HAL open science

Emergence in Cognitive Multi-Agent Systems

Andrei Olaru, Adina Magda Florea

To cite this version:

Andrei Olaru, Adina Magda Florea. Emergence in Cognitive Multi-Agent Systems. CSCS17, the 17th International Conference on Control Systems and Computer Science, MASTS Workshop, May 2009, Bucarest, Romania. pp.515–522. ⟨hal-00642555⟩

HAL Id: hal-00642555

https://hal.science/hal-00642555v1

Submitted on 18 Nov 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Autorisation HAL

HAL Authorization

Emergence in Cognitive Multi-Agent Systems

Andrei Olaru * **Adina Magda Florea** **

* *Computer Science Department, University Politehnica of Bucharest, Splaiul Independentei 313, 060042 Bucharest, Romania (e-mail: cs@andreiolaru.ro).*
** *Computer Science Department, University Politehnica of Bucharest, Splaiul Independentei 313, 060042 Bucharest, Romania (e-mail: adina@cs.pub.ro).*

Abstract: Emergence is a hot topic of the present research in the domain of complex systems, especially connected with the modelling and simulation of large multi-agent systems. Emergent properties have been obtained in systems inspired from nature – ants, wasps, etc – but also in simple artificial systems like Cellular Automata – obtaining gliding structures – and Transporting Agents – obtaining traffic directions. Most of the systems designed to manifest (truly) emergent properties are made of reactive or mostly reactive agents, with a limited set of actions that they apply as a direct consequence of certain stimuli. As more computational power is available and cognitive multi-agent systems tend to become a mature technology, the study of emergence in the context of such systems is a promising direction of development. Starting from the analysis of emergence in reactive systems, this paper is an attempt to define the different forms and requirements of emergent properties in the context of cognitive multi-agent systems and highlights different settings in which such properties may significantly enhance the overall performance of the system.

Keywords: Multi-agent systems, Emergence, Cognitive agents, Self-organisation

1. INTRODUCTION

Emergence is not a new subject. The concept of "the whole before the parts" was introduced by the ancient Greeks to explain those properties that do not result from the mere addition of the features of the parts, but seem to magically arise when the parts are put together and interact, forming a system.

The subject of emergence has been discussed over time from different points of view (Goldstein (1999)), especially in the context of complex system analysis (Amaral and Ottino (2004)), as emergence is what gives the system properties and behaviour that are not obvious or expected when looking at the individual parts. More important, complex systems cannot be analysed from a purely mathematical point of view, as the number of parts is very large and the equations describing the system are non-linear, therefore a mathematical solution is virtually impossible.

Several definitions of emergence exist (Boschetti et al. (2005)) but none is yet generally accepted. Many times emergence is defined by its effect – the formation of patterns in the structure or behaviour of certain systems.

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 defined 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 (Heylighen (2002)) – the organisation of the system emerges dynamically from the interactions between the agents, with no external or centralised control, resulting in organisation that is robust and redundant.

So far, self-organisation and emergence have been studied mainly in the context of multi-agent systems formed of reactive agents – agents that, using simple rules, react to external stimuli, with very limited reasoning, holding almost no knowledge about themselves, their neighbours or the environment. Reactive agents are used for the reason that they are simple and can be easily implemented on devices with limited capacity. By means of emergence and self-organisation, the function they fulfill as a complex system is more complex. The emergents in this case are structures or structured behaviour that the agents organise into.

Nowadays, the capabilities of even very basic computing devices have considerably increased, allowing for a much more complex internal structure for agents – the possibility to hold reasonable amounts of data and to have a more nuanced behaviour. Cognitive agents have knowledge about the surrounding environment, have goals they desire to fulfill, make plans and take action in order to fulfill them. The purpose of this paper is to discuss the concept of emergence in the context of multi-agent systems formed of cognitive agents. The main question that arises is what properties can emerge in this new context, properties that could not emerge in a system formed of reactive agents.

