is difficult to define.

Randles et al. [17] say that the vital part of an emergent behaviour system is the observer who is capable of detecting the features of the system and from whose perspective the macro-description can be created. If the micro-description is clear once the individual entities have been designed, there may be more that one macro-description to describe the resulting behaviour of the system [6], so there is a need to find a reasonably good theory [21] that is equally explanatory and predictive.

Some other authors claim that this approach has led to confusion and replace the idea of level by a framework of scope, resolution and state, in which emergent properties are determined by the relationship between the scope of macro-state and micro-state descriptions [19]. This approach would establish a normative definition of emergent properties and emergence that makes sense of previous descriptive definitions of emergence and sheds light on which classes of emergent properties are epistemic and which are ontological.

From the point of view of a system designer, emergence is important because it allows obtaining, as the system output, a behaviour or function that is of higher level (or complexity) than the specification of its components [7]. The difficulty rests in determining how an emergent function can be obtained and how the system can be lead to obtaining the desired emergent function [20].

We can certainly claim that the *anticipated behaviour* of an artificial *system with emergent behaviour* is a cornerstone in the design of such a system, as emergent properties may be either desired or undesired ones. As stated above, a design engineer would want to know how an emergent property can be obtained in relation to the design of the interacting components and how he/she can guarantee that the desired property, placed on a higher level of description and/or complexity, will effectively emerge.

There are several definitions of anticipation, among which we quote: "An anticipatory system is a system whose current state is determined by a future state. The cause lies in the future." [18]. An anticipatory system may be difficult to model, since the value of its current state would be determined by the value of a future state. However, from the point of view of the designer and looking at the anticipatory system behaviour, one can circumvent this problem as follows: a behaviourally anticipatory system can have, at time *t*, one or more predicted model(s) of itself and/or of its environment at time $t+n$. Hence, at time t, its next state can depend on the current image of the predicted model (at $t+n$) of its state and/or its environment. Since the current image, at time t, of the predicted model of itself and/or of the environment is well specified, the next state can be calculated without any contradiction with causality [15].

In what follows, we shall take this view of anticipatory systems, namely behaviourally, and try to shed light on the predictive abilities of multi-agent systems, by separately considering systems formed with reactive or cognitive agents. Our analysis is mainly driven by the designer point of view and puts forward a model in which we attempt to reconcile both level/observer and the scope/resolution view of emergent behaviour systems.

3. Emergent behaviour in reactive agent societies

3.1 Emergence in reactive agent systems

In the field of multi-agent systems, emergent behaviour has been analysed mostly in the context of systems composed of reactive agents. This is because, on one hand, they are inspired by natural systems composed of simple individuals that act reactively and interact mainly by means of their environment [13] and, on the other hand, because reactive agents are easier to implement and study. Moreover, the simplicity of reactive agents makes them adequate for very small devices, with low computational power, like low-power sensors that form sensor networks, or particle computers [12].

A *reactive agent system* may be formally defined as a tuple, where E is the set of states of the environment, P is the set of precepts the agent can get from the environment, Ac is the set of actions an agent can execute, and M is the set of messages the agent can send to other agents:

$RA = (E, P, Ac, M,$ perceive, act, interact, change$)$

perceive : $E \rightarrow P$, act : $P \times M$ $[\times M \times \ldots \times M] \rightarrow Ac$

interact : $P \rightarrow M [\times M \times \ldots \times M]$, change : $E \times A \times A \times \ldots \times A \rightarrow E$

In reactive agent systems, the emergent organisation is, most of the time, of a physical (in the sense of spatial or space-related) nature. This is not very unexpected, as the language of the individuals composing the system is also space-related: it describes movement, position, and direction. The emergent property or behaviour, although novel and possibly unexpected, is not of a different nature than the properties and behaviour of the individual entities, it is just of a higher level. The advantage of obtaining an emergent pattern is that the agents do not need to be aware of the structure in order to form it. Moreover, self-organised structures have the properties of flexibility, redundancy and robustness - no individual agent is absolutely necessary to the structure and reduced damage to the structure has no permanent effect on it, as the agents reorganise and form the same pattern again. It is important to observe, however, that the resulting structure is always implicitly described in the behaviour of the individual agents that form it [12]. Still, explicit structure and organisation awareness would indeed need agents with better capabilities of representation and memory.

