

# Collective Inspection of Regular Structures using a Swarm of Miniature Robots

N. Correll and A. Martinoli

Swarm-Intelligent Systems Group, Nonlinear Systems Laboratory, EPFL, 1015 Lausanne  
nikolaus.correll|alcherio.martinoli@epfl.ch

**Abstract.** We present a series of experiments concerned with the inspection of regular, engineered structures carried out using swarms of five to twenty autonomous, miniature robots, solely endowed with on-board, local sensors. Individual robot controllers are behavior-based and the swarm coordination relies on a fully distributed control algorithm. The resulting collective behavior emerges from a combination of simple robot-to-robot interactions and the underlying environmental template. To estimate intrinsic advantages and limitations of the proposed control solution, we capture its characteristics at higher abstraction levels using non-spatial, microscopic and macroscopic probabilistic models. Although both types of models achieve only qualitatively correct predictions, they help us to shed light on the influence of the environmental template and control design choices on the considered non-spatial swarm metrics (inspection time and redundancy). Modeling results suggest that additional geometric details of the environmental structure should be taken into account for improving prediction accuracy and that the proposed control solution can be further optimized without changing its underlying architecture.

## 1 Introduction

Swarm Intelligence (SI) is an emerging computational and behavioral metaphor for solving distributed problems. SI takes its inspiration from biological examples provided by social insects [4] and by swarming, flocking, herding, and shoaling phenomena in vertebrates [12]. The abilities of such natural systems appear to transcend the abilities of the constituent individual agents, while being mediated by nothing more than a small set of simple local interactions.

Biological systems show a full blend of antagonist mechanisms for coordinating their collective behavior. For instance, an environmental template (e.g., a temperature gradient or a pheromone gradient generated by the termite queen), a sort of centralized source of information, in combination with distributed building activities of insects allow for the construction of extremely sophisticated structures [4]. As a consequence, the resulting overall behavior of the colony is self-organized and extremely robust to noise since it combines *exploitation* (i.e., following the environmental blueprint and pre-programmed individual behavioral rules) with *exploration* (i.e., the resulting individual behavior is heavily influenced by noisy, local perception).

In this paper, we propose a distributed control algorithm that exploits similar SI-based principles in order to achieve robust inspection and/or coverage of a regular

structure and overcomes intrinsic limitations of individual robots in terms of processing power and sensorial capabilities. The environmental template is in this case represented by the engineered structure to be inspected, while robots emulate biological entities interacting among themselves and with the structure for carrying out the inspection task. In scenarios where, for instance, robots cannot be endowed with sophisticated positioning systems or navigation algorithms, the SI-based approach represents a valid alternative to deliberative coverage strategies which usually achieve deterministic completeness only in absence of uncertainty in robot position (see for instance [1,7]). Finally, although a SI-based approach leads to a certain inspection redundancy because of its intrinsically probabilistic swarm coordination, depending on the reliability of the sensors used for detecting the desired features targeted in the inspection, this property might be beneficial if an error-free feature detection cannot be guaranteed [1,8].

In order to evaluate the performance of the swarm as inspection system, we have defined two different swarm metrics: time to complete inspection and inspection redundancy. Although in a swarm robotic system, microscopic interactions are often characterized by an important stochastic component and are mathematically intractable, the performance of a swarm as a whole is statistically predictable. Therefore, following the efforts recently performed on distributed manipulation experiments [2,9,11], we propose microscopic and macroscopic models to understand and evaluate the influence of key system parameters on the desired metrics.

## 2 The Inspection Case Study


As a first challenging case study, we consider the inspection of jet turbine engines. In order to minimize failures, jet turbine engines have to be inspected at regular intervals. This is usually performed visually using borescopes, a process which is time consuming and cost intensive [10]. One possible solution for speeding up and automating the inspection process is to rely on a swarm of autonomous, miniature robots which could be released into the turbine without disassembling it. While this idea is intellectually appealing and could pave the way for other similar applications in coverage/inspection of engineered or natural, regular structures, it involves a series of technical challenges which dramatically limit possible designs of robotic sensors and hence emphasizes a SI-driven approach. For instance, the shielded, complex, narrow structure of a turbine imposes not only strong miniaturization constraints on the design but also prevents the use of any traditional global positioning and communication system. Furthermore, a limited on-board energy budget might prevent computation of a sophisticated deliberative planning strategy and dramatically narrows the sensor and communication range of our robots.

