Adaptive Task Allocation Inspired by a Model of Division of Labor in Social Insects

Eric Bonabeau[1], Andrej Sobkowski[2], Guy Theraulaz[3], Jean-Louis Deneubourg[4]

[1] Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA, bonabeau@santafe.edu
[2] ENST, 46 rue Barrault, 75634 Paris Cédex 13, France
[3] Laboratoire d'Ethologie et de Psychologie Animale, CNRS - UMR 5550, Université Paul Sabatier, 118 route de Narbonne, 31062 Toulouse Cédex, France, theraula@cict.fr
[4] Unit of Theoretical Behavioural Ecology, Service de Chimie-Physique, CP 231, Université Libre de Bruxelles, Boulevard du triomphe, 1050 Brussels, Belgium

Abstract. Social insects provide us with a powerful metaphor to create decentralized systems of simple interacting, and often mobile, agents. The emergent collective intelligence of social insects – swarm intelligence – resides not in complex individual abilities but rather in networks of interactions that exist among individuals and between individuals and their environment. In particular, a recently proposed model of division of labor in a colony of primitively eusocial wasps, based on a simple reinforcement of response thresholds, can be transformed into a decentralized adaptive algorithm of task allocation. An application of such an algorithm is proposed in the context of a mail company, but virtually any type of flexible task allocation can be described within the same framework.

1 Introduction

Evidence of the ecological success of social insects can be found almost everywhere [56]. Some of the main reasons for this success are to be looked for in the organization of insect societies, that allows, for example [45], division of labor, specialization, collective regulation, plasticity, mass action responses or the building of complex architectures. Theories of self-organization (SO) [35,43], originally developed in the context of physics and chemistry to describe the emergence of macroscopic patterns out of processes and interactions defined at the microscopic level, can be extended to social insects to show that complex collective behavior may emerge from interactions among individuals that exhibit simple behaviors: in these cases, there is no need to invoke individual complexity [11,21].

These recent developments do not only have consequences on the study of social insects, but provide us with powerful tools to transfer knowledge about social insects to the field of intelligent system design. In effect, a social insect colony is undoubtedly a decentralized problem-solving system, comprised of many relatively simple interacting entities. The daily problems solved by a colony (at least partly) by means of self-organization include finding food [3,4,5,7,10,20,26,27], building or extending a nest [19,22,53,54], efficiently dividing labor among individuals [9,49,56], efficiently feeding the brood [13,30], responding to external challenges, spreading alarm [1,31], etc. Many of these problems have counterparts in engineering and computer science. One of the most important features of social insects is that they can solve these problems in a very flexible and robust way: flexibility allows to adapt to changing environments, while robustness endows the colony with the ability to function eventhough some individuals may fail to perform their task. Finally, social insects have limited cognitive abilities: it is therefore simple to design agents, including robotic agents, that mimic their behavior at some level of description. In short, the modeling of social insects by means of SO can help design artificial distributed problem-solving devices that self-organize to solve problems.The paths to problem solving are not predefined but rather emergent in these systems, and result from interactions among individuals and between individuals and their environment as much as from the behavior of the individuals themselves. Therefore, using a swarm intelligent system to solve a problem requires a thorough knowledge not only of what individual behavior must be implemented but also of what interactions are needed to produce such or such global behavior. A possible path consists in studying how social insects collectively perform some specific tasks, modeling their behavior and using the model as a basis upon which artificial variations can be developed, either by tuning the model parameters beyond the biologically relevant range or by adding non-biological features to the model. This approach is similar to other approaches consisting in imitating the way "nature" (that is, physical or biological systems) solves problems [12,28]. Another possible and apparently very promising pathway is to use economic or financial metaphors to solve problems [36,37,38,51,55].

Some authors have already applied this social-insect-based approach to optimization (for example, ants can find the shortest path to a food source) [16,17,18,25], with at least one significant real-world implementation [2,32,50], or data analysis (ants can cluster dead bodies and sort larvae on the basis of local information [14]) [6,39,42]. The aim of this paper is to show how another model of natural swarm intelligence can be transformed into a useful distributed algorithm. This example is based on a recently proposed model of division of labor in a primitively eusocial species of wasp, and will be transformed into an algorithm allowing flexible task assignment.

