

Task Allocation via Self-Organizing Swarm Coalitions in Distributed Mobile Sensor Network

Kian Hsiang Low and Wee Kheng Leow

Department of Computer Science
National University of Singapore
3 Science Drive 2, Singapore 117543, Singapore
{lowkh, leowwk}@comp.nus.edu.sg

Marcelo H. Ang, Jr.

Department of Mechanical Engineering
National University of Singapore
10 Kent Ridge Crescent, Singapore 119260, Singapore
mpeangh@nus.edu.sg

Abstract

This paper presents a task allocation scheme via self-organizing swarm coalitions for distributed mobile sensor network coverage. Our approach uses the concepts of ant behavior to self-regulate the regional distributions of sensors in proportion to that of the moving targets to be tracked in a non-stationary environment. As a result, the adverse effects of task interference between robots are minimized and sensor network coverage is improved. Quantitative comparisons with other tracking strategies such as static sensor placement, potential fields, and auction-based negotiation show that our approach can provide better coverage and greater flexibility to respond to environmental changes.

Introduction

Sensor network has recently received significant attention in the areas of networking, embedded and multi-agent systems (Lesser, Ortiz Jr., & Tambe 2003) due to its wide array of real-world applications (e.g., disaster relief, environment monitoring). In these applications, the distributed sensing task is achieved by the collaboration of a large number of static sensors, each of which has limited sensing, computational, and communication capabilities.

One of the fundamental issues that arises in a sensor network is coverage. Traditionally, network coverage is maximized by determining the optimal placement of static sensors in a centralized manner, which can be related to the class of art gallery problems (O'Rourke 1987). However, recent investigations in sensor network mobility reveal that mobile sensors can self-organize to provide better coverage than static placement (Yadgar, Kraus, & Ortiz Jr. 2003). Existing applications have only utilized uninformed mobility (i.e., random motion or patrol) (Lesser, Ortiz Jr., & Tambe 2003). In contrast, our work here focuses on informed, intelligent mobility to further improve coverage.

Our network coverage problem is motivated by the following constraints that discourage static sensor placement or uninformed mobility: a) no prior information about the exact target locations, population densities or motion pattern, b) limited sensory range, and c) very large area to be observed. All these conditions may cause the sensors to be unable to

cover the entire region of interest. Hence, fixed sensor locations or uninformed mobility will not be adequate in general. Rather, the sensors have to move dynamically in response to the motion and distribution of targets and other sensors to maximize coverage. Inspired by robotics, the above problem may be regarded as that of low-level motion control to coordinate the sensors' target tracking movements in the continuous workspace. Alternatively, it can be cast as a high-level task allocation problem by segmenting the workspace into discrete regions (Fig. 1a) such that each region is assigned a group or *coalition* of sensors to track the targets within. We will now refer to mobile sensors as robots since they are the same in this paper's context.

This paper presents a two-level integrated approach (Fig. 1b) to distributed mobile sensor network coverage in complex, dynamic environments. At the lower level, each robot uses a reactive motion control strategy proposed by Low, Leow, & Ang, Jr. (2003) known as *Cooperative Extended Kohonen Maps* (EKMs) to coordinate their target tracking *within* a region without the need of communication. This strategy is also responsible for obstacle avoidance, robot separation to minimize task interference, and navigation between regions via beacons or checkpoints plotted by a motion planner. Low, Leow, & Ang, Jr. (2003) showed that cooperative EKMs perform more efficient target tracking than the well-known potential fields method (Parker 2002) in complex, unpredictable environments.