### 2.1 The Physical Setup emulating a Turbine Scenario

In order to validate our distributed control strategy using the real miniature robots *Alice II* [5], we simplify the turbine inspection scenario by unfolding the axis-symmetric geometry of the turbine into a flat representation with the blades as vertical

extrusions. This simplified turbine environment was implemented in a rectangular arena of  $1.10 \times 1m^2$  monitored by an external overhead camera. The Alice miniature robot is endowed with a PIC microcontroller (368 bytes RAM, 8Kb FLASH), has a length of 22mm, and a maximal speed of  $4 \frac{cm}{s}$ . It is endowed with 4 IR modules which can serve as very crude proximity sensors (up to 3cm) and local communication devices (up to 6 cm in range). Its energetic autonomy ranges between 5h and 10h, depending on the actuators and sensors used.

This platform was chosen not only because it allows us to perform and monitor experiments with large number of robots in the lab, but also because Alice robots force us to design control architectures that are simple and robust enough to deal with their limited capabilities.




**Fig. 1.** *Left:* Overview of the robot arena emulating a simplified turbine scenario. *Middle:* The Alice II Micro Robot. *Right:* Simplified setup for the calibration of model parameters.

## 2.2 The Behavior-Based Robot Controller

In this paper, we are not concerned with detection of flaws - which could be achieved, for instance, by a miniature on-board camera - but rather with the individual and group motion in the turbine scenario. For the sake of simplicity, we therefore assume that circumnavigating a blade in its totality is a good emulation of a scanning-for-flaws maneuver.

The behavior of a single robot is determined by a schema-based controller [3] that tightly links the platforms' actions to sensor perception while using as little representational knowledge of the world as possible. For a schema-based controller, behavioral responses are represented by vectors generated by local potential fields and behavioral coordination is achieved by vector addition. Sequencing of behaviors is achieved by a dynamic action-selection mechanism based on internal timers which are in turn set and reset according to schema response. The overall behavior of a robot can be summarized as follows (see Figure 2, left). The robot searches for blades throughout the turbine, combining schemes that drive the robot forward, avoid obstacles, and follow contours. Using exclusively its on-board sensors, the robot can discriminate between the three different types of objects within the arena: teammates, blades, and external walls. Teammates are distinguishable from any other objects by their active, modulated emission. Blades are distinguishable from



**Fig. 2.** *Left:* The high-level behavioral flowchart of the robot controller as a Finite State Machine (FSM). *Right:* The corresponding Probabilistic FSM used in the models, capturing details of interest of the schema-based controller.

external walls because the former have uniform white contours while the latter are characterized by a pattern of black and white stripes.

Teammates and external walls are systematically avoided while blades stimulate a contour-following behavior in the robot. A robot can start circumnavigating a blade at any point of its contour but can leave the blade exclusively at its tip. This allows the robot to bias its blade-to-blade trajectory without using any sophisticated navigation mechanisms, exploiting a specific feature of the environmental pattern. The robot can recognize a blade's tip again via a specific sensorial pattern generated by its on-board proximity sensors.

A further design parameter is introduced in the controller in order to achieve a full inspection of a blade. The blade can be left only if the timeout parameter  $T_{max}$  has expired. The corresponding timer is set when the robot attaches to the blade. If the circumnavigating behavior of the robot were perfect,  $T_{max}$  would perfectly control the number of blade tours performed by a robot before moving to the next blade. However, rounding the blade's tip without losing contact with the blade represents a challenging maneuver for a miniature robots such as the Alice, in particular because of sensor noise, heterogenous lighting conditions, and so on. We have therefore systematically assessed (see Section 3.2) the average probability for a robot to lose contact with a blade while rounding its tip ( $p_l$ ), the mean time needed to partially circumnavigate a blade (random attaching point) and systematically leave the blade at its tip ( $T_{hb}$ ), and the average time to completely circumnavigate a blade ( $T_{fb} = 2T_{hb}$ ). In this paper, we have chosen to set  $T_{max}$  so that  $T_{fb} < T_{max} < T_{hb} + T_{fb}$ . In this case, the robot either leaves the blade at the first encounter with the blade's tip with probability  $p_l$  or continues the circumnavigation with probability  $1 - p_l$  and leaves at the second encounter with the blade's tip.