2 A model of sociogenesis in social insects

2.1 Division of labor

Division of labor is an important feature of colonial life in many species of social insects [44,49]. One of the most striking aspects of division of labor is plasticity, a property achieved through the workers' behavioral flexibility: the ratios of workers performing the different tasks that maintain the colony's viability and reproductive success can vary (i.e., workers switch tasks) in response to internal perturbations or external challenges. An important question is to understand how this flexibility is implemented at the level of individual workers, which certainly do not possess any global representation of the colony's needs.

2.2 The notion of response threshold

We recently developed a simple model, based on the notion of response threshold [9,46,47,48,49]: individuals start to become engaged in task performance when the level of the task-associated stimuli exceeds their threshold. Differences in response thresholds may either reflect actual differences in behavioral responses, or differences in the way task-related stimuli are perceived. When individuals performing a given task are withdrawn (they have low response thresholds with respect to stimuli related to this task), the associated demand increases and so does the intensity of the stimulus, until it eventually reaches the higher characteristic response thresholds of the remaining individuals, that are not initially specialized into that task; the increase of stimulus intensity beyond threshold has the effect of stimulating these individuals into performing the task. This is exactly what was observed by Wilson [57] in experiments where he artificially reduced the minor:major ratio to below 1:1 and observed a change in the rate of activity within 1 hour of the ratio change: for small ratios, majors engage in tasks usually performed by minors and efficiently replace the missing minors.

Let us give a simple example: the task can be larval feeding, and the associated stimulus larval demand, which is expressed, e.g., through the emission of pheromones. The nature of task-related stimuli may vary greatly from one task to another, and so can information sampling techniques (which may involve direct interactions among workers (trophallaxis, antennation, etc.) [33], nest "patrolling" [41], or more or less random exposure to task-related stimuli). Several experiments support the idea of thresholds in honeybees [47,49] or ants [23,24].

Assume that one task only has to be performed. This task is associated with a stimulus or demand, the level of which increases if it is not satisfied (because the task is not performed by enough individuals, or not performed with enough efficiency). Let X_i be the state of an individual i (X_i=0 corresponds to inactivity, X_i=1 corresponds to performing the task), and θ_i the response threshold of i. An inactive individual starts performing the task with a probability P per unit time:

$$P(X_i = 0 \rightarrow X_i = 1) = \frac{s^2}{s^2 + \theta_i^2}. \qquad (1)$$

The probability that an individual will perform a task depends on s, the magnitude of the task-associated stimulus, that affects the probability of being exposed to it, and on θ_i, the probability of responding to task-related stimuli, given exposure [49]. The choice of the detailed nonlinear functional form (1) is rather arbitray, since any threshold function is expected to generate qualitatively similar results. An active individual gives up task performance and becomes inactive with probability p per unit time (that we take identical for both castes):

$$P(X_i = 1 \rightarrow X_i = 0) = p. \qquad (2)$$

1/p is the average time spent by an individual in task performance before giving up this task. It is assumed that p is fixed, and independent of stimulus. Individuals give up task performance after 1/p, but may become engaged again immediatly if stimulus is still large.Variations in stimulus intensity are due to task performance, which reduces stimulus intensity, and the natural increase of demand irrespective of whether or not the task is performed. The resulting equation for the evolution of stimulus intensity s is therefore (in discrete time):

