



Interaction patterns and task allocation in ant colonies

Deborah M. Gordon

Summary

Social insect colonies must accomplish many tasks, such as foraging, tending brood, constructing a nest, and so on. Task allocation is the process that adjusts the numbers of workers engaged in each task. This chapter discusses how information from other individuals is used in task decisions, and in particular, how workers use the pattern of interactions they experience, rather than the content of messages received. Empirical studies of harvester ants led to a mathematical model of task allocation in which environmental stimuli and interaction patterns both influence an individual's task. I outline the main results from this model, and describe recent empirical work that begins to examine how interaction patterns contribute to task allocation in harvester ants.

A colony's organization determines how the colony gets things done. Colonies must accomplish many tasks, such as foraging, tending brood, constructing a nest, and so on. Task allocation is the process that adjusts the numbers of workers engaged in each task. Such numbers change as external conditions and the needs of the colony vary. Colony organization operates, at a finer scale than that of task allocation, to determine how each task is accomplished. The organization of some tasks, such as recruitment to food, nest construction, and brood transport, has been well studied in some social insects. This chapter is about the broader scale of colony organization, that of task allocation.

One way to understand task allocation is to figure out how the behavior of individuals comes together to produce the behavior of colonies. The basic mystery of colony organization in social insects is that there is no one in charge. No individual is able to assess the global needs of the colony, or to count how many workers are engaged in each task and decide how many should be allocated differently. Instead, the capacities of individuals are fairly limited. Each worker

must make fairly simple decisions. There is abundant evidence, throughout physics, the social sciences, and biology, that such simple decisions by individuals can lead to predictable patterns in the behavior of groups [e.g. 1]. In the study of social insects, most of this evidence is concerned with the organization of particular tasks. For example, a tendency of individuals to move in the direction of certain chemicals can lead to the formation of foraging trails [2]. Evidence is growing that task allocation can be explained in a similar way, as the consequence of simple decisions by individuals.

Most early work on task allocation, especially in ants, was concerned with division of labor. Task allocation has two components: which individual performs which task, and how many individuals perform a task at any instant. Division of labor characterizes the first component, specifying which individuals perform a task. The notion of division of labor was introduced by the economist Adam Smith in the 18th century, to describe a workplace in which individuals are specialized, such that each one performs the same task over and over; the labor is “divided” among the individuals. The most extreme form, conceptually, is that each individual performs only its assigned task: that is, whenever worker x is active, it always performs task i . The second component of task allocation is what determines whether, at any instant, an individual is active and performs its task. Recent work suggests that complete specialization is rare in social insects: workers usually change tasks, both as they grow older and, in the short term, as conditions vary (reviewed in [3, 4]). Thus the study of task allocation has to resolve two questions: what determines an individual’s task, and what determines whether, at a particular instant, that individual is active.

What influences which task a worker performs and whether the worker is active? Clearly workers react to their immediate environments, because task allocation changes as conditions vary. When more food is available, more foragers appear; when the nest is damaged, more nest builders become active; when predators are nearby, workers remain inactive inside the nest. The conditions a worker encounters as it performs its task must affect the probabilities that it continues that task, becomes inactive, or takes up a different task.

But each worker’s task decisions are based on more than its own independent assessment of the environment. Workers also use signals from each other. Chemical signals are an obvious example. A worker exposed to alarm pheromone may stop performing its task and begin to circle around and perhaps to attack. The worker is responding to an interaction, in the form of transfer of an airborne chem-

ical signal, with another worker. In the same sense, people who rush out of a building when someone shouts “Fire!” are responding to an interaction with the shouter.

An individual worker’s task decisions thus depend on information from two sources: the environment, and other individuals. Information from other individuals may be packaged in chemical, mechanical, or visual messages, which can be understood by analogy to human language. For example, a chemical might say “Disturbance!” or “I belong to this colony” or “There is a queen here” or “I walked here”. However, specific chemical messages cannot underlie all task decisions. There would have to be a distinct chemical message to represent all the possible task transitions, such as “start doing task i ”, “stop doing task j ”; “change from task i to task j ”, for all tasks i and j , and a worker would have to be capable of using them all. This is asking social insect workers to use a very large vocabulary. But beyond the requirement for an impossibly large vocabulary, there is another improbable aspect of the one-message-per-task-decision scenario. How would workers know which command to give? For one ant to tell another, “stop doing task j ”, that ant would have to make some judgement about whether task j is needed, and this requires a global assessment of colony needs.

