Collective Intelligence in Multi-Agent Robotics: Stigmergy, Self-Organization and Evolution

EDUARDO IZQUIERDO-TORRES
University of Sussex, Brighton BN1 9QH, U.K.
eji21@sussex.ac.uk

January 12, 2004

Abstract

Nature has been able to evolve (several times) natural systems which produce complex spatio-temporal patterns from agents with very simple behaviours by exploiting the interactions between the agents and their environment. Surprisingly, the systematic use of these principles have been mostly neglected within the field of collective robotics. We study the use of these biological principles in nature and their application in artificial systems and hypothesize on the role between such principles and evolution. We conclude with future directions towards collective intelligence in multi-agent robotics.

1 Introduction

In social insects large numbers of simple agents collectively achieve remarkable feats through exploiting a few principles. They offer a spectacular existence proof of the possibility of using many simple agents rather than one or a few complex agents to perform complex tasks quickly and reliably. Furthermore, they exhibit living proof that evolutionary processes tend to exploit the available dynamics of the environment to minimize unnecessary complexity at the individual level. It is therefore surprising that the systematic use of these principles have been mostly neglected within the field of collective robotics. The principles we are referring to are that of stigmergy, self-organization and evolution.

How would it be best to put stigmergic self-organization at work in multi-agent robotics? The classical computational paradigm of robotics involves sensing the environment, then detecting features, then constructing or modifying a world model, then reasoning about the task and the world model in order to find some sequence of actions which might lead to success, then executing that action sequence, to finally update the world model again, etc. Such approach has turned out to be a practical impossibility regardless of the hardware resources available. To date, the behaviour-based architecture has proven much more useful in exploiting the stigmergic ideas. Several researches have used this approach to achieve collective cooperative behaviours using very simple individual level behaviours. However, contrary to assumptions held in such works, behaviour-based approaches are not perfectly fit to stigmergic principles.

The relevance of this topic to date is considerable, since there is an increased research interest in systems composed of multiple autonomous mobile robots exhibiting collective intelligence. At the same time, much work has gone into proving that social insects are excellent examples of such complex global behaviours and from their study a picture of stigmergic self-organization is emerging in the field of robotics. Under this light, complex tasks can be solved with very simple individual behaviours given that the interactions among individuals include the use of their environment. Current approaches to collective robotics have overlooked many of the important issues that have arisen from the study of such biological principles. Furthermore, it is argued herein that stigmergic ideas point particularly towards the natural exploitation of an evolutionary approach and a continuous-time dynamical system approach to the design of the group of robots.

The object of this essay will be to study the role of stigmergy and self-organization in collective intelligence of biological systems and in collective robotics to date and to point towards the future directions in designing multi-agent robotics. In next sections we will first provide brief definitions for collective intelligence, self-organization, and stigmergy. We will then look at examples from them in social insects. This will be followed by a review and discussion of research in collective robotics using these concepts. Finally, we will conclude with what we believe are the future directions.

2 Collective Intelligence and Self-Organization

According to Brooks [6] the fundamental decomposition of the intelligent systems is not into independent information processing units which must interface with each other via representations. Instead, the intelligent system is decomposed into independent and parallel activity producers which all interface directly to the world through perception and action [6]. Collective intelligence is, perhaps, the most extreme example of this view. The concept was originally coined by Wheeler in [21] and has been extensively used in the literature of social insects. In a more general context, Sulis [19] has defined collective intelligence as consisting of a large number of quasi-independent, stochastic agents, interacting locally both among themselves as well as with an active environment, in the absence of hierarchical organization, and yet which is capable of adaptive behaviour.

Many activities performed by social insects result in such collective behaviours. Researchers are often tempted to assume that such complex patterns at the colony level can be generated only by complex individuals [5]. However, with the use of theories of self-organization from physics and chemistry a different picture is emerging. Self-organization describes how microscopic processes give rise to macroscopic structures in out-of-equilibrium systems. This has more recently provided an explanation of a wide range of collective phenomena in animals, especially in social insects [5].

