

Insect Societies and Manufacturing

Vincent A. Cicirello and Stephen F. Smith

The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
{cicirello, sfs}@cs.cmu.edu

Abstract

In this paper, we present examples from the literature of successful problem solving systems that have been heavily influenced directly from biological studies of insect societies. Included among these systems are a few of our own. These examples are presented in light of manufacturing applications and only the tip of the iceberg is touched upon. It is our hope that the reader will find a new source of inspiration in problem solving tool design; and realize the potential of coordination mechanisms for multi-agent manufacturing systems inspired by social insects.

1 Introduction

“Could there be an Artificial Ant Colony?” is a question asked by Douglas R. Hofstadter in one chapter of his famous book *Gödel, Escher, Bach: an Eternal Golden Braid* [Hofstadter, 1979]. In that chapter, Hofstadter contemplates on a dialogue in which the character Anteater is telling Achilles about his friend Aunt Hillary. Aunt Hillary is an ant colony seen as a single intelligent entity. The discussion is whether or not collective intelligence can emerge from the interactions of hundreds of simple less-than-intelligent agents and whether or not we can learn anything from studying such natural multi-agent systems at work.

Colonies of social insects have collaboratively solved such problems as resource transportation, routing, task allocation for millenia in efficient manners and there is a great deal that we can learn from them. A little over twenty years have passed since Hofstadter’s book was first published and we now have many examples of researchers applying coordination mechanisms based on models of such natural multi-agent systems as ants [Hölldobler and Wilson, 1990], bees [Kirchner and Towne, 1994; von Frisch, 1967], and wasps [Theraulaz *et al.*, 1991] to solve fairly difficult problems.

In this paper, we present examples from the literature of successful problem solving systems that have been heavily influenced directly from biological studies of insect societies. These examples are presented in light of manufacturing applications and only the tip of the iceberg is touched upon. In what follows you will find no equations; those can be found in the cited papers themselves. But rather, it is our hope that you

will find a new source of problem solving tools with which to use in tackling your manufacturing problems.

The remainder of this paper is organized as follows. Section 2 presents applications of ant colony coordination including applications to combinatorial optimization problems (Section 2.1) as well as to other more dynamic problems (Section 2.2). In Section 3 we discuss less studied ideas for task allocation (Section 3.1) and task prioritization (Section 3.2) problems based on wasp behavior. And finally, we conclude in Section 4.

2 Coordination in Ant Colonies

An ant colony is an example of a highly distributed natural multi-agent system. It is comprised of hundreds (perhaps thousands) of completely autonomous simple agents (i.e., the ants) and is robust with respect to loss of individual agents and to changes in the environment. The colony continues to function efficiently despite the loss of a few individual ants to predators for example and adapts foraging activity to account for unexpected changes to the local environment. The colony as a whole manages to coordinate its activities – such as foraging, brood care, cemetery formation, etc. – without any direct communication among the individual members. Furthermore, the performance of some of these activities can be seen as highly effective and in some cases optimal or near-optimal. For example, consider the well-known double bridge experiment [Bonabeau *et al.*, 1999]. In this experiment, there are two paths between the nest and a food source. Over time the colony is able to converge upon the use of the shorter of the two paths. The typical biological explanation for this successful outcome is that of pheromone trail laying and following [Hölldobler and Wilson, 1990]. Ants are believed to deposit chemical along the route between the nest and a food source as well as to wander (somewhat at random) in the general direction of heavier concentrations of such pheromone. Shorter paths begin building up concentrations of pheromone at a faster and faster rate as more and more ants begin using the shorter paths; while concentrations on other longer paths decay over time. The explanation is that an ant taking the shorter path both ways will doubly reinforce the path before an ant which took the longer route has a chance to doubly reinforce its path. Thus more ants begin taking this shorter path sooner further reinforcing the short path.




Figure 1: In the Ant Colony System (ACS) of Dorigo et al. as well as in the many other variants of Ant Colony Optimization (ACO) for the Traveling Salesperson Problem (TSP), a population of “ants” stochastically build tours of the cities based on artificial pheromone and heuristic information. The pheromone quantities are updated according to tour quality and the entire process iterated.

