On Problems of Self-organization

Richard Potter

March 5, 2007

1 Introduction

Brook's manifesto [7] [8] of embodied, situated robots coupled with the environment has had a very strong influence in shaping much of the mobile robot research performed over the last thirty years. However, while Brooks' work laid the foundation for the study of single-robot systems, no such dominance has emerged in the area of multi-robot systems [11].

There are many situations where, a group, team or swarm can solve tasks too complex for an individual to accomplish, or can complete the tasks more efficiently. The group approach has further advantages including that it can be tolerant of the failure of individuals and can be more flexible in solving tasks. When considering mobile robots, a group approach can be cheaper, rather than building one large complex robot, due to technical limitations. Multi-robot systems can also provide a framework to study the nature of society.

A group requires some form of control to enure that the system acts coherently. Control can be imposed through a leader, plan or recipe that can be followed. Alternatively, decentralised approaches such as self-organisation [10] can be applied to the system.

Self-organisation occurs in many biological systems where the interactions of individuals, driven by negative and positive feedback cycles, result in rich adaptive, flexible and robust behaviours. The systems' agents can be cognitively simple, equipped with basic behavioural rules, limited localised perception and little or no means for direct communication. For many systems, Stigmergy [16], defined as communication mediated by changing the environment, is the only form of communication required.

The biological grounding of self-organisation via stigmatic communication correlates strongly with the embodied cognitive approach of Brooks and others, which is itself biologically inspired. However, much of the research for multi-robot systems has avoided self-organisation as an approach. Where it has been applied to the control of multiple robots, simple ant metaphors have been pervasive [3] [6] [17] [19] [20] [24] [31].

In this paper I will explore the problems of extending the behaviour-based robotics approach to multi-robot groups using self-organisation and stigmergy, examining in detail the reason why control systems have not moved beyond the simple ant metaphor.

2 Control approaches for multi-robot systems

Multi-robot system control has been studied for more than a decade but the application of self-organisation in such systems has been limited [12]. Instead, the majority of the work has focused on what Gerkey & Matarić [15] call *intentional cooperation* [11] [13] [23].

The argument for intentional cooperation is that compared to the self-organised approach, 'intentional cooperation is usually better suited to the kinds of real-world tasks that humans might want robots to do' and that such co-operation can 'better exploit the capabilities of heterogeneous robot teams' [15]. In such systems, robots can be allocated a task or subtask by a leader [13] or use a shared resource, blueprint or recipe to guide task solving [9]. The problems with these approaches mirror the differences in Brooks' work and the sense-model-plan-act work that proceeded it.

However, for a society to be co-ordinated by a leader, requires highly effective communication from the leader to each individual and vice versa. In addition, leaders must have complex cognitive abilities to acquire and disperse information to plan the group's moves in a realistic time frame. Such planning requires in-depth knowledge of the problem domain and the societies sub-units. An expensive developmental process would also be necessary to enable the leader to learn the information needed to solve the required task.

Blueprints and recipes suffer from similar problems. Blueprints only specify what to build, not how, so they also require considerable cognitive ability to be useful. Recipes state how to build, making them extremely rigid and only suitable for tasks with a high degree of sequential steps, which negates group work. In biological systems, encoding genetically and expressing blueprints or recipes in the developmental process would be extremely costly.

Evolution has instead preferred a *distributed coordination* [3] approach based on self-organisation for many insect and animal societies. Self-organising systems give rise to complex patterns through interactions internal to the system. These interaction are based on localised information and the behaviour that arises is an emergent property of such systems. The pattern of a self-organised system often results in an arrangement of objects in both time and space. This is seen in construction of termite mounds or other nest fabrications [1] [10].

Therefore, if the Brooks' approach is adhered to, then self-organisation appears to correlate naturally as a method of controlling multi-robot systems. The behavioural rules underpinning a self-organised system are simple but result in the generation of global patterns that are much more complex than the components and properties they emerge from. Additionally, agents act upon what they perceive directly in the environment.