This chapter discusses how information from other individuals is used in task decisions. I focus on how workers use the pattern of interactions they experience, rather than the content of messages received. Various aspects of interaction pattern have been shown to be important: the rate of interactions, that is, the number per unit time; the interval between interactions; or the total number of interactions. The common feature of these results is that each interaction that an individual experiences contributes to a pattern, and the pattern itself influences the individual’s task decisions.

“Interaction” covers many different kinds of behavior. An interaction with a time delay occurs when one individual responds to a chemical signal deposited on a substrate by another individual, such as a trail pheromone or scent mark. Or interactions may be temporally constrained; for example, an interaction occurs immediately when one individual touches another with its antennae. Daily and short-term temporal patterns of interactions, often in the form of brief antennal contact, occur in many species of ants [5–7]. For most species, we do not know the function of such interaction patterns.

An interaction may not transmit any message besides the incidence of interaction itself. For example, ants may use the rate at which they contact others as an indication of density. In this case the message one ant gets when it touches

another is simply that it met another ant. This appears to occur in the ant *Lasius fuliginosus* [7]. In laboratory experiments, ants were kept in groups either of 35 or 75 ants. To each host group, we introduced 15 ants of another colony. The rate of contact with non-nestmates depended on the relative proportions of nestmates and non-nestmates. The 15 non-nestmates circulated freely, so each of the 35 ants would tend to encounter one of the 15 non-nestmates at a higher rate than would each of the 75. If the host ants' response depended on the total amount of alien pheromone introduced by the 15 non-nestmates, then both the 35-ant and 75-ant host groups should have had the same response. However, the 35-ant groups responded differently from the 75-ant groups. This suggests that ants were responding to the rate of interaction with non-nestmates, that is, to the proportion of non-nestmates present. The cue to non-nestmate density may be the incidence of interaction.

Another example of a response to interaction pattern is in the work of Seeley and colleagues on the organization of foraging, in particular the allocation of bees to nectar foraging [8–10]. This work is concerned, not with the question of which bee does which task, but the second component of task allocation, how a bee decides at a given instant whether to pursue its task actively. Whether a nectar forager pursues its task actively, leaving the nest to gather nectar, depends on the time that elapses between its arrival and its interaction with a nectar unloader. When there is less nectar to unload, more unloaders are available, and the rate at which foragers encounter unloaders increases. When there is more nectar to unload, unloaders are more likely to be busy elsewhere, unloading nectar, and the rate at which foragers encounter unloaders decreases. The time a forager has to wait for an unloader is directly, negatively proportional to this encounter rate. In this case, the unloader does not give the forager a message “Go out and forage some more” or “Remain inactive”. Instead, the pattern of interaction provides a cue that influences the nectar forager's task decision. The probability a nectar forager will be active at any instant is, of course, also influenced by cues related to the forager's environment, such as the quality of the food sources it has recently experienced, and their distances from the hive. It is also, of course, influenced by the unique physiology of that forager [11]. The link between interaction pattern and the forager's decisions allows the colony to regulate its nectar intake according to colony need, just as the link between nectar quality and the forager's decisions allows the colony to regulate its nectar intake according to food quality.

Interaction patterns and group size

There are vast numbers of social insect species that have never been studied, and those that we have studied differ greatly. There is no reason to expect that every species will use the same features of interaction patterns, or that cues based on interaction patterns will have exactly the same function in every species. Instead, it seems likely that many different features of interaction patterns are used, and in many ways. In general, though, as we learn more about colony organization and more about other species, should we expect to discover more individual decisions based on interaction patterns?

