Annals of Mathematics and Artificial Intelligence **31:** 1–5, 2001.

Guest Editorial

From **Ants** to **A(ge)nts**: A Special Issue on Ant-Robotics

Israel A. Wagner [a,b] and Alfred M. Bruckstein [b]

[a] *IBM Haifa Research Lab, Matam, Haifa 31905, Israel*
E-mail: wagner@il.ibm.com
[b] *Department of Computer Science, Technion City, Haifa 32000, Israel*
E-mail: freddy@cs.technion.ac.il

1. Introduction

It was said long ago:

> *Go to the ant, thou sluggard; consider her ways, and be wise: Which having no guide, overseer, or ruler, Provideth her meat in the summer, and gathereth her food in the harvest.* (Proverbs vi 6–8)

Indeed, ants are most fascinating social insects. These tiny creatures live in societies as complex as ours, and are the most abundant and resilient creatures on the earth. The communities of these myopic creatures capable of short-range interactions achieve amazing feats of global problem solving and pattern formation. It should come as no surprise that ant-behavior became a source of algorithmic ideas for distributed systems where a computer, or a robot, is the "individual" and the network, or a swarm of robots, plays the role of the "colony". The flow of ideas from insects to machines is becoming increasingly popular, see, e.g., [2]. Basic lessons learned from the ants are achieving reliability through redundancy, relying on decentralization via some "individual utility" optimization and transforming the environment into a grand shared database by active marking. Research into ant-robotic or distributed a(ge)nt systems can be classified into analysis and synthesis topics. In analysis, a rule of behavior is postulated and one attempts to predict or, more often, simulate the global behavior, while in synthesis a global goal is given and one must find a proper local rule the achieves it. Naturally, most of the body of research existing today is concerned with the considerably simpler analysis issues. Systematic and mathematically proved methods of analysis are still rare, and synthesis examples are almost nonexistent.

Recently ideas about simple a(ge)nts that cooperate to achieve global goals were pursued by many, among them [3,4,10]. Only in a few cases mathematical results exhibiting explicit benefits of cooperation have been found. Experimental research is still the universal tool of analysis and discovery, and it should be so. However, the last years have seen many attempts at formalizing and proving results on multi-agent interaction. The mathematics needed is always interesting, and interdisciplinary in nature, and we are very excited to see these new developments. Mathematical approaches to ant-robotics are not readily available and often rely on combining insights from many different points of view, and this is increasingly difficult in our age of (often sterile) over-specialization.

2. This special issue

In this special issue of AMAI, we have collected several papers dealing with multi-a(ge)nt interaction. In our call for papers, we tried to emphasize formal approaches, but we believe that an overly formalistic viewpoint can leave out a wealth of beautiful phenomenologial/experimental results. True interdisciplinarity in our area must incorporate and combine the study of real animal societies and the engineering of systems with multiple agents acting under communication and sensing constraints. There are ten papers in this issue, which fall into three main categories: *robotic covering*, *path following* and *new paradigms* related to ant-robotics.

2.1. Robotic covering

Covering is an important problem in robotics, having applications from floor cleaning to lawn mowing and field demining. Covering has many variations, which are related to traversal, patrolling, mapping and exploration.

Multi-robot collaboration for robust exploration is the topic of a paper by Rekleitis, Dudek and Milios. They present a new sensing modality for multirobot exploration, based on using a pair of robots that observe each other, and act in concert to reduce odometry errors. The proposed approach improves the quality of the map by reducing the inaccuracies that occur over time from dead reckoning errors. Two different algorithms, based on the size of the environment, are introduced, with a complexity analysis and experimental results in both simulation and with real robots.

Ant-robotic methods are efficient in some applications but not in all applications. The paper by Koenig, Szymanski and Liu presents efficient and inefficient ant coverage methods. They study, both theoretically and in simulation, the behavior of ant robots for one-time or repeated coverage of terrain, as required for various applications. As ant robots cannot use conventional planning methods due to their limited sensing and computational capabilities, the study concentrates on navigation methods that are based on real-time (heuristic) search and leave markings in the terrain, similar to what real ants do. Their analysis is surprising. It shows that the cover time of ant robots that use one of the real-time search methods is guaranteed to be polynomial in the number of locations, whereas the cover time of ant robots that use the other real-time search method can be exponential in (the square root of) the number of locations even in simple terrains that correspond to (planar) undirected trees.