3 Properties of self-organising systems

From reviewing literature [10] [17] [30], a number of key principles of self-organising systems can be extracted. For self-organisation to provide a dynamic method of control, a system requires: a society of agents, accumulation of information, positive feedback, negative feedback, stable state(s) and bifurcations in behaviour.

Social grouping of agents

Agents in a self-organising system have some basic requirements: the ability to move through the environment, to act upon the environment and to interact with other embedded mobile agents. The final requirement is based upon the premise that there must be a cohesive society of agents in the environment.

Communication in self-organising systems is between subunits; either individual to individual, individual to group or group to group. Social grouping or clustering is required to assure that social amplification of signals will occur providing an impetus for change.

Grassé proposed the idea that sociality is a characteristic deeply rooted in the ethological heritage of a species [30]. There must therefore be a predisposition for aggregation or socialising and this predisposition effects a creature's behaviour even when its alone [25].

For groups of robots, social grouping is generally inherently extant regardless of whether the robot group is homogeneous or hetrogenous. Nevertheless, Matarić [21] worked on enabling robots to distinguish between members of their group and other objects in the environment, so called 'Kin recognition', which could be an important foundation for some self-organising behaviours such as flocking or herding.

The accumulation of information

Self-organisation is built around the accumulation of spatial and or temporal structures in an often initially homogeneous medium [30]. Typically this takes the form of adding or removing information from the local environment, which then alters the quality of the environment. Biological examples include depositing pheromones [16], comb construction in bees [5] or wall construction in ants [14].

Individuals will modulate their behaviour based on the accumulation of this information. Results of their actions then modulate the behaviour of others via alterations to this information. This places a requirement on the system that any information placed in the environment must persist long enough for others to interact with it. The spatial distribution of this information in the environment controls the spatial distribution of agents in the environment, as they interact with this information and their behaviour is affected by it.

Positive & negative feedback

For self-organisation to work the simple behavioural rules must encode, directly or indirectly, the mechanisms for positive and negative feedback. Positive feedback is required to cause the system to canalize from uncertain starting points towards a behavioural steady-state. By promoting the accumulation of information or structures in the environment, the feedback mechanism amplifies change in the direction towards a desired behavioural state from small perturbations.

Examples of positive feedback include pheromones secretion [1] [16], bee dance patterns [27] and direct ant contact when recruiting foragers [10].

In insect societies, the large number of individuals provide strong interactions between each other, forcing change. The feedback mechanism is essentially distributed. As the number of individuals falls, then the density of feedback can fall, requiring more forceful forms of communication and less temporal structures to ensure positive feedback

Negative feedback is also required to stabilize the system towards a steady state. This feedback is often implicit in environmental or physiological constraints such as availability of a resource or the amount of energy required to keep accumulating new information. Pheromone evaporation in termite and ant societies is one natural example.

These feedback mechanisms do not have to be singular, there can be many mechanisms providing positive and negative feedback.

Steady states & bifurcations

Any self-organised system needs to support the possible coexistence of several stable states. The system needs to converge on on a steady state from its initial conditions to enable the solving of a task. Bifurcations must also be present to enable the behaviour of the self-organised society to change dramatically and canalize towards another steady state.

This enables a insect society to switch patterns such as selecting a nest site, sorting the brood, building defensive walls and foraging behaviours in ants [14].

4 Governing interactions

Individuals in a self-organising society must interact and exchange information locally for the feedback mechanisms to work. Such information can be gathered from individuals or through from the environment including any work in progress.

Direct communication is possible even in the absence of language and is present in many insect societies. Ants returning from foraging are seen to recruit others by touching antennas with conspecifics and regurgitating some of the food source. Other forms of grabbing are used for trail-running and some species use head movements similar to waggle dancing in bees to recruit foragers [10].

Such interaction is obviously a one-to-one communication method and as social and environmental complexity grows the sheer number of interactions that occur prevent sufficient information being conveyed through direct communication. Instead, one-to-many communication, via indirect signals or cues, becomes necessary to communicate efficiently.