The paper is organised as follows. Section 2 is dedicated to the existing definitions of emergence. Section 3 gives examples of how emergence is used in several reactive agent systems. The improvements of cognitive agents relative to reactive agents are presented in Section 4. Section 5 discusses the main topic of this paper: emergence in the context of cognitive agent systems. The last two sections are dedicated to an example and conclusions.

2. DEFINITIONS OF EMERGENCE

Throughout recent literature there have been many attempts to define emergence; however, to the moment, none of the definitions is generally accepted. What we know is that emergence appears in the context of complex systems (Amaral and Ottino (2004)) – systems composed of a large number of interacting individual entities. 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. 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" (Standish (2001)).

A more elaborated definition is that "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" (De Wolf and Holvoet (2005)). In the definition above, the "emergent" is a general concept that denotes the result of the emergence, that can be a property, a structure, a behaviour or some other phenomenon. The problem with this definition is the notion of novelty, which is not well defined and may leave room for interpretation.

An emergent may be novel the first time it is observed, but the novelty fades as the emergent is manifested again, although the emergent does not change. In fact, the novelty is represented, besides the element of surprise (the emergent is not expected by the designer of the system), by the impossibility to explain the emergent phenomenon knowing the properties and features of the individuals composing the system. However, the ability to explain the emergent is not invariable either, as, when one day the emergent will be explained, it will no longer be unexplainable, but, if genuine, it will still remain the result of emergence (Boschetti et al. (2005)).

The essence in the criteria above is the difference between the language used for the description of the individuals and the language used for the description of the system (Standish (2001)). The microdescription (the description of the entities) is clear, as it is part of the design of the individual parts of the system. The macrodescription however is difficult to define. First, it needs an observer (Randles et al. (2007)) that is capable of detecting the features of the system and from whose perspective the macrodescription can be created. Second, if the microdescription is clear once the individual entities have been designed, there may be more that one macrodescription to describe the resulting behaviour of the system (Gershenson (2002)), so there is a need to find a reasonably good theory (Standish (2001)) – that is equally explanatory and predictive.

Considering the observations above, a better definition of emergence (Beurier et al. (2002)) would be that an emergent is, in the context of an interacting set of agents whose dynamics are expressed in a vocabulary D, a global phenomenon – static or dynamic, but nevertheless invariant – that is observed by the agents or by an external observer and can only be interpreted in a vocabulary D' that is different from D.

Other approaches (Boschetti et al. (2005)) view emergents as those derived phenomena that increase the predictability of the system, relative to the process or dynamics they derive from (Shalizi (2001)). This view is especially interesting as it does not define emergence relative to an observer, but as a property of the system.

It is important to point out some properties of emergent phenomena (De Wolf and Holvoet (2005)). Emergence occurs in the context complex systems – out of the interaction between a large number of parts. If there is no interaction, emergence cannot occur. Emergence can only be defined in relation with two levels – micro and macro. It is manifested at the higher level, arising from the interactions at the lower level. Also, the individuals do not have any explicit representation of the emergent phenomenon, i.e., the emergent phenomenon cannot be reduced to the specification of the individuals. Although emergents arise only as the systems evolves, once they exist they maintain a certain identity over time (Goldstein (1999)). Emergents arise without any centralised or exterior control and the emergent phenomena is robust and flexible, i.e., it is not influenced by damage on the system (Heylighen (2002)).

To our work, 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 (Gleizes et al. (1999)). The difficulty rests in determining how an emergent function can be obtained and how it can lead to obtaining the desired emergent function (Serugendo et al. (2006)).

3. 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 the one hand, they are inspired by natural systems composed of simple individuals that act reactively and interact mainly by means of their environment (Mano et al. (2006)) 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 (Mamei et al. (2004)).

A reactive agent system may be formally defined as a tuple

$RAs = (E, P, A, M, perceive, act, interact, change)$

$perceive : E \rightarrow P$

$act : P \times M\ [\times M \times ... \times M] \rightarrow A$

$interact : P \times M\ [\times M \times ... \times M] \rightarrow M\ [\times M \times ... \times M]$

$change : E \times A \times A \times ... \times A \rightarrow E$

where E is the environment, P the possible stimuli an

agent can get from the environment, A the actions an agent can apply and M the messages agents can send to each other. Direct messages may not necessarily be used.

