Covering a Continuous Domain by Distributed, Limited Robots

Eliyahu Osherovich, Alfred M. Bruckstein, and Vladimir Yanovski

Technion - Israel Institute of Technology, Computer Science Department, Haifa, Israel
{oeli, freddy, volodyan}@cs.technion.ac.il

Abstract. We present an algorithm for covering continuous domains by primitive robots whose only ability is to mark visited places with pheromone and to sense the level of the pheromone in their neighborhood. These pheromone marks can be sensed by all robots and thus provide a way for indirect communication between the robots. Apart from this, the robots have no means to communicate. Additionally they are memoryless, have no global information such as the domain map, own position, coverage percentage, etc. Despite the robots' simplicity, we show that they are able to cover efficiently any connected domains, including non-planar ones.

1 Introduction

We say that a domain is *covered* by a robot if each and every point of the domain was swept by the robot's effector. In fact, every time we want to build an automatic machine suitable for applications such as floor cleaning, snow removal, lawn mowing, painting, mine-field de-mining, unknown terrain exploration and so forth, we face the problem of complete covering of corresponding domains by our machine.

A particular solution of the covering problems depends, of course, on the capabilities of our robots and various environmental constraints. Hence a vast number of algorithms can be, and actually have been, developed to accommodate the numerous constraints of the covering problem.

In this paper we adopt the model used in [1], which assumes that our robots are anonymous, i.e., any two robots are the same, memoryless, i.e., they have no ability to "remember" anything from the past and have no means of direct communication. This model was originally inspired by ants and other insects that use chemicals called *pheromones* that are left on the ground and used for some kind of indirect communication and coordination tasks. Ant colonies, despite primitivism of single ants, demonstrate surprisingly good results in global problem solving and pattern formation [2,3,4,5,6]. Consequently, some ideas borrowed from these insects are becoming increasingly popular in ant-robotics and distributed systems [5,6,7,8,9,10]. Such robots are usually capable of performing quite complex distributed tasks while providing the benefits of being small, cheap, easy to produce and easy to maintain.

M. Dorigo et al. (Eds.): ANTS 2006, LNCS 4150, pp. 144–155, 2006.
© Springer-Verlag Berlin Heidelberg 2006

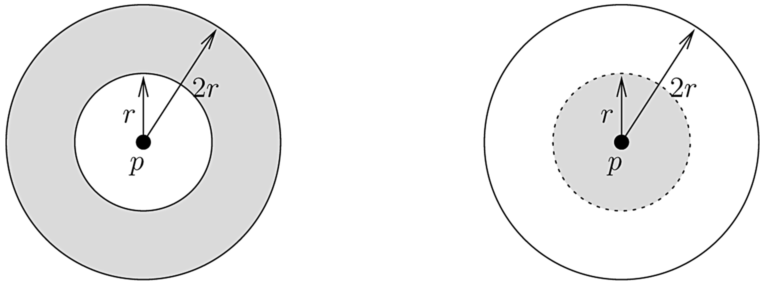

(a) Sensing area **(b)** Marking area

Fig. 1. Robot's sensing and marking areas

2 Agent Model

Mathematical formulation of the problem is as follows. The domain will be denoted by Ω. At the moment we consider only flat two-dimensional domains; further extensions will be given in Section 6.4. Given any two points $a, b \in \Omega$ we denote the *distance* between a and b as $\|a - b\|$. Again, we assume, initially, that the distance is the common Euclidean distance in two-dimensional space; extensions to other distance measures will be given in Section 6.4. The robot is able to sense the pheromone level at its current position p and in a closed ring of radii r and $2r$ around p denoted by $R(r, 2r, p)$. Additionally, our robot is able to set an arbitrary pheromone level in an open disk of radius r around its current location p denoted by $D(r, p)$, We assume that our time steps are discrete and denote by $\sigma(a, t)$ the pheromone level of point $a \in \Omega$ at time instance t.

3 The Mark-Ant-Walk (MAW) Algorithm

Initially, no point is marked with the pheromone and thus all σ values are assumed to be equal to zero: $\sigma(a, 0) = 0;\ \forall a \in \Omega$. A starting point is chosen (randomly) for the robot and then the MAW step rule is applied repeatedly. There is no explicit stopping condition for this algorithm; nevertheless, one can use the upper bound, provided later in this paper, on the cover time in order to stop robots after a sufficient time period that guarantees complete covering.

