

Controlling the Masses: Control Concepts for Multiagent Mobile Robotics

T D Barfoot, E J P Earon, G M T D'Eleuterio
Institute for Aerospace Studies
University of Toronto

Science fiction often depicts mobile robots as intelligent animats capable of operating in complex environments and interacting with humans as well as other robots. As science fiction turns into science, the possible applications for mobile robots will be limited only by the creativity of their designers. Motivation stems not only from fiction but from reality; biological life and intelligence are perhaps the largest motivators of all. Is it possible to produce mobile robots with intelligence comparable to biological life forms? If so, studying the simplest intelligent behaviour would likely be a good first step. Ants are one possibility. Individually they are relatively simple, but in large colonies interesting and useful behaviours emerge. Thus, it is one focus in mobile robotics today to develop teams of very simple mobile robots which work to some common goal. It is hoped such networks will benefit, for example, future space exploration missions.

1 Introduction

At the first Canadian Space Exploration Workshop held last summer, it was suggested that science should be the driver for future space missions. That is, missions should be designed around the types of science that Canada and the world want to see. Naturally, much effort to date has been concentrated on the exploration of other planets, for by studying other worlds we can learn more about our own. Of Mars, the late Carl Sagan wrote

But if Mars is lifeless, we have two planets, of virtually identical age, evolving next door to each other in the same solar system: life arises and proliferates on one, but not on the other. Why? This is the classic scientific circumstance of the experiment and the control.

An abundance of science still needs to be done on and below the surface of Mars and the other planets in our solar system and beyond. One of the most popular concepts in planetary science today is *networking*. Network science commonly refers to that requiring a distribution of (possibly simultaneous) measurements or a distribution of platforms. Consider, for example, seismology studies of an alien body that will require sending a signal from one point on the surface to be read at several other points in order to analyze the material characteristics of the body. Or, consider the deployment of a very-low frequency array (VLFA) on the Moon to allow for hitherto unattainable astrophysical observations using radio astronomy. Such an observatory will require a number of dipole units deployed over a region of a few hundred square kilometres. This concept was in fact studied by the 1993 session of the International Space University [11]. These and other examples of network science can be facilitated by a network of small mobile robots, similar to a colony of ants.

2 Network Science

Network science missions have been in the works since the mid-1980s with NASA's Mars Environmental Survey (MESUR) which had plans for as many as 16 stationary landers as well as ESA's MARSNET which also planned to have a small network of surface stations. For various budget-related reasons, these two missions were eventually merged to form INTERMARSNET which ESA did not elect to fly as it was too highly dependent on the US. By 1997, ESA turned its attention to Mars Express which built on past objectives of MARSNET/INTERMARSNET and planned for 4 landers. However, by 1998 it was downsized to 1 lander and a rover, similar to the NASA's Pathfinder mission with the highly successful Sojourner. Mars Express is currently due for launch in 2003. France recently put together a mission called NetLander which also plans for 4 landers which will use the communications package on Mars Express. NetLander is proposed to launch in 2005.

Not all network science missions have been aimed at Mars. The Lunar European Demonstration Approach (LEDA) and International Lunar Farside Observatory and Science Station (ILFOSS) [11] planned for networks of sensors on the Moon. It should be pointed out that the only network-like science mission that has actually made it into space is the Apollo program which used four seismic sensors on the Moon to study local internal lunar structure. Nevertheless, many scientists (as reflected in the Proceedings of the 1996 Planetary Surface Instruments Workshop [8]) believe that it is only through multiple simultaneous measurements that we will be able to fully understand complex planetary systems.

Fields that lend themselves to network science include, but are not limited to, seismology, meteorology, geology, magnetism, geophysics (subsurface heat flow, water detection), astrobiology, astronomy, mineralogy and cartography. The degree to which a network of sensors would benefit these fields is determined by the *time-constant* of each system. Fast time-constant systems change quickly and literally require simultaneous measurements (seismology, magnetism, geophysics, atmospheric circulation, meteorology) whereas slow time-constant systems change slowly and could actually have measurements taken hours or days apart (the network aspect is of benefit in allowing a large quantity of measurements to be

taken in different locations). Furthermore, there are intermediate tasks which may be accomplished by a network (e.g., it may be easier to navigate over some region with a network of rovers than a single one thus allowing a geology study to be completed more easily). These intermediate tasks are strictly not network science but science aided by a network.

3 Network Robotics

Having outlined the potentially extensive role that network science can play in planetary exploration, it is only proper that we address the technical and technological objectives that would have to be achieved to realize network science. What is required is, in essence, a method of deploying a distributed system of sensors and actuators. Rather than exhaustively exploring the entire spectrum of possibilities, we shall limit ourselves to *network robotics*, that is, a network of mobile robots or rovers.

One objective of multiagent mobile robotics is to come up with some control scheme which causes the group to autonomously achieve some predefined goal. Owing to the unstructured nature of mobile robotics tasks, it is widely believed that the key to accomplishing this goal is to make robots behave intelligently (in the biological sense).

It should perhaps be mentioned that there are typically two schools of thought when it comes to intelligent robotic behaviour. The first is traditional artificial intelligence (a.k.a. symbolic AI, deliberative AI, old AI). The second is machine learning (a.k.a. behaviour-based AI, new AI). The focus here will be on the latter approach as it seems to be more readily applicable to multiagent systems.

