

Adaptive Division of Labor in Large-Scale Minimalist Multi-Robot Systems

Chris Jones and Maja J. Matarić
Computer Science Department
University of Southern California
941 West 37th Place, Mailcode 0781
Los Angeles, CA 90089-0781
{cvjones|maja}@robotics.usc.edu

Abstract

A Large-Scale Minimalist Multi-Robot System (LMMS) is one composed of a group of robots each with limited capabilities in terms of sensing, computation, and communication. Such systems have received increased attention due to their empirically demonstrated range of capabilities and beneficial characteristics, such as their robustness to environmental perturbations and individual robot failure and their scalability to large numbers of robots. However, little work has been done in investigating ways to endow such a LMMS with the capability to achieve a desired division of labor over a set of dynamically evolving concurrent tasks, a necessity in any task-achieving LMMS. Such a capability can help to increase the efficiency and robustness of overall task performance as well as open new domains in which LMMS can be seen as a viable alternative to more complex control solutions. In this paper we present a method for achieving a desired division of labor in a LMMS, experimentally validate it in a realistic simulation, and demonstrate its potential to scale to large numbers of robots and its ability to adapt to environmental perturbations.

I. Introduction and Motivation

A Large-Scale Minimalist Multi-Robot System (LMMS) is a multi-robot system composed of a large number of robots, each having limited capabilities in terms of sensing, computational power, and communication range and bandwidth. We define a *minimalist robot* as one which maintains little or no state information, extracts limited, local, and noisy information from its available sensors, and lacks the capability for active communication with other robots. Due to these limited capabilities, the world in which a minimalist robot is situated is formally partially-observable and highly non-stationary, and it is therefore not practical to assume that such a robot is capable of reliably knowing a significant portion of the current global state of the environment or of overall task progress.

These limitations in sensing, communication, and computation preclude a minimalist robot from performing tasks requiring significant computation or communication capabilities. Nonetheless, minimalist robots have been shown to be highly effective at a number of collective tasks, such as multi-robot formation control (Fredslund & Matarić 2002), collection tasks (Goldberg & Matarić 2002), and robotic soccer (Werger 1999). A system composed of a large number of such minimalist robots has the potential of conferring advantages including increased robustness to individual robot failure as no single robot is critical to task performance, the prospect of scaling to increasingly larger numbers of robots as there are few bottlenecks in terms of complex communication, planning, or coordination requirements, and increased adaptability to changes in the environment since individuals act based on local information and are not tied to globally coordinated plans.

The aim of this work is to investigate a method by which to endow a LMMS with the capability to achieve a desired division of labor over a set of dynamically evolving concurrent tasks, a critical requirement of any task achieving large-scale multi-robot system. We define division of labor as the phenomenon in which individuals in a multi-robot system concurrently execute a set of tasks. The division of labor may need to be continuously adjusted in response to changes in the task environment or group performance. The broader scope of this work is in understanding ways in which to achieve robust, scalable, and efficient coordination in a LMMS.

This paper is organized as follows. In Section II we provide the relevant related work. In Section III we give a detailed description of the concurrent foraging task domain we use as validation of our division of labor mechanism. In Section IV we describe our experimental setup used for empirical evaluation of our work. In Section V we present the robot controller we use to produce a division of labor in a LMMS. In Section VI we describe and analyze experimental results, and in Section VII we draw conclusions.

II. Related Work

Here we summarize briefly the related work in physical LMMS using robots with similar capabilities to those on which our system is based. Mataric (1995) provides early work on group coordination in LMMS using a collection of simple basis behaviors. Theraulaz, Goss, Gervet & Deneubourg (1991) and Agassounon & Martinoli (2002) present minimalist methodologies for coordination in robot groups. Beckers, Holland & Deneubourg (1994) demonstrate the capabilities of minimalist multi-robot systems in object clustering and sorting. Kube & Zhang (1996) present an approach to box-pushing using a group of robots with simple sensors and reactive control. Werger & Mataric (1996) present a minimalist solution in the multi-robot foraging domain. Martinoli, Ijspeert & Mondada (1999) present work on the probabilistic modeling of robot behavior in the task regulation domain, demonstrating its performance as compared to experiments on physical and simulated robots. Werger (1999) presents coordinated behavior in a robot soccer team using a minimalist behavior-based control system. Krieger & Billeter (2000) present a decentralized task allocation mechanism for large mobile robot groups based on individual task-associated response thresholds in a collection domain. Holland & Melhuish (2000) use probabilistic behavior selection in minimalist robotic clustering and sorting. Goldberg & Mataric (2002) precisely define the foraging task for LMMS and provide a collection of general distributed behavior-based algorithms and their empirical evaluation. Fredslund & Mataric (2002) present work on the problem of achieving coordinated behavior in the context of formations using a distributed group of physical robots using only local sensing and minimal communication.