Anderson et al [1] defines a signal as 'an act of communication that has been selected for by natural selection' such as alarm pheromones in ants. Cues are defined as 'a structure or behaviour that conveys information, but only incidentally, and has not been shaped [directly] by natural selection'. A further class of signals, modularity signals, effect individuals responsive to cues and signals such as ants drumming on the nest walls to signal danger [14] or the tremble dance seen in bees to discourage foraging [27].

Some signals are global, such as defence and alarm pheromones, designed to reach a large proportion of the society quickly. With a small number of individuals, signals have a potential to reach a large proportion of the society as they are strong simple signs. As social and environmental complexity grows, the range of the signal becomes more and more localized. Agents out of range will not be able to exchange signals. The dances performed by bees are an example.

Cues can still be localized like signals but they can persist in time as well. Insect examples of cues include the number of brood, the amount of stored honey, the time to search for honey and pollen in a hive. As information from cues is incidental and can originate in the environment through modification, only the receiver of the cue can be acted upon by natural selection.

Stigmergy

One of the most powerful methods of indirect communication is Stigmergy [16], defined by Grassé to explain how nest building in termites is driven by 'incitement to work by the products of work'. Thus, the behaviour of an individual is modified by changes in the local environment, which were caused by the previous behaviour of that individual or others. Using Stigmatic communication, a society can trace a record of their activity in the environment that surrounds them.

However, as Shell & Matarić [28] point out, over time the term has become loosely defined and for many researchers self-organisation and stigmergy have become intertwined. The view taken in this paper is that stigmergy is a powerful mechanism for guiding interactions between subunits in a self-organising society but it is neither a required component of self-organisation, nor requires self-organisation.

Stigmatic communication forms a natural fit with the Brook's approach as a robot senses changes in its surroundings, enabling the use of the environment as shared memory for itself and conspecifics. The structure of such communication can be via the manipulation of material in the environment but can also include the sensing of conspecifics as well as hetrospecifics [5] [21].

Such communication, mediated via the environment, has been successfully used in several studies with multi-robot control [17] [20].

5 Problems of self-organisation

Even though self-organisation has been a marginal area of research in multi-agent systems, work has been performed examining foraging [2] [4] [20], box-pushing [19], clustering & sorting [17] and swarm formation [3] [6] [29] [31].

What can be seen from examining self-organising biological and robot systems is that there are a large number of design issues that must be considered for the system to function in a dynamic uncertain environment.

Firstly, patterns emerging from a self-organising system are dependent on many factors. The previous history of the system will affect the outcome and minor differences can effect the canalization leading to the production of different patterns. The presence of environmental cues such as temperature, moisture and light levels can also influence the pattern of a system, potentially triggering bifurcations [10].

Positive feedback mechanisms are relatively easy to encode in the individual as simple behaviour rules. However, negative feedback mechanisms are often in implicit in the form of environmental templates making it hard to encode or design such mechanisms. When the starting conditions are unknown, a multi-robot group may need the cognitive ability to recognise if the local environment is insufficient to support the current task.

The response of the system is sensitive to the density or number of individuals involved, as this controls the success of positive feedback mechanisms. This is especially true if information is temporal, in which case the information has to persist long enough to be consumed. If information is spatial, then then density of the information or the individuals must be large enough so information is found, sensed and consumed. There has been little research in approaches to organising stimuli in a space and time to aid self-organisation.

The success of communication strategies in an uncertain environment can be highly variable due to obstructions, sensor issues and other factors that limit communication range. Failure of individuals to respond to communication reduces interaction density. Balch et al [2] showed that many tasks can be solved with no or minimal communication. It is my opinion, that approaches which rely on the ability of each robot to communicate with each other, such as Shen et al [29], and Parker [23], limit their application to 'friendly', certain environments.