At the higher level, the robots use a dynamic *ant-based task allocation* scheme to cooperatively self-organize their coalitions in a decentralized manner according to the target distributions *across* the regions. This is the focus of the paper. It contrasts with the other works of biologically-inspired robot swarms (Balch & Arkin 1998; Matarić 1997) that emphasize control- rather than task-level cooperation. Our ant-based scheme addresses the following issues, which distinguish it from the other task allocation mechanisms:

Task Allocation for Multi-Robot Tasks Existing Multi-Robot Task Allocation (MRTA) algorithms (i.e., auction- and utility-based) (Gerkey & Matarić 2002; Parker 1998) generally assume that a multi-robot task can be partitioned into several single-robot tasks. But this may not be always possible or the multi-robot task can be more efficiently performed by coalitions of robots. Furthermore, the partitioned single-robot tasks are sometimes assumed to be indepen-

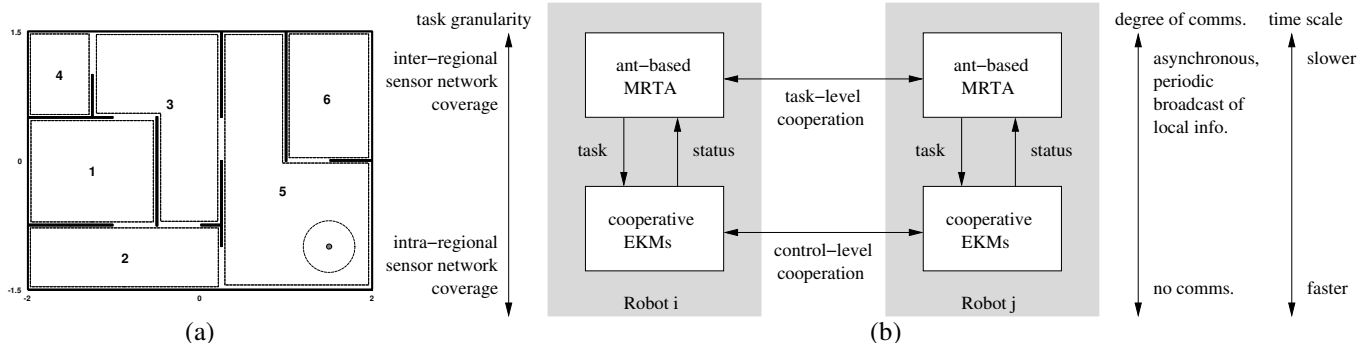


Figure 1: (a) A 4 m × 3 m environment that is divided into 6 regions. The circle at the bottom right represents the robot’s sensing radius of 0.3 m (drawn to scale). The environment is 42.44 times as large as the robot’s sensing area. (b) Two-level integrated framework for multi-robot cooperation (MRTA = Multi-Robot Task Allocation, EKM = Extended Kohonen Map).

dent, i.e., no interference would result. However, the robots are bound to interfere with each other’s ongoing activity either physically (e.g., space competition) or non-physically (e.g. shared radio bandwidth, conflicting goals). In the extreme case, when too many robots are involved, little or no work gets done as they totally restrict each other’s movement. Hence, task interference has an adverse effect on the overall system performance (Goldberg & Matarić 1997). Knowing that physical interference can be implied from robot density (Goldberg & Matarić 1997), our decentralized task allocation scheme dynamically distributes robots in real-time by estimating robot densities in different regions to minimize interference. Bucket brigade algorithm (Schneider-Fontañ & Matarić 1998) can also eliminate interference by assigning the robots to separate regions. However, it cannot respond in real-time to changes in regional distributions of targets due to target motion.

Coalition Formation for Minimalist Robots Existing multi-agent coalition formation schemes (Sandholm *et al.* 1999; Shehory & Kraus 1998) require complex planning, explicit negotiation, and precise estimation of coalitional cost. Hence, they may not be able to operate in real-time in a large-scale sensor network. Our task allocation method via self-organizing swarm coalitions is reactive, dynamic, and can operate with uncertain coalitional cost and resource-limited robots.