### 3 Microscopic and Macroscopic Models

The central idea of the probabilistic modeling methodology is to describe the experiment as a series of stochastic events with probabilities computed from the interactions' geometrical properties and systematic experiments with one or two real robots [11] or embodied agents [2]. Consistent with previous publications, we can use the

controller's FSM depicted in Figure 2 as a blueprint to devise the Probabilistic FSM (PFSM or Markov chain) representing an individual agent at the microscopic level or the whole swarm at the macroscopic level. At the microscopic level, a specific state represents the current mode of a certain individual, while a state at the macroscopic level defines the average number of individuals in that same mode. The state granularity can be arbitrarily chosen in order to capture details of the robot's controller and environment which influence the swarm performance metrics (in our case, the time to completion and the number of redundantly inspected blades). Although at first glance the microscopic model might appear unnecessary since the models presented are essentially linear (and therefore no major discrepancies between mean microscopic and macroscopic predictions can arise), numerical difficulties in assessing the end criterion at the macroscopic level for one of the chosen metrics (time to inspection completion) justify this choice. This is a well-known problem in such types of models when they are used to predict metrics based on discrete quantities (i.e., in our case a finite number of blades). The overall PFSM for the system is represented graphically in Figure 2, *right* using two coupled PFSMs, one representing the robot(s) and one representing the shared turbine environment.

### 3.1 Modeling Assumptions

As more extensively detailed in [2,11], the modeling methodology relies on three main assumptions. First, coverage of the arena by the group of robots is uniform and robots' trajectories and objects' positions in the arena do not play a role in the non-spatial metrics of interest. Second, a robot's future state depends only on its present state and how much time it has spent in that state (semi-Markov property). Third, agents change their state autonomously but synchronously to a common clock whose time step has been chosen to capture with sufficient precision all time delays considered in the system as well as changes in the metrics of interest.

### 3.2 Characterization of Models' Parameters

All our models are characterized by two categories of parameters: state-to-state transition probabilities and behavioral delays. In contrast with previous publications, we do not assume any coupling between these two categories of parameters, but we compute and calibrate them based on either systematic measurements (e.g.,  $p_i$ ) or the concept of *encountering rates* [6], keeping a clear separation between *geometric detection probabilities* and *encountering probabilities*. Here, geometric detection probability is the probability that a robot is within the detection area of a certain object. The detection area of an object is determined by its physical size, the sensory configuration and corresponding processing used by the robot to reliably detect it. After defining the contours of the detection area  $A_i$  for a given object  $i$ , we calculate its geometric detection probability  $g_i$  by dividing  $A_i$  by the whole arena area  $A_a$ . We can hence calculate the corresponding encountering probability, i.e. the probability of encountering the object  $i$  per time step, using the corresponding encountering rate  $r_i$  (in  $s^{-1}$ ). The conversion factor from geometric detection probabilities to encountering rates is given by the average robot speed  $v_r$  ( $4 \frac{cm}{s}$ ), its detection width

$w_r$ , and the detection area  $A_s$  of the smallest object in the arena (a robot is a possible type of object). The detection width is defined as twice the maximum detection distance of the smallest robot/object in the arena, measured from center of the robot to the center of the object, given here by  $2R_s = 4.2cm$  with  $A_s = R_s^2\pi$ . Equation 1 shows how to compute the encountering probability for the object  $i$  given the geometric detection probability  $g_i$ :

$$p_i = r_i T = \frac{v_r w_r}{A_s} g_i T \quad (1)$$

with  $T$  being the time step characterizing our time-discrete models. In this paper, we discretize the different average durations of interactions so that changes in our two arbitrary metrics are described with sufficient precision using  $T = 1s$ . Mean numerical values used for the model parameters are summarized in Table 1. Mean detection range  $R_s$ ,  $p_l$ , and  $T_{hb}$  were measured using real robots in a simplified setup (compare Figure 1, right). Encountering probabilities based on the geometric probability of each object were computed by Equation 1.

**Table 1.** Encountering probabilities for the three objects a robot might encounter, their detection area, and their mean interaction time. Each interaction experiment was repeated 40 times.

Object	Computed enc. rate	Detect. area	Mean interaction time
Wall	$p_w = 0.0458$	$420cm^2$	$T_w = 10s$
Blade	$p_b = 0.0106$	$96.48cm^2$	$T_{hb} = 10s, p_l = 0.4$
Robot	$p_r = 0.0015$	$14cm^2$	$T_r = 4s$

The mean blade interaction time  $T_b$  is a function of  $p_l$ ,  $T_{hb}$ , and  $T_{max}$ . With the  $T_{max}$  used in this paper, it can be calculated as follows:  $T_b = p_l T_{hb} + (1 - p_l)(3T_{hb}) = 22s$ .