Deadlock of the system is another important control issue that is ignored in the majority of studies. The system's sub-units work with locally sensed information so deadlocks and stagnation may occur, as information is limited. In Kube & Bonabeau's study [19] of box pushing, the robots would use biologically inspired patterns of behaviour if the box seemed stuck, first by making small changes to the system and then allowing for larger changes. Parker's agents would give up performing a task when it took too long [23]. Such mechanisms slow system performance but increase success, which has ramifications when the agents are in competition or in hostile territory. The systems ability to escape deadlock situations is enhanced by the randomness or fluctuations in individuals [19]. In mobile robots this randomness could arise due to the imperfect nature of sensory-motor circuits, but the use of both hetrogenous robots, controllers and deadlock mechanisms may also be required.

The biggest problem for applying self-organisation as a method of robot control is this large range of parameters and mechanisms, present even in simple systems. As its not clear for a given set of parameters what patterns may emerge, such emergent systems are difficult to analyse and design. It is highly unlikely that a designer could replicate a self-organising system that has multiple stable states or deviates from the simple ant metaphor. Evolution has shaped self-organising biological systems, by selecting signals, cues, the individual's receptiveness and responsiveness to communication amongst others. Therefore to develop self-organising system, evolutionary approaches must be used to evolve the individual agents, their means of interaction and behavioural rules. The system then must be evaluated as whole [22].

The simple ant metaphor revisited

One commonly held view of self-organising systems is that they are only useful with large numbers of agents, all of which have limited cognitive and communicative abilities. Anderson & McShea's influential study [1] showed that as the complexity of ant societies grows, the amount of direct communication falls between subunits. There is also a strong correlation between social complexity, the 'cognitive' complexity of individuals and the switching of communication from direct, to signals and to cue based forms. The use of this *simple ant metaphor* has been pervasive [3] [6] [17] [19] [20] [24].

Ant societies are some of the most studied self-organising biological systems. The metaphor of simple individuals responding to a few strong indirect signals, makes an attractive framework to study and implement self-organisation. Such systems are stable almost regardless of starting state [10] and they fit well within what modern robotics is capable of.

However, Anderson & McShea's study shows that signals and cues are also used by simpler societies and there is a mixture of cognitive ability and communication strategies in between simple and complex ant societies. The simple ant metaphor is one of many potential metaphors from studying biological systems. For example, Beavers live in small colonies of 2-11 individuals and their construction of dams is thought to be self-organised via stigmatic communication.

The simple ant metaphor is based on the premise that if individuals process only local information and are 'cognitively' simple then individually they have a higher chance of making poor decisions. But if they exist in a large society then then collectively they will make better decisions.

For robot groups there are two influences that may counter the use of simple agents. Firstly, very few multi-agent robot systems have more than a small number of agents (due to resource and technical limitations [11] [12]), thus agents in such systems must require some complexity. Secondly, multi-robot systems typically do not have a nest construct to return to so, forcing individuals to be more self-reliant and more cognitively capable, rendering much of the simple ant metaphor redundant.

As the cognitive ability of agents increases, the potential to absorb signal, cues and templates also increases. Biesmeijer & Shaa [5] shows an individual bee can access over twenty known information sources that come from five major groups: intrinsic sources (itself and its own experience), nest environment (dancing, information about food reserves), field (cues left by conspecifics, presence of non nest-mates). The expansion of information and the extended range of such agent may operate in, complicate the design of such self-organising systems.

A final observation is that the simple ant metaphor is often applied in a system with one steady steady state. No known robot control systems have been implemented utilizing the emergence of multiple patterns and bifurcations in systems behaviour. The effect of environmental cues, termed pre-existing environmental templates by Camazine et al [10], has also been ignored in most studies as a method of starting a pattern or switching patterns. Bifurcations and environment cues have a huge influence on biological systems.

Using a metaphor that does not support cognitive agents and does not support multiple steady states seems highly restrictive. Nevertheless, without the any formalism, modularity or evolution in multi-robot systems, the simple ant metaphor will dominate due to its practicality of application.

