

Cooperative transport by ants and robots

C. Ronald Kube^{a,1} and Eric Bonabeau^{b,2}

^a*Edmonton Research Centre, Syncrude Canada Limited 9421-17 Avenue,
Edmonton, Alberta, Canada T6N 1H4*

^b*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA*

Abstract

In several species of ants, workers cooperate to retrieve large prey. Usually, one ant finds a prey item, tries to move it, and, when unsuccessful for some time, recruits nestmates through direct contact or chemical marking. When a group of ants tries to move large prey, the ants change position and alignment until the prey can be moved toward the nest. A robotic implementation of this phenomenon is described. Although the robotic system may not appear to be very efficient, it is an interesting example of decentralized problem-solving by a group of robots, and it provides the first formalized *model* of cooperative transport in ants.

1 Introduction

Social insect societies –ants, bees, termites and wasps– are distributed systems in which colony-level behavior emerges out of interactions among individual insects [4]. In addition to being a decentralized system, a colony of insects exhibits *flexibility* and *robustness*, two features that would be desirable in an artificial system. A recent trend in both *Artificial Intelligence* and *Operations Research* consists of viewing the social insect metaphor as a new paradigm for designing artificial problem-solving devices and optimization algorithms [3]. In *Autonomous Robotics*, *swarm-based robotics* relies on the same metaphor for the design of distributed control algorithms for swarms of robots.

There has been an upsurge of interest in swarm-based robotics in recent years [9] as it provides an interesting alternative to more classical approaches in robotics. Some tasks may be inherently too complex or impossible for a single robot to perform. For example, in the case study described in this paper,

¹ email: kube@cs.ualberta.ca

² email: bonabeau@santafe.edu

pushing a box requires the “coordinated” efforts of at least two individuals. Speed up can result from using several robots. Designing, building, and using several simple robots may be easier than designing, building and using a single complex robot. It may also be cheaper because of the robots’ simplicity. A swarm of simple robots may also be more flexible without the need to reprogram the robots, and more reliable and fault-tolerant because one or several robots may fail without affecting task completion. Furthermore, theories of self-organization teach us that randomness or fluctuations in individual behavior, far from being harmful, may in fact greatly enhance the system’s ability to explore new behaviors and find new solutions. In addition, self-organization and decentralization, together with the idea that interactions among agents need not be direct but can rather take place through the environment, point to the possibility of significantly reducing communications between robots: explicit robot-to-robot communications rapidly become a big issue when the number of robots increases; this issue can be to a large extent eliminated by suppressing such communications! Also, central control is usually not well suited to dealing with a large number of agents, not only because of the need for robot-to-controller-and-back communications, but also because failure of the controller implies failure of the whole system.

The current success of collective robotics is the result of several factors:

- (1) The relative failure of the *Artificial Intelligence* program, which classical robotics relied upon, has forced many computer scientists and roboticists to reconsider their fundamental paradigm. This paradigm shift has led to the advent of connectionism, and to the view that sensori-motor intelligence is as important as reasoning and other higher-level components of cognition. Swarm-based robotics relies on the anti-classical-AI idea that a group of robots may be able to perform tasks without explicit representations of the environment and of the other robots and that planning may be replaced by reactivity.
- (2) The remarkable progress of mobile robotics during the last decade has allowed many researchers to experiment with mobile robots, which have not only become more efficient and capable of performing many different tasks, but also cheap(er).
- (3) The field of *Artificial Life*, where the concept of emergent behavior is emphasized as being essential to the understanding fundamental properties of the living, has done much to propagate ideas about collective behavior in biological systems, particularly social insects; facts and theories that were unknown to roboticists reached them.

Using a swarm of robots has some drawbacks. For example, stagnation is one: because of the lack of a global knowledge, a group of robots may find itself in a deadlock, where it cannot make any progress. Another problem is

to determine how these so-called “simple” robots should be programmed to perform user-designed tasks. The pathways to solutions are usually not pre-defined but emergent, and solving a problem amounts to finding a trajectory for the system and its environment so that the states of both the system and the environment constitute the solution to the problem: although appealing, this formulation does not lend itself to easy programming. Until now, we implicitly assumed that all robots were identical units: the situation becomes more complicated when the robots have different characteristics, respond to different stimuli, or respond differently to the same stimuli, and so forth; if the body of theory that roboticists can use for homogeneous groups of robots is limited, there is virtually no theoretical guideline for the emergent design and control of heterogeneous swarms.

Many potential applications of swarm-based robotics require miniaturization. Very small robots, micro- and nano-robots, which will by construction have severely limited sensing and computation, may need to *operate in very large groups or swarms to affect the macroworld* [34]. Approaches directly inspired or derived from swarm intelligence may be the only way to control and manage such groups of small robots. As the reader will perhaps be disappointed by the simplicity of the tasks performed by state-of-the-art swarm-based robotic systems such as the one presented in this paper, let us remind her that it is in the perspective of miniaturization that swarm-based robotics becomes meaningful. In view of the great many potential applications of swarm-based robotics, it seems urgent to work at the fundamental level of what algorithms should be put into these robots: understanding the nature of coordination in groups of simple agents is a first step toward implementing useful multirobot systems.

In swarm-based robotics, cooperative transport—particularly cooperative box-pushing—has been an important benchmark for testing new types of robotic architecture. One of the swarm-based robotic implementations of cooperative transport is so closely inspired by cooperative prey retrieval in social insects that it is a genuine *model* of the phenomenon, thereby providing a unique example of a truly bidirectional exchange between biology and robotics. Ants of various species are capable of collectively retrieving large prey that are impossible for a single ant to retrieve. Usually, a single ant finds a prey item and tries to move it alone; when successful, the ant moves the item back to the nest. When unsuccessful, the ant recruits nestmates through direct contact or trail laying. If a group of ants is still unable to move the prey item for a certain time, specialized workers with large mandibles may be recruited in some species to cut the prey into smaller pieces. Although this scenario seems to be fairly well understood in the species where it has been studied, the mechanisms underlying cooperative transport—that is, when and how a group of ants move a large prey item to the nest—remain unclear. No formal description of the biological phenomenon has been developed, and, surpris-

ingly, roboticists went further than biologists in trying to model cooperative transport: perhaps the only convincing model so far is one that has been introduced and studied by roboticists [31] and, although this model was not aimed at describing the behavior of real ants, it *is* biologically plausible. This paper first describes empirical observations of cooperative transport in ants, and then describes the work of Kube and Zhang [27–29,31].

2 Cooperative Prey Retrieval in Ants³

A small prey or food item is easily carried by a single ant. But how can ants “cooperate” to carry a large item? Cooperative prey (or large food item) retrieval and transport has been reported in several species of ants [46,42,50]: weaver ants *Oecophylla smaragdina* [25] and *Oecophylla longinoda* [23,54], army ants *Eciton burchelli* [18] African driver ants *Dorylus* [20,36], and other species such as *Pheidole crassinoda* [45], *Myrmica rubra* [47], *Formica lugubris* [47], *Lasius neoniger* [49], the desert ants *Aphaenogaster* (ex-*Novomessor*) *cockerelli* and *Aphaenogaster albisetosus* [24,32], *Pheidologeton diversus* [36], *Pheidole pallidula* [13,14], *Formica polyctena* [10,11,35,51], *Formica schau-fussi* [42,41,50] and the ponerine ants *Ectatomma ruidum* [39] and possibly *Paraponera clavata* [6]. This cooperative behavior can be quite impressive. For example, Moffett [36] reports that a group of about 100 ants *Pheidologeton diversus* was able to transport a 10-cm earthworm weighing 1.92 g (more than 5000 times as much as a single 0.3-mg to 0.4-mg minor worker) at 0.41 cm/s on level ground. By comparison, ants engaged in solitary transport of food items on the same trail were carrying burdens weighing at most 5 times their body weight at about 1 cm/s: this means that ants engaged in the cooperative transport of the earthworm were holding at least 10 times more weight than did solitary transporters, with only a modest loss in velocity [36].

We believe that the phenomenon of cooperative transport is much more common in ants than these few studies suggest: to the best of our knowledge, these studies are the only ones that report detailed observations of cooperative prey transport. This phenomenon involves several different aspects:

- (1) Is there an advantage to group transport as opposed to solitary transport? Is worker behavior in group transport different than in solitary transport?
- (2) When and how does an ant know that it cannot carry an item alone because it is either too large or too heavy?
- (3) How are nestmates recruited when help is needed?
- (4) How do several ants cooperate and coordinate their actions to actually transport the item?

³ Portions of section 2 have been modified from [3].