### 3.3 Mathematical Description of the Macroscopic Model

From Figure 2, *right* we can derive a set of difference Equations (DE) to capture the dynamics of the whole system at the macroscopic level. We formulate one DE per considered state and exploit equations stating the conservation of the number of robots and the number of blades to replace one of the DEs.

Given  $M_0$  blades and  $N_0$  robots, the number of robots covering virgin and inspected blades  $N_v$  and  $N_i$ , the number of robots in obstacle avoidance  $N_a$ , and the number of robots in search mode  $N_s$  are given by Equation 2-5 (compare also Figure 2); the number of virgin blades  $M_v$  and the number of inspected blades  $M_i$  are calculated by Equation 6-7:

$$N_s(k+1) = N_s(k) - \Delta_v(k) - \Delta_i(k) - \Delta_r(k) - \Delta_w(k) + \Delta_v(k - T_b) + \Delta_i(k - T_b) + \Delta_r(k - T_r) + \Delta_w(k - T_w) \quad (2)$$

$$N_a(k+1) = N_a(k) + \Delta_r(k) + \Delta_w(k) - \Delta_r(k - T_r) - \Delta_w(k - T_w) \quad (3)$$

$$N_v(k+1) = N_v(k) + \Delta_v(k) - \Delta_v(k - T_b) \quad (4)$$

$$N_i(k+1) = N_0 - N_s(k+1) - N_a(k+1) - N_v(k+1) \quad (5)$$

$$M_v(k+1) = M_v(k) - \xi_b \Delta_v(k - T_b) \quad (6)$$

$$M_i(k+1) = M_0 - M_v(k+1) \quad (7)$$

with  $k$  representing the current time step (and absolute time  $kT$ );  $k = 0 \dots n$ ,  $n$  being the total number of iterations (and therefore  $nT$  the end time of the experiment). The  $\Delta$ -functions define the coupling between state variables of the model and can be calculated as follows:

$$\Delta_v(k) = p_b(M_v(k) - N_v(k))N_s(k) \quad (8)$$

$$\Delta_i(k) = p_b(M_i(k) + N_v(k))N_s(k) \quad (9)$$

$$\Delta_r(k) = p_r(N_0 - 1)N_s(k) \quad (10)$$

$$\Delta_w(k) = p_w N_s(k) \quad (11)$$

Here,  $p_b$ ,  $p_r$ , and  $p_w$  represent the encountering probabilities of blades, robots, and wall.  $T_b$ ,  $T_r$ ,  $T_w$  define the average time needed for circumnavigating a blade, avoiding a teammate, and avoiding a wall respectively. Due to the probability of leaving a blade before it has been completely covered, we introduce a parameter  $\xi_b$ , being the percentage of a blade a robot covers on average at each new interaction with it. Similarly to  $T_b$ ,  $\xi_b$  can be calculated from  $p_l$ ,  $T_{hb}$ , and  $T_{fb}$ :  $\xi_b = p_l T_{hb} / T_{fb} + (1 - p_l) T_{fb} / T_{fb} = 0.8$ .

The initial conditions are  $N_s(0) = N_0$  and  $N_a(0) = N_v(0) = N_i(0) = 0$  for the robotic system (all robots in search mode) while those of the environmental system are  $M_v(0) = M_0$  and  $M_i(0) = 0$  (all blades virgin). As usual for time-delayed DE, we assume  $\Delta_x(k) = N_x(k) = M_x(k) = 0$  for  $k < 0$ .

For instance, we can interpret the first DE (Equation 2) as follows. The average number of robots in the searching state is decreased by those that start to cover a virgin blade or an inspected blade and those that start avoiding either a teammate or a wall; it is increased by all robots resuming searching after either an inspection or an obstacle avoidance maneuver, each of them being characterized by a specific duration. The other state equations can be interpreted in a similar way.

### 3.4 Swarm Performance Metrics

We consider two metrics for evaluating the performance of the swarm: time to completion  $nT$  and coverage redundancy. The task is completed if all blades are inspected ( $M_v(n) = 0$ ), while the coverage redundancy  $R$  is given by  $R = M_r(n) - M_0$  with

$$M_r(n) = \sum_{k=0}^n \rho_v \Delta_v(k - T_b) + \rho_i \Delta_i(k - T_b) \quad (12)$$

the total number of inspected blades. Here, the  $\rho$ -terms reflect the mean redundant coverage on encountering a virgin or an inspected blade respectively. Hence,  $\rho_i$  and

$\rho_v$  can be computed as weighted sums in terms of coverage (partial circumnavigation with probability  $p_l$  and over-complete circumnavigation with probability  $1 - p_l$ ):  $\rho_i = p_l T_{hb}/T_{fb} + (1 - p_l)(T_{hb} + T_{fb})/T_{fb} = 1.1$  and  $\rho_v = (1 - p_l)T_{hb}/T_{fb} = 0.3$ .