In the more biological context, Bonabeau and his colleagues [5] describe self-organization as a set of dynamical mechanisms whereby structures appear at the global level of a system from interactions among its lower-level components. The rules specifying the interactions among the systems constituents units are executed on the basis of purely local information [12]. Four basic ingredients and three characteristic signatures have been identified. The ingredients are positive feedback, negative feedback, the amplification of fluctuations and the presence of multiple interactions; the signatures are the creation of spatiotemporal structures in an initially homogeneous medium, the possible attainability of different stable states, and the existence of parametrically determined bifurcations [7].

These local interactions within self-organized systems can be based on two very different methods of communication. First, by means of symbolic or non-symbolic signalling from one agent to another, referred to as direct communication. Second, from stimuli obtained through the environment, referred to as indirect communication. However, whereas direct information transfers tends to be conspicuous, since natural selection has shaped signals to be strong and effective displays, indirect information transfers are often more subtle and based on incidental stimuli in an organism's environment. The lack of prominence of this last method means that many interactions within animal groups are easily overlooked [7].

3 Stigmergy in Biological Systems

The indirect form of communication among individuals was first described by French entomologist Pierre-Paul Grassé in the 1950s and denominated stigmergy (from the Greek sigma: sting and ergon: work) [10]. Stigmergy has helped researchers understand the connection between individual and collective behaviour, showing that an alternative theory could explain the paradox of coordination in social insects: Although the behaviour of the colony as a whole looks wonderfully organized and coordinated, it seems that every insect is pursuing its own agenda without paying much attention to its nest mates [20].

The basic principle of stigmergy is extremely simple: Traces left and modifications made by individuals in their environment may feed back on them [20]. The colony records its activity in part in the physical environment and uses this record to organize the collective behaviour. Various forms of storage are used: gradients of pheromones, material structures, or spatial distribution of colony elements. Such structures materialize the dynamics of the colony's collective behaviour and constrain the behaviour of the individuals through a feedback loop.

However, stigmergy alone is not sufficient to explain collective intelligence, as it only refers to animal-animal interactions. Therefore, it has to be complemented with an additional mechanism that makes use of these interactions to coordinate and regulate the collective task in a particular way. At least two mechanisms have been identified: self-organization and self-assembly. Only the first will be treated here, also known as quantitative stigmergy, or stigmergic self-organization [20]. This mechanism provides the basic ingredients in social insects; the resultant process produces outcomes that display the characteristic signatures. Stigmergic self-organization is distinguished form the purely physical mentioned above in that it involves agents that can alter the environment.

One of the best examples of this mechanism was studied by Grassé: the building behaviour of termites. Grassé showed that the coordination and regulation of building activities do not depend on the workers themselves but are mainly achieved by the nest structure: A stimulating configuration triggers a building action of a termite worker, transforming the configuration into another configuration that may trigger in turn another (possibly different) action performed by the same termite or any other termite in the colony [4].

The use of stigmergy is not confined to building structures. It also occurs in cooperative foraging strategies such as trail recruitment in ants, where the interactions between foragers are mediated by pheromones put on the ground in quantities determined by the local conditions of the environment. For example, trail recruitment in ant species are able to select and preferentially exploit the richest food source in the neighbourhood or the shortest path between the nest and a food source [2]: foragers are initially evenly distributed between the two sources, but one of the sources randomly becomes slightly favoured, and this difference may be amplified by recruitment, since the more foragers there are at a given source, the more individuals recruited to that source [5].

Michener describes in [14] many activities in bee colonies that result in nest structures, conditions of brood or stored food, to which other bees respond. Referring to this as indirect social interactions, where the construct is made for other primary objectives, not for signalling, although the information content becomes essential for colony integration. In nectar source decision making in honey bees it is less clear if only direct communication through the recruitment dances of the bees produce the self-organizing behaviour, or if also the indirect communication given by the waiting time for downloading the honey is affecting the collective behaviour [18].

As a consequence of stigmergy and self-organization, complex behaviours which had been explained on the basis of certain rules of interaction among individuals were later accounted for even simpler behaviours in the context of environmental stimuli. This is exemplified by the history of theories explaining paper wasps constructions. Given that the structures built by these animals are highly deterministic, early analysis of wasp behaviour was though to be guided by a blueprint, a mental image. Later, authors began to outline the behaviour as an inherited building program. This involved no interaction with the structure, just a sequence of actions. A cycle of inspection was later added to their repertoire, invoking different mechanisms to process the collected information. In addition to these, many other individual-level intelligent properties had been invoked to explain the colony-level phenomena, such as counting abilities, etc. With the application of stigmergic algorithms, Karzai [13] and others have shown that natural-like, complex multicomb structures could be generated.