2.1 Combinatorial Optimization

Based on this ant foraging analogy, Dorigo et al. have founded a new paradigm of evolutionary computation that has come to be known as Ant Colony Optimization (ACO) [Dorigo and Di Caro, 1999]. ACO is perhaps the most successful application inspired by ants – and perhaps by any insect society. Their initial Ant System (AS) [Colomi et al., 1992b; 1992a; Dorigo et al., 1996] showed promising results with the Traveling Salesperson Problem (TSP). And they later refined their approach in their Ant Colony System (ACS) [Dorigo and Gambardella, 1997a; 1997b] and through the addition of some local search procedures this ACS is competitive with some of the best known heuristics for the TSP. The general idea behind the class of ACO algorithms for the TSP is as follows (also see Figure 1). Begin by laying some initial quantities of artificial pheromone on the edges between the cities in your TSP. Now set loose a population of artificial ants. Each ant builds a tour of the cities by stochastically following the pheromone and heuristic information captured in the city-to-city distances. Next the pheromone levels are updated based on the quality of the tours found by this population. After this pheromone update is performed, the population of ants is set loose again. And this continues until a single path is converged upon or for some large number of iterations.

Since its development, researchers have applied ACO to a number of combinatorial optimization problems. One area particularly important to manufacturing where there has been a lot of work is scheduling. There is at least one ACO-based approach to almost any scheduling problem you can think of. Some of these include: the sequential ordering problem [Gambardella and Dorigo, 1997], job shop scheduling [van der Zwaan and Marques, 1999], flow shop scheduling [Stützle, 1998], vehicle routing [Bullnheimer et al., 1999; Gambardella et al., 1999], bus driver scheduling [Forsyth and Wren, 1997], tardiness scheduling problems [Bauer et al., 1999; den Besten et al., 2000], and resource-constrained project scheduling [Merkle et al., 2000]. In most of these cases, the scheduling problem at hand is reduced to a TSP-like problem in which the problem is to find some optimal




Figure 2: In the Ant Colony Control (AC²) system of [Cicirello and Smith, 2001a], ant-like agents are assigned to jobs as they arrive at the factory. These “ants” make all routing decisions for their respective jobs stochastically biased according to quantities of artificial pheromone.

path through a graph of some sort.

Scheduling problems need not be the only manufacturing related problems for which ACO is useful. An ACO-based algorithm can potentially be designed for any combinatorial optimization problem – although easier for some than for others. For instance, ACO seems particularly well suited to problems that require finding some optimal ordering of something. Many combinatorial problems often appear in the domain of manufacturing. For example, the largest common subgraph problem (LCSP) was used by Cicirello and Regli [Cicirello and Regli, 1999; 2001] for retrieving CAD models based on similarity. They used a hill climbing based approach to solve the problem. But perhaps ACO can be of use here (to my best knowledge ACO has not yet been used for the LCSP). If you arbitrarily order the nodes of one graph, the problem reduces to finding an ordering of the nodes of the second graph such that the implied node mapping of the respective orderings defines the largest common subgraph. And thus we have the problem in a form for which some variant of ACO should be well suited.

2.2 Dynamic Problems

The problems discussed so far that can be solved with ant-based algorithms have all been static problems. That is, the problem is defined at the beginning and does not change during problem solving. However, ant-based algorithms are not limited to static problems. For example, based on ACO Schoonderwoerd et al. [Schoonderwoerd et al., 1997a; 1997b] have developed their Ant-Based Control (ABC) system for optimizing the routing tables in circuit-switched networks. And similarly, Di Caro and Dorigo [Di Caro and Dorigo, 1998a; 1998b] incorporate ACO techniques into their AntNet system for the optimization of package-switched network routing tables. These are both systems in which the state of the world is extremely dynamic.