I think that individual decision rules based on interaction patterns will prove to be widespread in social insects. Interaction patterns are an obvious device for transmitting the effects of group size to individuals without requiring individuals to count. There is a basic spatial relation between group size and interaction pattern. If some objects are jostling around in an enclosed space, the number of times they will bump into each other, and the interval that elapses between encounters, will depend on the number or density of objects present. The pattern of encounters each individual experiences depends on the size of the group, so that each individual's experience tracks group size without any need for the individual to assess the overall size of the group.

Interaction patterns provide a way for individuals to track group size, and group or colony size is crucial to colony organization. A colony's size determines in many ways its needs and capacities. In many species, colonies grow larger as they grow older. Colony size influences how large a space the colony needs for a nest, how much food the colony requires, and how it competes with its neighbors. Colony size also affects the composition of task groups within colonies, and thus the resources available to perform each task; larger colonies have more foragers, and so on. Task group size, the numbers performing a task, determines how much of the task gets done. Of course, individuals may differ in the speed and quality of their work, but in general more workers on a task means more of that task is accomplished.

The main problem that colony organization must solve is to find ways that simple individuals, in the aggregate, can do what is necessary for colonies to live, grow, and reproduce. Because group size influences so strongly what colonies need, and what they can do, it seems likely that decision rules will appear that use simple cues which reflect group size. Interaction patterns can

translate group size, which is difficult for individuals to assess, into simpler cues.

Group size and interaction patterns in harvester ants

My research on interaction patterns began from a surprising finding involving colony size in harvester ants (*Pogonomyrmex barbatus*), a species of seed-eating ants I study in the desert of southeastern Arizona. Perturbation experiments showed that the numbers engaged in one task, such as foraging, depended on the numbers engaged in another task, such as nest maintenance work [12]. These relations among distinct worker groups depend on colony size [13].

The ants clearly used environmental cues in task decisions: when more food became available, more ants foraged. When I introduced foreign objects on the nest mound, more ants appeared to clear them away. This was expected; if colonies did not adjust numbers engaged in certain tasks to current conditions, they would be unable to track their environments or respond to disturbances.

The experiments revealed a less obvious feature of task allocation, the relation of different task groups. It was conceivable that task groups could be independent of each other; for example, foragers might decide whether to forage regardless of events affecting only nest maintenance workers. The experiments showed that this is not the case. Further work with marked individuals [14] showed one way that numbers in one task affect numbers in another: ants switch tasks. When more foragers are needed, ants previously engaged in other tasks switch to foraging. This switching increases the numbers foraging, but decreases the numbers available to perform the tasks the new foragers came from. However, task switching does not occur in all directions; ants that have switched from nest maintenance to another task do not switch back to nest maintenance. Instead, when more nest maintenance workers are needed, they are recruited from ants previously working only inside the nest.

Task switching did not account for all of the interactions between task groups. In some cases, a change in numbers performing one task causes the workers in a different task group to change their active-inactive status. For example, experiments that increase the numbers of nest maintenance workers recruited from inside the nest led to a decrease in the number of workers foraging outside the nest. This is not because foragers switch to nest maintenance; switching does not

occur in this direction. Instead, the foragers remain inactive inside the nest while other ants are recruited to perform nest maintenance. This may reflect a cycle in nest maintenance and foraging that seems to be common during the summer rains: after a storm, there is an increase in nest maintenance work to repair the damage done by flooding, and then follows a burst of foraging activity, as ants go out to collect newly exposed seeds, and nest maintenance subsides. It may be that perturbation experiments trigger the usual response to this cycle; if it's a day when intensive nest maintenance work is required, then foraging is low priority and the foragers tend to remain inside the nest.

These results show that task decisions by individuals, which task to perform and whether to be active, must be influenced by interactions among workers. Ants engaged in different tasks outside the nest are spatially segregated from each other (which is what made it possible to design experiments that directly affected only the workers engaged in one task). But workers of all tasks performed outside the nest mix as they come in and out of the nest, with many opportunities to interact.