Cooperation of Resource-Limited Robots Robots with limited communication and sensing capabilities can only extract local, uncertain information of the environment. As such, distributed methodologies are required to process and integrate the noisy, heterogeneous information to improve its quality so that it can be effectively utilized to boost the task performance. Furthermore, if the robots have limited computational power, their cooperative strategies cannot involve complex planning or negotiations. Existing task allocation mechanisms have either assumed perfect communications, high computational power, centralized coordination or global knowledge of the task and agents (Gerkey & Matarić 2002; Krieger, Billeter, & Keller 2000; Parker 1998; Simmons *et al.* 2000). For example, recent coverage applications (Yadgar, Kraus, & Ortiz Jr. 2003) employ coalition leaders, one in each region, to negotiate with each other.

This negotiation is conducted iteratively using an auction-based mechanism and attempts to balance the proportion of robots to that of the targets across all regions. To do so, each coalition leader must be able to obtain the exact number of robots and targets in its region as well as the task performance of these robots. Furthermore, it has to synchronize its negotiation with the coalition leaders in other regions via long-range communication. Note that this negotiation can be conducted entirely by a central coordinator but it requires even more resources. In contrast, our proposed method does not require such expensive resources, thus catering to resource-limited robots. The robots endowed with our ant-based scheme require only local sensing information and short-range communication. The robot coalitions can also be self-organized asynchronously without negotiation.

Ant-Based Multi-Robot Task Allocation

Many multi-robot tasks, e.g., foraging, transportation, manipulation, sensing and exploration, have been inspired by social insects (Bonabeau, Dorigo, & Théraulaz 1999), in particular, ants. Our MRTA scheme encapsulates three concepts of ant behavior to self-organize the robot coalitions according to the target distributions across regions: (a) encounter pattern based on waiting time, (b) self-organization of social dominance, and (c) dynamic task allocation.

Encounter Pattern Based on Waiting Time

Encounter patterns provide a simple, local cue for ants with sensory and cognitive limitations to assess regional densities of ants and objects of interest, which are crucial to regulating the division of labor (Gordon 1999). Instead of relying on global communication to relay target positions and density estimation (Jung & Sukhatme 2002), our scheme uses encounter patterns to predict target density via local sensing. Regional robot density is captured in a similar way using local communication.

An encounter pattern can be derived from a series of waiting time or interval between successive encounters. This simple form of information processing has accounted for the complex adaptive process of task allocation in ant colonies (Hirsh & Gordon 2001). In our coverage task, the waiting time of a robot is defined in terms of its encounters with the other robots and targets. A robot encounter is defined

as a reception of a message from another robot in the same region. A target encounter is defined as an increase in the number of targets tracked between the previous and the current time steps. For a robot i in region r , the waiting time for other robots $w_{ir}(k)$ and targets $w'_{ir}(k)$ is the time interval between the $(k-1)$ th and k th encounters. Note that each waiting time is subject to stochastic variation. Hence, multiple samplings of waiting time have to be integrated to produce an accurate estimation of the regional density. The average waiting time $W_{ir}(k)$ between the $(k-1)$ th and k th robot encounters for a robot i in region r is computed as:

$$W_{ir}(k) = \frac{1}{n}w_{ir}(k) + \frac{n-1}{n}W_{ir}(k-1) \quad (1)$$

$$n = \min(k, n_{max})$$

where n_{max} is the maximum number of encounters that is monitored. This limit allows the robot to forget the early samplings of waiting time, which have become obsolete. The average target waiting time $W'_{ir}(k)$ is updated in the same manner. Both waiting times are updated according to the changing environment, and are inversely proportional to the robot and target densities in region r . The target density directly reflects the task demand of the region. The robot density reflects the amount of physical interference in the region, which is inversely proportional to the task demand. Therefore, the task demand $S_{ir}(k)$ of a region r can be determined by robot i as the ratio of the waiting times:

$$S_{ir}(k) = \frac{W_{ir}(k)}{W'_{ir}(k)}. \quad (2)$$

The task demand $S_{ir}(k)$ will be used in the self-organization of social dominance as well as in dynamic task allocation.