Table 1. MAW step rule

Mark-Ant-Walk step rule (current time is t and agent location is p)	
(A)	$x :=$ a point from $R(r, 2r, p)$ with *minimal* value of $\sigma(x, t)$
	/* *In case of a tie - make an arbitrary decision* */
(B)	If $\sigma(p) \leq \sigma(x) : \forall u \in D(r, p) \quad \sigma(u) = \sigma(x) + 1$
	/* *we mark open disk of radius r around current location* */
(C)	$t := t + 1$
(D)	move to x

4 Related Work

Covering of discrete domains (graphs) is an old problem and thus it has a number of solutions with a sound mathematical background. Probably, the most known examples are the Breadth-First Search (BFS) and the Depth-First Search (DFS) algorithms for graph traversal. Both algorithms provide excellent results in terms of time complexity.

A step toward an odor-oriented model was taken in [11,12] where *pebbles* were used to assist the search. Pebbles are tokens that can be placed on the ground and later removed. The idea of pebbles was further developed in [13] where they were used for unknown graph exploration and mapping. Two different algorithms that fit our paradigm entirely, i.e. fully distributed autonomous agents that mark the ground with pheromones, were suggested for efficient and robust graph covering. One, called the Edge-Ant-Walk, marks the graph edges [14]. Another one, called the Vertex-Ant-Walk, leaves marks on graph vertices instead [1,15]. Both algorithms provided significant improvement over DFS in robustness terms along with quite efficient cover time.

Random walks are defined for both discrete and continuous domains and provide unrivaled robustness and scalability; however, they cannot guarantee complete coverage, providing only expected time. We would like to concentrate on solutions that can guarantee complete coverage after a limited time period.

One possible approach is to introduce an *artificial potential field* in order to accomplish the robot motion planning task (e.g. [16,17]). This approach can easily be adopted by our robots where the potential is represented by the odor level. However, it assumes that the potential field is constructed prior to the start of robot motion and thus requires a global knowledge of the domain boundaries and obstacles, which is unavailable in our model. Some authors used trails that mark the path travelled by the agent so far and performed some kind of peeling/milling. This approach often fails with non-convex domains and thus the whole domain may be approximated as a union of convex non-overlapping cells [18,19,20], ,however, this approach, in fact, takes us back to a graph whose vertices are associated with the cells and edges between vertices that are defined according to the corresponding inter-cell connectivity. Another representative of trail-based algorithms is the Mark-And-Cover (MAC) algorithm [21], which is actually an adaptation of the DFS to continuous domains. This algorithm provides efficient and effective coverage with excellent provable cover time. Additionally, the agent model used in the paper fit our paradigm entirely. Nevertheless, the problem of the MAC algorithm, and probably all trail-based algorithms, is their sensitivity to noise and agents failure. Moreover, trails of one agent may hamper performance of another agent. Another shortcoming of these algorithms is seen in the situation when the domain is required to be covered repeatedly, e.g., in surveillance tasks or in the scenario described in [22] where autonomous agents are used to de-mine minefields using imperfect sensors, i.e. the probability of a mine detection is less than 1. Our algorithm guarantees that the whole domain is covered repeatedly time after time. Furthermore, the time between two successive visits at any point is bounded (see Section 6.1).

5 MAW - Formal Proof of Correctness and Upper Time Bound

Let us show that a single robot governed by the MAW rule covers any connected bounded domain in a finite number of steps. The outline of the proof is as follows.

First, we prove that at any time instance, any two points that are close enough, i.e., their distance from each other is less than or equal to r, must have pheromone levels that differ by one at most. We call this *the proximity principle.* It has also been used in several other research studies, e.g., [1,15,14].

Second, we look at the diameter d of the domain that is defined as the length of the longest geodesic line embedded in the domain, i.e., $d = \sup_{a,b \in \Omega} \|a - b\|$. Assuming that d is finite, we easily conclude with the aid of the proximity principle that at any time t for any two points $a, b \in \Omega$, the difference between the pheromone levels of these two points is limited by $\lceil d/r \rceil$. This, in turn, means that once the value of $\lceil d/r \rceil + 1$ is reached at any point, no unmarked point remains and thus the whole domain has been covered. Finally, we show that we eventually reach value of $\lceil d/r \rceil + 1$. A formal proof is given below.