Related to the domain of manufacturing, in [Cicirello and

Smith, 2001a] we presented an ACO-inspired algorithm for shop floor routing in a dynamic factory setting (see Figure 2). As they arrive in the factory, jobs are assigned to ants. These ants are in charge of making all routing decisions for their respective jobs. For single job type factories, our system is able to effectively balance machine load. For multiple job types, multipurpose machines, and sequence-dependent setups, by using multiple ant colonies (one for each job type) the system is able to converge to specialized product lines.

Although not based on ACO, in [Vaughan *et al.*, 2000a; 2000b], Vaughan *et al.* present a method based on ant trail following behavior and the dance behavior of bees for coordinating resource transportation by a team of robots in uncertain environments. Vaughan *et al.* list robotic supply columns in FMS as a potential application of their system. In their system, robots begin by searching randomly for the goal location. Upon attaining the goal, they return to the nest with a quantity of resource contained there. As a robot reaches the nest it announces to its teammates the path that it took. This is analogous to the way a foraging bee conveys the location of a food source to the hive upon returning with food through its dance language [Kirchner and Towne, 1994; von Frisch, 1967]. This path information is used by the robots to update a set of what Vaughan *et al.* call “crumbs”. Now when foraging, the robots try to go in the average direction of all nearby “crumbs” in an analogous way to ants following pheromone trails.

Ant trail following behavior is not the only behavior of ants that has served as a useful basis for multi-agent coordination mechanisms. Deneubourg *et al.* [Deneubourg *et al.*, 1991] used a model of the brood sorting behavior of ants to coordinate a team of robots in the task of sorting different objects by color. Beckers *et al.* [Beckers *et al.*, 1994] applied a model of ant cemetery formation to coordinate multiple robots in the clustering of objects. Parunak and Brueckner [Parunak and Brueckner, 1999] have used methods based on ant coordination mechanisms for the missionaries and cannibals problem.

3 Coordination in Wasp Colonies

In [Theraulaz *et al.*, 1991], Theraulaz *et al.* present a model for the self-organization that takes place within a colony of wasps. Interactions between members of the colony and the local environment result in dynamic distribution of tasks such as foraging and brood care. In addition, a hierarchical social order among the wasps of the colony is formed through interactions among individual wasps of the colony. This emergent social order is a succession of wasps from the most dominant to the least dominant. Theraulaz *et al.* model these two aspects of wasp behavior as distinct behaviors without making any explicit connection between the two. We have used these two independent models of wasp behavior to develop coordination mechanisms for a multi-agent factory control system where each machine is modeled as a wasp nest comprised of wasp-like agents which take part in routing and scheduling activities as seen in Figure 3.

3.1 Applications of Task Allocation Behavior

The model of [Theraulaz *et al.*, 1991] describes the nature of interactions between an individual wasp and its local en-




Figure 3: Machines in the factory are modeled as wasp nests comprised of colonies of wasp-like agents in charge of scheduling and routing activities. (Figure originally appeared in [Cicirello and Smith, 2001d].)

vironment with respect to task allocation. They model the colony’s self-organized allocation of tasks using what they refer to as response thresholds. An individual wasp has a response threshold for each zone of the nest. Based on a wasp’s threshold for a given zone and the amount of stimulus from brood located in that zone, a wasp may or may not become engaged in the task of foraging for that zone. A lower threshold for a given zone amounts to a higher likelihood of engaging in activity given a stimulus. Bonabeau, Theraulaz, and Deneubourg discuss in [Bonabeau *et al.*, 1998] a model in which these thresholds remain fixed over time. But in [Theraulaz *et al.*, 1998], a threshold for a given task decreases during time periods when that task is performed and increases otherwise. Bonabeau *et al.* [Bonabeau *et al.*, 1997] demonstrate how this model leads to a distributed system for allocating mail retrieval tasks to a group of mail carriers.

In [Cicirello and Smith, 2001d; 2001b], we present an approach to dynamic shop floor routing based on this computational model of wasp task allocation (see Figure 4 for an illustration of the approach). In factory environments comprised of multipurpose machines and sequence-dependent setups, our system successfully allocates jobs to the various machines, constrained by the job mix, so as to limit the amount of required setup time accrued by the system and thus optimize system throughput performance based on current product demands.