The most surprising result of these experiments was an effect of colony age, which is related to colony size. Harvester ant colonies grow larger as they grow older, being founded by a single queen and reaching a stable size of about 10 000 workers by the time they are 5 years old and begin to reproduce [15, 16]. Colonies live for 15–20 years, until the queen dies [17]. Young, 2-year-old colonies, with 2–4000 workers, responded differently to perturbations than colonies 5 years or older with about 10 000 workers [13]. Older colonies were more consistent in response to repeated experiments, and their reactions to perturbations were more likely to emphasize foraging at the expense of other tasks. Since ants live at most a year [18], this is not due to the experience of older ants. The simplest explanation is an effect of colony size. This led me to consider what kind of decision rules individuals might use that could lead to different outcomes in a large colony and a small one.

Decision rules based on interaction patterns seem a good candidate to explain how colony size affects task allocation. Ants in a small colony might use the same rules as ants in a large one to decide which task to perform, and whether to perform it actively. But colony size could influence interaction patterns. In a small colony, ants might experience fewer interactions, or a lower interaction rate, or a longer interval between interactions, than ants in a large colony. Thus

in the aggregate, smaller colonies might tend to show different outcomes, in allocation of workers to various tasks, than larger ones.

Questions about interaction patterns

Investigating the relation of group size, task allocation, and interaction patterns raises a set of related questions. Some are empirical: What determines interaction rates in ants? How do such patterns depend on group size? Do species vary in interaction patterns? Do ants use interaction patterns in task decisions?

Some questions are theoretical: How do interaction patterns depend on the ways that ants move around in space? Can we predict how spatial behavior will affect interaction rate? What is the simplest possible decision rule based on interaction patterns that could still give the observed features of task allocation? This exercise in model making generates another set of empirical questions: How well does the behavior of real ants match that of the model?

Many of these questions have been addressed elsewhere, including species-specific encounter patterns [7], the effect of group size on encounter pattern [7, 19], the relation between spatial behavior and interaction pattern [19, 20], and models involving interaction patterns [21, 22], which are reviewed in [4, 23]. Here I will discuss some theoretical and empirical work that asks: Do ants use interaction patterns in making task decisions?

A model of task allocation

In one model of task allocation, an individual's decision is based on two kinds of stimuli: the rate of interaction with other workers, and the state of the environment relevant to a particular task. The combination of interaction rate and environmental stimuli is one feature of the model that makes it more realistic than previous ones that incorporate only interactions (e.g. [21]) or those that incorporate only environmental cues. There is overwhelming evidence that such decisions are affected by cues from the environment. If this were not the case, social insect colonies would not respond to changes in environmental conditions by reallocating workers to various tasks. Yet we know that this occurs frequently; worker allocation is adjusted in response to changes in food supply [8], predation

[24, 25], and nest condition [12, 25]. Thus the cues individuals use in decisions that generate task allocation must involve both interaction with others and environmental stimuli. Both factors, interactions and environmental cues, are components of other models of task allocation. For example, response to environmental stimuli, or to the extent to which other workers have successfully accomplished a task, is a component of the models of Jeanne [25, 26]. Numbers of interactions, rate of interaction or time elapsed since the last interaction are a component of several recent models of the organization of certain tasks (e.g. [2, 10]).

How the model works

The model of Pacala et al. [22] predicts the numbers allocated to each of Q tasks. The vector representing the numbers engaged in each of Q tasks is $\underline{X} = \{ X_1, X_2, \dots, X_Q \}$. An individual's state is defined in the following way: either it is actively engaged in one of the Q tasks, or else it is inactive. If it is actively engaged in one of the tasks, it is either successful or unsuccessful. In the model, an individual can detect the task of another individual that it meets; empirical evidence concerning this capacity in harvester ants is discussed below.

Success is represented as s_i , the probability that an individual engaged in task i is in the successful state. This depends on current environmental conditions; for example, few foragers will be successful when there is little food available. In the model, an individual switches from task i to task j when she is unsuccessful at task i and encounters another ant that is successful at task j . It is well known that social insect workers can detect the success of other workers at some tasks, such as foraging. For example, harvester ant foragers are more likely to leave the nest to forage when the rate of successful returning foragers is high [17]. In the model, unsuccessful individuals engaged in task i will quit task i , and become inactive, at a rate q_i . This quitting rate depends in part on the environment, which determines how likely an individual is to be successful or unsuccessful. The model assumes a unique rate q_i for each task. Differences among tasks in the quitting rate might depend on the cost of performing a task. A high-cost task might not be worthwhile to pursue when workers engaged in it cannot succeed. For example, harvester ants obtain water by metabolizing the lipids in the seeds they collect. Foragers are using water when they travel in the hot desert sun, and if they are unsuccessful at collecting seeds, they may be losing more water than