Self-Organization of Social Dominance

The division of labor in an ant colony is strongly influenced by its social dominance order (Camazine *et al.* 2001), which self-organizes to match the task demands of the colony and the changing environment. Our scheme is inspired by this concept to move robots out of a region that has a lower target-to-robot density ratio than the other regions. Instead of fixing the dominance order (Goldberg & Matarić 1997), the social dominance of the robots in each coalition is self-organized according to their individual task performance. Robots in the same coalition engage in dominance contests at a regular interval τ if they are within communication range. The winner increases its tendency to stay in the current region while the loser increases its tendency to leave the current region and join another coalition in other regions. When robot i encounters robot j in region r , the probability of robot i winning a contest against robot j is defined as:

$$P(\text{robot } i \text{ winning}) = \frac{n_i^2 S_{ir}^2}{n_i^2 S_{ir}^2 + n_j^2 S_{jr}^2} \quad (3)$$

where S_{ir} and S_{jr} are respectively the task demand of region r determined by robot i and robot j , and n_i and n_j are the number of targets currently under observation by robot i and robot j respectively. Equation 3 implies that robot i would

most likely win the contest if it observes more targets than robot j . However, if both are tracking the same number of targets, then their individual evaluation of the task demand can be used to differentiate them. This will distinguish a robot that has been observing the targets for a long time from another that just encounters the same number of targets.

To inject the influence of social dominance on the self-organization of robot coalitions, each time a robot i wins a contest (Eq. 3), it increases its tendency of staying in the current region, which is represented by the response threshold $\theta_i(t)$ to be used for dynamic task allocation:

$$\theta_i(t) = \theta_i(t-1) + \delta \quad (4)$$

where δ is small constant. Conversely, each time the robot loses, it decreases its tendency of staying in the region:

$$\theta_i(t) = \theta_i(t-1) - \delta. \quad (5)$$

θ_i varies in the range [0,1] to prevent robots from being overly submissive or dominating.

Dynamic Task Allocation

The distributed task allocation algorithm in ants can efficiently arrange the ants in proportion to the amount of work in the changing environment (Tofts 1993). In a similar manner, our scheme aims to self-organize the robot coalitions according to the target distributions across the regions.

Our dynamic task allocation scheme is based on the notion of response thresholds (Bonabeau, Dorigo, & Théraulaz 1999). In a threshold model, robots with low response thresholds respond more readily to lower levels of task demand than do robots with high response thresholds. Performing the task reduces the demand of the task. If robots with low thresholds perform the required tasks, the task demand will never reach the thresholds of the high-threshold robots. However, if the task demand increases, high-threshold robots will engage in performing the task.

MRTA strategies that utilize fixed response thresholds (Jung & Sukhatme 2002; Krieger, Billeter, & Keller 2000) are incapable of responding effectively to dynamic environments (Bonabeau, Dorigo, & Théraulaz 1999). In contrast, the thresholds in our scheme are continuously updated by the self-organizing process of social dominance.

To be effective in task allocation, a robot must at least have some knowledge of the task demands in its neighboring regions in order to make rational task decisions. To do so, robot i maintains a memory of the task demand S_{ir} of each region r (initialized to 0) and the amount of time T_{ir} that it previously spent in region r . T_{ir} can be used as a certainty measure of S_{ir} . In addition to computing S_{ir} using Equation 2, S_{ir} can also be updated when robot i receives a message from a neighboring robot j with S_{jr} less than S_{ir} . Then S_{ir} and T_{ir} are updated to take the values S_{jr} and T_{jr} respectively. In this manner, the task demands of the regions are kept in memory. Robot i can then predict which region has the greatest task demand and join that region. At every time interval of τ , if S_{ir} receives no update, the certainty value T_{ir} is decreased by τ while the task demand S_{ir} is increased by a small amount, such that its magnitude reflects the robot's motivation to explore.