$$s(t+1) = s(t) + \delta - \frac{\alpha N_{act}}{N}, \tag{3}$$

where N_{act} is the number of active individuals, N is the total number of potentially active individuals in the colony, δ is the increase in stimulus intensity per unit time, and α is a scale factor measuring the efficiency of task performance. The amount of work performed by active individuals is scaled by N to reflect the intuitive idea that the demand is an increasing function of N, that we take linear here. This corresponds to the idea that brood should be divided by 2 when colony size is divided by 2: in other words, colony requirements scale (more or less) linearly with colony size. This simple fixed threshold model shows remarkable agreement with experimental results in the case where there are two castes characterized by two different values of the response threshold: when "minors", with low response thresholds, are removed from the simulated colony, "majors", with higher response thresholds, start to perform tasks usually performed by minors. Fig. 1 shows the fraction of majors engaged in task peformance as a function of the fraction of majors in the colony (θ_1=8, θ_2=1, α=3, δ=1, p=0.2; exact analytical result). This curve is very similar to the one observed by Wilson [57]. This simple model with one task can also be easily extended to the case where there are two or more tasks to perform. Each individual has a set of thresholds, each threshold being associated to the stimuli of a specific task or group of tasks.

Fig. 1. Fraction of majors engaged in task performance as a function of the proportion of majors in the colony.

2.3 Simple learning

Theraulaz et al. [52] have developed a model which is very close to the one described in the previous paragraph, but with learning, that is, with variable thresholds, updated in a self-reinforcing way: the more individual i performs a task, the lower its response threshold with respect to stimuli associated with this task, and vice-versa:

$$\theta_i \rightarrow \theta_i - \xi \tag{4}$$

if i performs the task, and

$$\theta_i \rightarrow \theta_i + \varphi \tag{5}$$

if i does not perform the task in a given period of time. ξ and φ are the learning and forgetting coefficients, respectively. The evolution of θ_i is furthermore assumed to be restricted to an interval $[\theta_{min}, \theta_{max}]$. The decisions to become engaged in, or give up, task performance are still given by equations (1) and (3), and the evolution of the demand s is given by (3). Their model is therefore a model of the genesis of specialization. Fig. 2.1 shows an example of specialization: thresholds are represented as a function of time for 10 individuals (initial distribution of thresholds uniform in $[\theta_{min} = 1, \theta_{max} = 1000]$, α=4, δ=1, p=0.2, ξ=10, φ=1; results obtained by numerically integrating equations derived from transition probabilities). Some of these individuals become highly responsive to task-associated stimuli, while others become only weakly responsive. Fig. 2.2 shows the demand, which can be seen to be kept under control, at a low level. One advantage of this strategy over the fixed threshold strategy is that in the case where individuals with low thresholds are removed, the demand is quickly taken back to a much lower level because individuals with previously high thresholds lower their thresholds and become more responsive to task-associated stimuli. Flexibility and robustness can also be observed when task performance becomes less efficient (lower value of α). For example, Fig. 3.1 shows how thresholds evolve when α suddenly decreases from 5 to 2 at t=200. One can see that the 7 individuals that had high thresholds before t=200 quickly become responsive to task-associated stimuli after t=200, because the demand increases (as is shown in Fig. 3.2). But this increase stimulates individuals in performing the task, so that the demand eventually decreases to a level comparable to its level before t=200 (Fig. 3.2). By contrast, a strategy based on fixed response thresholds is unable to bring the demand back to its non-perturbed level (Fig. 3.3; thresholds are quenched in a random uniform distribution in [1,1000]): this is because the demand has to exceed the average threshold of the group in order to be significantly reduced.

Fig. 2.1. Dynamics of thresholds

Fig. 2.2. Dynamics of demand

Because learning is involved in this model, the relevant timescale is naturally larger than in the previous paragraph. Furthermore, the fact that the initial condition is a uniform distribution of thresholds indicates that this model may be more appropriate for primitively eusocial insects, that is, those eusocial insects in which there is little or no morphological differentiation among individuals: this feature is particularly appealing when it comes to problem-solving, since this model can explicitly account for the (self-organized) genesis of division of labor, and does not require predefined response thresholds. The original model [52] is slightly more complicated than simply a model of division of labor, since it is coupled to a model of self-organizing hierarchies. It is a simplified version of this model which will serve as a basis for the implementation of an adaptive task allocation algorithm.

Fig. 3.1. Dynamics of thresholds (perturbed).

Fig. 3.2. Dynamics of demand (perturbed).