they gain. At this point, they might quit foraging and remain inactive inside the nest until more food becomes available. However, ants working inside the nest to enlarge a chamber may not be using as much energy. Suppose they hit a rocky patch of soil in which few individuals can succeed in scraping away at the wall of the chamber. These unsuccessful nest workers may continue longer before they quit than would unsuccessful foragers. A high-cost task, less worthwhile to pursue when workers engaged in it cannot succeed, might have a higher q .

The rate at which an individual engaged in task i encounters individuals engaged in task j is proportional to D_{ij} , the average local density (numbers per unit area) of task- j individuals in the vicinity of a task- i individual. D_{ij} is averaged over all individuals engaged in task i . Local densities depend on the movement patterns of individuals. These may vary from task to task, and the resulting interaction rates among workers engaged in different tasks could vary accordingly in ways that might be complicated to predict [20]. To imagine how this might work, one can think of interactions as simple collisions (though in fact there could be a time delay so that the two participants in the interaction are in the same place at different times). One extreme in the process that determines collision rate would be when all individuals mix completely, and D_{ij} is proportional to X_j/N . That is, the higher the proportion of task- j individuals, the more likely any task- i individual is to meet one. The other extreme would occur when individuals engaged in different tasks are spatially segregated. Then D_{ij} would be proportional to X_j . A task- i individual that happened to fall in among the task- j individuals would tend to experience more contacts, the more task- j individuals there were. In fact, the processes that determine interaction rates probably fall between these two extremes. Such processes may differ among tasks, since tasks differ in use of space.

The basic model is

$$\frac{dX_i}{dt} = -X_i (1 - s_i) q + [\sum X_i s_i D_{ij} (1 - s_j) - \sum X_j s_j D_{ji} (1 - s_i) + (N - \sum X_j) D_{ii} s_i]$$

where all sums are over all Q tasks, and D_{ii} is the average local density of task- i individuals around an inactive individual.

The change in numbers engaged in task i (the left-hand side) depends on four terms: (i) individuals that are unsuccessful at task i and quit, thereby becoming inactive; (ii) individuals from other tasks that encounter successful task- i ants and switch tasks to perform task i ; (iii) individuals that are unsuccessful at task i ,

encounter successful individuals from another task, and switch from task i to the other task; and (iv) inactive individuals that encounter successful task- i individuals, and switch to perform task i .

Investigating the dynamics of this model led to the following results.

Optimal task allocation

First, if success rates are related to fitness, then the model leads to some optimality predictions. The predicted optimum distribution of workers among tasks will depend on whether we add up the fitnesses of each individual in the colony, such as the net amount of resource brought into the colony by each individual, or instead use some measure of the fitness of the whole colony, such as the total amount of resource brought into the colony given the colony's costs in obtaining these resources. If the former, then the optimal outcome of the model is the ideal free distribution: no individual should perform a task if it could have higher success at some other task. In the latter case, when colony fitness is maximized, then all tasks performed actively at any moment must yield equal marginal benefits, and the benefit of performing each task must at least equal the cost of performing it.

Stability of task allocation

The second result is that the dynamics of the model lead to equilibria that resemble the optimal outcomes described above. The conditions when the allocation of workers into various tasks is stable depend on the way in which the performance of each task is assumed to modify the environment, and thus on the ensuing success rate of individuals engaged in that task. Three possible relations of task performance and success rates were considered:

- (1) As *per capita* returns from task i increase, so does s_i , the proportion of task- i workers that are successful; s_i is a monotonically increasing function of the *per capita* returns from task i . In addition, the *per capita* returns are independent of X_i . An example of this is when foragers are in an environment with very abundant food: each individual returning with food is a successful for-

ager, and no matter how many individuals forage, the amount of food available does not decrease.