Our distributed MRTA scheme uses a stochastic problem solving methodology. It is performed at intervals of τ to allow for multiple samplings of waiting time during each interval. The probability of a robot i to stay in its current region c is defined as:

$$P(\text{stay}) = \frac{S_{ic}^2}{S_{ic}^2 + (1 - \theta_i)^2 + T_{ic}^{-2}}. \quad (6)$$

On the other hand, the probability of a robot i to leave region c to go to region r is defined as:

$$P(\text{leave}) = \frac{S_{ir}^2}{S_{ir}^2 + \theta_i^2 + T_{ir}^{-2} + d_{cr}^2} \quad (7)$$

where d_{cr} is the pre-computed collision-free distance between region c and region r , which can be viewed as the cost of task switching. Note that a robot that loses in the dominance contest in a coalition does not always leave the region. If it experiences a higher task demand in its region than in other regions, it will have a high tendency of remaining in its coalition.

From Equations 6 and 7, if the robot does not respond to any regions, it will not switch task and will remain in the current coalition. The robot may also respond to more than one region. This conflict is resolved with a method that is similar to Equation 3. The probability of a robot i choosing a region r that it has responded to is:

$$P(\text{choose}) = \frac{(S_{ir} \ln T_{ir})^2}{\sum_r (S_{ir} \ln T_{ir})^2}. \quad (8)$$

If the robot i chooses region r that is not the current region c , then it will employ the reactive motion control strategy to move through the checkpoints plotted by the planner to region r . The generation of checkpoints is performed by the approximate cell decomposition method for motion planning (Low, Leow, & Ang, Jr. 2002).

Experiments and Discussion

This section presents quantitative evaluations of the ant-based MRTA scheme for distributed mobile sensor network coverage in a complex, unpredictable environment. The experiments were performed using Webots, a Khepera mobile robot simulator, which incorporated 10% white noise in its sensors and actuators. 12 directed distance sensors were modelled around its body of radius 3 cm. Each sensor had a range of 17 cm, enabling the detection of obstacles at 20 cm or nearer from the robot's center, and a resolution of 0.5 cm to simulate noise. Each robot could also sense targets and kin robots at 0.3 m or nearer from its center and send messages to other robots that were less than 1 m away via short-range communication. A 4 m \times 3 m environment (Fig. 1a) was used to house the Khepera robots and targets, which were randomly scattered initially. The number of robots varied from 5, 10 to 15, which corresponded to total robot sensing area of 11.8%, 23.6%, and 35.3% of the environment size. The mobile targets were forward-moving Braitenberg obstacle avoidance vehicles (Braitenberg 1984) that changed their direction and speed with 5% probability.

Sensor Network Coverage

The first performance index determines the overall sensor network coverage of the robots (Parker 2002):

$$\text{sensor network coverage} = \sum_{t=1}^T 100 \frac{n(t)}{NT} \quad (9)$$

where N is the total number of targets, n is the number of targets being tracked by the robots at time t , and the experiment lasts T amount of time. N and T are fixed respectively as 30 targets and 10000 time steps at intervals of 128 ms.

Using this index, a quantitative test was conducted to compare the network coverage of the robots adopting five distributed tracking strategies: (1) potential fields, (2) cooperative EKMs, (3) static placement, (4) auction-based negotiation, and (5) ant-based MRTA. Unlike the latter three strategies, potential fields and cooperative EKMs are reactive motion control techniques that do not involve explicit task allocation. With static placement, static sensors are placed at least 0.6 m apart to ensure no overlap in coverage. With auction-based negotiation and ant-based MRTA, the robots are fitted with cooperative EKMs to coordinate their target tracking within a region, avoid obstacles, and navigate between regions.

Test results (Fig. 2a) show that ant-based MRTA provides better coverage than the other strategies. The differences in coverage between any two strategies have been verified using t -tests ($\alpha = 0.1$) to be statistically significant. Notice that 5 mobile robots endowed with our method can track better than 10 static sensors. Although auction-based negotiation uses complex negotiation, longer communication range, and more information about the robots and targets, it does not perform better than our ant-based scheme. This will be explained in the section of degree of specialization.