- (5) How do ants ensure that there is the right number of individuals involved in carrying the item?
- (6) How does a group of transporting ants handle deadlocks and, more generally, situations where the item to be transported is stuck, either because of antagonistic forces or because of the presence of an obstacle or heterogeneities in the substrate?

All these questions, that have been more or less satisfactorily dealt with in the above-mentioned studies, are of enormous interest in view of implementing a decentralized cooperative robotic system to transport objects the locations and sizes of which are unknown.

2.1 Solitary Transport Versus Group Transport

In *Pheidologeton diversus*, single worker ants usually carry burdens (grasping them between their mandibles, lifting them from the ground and holding them ahead as they walk forward) rather than drag them [36]. By contrast, in cooperative transport, one or both forelegs are placed on the burden to aid in lifting it, mandibles are open and usually lay against the burden without grasping it. The movement patterns of group-transporting ants corresponding to their positions around the perimeter of a burden with reference to the direction of transport are also different than those of ants engaged in solitary transport: workers at the forward margin walk backward, pulling the burden, while those along the trailing margin walk forward, apparently pushing the burden; ants along the sides of the burden shuffle their legs sideways and slant their bodies in the direction of transport [36].

By contrast, Sudd [45,47] observes that individual *Pheidole crassinoda*, *Myrmica rubra*, and *Myrmica lugubris* ants appear to exhibit the same behavioral patterns in solitary and group transport: in group transport, all three species used the same method as when they work alone, including realignment and repositioning. This, however, does not exclude cooperative behavior: group transport in these species is particularly interesting because the same individual behavior is functional either in isolation or in group, and may even lead to increasing returns (up to a maximum group size: see section 2.2) despite the lack of direct response of individuals to the presence of their nestmates.

In general, whether ants behave similarly or differently when engaged in solitary and group transport, group transport is more efficient than solitary transport for large prey. Ants can dismantle a large food item into small enough pieces to be carried by individual ant workers. Moffett [36] observed that a large piece of cereal, which would have required 498 solitary *Pheidologeton diversus* transporters if broken down into small enough pieces, could be trans-

ported collectively by only 14 ants. More generally, he observed that the weight carried by ant increases with group size: the total weight carried by a group of N workers increases as $W \propto N^{2.044}$, which means that the weight carried by each ant increases on average as $N^{1.044}$. Franks [18] made similar observations on *Eciton burchelli*: let W_i be the dry weight of transported items and W_a the total dry weight of the group of transporting ants, the relationship between both is $W_i \propto W_a^{1.377}$, which, assuming that W_a is proportional to N , implies that the dry weight carried by ant increases as $N^{0.377}$. Franks [18] also observed that items were always retrieved at a standard speed, relatively independent of group size: he hypothesized that the increased efficiency of group transport with group size results from the group's ability to overcome the rotational forces necessary to balance a food item. Along the same lines, we already mentioned Moffett's [36] experiment in which he showed that group-transporting ants could carry more than 10 times more weight than did solitary transporters at a speed only divided by 2. He found that the velocity of transport decreases as a function of the number of *Pheidologeton diversus* carriers, but decreases significantly only for large group sizes (>12 carriers). The transport efficiency per ant, measured by the product of burden weight by transport velocity divided by the number of carriers, increases with group size up to a maximum for groups of 8 to 10 ants, and then declines [36].

As emphasized by Traniello and Robson [50], transport efficiency may not be the only and primary reason for group transport in ants. In *Aphaenogaster cockerelli* group retrieval of prey decreases interference competition with sympatric species [24,32], and in *Lasius neoniger* the rapid formation of a cooperative retrieval group is crucial to foraging success [49].

2.2 From Solitary to Group Transport

All reports of how the decision is made to switch from solitary to group transport describe variants of the same phenomenon. A single ant first tries to carry the item, and then, if the item resists motion, to drag it (although dragging is rare in *Pheidologeton diversus*). Resistance to transport seems to determine whether the item should be carried or dragged [13,14,45,47]. The ant spends a few seconds testing the resistance of the item to dragging before realigning the orientation of its body without releasing the item: modifying the direction of the applied force may be sufficient to actually move the item. In case realignment is not sufficient, the ant releases the item and finds another position to grasp the item. If several repositioning attempts are unsuccessful, the ant eventually recruits nestmates. Recruitment per se is examined in the next section. Sudd [47] reports that the time spent attempting to move the item decreases with the item's weight: for example, an ant may spend up to 4 minutes for items less than 100 mg, but only up to 1 minute for items more than 300

mg. Detrain and Deneubourg [13,14] have shown that in *Pheidole pallidula*, it is indeed resistance to traction, and not directly prey size, that triggers recruitment of nestmates, including majors, to cut the prey: they studied recruitment through individual trail laying for prey of different sizes (fruit flies versus cockroaches), or of the same size but with different levels of retrievability (free fruit flies versus fruit flies covered by a net). A slow recruitment to free fruit flies was observed, in connection to weak individual trail laying; in contrast, strong recruitment and intense individual trail laying were observed when large prey or small but irretrievable prey were offered. It is therefore the ability or inability of an individual or a group that governs recruitment.

2.3 Recruitment of Nestmates

Hölldobler et al. [23] studied recruitment in the context of cooperative prey retrieval in two *Aphaenogaster* (ex-*Novomessor*) species: *Aphaenogaster albisetosus* and *Aphaenogaster cockerelli*. They show that recruitment for collective transport falls within two categories: short-range recruitment (SRR) and long-range recruitment (LRR). In SRR, a scout releases a poison gland secretion in the air immediately after discovering a large prey item; nestmates already in the vicinity are attracted from up to 2 m. If SRR does not attract enough nestmates, a scout lays a chemical trail with a poison gland secretion from the prey to the nest: nestmates are stimulated by the pheromone alone (no direct stimulation necessary) to leave the nest and follow the trail toward the prey.

Hölldobler [25] reports short-range, and more rarely long-range (rectal gland-based), recruitment in *Oecophylla smaragdina* in the context of prey retrieval, during which secretions from the terminal sternal gland and alarm pheromones from the mandibular glands interact. This short-term recruitment attracts nestmates located in the vicinity, which quickly converge toward the intruder or prey item, which is retrieved into the nest when dead. In a series of experiments with 20 freshly killed cockroaches placed at randomly selected locations in a colony's territory, the prey were discovered within several minutes (average: 8.05 min.); ants in the vicinity were attracted by short-range recruitment signals; 5 to 8 ants grasped the prey item and held it on the spot for several minutes (average: 11.6 min.) before jointly retrieving it to the nest. This last phase involved 5.3 ants on average. In *Oecophylla longinoda*, even when the prey were pinned to the ground and the ants were unable to retrieve it, long-range recruitment was not used [23]. By contrast, long-range recruitment was observed in *Oecophylla smaragdina* when the cockroaches were pinned to the substrate and several workers had attempted without success to remove the prey: recruiting ants moved back to the nearest leaf nest (although there was only one queen, as is usual in this species, the nest of the considered colony

was composed of 19 separate leaf nests, which is also common in the species) where they recruit nestmates which soon moved out of the leaf nest toward the prey. From 25 to 59 could be recruited, whereas between 9 and 19 ants were involved in actually retrieving the prey to the nest once the prey were eventually retrieved. This indicates that the ants do not estimate the size or weight of the prey but rather adapt their group sizes to the difficulty encountered in first moving the prey. Hölldobler [25] reports that the recruited ants were gathering around the prey, seeking to get access, and sometimes grasped nestmates that were already working at the prey, thereby forming a pulling chain, a common behavior in weaver ants. The prey were usually first transported to the leaf nest from which helpers had been recruited.

2.4 *Coordination in Collective Transport*

Coordination in collective transport seems to occur through the item being transported: a movement of one ant engaged in group transport is likely to modify the stimuli perceived by the other group members, possibly producing, in turn, orientational or positional changes in these ants. This is an example of stigmergy [22], the coordination of activities through indirect interactions. Here, stigmergy is a promising step toward a robotic implementation, because it suggests that a group of robots can cooperate in group transport *without direct communication* among robots; moreover, robots do not have to change their behaviors depending on whether or not other robots are engaged in the task of carrying (or dragging, or pulling, or pushing) the item. The coordination mechanism used by ants in cooperative transport is not well understood, and has never really been modeled. The swarm of robots described in section 3 is just such a model, which shows that the biology of social insects and swarm-based robotics can both benefit from each other.