Fig. 3.3. Dynamics of demand (perturbed, fixed thresholds)

3 Adaptive task allocation inspired by division of labor in wasps: the example of mail retrieval

Imagine that a group of mailmen (hereafter called agents), belonging to some express mail company, have to pick up letters in a large city. Customers should not have to wait more than a given amount of time: the aim of the mail company is therefore to allocate the agents to the various demands that appear in the course of the day so as to keep the global demand as low as possible. The probability that individual i, located in zone z(i), responds to a demand S_j in zone j is given by

$$P_{ij} = \frac{S_j^2}{S_j^2 + \alpha\theta_{i,j}^2 + \beta d_{z(i),j}^2} \qquad (6)$$

where $\theta_{i,j}$ ($\in [\theta_{min}, \theta_{max}]$) is the response threshold of agent i to a demand from zone j, $d_{z(i),j}$ is the distance between z(i) and j (this distance can either be euclidian or include factors such as one-ways, lights, usual traffic jams, etc.), and α and β are two positive coefficients that modulate the respective influences of θ and d. Each time an agent allocates himself to zone j to retrieve mail, his response thresholds are updated in the following way:

$$\theta_{i,j} \to \theta_{i,j} - \xi_0 \tag{7}$$

$$\theta_{i,n(j)} \to \theta_{i,n(j)} - \xi_1 \tag{8}$$

$$\theta_{i,k} \to \theta_{i,k} + \varphi \text{ for } k \neq j,\ k \notin \{n(j)\} \tag{9}$$

where $\{n(j)\}$ is the set of zones surrounding j, ξ_0 and ξ_1 are two learning coefficients corresponding to zone j and its neighboring zones respectively ($\xi_0 > \xi_1$), and φ is the forgetting coefficient applied to response thresholds associated with other zones. This procedure allows to increase spatial fidelity not only purely locally but also with respect to other zones. We have performed simulations with a grid of 5x5 zones (we consider four neighbors for the update of equation (8), with periodic boundary conditions), and 5 agents; at every iteration, the demand increases in 5 randomly selected zones by an amount of 50, α=0.5, β=500, θ_{min}=0, θ_{max}=1000, ξ_0=150, ξ_1=70, φ=10. Agents are swept in a random order, and decide to respond to the demand from a particular zone according to equation (6). If no agent responds after 5 sweepings, the next iteration starts. If an agent responds, this agent will be unavailable for an amount of time that we take to be equal to the distance separating his current location from the zone where the demand comes from. Once the agent decides to allocate himself to that zone, the associated demand in that zone is maintained at 0 (since any demand emerging between the time of the agent's response and his arrival in the zone will be satisfied by the same agent). Fig. 4.1 shows how the demand fluctuates around a finite value in "normal conditions", and how the demand increases but is still kept under control when one individual fails to perform his task (at t=2000). Fig. 4.2 shows how the threshold of an individual with respect to a single zone can vary as a function of time. A special behavior can be observed after the removal of another individual. The removed individual was a specialist of a given zone: another individual lowers his threshold with respect to that zone and becomes in turn a new specialist of that zone. This is what is observed in Fig. 4.2. However, because the workload may be too high to allow agents to settle into a given specialization, response thresholds may oscillate in time. All these features point to the flexibility and robustness of this algorithm.

Fig. 4.1. Dynamics of demand (perturbed).

Fig. 4.2. Dynamics of one threshold (perturbed).

Although we have briefly presented how the algorithm works on one specific example, it can certainly be modified to apply to virtually any kind of task allocation: the demand S_j can be the abstract demand associated to some task j, $\theta_{i,j}$ is a response threshold of actor i with respect to the task-associated stimulus S_j. Finally, $d_{z(i),j}$ is an abstract distance between i and task j which can, for example, represent the ability or lack of ability of i to deal with task j: if i is not the most efficient actor to perform task j, it will not respond preferentially to S_j, but if no other actor is in a position to respond, it will eventually perform the task. It is certainly possible to design a scheme in which d can vary depending on the efficiency of i in performing task j.