Total Coalitional Cost

The second performance index determines the total coalitional cost (Shehory & Kraus 1998) of the robots. Given a set of connected regions where coverage tasks are to be performed, and a set A of M robots, the task allocation algorithm assigns a robot coalition $C_r \subseteq A$ to the coverage task in region r such that (a) $\bigcup_r C_r = A$, (b) $\forall r \neq s, C_r \cap C_s = \emptyset$, and (c) each C_r has a positive cost $|(n_r/N) - (m_r/M)|$ where m_r and n_r are the number of robots and targets in region r respectively and N is the total number of targets. The objective is to minimize the total coalitional cost (Shehory & Kraus 1998):

$$\text{total coalitional cost} = \sum_r \left| \frac{n_r}{N} - \frac{m_r}{M} \right|. \quad (10)$$

This index varies within the range [0,2]. A coalitional cost of 0 implies that the robot distribution over all regions is exactly proportional to the target distribution. In this manner, interference between robots is at its minimum, which will improve overall coverage. High costs imply the opposite.

Test results (Fig. 2b) show that auction-based negotiation and ant-based MRTA have the lowest coalitional costs. Hence, we can conclude from Figures 2a and 2b that, with

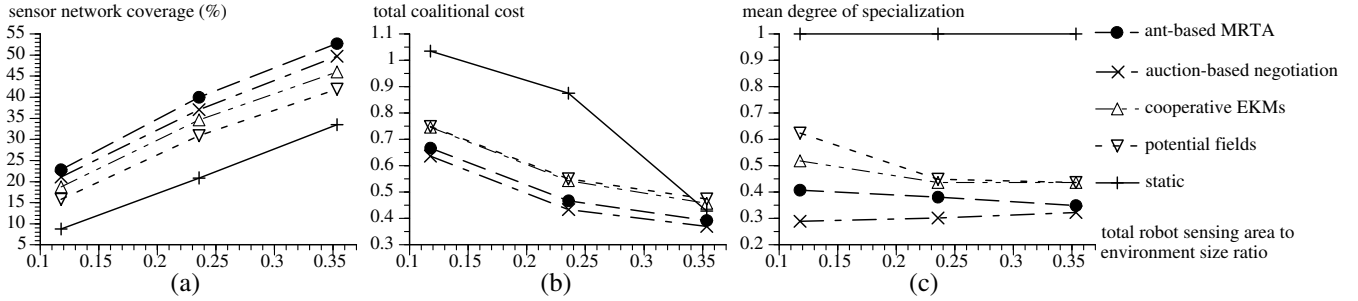


Figure 2: Comparison of performance using different motion control and task allocation strategies: (a) Sensor network coverage, (b) total coalitional cost, and (c) mean degree of specialization.

a lower cost, a higher coverage can be achieved. Although auction-based negotiation achieves slightly lower coalitional cost than ant-based MRTA, its coverage is lower. This will be explicated in the next section. Coalitional cost has been validated using t -tests ($\alpha = 0.1$) to be significantly different for various strategies except those without explicit task allocation (i.e., potential fields and cooperative EKMs). This is expected since they do not perform coalition formation, which account for their higher costs.

Coalitional cost is higher with fewer robots because with less robots, it is more difficult to achieve the same proportion of robots to that of the targets over all regions.

Degree of Specialization

To achieve low coalitional cost, the robot coalitions must be highly responsive, i.e., they can self-organize rapidly according to the changing distributions of targets across regions. In a temporally varying environment, an ant colony has to increase its responsiveness to cope with frequent changes in task demands by employing more generalist ants, which perform a range of tasks (Wilson & Yoshimura 1994). Similarly, we will like to examine the effect of our non-stationary task environment, induced by moving robots and targets, on the degree of specialization in the robots. Based on Shannon-Wiener information variable H (Lehner 1979), the third performance index quantifies the degree to which a robot specializes in a region:

$$\text{degree of specialization} = 1 - H$$

$$H = -\sum_r p_r \log_R p_r \quad (11)$$

where p_r is the proportion of time a robot stays in region r for the task duration of T , and R is the total number of regions. This index varies within the range $[0,1]$. A degree of 1 implies the robot specializes in tracking only one region whereas a degree of 0 means the robot spends equal proportion of time tracking in all R regions.