Stigmergy seems indeed at the root of several collective behaviours of social insects, especially in their building activities. This is certainly a powerful principle, as social insect constructions are remarkable for their complexity, size and adaptive value. However, it is possible to extend the idea easily to other domains; it can be seen as an even more impressive and general account of how simple systems can produce a range of apparently highly organized and coordinated behaviours and behavioural outcomes, simply by exploiting the influence of the environment [12].

4 Stigmergy in Collective Robotics

Most of the work in robotics so far has focused on control of a single agent, but increased efforts have begun to address systems composed of multiple autonomous mobile robots. There are several reasons for this increased interest, among others: (a) some tasks may be inherently too complex for a single robot to accomplish, (b) performance, robustness and flexibility benefits, (c) to yield insight into fundamental problems in the life and social sciences.

In collective robotics direct communication has been most commonly attempted in order to achieve coordination [8]. However, this technique requires that the sending robot must encode and transmit a message about what is to be done, and where it is to be done; this message is local in time and space, and so only those robots close enough and not otherwise engaged will be free to receive the message; they must then decode the message, and either remember it for long enough to get to the place and carry out the action, or remember it for even longer while they carry out some other important task [12].

A stigmergic communication requires no encoding or decoding, no knowledge of place, no memory and it is not transient; all that it requires is that a robot passes near enough to the location where the communication was placed to be affected by it [12]. In this sense, stigmergy can be regarded as the general exploitation of the environment as an external memory resource. With a very few notable exceptions there are not many approaches to collective behaviour using stigmergic self-organization, these will be revised ahead.

Deneubourg et al in [9] studied the performance of a distributed sorting algorithm, inspired by how ant colonies sort their brood. For this he uses simulated robot teams that move randomly, do not communicate, have no hierarchical organisation, have no global representation, can only perceive objects in front of them, but can distinguish between objects of two or more different types with certain error. The probability that they pick up or put down an object is modulated as a function of how many of the same objects they have met in the past. This generates a positive feedback sufficient to coordinate the group of robots in sorting the objects into clusters.

Despite successfully modelling brooding in ant colonies based on stigmergic principles, the major caveat of Deneubourgs experiments was the use of simulation, as opposed to exploiting the use of real environments. This was tackled successively by Beckers et al in [2]. They extended Deneubourg's work, this time using physical-agents. They were able to achieve clustering with an even simpler algorithm, using physical robots that were unable to detect whether or not they were moving any objects, that had no memory, and could only sense the local density of objects as a boolean variable.

Additionally, they noticed that stigmergy increased the robustness of the collective system given that the failure of anyone agent did not cause major dysfunction. More so, adding additional robots without reconfiguring any of the already working in the task resulted in speed gains. An important aspect which emerged out of using physical agents was the notion of a critical group size threshold in which the efficiency of the work becomes affected, due to an exponential increase in the number of interactions between robots.

Other, more recent, Holland and Melhuish in [12] applied a very similar robotic approach but exploring several other sorting tasks. They showed that both segregation and also crude annular sorting of two types of object differing only in visual appearance can be achieved by a system of simple mobile robots that can sense only the type of object they are carrying and have no capacity for spatial orientation or memory.

Nevertheless, both of these experiments [2, 12] used hand designed behaviour based agents based on Brooks' subsumption architecture. The robots have few built-in behaviours, and only one is active at any time. The behaviour-based approach is based upon the idea of providing the robot with a range of simple basic behaviours and letting the environment determine which basic behaviour should have control at any given time. This choosing of the basic behaviours by the designers greatly constraints the possibilities of agent-environment interactions which could, in principle, arise from stigmergy or self-organization.

5 Stigmergy and Evolution

The fact that behaviour, from a point of view of an observer, is the result of a dynamical interaction between the agent and the environment can explain why it is difficult to break down a global behaviour into a set of basic behaviours which may well be simple from the point of view of the agent. An evolutionary approach, by relying on an evaluation of the system as a whole and of its global behaviour, releases the designer from the burden of deciding how to break the desired behaviour down into simple basic behaviours [15].