2.5 *Number of Ants Engaged in Group Transport*

Apparently, the number of ants engaged in transporting an item is an increasing function of the item's weight, which indicates that group size is adapted to the item's characteristics. For example, Moffet [36] reports how the number of *Pheidologeton diversus* carriers varies with burden weight. Inverting the relationship described in section 2.1, we obtain $N \propto W^{0.489}$. The fit to the data is remarkable, suggesting that the adaptation of group size is accurate. Using the same notations as in section 2.1, Franks [18] finds that $W_a \propto W_i^{0.726}$ for *Eciton burchelli*. However, as mentioned in the previous section, Hölldobler's [25] observations suggest that the ants adapt group size to the difficulty encountered in first moving prey: decisions rely on how difficult it is to carry the prey, and

not simply on weight. A prey item that resists (either actively or passively) stimulates the ant(s) to recruit other ants. Success in carrying a prey item in one direction is followed by another attempt in the same direction. Finally, recruitment ceases as soon as a group of ants can carry the prey in a well-defined direction: in that way, group size is adapted to prey size.

In addition to the size of the cooperative transport group, it seems that the composition of the group is not random: for example, in army ants (*Eciton burchelli*), groups have a specific distribution of submajors that comprise a specialized transport caste [18]. Of course the situation is less clear in monomorphic species, that is, species in which there is a single physical worker caste, but some individuals may be specialized in group transport.

2.6 Deadlock and Stagnation Recovery

Sometimes, the item's motion can no longer progress either because forces are applied by ants in opposite directions and cancel one another, or because the group has encountered an obstacle or any significant heterogeneity on the substrate. We have already mentioned that a single ant, who first discovers a food item, tries to transport it alone: the ant first tries to carry it, then to drag it; an unsuccessful ant tries another direction and/or another position and then, if still unsuccessful, gives up the prey temporarily to recruit nestmates. The same phenomenon occurs when ants are engaged in group transport: if, for any reason, the item is stuck, ants exhibit realigning and repositioning behaviors [45,47]. The frequency of spatial rearrangements, which may result from the ants' response to the reactive forces communicated through the item being transported [47], increases with time, and so does the rate of transport. As is the case for solitary transporters, realignments tend to occur before, and are much frequent than, repositionings: only when realignment is not sufficient do ants try to find other slots around the prey.

Along the same lines, Moffett [36] reports that ants (*Pheidologeton diversus*) gather around food items at the site of their discovery, gnawing on them and pulling them; during the first ten minutes or so, the item is moved about slowly in shifting directions, before ants "sort out" their actions and actual transport can begin. During these ten minutes, a lot of spatial rearrangements take place.

Personal observations of weaver ants *Oecophylla longinoda* confirm the existence of such spatial rearrangements in this species too.

Van Damme and Deneubourg [51] studied cooperative transport of *Tenebrio molitor*'s larvae (a worm) in the ant *Formica polyctena*, and found that after a period of unsuccessful attempts to transport the larvae individually or in




Fig. 1. Distance over which a larva of *Tenebrio molitor* has been transported by *Formica polyctena* ants as a function of time. Eight experiments are shown. After Van Damme and Deneubourg [51], reprinted by permission.

group, transport suddenly becomes successful, one possible reason being that the forces applied by the various individuals engaged in cooperative transport become aligned. Figure 1 shows the distance over which a larva has been transported as a function of the time elapsed since the larva was discovered. Distance is positive when progress has been made toward the nest and negative otherwise. It can be clearly seen that a “phase transition” occurs at some point (which, however, cannot be predicted: it varies from experiment to experiment), when group transport suddenly becomes successful. After that transition, transport proceeds smoothly until the larva reaches the nest.

3 Cooperative Transport by Robots

From the previous section, we understand better, although not perfectly, how ants cooperate in collective prey transport. In this section we introduce cooperative transport by robots, more precisely cooperative box-pushing. Box-pushing requires a cooperative effort from at least two robots to move a box along some trajectory [8,37,15,27,38,44,33]. Of the multi-robot tasks including foraging and formation marching, box-pushing has generally used a combination of centralized planning and conflict resolution with explicit communication between robots to coordinate their actions.

In the following sections we describe a series of work, by Kube and Zhang [27–29,31] and Kube [30], consistent in its ant-based approach to the problem of cooperative transport by a group of robots. The initial task under study was undirected box-pushing, in which a group of robots found a box and

pushed it in a direction that was dependent upon the initial configuration. The task evolved into directed box-pushing, with the robots pushing the box from an initial position towards a fixed goal position. Finally, the transport task, a variant of the directed box-pushing task in which multiple goals were sequenced, is presented in which the robots to push the box from one location to the next.

The initial simulation model was implemented in a group of five physical robots [27]. Then, inspired by Sudd’s observations of group prey retrieval [45,47], stagnation recovery behaviors were added [29] and an approach to task modeling [31]. Currently the system consists of a group of homogeneous robots capable of transporting large boxes between arbitrary goal positions.

3.1 *From Social Insects to Robots*

Social insects are nature’s proof-by-example of a decentralized multiagent system whose control is achieved through locally sensed information, as Section 2 clearly suggests. In earlier work, we began with a simple simulation of a swarm of robots designed to locate and push a box and then implemented a subset of the behaviors in five physical robots [27,28].

A robot’s box-pushing controller was modeled as three sensors connected to two actuators through a set of five behaviors. A goal sensor was used to locate the box while a robot sensor provided information on the closest robot and an obstacle sensor warned of objects in close proximity. Left and right wheel motors used for steering were the two actuators. A modified fixed priority subsumption architecture [7] for behavior arbitration was used with the five behaviors, listed in ascending order:

- (1) FIND is the default motion behavior moving the robot forward along a gradual arc.
- (2) FOLLOW causes the robot to follow the closest robot within view.
- (3) SLOW switches the two speed wheel motors from medium to low.
- (4) GOAL moves the robot towards the box.
- (5) AVOID moves the robot away from an obstacle.

A simple taxis-based stimulus-response mechanism maps sensors to actuators. Inspired by Braitenberg’s Vehicles [5] and observations of social insects, sensors provide input to behaviors which map primitive discrete motion commands to left and right wheel motors.

In a single simulation timestep, each behavior takes its connected sensors and calculates an appropriate motor response with the highest priority behavior taking effect. The result is a controller with the FOLLOW and GOAL behaviors




Fig. 2. A FOLLOW and GOAL behavior produce coordinated motion towards a box. While the AVOID behaviour causes robot “A” to disperse around the box until an empty spot is found.

producing coordinated motion toward the box and the AVOID behavior causing robots to disperse around its perimeter (see Figure 2).

Keeping robots together using a FOLLOW behavior had an advantage, in the nondirected box-pushing task, of distributing the robots around the same half of the box. *Behavior preferences* allowed a behavior to narrow an attached sensor’s field-of-view. For example, the initial setting of the robot-sensor allowed the robot to see in a forward facing hemisphere. This setting resulted in robots breaking from a swarm when they sensed a closer robot who may happen to be moving in the opposite direction. By narrowing the view of the sensor while in the following state passing robots are ignored. In this way a behavior could adapt its sensings to suit the immediate purpose.

Based on the simulation results five physical robots were built with controllers containing two behaviors: AVOID and GOAL. The behaviors were implemented as combinational logic which mapped a left and right obstacle sensor to left and right wheel motors respectively, causing the robot to move away from obstacles. Left and right box sensors were mapped to right and left wheel motors causing the robot to turn towards a brightly lit box. The resulting controllers allowed the robots to locate the box, converge and push in a number of directions depending on the number of robots per side. The box was weighted such that at least two robots were required to move the box as it was pushed toward the edge of the test area.

The coordination demonstrated was possible by using a common goal and behaving using a simple “noninterference principle” [40]. The result demonstrated a simple cooperation without direct communication, although indirect communication occurs through *stigmergy* (see [21]).

Further simulation results showed that the success rate for nondirected box-

pushing⁴ increases as a function of the number of robots up to a point that appeared dependent on the size of the box. However, the system would stagnate or deadlock when an equal number of pushing robots surrounded the box resulting in an even distribution of box forces. To solve the stagnation problem we turned our attention back to cooperative prey retrieval by ants.

4 Stagnation Recovery and Mass Effect

A detailed study of cooperative prey retrieval in ants by Sudd uncovered several strategies used to combat stagnation [45,47]. If during transport the food item becomes stuck ants will realign their body orientation without releasing their grasp, as was described in more detail in Section 2. This has the effect of changing the direction of the pulling or pushing forces and was often sufficient to resume motion. If after several minutes realignment fails, the ant will release their prey and reposition themselves along the perimeter. Repositioning seems to result in larger cumulative changes in the forces acting on the transport item and was often successful in resuming motion. Could realignment and repositioning behaviors be used for stagnation recovery in box-pushing robots?

Our simulation experiments compared box-pushing controllers which included stagnation recovery behaviors [29]. The results demonstrated that the application of random pushing motions by either realigning the pushing angle or repositioning the pushing force was an effective technique against stagnation (see Figures 3 and 4). The results also showed that the task success rate and efficiency improved as a function of the number of robots. However, efficiency measured as the number of simulation timesteps, improved to a point that appeared dependent on the number of robots able to fit along the box.