- (2) As *per capita* returns from task i increase, so does s_i , the proportion of task- i workers that are successful; as in case (1), s_i is a monotonically increasing function of the *per capita* returns from task i . But the *per capita* returns are a decreasing function of X_i . For example, suppose food is scarce. Initially the more foragers are active, the more food will be retrieved, and the more foragers will be successful—but eventually the foragers will deplete the food, and the returns in food will go down as the numbers foraging go up.
- (3) Individuals somehow interfere with each other, with the consequence that successful individuals that meet unsuccessful ones sometimes change state to become unsuccessful themselves. In this case, as in the previous one, *per capita* returns are a decreasing function of X_i . But unlike the previous two cases, s_i is a monotonically increasing function of the marginal rate of return to the group as a whole, not the *per capita* rate of return.

For (1) and (2) the stable equilibrium resembles the predictions for the optimal distribution of workers into tasks to maximize individual fitness, that is, individual returns. For (1) or (3) the stable equilibrium also resembles the predictions for the optimal distribution to maximize the fitness of the colony as a whole.

The equilibria are unstable when the success rate for any task not currently performed is greater than the success rate common to all tasks currently performed. Thus if the environment changes, causing the success rate of particular tasks to change, then the distribution of workers in various tasks will be altered so that all equally successful tasks are being performed.

Colony size and ability to track a changing environment

The third result is a prediction for how quickly task allocation can track a changing environment. Tracking speed depends on colony size. Everything else being equal, large colonies can track a changing environment more rapidly than small ones. This is because the larger the colony, the higher the rate of interaction. High interaction rate means a high rate of dissemination of information about the environment, as successful individuals interact with unsuccessful ones.

However, there is a disadvantage to large colony size. In a large colony, a successful individual might interact with large numbers of unsuccessful or inactive individuals that would then be recruited to the task of the successful one. As a consequence there might be more workers performing the task than the environment warrants. For example, suppose a successful forager comes back to a large colony and meets large numbers of unsuccessful or inactive workers inside the nest. This would lead many workers to switch to forage, possibly more than would be warranted by the amount of food available. In general, workers in large colonies might be more likely to be active in unprofitable environments.

If real colonies operate the way the model does, then one way they could overcome this disadvantage of large colony size would be to curtail interaction rates. The ant *Lasius fuliginosus* seems to do this [7]. We manipulated density experimentally. When density, number of ants per unit area, is high, ants curtail the rate at which they engage in antennal contact with others. The ants can avoid interaction before it occurs. We found that in this species one ant can perceive another at a distance of about a body length [7], and thus an ant can move out of the range of contact with an approaching ant.

Interaction patterns in task allocation

Recent empirical work shows that interaction patterns affect the task decisions of harvester ants [28]. Observers followed marked ants for about 20 min, and recorded the task of the focal ant and those of all ants it met. An encounter was defined as a brief contact by the antennae of the focal ant and any part of the body of another ant; the great majority of such contacts were between the antennae of the two ants. Whether an ant switched tasks, and whether it remained inactive or pursued a task actively, depended on its recent encounter history. There was a statistically significant association between the number and type of contacts a focal ant experienced and its subsequent tasks.

This study focussed on midden work, the sorting and piling of the colony refuse pile, or midden. Ants not engaged in midden work, that encountered high numbers of midden workers, were likely to begin midden work. Ants that were walking or standing around, apparently not engaged in any task, were likely to remain so if they met low numbers of midden workers.

In this laboratory study, interaction patterns provide positive feedback to perform midden work. In the field, midden material is removed by other insects. The absence of midden material to work with probably provides negative feedback. In general, however, there is no reason to suppose that interaction patterns generally provide positive or negative feedback; either one is equally probable a priori. There is, however, abundant reason to believe that interaction patterns alone do not determine task decisions. First, if an ant's decision whether to perform midden work depended only on its interaction with midden workers, our results suggest that once a few ants begin midden work, all ants will end up joining in. Second, if ants were to rely solely on interaction patterns, they would be unable to respond to any environmentally induced changes in their colony's need for midden work. Moreover we never observe all ants in a colony engaged in midden work. This means that some process of negative feedback, probably linked to the quantity of midden material present, influences worker decisions about midden work.