Some systems also use a very simple set of states S, and the domain of the *act* and *interact* functions changes from $P \times M$ and P to $P \times S \times M$ and $P \times S$, respectively. The state is calculated by

$state : S \times P \times M\ [\times M \times ... \times M] \rightarrow S$
as a function of previous state and the current stimulus.

In many previous studies (Bourjot et al. (2003); Beurier et al. (2002); Mamei et al. (2004); Picard (2005); Randles et al. (2007); Unsal and Bay (1994)), agents are capable of moving according to very simple laws and in reaction to simple stimuli. The result is usually the formation of a certain geometrical or geometry-related structure or behaviour: arrangement of agents in a circular or almost circular shape (Beurier et al. (2002); Mamei et al. (2004)); detection of areas in an image (Bourjot et al. (2003)); gathering of resources in a single area (Randles et al. (2007)); foraging of food or transportation of loads (Unsal and Bay (1994)); emergence of a traffic direction (Picard (2005)). These types of emergents are enumerated in Table 1.

Table 1. Examples of emergents in reactive multi-agent systems

Agent capabilities	Emergent property
fixed, binary states	conservation, gliders
mobile, binary states, attraction and repulsion	shape formation
mobile, multiple states, attraction and repulsion	multi-level shape formation
mobile, reinforcement	area coverage
mobile, transportation	accumulation of resources
mobile, transportation limited transit space	emergence of a traffic direction

In all the examples above a certain structure is formed, in space or in both space and time (as in the case of dynamic, stable, behaviour). The structure is emergent because the concept of structure – layout or order among multiple agents – is not defined explicitly in the description of individual agents. The emergent structure arises out of the movement of agents and/or their change of state. Changes of state and movement of an agent usually occur as a result of the state of neighbour agents and the attraction or repulsion to neighbour agents. Therefore, the structure arises out of certain tensions between individuals. The same happens with inanimate natural complex systems (Gleizes (2004); Heylighen (2002)): due to the forces between the parts of the system, the parts organise themselves into a certain structure.

Many examples of emergence in multi-agent systems relate to the notion of self-organisation, and the two concepts are often considered synonymous because of common features like decentralised control, flexibility and robustness (Heylighen (2002)). Although the two concepts are different, few examples of emergence without self-organisation exist (De Wolf and Holvoet (2005)) and none of them is in the field of multi-agent systems.

In reactive agent systems the emergent organisation is of 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 (Mamei et al. (2004)). Still, explicit structure and organisation awareness would indeed need agents with better capabilities of representation and memory.

In conclusion of this section, considering the above mentioned work, most present multi-agent systems use emergence to obtain robust spatial structures or space-related behaviour by using simple – reactive – agents, which are only capable of actions in response to the conditions in their immediate environment and/or neighbourhood. The emergent organisation is not of a different nature than the properties and behaviour of the agents, but it is of a higher level.

4. COGNITIVE AGENTS VERSUS REACTIVE AGENTS

An agent is considered to be an "encapsulated computer system that is situated in some environment, and that is capable of flexible, autonomous action in that environment in order to meet its design objectives" (Jennings and Wooldridge (1999)). According to this definition, an agent is capable of acting autonomously according to the objectives that it pursues. Therefore, a cognitive agent is considerably closer to this definition, as a reactive agent may have no explicitly specified objectives and its actions may not be meaningful in the absence of a larger system of agents surrounding it. Moreover, reactive agents have no plans, their actions being immediate responses to stimuli.

There are many advantages that a cognitive agent has over a reactive agent. First, it is proactive. Even if there are no signals, perceptions or stimuli from the environment, a cognitive agent may act by itself, taking action according to its objectives.

Second, a cognitive agent is aware of its situation and may reason about it. It is aware of what it is suppose to fulfill as final goal and is capable of making plans and taking action towards the realisation of its goal. As opposed to a reactive agent that acts blindly according to its simple set of rules, the cognitive agent can use its experience and information about the environment and the consequences of past actions to develop a better plan every time a similar situation occurs.