The realignment behavior produced a small random change in pushing angle while the reposition behavior caused the robot to change the point of contact with the box. The box would translate or rotate if the resultant force or torque exceeded a threshold. Stagnation was detected by a robot if it was in contact with the box after an elapsed period without also detecting forward motion.

Ordering of the realignment and reposition behaviors was accomplished with timeout thresholds. For example, realignment became active at $t_c + X$ where t_c is the time the robot contacted the box. Reposition became active at $t_c + 4X$ with t_c reset each time the robot moved.

⁴ Where success was defined as pushing the box 200 units in 2000 simulated timesteps.




Fig. 3. The results comparing box-pushing controllers with and without stagnation recovery behaviors. Success was defined as pushing the box 200 units from an initial position within 2000 simulation timesteps. Each data point represents the average of 25 trials with percentage successful shown as a function of the number of robots. Reprinted with permission ©IEEE.

Figure 4 and 3 show the success percentage (reliability) and execution time (efficiency) of the controllers as a function of group size for the four different strategies: (1) Without stagnation recovery; (2) Realignment only; (3) Reposition only; and (4) Realignment & Reposition behaviors. Controller reliability was improved by including stagnation recovery. For small group size strategy (2) is best while (3) is more successful when the group size is large. Strategy (4) falls between (2) and (3) as expected since (2) then (3) are activated in sequence. Figure 4 compares controller efficiency as a function of simulated time. When the group size is small strategy (1) is best *provided it is successful*. When the group size is large strategy (4) is both the fastest and most successful with (2) and (3) having intermediate performance.

What do we learn from these results for real ants, assuming that this is a good model of cooperative transport in ants? There are two interesting results for biologists:

1. Although adding one or two mechanisms for stagnation recovery increases retrieval time, it also increases the probability of success. In species for which




Fig. 4. The task efficiency results of four controllers with and without stagnation recovery behaviors. Shown are the average number of simulated timesteps taken to move the box 200 units from its initial position as a function of the number of robots. Reprinted with permission ©IEEE.

competition with other colonies or other species is not a critical factor, the probability of success should be favored: we expect stagnation recovery mechanisms in such species. On the other hand, if retrieval time is crucial, for example because of competition, then reliability is less critical but the speed of retrieval is essential: in species facing strong competition, we expect a less sophisticated or even no stagnation recovery mechanism. This prediction can be tested.

2. The probability of success in retrieving the prey depends on group size, which itself depends on prey size (see Section 2). What the results tell us is that we might observe realignment only for small prey (small groups) and both realignment and repositioning for larger prey (larger groups). It is also possible that both realignment and repositioning may be observed for all group sizes, but it is unlikely that repositioning only could be observed because it is significantly less reliable for small group sizes. These predictions, again, can be tested.

5 A Task Model for Directed Box-Pushing

The directed box-pushing task required a new approach to task modeling. In undirected box-pushing the state of the robot’s controller was determined by selecting the highest priority behavior from a small select set. As a goal direction to push the box was added, it became apparent that accomplishing the task would involve a series of phases or steps often executed in a repetitive manner. Success would rely on the redundancy of mass effect. In each step, some of the previous behaviors would not be useful and could cause interference. For example, if a robot was correctly positioned for pushing the box towards the goal, then obstacle avoidance was not needed and if accidentally triggered would cause the progression to halt.

A termites nest, with its well defined mushroom shape, is constructed through a series of building steps. Each construction phase is thought to be governed by a building program with step transition specified as stimulus cues. In fact, this communication through the environment is the basis of Grassé’s Stigmergy Theory [21]. Thus describing a task as a series of steps with the transition between the steps specified as locally sensed cues formed the basis of our approach to task modelling [31].

In this section we briefly describe the directed box-pushing model and present new results of experiments using four different box types transported between two goal positions.

5.1 Coherent Behavior

In order to get coherent behavior from a group a robots each robot has an identical task controller which is composed of subtask controllers designed to accomplish each step of the task. The controllers are finite state machines (Q-machines) with state transition specified using binary sensing predicates, which we call *perceptual cues*.

The transport task is defined by three states: FIND-BOX, MOVE-TO-BOX and PUSH-TO-GOAL. Each state is implemented as a subtask controller designed to accomplish a single function. The two perceptual cues used to determine the state of the transport system are ?BOX-DETECT (BD) and ?BOX-CONTACT (BC). The states are shown in Table 1.

Each subtask controller is a finite state machine with states represented as *primitive actuation* (PA) behaviors. PA behaviors use motion primitives to move the robot. Each of the three subtask controllers are specified using states as shown in Tables 2 - 4.

Transport Controller		
Perceptual Cue (Input)		Task State (Output)
BD	BC	Subtask Controller
0	X	FIND-BOX
1	0	MOVE-TO-BOX
1	1	PUSH-TO-GOAL

Table 1

Task execution state is determined by two perceptual cues: ?BOX-DETECT (BD) and ?BOX-CONTACT (BC).

FIND-BOX Subtask Controller			
Perceptual Cue (Input)			Behavior State (Output)
?TOUCH	?CONTACT-	?AVOID-	PA Behavior
0	0	0	RANDOM-WALK
0	0	1	AVOID
0	1	X	CONTACT
1	X	X	BACK-OFF

Table 2

The FIND-BOX Q-machine is the subtask controller used to locate the box to be manipulated. Input is from the listed perceptual cues which define the output behaviour state specified as a primitive actuation (PA) behaviour. The “X” in the input table indicates a *don’t care* term. The perceptual cues corresponding to the dashed labels are: ?CONTACT-DETECT; ?AVOID-DETECT.

5.2 Perceptual Cues for Box-Pushing

Transporting a box from an unknown initial position towards a final goal destination was modeled using three types of perceptual cues. Obstacle avoidance cues were used to detect an obstacle and trigger avoidance behaviors. Box detection cues were used to locate and track a moving box, as well as, to control state transitions among the task step controllers. And a goal detection cue was used to indicate proper robot orientation, with respect to the goal, for a pushing behavior. The cues are designed with a given set of motor actions in mind. The design and implementation of each perceptual cue involve the following steps:

- (1) **Sensor Placement** Given a sensor type, determine the position, orientation and number of sensors to be used in the sensor system.
- (2) **Data Collection** For a given environment, collect data from the sensor that represents the condition under which the task is performed.

MOVE-TO-BOX Subtask Controller				
Perceptual Cue (Input)				Behavior State (Output)
?TOUCH	?CONTACT-	?AVOID-	?BOX-	PA Behavior
0	0	0	1	SEEK-BOX
0	0	1	X	AVOID
0	1	X	X	CONTACT
1	X	X	X	BACK-OFF

Table 3

The MOVE-TO-BOX Q-machine is the subtask controller that moves the robot towards any side of the brightly lit box to be manipulated. Input is from the listed perceptual cues which define the output behavior state specified as a primitive actuation (PA) behavior. The “X” in the input table indicates a *don’t care* term. The perceptual cues corresponding to the dashed labels are: ?CONTACT-DETECT; ?AVOID-DETECT; and ?BOX-DETECT.

PUSH-TO-GOAL Subtask Controller	
Perceptual Cue (Input)	Behavior State (Output)
?SEE-GOAL	PA Behavior
0	REPOSITION
1	PUSH-BOX

Table 4

The PUSH-TO-GOAL Q-machine is the subtask controller that either pushes the box towards a goal destination or repositions the robot on another position of the box to be manipulated. Input from the ?SEE-GOAL perceptual cue, which determines pushing angles, can vary the acceptable pushing angles.

- (3) **Data Analysis** Determine what features of the data may be used to meet the perceptual cue’s specification.
- (4) **Algorithm Design** Design an algorithm to extract the desired feature.
- (5) **Algorithm Verification** Specify the tests to verify that the cue performs as designed.

5.2.1 Obstacle Detection Cues

The purpose of the obstacle detection cues are to provide obstacle distance information to the robot. Three discrete thresholds are used corresponding to the distances of: less than 25 cm, less than 12.5 cm, and in physical contact with the robot. Active infrared emitter/detector pairs are used to provide non-contact obstacle information for the left and right front of the robot. Contact obstacle detection is determined using a single bit contact switch. The obstacle

detection cues are defined as:⁵

?OBSTACLE Return right and left true flags indicating the corresponding obstacle sensor has exceeded the input threshold.

?TOUCH Return a true flag if the front contact switch is pressed.