6 Conclusion

Self-organisation is a powerful biologically inspired control system that complements the behaviour-based robotics approach. It requires only simple behavioural rules amplified through interactions for emergent global patterns to arise that are robust and flexible. In addition, stigmatic communication can be used to channel nearly all communication required. Work has shown that small groups of robots can use self-organisation successfully using simple indirect communication techniques [19] [26].

Self-organisation is scalable. In biological systems it can work for systems of a few individuals as well as in systems with a huge number of individuals [1]. Self-organised robot systems have shown robust behaviour in the presence of fluctuating numbers of agents [3] [17] [19] .

The principle problem of applying self-organisation is to ensure that coordinated and coherent system-level behaviours arise as a result of the interaction of individual agents. However, decomposing a self-organised system to individual behaviours and their interactions via signals, cues and feedback mechanisms is an extremely difficult problem. Taking further inspiration from biological systems, where nature has shaped the patterns via evolution, artificial evolution must therefore be used.

The central issue with evolving self-organising systems, is the identification and parametrising the elements of the system that underpin self-organisation. This is compounded by the time consuming nature of evolutionary algorithms where the outcome of the system is only known after careful study. However, recent work in evolving self-organised systems has begun to generate real results. Quinn et al [26] has evolved controllers for simple robots that moved in an organised formation using localized information about their immediate neighbours. The Swarm-bot project [3] also intends to evolve controllers for real robots via evolution techniques to achieve motion with a metamorphic collection of robots.

My hypothesis is that even with evolutionary techniques, self-organisation for multi-robot systems will only become a powerful tool if researchers move beyond simple ant metaphors. This requires careful study into formalizations and methodologies of application. Constructing simple and robust building blocks where subunits are known to solve sub-tasks correctly is one solution. Once a pattern is formed it can then be evolved further. Simple foraging techniques such as those seen in ant societies could be developed and then evolved to the more sophisticated patterns seen in bee foraging [5].

An alternative solution would be to not view self-organisation as a separate approach from the intentional co-operation strand. Hybrid approaches are possible. The *ALLIANCE* work of Parker [23] has decentralized behaviour based control with many features similar to self-organising system. While it relies on the ability of each agent to communicate directly with all other conspecifics, it does utilise a method of switching behaviours which could be used to allow multiple patterns in a single multi-robot system.

References

[1] Anderson C. & McShea D.W. *Individual versus social complexity, with particular reference to ant colonies.* Biology Review, 76:211-237, 2001.

[2] Balch T. & Arkin R.C. *Communication in reactive multiagent robotic systems.* Autonomous Robots, 1:1-25, 1994.

[3] Baldassarre G., Parisi D. & Nolfi S. *Distribute coordination of simulated robots based on self-organisation.* Artificial Life, 12:289-311, 2006.

[4] Beckers R., Holland O.E. & Deneubourg J.-L. *From local action to global tasks: Stigmergy and collective robotics*, Artificial Life IV, Proc of the Fourth International Workshop on Synthesis and Simulation of Living Systems, 181-189, 1994.

[5] Biesmeijer J.C. & Slaa E.J. *Information flow and organization of stingless bee foraging.* Apidologie 35, 143-157. 2004.

[6] Bojinov H., Casal A. & Hogg T. *Emergent structures in modular self-reconfigurable robots.* Proc. of the 2000 IEEE International Conference on Robotics & Automation, 1734-1741, 2000.

[7] Brooks R.A. *A robust layered control system for a mobile robot.* Robotics and Automation, Vol. 2, 1:14-23, 1986.

[8] Brooks R. A. *Intelligence without reason*, Proc. of 12th Int. Joint Conf. on Artificial Intelligence, 569-595, 1991.

[9] Burgard W., Fox D., Moors M., Simmons R. & Thrun S. *Collaborative multi-robot exploration.* Proc. IEEE Int. Conf. on Robotics & Automation (ICRA), Vol. 1, 476-48, 2000.