But this is not the only potential use for this wasp model of task allocation. One of the goals of Theraulaz *et al.* when they originally presented this model in [Theraulaz *et al.*, 1991] was to point out its potential use in coordinating teams of robots. Given a group of agents, whether they be hardware or software agents, with similar or overlapping capabilities, this model can be of great use in the self-coordination of task assignment amongst the agents. Manufacturing systems based

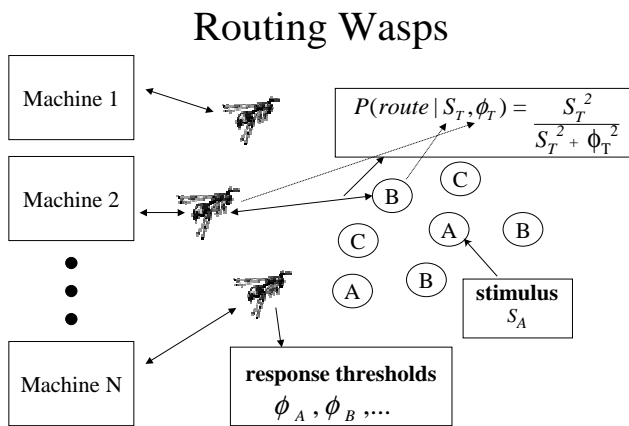


Figure 4: Routing wasps are wasp-like agents used for distributed assignment of jobs to machines based on the self-allocation of tasks among real wasps. Each wasp-like agent is in charge of routing jobs to its respective machine and maintains response thresholds for the different types of jobs its machine can process. (Figure originally appeared in [Cicirello and Smith, 2001b].)

on a multi-agent design paradigm can greatly benefit from such a model in distributing tasks among its various component agents.

3.2 Applications of Dominance Hierarchies

The model of [Theraulaz *et al.*, 1991] also describes the nature of wasp-to-wasp interactions that take place within the nest. When two individuals of the colony encounter each other, they may with some probability interact with each other. If this interaction takes place, then the wasp with the higher social rank will have a higher probability of dominating in the interaction. Social rank is modeled by what they call the individual's force variable F . And the successful wasp is chosen stochastically based on the force variables of the two wasps. The value of F is increased for the successful wasp and decreased for the unsuccessful. Through such interactions as these, wasps within the colony self-organize themselves into a dominance hierarchy. In [Theraulaz *et al.*, 1995], Theraulaz, Bonabeau, and Deneubourg discuss in greater detail the self-organization of dominance hierarchies among wasps. For example, they discuss a number of ways of modeling the probability of interaction during an encounter which range from always interacting to interacting based upon certain tendencies of the individuals.

To my knowledge, there are currently only two applications of this part of Theraulaz *et al.*'s model. The first we present in [Cicirello and Smith, 2001d; 2001c] where we use the concept of force to model job priority in a dynamic factory scheduling environment, and exploit this wasp-to-wasp interaction model to prioritize jobs in a given queue (see Figure 5 for illustration). High priority jobs correspond to high ranking wasps in the social hierarchy of the nest. In [Cicirello and Smith, 2001d] we were trying to optimize system throughput under the constraint of sequence-dependent setups. Currently, we are extending this model to the weighted tardiness