Lemma 1
The difference between marker values of close points is bounded.

$$\forall t; \forall a, b \in \Omega \, : \, \textit{if} \, \|a - b\| \leq r \, \textit{ then } \, |\sigma(a, t) - \sigma(b, t)| \leq 1$$

Proof: We shall prove the lemma by mathematical induction on the step number. The lemma is clearly true at $t = 0$. Assuming it is also true at time $t = n$, we shall show it remains true at time $t = n + 1$. Let us look at two points $a, b \in \Omega$, such that $\|a - b\| \leq r$. In the trivial case neither a nor b changes its marker value at the $(n + 1)$th step; therefore, the lemma holds according to the induction hypothesis. If both a and b change their values, then $\sigma(a, t + 1) = \sigma(b, t + 1)$ since the algorithm assigns the same values to all the points it changes. Hence the only interesting case is when only one point (say a) changes its marker value Assuming the current agent's location is p_t we conclude that $a \in D(r, p_t)$, otherwise it could not change its marker value. And therefore, $\|a - p_t\| < r$. b, however, does not change its marker value and thus $\|b - p_t\| \geq r$. Combining these constraints we get $r \leq \|b - p_t\| \leq 2r$ or, equivalently, $b \in R(r, 2r, p_t)$. Now let us recall how the new marker value of a is determined. First, we look for the minimal marker value among all points in $R(r, 2r, p_t)$. Assume that this value is attained at some point $x \in R(r, 2r, p_t)$. The new marker value of a is then set if and only if $\sigma(p_t, t) \leq \sigma(x, t)$:

$$\sigma(a, t + 1) = \sigma(x, t) + 1. \tag{1}$$

Since both points x and b belong to $R(r, 2r, p_t)$, we have

$$\sigma(b, t) \geq \sigma(x, t), \tag{2}$$

because of the way the point x was chosen. Now, on the one hand, we have:

$$\begin{cases} |\sigma(a, t) - \sigma(b, t)| \leq 1 \\ \sigma(b, t) \geq \sigma(x, t) \end{cases} \Rightarrow \sigma(a, t) \geq \sigma(x, t) - 1; \tag{3}$$

and on the other hand:

$$\begin{cases} |\sigma(a,t) - \sigma(p_t,t)| \leq 1 \\ \sigma(p_t,t) \geq \sigma(x,t) \end{cases} \Rightarrow \sigma(a,t) \leq \sigma(x,t) + 1. \quad (4)$$

Combining inequalities (3) and (4), we get

$$|\sigma(a,t) - \sigma(x,t)| \leq 1. \quad (5)$$

Using the system of inequalities (3), we conclude that

$$0 \leq \sigma(b,t) - \sigma(x,t) \leq 2. \quad (6)$$

Combining the above inequality with the fact that $\sigma(a,t+1) = \sigma(x,t) + 1$ and $\sigma(b,t+1) = \sigma(b,t)$, we get the desired result: $|\sigma(a,t+1) - \sigma(b,t+1)| \leq 1$. Thus the lemma is proven. ∎

Lemma 2
The difference between marker values of any two points is bounded at all times.

$$\forall t; \forall a,b \in \Omega : |\sigma(a,t) - \sigma(b,t)| \leq \left\lceil \frac{d}{r} \right\rceil$$

where d - diameter of Ω.

Proof: Follows immediately from Lemma 1. ∎

Our next step will be to show that the maximal marker value tends to ∞ as t goes to ∞. First, we prove that marker values can only grow and never decrease.

Lemma 3
Marker values of any point form a non-decreasing series; that is

$$\forall t; \forall u \in \Omega : \sigma(u,t+1) \geq \sigma(u,t).$$

Proof: Let us assume the contrary, i.e., there exists a point $u \in \Omega$ and time instance t such that the pheromone level of u decreases during the t-th step: $\sigma(u,t+1) < \sigma(u,t)$. Let us now look at point p_t – the location of the agent at time t. Obviously $u \in D(r,p_t)$ (otherwise it could not change its value), hence $\|u - p_t\| < r$. Assume that the minimal marker value among all points in $R(r,2r,p_t)$ was attained at some point x. We know also that $\sigma(p_t,t) \leq \sigma(x,t)$; otherwise, the robot does not change the pheromone values. Thus we have

$$\begin{cases} \sigma(x,t) + 1 = \sigma(u,t+1) < \sigma(u,t) \\ \sigma(p_t,t) \leq \sigma(x,t) \\ \|u - p_t\| < r \end{cases} \quad (7)$$

This implies

$$\begin{cases} |\sigma(u,t) - \sigma(p_t,t)| > 1 \\ \|u - p_t\| < r \end{cases} \quad (8)$$

which contradicts Lemma 1. ∎

At this point we are ready to prove the main result of this work.