Third, cognitive agents are able to recognise when the objectives are completed. Usually, in the case of reactive agent systems, only an external observer can have an overall perspective on the state of the system, as the agents continue to act without knowing if their function has been fulfilled.

Structurally, if a reactive agent contains only a set of simple rules about how it should react to external stimuli (and, in some examples, a simply defined state), a cognitive agent has a more complex internal structure. A common model for cognitive agents is the BDI model of agency, based on the three components a cognitive agent should have: Beliefs about the environment and, possibly, about other agents; Desires or goals that it wants to achieve; Intentions, or actions that, according to the plans it develops, it will execute and will potentially lead to the realisation of its goals (Rao and Georgeff (1995)).

The BDI architecture considers agents with the following structure:

$Agent = (Beliefs, Desires, Intentions, Plans, Percepts, Actions)$

Usually, $Plans \subseteq Beliefs$ and a certain indication of $State$ is also likely to be included. The functioning of the agent is governed by the functions:

$revise_beliefs : Beliefs \times Percepts \rightarrow Beliefs$, that revises beliefs according to recent perception;

$get_options : Beliefs \times Desires \rightarrow Intentions$, that decides on the appropriate goals to follow next;

$plan : Intentions \times Beliefs \times Subplans \rightarrow Plans$, that builds plans towards achieving the desired goals;

$act : Plans \times State \rightarrow Actions$, that executes the plans as sequences of actions.

5. EMERGENCE IN COGNITIVE AGENT SYSTEMS

Reactive agent systems use emergence to obtain, as a system, a result that is of higher level than the result that can be obtained by any of the agents, individually. The difference in level is the outcome of the interactions between the agents. The global function is obtained without the knowledge of the agents, and, because they do not need to know about (and therefore be able to conceive) the global function, the agents can have a simpler implementation and require lower computational capabilities.

In reactive agent systems the emergent is strongly related, in nature, to the possible actions of agents and, more important, to the interactions between agents: if the system is a Cellular Automaton, where cells have two states (active or inactive), then the result is a certain invariable structure or sequence of active cells (as in the glider); 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; finally, 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: knowledge and experience, plans, objectives. The emergents in cognitive multi-agents systems must be of this nature – some form of structures related to the knowledge, plans and objectives of agents, based on interactions comprising these elements.

5.1 Belief-related Emergence

Emergent organisations in reactive agent systems consist of a correlation in agent's position relative to a certain point, or a correlation between position and state (Beurier et al. (2002); Mamei et al. (2004)).

In the cognitive case, instead of simple properties like position and state, agents hold knowledge bases, that usually hold information about their environment and about other agents. There are two aspects of organisation. First, agents may group or disperse according to the knowledge they have. Second, agents may exchange knowledge. Several types of behaviour may emerge.

Attraction between agents with similar knowledge leads to groups that are dedicated to certain types of information. If the information in the system is on different domains, agent groups will hold specialised data. Information exchange will make agents have more diversified knowledge on their domain. If data from other domains is considered as useless and discarded, then each agent in the system will become a specialist in the domain of the group it is part of. If, however, information is on the same domain, agents might prefer collaboration with agents with similar views in the field, and groups will enter into a competition of different views on the same area of expertise. The result will be a system that will try to solve tasks in different manners, serving for comparison or redundancy.

A possible application of such a multi-agent system would be in social networking: agents, assigned to users, search for other users with as similar features as possible. They join groups that share their opinions and/or are interested in the same topics. Even if agents have, as objective, finding potential new friends for the user they represent, the global result is groups of users that share the same interest.

Repulsion between agents holding similar knowledge will lead to distribution. In this case, attraction towards agents with different knowledge might be used for facilitating information sharing, and then, after information is exchanged (should the agents consider the other agent's knowledge useful), repulsion will lead to the transportation of the copied data to another area of the system. The result will be distributed knowledge storage and replication.

For example, this behaviour may be used in a decentralised system for the distribution of some physical resource – food, materials, etc. Transporter agents find agents that have other resources than themselves, approach them, exchange resources that they consider of interest, and then go and distribute them at a further location.