In collective robotics, currently, we know of only two works that use an evolutionary approach in a cooperative multi-agent environment. Quinn et al in [17] successfully evolved a group of physical robots to perform a formation movement task from random starting positions, equipped only with infrared sensors. On the other hand, Baldassarre et al [1] evolved a group of simulated robots for the ability to aggregate and move together toward a light target. Other two works are known to have used an evolutionary approach in real multi-robot systems but neither cooperative nor coordinated behaviours were required.

Overall, the results presented by these researches [1, 17] demonstrate that evolutionary techniques, by exploiting the self-organizing behavioural properties that emerge from the interactions between the robots and the environment, are a powerful method for synthesizing collective behaviour.

The question of how stigmergic self-organization and evolution interact is largely open, not only in robotics but in insect societies as well. Nonetheless, natural selection operating on parameters that modulate individual and colony-level properties has certainly picked the forms of stigmergic self-organization that we see in social insects.

In fact, specific stigmergic strategies may have appeared in the first place because of the underlying simplicity of their behavioural mechanisms and because of the relatively weak conditions required for emergence. Furthermore, evolution will favour self-organizing strategies that take advantage of existing biological and environmental implementations or mechanisms. Yet another reason why evolution will favour the design of stigmergic principles is that the same-individual behaviours may be used to generate different collective response in different environments.

Therefore, an evolutionary approach to collective robotics has the potential to benefit extensively and naturally from stigmergic principles. However, to date, this field has focused almost exclusively on single robot systems (see [16] for a good survey of evolutionary robotics research).

6 Towards Collective Intelligence in Robotics

Due to its elegant simplicity, stigmergy seems to provide the most general explanation for decentralized control of complex collective behaviour in many social insects. Multi-agent groups require coordination for their efforts and it seems this concept is a powerful means to coordinate activity over great spans of time and space.

As collective intelligent systems manage such an overwhelming number of complex interactions among agents and environment to solve fairly simple tasks, there is an imminent need to head towards only those approaches which naturally exploit the benefits of these interactions.

6.1 Embodied Approach

Embodiment stresses the importance of the physical aspects of a system (physical shape, gravity, friction, inertia, idiosyncratic characteristics of each sensor and actuator, etc). Stigmergy calls upon the exploitation of exactly such physical properties of the interaction between agents and environments.

As we saw earlier, embodiment has been taken into account among the stigmergic-robotics researchers, and the advantages shown significant. It is mentioned here because it is a crucial factor approaches that follow; however, it will not be treated further here.

6.2 Evolutionary Approach

The approaches taken to study stigmergic self-organization in collective robotics have been based mostly on behavioural decomposition (as studied in section 3). However, it is generally accepted that the design of robust mobile control systems is highly complex because of the extreme difficulty of foreseeing all possible interactions between separate parts of the robot itself [11]. The complexity can only increase for designing groups of interacting robots. In the evolutionary approach this is often the result of a self-organizing process. Furthermore, it can be said that evolution will naturally favour stigmergic self-organizing exploiting solutions to collective tasks on the basis that it requires much simpler behaviours from the agents.

6.3 Continuous-Time Dynamical System Approach

When relying on evolution for the design of a control system we must choose appropriate building blocks [11]. High level semantics incorporate the human designer's prejudices, this in turn will constrain the exploitation of the biological principles we are after. Since the collective of robots and the environment can be described as a dynamical system because the sensory state of each robot at any given time is a function of both the environment and of the rest of the robot's previous actions (including itself); and since stigmergic and self-organizing principles rely on these exact interactions to coordinate and regulate the collective behaviour, the building blocks should be the primitives of such dynamical system. As the agents are situated and embodied in their environment the dimension of time cannot be ignored; thus a continuous-time dynamical perspective.

The main issue that arises from studying stigmergic self-organization has been to make use of this idea exclusively in designing the group of robots. Because, even though stigmergic ideas point towards the importance of reactive behaviours, a purely reactive agent would be a degenerate case, and would be constantly pushed around by its environment. A dynamical systems perspective on autonomous agents emphasizes the importance of internal state to an agents operation. This way, unlike a reactive agent, an agent can initiate behaviour independently from its immediate circumstances and organize its behaviour in anticipation of future configurations of its environment [3].