Theorem 1
The domain Ω will be covered within a finite number of steps.

PROOF: Imagine that the domain Ω is tessellated into n cells so that every such cell can be inscribed into a circle of diameter less than r. Let us examine the following sum:

$$S_t = \sum_{i=1}^{n} m_t^i - \sigma(p_t, t), \tag{9}$$

where m_t^i is the minimal marker value over the ith cell at time t and $\sigma(p_t, t)$ is the marker value at the agent's location p_t at time instance t. With the aid of Lemma 3 one can easily verify that

$$S_{t+1} > S_t. \tag{10}$$

Given that $S_0 = 0$, we easily conclude that

$$S_t \geq t \Rightarrow \sum_{i=1}^{n} m_t^i \geq t \quad \forall t, \tag{11}$$

which leads us to the conclusion that after $n\lceil \frac{d}{r} \rceil + 1$ steps, at least one of the m_{nd+1}^i values will be greater than $\lceil \frac{d}{r} \rceil$ and thus the whole domain will be covered. ∎

In order to find an approximation to n , we can tile the domain with regular hexagons of side length $r/2$. In order to guarantee full coverage by the hexagons we look at the "augmented" domain $\bar{\Omega}$, which results from Ω that has undergone morphological dilation with a disk of radius r. Using a development similar to the one shown in [21], we get the following bound on the area of $\bar{\Omega}$

$$A_{\bar{\Omega}} \leq A_\Omega + rP_\Omega + \pi r^2, \tag{12}$$

where A_Ω and P_Ω are the area and the perimeter of Ω, respectively. Thus we have

$$n \leq \frac{A_\Omega + rP_\Omega + \pi r^2}{\frac{3\sqrt{3}}{8} r^2}, \tag{13}$$

where $\frac{3\sqrt{3}}{8} r^2$ represents the area of a hexagon of side length $r/2$.

6 Extensions

6.1 Repetitive Coverage

In some scenarios we might be interested in repetitive coverage of the domain, e.g., the aforementioned scenario of minefield de-mining with imperfect sensors [22] or tasks such as surveillance and patrolling. In all cases we would like to bound the time between two successive visits.

Lemma 4
For any two time instances t_1 and t_2, if only $t_2 > t_1$ then the following inequality must hold:

$$S_{t_2} - S_{t_1} \geq t_2 - t_1.$$

PROOF: The proof is very simple. We can always write $t_2 = t_1 + n$ for some natural n and prove the lemma by mathematical induction. For $n = 1$ the lemma holds due to Equation (10). Assuming that the lemma holds for some n, we can easily conclude that the lemma holds for $n + 1$ as well. ∎

Theorem 2
For any point $a \in \Omega$, the time period between two successive visits of the robot is bounded by $2n\left(\lceil \frac{d}{r} \rceil + 1\right)$.

PROOF: If we show that after a sufficient time period the pheromone level changes at all locations in the domain Ω, we can obviously be sure that all points were re-visited by the robot during this time period. Let us look at time instance t_s when the robot covers our point of interest a. We denote by $\sigma_{max}(t_s)$ the maximal pheromone level over Ω at that time. If we show that at some time instance t_e the minimal pheromone level denoted by $\sigma_{min}(t_e)$ becomes greater than the maximal value that was at time t_s: $\sigma_{min}(t_e) > \sigma_{max}(t_s)$, then we can easily conclude that during the time period $t_e - t_s$ the pheromone level changed at all points and thus all points (including a) were re-covered by the robot. Let us examine S_{t_s} and S_{t_e} as defined in the Equation (9). On the one hand:

$$S_{t_s} = \sum_{i}^{n} m_{t_s}^i - \sigma(p_{t_i}, t_i) \geq \sum_{i}^{n} m_{t_s}^i \geq \sum_{i}^{n} \sigma_{min}(t_s) = n\,\sigma_{min}(t_s) \tag{14}$$