In both cases above, the exchange and replication of information guarantees robustness of groups, in the first case, and of the whole system, in the second.

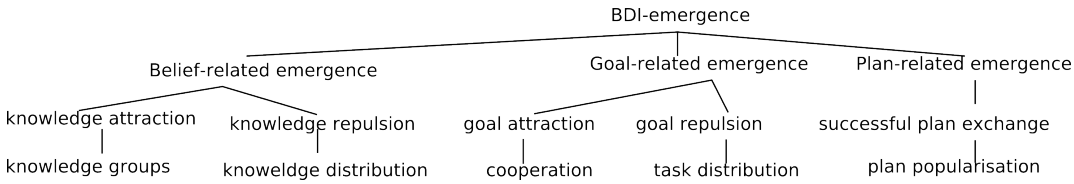

Fig. 1. Expected emergents in different contexts of emergence

Another possibility of cognitive self-organisation is when agents are not mobile entities, but confined to certain locations (like web servers), but are able to choose the agents that they communicate and collaborate with, for the purpose of solving certain tasks. As not all agents are the same (different servers have different capabilities and may be experience different loads), an agent might tend to choose for collaboration agents that it has collaborated with before, with good results. As features of the agents do not change all the time, collaboration patterns may arise, leading to more efficient global computation.

5.2 Goal-related Emergence

Cognitive agents have goals (or desires), and therefore know what they intend to achieve. These objectives define the final target of an agent's sequence of actions. By comparison, reactive agents, although they have rules, do not explicitly "know" what is the final goal which they are using the rules for.

Interaction based on the goals of the agents 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.

Attraction of an agent towards another agent with a similar objective, or with an objective that is considered as more worthy, leads to the forming of groups of agents that have common goals. A common goal allows for collaboration towards achieving that goal and a more efficient solution is obtained.

An example would be a task processing multi-agent system, where tasks have different priorities. Agents want to process tasks with higher priority. When an agent takes an urgent task, surrounding agents that work on tasks with lower priority will change their objectives in order to collaborate in the processing of this high priority task, that is more attractive to them. The emergent effect will be an improvement in the efficiency of task processing.

5.3 Planning-related Emergence

As reactive agents have rules, cognitive agents usually feature a plan library, used for building plans in order to attain goals. These sub-plans are what characterises the behaviour of an individual agent, as the 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.

Moreover, a measure of success may be added to each agent. Consequently, agents might prefer to copy the behaviour of more successful agents. Or, in case several plans for a certain objective exist, agents will keep the plan provided by the agent which is most successful, this way copying its behaviour and general approach to problems. Successful and, possibly, more experienced agents will become leaders and role models for the less-fortunate agents. Hierarchies will emerge, and in general the agents will act more efficiently, as they will use the best plans.

A foraging application is a good example. In a land full of obstacles, an agent's objective is to find food and bring it back to base. There are many paths towards food repositories. Some agents will find shorter paths, and they will return to base more quickly and therefore be rewarded with a higher score and, possibly, with some upgrades. When agents meet, they exchange information from their plan libraries (of course, they might also exchange knowledge about the location of food). Less successful agents will copy the partial paths of agents with higher score, resulting in a higher potential score for themselves, potential upgrades and thus better means to find food and fulfill their goal. Although the personal objective of agents is to find and retrieve food, the emergent will be, primarily, a hierarchy among agents (based on success), and, secondly, a reinforcement of the best paths. This result is similar to results obtained by stigmergy, however the knowledge spreads much faster.

5.4 BDI-Emergence

We conclude this section with some important observations. First, a cognitive multi-agent system may (and probably should) combine in the interaction between agents all three elements – beliefs, goals (desires) and plans (intentions) (several types of emergents are presented in Table 2 and Figure 1).

In order to obtain true emergents, the goals of the agents must not be the same as the desired global effect or behaviour. That is, the agents must be unaware of the global objective, whose fulfilment can only be acknowledged by an external observer. Moreover, emergence is used to obtain a property that could not be conceived (due to capability limitations and more simple implementation) by individual agents. The goals of the agent should not contain the global goal, nor should *Subplans* contain the algorithm for achieving it.