3.2 Predictive ability of reactive emergent behaviour

Following the definitions of emergence presented in Section 2, it is natural to say that the key feature of a system with emergent behaviour is that the emergent becomes apparent or observable. The question to be asked is "apparent to whom/what?" or "observed by whom/in what way?" Sharing the view presented in [17], we state that an observer is an important part of any system seeking to engineer or make use of system emergence. Therefore the *observer*, or the *system designer*, situated outside the agent society, is certainly the one to whom the emergent should be apparent or observed. This is a common feature of both reactive and cognitive agent systems with emergent behaviour. However, as opposed to [17] and as it will be shown in the next section, we claim that the cognitive agent itself, based on its high-level representation capabilities, is able to detect and exploit in a predictive manner the emergent of the system or, at least, the emergent behaviour from its self point of view.

Figure 1 describes a view of reactive agents in which the agent society and the system model are decoupled. The agents are simple entities, as stated in the previous section, while the system model is accessible to the observer or system designer. The designer observes the behaviour of the entire system and compares it with the predicted (or desired) behaviour that is anticipated based on the system model. According to the results of this evaluation, the observer will be likely to operate modifications in the system model so as to correct the observed behaviour in case it is not according to the predicted one.

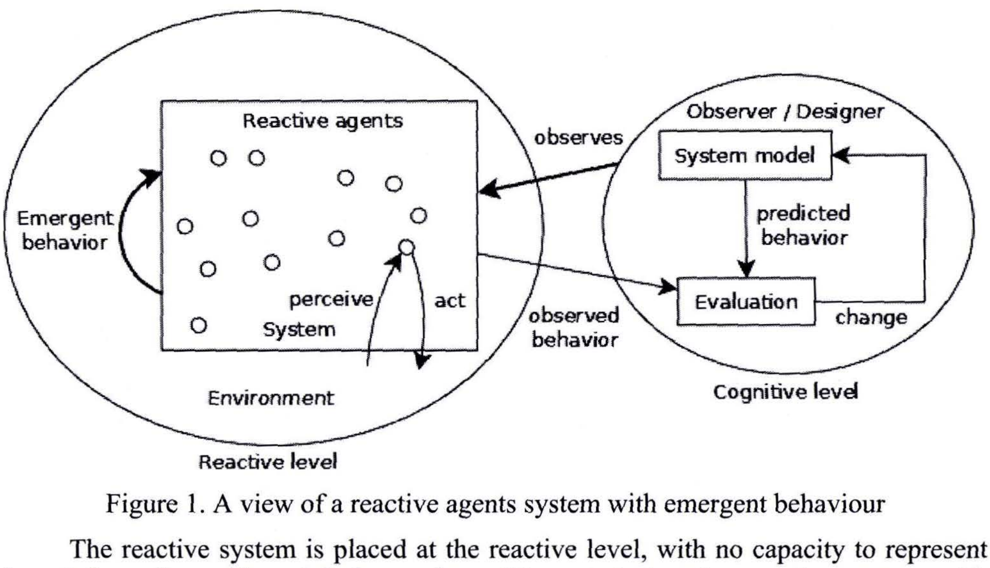

Figure 1. A view of a reactive agents system with emergent behaviour

The reactive system is placed at the reactive level, with no capacity to represent knowledge of a self model. It can be safely stated that the reactive system with emergent behaviour has *no anticipatory capabilities*, as it can not take into account its emergents while continuing to function. The designer or observer is placed at the knowledge level. Based on the system model, he/she is able to predict the emergent behaviour of the reactive system; however, as with any prediction, he/she may be wrong.

Some comments have to be made. First, it is supposed that the system is observable by the designer. Second, it is supposed that the designer has full control over the agents in the society and that he/she is able, by modifying the model, to alter the agents' behaviour accordingly. This last supposition is true in general, as systems with reactive agents are usually closed, conceived and implemented by one designer.