5.2.2 *Box Detection Cues*

Three perceptual cues are used for box detection:

?BOX-DIRECTION Return right and left true flags indicating the corresponding box sensor has exceeded the input threshold.

?BOX-DETECT Return a true flag if either left or right box sensors exceed a given input threshold.

?BOX-CONTACT Return a true flag if ?TOUCH is true AND either right or left box sensors exceed a given input threshold.

Box detection is simplified by using a bright light placed at the center of the box. The box detection cue asks the question: Can the robot see the box-light? The answer depends on the robot's distance from the box and the orientation of its two forward pointing sensors with respect to the box. An adjustable threshold varies the range at which the box-light is detectable and is dynamically determined as a function of ambient light. Recognizing physical contact with the box is a combination of two different types of sensing, touch and light intensity. This combination of stimulus is unique in the task's environment simplifying box recognition.

5.2.3 *Goal Detection Cue*

The goal detection cue asks the question: Can the robot see the goal? The answer is a function of the robot's orientation with respect to the goal indicator, which in this instance is a spotlight placed near the ceiling. The goal detection cue is defined as:

?SEE-GOAL Return a true flag if a signal peak greater than the input threshold is detected within the user defined field-of-view.

The final design consists of a narrow field-of-view sensor which is swept by a motor in an upward pointing arc. If a signal peak occurs, caused by the spotlight, within an adjustable window the goal is detected. The box detection sensors which face horizontally are shielded from light sources above the robot,

⁵ Perceptual cues will be identified by their leading question mark.

while the goal detection sensors face upward and therefore the goal-light does not interfere with the box-light.

At any given moment a robot is controlled by a single PA behavior. These behaviors issue discrete actions which affects the robot's orientation. As in the simulated robots, orientation is based on a taxis mechanism.

5.3 *Taxis-based Discrete Action*

Jander defines insect orientation as “the capacity and activity of controlling location and attitude in space and time with the help of external and internal references i.e. stimuli.” [26]. In insects the behavioral act of orientation is controlled either externally, and results in a directional orientation using a taxis mechanism, or internally under kinesthetic control. Taxis is defined by Webster's as a reflex translational or orientational movement by a freely motile organism in relation to a source of stimulation [52]. Inspired by Braitenberg's Vehicles [5] robot actions are based on taxis orientation or kinesthetic orientation as fixed motion patterns. The resulting action is used to create motor behaviors used in a reactive controller. The only required knowledge about the perception side of the robot is that it corresponds to a left and right division of the mobility system used to produce the actions. In other words, the input to the action model is a stimulus as measured by a sensor and does not depend on either the stimuli's modality or magnitude.

In box-pushing motion is restricted to translation and rotation in two dimensions. All robot motor actions, therefore, result in changes in position and orientation with respect to a given coordinate frame. To facilitate a quick response to changes in sensor data, a reactive control system is used for motor actions.

A mobility base was built and used differential steering as its means for achieving changes in translation and rotation. Discrete motion primitives were developed to be used as the underlying mechanism for all actions taken by the system. Perceptual processes are designed independently, but rely on the taxis model and its differential steering method for mobility.

Primitive actuation behaviors are classified into three groups: positive taxis or goal driven, which provide a change in orientation or translation *towards* a stimulus; negative taxis or avoidance driven, which effect a change in orientation or translation *away* from a stimulus; and kinesthetically driven, which execute a fixed action sequence in response to stagnating or deadlock conditions.

A wheel motor is controlled using two parameters: speed and direction of ro-

Positive and Negative Taxis Mappings				
Stimulus		Negative Taxis		Positive Taxis
L	R	AVOID	CONTACT	SEEK-BOX
0	0	<i>null</i>	<i>null</i>	<i>null</i>
0	1	<i>left-turn</i>	<i>left-rotate</i>	<i>right-turn</i>
1	0	<i>right-turn</i>	<i>right-rotate</i>	<i>left-turn</i>
1	1	<i>right-turn</i>	<i>right-rotate</i>	<i>forward</i>

Table 5

The positive and negative taxis behavior mappings. Behaviors that cause directional changes based on external stimuli expect a stimulus from the left and right sides of the robot similar to stimulus sensing found in insects. The “null” output means the behavior doesn’t produce a motion command.

tation. Speed is proportional to the applied input voltage and a fixed speed setting is used in all motion commands except while applying a pushing force. Continuous motion is accomplished by issuing a series of discrete motion commands, each of which moves the robot a small incremental amount. The commands have the general form: *begin(action), wait Δt , end(action)*.

A positive taxis or goal driven behavior moves the robot towards a given external stimulus. Input to the behavior takes the form of a left and right divided stimulus pair which may correspond to left and right sensors on the robot. The input variables to the behavior are boolean and indicate the presence or absence of the stimulus within a given range and field-of-view. Output from the behavior is a motion command selected from a set of four commands representing the possible number of input combinations. In the case of a behavior with a single input variable, 0 is mapped to the *null* motion command and 1 is mapped to the *forward* command. For the box-pushing task two goal driven behaviours are:

- SEEK-BOX - moves the robot towards a box.
- PUSH-BOX - pushes the box by increasing motor voltage.

In the same manner negative taxis or avoidance driven behavior repels a robot from a given stimulus. For the box-pushing task the two avoidance driven behaviors are:

- AVOID - *turns* the robot away from obstacles.
- CONTACT - *rotates* the robot away from obstacles.

The motor behaviours which cause changes in orientation are summarized in Table 5.

Kinesthetic orientation is used to produce motion in the absence of external stimuli and for stagnation recovery movements. In the case of both positive and negative taxis, orientation of the robot is under control of external stimuli. At any time the motor behavior relies on an external stimulus to decide the correct response in orientation. However, many behavioral acts in both insects and robots lack the external stimulus needed to guide the orientation mechanism. Rather a correct behavioral response might simply be a fixed pattern of motor activity stored in memory and released under suitable conditions. For example, a spider can return to a given location by “remembering and kinesthetically controlling its movements,” a skill also found in bees and ants [26].

In the absence of stimuli, a fixed pattern of motor activity can serve as a strategy while foraging for food or searching for a goal. For instance, when an ant leaves its nest to search for food it leaves in a straight line until it encounters either food or an odor trail which it then follows using a positive odor-taxis mechanism [53]. In box-pushing, a search strategy called RANDOM-WALK is used which keeps the robot moving in a forward direction by issuing a sequence of motion primitives. Continuous motion by the robot in the absence of any external stimulus is thus accomplished.

Recovery from deadlock or stagnation is the second use of kinesthetic orientation. During the execution of a task by robots using reactive control strategies, the absence of a plan can result in a condition in which the execution of the task gets stuck or is said to stagnate. For example, a dead end is reached by a robot trying to navigate to a given goal as in Arkin’s box canyon problem [12]. The problem is similar to finding a local maximum, encountered by hill-climbing algorithms, when the goal is to find the global maximum. In nondirected box-pushing the net force applied by the robots may equal zero if the robots are evenly distributed around the perimeter of the box. In such a case, a robot might attempt indefinitely to push the box unsuccessfully. Kinesthetic orientation, in the form of fixed action sequences and triggered by either the presence or absence of a controlling stimulus, is one solution suitable to the stagnating conditions in the box-pushing task.

6 Group Size in Cooperative Transport

Stigmergy, a term coined by French biologist P. Grassé, which means to incite work by the effect of previous work [21] is a principle finding its way from the field of social insects to collective robotics [2,48]. With their limited repertoire of behavioral acts social insects display an amazing competence in building nest structures. From the simple nests produced by the blind bulldozing of ants [19] to the termite homes that stand over a meter tall [46] all of which result from common task coordination that does not appear to depend on

interaction between the agents, but rather on the object they act upon. In this section, the results are presented for the integrated models of the previous sections. This global action is demonstrated in the collective transport task.

Global action is the effect produced when a set of identical mobile robots execute the common task of pushing an object towards an arbitrarily specified goal position. Coordination is achieved without resort to direct inter-robot communication or robot differentiation. Instead, context sensitive subtask controllers decompose the box transport task into three phases. The phases describe *what* is to be achieved, in terms of the externally observable events described by box position, without specifying *how* the task is to be accomplished by way of a unique path.

6.1 *Experimental System*

The experimental setup used to gather the data presented in the sequel consisted of a robot environment, in which various boxes were placed along with two spotlights used to indicate final goal positions, and a set of identical mobile robots complete with sensors and Q-machine task controllers. In total over 100 box-pushing trials were run using from one to 11 robots, four different box types and in three different venues. The final set of experiments were recorded on over four hours of video tape with an individual trial lasting between 30 seconds and five minutes. Described briefly is both the robot environment and hardware used.