In short, emergence in cognitive multi-agent systems can be defined as the phenomenon of achieving a superior goal by using agents that interact following (different) individual objectives.

Table 2. Examples of emergents in cognitive multi-agent systems

Agent capabilities	Emergent property
mobile, exchange knowledge, attraction	knowledge groups
mobile, exchange knowledge, rejection	knowledge distribution
fixed, choice of collaborators	collaboration patterns
mobile, exchange of plans	reinforcement of best plans
fixed, attraction to other agents' goals	improved task processing

6. EXAMPLE

A simple example of cognitive multi-agent system will be considered and compared to a similar reactive multi-agent system - the "Game Of Life" Cellular Automaton (Gardner (1970)). The reactive agent system will be named CA and the cognitive agent system will be named KCA (Knowledge Cellular Automaton).

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.

$CA = (E, P, A, M, percieve, act, interact, change)$

$E = \{0, 1, ..., n\} \times \{0, 1, ..., m\} \times \{dead, alive\}$

$P = \{dead, alive\} \;\; A = \{die, live\} \;\; M = \{0, 1\}$

$perceive_{x,y} = s$, where $(x, y, s) \in E$

$interact_{x,y} = 1 \; if \; perceive_{x,y} = alive, \; 0 \; otherwise$

$change_{x,y} = act_{x,y}$

$act_{x,y} = die \; if (s = dead \wedge n < 3) \vee (n > 3)$

$act_{x,y} = live \; if (s = alive \wedge n = 2) \vee (n = 3)$, where $s = perceive_{x,y}$ and $n = \sum received_messages$.

In the case of the CA, the most interesting emergents are the patterns formed by living cells, the typical ones being the gliders – dynamic structures of cells that advance across the grid. However, more relevant to our example 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. These properties are not implemented in the cells, but they result from the rules and the initial distribution.

The KCA is specified as a grid that contains in each cell an agent with a certain data storage capacity. Each data element is initially injected in the system from the exterior, into a certain agent. Elements from the environment might also request data from the agents and the agents must try to fulfill the requests. A good implementation of a such system would be a web storage network, where the data is represented by individual web pages and the agents reside on web servers.

There are some specifications of the KCA. An agent can only share information with the agents in the 8 neighbouring cells. The goal of an individual agent is to store data, preferentially pieces of data with a certain degree of relation between them (for example the agents will prefer storing web sites of the same genre). There is a limit on the storage capacity of the agents. As each agent must always be able to receive data injected from the exterior, it might choose to discard some of the data it holds, in case its storage capacity is almost full. Discardable items are items that are not interesting to the agent and that are already stored by other neighbouring agents. Each agent occasionally informs the surrounding agents of the information it has and of the data it would like to obtain.

The agents in the KCA may be modelled by a simplified BDI architecture:

$KA = (Beliefs, Capacity, Data, Goals, Intentions)$

The $Beliefs$ contain $(identifier, owner)$ associations for data elements held by this agent and by the surrounding agents. The $Goals$ of an agent are the following: to keep the used storage capacity around 75%; to fulfill requests for data (coming from the exterior or from neighbour agents); to send to neighbour agents information and intentions that the agent has. The possible actions that the agent can perform are to request data from a neighbour, to store received data or information, to dispose of data and to send its beliefs and intentions to neighbouring agents. An agent will inform its neighbours only of the data held by it and by the 8 surrounding agents, therefore all agents will be aware of the data existing at a maximum distance of 2 cells. An agent will tend to consider intentions of a neighbour agent as own intentions, if they are considered as relevant.