In the multi-robot literature, there is work on more communication and computationally complex forms of task regulation in multi-robot systems through the use of publish/subscribe and market-based methods (Gerkey & Mataric 2002) and systems in which significant global state is made known to all robots (Parker 1998).

There is related work in the area of research that studies and simulates insect colonies and their behaviors. Theraulaz, E. & Deneubourg (1998) describe how the adaptability of complex social insect societies is increased by allowing members of the society to dynamically change tasks (behaviors) when necessary. Giving that ability to robots allows a LMMS to operate in domains requiring the simultaneous regulation of many tasks. Bonabeau, Theraulaz & Deneubourg (1996) describe a model of a task regulation mechanism in insect societies through the use of response thresholds for task-related stimuli. Theraulaz et al. (1998) extended that model by introducing an adaptive threshold that changes over time based on

individual task performance.

The division of labor mechanism we present can be considered an instance of a response threshold model as presented in Bonabeau et al. (1996), Krieger & Billeter (2000), Theraulaz et al. (1998), and Agassounon & Martinoli (2002). However, our task domain and division of labor mechanism differ in that the task-related stimuli are perceived locally by the individual robots and are not altered as a result of task performance. Furthermore, the individual robots are initially homogeneous, as opposed to Krieger & Billeter (2000) in which robot are initially assigned different response thresholds, and the robots do not learn or become specialized through adaptive response thresholds as is the case in Theraulaz et al. (1998) and Agassounon & Martinoli (2002).

III. Concurrent Foraging Task Domain

In order to experimentally study a mechanism for providing a LMMS with division of labor capabilities, we investigated division of labor in a concurrent foraging task domain. *Concurrent foraging*, a variation on traditional foraging, consists of an arena populated by multiple types of objects to be collected. Each robot is equally capable of foraging all object types, but can only be allocated to foraging for one type at any given time. Additionally, all robots are engaged in foraging at all times (i.e., a robot cannot be idle). A robot may switch the object type according to its control policy, when it determines it is appropriate to do so. It is desirable for a robot to avoid thrashing (i.e., wasting time and energy) by needlessly switching the object type for which it is foraging.

A. Task Description

Our experimental domain of concurrent foraging requires multiple object (puck) types to be foraged from a circular arena. Initially, the arena is randomly populated by two types of pucks: Puck_{Red} and Puck_{Green}, which are distinguishable by their color.

In this task, the robots move around an enclosed arena and pick up encountered pucks. When a robot picks up a puck, the puck is consumed (i.e., it is immediately removed from the environment, not transported to another region) and the robot carries on foraging for other pucks. Immediately after a puck is consumed, another puck of the same type is placed in the arena at a random location. The reason for this replacement is to maintain constant puck density in the arena through the course of an experiment. In some situations, the density of pucks can have an affect on the division of labor performance. This is an important consideration in mechanisms for division of labor in LMMS for many domains; however, in this work

we want to limit the number of experimental variables impacting system performance. Therefore, we reserve the investigation on the impact of varying puck densities on division of labor in LMMS for future work.

The division of labor portion of the task requires the robots to split their numbers by having some forage for Puck_{Red} pucks and others for Puck_{Green} pucks. For the purpose of our experiments, we desire a division of labor such that the proportion of robots foraging for Puck_{Red} pucks is equal to the proportion of Puck_{Red} pucks present in the foraging arena (e.g., if Puck_{Red} pucks make up 30% of the pucks present in the foraging arena, then 30% of the robots should be foraging for Puck_{Red} pucks). In general, the desired division of labor could take other forms. For example, it could be related to the relative reward or cost of foraging each puck type without change to our approach.