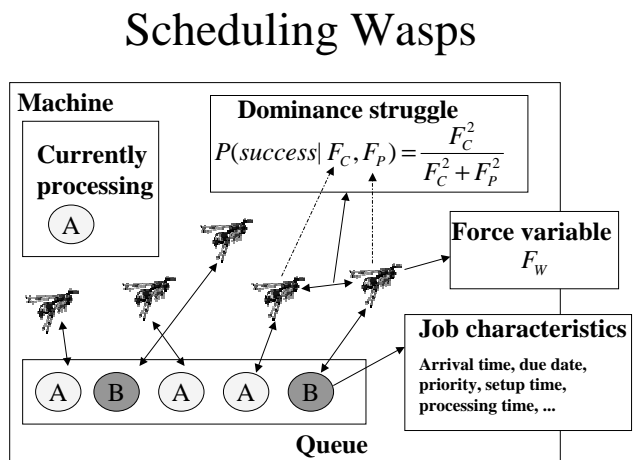


Figure 5: Scheduling wasps are wasp-like agents that use a model of the self-organized social hierarchies of real wasps to prioritize jobs in a machine's queue. Each wasp-like agent is in charge of a job in the queue and takes part in dominance contests with other scheduling wasps. The outcome of these contests determines the priorities of the respective jobs. (Figure originally appeared in [Cicirello and Smith, 2001c].)

objective [Cicirello and Smith, 2001c].

The second application we present in [Cicirello and Smith, 2001b] as a refinement to our shop floor routing approach. If two or more machines want the same job, they take part in a tournament of dominance contests to decide the winner. So in a sense, despite no connection between the two biological models of wasp behavior for task allocation and social hierarchy formation, we have combined aspects of both into an effective system for task assignment in a factory environment.

4 Conclusion

In this paper we presented a brief overview of example systems that have at their foundation the coordination mechanisms of insect societies. Some of these systems deal simply with statically defined manufacturing related problems; but others take a more active decision-making role in dynamic factory environments. From these examples, it is clear that computational models of complex biological systems offer promise as a means of designing and coordinating manufacturing processes. This makes sense since manufacturing systems and processes inherently require effective interaction among multiple agents in dynamic environments and the mechanisms employed by insect societies have survived the evolutionary cycle of survival of the fittest for millenia. We hope that you are now convinced that social insect coordination models deserve a place in the toolbox of the AI researcher working in the domain of manufacturing; or at the very least, we hope that you believe the answer to Hofstadter's question which began this paper is "Yes."

Acknowledgments

The drawing of the wasp that appeared in Figure 3 is courtesy of Deborah Cicirello.

This work has been funded in part by the Department of Defense Advanced Research Projects Agency and the U.S. Air Force Rome Research Laboratory under contracts F30602-97-2-0066 and F30602-00-2-0503 and by the CMU Robotics Institute. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force or U.S. Government.