4 Conclusion

The aim of this paper was to draw a clear parallel between a mode of problem-solving in social insects and a distributed, reactive, algorithmic approach. Some of the mechanisms underlying the regulation of division of labor in social insects, particularly those mechanims that can explain the workers' behavioral flexibility, have first been modeled, and the model then served as a basis to regulate the allocation of tasks in an artificial system. Not only have we shown that it is possible to implement in artificial systems, with appropriate modifications, strategies used by social insects to solve problems, we have also shown how modeling is the necessary intermediate step between understanding natural behavior and designing artificial problem-solving devices. The algorithm that we used for our toy example can of course be greatly improved. One possible idea to improve the flexibility of the algorithm would be to allow some kind of partial "reset" function, for example by means of oscillations (as have been observed in ants [15,29,34]): in effect, oscillations allow the system to take new decisions, which is particularly desirable if the environment has been modified within the timescale of one oscillation [8]. If thresholds are regularly allowed to relax partially, the set of thresholds can converge to a new state, possibly more adapted to the environment. Such a process should be especially efficient when the system starts to fail to keep the demand under control. How often and to what extent should the thresholds relax depends on the magnitude and timescale of environmental perturbations.

Finally, let us mention one issue which is common to all adaptive algorithms, that is algorithms that are designed explicitly to deal with dynamically varying problems: there is no standard way of evaluating the performance of such algorithms, because benchmark problems are static problems (such as the traveling salesman problem [40]). It would be most important for adaptive algorithms to have their own benchmarks, as these algorithms may not be very good heuristics in the static case.

Acknowledgements

This work was supported in part by a grant from the GIS (Groupe d'Intérêt Scientifique) Sciences de la Cognition. E. B. is supported by the Interval Research postdoctoral fellowship at the Santa Fe Institute.

References

1. Adler, F. R. & Gordon, D. M. (1992). Information collection and spread by networks of patrolling ants. Am. Nat. **140**, 373.
2. Appleby, S. & Steward, S. (1994). Mobile software agents for control in telecommunications networks. British Telecom Technol. Journal **12**, 104-113.
3. Aron, S., Deneubourg, J.-L., Goss, S. & Pasteels, J. M. (1990). Functional self-organisation illustrated by inter-nest traffic in the Argentine ant Iridomyrmex humilis. In Biological Motion (Alt, W. & Hoffman, G., eds), Lecture Notes in Biomathematics, Springer Verlag.
4. Beckers, R., Deneubourg, J.-L. & Pasteels, J.-M. (1989). Le recrutement de masse : un exemple de communication trop efficace ? Actes Coll. Ins. Soc. **4**, 219-226.
5. Beckers, R., Deneubourg, J.-L., Goss, S. & Pasteels, J.-M. (1990). Collective decision making through food recruitment. Ins. Soc. **37**, 258.
6. Bilchev, G. & Parmee, I. C. (1995). The ant colony metaphor for searching continuous design spaces. In: Lecture Notes in Computer Science (Fogarty, Y., ed.) 993, 25-39, Springer-Verlag.
7. Bonabeau, E. (1995). Marginally stable swarms are flexible and efficient. J. Phys. I France **6**, 309-320.
8. Bonabeau, E. & Cogne, F. (1996). Oscillation-enhanced adaptability in the vicinity of a bifurcation: the example of foraging in ants. In: From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, (Maes, P., Mataric, M., Meyer, J.-A., Pollack, J. & Wilson, S., eds.), 537-544, MIT Press.
9. Bonabeau, E., Theraulaz, G. & Deneubourg, J.-L. (1996). Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies. Proc. Roy. Soc. London B **263**, 1565-1569.
10. Bonabeau, E. (1997). Comment on "Phase transitions in instigated collective decision making". Adaptive Behavior **5**, 99-105.
11. Bonabeau, E., Theraulaz, G., Deneubourg, J.-L., Aron, S. & Camazine, S. (1997). Self-organization in social insects. Trends in Ecol. Evol. (in press).
12. Bounds, D.G. (1987). New Optimization Methods from Physics and Biology. Nature **329**, 215-219.
13. Camazine, S. (1991). Self-organizing pattern-formation on the combs of honeybee colonies. Behav. Ecol. Sociobiol. **28**, 61-76.
14. Chrétien, L. (1996). Organisation spatiale du matériel provenant de l'excavation du nid chez Messor barbarus et des cadavres d'ouvrières chez Lasius niger (Hymenopterae: Formicidae). PhD dissertation, Département de Biologie Animale, Université Libre de Bruxelles.
15. Cole, B. J. (1991). Short-term activity cycles in ants: generation of periodicity by worker interaction. Am. Nat. **137**, 244.
16. Colorni, A., Dorigo, M. & Maniezzo, V. (1991). Distributed optimization by ant colonies. Proc. First Europ. Conf. on Artificial Life (Varela, F. & Bourgine, P., eds), pp. 134-142, MIT Press.
17. Colorni, A., Dorigo, M. & Maniezzo, V. (1992). An investigation of some properties of an ant algorithm. Proc. of PPSN 92 (Parallel Problem Solving from Nature Conference) (Männer, R. & Manderick, B., eds), pp. 509-520, Elsevier Publishing.
18. Colorni, A., Dorigo, M., Maniezzo, V. & Trubian, M. (1993). Ant system for job-shop scheduling. Belg. J. Oper. Res., Stat. and Comput. Sci. **34**, 39-53.
19. Deneubourg, J.-L. (1977). Application de l'ordre par fluctuations à la description de certaines étapes de la construction du nid chez les termites. Ins. Soc. 24: 117-130.