Harvester ants of different task groups differ in their cuticular hydrocarbon profiles [27]. Ants may use the cuticular hydrocarbon odour to distinguish the task of the ants they meet.

What feature of interaction pattern do the ants use?

We found a statistical effect of number of encounters on task decisions. We do not know which feature of an ant's interaction history, correlated with number, influences its behavior. We counted number of encounters in a fairly constant period of time. Thus numbers of encounters are positively correlated with rate of encounter, and negatively correlated with the interval between encounters. The effect of number may occur because ants respond to the rate, in numbers per unit time, or to the interval between encounters. We tested whether there was any effect of the sequence of encounters; did the probability an ant performs task i depend on whether its last encounter was with an ant of task i , its second-to-the-last encounter, and so on? We found no such effect.

The question of which aspect of its recent interaction history influences an ant's task decisions is an empirical one, but it is also an interesting theoretical question. What feature of the pattern should an ant use? There must be some time interval that is relevant. An ant's decision about its task this minute almost cer-

tainly does not depend on the whole pattern of every interaction it has experienced since the moment it eclosed. Our observations show that ants change task state several times, at least from active to inactive or back, in 20 min. This suggests the relevant interval is on the scale of tens of minutes or less. Let us call this the ant's forgetting time. During this forgetting time, the ant might count interactions. It would then respond to the total number of interactions it has counted, as far back as it can remember. Another alternative is that the ant might calculate an interaction rate. That is, it would count number of interactions, for some interval less than or equal to its forgetting time, and divide this by the amount of time elapsed. A third alternative is that the ant tracks the interval between successive encounters, measuring somehow the time elapsed since the last time it met another ant of task i .

Asking which of these alternatives is more likely branches into two questions. The first is about the capacities of ants. Which alternative requires the least cognitive power? Measuring intervals seems to require the least calculation, and there is much work to suggest that many animals, including honey bees and ants, can assess time intervals (Fourcassié et al., this volume). However, the result that the last encounter rarely determines task decisions indicates that more than one encounter, and thus more than one interval, is involved. Tracking many intervals, or counting many encounters, or calculating a rate involving many encounters, may be equally difficult operations. At least, it is not obvious that one is much easier than the rest.

A second way to ask which of these alternatives is more likely is to ask what would be most effective for the colony. One consideration is that the larger the number of interactions involved in the decision, the less likely are decisions to be affected by sampling error. An ant moves around and encounters other ants. In the short term, who each ant meets is greatly affected by individual variation in movement. A larger sample may more accurately reflect the actual proportion of ants engaged in each task; if the sample were large enough that the ant met every other ant in the colony, its interaction record would be equivalent to the distribution of ants into each task.

Future work

Task allocation, how colonies change the allocation of workers in response to changing conditions and changing colony needs, has been studied in relatively few species of social insects. The basic empirical question for any species is how different tasks are related. How do numbers engaged in one task depend on numbers engaged in another? The next question is how individual decisions create shifts in task allocation by the colony. Under what conditions do individuals switch tasks, and which transitions are most likely? Under what conditions do individuals decide to become active, or to remain inactive? Once we know something about how task allocation operates in a colony, it is possible to construct models or hypotheses. The problem is then to determine what rules individuals use that, in the aggregate, produce colony behavior. Such rules will probably involve information from the environment and from interaction with other individuals. The ways that colony organization produces task allocation may be as various as the tens of thousands of social insect species. Or there may be common features, such as perhaps the use of interaction patterns, across broad taxonomic groups. Only empirical work on a variety of species will make it possible to discover the general principles of task allocation in social insects, and to understand the relation between task allocation, interaction patterns, and colony size.