A specific case is presented in Figure 2. The figure shows a section of a data storage system, containing 12 agents, named from A to L. The system currently holds copies of 6 data chunks D1-D6. The first four are related. The last two are from a different domain. Each agent can store at most 4 data chunks. We presume that all agents currently have correct information about the data held by surrounding agents. The $Intentions$ of the agents will be the following: agents A, B, C and L will request a random piece of data from their neighbours, as they do not already store any data; agents D and I are at an optimal 75% used capacity, so they will not try to get more data; agent E is at full capacity, so it will remove D4, as it is already held by two of its neighbours; agent F will request D2, as it is related to D1 and D3; agent G may request any of D1-D4, as D5 and D6 are not related to any other known data; agent K will desire to get D2 (it knows from agent J that agent I has it), as being most related to D1, but no neighbour has it. However, agents F and J will share the goal of getting D2 (agent F already has it). Therefore J will get to request D2 from I, and at a later time agent K will retrieve it.

The example of the KCA system combines different types of emergents. There exists belief-related emergence: agents

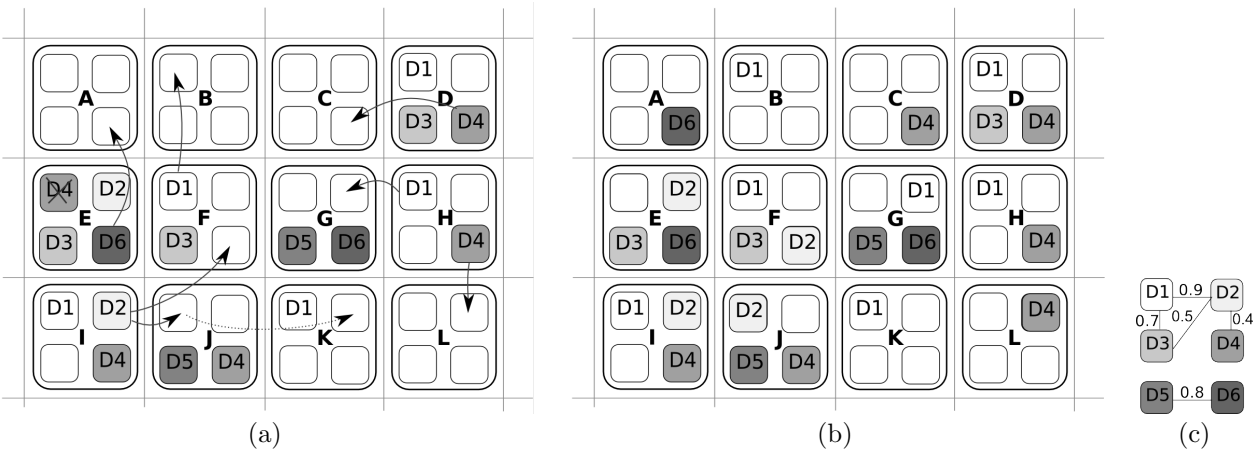

Fig. 2. Agents in a KCA, named A to L, before (a) and after (b) the data transfers indicated by the arrows. Each agent can hold at most 4 of the data chunks present in the system (D1-D6). To the right (c), the degree of relation between the pieces of data.

are attracted to related knowledge but will also discard knowledge that is superfluous in the area. Also, goal-related emergence: if an agent needs certain data surrounding agents might share this goal and cooperate with it to help the transfer. Moreover, plan-related emergence added in the context of real web server agents, as some of the agents might be more responsive than others. An agent will remember (as a belief) the response time of neighbour agents and common neighbours will use the information to choose which agent to contact if they need certain data held by more agents.

As the expected emergent properties of the CA relate to the conservation of the cell population, the KCA is expected to present emergents of similar nature. Uniform distribution and replication of data on the grid is not directly implemented in the individual agents, as their goals are mostly selfish and only describe the desired properties of their own knowledge set. Some glider-like dynamic structures may also result, propagating knowledge through the grid. It is important to point out that, although these are the emergents one might expect, emergence can currently be verified only by experiment.

Compared to the CA, the cognitive agent system is not only more efficient, as the agents can remember the information that they have received and may use it in building future plans, but the cognitive approach is more powerful and capable.

7. CONCLUSION

As computational power increases and cognitive agent systems are improved, it is necessary to discuss the notion of emergence in the cognitive context. The purpose of this paper was to define the different forms and requirements of emergence in cognitive agent systems. A simple, but relevant example has been given, combining different types of emergence and emphasising the advantages of the cognitive approach versus the reactive one.