References

- [Bauer *et al.*, 1999] A. Bauer, B. Bullnheimer, R. F. Hartl, and C. Strauss. An ant colony optimization approach for the single machine total tardiness problem. In *CEC99: Proceedings of the Congress on Evolutionary Computation*, pages 1445–1450, July 1999.
- [Beckers *et al.*, 1994] R. Beckers, O. E. Holland, and J. L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In R. A. Brooks and P. Maes, editors, *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 181–189, 1994.
- [Bonabeau *et al.*, 1997] E. Bonabeau, A. Sobkowski, G. Theraulaz, and J. L. Deneubourg. Adaptive task allocation inspired by a model of division of labor in social insects. In D. Lundh and B. Olsson, editors, *Bio Computation and Emergent Computing*, pages 36–45. World Scientific, 1997.
- [Bonabeau *et al.*, 1998] E. Bonabeau, G. Theraulaz, and J. L. Deneubourg. Fixed response thresholds and the regulation of division of labor in insect societies. *Bulletin of Mathematical Biology*, 60:753–807, 1998.
- [Bonabeau *et al.*, 1999] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Santa Fe Institute Studies in the Sciences of Complexity. Oxford University Press, 1999.
- [Bullnheimer *et al.*, 1999] B. Bullnheimer, R. F. Hartl, and C. Strauss. An improved ant system algorithm for the vehicle routing problem. *Annals of Operations Research*, 89:319–328, 1999.
- [Cicirello and Regli, 1999] V. A. Cicirello and W. C. Regli. Resolving non-uniqueness in design feature histories. In W. F. Bronsvort and D. C. Anderson, editors, *Proceedings of the Fifth Symposium on Solid Modeling and Applications*, pages 76–84, New York, June 9-11 1999. ACM SIGGRAPH, ACM Press.
- [Cicirello and Regli, 2001] V. A. Cicirello and W. C. Regli. Machining feature-based comparisons of mechanical parts. In *SMI 2001: International Conference on Shape Modelling and Applications*. IEEE Computer Society Press, May 2001.
- [Cicirello and Smith, 2001a] V. A. Cicirello and S. F. Smith. Ant colony control for autonomous decentralized shop floor routing. In *ISADS-2001, Fifth International Symposium on Autonomous Decentralized Systems*. IEEE Computer Society Press, March 2001.
- [Cicirello and Smith, 2001b] V. A. Cicirello and S. F. Smith. Improved routing wasps for distributed factory control. In *IJCAI-01 Workshop on AI and Manufacturing*, August 2001.
- [Cicirello and Smith, 2001c] V. A. Cicirello and S. F. Smith. Randomizing dispatch scheduling heuristics, 2001. Submitted to the AAAI Fall Symposium: Using Uncertainty within Computation.
- [Cicirello and Smith, 2001d] V. A. Cicirello and S. F. Smith. Wasp nests for self-configurable factories. In *Agents 2001, Proceedings of the Fifth International Conference on Autonomous Agents*. ACM Press, May-June 2001.
- [Colomi *et al.*, 1992a] A. Colomi, M. Dorigo, and V. Maniezzo. Distributed optimization by ant colonies. In *Proceedings of the First European Conference on Artificial Life*, pages 134–142. Elsevier Publishing, Paris, France 1992.
- [Colomi *et al.*, 1992b] A. Colomi, M. Dorigo, and V. Maniezzo. An investigation of some properties of an “ant algorithm”. In *Proceedings of the Parallel Problem Solving from Nature Conference*, pages 509–520. Elsevier Publishing, Brussels, Belgium 1992.
- [den Besten *et al.*, 2000] M. den Besten, T. Stützle, and M. Dorigo. Ant colony optimization for the total weighted tardiness problem. In *Proceedings of the Parallel Problem Solving from Nature Conference*, 2000.
- [Deneubourg *et al.*, 1991] J. L. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain, and L. Chrétien. The dynamics of collective sorting robot-like ants and ant-like robots. In *From Animals to Animals: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 356–363. MIT Press, 1991.
- [Di Caro and Dorigo, 1998a] G. Di Caro and M. Dorigo. AntNet: Distributed stigmergetic control for communications networks. *Journal of Artificial Intelligence Research*, 9:317–365, 1998.
- [Di Caro and Dorigo, 1998b] G. Di Caro and M. Dorigo. Mobile agents for adaptive routing. In *Proceedings of the 31st Hawaii International Conference on Systems Science*, pages 74–83, January 1998.
- [Dorigo and Di Caro, 1999] M. Dorigo and G. Di Caro. The ant colony optimization meta-heuristic. In D. Corne, M. Dorigo, and F. Glover, editors, *New Ideas in Optimization*. McGraw-Hill, 1999.
- [Dorigo and Gambardella, 1997a] M. Dorigo and L. M. Gambardella. Ant colonies for the traveling salesman problem. *BioSystems*, 43:73–89, 1997.
- [Dorigo and Gambardella, 1997b] M. Dorigo and L. M. Gambardella. Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1):53–66, 1997.
- [Dorigo *et al.*, 1996] M. Dorigo, V. Maniezzo, and A. Colomi. Ant system: Optimization by a colony of cooperating