In what follows, we shall illustrate the ideas for a simple model of reactive agents, namely the "Game Of Life" Cellular Automaton [5] - *CA*. The *CA* is specified as a grid in which each cell can be "alive" or "dead." Alive cells die if the vicinity is under populated or overcrowded. Dead cells become alive if the number of alive neighbours is just right. In a reactive agent system modelling the *CA*, each agent manages one cell in the environment - a set of $m \times n$ cells in a state of "alive" or "dead." An agent at position $x; y$ is able to perceive, change, and broadcast the state of its cell. Using the general definition of reactive agents given in the previous section, we can describe the game of life in the following way:

$CA = (E, P, Ac, M, \text{perceive, act, interact})$
$E = \{0, 1, \dots n\} \times \{0, 1, \dots, m\} \times \{\text{dead, alive}\}$
$P = \{\text{dead, alive}\}$
$Ac = \{\text{die, live}\}$
$M = \{0, 1\}$

The current behaviour of the *CA* system may be specified by a set of simple rules, implemented in each reactive agent – the rules are labelled according to whether they are rules for being or becoming alive (RA), or dead (RD), respectively.

RA_1: *if* $\text{perceive}_{x,y} = \text{alive}$ *then* $\text{interact}_{x,y} = 1$
RA_2: *if* $\text{perceive}_{x,y} = \text{alive}$ *and* $s = 2$ *then* $\text{act}_{x,y} = \text{live}$
RA_3: *if* $s = 3$ *then* $\text{act}_{x,y} = \text{live}$

RD_1: *if* $\text{perceive}_{x,y} = \text{dead}$ *then* $\text{interact}_{x,y} = 0$
RD_2: *if* $\text{perceive}_{x,y} = \text{dead}$ *and* $s < 3$ *then* $\text{act}_{x,y} = \text{die}$
RD_3: *if* $s > 3$ *then* $\text{act}_{x,y} = \text{die}$

$$\text{where } s = \sum_{x,y.neighbor} \text{interact}_{x,y}$$

In the case of the *CA*, the most quoted emergents are the patterns formed by living cells, the typical ones being the gliders - dynamic structures of cells that advance across the grid. However, there is another emergent that is not usually mentioned: the conservation of the cell population. Depending on the rules for the cells and on the initial number and distribution of the cells, the population may eventually become extinct, may remain constant in size or may grow to cover the whole grid, uniformly. Other significant factor that may alter the system emergent behaviour is the predefined order of applying the rules (rule selection strategy). In the example above, if $s = 3$, both RA_1 and RA_3 can be applied if the agent is alive, or RD_1 and RA_3 if the agent is dead.

The order of applying these rules has influence on the emergent properties and can be altered by the designer only.

Therefore, the emergent properties are not implemented in the cells, but they result from the rules and associated parameters, like the initial distribution, threshold for living or dying, etc. All those parameters are fixed by the observer/designer, who holds the predictive power over the system behaviour. If the observed behaviour is not similar to the predicted one, the designer will change the rules, rule strategy and/or initial parameters.

4. Emergent behaviour in cognitive agent societies

4.1 Emergence in cognitive agent systems

In reactive agent systems, the agents are simple computing units, with no capacity to represent knowledge or reason about their behaviour. In such a system, the emergent is strongly related to the possible (simple) actions of agents and to the (simple) interactions between agents. If the system is formed of agents that can move around, change state, and are attracted or rejected by other agents, the result is a certain structure formed by the agents, with a certain distribution of agents in a certain state. If the system contains agents that can move and have a notion of direction, the result may be a certain behaviour and a global invariable direction in a certain space.

Compared to reactive agents, *cognitive agents* have several additional features: capacity to represent, at a symbolic level, knowledge, beliefs, plans, and goals or desires, ability to reason, act autonomously and pro-actively for accomplishing their goals. In the cognitive case, instead of simple properties like position and state, agents hold knowledge and belief bases, which usually contain information about the agent self, the environment, and other agents (called sometimes acquaintances), and are capable of high level interactions, such as speech-act performatives, negotiation, cooperative plan formation, etc.