According to Lemma 2

$$\sigma_{min}(t_s) \geq \sigma_{max}(t_s) - \left\lceil \frac{d}{r} \right\rceil. \tag{15}$$

Combining Equations (14) and (15) we get

$$S_{t_s} \geq n\left(\sigma_{max}(t_s) - \left\lceil \frac{d}{r} \right\rceil\right). \tag{16}$$

On the other hand we want to know the time instance t_e that guarantees that $\sigma_{min}(t_e) \geq \sigma_{max}(t_s) + 1$. Instead of estimating t_e directly from $\sigma_{min}(t_e)$, we shall look for t_e that guarantees the existence of σ value greater than or equal to $\sigma_{max}(t_s)+1+\lceil \frac{d}{r}+1 \rceil$, which guarantees by Lemma 2 that $\sigma_{min}(t_e) \geq \sigma_{max}(t_s)+1$. Now, in the same way as the proof of Theorem 1, we can say that once $S_{t_e} \geq n(\sigma_{max}(t_s) + 1 + \lceil \frac{d}{r} \rceil + 1)$, we have $\sigma_{min}(t_e) \geq \sigma_{max}(t_s) + 1$. Thus we have

$$S_{t_e} - S_{t_s} \leq n\left(\sigma_{max}(t_s) + 1 + \left\lceil \frac{d}{r} \right\rceil + 1\right) - \left(\sigma_{max}(t_s) - \left\lceil \frac{d}{r} \right\rceil\right) = 2n\left(\left\lceil \frac{d}{r} \right\rceil + 1\right). \tag{17}$$

According to Lemma 4 we have

$$t_e - t_s \leq S_e - S_s \leq 2n\left(\left\lceil \frac{d}{r} \right\rceil + 1\right), \tag{18}$$

which completes the proof. ∎

6.2 Noise Immunity

Until now we always assumed that there is no noise in the input, i.e., the robot starts with a domain that does not contain any pheromone marks. Unfortunately, in the real life such a clear environment is not always available hence, we shall consider situation when the initial pheromone level is not zero. Unlike trail-based algorithms that cannot cope with noise our algorithm, can easily overcome this problem as demonstrated by the experiments in Section 7.3.

6.3 Multiple Robots

As a natural extension we would like to analyze how the MAW algorithm can be applied to multi-robot environments. First of all, we must address problems such as collisions both between the robots themselves (if we deal with physical robots and not programs) and between different pheromone levels when two (or more) robots try to mark the same point in the domain.

At the moment we assume that the clock phases of all robots are slightly different so that no two robots are active at the same time. Thus each robot sees other robots as regular stationary obstacles and acts accordingly. This approach also resolves the problem of different pheromone levels that might be assigned to the same point by different robots, since only one robot is active at given time.

Let us find the upper bound for complete coverage provided we have k robots. Using the same notation as in Equation (9) we have:

$$S_t = \sum_{i=1}^{n} m_t^i - \sum_{j=1}^{k} \sigma(p_t^j, t), \tag{19}$$

where p_t^j denotes the location of the j-th robot at time t. Using exactly the same reasoning as before, we again obtain:

$$S_t \geq t, \tag{20}$$

which leads us to the same upper bound we got for a single robot. Hence adding more robots does not necessarily guarantees better coverage time. However our simulations (see Section 7) demonstrate that there is a substantial improvement when we use more robots.

6.4 Using Other Metrics

Until now we always used the usual notion of the distance, nevertheless, it is easy to verify that all the proofs remain valid if we change the Euclidean (L_2) distance to another one. For example, we used L_∞ in our simulations. Since corresponding effector shape is a square in this case.

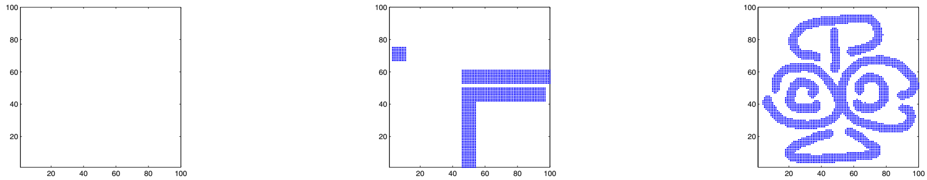

(a) Domain A **(b)** Domain B **(c)** Domain C

Fig. 2. Simulation domains

7 Simulations and Experiments

7.1 General Notes

Experiments were conducted on the domains shown in Figure 2. All domains are of size 100×100 pixels and marking radius in all experiments was set to 3, i.e., each step robot marks a square of 5×5 pixels. Figure 3 demonstrates some stages of covering Domain B by ten robots.