To compute the time to completion  $nT$ ,  $M_v(n) = 0$  is an easy condition to apply in the experiment and in the microscopic model. However, in the macroscopic model, this represents a limit condition as  $\lim_{k \rightarrow \infty} M_v(k) = 0$ . Therefore, we solved the DEs numerically for  $M_r(n) = \mu$ , with  $\mu$  the mean resulting from the microscopic model.


## 4 Results and Discussion

In this section we will present the results collected using real robots and compare them to model predictions with regard to the desired swarm metrics. In contrast to previous experiments in the distributed manipulation class where obstacles were either moving [2] or much smaller [11] than the immovable blades considered in this case study, the swarm performance metrics in the inspection task appear to be more closely influenced by the distribution of the robots, while still being non-spatial. Furthermore, objects in the arena are not all axis-symmetric as in the previous experiments. Finally, the structure created by the blades themselves also plays a relevant role in the prediction accuracy of the models as it biases not only the uniformity of the robot's spatial distribution at a given moment but also its evolution over time. All these features are not taken into account in our current modeling methodology and therefore generate discrepancies between models and experimental results. We will discuss them in more detail in Section 4.2.

### 4.1 Swarm Performance Metrics

We estimated time to completion to cover 16 blades and the total number of covered blades by doing 10'000 runs of the microscopic model for team sizes of 1 to 20 robots and using the parameters summarized in Table 1. To validate model predictions, we ran 20 experiments each for team sizes of 5, 10, 16, and 20 robots. In order to come closer to our assumption of spatial uniformity (see Section 3.1), robots were initially distributed randomly in the arena. Figure 3 shows experimental and modeling results for both swarm performance metrics (time to complete the inspection, number of redundantly inspected blades). We observe that model predictions for the overall task dynamics in both metrics are only qualitatively correct for small teams. For larger team sizes however, quantitative prediction improves. We believe that this is because an individual robot's trajectory does not satisfy our assumption of spatial uniformity. Increasing the team size instead weakens the effect of individual trajectories and increases the quality of prediction. It is worth noting that, in contrast to experiments reported in [6] where robots never left a blade prematurely during a circumnavigation maneuver, in the experiments reported here the inspection redundancy before task completion increases with the number of robots (see Figure 3, *right*).






**Fig. 3.** Comparison of microscopic and macroscopic model predictions (basically superimposed) and experimental results. Experimental results are represented by the median and standard deviation because of long tail distributions while microscopic models by mean and standard deviation. *Left:* Time to completion (16 blades) vs. team size (1 to 20 robots, 5, 10, 16, 20 robots respectively). *Right:* Number of blades redundantly covered until inspection completion vs. team size.

#### 4.2 Influence of Boundary Conditions

As described in Section 2.2, the robots exploit the regularity of the blade pattern to effectively traverse the turbine moving to the next blade only when they reach the tip of the current blade. This behavior in combination with the geometrical structure of the pattern biases the distribution of the robots in space and over time.



**Fig. 4.** Distribution of the robots on the arena (dark areas correspond to frequently visited regions). Each experiment lasted 2h with 20 robots. *Left:* Default setup with the default inspection algorithm. *Center:* Set-up rotated of 180 degrees. *Right:* Rotated setup and robots characterized by a random walk behavior.

In order to investigate the influence of the blade configuration, we performed additional experiments with different configurations, the original setup turned by 180 degrees in order to cancel out external environmental influences (light, arena slope etc.), and the rotated setup but with randomly moving robots, avoiding obstacles. The resulting cumulative distributions for a 2h period and 20 robots are depicted in Figure 4.

In addition to local high density zones where traffic jams occur, we note an overall tendency: although the robots are scattered randomly in the arena at the beginning of the experiments, fairly independently from their behavior, over time they tend to move in direction of the blades's tips and thus visit some areas of the arena more often than others. In Figure 4 *left*, we observe that frequency of coverage by the robots increases from the left to the right and from the right to the left in Figure 4, *middle* and *right*. In the realistic simulation experiments reported in [6], this effect did not appear as it was compensated by a wrap-around zone at the arena borders (mimicking the cylindrical turbine environment). This is no longer the case in the physical setup presented here since the turbine structure is delimited by walls. As a consequence, in a fully enclosed arena, the probability to encounter a blade is dependent upon its level along the main axis of the structure and the elapsed experimental time ( $p_b \rightarrow p_b(x, y, t)$ ), thus violating our assumptions of spatial uniformity (see Section 3.1).