[10] Camazine S., Deneubourg J.-L., Franks N.R., Sneyd J., Theraulaz G. & Bonabeau E. *Self-organization in bilogical systems.* Princeton University Press, 2001.

[11] Cao Y.N., Fukunaga A.S. & Kahng A.B. *Cooperative mobile robots: Antecedents and directions.* Autonomous Robots, 4:1-23, 1997.

[12] Dudek G. *A taxonomy for multi-agent robotics.* In Robot Teams: From Diversity to Polymorphism, AK Peters Ltd., 2002.

[13] Feddema J.T., Lewis C. & Schoenwald D.A. *Decentralized control of cooperative robotic vehicles: Theory and application.* Robotics and Automation, Vol. 18, 5:852-864, 2002.

[14] Franks N.R., Wilby A. & Silverman B.W. & Tofts C. *Self-organisation nest construction in ants: Sophisticated building by blind bulldozing.* Animal Behaviour, Vol. 44, 2:357-375, 1992.

[15] Gerkey B.P. & Matarić M.J. *A formal analysis and taxonomy of task allocation in multi-robot systems.* Int. Journal of Robotics Research, Vol. 23, 9:939-954, 2004.

[16] Grassé P.-P. *La reconstruction du nid et les coordinations inter-individuelles chez Ellicositermes Natalensis et cubitermes sp. la theorie de la stigmergie: Essai d'interpretation du comportment des termites constructeurs.* Insectes Sociaux, Vol. 6, 1:41-80, 1959.

[17] Holland O. & Melhuish C. *Stigmergy, self-organization and sorting in collective robotics.* Artificial Life, 5:172-202, 1999.

[18] Klavins E. *Communication complexity of multi-robot systems.* Proc. 5th Int. Workshop on the Algorithmic Foundations of Robotics, France, 2002.

[19] Kube C.R. & Bonabeau E. *Cooperative transport by ants and robots.* Robotics and Autonomous Systems, 30:85-101, 2000.

[20] Labella T.H., Dorigo M. & Deneubourg J.-L. *Division of labor in a group of robots inspired by ants' foraging behavior*, ACM Transactions on Autonomous and Adaptive Systems, Vol. 1, 1:4-25, 2006.

[21] Matarić M. *Kin recognition, similarity, and group behavior.* Proc. 15th Annual Cognitive Science Society Conference, pp. 705–710, 1993.

[22] Nolfi S. & Floreano D. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines.* MIT Press, 2000.

[23] Parker L.E. *ALLIANCE: an architecture for fault tolerant multirobot cooperation.* Robotics and Automation, Vol. 14, 2:220-240, 1998.

[24] Parnauk H.V.-D. *Making swarming happen.* Proc. Conf. on Swarming and Network Enabled Control, C4ISR, 2003.

[25] Peters P.J. *Orb web construction: interaction of spider (Araneus diadematus Cl.) and thread configuration.* Animal Behavior, 18:478-484, 1970.

[26] Quinn M., Smith L., Mayley G. & Husbands P. *Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors.* Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences, Vol. 361, 2321-2344, 2003.

[27] Seeley T. D., Camazine S. & Sneyd J. *Collective decision making in honey bees: How colonies chose among nectar sources.*, Behavioral Ecology and Sociobiology, Vol. 28, 4:277-290, 1991.

[28] Shell D. A. & Matarić M. J. *On the use of the term Stigmergy.* Proc. 2nd Int. Workshop on the Mathematics and Algorithms of Social Insects, 193, 2003.

[29] Shen W.-D., Chuong C.-M. & Will P. *Simulating self-organization for multi-robot systems.* Proc. Int. Conf. on Intelligent and Robotic Systems, 2776-2781, 2002.

[30] Theraulaz G. & Bonabeau E. *A brief history of Stigmergy.* Artificial Life, Vol. 5, 97-116, 1999.

[31] Wessnitzer J., Adamatzky A. & Melhuish C. *Towards self-organized robot formations: A decentralised approach.* Proc. 6th European Conf. on Advances in Artificial Life, 573-581, 2001.