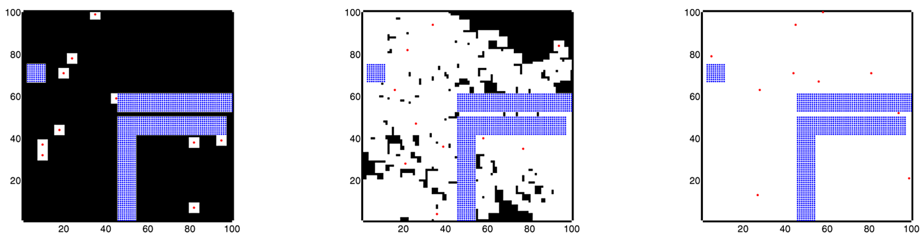

(a) Coverage: step 1 **(b)** Coverage: step 50 **(c)** Coverage: step 125

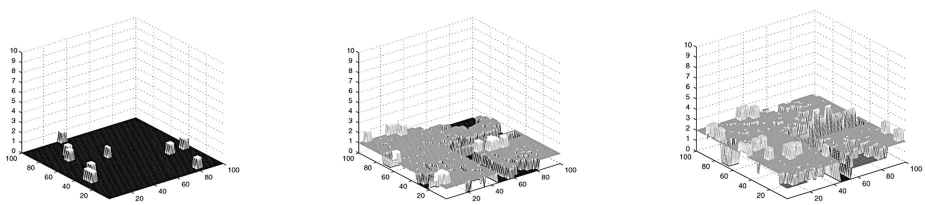

(d) Odor map: step 1 **(e)** Odor map: step 50 **(f)** Odor map: step 125

Fig. 3. MAW progress on Domain B

7.2 Comparing MAW to Other Algorithms

In this experiment we studied performance of three different algorithm: MAW, MAC [21], and Random Walk. All algorithms used the same square effector

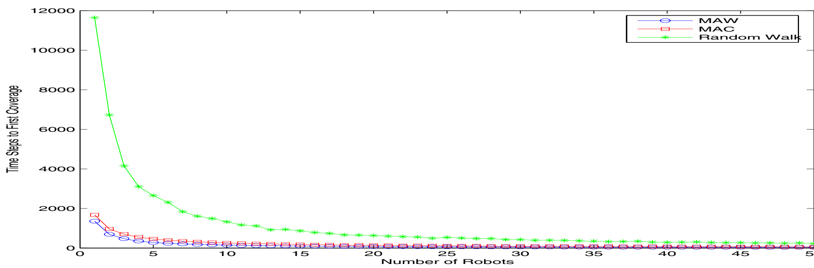

(a) Cover Time: Domain A

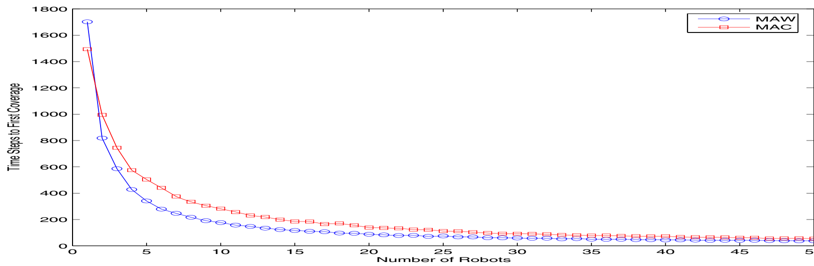

(b) Cover Time: Domain B

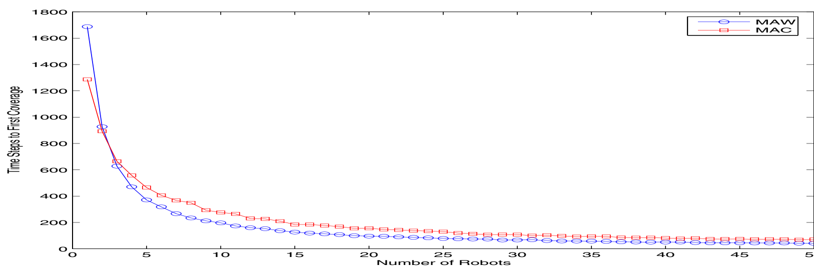

(c) Cover Time: Domain C

Fig. 4. Cover Time

of size 5×5 pixels; additionally, the steps of the Random Walk algorithm were restricted to be in interval $[r, 2r]$ just like the steps in the MAW algorithm.