6.1.1 *Robot Environment*

The ideal test environment would be a large open space without walls leaving the robots free to push the box along any desired path. Since this environment was not available a smaller and more restrictive area defined by walls was used. A permanent space large enough in which to conduct experiments was often difficult to find, resulting in the creation of a portable testing environment consisting of: 11 robots, two spotlights on stands for goal position indicators, the box to be manipulated, and a video camera to record the results. However, the majority of the experiments were conducted in the area depicted in Figure 5 which became available towards the end of this study.

6.2 *Robot Hardware*

The system is composed of a set of homogeneous two-wheeled robots, each weighing 1.3 kilograms and measuring approximately 18 centimeters in height




Fig. 5. In each trial the box was placed at initial position three meters from the goal line and the robots were placed at one of the indicated starting positions labelled $P_1 - P_5$. After Kube and Zhang [31]. Reprinted by permission ©Kluwer Academic Publishers.

and diameter as shown in Figure 6. A battery allows for 45 minutes of operation with a 10 minute recharge time. A Motorola 68HC11 microcontroller with 8K of RAM and programmed in Forth is used to map sensor output to one of nine motion primitives. A minimum number of sensors (6) was sought in implementing the perceptual cues. Additional sensors would allow a more omnidirectional field-of-view in the case of obstacle and box sensing and better pushing orientation in the case of box contact sensing, but the objective was to determine what could be accomplished with the minimal number of sensing bits. The hardware proved to be robust with few breakdowns.

6.2.1 Directed Box-Pushing

Increasing the number of robots from two to six did not affect the successful outcome of the transport experiments. This is an analogous result to the simulation results (shown in Figure 10) in which successful task completion remained high despite an increase in the number of robots. However, no claim is being made that task completion time is not affected, since completion times were found to vary as the number of robots increased and were dependent on available resources. In each of the 58 successful trials recorded the box was pushed from an initial starting position, located approximately in the center of a five by four meter area, towards the goal area indicated in Figure 5 and ending in quadrant I at a distance of at least 2.5 meters. The robots were started in each trial from positions P_1 to P_5 in quadrants II-IV shown in Figure 5. Successful trials would run between 32 and 214 seconds and were executed in three phases.




Fig. 6. Each of the robots are equipped with two forward pointing infrared obstacle sensors, one touch sensor, two CdS box-tracking photocells, and a destination sensor, all mounted on a differentially steered base.

The first phase began when the robots were powered on, the box-light was off and the goal-light was on. System initialization consists of taking ambient light readings used to set the box-detection threshold. The robots began executing FIND-BOX and quickly dispersed themselves in the area. Shortly thereafter, the box-light was turned on and those robots that were facing the box and sufficiently close would move towards and make contact with a boxside using the MOVE-TO-BOX controller. Depending on an individual robot's position, with respect to the box when box-detection occurred, the distribution of robots around the box would vary and mark the beginning of the second phase.

In the second phase, some of the robots incorrectly positioned for pushing, as determined by the PUSH-TO-GOAL controller, began moving counterclockwise around the box perimeter searching for an open spot on a correct side. This behavior is the result of several cycles through the transport Q-machine consisting of in turn FIND-BOX, MOVE-TO-BOX and PUSH-TO-GOAL subtask controllers and can be explained as follows. Once contact is made with a boxside the ?SEE-GOAL perceptual cue determines that the robot is on the wrong side for pushing. The PUSH-TO-GOAL controller then executes the REPOSITION behavior which moves the robot away from the box in a counterclockwise direction. If the box is within view, determined by the ?BOX-DETECT cue, MOVE-TO-BOX brings the robot into contact with a new position on the box providing it is unoccupied. The obstacle avoidance behaviors keep a robot away from occupied positions on a boxside. If the box is not within view then FIND-BOX executes and searches for the box.

The PUSH-TO-GOAL controller evaluates the new position and the cycle repeats.

The third and final phase is characterized by the box moving towards the goal position. Once a net force sufficient to move the box occurs the box begins to translate and possibly rotate. During the box movement phase a robot continuously determines if it remains on the correct side for pushing using the ?SEE-GOAL cue. A robot located at the edge of the pushing swarm may suddenly lose site of the goal and begin repositioning. The resulting drop in pushing force may be sufficient to halt the box movement until another robot joins the group effort. The dynamics of both the box and robots is such that the path taken by the box towards the goal is seldom straight. Rather, box movement can be said to converge towards the goal since its trajectory is the net result of several force vectors applied by individual robots. A typical box path might begin at position P_0 proceed towards P_6 and then move to P_7 as illustrated in Figure 5.

6.2.2 Pushing Different Box Types

To evaluate the controller’s sensitivity to object geometry, 38 successful trials were performed using six robots and four different box types. The initial box, BOX A, tested was 42 centimeters square and large enough for two 18 centimeter robots on a side. A second 84 centimeter square box, BOX C, was built by extending the initial box with a second frame. This increased the box dimensions, but used the same base on which the box slid along the floor. A third 84 centimeter box, BOX B, was built on a new base which increased the number of points in contact with the floor and therefore its sliding friction. The fourth box, BOX D, was round with a diameter of 84 centimeters and the results of the 39 trials can be summarized as follows:

- BOX A. A total of 10 trials were successful in pushing BOX A from the initial position to the goal positions in quadrant I (see Figure 9). The robots started from positions P_{1-5} . In general as the number of robots increased the task took longer to complete as the robot interference was high since the limited box side space created competition among the robots.
- BOX B. A total of eight trials were successful in transporting BOX B from its initial position using 6 robots starting from position P_4 and ending at positions P_{5-7} .
- BOX C. A total of seven successful trials were recorded in which BOX C was moved to the goal area by six robots starting from positions P_{2-4} . This box had the highest failure rate among the four boxes used and was due to a robot getting caught on the frame.
- BOX D. A total of 14 trials using a round box, BOX D, and four to six robots were successful in moving the box between two goal positions. The

round box was the last box built and experienced the most success of the four types tested. The lack of corners provided the robots with a uniform contact surface to push against unlike the square boxes which had sharp points at its corners.

6.2.3 Changes in Goal Position

The initial success of the directed box-pushing task led to the following extension which increased the task difficulty. Pictured in Figure 7 are two goal positions labelled P_A and P_B . The robots begin from position P_4 and a goal-light at position P_A is illuminated causing the robots to push the box towards P_A . Once reached the goal-light at P_A is turned off and the goal-light at P_B is switched on. The robots reposition around the box and begin pushing towards the goal at P_B . Figure 8 is a sequence of three images taken from a video segment in which two goals were used. A total of eight successful trials using three different goal positions were recorded using a single box.




Fig. 7. A schematic of the lab environment used to test the transport of a round box between two goal positions. Shown are the initial positions of the five robots and the box. The first step is to move the box from its initial position to the goal located at P_A . The second step moves the box from P_A to position P_B . The goal positions are indicated with a bright spotlight positioned at a height of 2.5 meters. To sequence the task steps the spotlight at position P_A is turned off and the light at P_B is turned on when the box reaches P_A .

In the following discussion some interesting secondary results are presented which compare execution times as a function of system size in the first experiment and as a function of object geometry in the second experiment along with the following caveat. In experiments involving physical mobile robots, holding the many system variables invariant is near impossible making comparisons based on execution runtimes tenuous at best. In this experimental-



Fig. 8. Shown are five robots pushing a round box from its initial position first towards a goal-light in the right of the picture and then towards a goal-light on the left of the picture. The mpeg video from which this sequence was taken is available at <http://www.cs.ualberta.ca/~kube/>

ist approach to robotics “things change” is axiomatic. Coefficients of friction change because the floor gets dusty, force is reduced because batteries run down, motors wear reducing repeatability, wheels slip in response to changes in load and the list goes on. However, in general there still seems to be a trend in the data making it worth presenting.

6.2.4 System Size

The mean execution time for moving the smaller 42 centimeter square box from its initial position to the goal positions were compared for two to six robots as shown in Figure 9. Starting positions for the robots were varied and included $P_{1,3-5}$ with the final end position of the box recorded for timing to be $P_{5,7}$. Indicated in each plot are the number of trials used to compute the mean. The large variance in runtimes was due to robot start positions $P_{1,5}$ which could result in long repositioning phases ⁶. In general, execution times increased as a function of the number of robots due to the increase in robot interference competing for the limited box space. A much larger number of trials is needed for any statistical conclusions.

6.2.5 Convex Object Geometry

Our previous simulation study had shown that in a box-pushing task performance, as measured by completion time or success rate, ⁷ could be improved if stagnation recovery behaviors were added to the controller to avoid deadlock from occurring when the robots applied an equal distribution of forces to the box [29]. What was also noted was the sudden drop in performance as the

⁶ Both the maximums indicated in the case of three and five robots occurred from P_5 .