References

- 1 Langton CG (1989) Artificial Life. In: CG Langton (ed): *Artificial Life*. SFI Studies in the Sciences of Complexity, vol 6. Addison-Wesley Publishing Company, Redwood City CA
- 2 Deneubourg JL, Goss S (1989) Collective patterns and decision-making *Ecol Ethol Evol* 1: 295–311
- 3 Gordon DM (1989) Caste and change in social insects. In: P Harvey, L Partridge (eds): *Oxford surveys in evolutionary biology*. Oxford University Press, Oxford
- 4 Gordon DM (1996) The organization of work in social insect colonies. *Nature* 380: 121–124
- 5 Franks NR, 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
- 6 Cole BJ (1991) Short-term activity cycles in ants: generation of periodicity by worker interaction. *Amer Naturalist* 137: 244–259
- 7 Gordon DM, Paul REH, Thorpe K (1993) What is the function of encounter patterns in ant colonies? *Anim Behav* 45: 1083–1100
- 8 Seeley TD (1989) Social foraging in honey bees: how nectar foragers assess their

- colony's nutritional status. *Behav Ecol Sociobiol* 24: 181–199
- 9 Seeley TD, Camazine S, Sneyd J (1991) Collective decision-making in honey bees: how colonies choose among nectar sources. *Behav Ecol Sociobiol* 28: 277–290
 - 10 Seeley TD, Tovey CA (1994) Why search time to find a food-storer bee accurately indicates the relative rates of nectar collecting and nectar processing in honey bee colonies. *Anim Behav* 47: 311–316
 - 11 Robinson GE, Page REJr, (1989b) Genetic basis for division of labor in an insect society. In: *The Genetics of Social Evolution*. Westview Press, Boulder, CO
 - 12 Gordon DM (1986) The dynamics of the daily round of the harvester ant colony. *Anim Behav* 34: 1402–1419
 - 13 Gordon DM (1987) Group-level dynamics in harvester ants: young colonies and the role of patrolling. *Anim Behav* 35: 833–843
 - 14 Gordon DM (1989) Dynamics of task switching in harvester ants. *Anim Behav* 38: 194–204
 - 15 Gordon DM (1992) How colony growth affects forager intrusion in neighboring harvester ant colonies. *Behav Ecol Sociobiol* 31: 417–427
 - 16 Gordon DM (1995) The expandable network of ant exploration. *Anim Behav* 50: 995–1007
 - 17 Gordon DM (1991) Behavioral flexibility and the foraging ecology of seed-eating ants. *Amer Naturalist* 138: 379–411
 - 18 Gordon DM, Hölldobler B (1987) Worker longevity in harvester ants. *Psyche* 94: 341–46
 - 19 Gordon DM (1995) The development of an ant colony's foraging range. *Anim Behav* 49: 649–659
 - 20 Adler FR, Gordon DM (1992) Information collection and spread by networks of patrolling ants. *Amer Naturalist* 40: 373–400
 - 21 Gordon DM, Goodwin B, Trainor LEH (1992) A parallel distributed model of ant colony behaviour. *J Theor Biol* 156: 293–307
 - 22 Pacala SW, Gordon DM, Godfray HCJ (1996) Effects of social group size on information transfer and task allocation. *Evol Ecol* 10: 127–165
 - 23 Gordon DM (1994) How social insect colonies respond to variable environments. In: LA Real (ed): *Behavioral mechanisms in evolutionary biology*. University of Chicago Press, Chicago
 - 24 Nonacs P, Dill LM (1990) Mortality risk versus food quality trade-offs in a common currency: ant patch preferences. *Ecology* 71: 1886–1892
 - 25 Jeanne RL (1986) The organization of work in *Polybia occidentalis*: the costs and benefits of specialization in a social wasp. *Behav Ecol Sociobiol* 19: 333–341
 - 26 Jeanne RL (1996) Regulation of nest construction behaviour in *Polybia occidentalis*. *Anim Behav* 52: 473–488
 - 27 Wagner D, Brown MJF, Broun P, Cuevas W, Moses LE, Chao DL, Gordon DM (1998) Task-related differences in the cuticular hydrocarbon composition of harvester ants, *Pogonomyrmex barbatus*. *J Chem Ecol* 24: 2021–2037
 - 28 Gordon DM, Mendiabadi N (1999) Encounter rate and task allocation in harvester ants. *Behav Ecol Sociobiol* 45: 370–377