In our previous work [14], we have analysed the emergent capabilities of cognitive multi-agent systems and we have shown that such systems are capable of complex emergent behaviour, based and triggered by the mere high-level properties that they have. The emergents in cognitive multi-agents systems must be related to these properties: some form of structures related to agent's knowledge and beliefs, plans and goals of the agent or of the acquaintances, structures that are formed based on the high-level interactions existing in the agent society. In a cognitive multi-agent system, there are several aspects of self-organisation. First, agents may group or disperse according to the knowledge or beliefs they have. Second, agents may exchange knowledge. Several types of behaviour may emerge according to the type of exchanged knowledge (see [14] for details). As cognitive agents have explicit goals (or desires), they know what they intend to achieve. These goals define the final target of an agent's sequence of actions. Interaction based on the agents goals leads to the essential elements of collaboration or competition – working together for a common goal or working separately towards achieving better individual or group performance.

As reactive agents have rules, cognitive agents usually feature plans or a plan library, used for building plans in order to attain goals. These plans are what characterises the behaviour of an individual agent, as plans lead to the establishment of immediate intentions that are translated into agent's actions. Being able to exchange parts of the plan libraries, agents will, in fact, transfer behaviour. As plan templates change as a result of past experiences, behaviour transfer leads to better agents that might know how to solve a problem even if it is new to them, because other agents have provided them with the solution to that sort of problem.

Therefore, based on highly structured properties, but using the same principles as in the reactive case, cognitive agents are able to manifest advanced emergent behaviour. Agents, although cognitive, may have limited knowledge and perception about the whole system - however, by using local goals and behaviour, complex and higher level features may emerge at the system level.

4.2 Predictive ability of cognitive emergent behaviour

An agent society formed of cognitive agents with emergent behaviour can be seen as having anticipatory abilities at two levels: the level of the *observer* and the level of the *cognitive agent* (Figure 2). At the level of the observer, the view of the system is similar with the one presented in Section 3.2 (Figure 1). The designer is holding the system model, in this case a complex and more elaborated one than in the case of reactive agents systems, and is able to modify/tune the model in order to bring as close as possible the observed behaviour to the predicted behaviour. As opposed to a reactive system, in case of a cognitive agent society, the system may be open, namely agents can enter of leave the society at any time during the system functioning and, moreover, can be designed and implemented by different bodies/designers. The *openness of the system* has a significant influence on the ability of a designer to control the predicted behaviour. Suppose that the designer is able to perceive the entire system behaviour (the behaviour of all agents and the emergents resulting from the agent interactions). If the emergents are not conform to his/hers predicted behaviour, the designer can only modify the functioning of the agents he designed, which may not always be all the agents in the system.

The second level of predictive abilities is, in case of cognitive agents societies, related to the agents themselves. A depicted in Figure 2, a cognitive agent has a model of the system: a *self model* and an *acquaintances model* (agents he knows in the system and with which it can interact). This model may not be a true model of the system or of the world, as it is based on the agent's beliefs about the world; however, the agent is capable of using this model and try to predict both its behaviour and the behaviour of the other agents, predictive capability which, we claim, is essential if the system is able to produce emergents. A cognitive agent is able to detect and exploit in a predictive manner the emergent of the system or, at least, the emergent behaviour from its self point of view.