## 5 Conclusion and Future Work

In this paper, we proposed a simple distributed control algorithm to inspect a regular structure with a swarm of miniature robots. We investigated the overall collective behavior of the swarm by generating the corresponding non-spatial models. The results show that the time to inspection completion decreases sublinearly with the number of robots, that the inspection redundancy until task completion increases with the swarm size if robots prematurely leave the blade they are inspecting, and that the swarm's center of mass can actually be moved along the regular structure by exploiting local interactions between the robots and the structure itself.

We have also shown that our modeling methodology provides useful insight in this experiment as has been the case in previous distributed manipulation experiments. In particular, we believe that the control design parameter  $T_{max}$  might be further optimized to obtain more interesting  $T_b$  and  $\xi_b$  pairs for a constant  $p_l$ . Indeed, this latter parameter might be more difficult to control since it defines the interaction of a miniature robot with limited capabilities with the environmental structure. Finally, we also demonstrated that the current modeling methodology reaches its limitations in case studies where the structure of the environment plays a major role on the robot distribution. In order to be able to exploit models as tools for performance prediction and design, we need to incorporate more geometric details of the environmental template. We believe that this will enable the design of more customized algorithms as a function of the regular engineering structure to be inspected.

### Acknowledgements

This work has been initiated at the California Institute of Technology, primarily supported by the NASA Glenn Center and in part by the Caltech Center for Neuromorphic Systems Engineering under the US NSF Cooperative Agreement ERC-9402726. Both authors are currently sponsored by a Swiss NSF grant (contract Nr.

PP002-68647/1). We would like to thank M. Asadpour and G. Caprari for their crucial support while working with the Alice robots.

## References

1. E. Acar, H. Choset, Y. Zhang, and M. Schervish. Path planning for robotic demining: Robust sensor-based coverage of unstructured environments and probabilistic methods. *Int. Journal of Robotics Research*, 22(7-8), 2003.
2. W. Agassounon, A. Martinoli, and K. Easton. Macroscopic modeling of aggregation experiments using embodied agents in teams of constant and time-varying sizes. *Autonomous Robots*, 17(2-3):163–192, 2004. Special issue on Swarm Robotics, M. Dorigo and E. Sahin, editors.
3. R. Arkin. *Behavior-Based Robotics*. The MIT press, Cambridge, MA, USA, 2000.
4. E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. SFI Studies in the Science of Complexity, Oxford University Press, New York, NY, USA, 1999.
5. G. Caprari and R. Siegwart. Design and control of the mobile micro robot alice. In *Proceedings of the 2nd International Symposium on Autonomous Minirobots for Research and Edutainment*. Brisbane, Australia, pages 23-32, 2003.
6. N. Correll and A. Martinoli. Modeling and optimization of a swarm-intelligent inspection system. In *Proceedings of the 7th Symposium on Distributed Autonomous Robotic System (DARS)*. Toulouse. Distributed Autonomous Systems 6. To Appear, 2004.
7. Y. Gabriely and E. Rimon. Spiral-stc: An on-line coverage algorithm of grid environments by a mobile robot. In *Proceedings of the 2002 IEEE Conference on Robotics & Automation*. Washington, DC, USA, pages 954-960, 2001.
8. Douglas Gage. Many-robot mcm search systems. In A. Bottoms, J. Eagle, and H. Bayless, editors, *Proc. of the Autonomous Vehicles in Mine Contermeasure Symposium*. Monterey, CA, USA, pages 9.55-9.63, 1995.
9. K. Lerman and A. Galstyan. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, 2(13):127–141, 2002.
10. K. Martin and C.V. Stewart. Real time tracking of borescope tip pose. *Image and Vision Computing*, 10(18):795–804, July 2000.
11. A. Martinoli, K. Easton, and W. Agassounon. Modeling of swarm robotic systems: A case study in collaborative distributed manipulation. *Int. Journal of Robotics Research*, 23(4):415–436, 2004. Special issue on Experimental Robotics, P. Dario and B. Siciliano, editors. Invited paper.
12. J. K. Parrish and W. M. Hammer. *Animal Groups in Three Dimensions*. Cambridge University Press, UK, 1997.