- agents. *IEEE Transactions on Systems, Man, and Cybernetics – Part B: Cybernetics*, 26(1):29–41, February 1996.
- [Forsyth and Wren, 1997] P. Forsyth and A. Wren. An ant system for bus driver scheduling. Technical Report 97.25, University of Leeds, School of Computer Studies, July 1997. Presented at the 7th International Workshop on Computer-Aided Scheduling of Public Transport, Boston, July 1997.
- [Gambardella and Dorigo, 1997] L. M. Gambardella and M. Dorigo. HAS-SOP: Hybrid ant system for the sequential ordering problem. Technical Report IDSIA 97-11, Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano, Switzerland, 1997.
- [Gambardella *et al.*, 1999] L. M. Gambardella, É. Taillard, and G. Agazzi. MACS-VRPTW: A multiple ant colony system for vehicle routing problems with time windows. Technical Report IDSIA-06-99, Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano, Switzerland, 1999.
- [Hofstadter, 1979] D. R. Hofstadter. *Gödel, Escher, Bach: an Eternal Golden Braid*. Basic Books, New York, 1979.
- [Hölldobler and Wilson, 1990] B. Hölldobler and E. O. Wilson. *The Ants*. The Belknap Press of Harvard University Press, 1990.
- [Kirchner and Towne, 1994] W. H. Kirchner and W. F. Towne. The sensory basis of the honeybee's dance language. *Scientific American*, 270(6):74–80, June 1994.
- [Merkle *et al.*, 2000] D. Merkle, M. Middendorf, and H. Schmeck. Ant colony optimization for resource-constrained project scheduling. In *GECCO-2000: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 893–900. Morgan Kaufmann, July 2000.
- [Parunak and Brueckner, 1999] H. V. D. Parunak and S. Brueckner. Synthetic pheromones for distributed motion control. In *Advances in Enterprise Control Symposium Proceedings*, pages 151–160. DARPA-ISO, San Diego, CA, November 15-16 1999.
- [Schoonderwoerd *et al.*, 1997a] R. Schoonderwoerd, O. Holland, and J. Bruten. Ant-like agents for load balancing in telecommunications networks. In *Agents '97, Proceedings of the First International Conference on Autonomous Agents*, pages 209–216. ACM Press, 1997.
- [Schoonderwoerd *et al.*, 1997b] R. Schoonderwoerd, O. Holland, J. Bruten, and L. Rothkrantz. Ant-based load balancing in telecommunications networks. *Adaptive Behavior*, 5(2):169–207, 1997.
- [Stützle, 1998] T. Stützle. An ant approach to the flow shop problem. In *Proceedings of EUFIT'98*, pages 1560–1564, 1998.
- [Theraulaz *et al.*, 1991] G. Theraulaz, S. Goss, J. Gervet, and J. L. Deneubourg. Task differentiation in polistes wasp colonies: A model for self-organizing groups of robots. In *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 346–355. MIT Press, 1991.
- [Theraulaz *et al.*, 1995] G. Theraulaz, E. Bonabeau, and J. L. Deneubourg. Self-organization of hierarchies in animal societies: The case of the primitively eusocial wasp *polistes dominulus christi*. *Journal of Theoretical Biology*, 174:313–323, 1995.
- [Theraulaz *et al.*, 1998] G. Theraulaz, E. Bonabeau, and J. L. Deneubourg. Response threshold reinforcement and division of labour in insect societies. *Proceedings of the Royal Society of London B*, 265(1393):327–335, February 1998.
- [van der Zwaan and Marques, 1999] S. van der Zwaan and C. Marques. Ant colony optimisation for job shop scheduling. In *Proceedings of the '99 Workshop on Genetic Algorithms and Artificial Life GAAL'99*, Lisbon, Portugal, March 1999.
- [Vaughan *et al.*, 2000a] R. T. Vaughan, K. Støy, G. S. Sukhatme, and M. J. Matarić. Whistling in the dark: Cooperative trail following in uncertain localization space. In *Agents 2000, Proceedings of the Fourth International Conference on Autonomous Agents*. ACM Press, 2000.
- [Vaughan *et al.*, 2000b] R. T. Vaughan, K. Støy, G. S. Sukhatme, and M. J. Matarić. Blazing a trail: Insect-inspired resource transportation by a robotic team. In *Proceedings, 5th International Symposium on Distributed Robotic Systems*, October 2000.
- [von Frisch, 1967] K. von Frisch. *The Dance Language and Orientation of Bees*. The Belknap Press of Harvard University Press, 1967. Translated by L. E. Chadwick.