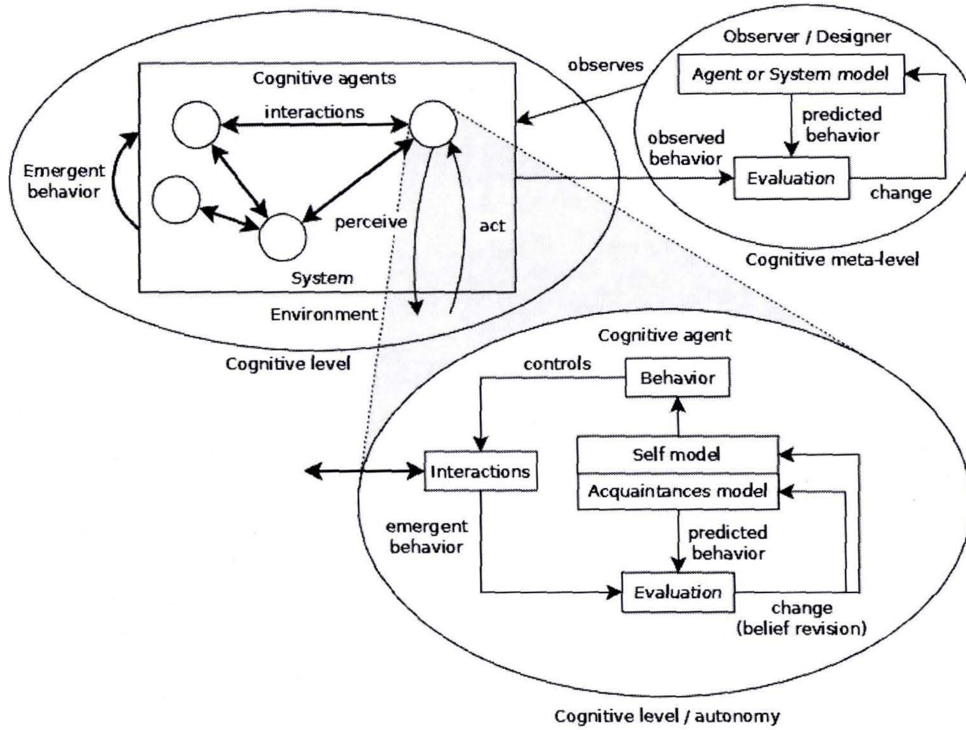

Figure 2. A view of a cognitive agent system with emergent behaviour

In order to illustrate our view, we consider one scenario of emergent behaviour for establishing cooperation relationships in a cognitive multi-agent system with self-interested agents. In this system, the agents try to establish preferred cooperation relationships with some other agents in the system, based on previous interactions. A task allocated to one agent may be a complex task, to be split into simple tasks by the planning component of the agent (the agent can use a plan library). A simple task requires a certain *competency* to be solved. Each agent has a set of competencies, a subset of the entire competence set required to solve all tasks. When receiving a task to solve, an agent decomposes it into simple tasks and makes a plan, choosing which tasks to be solved by itself and which tasks to distribute to other agents with which it collaborates.

We will not focus here on how tasks are decomposed and planned, but rather on how an agent chooses its collaborators. This decision is based on previous interactions. For a certain amount of time, the agent remembers which agents it collaborated with, for what types of tasks, and what was the result of the collaboration. This is where prediction comes in. Previous experience of the collaboration with another agent may give an indication of how likely it is for that agent to accept a new request.

Our scenario may be described as follows:

- *World (system) model* – available to the observer

  $T = \{T_1, \ldots, T_p\}$ – set of tasks

  $A = \{A_1, \ldots, A_n\}$ – set of agents

  $C = \{C_1, \ldots, C_m\}$ – set of all required competencies to solve tasks in T

  TaskC: $T \to C$ – function to map simple tasks to required competencies

  AgentC: $A \to P(C)$ – agent's competencies (an agent may have several competencies)

- *Agent self and acquaintances model* – available to the agent

  Self-beliefs

  SelfCap = $\{C_{i1}, C_{i2}, \ldots\} \subset C$ – self model about its own competencies

  $$\mathrm{SelfCap}_{Ai} = \mathrm{AgentC}(A_i)$$

  World beliefs

  AgentC*: $A \setminus \{A_i\} \to P(C)$ – model of the other agents in the system (this function has a * because it may not be identical to the function AgentC in the world model – the agent may sometimes hold wrong beliefs).

Any agent can communicate directly with any other agent in the system, i.e., the agents can be considered as nodes in a complete graph. The messages exchanged by the agents are of several types: *Request(task, deadline)* – request for performing task within deadline; *modifyRequest(task, deadline)* – reply with a modified request, that is a task will be completed within a different deadline than initially requested; *Accept(task, deadline)* – acceptance to complete a task within the requested deadline; *Reject(task, justification)* – rejection of the request for a task, given the justification; *Reject(task, justification, recommendedAgent)* – rejection of the request for a task, also containing a recommendation for another agent that might be able to complete the task instead.