As was stated earlier, due to their minimalist capabilities, individual robots do not have direct access to global information such as the size and shape of the foraging arena, the initial or current number of pucks to be foraged (total or by type), or the initial or current number of foraging robots (total or by foraging type). Also, it cannot be assumed that any robot or subset of robots will always be operational or the proportion of pucks will remain constant over time.

IV. Simulation Environment

All simulations were performed using Player and Stage. Player (Gerkey, Vaughan, Støy, Howard, Sukhatme & Mataric 2001) is a server that connects robots, sensors, and control programs over the network. Stage (Vaughan 2000) simulates a set of Player devices. Together, the two represent a high-fidelity simulation tool for individual robots and robot teams which has been validated on a collection of real-world robot experiments using Player and Stage programs transferred directly to physical Pioneer 2DX mobile robots.

The experimental arena (a close-up is shown in Figure 1) was used in all experiments presented in this paper. The arena, shown populated with robots and pucks, is circular and has an area of approximately 315 square meters.

V. The Robots

The robots used in the experimental simulations are realistic models of the ActivMedia Pioneer 2DX mobile robot. Each robot, approximately 30 cm in diameter, is equipped with a differential drive, an odometry system




Fig. 1. Close-up of the arena used in experiments; example pucks and robots are shown.

using wheel rotation encoders, 8 evenly spaced sonars covering the front 180 degrees used for obstacle avoidance, and a forward-looking Sony color camera with a 60-degree field-of-view and a color blob detection system (used for puck and robot detection and classification through color). Each robot is also equipped with a 2-DOF gripper on the front, capable of picking up a single puck at a time. There is no capability available for explicit, direct communication between robots nor can pucks and other robots be uniquely identified.

A. Behavior-Based Controller

All robots have identical behavior-based controllers consisting of the following mutually exclusive behaviors: Avoiding, Wandering, Visual Servoing, Grasping, and Observing. Descriptions of the behaviors used in the division of labor implementation are given below.

- The **Avoiding** behavior causes the robot to turn to avoid obstacles in its path.
- The **Wandering** behavior causes the robot to move forward and after a random length of elapsed time turn left or right through a random arc for a random period of time.
- The **Visual Servoing** behavior causes the robot to move toward a detected puck of desired type. If the robot's current foraging state is Robot_{Red} the desired puck type is Puck_{Red}, and if the robots current foraging state is Robot_{Green} the desired puck type is Puck_{Green}.
- The **Grasping** behavior causes the robot to use the gripper to pickup and consume a puck within the gripper's grasp.
- The **Observing** behavior causes the robot to take an

image from its camera and record the detected pucks and robots to their respective histories. The robot then updates its foraging state based on the puck and robot histories. A description of the histories is given in Section V-B and a description of the foraging state update procedure is given in Section V-C.

Each behavior listed above has a set of activation conditions based on relevant sensor inputs and state values. When met, the conditions cause the behavior to become active. A description of when each activation condition is active is given below. The activation conditions of all behaviors are shown in Table I.

- The **Obstacle Detected** activation condition is true when an obstacle is detected by the sonar within a distance of 1 meter. Pucks are not detectable by the sonar, so are therefore not considered obstacles.
- The **Puck_{Det} Detected** activation condition is true if the robot's current foraging state is Robot_{Det} and a puck of type Puck_{Det} (where Det is Red or Green) is detected by the color camera within a distance of approximately 5 meters and within ± 30 degrees of the robot's direction of travel.
- The **Gripper Break-Beam On** activation condition is true if the break-beam sensor between the gripper jaws detects an object;
- The **Observation Signal** activation condition is true if the distance traveled by the robot according to odometry since the last time the **Observing** behavior was activated is greater than 2 meters.

B. State Information

The robots maintain three types of state information: foraging state, observed puck history, and observed robot history. The foraging state identifies the type of puck the robot is currently involved in foraging. A robot with a foraging state of Robot_{Red} refers to a robot engaged in foraging Puck_{Red} pucks and a foraging state of Robot_{Green} refers to a robot engaged in foraging Puck_{Green} pucks.

Each robot is outfitted with a colored beacon observable by nearby robots which indicates the robot's current foraging state. The color of the beacon changes to reflect the current state – a red beacon for a foraging state of Robot_{Red} and a green beacon for Robot_{Green}. Thus, the colored beacon acts as a form of local, passive communication conveying the robot's current foraging state. All robots maintain a limited, constant-sized history storing the most recently observed puck types and another constant-sized history storing the foraging state of the most recently observed robots. Neither of these histories contain a unique identity or location of detected pucks or robots, nor does it store a time stamp of when any given observation was made. The history of observed pucks is limited to the last MAX-PUCK-HISTORY pucks observed

and the history of the foraging state of observed robots is limited to the last MAX-ROBOT-HISTORY robots observed.