For a good review of machine learning in general see [9] and for one of multiagent systems see [12]. Multiagent robotics researchers have had good success on problems such as box pushing [7], heap formation [5], structure building [2], multirobot task division [10], game playing [4], navigation [1], foraging [6], herding [3]. Some of these studies were done in hardware and some in simulation. Much of the motivation behind this work comes from observations of biological systems, particularly social insects (e.g., ant, termites, wasps) and mammals (e.g., naked mole rats). Although individuals in these species act autonomously, interesting global behaviours result. Multiagent systems are sometimes described as self-organizing or emergent because the global behaviour often arises unexpectedly; it is sometimes not explicitly designed into the system but is discovered as a means to the desired end.

There are a number of important issues to be considered in multiagent learning. Among these are the issues of centralized vs. decentralized control, local vs. global goals, local vs. global observers (to provide feedback), direct vs. indirect communication amongst agents, simulation vs. hardware studies, scaling, and learning vs. fixed control.

The common idea which ties together all work in machine learning is the notion of setting up the problem in terms of a model class (e.g., decision tree, neural network, CMAC, cellular automaton) and a search algorithm (e.g., genetic algorithm, reinforcement learning, error back-propagation). The search algorithm is intended to hunt for a specific instance of the model class which produces the desired behaviour (whatever that may be). At this

point, there is no generalized approach to multiagent robotics. Each model class and learning algorithm combination is different. Even once a solution is found, its analysis can be quite difficult so the obvious question is: what good is machine learning? The quick answer, afforded in a paper of this length, is that very often machine learning performs extremely well.

4 The Future of Network Robotics

The bulk of research in multiagent robotics has occurred in the last decade so one could say that the field is quite young. There has been some progress forward but most research in this area lacks a mathematical framework to explain how simple local interactions can produce a desirable global effect. A great number of fascinating examples have been synthesized but their analysis is typically difficult. The use of global goals rather than local ones is not particularly widespread (likely due to the increased complexity of the problem) but will hopefully become more popular in the future such that real team solutions to problems may be found.

The study of multiagent systems is already quite interdisciplinary, and is certainly not restricted in application to network space science. It draws on computer science, engineering, biology, information theory, operations research. Thus as progress is made in these disciplines, it will hopefully mean progress for multiagent research. The field has by no means stagnated but is rather proliferating at great speed (e.g., software agents for data searching, multiagent approaches to economics, ecological modeling, planetary space exploration). It is safe to say that even more fascinating multiagent examples will be synthesized as time passes and computers improve. Looking to biology, one sees the possibilities for network robotics, especially in insect societies. Should we be able to construct an artificial system as robust and functional as a colony of ants a great deal of could be learned not to mention the engineering applications, particularly for space exploration as communication with Earth becomes less feasible the further we roam.

It is hoped some general principles of multiagent systems will be found but the field will progress regardless. It is possible that no general approach exists and that each system must be studied on its own; this possibility seems remote. There are already inklings of a general theory floating around (e.g., universality classes, synergetics, far-from-equilibrium systems, edge of chaos, self-organized criticality). The field is in need of somebody to bring everything together in a set of general laws of self-organization as Maxwell did for electricity and magnetism. For now, we may have to settle for the synthesis of interesting systems for both pure research and practical development.

References

- [1] Ronald C. Arkin and Kaled S. Ali. Integration of reactive and telerobotic control in multi-agent robotic systems. In D. Cliff, P. Husbands, J. A. Meyer, and S. W. Wilson,

- editors, *Simulation of Adaptive Behaviour: from Animals to Animats 3*. MIT Press, 1994.
- [2] Eric Bonabeau, Guy Theraulaz, and François Cogne. The design of complex architectures by simple agents. Working Paper 98-01-005, Santa Fe Institute, 1998.
 - [3] Jessica K. Hodgins and David C. Brogan. Robot herds: Group behaviors for systems with significant dynamics. In Rodney A. Brooks and Pattie Maes, editors, *Artificial Life IV, Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press, 1994.
 - [4] Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuk Noda, and Eiichi Osawa. Robocup: The robot world cup initiative. In *Proceedings of the First International Conference on Autonomous Agents*, Marina del Ray, 1997. ACM Press.
 - [5] Marinus Maris and René te Boekhorst. Exploiting physical constraints: Heap formation through behavioural error in a group of robots. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Senri Life Science Center, Osaka, Japan, Nov. 4-8 1996.
 - [6] Maja J. Matarić. Learning to behave socially. In D. Cliff, P. Husbands, J. A. Meyer, and S. W. Wilson, editors, *Simulation of Adaptive Behaviour: from Animals to Animats 3*. MIT Press, 1994.
 - [7] Maja J. Matarić. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16:321–331, 1995.
 - [8] C. Meyer, A. Treiman, and T. Kostiuk. Proceedings of the 1996 planetary surface instruments workshop. LPI Technical Report 95-05, Lunar and Planetary Institute, NASA, 1996. Workshop held May 12-13, 1995, Houston, Texas.
 - [9] Tom M Mitchell. *Machine Learning*. WCB/McGraw-Hill, 1997.
 - [10] G. Theraulaz, S. Goss, J. Gervet, and J. L. Deneubourg. Task differentiation in polistes wasp colonies: a model for self-organizing group of robots. In J. A. Meyer and S. Wilson, editors, *Simulation of Adaptive Behaviour: from Animals to Animats*. MIT Press, 1990.
 - [11] International Space University. International lunar farside observatory and science station. WWW, 1993. <http://exploration.jsc.nasa.gov/explore/Data/Lib/DOCS/EIC039.HTML>.
 - [12] Gerhard Weiss, editor. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, Cambridge, Massachusetts, 1999. QA 76.76.I5 M85 1999.