Figure 2c shows the mean degree of specialization of all the robots, which is lower for auction-based negotiation and ant-based MRTA. Hence, we can conclude from Figures 2b and 2c that a larger number of generalist robots leads to a lower coalitional cost. Although auction-based negotiation achieves lower degree of specialization and coalitional cost than ant-based MRTA, its coverage is lower. This is because reducing the degree of specialization will incur more time in




Figure 3: Comparison of proportions of robots within different ranges of degrees of specialization.

task switching and consequently decrease the time for performing the task (Spencer, Couzin, & Franks 1998). In our test, this means that a robot will switch between several regions, thus incurring longer time in travelling between regions and target searching, and spending less time in target tracking. This accounts for poorer coverage of auction-based negotiation.

For ant-based MRTA, the mean degree of specialization is slightly higher with a smaller number of robots because each robot receives fewer messages from the other robots. As a result, the robots are less certain about the task demands in other regions. This causes the robots to be more specialized and less inclined to explore other regions. With cooperative EKMs or potential fields, fewer robots result in higher mean degree of specialization because the robots interfere less with each other and stay longer in a particular region.

Figure 3 shows the proportions of robots within different ranges of degrees of specialization for the case of 10 robots. Using ant-based MRTA and auction-based negotiation for explicit task allocation, most of the robots have degrees of specialization < 0.6 . The other two methods without explicit task allocation have comparatively larger number of robots with degrees of specialization ≥ 0.6 . Hence, the methods with explicit task allocation are less rigid to changes in regional task demands and incur lower coalitional cost.

Summary of Test Results

Compared to the other schemes, ant-based MRTA and auction-based negotiation have lower degree of specialization, coalitional cost, and higher coverage. But the degree of specialization cannot be too low as the cost of generalization would then exceed its benefits. This explains the higher coverage of ant-based MRTA over auction-based negotiation. Also, strategies without explicit task allocation can perform

better than static placement by utilizing robot mobility to track the targets.

Our approach has also been tested on the coverage of evasive targets that avoid the tracking robots. In this test of 15 robots, the coverage of the static sensors has dropped to 10% but our ant-based scheme can still maintain a 53% coverage. The coverage of the other schemes have dropped slightly.

Our scheme is robust to robot failures, which is crucial for operating in dynamic, uncertain environments. For example, in the event that 5 robots fail, our scheme can still outperform a fully operational static sensor network (Fig. 2a).

We have also implemented a deterministic version of ant-based MRTA but it performs slightly worse than the stochastic version. This is because robots endowed with the deterministic scheme will move to a region only if information about the region flows out from it. If this region does not contain any robot, no information about this region can be transmitted to robots in other regions. Thus, no robot will move into this region even if it contains many targets to be tracked. Our stochastic scheme resolves this problem.

Conclusion

This paper describes a task allocation scheme via self-organizing swarm coalitions for distributed mobile sensor network coverage. It uses concepts of ant behavior to self-regulate the division of labor in proportion to the regional task demands in the changing environment. Hence, the sensor network functions like an adaptive multiagent system. Quantitative comparisons with auction-based negotiation, potential fields and static deployment have shown that our approach can provide better coverage. Our work also demonstrates that in a temporally varying task environment, the self-organization of coalitions enables the robots to respond rapidly to frequent changes in regional task demands.

References

Balch, T., and Arkin, R. C. 1998. Behavior-based formation control for multirobot teams. *IEEE Trans. Robot. Automat.* 14(6):926–939.

Bonabeau, E.; Dorigo, M.; and Théraulaz, G., eds. 1999. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford Univ. Press.

Braitenberg, V. 1984. *Vehicles: Experiments in Synthetic Psychology*. MIT Press.

Camazine, S.; Deneubourg, J.-L.; Franks, N. R.; Sneyd, J.; Theraulaz, G.; and Bonabeau, E. 2001. *Self-Organization in Biological Systems*. Princeton Univ. Press.