While moving about the arena, each robot keeps track of the approximate distance it has traveled by using odometry measurements. At every interval of 2 meters traveled, the robot makes an observation. An observation consists of the robot taking the current image from its color camera and, using simple color blob detection, classifying all currently visible pucks and robots through their respective colors and adding them to their respective histories. This procedure is nearly instantaneous; therefore, the robot's behavior is not outwardly affected. The area in which pucks and other robots are visible is within 5 meters and ± 30 degrees in the robot's direction of travel. Observations are only made after traveling 2 meters because updating too frequently leads to over-convergence of the estimated puck and robot type proportions due to repeated observations of the same pucks and/or robots. On average, during our experiments, a robot detected 2 pucks and robots per observation.

C. Foraging State Transition Functions

After it makes an observation, the robot re-evaluates its current foraging state given the newly updated puck and robot histories. We have experimented with several functions, given below, to determine the conditions at which a robot should change its current foraging state.

The first method is a simple step transition function. The condition in which a robot with a current foraging state of Robot_{Green} will change its foraging state to Robot_{Red} is given by Change(Green-Red), shown in Equation 1. Similarly, the condition in which a robot with a current foraging state of Robot_{Red} will change its foraging state to Robot_{Green} is given by Change(Red-Green), shown in Equation 2. In Equations 1 and 2, RR is the proportion of Robot_{Red} entries in the Robot History and RP is the proportion of Puck_{Red} entries in the Puck History.

$$\text{Change}(\text{Green} - \text{Red}) = \begin{cases} \text{Yes}, & \text{if } RR \leq RP, \\ \text{No}, & \text{otherwise} \end{cases} \quad (1)$$

$$\text{Change}(\text{Red} - \text{Green}) = \begin{cases} \text{Yes}, & \text{if } RR \geq RP, \\ \text{No}, & \text{otherwise} \end{cases} \quad (2)$$

The second method we explored uses a probabilistic transition function. The probability that a robot with a current foraging state of Robot_{Green} will change its foraging state to Robot_{Red} is given by the probability, P(Green-Red), shown in Equation 3. Similarly, the probability that a robot with a current foraging state of Robot_{Red} will change its foraging state to Robot_{Green} is given by

Obstacle Detected	Puck _{Det} Detected	Gripper Break-Beam On	Observation Signal	Active Behavior
X	X	X	1	Observing
1	X	X	X	Avoiding
0	1	0	0	Visual Servoing
0	X	1	0	Grasping
0	X	X	X	Wandering

TABLE I

BEHAVIOR ACTIVATION CONDITIONS. BEHAVIORS ARE LISTED IN ORDER OF DECREASING RANK. HIGHER RANKING BEHAVIORS PREEMPT LOWER RANKING BEHAVIORS IN THE EVENT MULTIPLE ARE ACTIVE. X DENOTES THE ACTIVATION CONDITION IS IRRELEVANT FOR THE BEHAVIOR.

the probability, $P(\text{Red-Green})$, shown in Equation 4. In Equations 3 and 4, RR is the proportion of Robot_{Red} entries in the Robot History and RP is the proportion of Puck_{Red} entries in the Puck History.

$$P(\text{Green} - \text{Red}) = 100^{(RP-RR)} / 100.0 \quad (3)$$

$$P(\text{Red} - \text{Green}) = 100^{((1.0-RP)-(1.0-RR))} / 100.0 \quad (4)$$

VI. Experimental Results

We experimentally validated our LMMS division of labor mechanism using both the step and probabilistic transition functions presented in Section V-C in the realistic simulated environment described in Section IV. All experiments used 20 robots and 50 pucks and all presented results have been averaged over 20 experimental runs.

To test the adaptability of the division of labor mechanism to external perturbations in puck type proportions, they were dynamically changed at various times during the experimental trials. The experiments begin with 30% Puck_{Red} and 70% Puck_{Green} pucks. At time 6000 seconds, the relative proportion of pucks are changed to 80% Puck_{Red} and 20% Puck_{Green} pucks, and at time 12000 the relative proportions are changed to 50% Puck_{Red} and 50% Puck_{Green} pucks. The total number of pucks remains constant throughout the experiment.