A systematic approach for practical covering is the topic of the paper "Spanning-tree based coverage of continuous areas by a mobile robot", by Gabriely and Rimon. They consider the problem of covering a continuous planar area by a square-shaped tool attached to a mobile robot, using a tool-based approximation of the work-area. The algorithm, called *Spanning Tree Covering* (STC), subdivides the work-area into disjoint cells and then follows a spanning tree of the graph induced by the cells, while covering every point precisely once. Three versions of the STC algorithm are analyzed. An off-line version finds optimal path in time $O(N)$, N being the number of cells comprising the area. A second version of STC is on-line, where the robot uses its sensors to detect obstacles and construct a spanning tree of the environment while covering the work-area, completing an optimal covering path in time $O(N)$, but requiring $O(N)$ memory. The third version of STC is *ant*-like. In this version, too, the robot has no apriori knowledge of the environment, but it may leave pheromone-like markers during the coverage process. The ant-like STC algorithm runs in time $O(N)$, and requires only $O(1)$ memory. Simulation results are presented for the three STC algorithms, demonstrating their effectiveness in cases where the tool size is significantly smaller than the work-area characteristic dimension.

A specific advantage of ants is their robustness against changes in the environment. Yanovski, Wagner and Bruckstein address the robustness of ant algorithms in their paper "Vertex ant walk – A robust method for efficient exploration of faulty graphs", where they investigate a decentralized exploration method, called Vertex Ant Walk (VAW), executed by several simple, memoryless agents on a faulty network. A model for the network is a directed graph, and the algorithm is analyzed assuming failures of network edges and nodes. The algorithm is also shown to be *self-stabilizing* in the sense that it can be started with any initialization, and *scalable* – new agents can be added while other agents are already running.

The group of papers about covering ends with an overview paper by Choset, entitled "Coverage for robotics – A survey of recent results". This paper surveys recent results in *coverage path planning*, a new path planning approach that determines a path for a robot to pass over all points in its free space. The paper conjectures that most complete algorithms use an exact cellular decomposition, either explicitly or implicitly, to achieve coverage. Therefore, the survey organizes the coverage algorithms into four categories: heuristic, approximate, partial-approximate and exact cellular decompositions. Finally it describes some provably complete multi-robot coverage algorithms.

2.2. Path following

Ants, and other social insects, are known to use pheromones to navigate and exchange information and achieve amazing pattern formations (e.g., [1]) . One particularly spectacular example are the trails created by an ants between food sources and their nests. Mathematical problems and puzzles of pursuit, and cyclic pursuit in particular, have attracted interest for many years (e.g., [5–7]). They are, in our opinion, quite relevant to ant-robotics due to the fact that simple local rules of interaction are involved and they lead to global results.

Two papers by Richardson describe new analytic findings about a(ge)nt pursuit. In the first one, "Non-mutual captures of cyclic pursuit", Richardson answers an open question regarding k-dimensional pursuit. He shows that it is possible for bugs to capture their prey without *all* bugs simultaneously doing so, even for non-collinear initial positions. The set of initial conditions which give rise to non-mutual captures is a submanifold in the manifold of all possible initial conditions. Hence, if the initial positions are picked randomly according to a smooth probability distribution, then the probability that a non-mutual capture will occur is zero.

In a related paper, titled "Stable polygons of cyclic pursuit", Richardson studies the stability of regular geometries in cyclic pursuit, and shows that in all dimensions the only stable regular n-bug shapes are the regular two dimensional n-gons, $n \geqslant 7$, in which each vertex chases its neighboring vertex in some fixed orientation. He also analyzes the three bug cyclic pursuit in detail, proving that, except for the equilateral initial position, the triangle formed is asymptotically degenerate with the minimum interior angle tending to zero while the vertex at which the minimum is located rotates among the vertices infinitely often.

2.3. New paradigms

As ant-robotics is a relatively new subject which is still in its "annealing" phase, new and very interesting paradigms are evolving which certainly deserve attention and study.

The paper "From insect to Internet: Situated control for networked robot teams" by Werger and Mataric studies the exploitation of *abstract situatedness* – situatedness in non-physical environments. They focus on the problem of *role assumption*, distributed task allocation in which each agent selects its own task-performing role. The paper details the authors' general *Broadcast of Local Eligibility* (BLE) technique for role-assumption in such behavior-space-situated systems, and provides experimental results from the *Cooperative Multi-robot Observation of Multiple Moving Targets* (CMOMMT) target-tracking task.