In perspective, it is necessary to develop examples and simulations that will establish more precisely the conditions in which emergents arise, towards an appropriate formalisation of emergence as a phenomenon.

REFERENCES

Amaral, L. and Ottino, J. (2004). Complex networks: Augmenting the framework for the study of complex systems. *The European Physical Journal B-Condensed Matter*, 38(2), 147–162.

Beurier, G., Simonin, O., and Ferber, J. (2002). Model and simulation of multi-level emergence. *Proceedings of IEEE ISSPIT*, 231–236.

Boschetti, F., Prokopenko, M., Macreadie, I., and Grisogono, A. (2005). Defining and detecting emergence in complex networks. *Lecture notes in computer science*, 3684, 573–580.

Bourjot, C., Chevrier, V., and Thomas, V. (2003). A new swarm mechanism based on social spiders colonies: From web weaving to region detection. *Web Intelligence and Agent Systems*, 1(1), 47–64.

De Wolf, T. and Holvoet, T. (2005). Emergence versus self-organisation: Different concepts but promising when combined. *Engineering Self Organising Systems: Methodologies and Applications*, 3464, 1–15.

Gardner, M. (1970). Mathematical games: The fantastic combinations of john conway's new solitaire game 'life'. *Scientific American*, 223(4), 120–123.

Gershenson, C. (2002). Complex philosophy. *Proceedings of the 1st Biennial Seminar on Philosophical, Methodological & Epistemological Implications of Complexity Theory*, 1–7.

Gleizes, M.P. (2004). Learning self-organisation and emergence. Master 2 IA Lecture.

Gleizes, M., Camps, V., and Glize, P. (1999). A theory of emergent computation based on cooperative self-organization for adaptive artificial systems. In *Fourth European Congress of Systems Science*.

Goldstein, J. (1999). Emergence as a construct: History and issues. *Emergence*, 1(1), 49–72.

Heylighen, F. (2002). The science of self-organization and adaptivity. *The Encyclopedia of Life Support Systems*, 1–26.

Jennings, N. and Wooldridge, M. (1999). Agent-oriented software engineering. *Lecture Notes in Computer Science*, 4–10.

Mamei, M., Vasirani, M., and Zambonelli, F. (2004). Self-organizing spatial shapes in mobile particles: The

TOTA approach. In S. Brueckner, G.D.M. Serugendo, A. Karageorgos, and R. Nagpal (eds.), *Engineering Self-Organising Systems, Methodologies and Applications (after ESOA 2004 workshop)*, volume 3464 of *Lecture Notes in Computer Science*, 138–153. Springer.

Mano, J.P., Bourjot, C., Lopardo, G.A., and Glize, P. (2006). Bio-inspired mechanisms for artificial self-organised systems. *Informatica (Slovenia): Special Issue: Hot Topics in European Agent Research II Guest Editors: Andrea Omicini*, 30(1), 55–62.

Picard, G. (2005). 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.

Randles, M., Zhu, H., and Taleb-Bendiab, A. (2007). A formal approach to the engineering of emergence and its recurrence. *Proc. of EEDAS-ICAC*, 1–10.

Rao, A.S. and Georgeff, M.P. (1995). BDI agents: From theory to practice. In V.R. Lesser and L. Gasser (eds.), *Proceedings of the ICMAS-95, First International Conference on Multiagent Systems, June 12-14, San Francisco, California, USA*, 312–319. San Francisco, CA, The MIT Press.

Serugendo, G.D.M., Gleizes, M.P., and Karageorgos., A. (2006). Self-organization and emergence in MAS: An overview. *Informatica*, 30(1), 45–54.

Shalizi, C. (2001). *Causal architecture, complexity and self-organization in the time series and cellular automata*. University of Wisconsin–Madison.

Standish, R. (2001). On complexity and emergence. *Arxiv preprint nlin.AO/0101006*, 1–6.

Unsal, C. and Bay, J. (1994). Spatial self-organization in large populations of mobile robots. *Proceedings of the 1994 IEEE International Symposium on Intelligent Control*, 249–254.