The plots in Figures 2 and 3 show a comparison between the performance of the step and the probabilistic transition functions using MAX-PUCK-HISTORY and MAX-ROBOT-HISTORY values of 2, 10, and 50. Both the step and probabilistic transition functions converge to a stable division of labor. The use of the step transition function leads to faster convergence to a stable division of labor at the expense of more high frequency oscillation.

When the puck proportions are equal, both transition functions converge to the desired division of labor. As the puck type proportions become more skewed, with one




Fig. 2. The proportion of Puck_{Red} pucks and robots foraging for Puck_{Red} pucks over time when using the step transition function with different puck and robot history lengths.

puck type becoming significantly more prevalent than the other, the systems using shorter history lengths tend to converge slightly off of the desired division of labor. This is in part due to the insufficient granularity in their puck and robot histories.

Another factor in evaluating the efficiency of these methods is through the frequency by which individual robots switch between tasks. In some task domains, switching between tasks can be very expensive and therefore it should be avoided. Figures 4 and 5 show the cumulative number of times the robots change state during the course of the experiments. The data points are obtained by summing the total number of forage state changes over the course of the previous 50 seconds of the experiment (it is possible that a single robot could change foraging state




Fig. 3. The proportion of Puck_{Red} pucks and robots foraging for Puck_{Red} pucks over time when using the probabilistic transition function with different puck and robot history lengths.

more than once during this interval). As the plots show, the shorter the puck and robot history lengths, the more foraging state changes occur. Also, for puck and robot histories of the same length, the use of the probabilistic transition function leads to fewer foraging state changes than the step transition function.

In general, shorter puck and robot history lengths result in faster convergence to the desired division of labor but lead to higher frequency oscillations due to more frequent changes in individual robot foraging state. For increasingly small histories, convergence to a division of labor short of the desired is exhibited as the puck type proportions become skewed. The probabilistic transition function results in more stable convergence with fewer robot state changes and fewer high frequency oscillations in the division of labor. What transition function and history lengths are to be for an arbitrary task environment depends on factors such as how quickly the environment changes, how often the environment changes, and the expense of changing tasks.

VII. Conclusions

We have presented a Large-Scale Minimalist Multi-Robot System (LMMS), composed of 20 simulated mobile robots, in which the individual robots maintain a minimal amount of state information, extract a limited amount of information from available sensors, and cannot actively




Fig. 4. The number of foraging state changes when using the step transition function with different puck and robot history lengths.

or directly communicate with other robots in the system. Using this LMMS, we have demonstrated a method by which to achieve a desired division of labor in a concurrent foraging task domain, experimentally validated it in a realistic simulation, and demonstrated its robustness and adaptability to environmental perturbations.

VIII. Acknowledgments

This work is supported in part by DARPA TASK Grant F30602-00-2-0573 and in part by NSF ITR Grant EIA-0121141.

IX. REFERENCES

- Agassounon, W. & Martinoli, A. (2002), Efficiency and robustness of threshold-based distributed allocation algorithms in multi-agent systems, *in* 'First Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems', ACM Press, Bologna, Italy, pp. 1090–1097.
- Beckers, R., Holland, O. & Deneubourg, J. (1994), From local actions to global tasks: Stigmergy and collective robotics, *in* R. Brooks & P. Maes, eds, 'Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems', MIT Press, Cambridge, MA, pp. 181–189.
- Bonabeau, E., Th´eraulaz, G. & Deneubourg, J. (1996), 'Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies', *Proceedings Royal Society of London B* **263**, 1565–1569.
- Fredslund, J. & Matari´c, M. J. (2002), 'A general, local algorithm for robot formations', *IEEE Transactions on Robotics and Automation, Special Issue on Multi-Robot Systems* **18**(5), 837–846.