Initially, the agents do not have any experience at all, so they randomly contact other agents and request them to do tasks. The replies are favourable or not, and the profiles of the other agents start to take shape. The replies may be considered as feedback for the links between agents. Although the activation of links is initially chaotic, over time the system self-organizes and collaboration patterns emerge. That is, each agent will prefer a certain agent(s) over all the others for a certain type of tasks. This preference is recorded in the agent's beliefs as a recent history of successful requests regarding a certain type of task sent to the other agent. When a new request is carried out, the reply will be scored with a number R in the interval [0, 1] characterizing the success of the request. A rejection will have a score of 0. Acceptance scores 1. Between them, there are modifications and rejections with recommendations. When the reply is received, the profile of the other agent is updated for the capability needed by the requested task. In order to achieve this, each agent keeps an acquaintances' profile, which is updated according to the equation (1):

Acquaintance profile P: $A \times C \rightarrow \mathbf{R}$

$$P_{Ai}(A_j, C_t) = P_{Ai}(A_j, C_t)(1 - \alpha) + \alpha R(C_t) \quad (1)$$

that is, the profile that agent A_i has on agent Aj (to which it sent the request) for the competency C_t needed by the requested task, is updated with the score R for the last request, keeping the old indication as a fraction (1 - α), $0<\alpha\leq 1$, where α is the adaptation or the prediction speed.

When a new request will be carried out for a task needing C_t, the profiles corresponding to this competency will be considered as a good indication of how an agent will respond to this kind of request. Therefore, by anticipating the response, the requesting agent will make the request to the agent with the best profile for that competency. The individual success of each request may be measured as a function of the number of trials needed for the request to be accepted by another agent, and of the modification that have been performed on the request in order to get it accepted. A global measure of satisfaction may be computed for each agent, as a function of the individual success measures for the requests made over a certain period of time. As the system stabilizes and the collaboration patterns between agents become clear, the satisfaction measures for the agents will also stabilize to certain values.

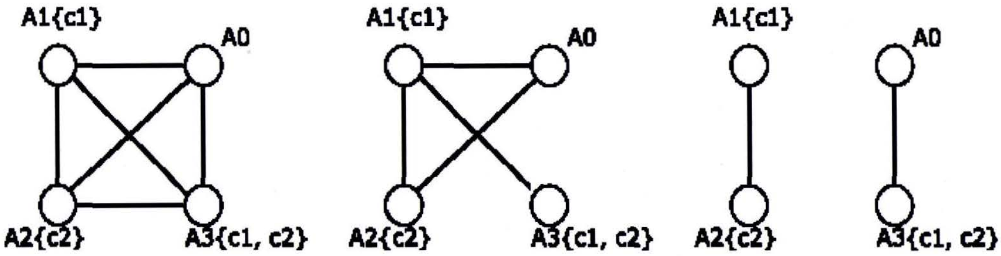

a) b) c)

Figure 3. The system in a) initial state, b) stable state, c) after the stabilization following a perturbation.

Figure 3 presents an intuitive view of the described interaction and belief update of the agents. Stronger line segments means stronger collaboration profiles, while faded line segments weaker collaboration. The system in Figure 3 is composed of 4 agents, A0 to A3, having the following competencies: A0 has no competency, A1 has competency c1, A2 has competency c2, and A3 has both competencies c1 and c2. Initially, agents have no special preference for any other agent. As agents start to receive tasks, plan and make requests for task execution to other agents, profiles start to take shape. For example, A1 can do by itself tasks involving competency c1 and discovers that it can use A2 for tasks involving c2. As tasks involving c2 are more difficult, A2 may refuse them more often, so A1 sometimes makes requests to A3. A2 uses A1 for tasks involving c1. A3 can do all tasks by itself. A0 cannot do any task by itself, however it discovers that A1 and A2 respond favourably to requests regarding competencies c1 and c2. A0's profile for A3 is not so favourable. One reason would be because A0 discovered A1 and A2 first and their response is satisfactory. Another reason may be that A3 responded negatively for a few times in a row. Nevertheless, it just happens that the initial experience favoured A1 and A2, and the system started to self-organize in this fashion. Consequently, every time A0 needs to make a request, it makes it to one of A1 or A2 (depending of the required competency).