20. Deneubourg, J.-L., Aron, S., Goss, S., Pasteels, J.-M. & Duerinck, G. (1986). Random behaviour, amplification processes and number of participants: how they contribute to the foraging properties of ants. Physica D **22**, 176-186.
21. Deneubourg, J.-L. & Goss, S. (1989). Collective patterns and decision making. Ethol. Ecol. & Evol. **1**, 295-311.
22. Deneubourg, J.-L. & Franks, N. R. (1995). Collective control without explicit coding: the case of communal nest excavation. J. Ins. Behav. **8**, 417-432.
23. Detrain, C. & Pasteels, J. M. (1991). Caste differences in behavioral thresholds as a basis for polyethism during food recruitment in the ant Pheidole pallidula (Nyl.) (Hymenoptera: Myrmicinae). J. Ins. Behav. **4**, 157-176.
24. Detrain, C. & Pasteels, J. M. (1992). Caste polyethism and collective defense in the ant Pheidole pallidula: the outcome of quantitative differences in recruitment. Behav. Ecol. Sociobiol. **29**, 405-412.
25. Dorigo, M., Maniezzo, V. & Colorni, A. (1996). The Ant System: Optimization by a colony of cooperating agents. IEEE Trans. Syst. Man Cybern. B **26**, 1-13.
26. Edelstein-Keshet, L. (1994). Simple models for trail-following behaviour. Trunk trails versus individual foragers. J. Math. Biol. **32**, 303-328.
27. Edelstein-Keshet, L., Watmough, J. & Ermentrout, G. B. (1995). Trail following in ants: individual properties determine population behaviour. Behav. Ecol. Sociobiol. **36**, 119-133.
28. Farmer, J. D., Packard, N. H. & Perelson, A. S. (1986). The immune system, adaptation, and machine learning. Physica D **22**, 187-204.
29. Franks, N. R., Bryant, S., Griffiths, R. & Hemerik, L. (1990). Synchronization of the behaviour within nests of the ant Leptothorax acervorum (Fabricius) — I. Discovering the phenomenon and its relation to the level of starvation. Bull. Math. Biol. **52**, 597-612.
30. Franks, N. R. & Sendova-Franks, A. B. (1992). Brood sorting by ants: distributing the workload over the work surface. Behav. Ecol. Sociobiol. **30**, 109-123.
31. Frehland, E., Kleutsch, B. & Markl, H. (1985). Modeling a two-dimensional alarm process. BioSystems **18**, 197.
32. Geake, E. (1994). How simple ants can sort out BT's complex nets. New Scientist **141**, 20-20.
33. Gordon, D. M. (1996). The organization of work in social insect colonies. Nature **380**, 121-124.
34. Goss, S. & Deneubourg, J.-L. (1988). Autocatalysis as a source of synchronized rythmical activity in social insects. Ins. Soc. **35**, 310-315.
35. Haken, H. (1983). Synergetics. Berlin: Springer Verlag.
36. Huberman, B., ed. (1988). The Ecology of Computation. Elsevier North-Holland.
37. Huberman, B. A., Lukose, R. M. & Hogg, T. (1997). An economics approach to hard computational problems. Science **275**, 51-54.
38. Kephart, J. O., Hogg, T. & Huberman, B. A. (1989). Dynamics of computational ecosystems. Phys. Rev. A **40**, 404-421.
39. Kuntz, P. & Layzell, P. (1995). A new stochastic approach to find clusters in vertex set of large graphs with applications partitioning in VLSI technology. Tech. Rep. LIASC, Ecole Nationale Supérieure des Télécommunications de Bretagne.
40. Lawler, E. L., Lenstra, J. K., Rinnooy-Kan, A. H. G. & Shmoys, D. B. (eds) (1985). The travelling salesman problem. Wiley.
41. Lindauer, M. (1952). Ein Beitrag zur Frage der Arbeitsteilung im Bienenstaat. Z. Vgl. Physiol. **34**, 299-345.
42. Lumer, E. & Faieta, B. (1994). Diversity and adaptation in populations of clustering ants. Proc. of Third Intl. Conf. on Simulation of Adaptive Behavior, pp. 499-508, MIT Press.
43. Nicolis, G. and Prigogine, I. (1977) Self-organization in non-equilibrium systems. Wiley.