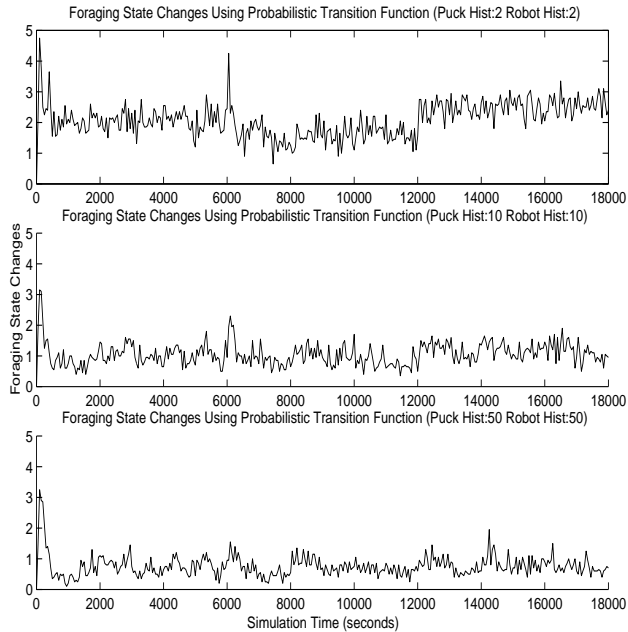


Fig. 5. The number of foraging state changes when using the probabilistic transition function with different puck and robot history lengths.

Behavior: From Animals to Animats', MIT Press, Paris, France, pp. 346–355.

Vaughan, R. (2000), 'Stage: A multiple robot simulator', *Institute for Robotics and Intelligent Systems Technical Report IRIS-00-394, Univ. of Southern California*.

Werger, B. B. (1999), 'Cooperation without deliberation: A minimal behavior-based approach to multi-robot teams', *Artificial Intelligence* **110**, 293–320.

Werger, B. & Matarić, M. (1996), Robotic "food" chains: Externalization of state and program for minimal-agent foraging, in P. Maes, M. Matarić, J. Meyer, J. Pollack & S. Wilson, eds, 'From Animals to Animats 4, Fourth International Conference on the Simulation of Adaptive Behavior (SAB-96)', Cape Cod, Massachusetts, pp. 625–634.

Gerkey, B. P. & Matarić, M. J. (2002), 'Sold!: Auction methods for multi-robot coordination', *IEEE Transactions on Robotics and Automation, Special Issue on Multi-Robot Systems* **18**(5), 758–768.

Gerkey, B., Vaughan, R., Støy, K., Howard, A., Sukhatme, G. & Matarić, M. (2001), Most valuable player: A robot device server for distributed control, in 'IEEE/RSJ International Conference on Intelligent Robots and Systems, (IROS-01)', Maui, Hawaii, pp. 1226–1231.

Goldberg, D. & Matarić, M. J. (2002), Design and evaluation of robust behavior-based controllers for distributed multi-robot collection tasks, in T. Balch & L. E. Parker, eds, 'Robot Teams: From Diversity to Polymorphism', AK Peters, pp. 315–344.

Holland, O. & Melhuish, C. (2000), 'Stigmergy, self-organization, and sorting in collective robotics', *Artificial Life* **5**(2), 173–202.

Krieger, M. B. & Billeter, J.-B. (2000), 'The call for duty: Self-organized task allocation in a population of up to twelve mobile robots', *Robotics and Autonomous Systems* **30**(1-2), 65–84.

Kube, C. & Zhang, H. (1996), The use of perceptual cues in multi-robot box-pushing, in 'IEEE International Conference on Robotics and Automation', Minneapolis, Minnesota, pp. 2085–2090.

Martinoli, A., Ijspeert, A. J. & Mondada, F. (1999), 'Understanding collective aggregation mechanisms: From probabilistic modeling to experiments with real robots', *Robotics and Autonomous Systems* **29**, 51–63.

Matarić, M. (1995), 'Designing and understanding adaptive group behavior', *Adaptive Behavior* **4**:1, 51–80.

Parker, E. L. (1998), 'Alliance: An architecture for fault tolerant multi-robot cooperation', *IEEE Transactions on Robotics and Automation* **14**(2), 220–240.

Thérault, G., E., B. & Deneubourg, J. (1998), 'Threshold reinforcement and the regulation of division of labour in insect societies', *Proceedings Royal Society of London B* **265**, 327–335.

Thérault, G., Goss, S., Gervet, J. & Deneubourg, J. (1991), Task differentiation in polistes wasp colonies: A model for self-organizing groups of robots, in J. Meyer & S. Wilson, eds, 'International Conference on Simulation of Adaptive