⁷ Success was defined to be the movement of the box by 200 units in under 2000 simulation timesteps.




Fig. 9. The mean execution time of moving a 42^2 centimeter box 2.5 meters towards a goal position (P5, P6, P7) as a function of the number of robots. For each plot the number of trials as well as the minimum and maximum run times are indicated. A boxside is approximately twice the robot’s diameter and increasing the number of robots increases the robot interference as they compete for the limited space available.




Fig. 10. The effects of doubling box contact space on the task success rate. The results from two simulation experiments in which the only parameter changed was the robot’s diameter, with the size of the box side fixed at 90 units. Robot diameters of 20 and 10 were compared for a task in which a box was moved 200 units from its initial position. Each data point is the average of 25 simulation runs each with a different random initial configuration.

size of the system grew for controllers without stagnation recovery. This was conjectured to be due to the number of robots able to fit on a box side. To test this hypothesis, simulations were run for the same behavior controller and the




Fig. 11. The effects of doubling box contact space on execution time. The results from two simulation experiments showing execution time versus system size. The only parameter varied was the size of the robot; the size of the box side was held constant at 90 units.

robot diameter (RD) was tested for $RD = 10$ and compared with the results using $RD = 20$. The results are shown in Figures 10 and 11. If the diameter of the robots were reduced, for a fixed box side, the performance increases, which leads to the conjecture that for a given task, performance is dependent on some yet to be determined *task density* function.

In Figure 12 the mean execution times were compared for the four box types and six robots starting from the same initial position. In general, it appears that as the available contact space increases more robots are able to participate in pushing at the same time reducing the time taken to complete the task. However, due to the sparseness of the data additional experiments would allow statistical conclusions.

7 Discussion: From Ants to Robots and Back

Visualize a room in which a group of robots sit in one corner and a large box sits approximately center with a spotlight placed in another corner. The robots begin moving and soon disperse into the room. Soon after the box-light comes on the robots begin moving towards it and eventually come into contact with a side. Then some of the robots beginning to reposition themselves by moving around the box in a counterclockwise fashion, while others which are correctly positioned begin to push the box towards the spotlight. The box begins to move in the direction of the lit corner of the room, but the path is not quite straight and veers to the right and the box movement stops. Again some of the




Fig. 12. The mean execution time of moving a box towards the goal as a function of box type. Box A is a 42 centimeter square box, Box B and C are 84 centimeter square boxes with B having a higher sliding friction than Box C, Box D is an 84 centimeter diameter round box. All box types are approximately the same weight and can be pushed by at least two robots. For each plot the number of trials as well as the minimum and maximum run times are indicated. All trials used six robots. Robot interference is minimized by increasing the available contact space around the box.

robots begin their counterclockwise repositioning and assume a new position more correctly oriented for pushing. Finally, the box begins to move in a new direction towards the goal-light. The dynamics of the swarm of robots is very reminiscent of the dynamics of ants represented in Figure 1: given that the implementation of individual robot behavior is based on ant behavior, it is encouraging that the robotic *model* produces the same kind of emergent collective behavior as ants, and it suggests that the robotic implementation constitutes a plausible model of cooperative transport in (some species of) ants.

Once the box reaches the goal position the spotlight turns off and a second goal light on the opposite corner of the room is illuminated. Now all the robots begin repositioning, eventually making it to the opposite side and begin to once again push the box towards the new goal destination. Robots leave the task, seemingly at random, and wander off only to return and join the group effort in transporting the box towards its goal. The experiments are repeated, this time with boxes of different shapes and sizes and the number of robots in a group are varied. Our video recordings shows, and those that have seen them agree, that the robots make a coordinated effort in pushing the box in a direction that converges towards the indicated goal position.

The results show in the many successful trials of directed box-pushing that

a coordinated group effort is possible without use of direct communication or robot differentiation. Rather a form of indirect communication takes place through the environment by way of the object being manipulated. For directed box-pushing, the control strategy was shown to be insensitive to system size, some convex object geometries and changing goal positions. The results of experiments with physical robots presented here, adds support to Arkin’s simulation studies which showed that cooperation in some tasks are possible without direct communication [1].

The data presented here also agrees in certain aspects with other studies in which stigmergy is used as the task coordinating mechanism. Stigmergy as proposed by Grassé is a model used to explain the regulation of building behavior in termites [21]. Stigmergy theory holds that transitions between a sequence of construction steps is regulated by the effect of previous steps. In more general terms, the theory has been used to explain and describe the process by which task activity can be regulated using only local perception and indirect communication through the environment as applied to algorithms for coordinating distributed building behaviour [48] and foraging tasks by multi-robot systems [2]. In the box-pushing task the results support the use of indirect communication through the environment as proposed by stigmergy theory. However, Downing and Jeanne found that stigmergy theory does not explain the use of additional cues, not dependent on previous steps, in regulating task execution in nest construction by paper wasps [16]. For collective robotics this means that perceptual cues can also be formed from stimuli other than that which are immediately available from the task itself. For example, in directed box-pushing the box-detection cues are adaptive to the ambient light level of the environment by specifying box-detection as a multiple of the ambient light level.

Coherent behavior from a collective system of robots must also account for task resource management. Coordination improves by minimizing antagonistic actions that can result from conflicts over limited resources. In box-pushing antagonistic forces are mitigated by increasing the available boxside space while enforcing a noninterference behavior. The data on transporting small boxes versus large boxes by the same number of robots confirms the observations made during task execution. For box-pushing, this result implies that group size is important for a fixed resource size in a given task and agrees with the result obtained by Beckers *et al.* [2] for a foraging task in which one to five robots were used to gather 81 objects randomly distributed in their environment then placing them into one large pile. Their study showed that group size was a critical factor in determining task efficiency and that increasing the number of robots used without increasing the available task resources increased task execution time due to the increase in inter-robot interference. In general, increasing task resources minimizes inter-robot interference. Thus, reducing robot interference increases group coordination and consequently leads to a

more efficient coherence as demonstrated by the decreasing execution times.

The coherent behavior displayed for the transport task can also be attributed to the common goal shared by the individual robots along with an identical set of interaction rules. This is the same effect noted by Seeley while considering the collective decision making in honey bees [43]. As an explanation for how a swarm of honey bees could reach the same decision on the profitability of several food sources, Seeley hypothesized that each bee's nervous system was calibrated in a similar manner. Since all members of the colony share the same rules for adjusting response thresholds, the bees can operate independently yet generate a collective response to various nectar sources. Thus common goals and common rules of interaction allow a decentralized decision making process to produce a coherent global response. By way of the social insects, nature is showing us how to build decentralized and distributed systems that are autonomous and capable of accomplishing tasks through the interaction of many simple and highly redundant agents. From their local perception to the mass effect that results in a global action these biological systems serve to elucidate the mechanisms thought to be at the heart of self-organizing behavior.

In return, the robotic system described in this paper tells us a lot about cooperative transport in ants. We have seen that the model makes predictions about the kind of stagnation recovery mechanism (if any) to be expected depending on ecological conditions. At a more fundamental level, because the model is able to reproduce many of the collective features of cooperative transport in ants and because it is based on plausible assumptions, it suggests that these assumptions may be sufficient to explain the behavior observed in ants. Many of the predictions of the model can now be tested empirically. Of course it can be argued that the actual robotic implementation was not needed: simulations were just as good. This is only partially true because the robotic implementation shows that the assumptions the model is based on can produce the expected behavior in the real world, that is, with real constraints, a result that is far from obvious as many factors (friction, heterogeneity, etc.) play a role in cooperative transport.

Acknowledgments. E. B.'s work is supported through the Interval Research Fellowship.