The designer may choose to modify the behaviour of the agents he owes in the open system. One possible modification may be to finally tune the profile update equation (1) to equation (2), making the agent a "reinforcement learning"-like agent, where δ is the discount factor.

$$P_{Ai}(A_j, C_t) = P_{Ai}(A_j, C_t)(1 - \alpha) + \alpha \cdot (R(C_t) + \delta \max_{c_k, AgentC^*(A_j)} P_{Ai}(A_j, C_t)) \quad (2)$$

When the self-organization of the system stabilizes to some preferred collaboration relationships, the desires of the agents may not be in conformity with their actual measure of satisfaction. The reason is that the system might have stabilized in an attractor that is less fit to the external conditions, that is, the collaborative patterns that have emerged, although stable, are not the best way to collaborate given the top-level tasks that the agents have to solve. In order to pull the system out of the attractor it is in, the system must be perturbed strong enough to reach the basin of a different attractor, hopefully a better one. This may be done by the agents themselves: given a stable state, an agent that is unsatisfied with its collaborators may perturb its local, stable, collaborative pattern. That can be done by not choosing the agent that experience says it is the best, but by choosing another one, possibly randomly. Several situations might result. If the perturbing agent is the only one in the system that is unsatisfied, it is most likely that its action will receive negative feedback and the stable pattern will be restored. However, if several agents are in the same situation, potentially in the same area of the system, their "random" requests may destabilize the local pattern and lead, by means of positive feedback and cascade reaction, to a reorganization of a larger area of the system and, potentially, to the reorganization of the whole system. Starting from equation (2), the agent may be endowed with the capacity to switch from exploitation - equation (3) of selecting the best collaborator, to exploration, by randomly selecting an agent if it is not satisfied. Alternately, the agent may be designed from the beginning to achieve a certain balance between exploitation and exploration, as in equation (4).

Exploitation

$$\text{SelectAgent}_{Ai}(C_t) = \arg\max_{0 \le j \le n} P(A_j, C_t) \quad (3)$$

Exploitation and exploration

$$\text{SelectAgent}_{Ai}(C_t) = \text{explore}(\arg\max_{0 \le j \le n} P(A_j, C_t), \text{freq}(A_j, C_t)) \quad (4)$$

For the example presented in Figure 3, as each of A1 and A2 are working for their own tasks, for the other's tasks, and also for the tasks of A0, none of the three is very satisfied with the performance. Indeed, the external observer would notice that A3 is only working for its own tasks. Therefore, one of A0, A1, or A2 will decide to perturb the system so that it will go out of the attractor it is currently in. Suppose it is A0 that first takes the decision. At some point where it needs to make a request for a task involving competency c1, it no longer makes the request to A1, but to A3. A3 accepts and its profile in A0's table is improved. As A3 will be busier now, it might reject requests from A1, so more requests from A1 regarding competency c2 will now go to A2. If the perturbation was strong enough, the system will enter another attractor where agents A1 and A2 only work for each other, and agent A3 will work for A0.

5. Conclusions

We have presented an overview of artificial agent societies with emergent behaviour, starting from several definitions of emergence that exist in the literature, and have highlighted that anticipation and the predictive ability of such systems are essential properties, especially from the point of view of satisfying the desired complex behaviour expected to emerge. The main contribution of this paper is the analysis and insight achieved by identifying several levels of predictive capabilities in such systems, putting thus a ground for the system designer.

In both reactive and cognitive systems, the predictive ability is in the "power" of the system designer, which holds the system model. In case of a simple reactive agent system, the agents do not have anticipatory capabilities by themselves when emergence is considered. As apposed to the reactive case, for a cognitive agent system with emergent behaviour, we have shown that the cognitive agent itself, based on its high-level representation and reasoning capabilities, is able to detect and exploit in a predictive manner the emergent of the system or, at least, the emergent behaviour from its self point of view.

Our future work is mainly focus on developing and testing several scenarios, for both the reactive and cognitive case, based on the models presented in this paper. Another interesting direction to follow is to consider the wide spectrum of currently existing artificial agent societies, in which the agents are either heterogeneous, i.e., a mixture of reactive and cognitive ones, or span all hybrid levels of design between reactive and cognitive.