In general, the approaches which exploit the benefit of stigmergy and self-organization have proved successful for single-robots, and it only makes sense that when increasing the complexity of the interactions, as is the case in collective robotics, these approaches remain useful. Furthermore, the use of stigmergic self-organizing principles place an even greater importance on the issues on which these approaches are centred.

7 Conclusions

Stigmergic self-organization can not be regarded as a complete theory of collective behaviour, but it is an important concept that can help to provide a simple explanation for several aspects of it. Furthermore, the use of such concepts in achieving collective intelligence through robotics points strongly to the use of (a) embodied agents, (b) a continuous-time dynamical systems approach to agent-environment interaction and (c) an evolutionary approach to robot controller design.

References

[1] G Baldassarre, S Nolfi, and D Parisi. Evolving mobile robots able to display collective behaviors. *Artificial Life 9: 255-267*, 2003.

[2] R Beckers, OE Holland, and JL Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. *In Art. Life IV, 4 th Int. Worksh. on the Synth. and Simul. of Living Sys*, 1994.

[3] RD Beer. A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence 72:173-215*, 1995.

[4] E Bonabeau. Editor's introduction: Stigmergy. *Artificial Life 5: 95-96*, 1999.

[5] E Bonabeau, G Theraulaz, JL Deneubourg, S Aron, and S Camazine. Self-organization in social insects. *Trends in Ecology and Evolution 12:188-193*, 1997.

[6] RA Brooks. Intelligence without reason. In John Myopoulos and Ray Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 569–595, Sydney, Australia, 1991. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA.

[7] S Camazine, J Deneubourg, N Franks, J Sneyd, G Theraulaz, and E Bonabeau. *Self-Organization in Biological Systems.* Princeton University Press, 2001.

[8] YU Cao, AS Fukunaga, and AB Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots 4, 7-27*, 1997.

[9] JL Denebourg, S Goss, N Franks, A SendovaFranks, C Detrain, and L Chretien. The dynamic of collective sorting robot-like ants and ant-like robots. *In J.A. Meyer, S.W. Wilson First Conf. on Simulation of Adaptive Behavior: From Animal to Animats MA:MIT Press, pp.356-365*, 1991.

[10] PP Grassé. La reconstruction du nid et les coordinations interindividuelles chez bellicositermes natalensis et cubitermes sp. la theorie de la stigmergie: essai d'interpretation du comportoment des termites constructeurs. *Insects Sociaux 6:41-83*, 1959.

[11] I Harvey. *The artificial evolution of adaptive behaviour.* DPhil Tesis. University of Sussex, 1995.

[12] O Holland and C Melhuish. Stigmergy, self-organization and sorting in collective robotics. *Artificial Life 5: 173-202*, 1999.

[13] I Karsai. Decentralized control of construction behavior in paper wasps: An overview of the stigmergy approach. *Artificial Life 5: 97-116*, 1999.

[14] CD Michener. *The Social Behavior of the Bees: a comparative study.* Harvard University Press, 1974.

[15] S Nolfi. Evolutionary robotics: Exploiting the full power of self-organization. *Connection Science, 10(3-4)*, 1998.

[16] S Nolfi and D Floreano. *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines.* MIT Press, 2000.

[17] M Quinn, L Smith, G Mayley, and P Husbands. Evolving teamwork and role-allocation with real robots. *Artificial Life 8, pp 302-311*, 2002.

[18] TD Seeley, S Camazine, and J Sneyd. Collective decision making in honey bees: how colonies choose among nectar sources. *Behavioral Ecology and Sociobiology 28:277-290*, 1991.

[19] W Sulis. Fundamental concepts of collective intelligence. *Nonlinear Dynamics, Psychology, and Life Sciences, Vol. 1, No. 1*, 1997.

[20] G Theraulaz and E Bonabeau. A brief history of stigmergy. *Artificial Life 5: 97-116*, 1999.

[21] WM Wheeler. The ant colony as an organism. *Journal of Morphology, 22,2:307-325*, 1911.