References

- [1] Arkin, R.C. Cooperation without Communication: Multiagent Schema-Based Robot Navigation. *Journal of Robotic Systems* **9** (1992): 351–364.
- [2] Beckers, R. and O.E. Holland and J.L. Deneubourg. ¿From Local Actions to

- Global Tasks: Stigmergy and Collective Robotics. *Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems Artificial Life IV* (1994): 181–189.
- [3] Bonabeau, E., M. Dorigo, and G. Théraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford, UK: Oxford University Press, 1998.
- [4] Bonabeau, E., G. Théraulaz, J.-L. Deneubourg, S. Aron, and S. Camazine. Self-Organization in Social Insects. *Trends in Ecol. Evol.* **12** (1997): 188–193.
- [5] Braitenberg, V. *Vehicles: Experiments in Synthetic Psychology*. MIT Press, 1984.
- [6] Breed, M. D., J. H. Fewell, A. J. Moore, and K. R. Williams Graded recruitment in a ponerine ant. *Behav. Ecol. Sociobiol.* **20** (1987): 407–411.
- [7] Brooks, Rodney A. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation* **RA-2** (1986): 14–23.
- [8] Caloud, P., W. Choi, J.C. Latombe, C. Le Pape and M. Yim. Indoor Automation with many mobile robots. *IEEE/RSJ International Conference on Intelligent Robots and Systems* (1990): 67–72.
- [9] Cao, Y. U., A. S. Fukunaga, and A. B. Kahng. Cooperative mobile robotics: antecedents and directions. *Autonomous Robots* **4** (1997): 7–2
- [10] Chauvin, R. Sur le transport collectif des proies par *Formica polyctena*. *Insectes Sociaux* **25** (1968): 193–200.
- [11] Chauvin, R. Les lois de l’ergonomie chez les fourmis au cours du transport d’objets. *C. R. Acad. Sc. Paris D* **273** (1971): 1862–1865.
- [12] Clark, R.J., R.C. Arkin and A. Ram. Learning Momentum: On-line Performance Enhancement for Reactive Systems. *IEEE International Conference on Robotics and Automation* (1992): 111–116.
- [13] Detrain, C., J.-L. Deneubourg. Origine de la diversité des réponses collectives des fourmis lors de la récolte de proies. *Actes Coll. Ins. Soc.* **10** (1996): 57–65.
- [14] Detrain, C., J.-L. Deneubourg. Scavenging by *Pheidole pallidula*: A key for understanding decision-making Systems in ants. *Anim. Behav.* **53** (1997): 537–54.
- [15] Donald, B., J. Jennings and D. Rus. Analyzing teams of cooperating Mobile robots. *IEEE International Conference on Robotics and Automation* (1994): 1896–1903.
- [16] Downing, H. A. and R. L. Jeanne. Nest construction by the paper wasp, *Polistes*: a test of stigmergy theory. *Animal Behaviour* **36** (1988): 1729–1739.
- [17] Downing, H. A. and R. L. Jeanne. The regulation of complex building behaviour in the paper wasp, *Polistes fuscatus* (Insecta, Hymenoptera, Vespidae). *Animal Behaviour* **39** (1990): 105–124.

- [18] Franks, N. R. Teams in social insects: group retrieval of prey by army ants (*Eciton burchelli*, Hymenoptera: Formicidae). *Behav. Ecol. Sociobiol.* **18** (1986): 425–429.
- [19] Franks, N. R., A. Wilby, B. W. Silverman and C. Tofts. Self-organizing nest construction in ants: sophisticated building by blind bulldozing. *Animal Behaviour* **44** (1992): 357–375.
- [20] Gotwald, W. H. Death on the march. *Rotunda* **Fall/Winter** (1984/1985): 37–41.
- [21] Grassé, P. La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. La théorie de la stigmergie: essai d'interpré. *Insectes Sociaux* **6** (1959): 41–81.
- [22] Grassé, P.-P. *Termitologia, Tome II. Fondation des Sociétés. Construction*. Paris: Masson, 1984.
- [23] Hölldobler, B., and E. O. Wilson. The multiple recruitment systems of the African weaver ant *Oecophylla longinoda* (Latreille). *Behav. Ecol. Sociobiol.* **3** (1978): 19–60.
- [24] Hölldobler, B., R. C. Stanton, and H. Markl. Recruitment and food-retrieving behavior in *Novomessor* (Formicidae: Hymenoptera). I. Chemical signals. *Behav. Ecol. Sociobiol.* **4** (1978): 163–181.
- [25] Hölldobler, B. Territorial behavior in the green tree ant (*Oecophylla smaragdina*). *Biotropica* **15** (1983): 241–250.
- [26] Jander, R. Insect Orientation. *Annual Review of Entomology* **8** (1963): 95–114.
- [27] Kube, C. R., and H. Zhang. Collective robotic intelligence. In *Proceedings Second International Conference on Simulation of Adaptive Behavior: From Animals to Animats 2*, 460–468. Cambridge, MA: MIT Press, 1992.
- [28] Kube, C. R., and H. Zhang. Collective robotics: from social insects to robots. *Adaptive Behavior* **2** (1994): 189–218.
- [29] Kube, C. R., and H. Zhang. Stagnation recovery behaviors for collective robotics. In *Proceedings 1994 IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, 1883–1890. Los Alamitos, CA: IEEE Computer Society Press, 1995.
- [30] Kube, C. R. *Collective Robotics: From Local Perception to Global Action*. Ph.D., University of Alberta, 1997.
- [31] Kube, C. R., and H. Zhang. Task modelling in collective robotics. *Auton. Robots* **4** (1997): 53–72.
- [32] Markl, H. and B. Hölldobler. Recruitment and food-behavior in *Novomessor* (Formicidae, Hymenoptera). II. Vibration signals. *Behav. Ecol. Sociobiol.* **4** (1978): 183–216.

- [33] Mataric, M.J., M. Nilsson and K. T. Simsarian. Cooperative Multi-Robot Box-Pushing. *IEEE International Conference on Intelligent Robots and Systems* (1995): 556–561.
- [34] Melhuish, C., O. Holland, and S. Hoddell. Collective sorting and segregation in robots with minimal sensing. (1998): Preprint.
- [35] Meyer, J.-A. Sur les lois régissant l’accomplissement d’une tâche collective complexe chez *Formica polyctena*. *Compte Rendus de l’Academie des Sciences Paris D* **270** (1970): 2111–2114.
- [36] Moffett, M. W. Cooperative food transport by an Asiatic ant. *National Geog. Res.* **4** (1988): 386–394.
- [37] Noreils, F.R. An Architecture for Cooperative and Autonomous Mobile Robots. *IEEE International Conference on Robotics and Automation* (1992): 2703–2710.
- [38] Parker, L.E. ALLIANCE: An architecture for fault tolerant, cooperative control of heterogeneous mobile robots. *IEEE/RSJ/GI International Conference on Intelligent Robots and Systems* (1994): 776–783.
- [39] Pratt, S. C. Recruitment and other communication behavior in the ponerine ant *Ectatomma ruidum*. *Ethology* **81** (1989): 313–331.
- [40] Premvuti, S. and Yuta, S. Consideration on the Cooperation of Multiple Autonomous Mobile Robots. *IEEE International Workshop on Intelligent Robots and Systems* (1990): 59–63.
- [41] Robson, S. K., and J. F. A. Traniello. Key individuals and the organisation of labor in ants. In *Information Processing in the Social Insects* (Detrain, C., J.-L. Deneubourg, and J. M. Pasteels). Birkhauser, in press.
- [42] Robson, S. K., and J. F. A. Traniello. Resource assesment, recruitment behavior, and organization of cooperative prey retrieval in the ant *Formica schaufussi* (Hymenoptera: Formicidae). *J. Insect Behav.* **11** (1998): 1–22.
- [43] Seeley, T.D., S. Camazine and J. Sneyd. Collective decision-making in honey bees: how colonies choose among nectar sources. *Behavioral Ecology and Sociobiology* **28** (1991): 277–290.
- [44] Stilwell, D.J. and J. S. Bay. Toward the Development of a Material Transport System using Swarms of Ant-like Robots. *IEEE International Conference on Robotics and Automation* (1993): 766–771.
- [45] Sudd, J. H. The transport of prey by an ant *Pheidole crassinoda*. *Behaviour* **16** (1960): 295–308.
- [46] Sudd, J. H. How insects work in groups. *Discovery* **24** (1963): 15-19.
- [47] Sudd, J. H. The transport of prey by ant. *Behaviour* **25** (1965): 234–271.
- [48] Theraulaz, G. and E. Bonabeau. Coordination in Distributed Building. *Science* **269** (1995): 686–688.

- [49] Traniello, J. F. A. Social organization and foraging success in *Lasius neoniger* (Hymenoptera: Formicidae): behavioral and ecological aspects of recruitment communication. *Oecologia* **59** (1983): 94–100.
- [50] Traniello, J. F. A., and S. N. Beshers. Maximization of foraging efficiency and resource defense by group retrieval in the ant *Formica schaufussi*. *Behav. Ecol. Sociobiol.* **29** (1991): 283–289.
- [51] Van Damme, T., and J.-L. Deneubourg. Cooperative transport in *Formica polyctena*. *Submitted*.
- [52] *Webster's Ninth New Collegiate Dictionary* Merriam-Webster Inc., 1985.
- [53] Wilson, E.O. and Hölldobler, B. *The Ants*. The Belkap Press of Harvard University Press, 1990.
- [54] Wojtusiak, J., E. J. Godzinska, and A. Dejean. Capture and retrieval of very large prey by workers of the African weaver ant *Oecophylla longinoda*. *Tropical Zool.* **8** (1995): 309–318.