Another paradigmatic insight into task allocation as observed in the behavior of ant colonies is suggested by Hirsh and Gordon in their paper "Distributed problem solving in social insects". They review the highly parallel and distributed form of information processing involved in task allocation, discuss its potential sophistication, its actual performance in various groups of social insects, its general strengths and liabilities, and finally, the adaptations that compensate for these liabilities. As a computational model they use Finite State Machines (FSM), a model which makes their result accessible for computer scientists who seek for inspiration in the world of insects.

A rather different point of view is suggested by De Schutter, Theraulaz and Deneubourg. In their paper "Animal-robots collective intelligence" they argue in favor of investigating mixed groups of interacting animals and robots. Using the fact that it is possible to interact with animals not only by mimicking their behaviors but also by making specially designed artifacts, and the theory of self-organization in animal societies that shows that very simple, but numerous, interactions between individuals may ensure complex performances and produce *Collective Intelligence* (CI) at the level of the group, the authors develop an experiment using an artifact interacting within a collective intelligence system in the wild (gull flocks), and similar cases for robots. Using the experiments, they discuss the expected difficulties in mixing robots and animals in CI systems, as well as the insights which can be gained from such systems.

All in all, we hope that this special issue will stimulate new lines of thought and interactions among entomologists, computer scientists, and engineers, and will make us all "go to the ant, consider its ways and be wise".

3. Special thanks

We deeply thank the following people for reading and reviewing papers for this special issue: Sigurd Angenent, Estie Arkin, Tucker Balch, David Carmel, Howie Choset, Nir Cohen, Jean-Louis Deneubourg, Marco Dorigo, Gregory Dudek, Susan Hert, Owen Holland, Eric Klavins, Sven Koenig, Yaxin Liu, Evangelos E. Milios, Jacques M. Pasteels, Elon Rimon, Barry Werger, Vladimir Yanovski, and two additional referees who preferred to remain anonymous.

Many thanks to the Editor-in-Chief of AMAI, Professor Martin C. Golumbic, for inviting us to organize and guest edit this special issue and for his handling the refereeing of our own paper.

References

[1] F.R. Adler and D.M. Gordon, Information collection and spread by networks of patrolling ants, The American Naturalist 140(3) (1992) 373–400.
[2] E. Bonabeau and G. Theraulaz, Swarm smarts, Scientific American (March 2000) 54–61.
[3] V. Braitenberg, *Vehicles* (MIT Press, Cambridge, MA, 1985).
[4] R.A. Brooks, Elephants don't play chess, in: *Designing Autonomous Agents*, ed. P. Maes (MIT Press/Elsevier, 1990) pp. 3–15.
[5] A.M. Bruckstein, N. Cohen and A. Efrat, Ants, crickets and frogs in cyclic pursuit, CIS Report # 9105, Technion, Israel (1991).
[6] A.M. Bruckstein, Why the ant trails look so straight and nice, The Mathematical Intelligencer 15(2) (1993) 59–62.
[7] A.M. Bruckstein, C.L. Mallows and I.A. Wagner, Probabilistic pursuits on the integer grid, American Mathematical Monthly 104(4) (1997) 323–343.
[8] D.M. Gordon, *Ants at Work: How an Insect Society is Organized* (W.W. Norton and Company, 2000).
[9] B. Holldobler and E.O. Wilson, *The Ants* (Belknap Press, 1990).
[10] M. Resnick, *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds* (MIT Press, Cambridge, MA, 1994).
[11] I.A. Wagner, M. Lindenbaum and A.M. Bruckstein, Smell as a computational resourse – A lesson we can learn from the ant, in: *Proceedings of the 4th Israeli Symposium on Theory of Computing and Systems* (ISTCS), Jerusalem, June 10–11 (1996) pp. 219–230. Web-accessible through: `http://www.cs.technion.ac.il/~wagner`.
[12] I.A. Wagner, M. Lindenbaum and A.M. Bruckstein, Distributed covering by ant-robots using evaporating traces, IEEE Transactions on Robotics and Automation 15(5) (1999) 918–933.
[13] I.A. Wagner, M. Lindenbaum and A.M. Bruckstein, Efficient graph search by a smell-oriented vertex process, Annals of Mathematics and Artificial Intelligence 24 (1998) 211–223.