44. Oster, G. & Wilson, E. O. (1978). Caste and ecology in the social insects. Princeton University Press, Princeton, NJ.
45. Page, R. E. & Mitchell, S. D. (1990). Self-organization and adaptation in insect societies. Phil. Sci. Assoc. (Fine, A., Forbes, M. & Wessels, L., eds) **2**, 289-298.
46. Page, R. E. & Robinson, G. E. (1991). The genetics of division of labour in honey bee colonies. Adv. Ins. Physiol. **23**, 117-169.
47. Robinson, G. E. (1987). Modulation of alarm pheromone perception in the honey bee: evidence for division of labour based on hormonally regulated response thresholds. J. Comp. Physiol. A **160**, 613-619.
48. Robinson, G. E. & Page, R. E. (1988). Genetic determination of guarding and undertaking in honey-bee colonies. Nature **333**, 356-358.
49. Robinson, G. E. (1992). Regulation of division of labor in insect societies. Annu. Rev. Entomol. **37**, 637-665.
50. Schoonderwoerd, R., Holland, O., Bruten, J. & Rothkrantz, L. (1997). Ant-based load balancing in telecommunications networks. Preprint.
51. Seife, C. (1997). Hedging bets on hard problems. Science **275**, 33-33.
52. Theraulaz G., Goss S., Gervet J. & Deneubourg J.-L. (1991). Task differentiation in Polistes wasp colonies: a model for self-organizing groups of robots. In: From Animals to Animats, Proc. of the 1st Intnl. Conf. on Simulation of Adaptive Behavior (Meyer, J.A. & Wilson, S.W., eds), 346-355, MIT Press.
53. Theraulaz, G. & Bonabeau, E. (1995a). Coordination in distributed building. Science **269**, 686-688.
54. Theraulaz, G. & Bonabeau, E. (1995b). Modeling the collective building of complex architectures in social insects with lattice swarms. J. theor. Biol. **177**, 381-400.
55. Waldspurger, C. A., Hogg, T., Huberman, B. A., & Kephart, J. O. (1992). Spawn: a distributed computational economy. IEEE Trans. Softw. Engineer. **18**, 103-117.
56. Wilson, E.O. (1985). The sociogenesis of insect colonies. Science **228**, 1489-1495.
57. Wilson, E. O. (1984). The relation between caste ratios and division of labour in the ant genus *Pheidole* (Hymenoptera: Formicidae). Behav. Ecol. Sociobiol. **16**, 89-98.