Gerkey, B. P., and Matarić, M. J. 2002. Sold!: Auction methods for multi-robot coordination. *IEEE Trans. Robot. Automat.* 18(5):758–768.

Goldberg, D., and Matarić, M. J. 1997. Interference as a tool for designing and evaluating multi-robot controllers. In *Proc. AAAI*, 637–642.

Gordon, D. M. 1999. Interaction patterns and task allocation in ant colonies. In Detrain, C.; Deneubourg, J. L.; and Pasteels, J. M., eds., *Information Processing in Social Insects*. Basel, Switzerland: Birkhäuser Verlag. 51–67.

Hirsh, A. E., and Gordon, D. M. 2001. Distributed problem solving in social insects. *Annals Math. Artif. Intell.* 31(1-4):199–221.

Jung, B., and Sukhatme, G. S. 2002. Tracking targets using multiple robots: The effects of environment occlusion. *Autonomous Robots* 13(3):191–205.

Krieger, M. J. B.; Billeter, J.-B.; and Keller, L. 2000. Ant-like task allocation and recruitment in cooperative robots. *Nature* 406:992–995.

Lehner, P. N. 1979. *Handbook of Ecological Methods*. New York: Garland STPM.

Lesser, V.; Ortiz Jr., C. L.; and Tambe, M. 2003. *Distributed Sensor Networks: A Multiagent Perspective*. Kluwer Academic Publishers.

Low, K. H.; Leow, W. K.; and Ang, Jr., M. H. 2002. A hybrid mobile robot architecture with integrated planning and control. In *Proc. AAMAS*, 219–226.

Low, K. H.; Leow, W. K.; and Ang, Jr., M. H. 2003. Action selection for single- and multi-robot tasks using cooperative extended Kohonen maps. In *Proc. IJCAI*, 1505–1506.

Matarić, M. J. 1997. Behaviour-based control: Examples from navigation, learning, and group behaviour. *J. Expt. Theor. Artif. Intell.* 9(2-3):323–336.

O’Rourke, J. 1987. *Art Gallery Theorems and Algorithms*. Oxford Univ. Press.

Parker, L. E. 1998. ALLIANCE: An architecture for fault tolerant multi-robot cooperation. *IEEE Trans. Robot. Automat.* 14(2):220–240.

Parker, L. E. 2002. Distributed algorithms for multi-robot observation of multiple moving targets. *Autonomous Robots* 12(3):231–255.

Sandholm, T.; Larson, K.; Andersson, M.; Shehory, O.; and Tohmé, F. 1999. Coalition structure generation with worst case guarantees. *Artif. Intell.* 111(1-2):209–238.

Schneider-Fontañ, M., and Matarić, M. J. 1998. Territorial multi-robot task division. *IEEE Trans. Robot. Automat.* 14(5):815–822.

Shehory, O., and Kraus, S. 1998. Methods for task allocation via agent coalition formation. *Artif. Intell.* 101(1-2):165–200.

Simmons, R.; Apfelbaum, D.; Burgard, W.; Fox, D.; Moors, M.; Thrun, S.; and Younes, H. 2000. Coordination for multi-robot exploration and mapping. In *Proc. AAAI*, 852–858.

Spencer, A. J.; Couzin, I. D.; and Franks, N. R. 1998. The dynamics of specialization and generalization within biological populations. *J. Complex Systems* 1(1):115–127.

Tofts, C. 1993. Algorithms for task allocation in ants (a study of temporal polyethism: Theory). *Bull. Math. Biol.* 55(5):891–918.

Wilson, D. S., and Yoshimura, J. 1994. On the coexistence of specialists and generalists. *Am. Nat.* 144(4):692–707.

Yadgar, O.; Kraus, S.; and Ortiz Jr., C. L. 2003. Hierarchical information combination in large-scale multiagent resource management. In Huget, M.-P., ed., *Communications in Multiagent Systems*. LNAI 2650, 129–145.