References

1. L. Amaral and J. Ottino, "Complex networks: Augmenting the framework for the study of complex systems", The European Physical Journal B-Condensed Matter, vol. 38, no. 2, pp. 147-162, 2004.
2. G. Beurier, O. Simonin, and J. Ferber, "Model and simulation of multi-level emergence" Proceedings of IEEE ISSPIT, pp. 231-236, 2002.
3. C. Bourjot, V. Chevrier, and V. Thomas, "A new swarm mechanism based on social spiders colonies: From web weaving to region detection", Web Intelligence and Agent Systems, vol. 1, no. 1, pp. 47-64, 2003.

4. T. De Wolf and T. Holvoet, "Emergence versus self-organisation: Di_erent concepts but promising when combined", Engineering Self Organising Systems: Methodologies and Applications, vol. 3464, pp. 1-15, 2005.
5. M. Gardner, Mathematical games: The fantastic combinations of John Conways new solitaire game life. Scientific American 223 (4), pp. 120-123, 1970.
6. C. Gershenson, "Complex philosophy", Proceedings of the 1st Biennial Seminar on Philosophical, Methodological & Epistemological Implications of Complexity Theory, pp. 1-7, 2002.
7. M. Gleizes, V. Camps, and P. Glize, "A theory of emergent computation based on cooperative self-organization for adaptive artificial systems", in Fourth European Congress of Systems Science, 1999.
8. J. Goldstein, "Emergence as a construct: History and issues", Emergence, vol. 1, no. 1, pp. 49-72, 1999.
9. F. Heylighen, "The science of self-organization and adaptivity", The Encyclopedia of Life Support Systems, pp. 1-26, 2002.
10. G. Lewes, Problems of Life and Mind, vol. 2, K. Paul, Ed. Trench, Turbner, and Co, London, 1875.
11. K. Mainzer and R. Landauer, "Thinking in complexity: The complex dynamics of matter, mind, and mankind", American Journal of Physics, vol. 63, p. 766, 1995.
12. M. Mamei and F. Zambonelli, "Spatial computing: the TOTA approach", Self-star Properties in Complex Information Systems Conceptual and Practical Foundations, vol. 3460, pp. 307-324, 2005.
13. J.-P. Mano, C. Bourjot, G. Leopardo, and P. Glize, "Bio-inspired mechanisms for artificial self-organised systems", Special Issue: Hot Topics in European Agent Research II Guest Editors: Andrea Omicini, vol. 30, pp. 55-62, 2006.
14. A. Olaru and A. Florea, "Emergence in cognitive multi-agent systems", Proc. of CSCS17, 17th International Conference on Control Systems and Computer Science, Bucharest, Romania, pp. 515-522, 2009.
15. T. Oren and F. Longo, "Emergence, anticipation and multisimulation: bases for conict simulation", Proc. of EMSS 2008, 20th European Modeling and Simulation Symposium (Part of The 5th International Mediterranean and Latin American Modeling Multiconference), pp. 546-555, 2008.
16. G. Picard and F. Toulouse, "Cooperative agent model instantiation to collective robotics", in Engineering Societies in the Agents World V: 5th International Workshop, ESAW 2004, Toulouse, France, October 20-22, 2004: Revised Selected and Invited Papers. Springer, 2005.
17. M. Randles, H. Zhu, and A. Taleb-Bendiab, "A formal approach to the engineering of emergence and its recurrence", be presented at EEDAS-ICAC, pp. 1-10, 2007.
18. R. Rosen, Anticipatory systems. Pergamon Press New York, 1985.
19. A. Ryan, "Emergence is coupled to scope, not level: Research articles", Complexity, vol. 13, no. 2, pp. 67-77, 2007.

20. G. D. M. Serugendo, M.-P. Gleizes, and A. Karageorgos., "Self-organization and emergence in MAS: An overview", Informatica, vol. 30, no. 1, pp. 45-54, 2006.
21. R. Standish, "On complexity and emergence", Arxiv preprint nlin. AO/0101006, pp.1-6, 2001.
22. C. Unsal and J. Bay, "Spatial self-organization in large populations of mobile robots", Proceedings of the 1994 IEEE International Symposium on Intelligent Control, pp. 249-254, 1994.