In all experiments the robots were modeled as points and multiple robots were allowed to occupy the same location. We always measured the number of time steps until the robots covered the domain for the first time, averaged over 100 runs.

As we can see the MAW algorithm is a clear winner when we use three or more robots. For fewer robots the MAC algorithm performs better on complex domains. Note that the MAW algorithm in general performs better than the theoretical upper bound we got in Section 5. Cover time of the Random Walk was omitted from Figures 4b and 4c because the values were so big that the difference between the MAC and the MAW algorithms became invisible on this scale.

7.3 MAW in Noisy Environments

In this experiment we ran one robot on the Domain A, each time changing the amount of noisy pixels. Noise values are uniformly distributed in interval $[1, 10]$. Figure 5 shows cover time as a function of the amount of noisy pixels.

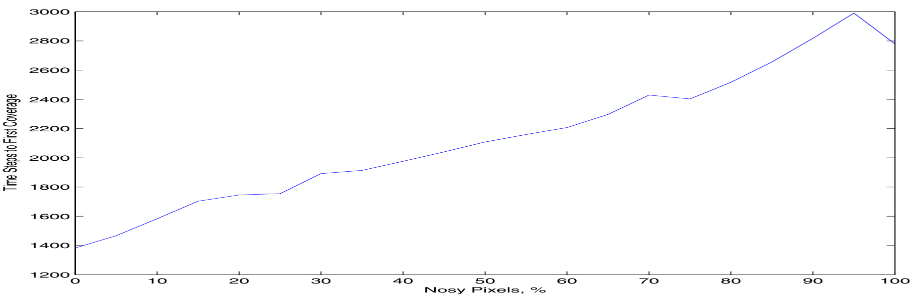

Fig. 5. Noisy environment

Note that noise does not affect the Random Walk on the one hand and it destroys completely the MAC algorithm on the other hand, making it unable to cover the domain completely.

8 Conclusions

In this paper we presented a new ant-inspired algorithm for continuous domain covering. We provided also a formal proof of complete coverage and upper time bounds for complete coverage and the time interval between two successive visits of the robot. Additionally a formal proof provided for multi-robot environments. Algorithm performance and noise immunity were verified by computer simulations.

References

1. Yanovski, V., Wagner, I.A., Bruckstein, A.M.: Vertex-ant-walk - A robust method for efficient exploration of faulty graphs. Annals of Mathematics and Artificial Intelligence **31**(1-4) (2001) 99–112
2. Bruckstein, A.M.: Why the ant trails look so straight and nice. The Mathematical Intelligencer **15**(2) (1993) 58–62
3. Hölldobler, B., Wilson, E.O.: The Ants! Harvard University Press (1990)
4. Schöne, H.: Spatial orientation : the spatial control of behavior in animals and man. Princeton University Press, Princeton, N.J. (1984)
5. Dorigo, M., Maniezzo, V., Colorni, A.: Ant system: Optimization by a colony of cooperating agents. IEEE Trans. on Systems, Man, and Cybernetics–Part B **26**(1) (1996) 29–41
6. Dorigo, M., Di Caro, G., Gambardella, L.M.: Ant algorithms for discrete optimization. Artificial Life **5**(2) (1999) 137–172

7. Wagner, I.A., Bruckstein, A.M.: From ants to a(ge)nts: A special issue on ant-robotics. Annals of Mathematics and Artificial Intelligence **31**(1-4) (2001) 1–5
8. Bonabeau, E., Théraulaz, G.: Swarm smarts. Scientific American **282**(3) (2000) 72–79
9. Russell, R.A.: Ant trails - an example for robots to follow? In: ICRA. (1999) 2698
10. Koenig, S., Liu, Y.: Terrain coverage with ant robots: a simulation study. In: AGENTS '01: Proceedings of the fifth international conference on Autonomous agents, New York, NY, USA, ACM Press (2001) 600–607
11. Blum, M., Sakoda, W.: On the capability of finite automata in 2 and 3 dimensional space. In: Ann. Symp. on Foundations in Computer Science. (1977) 147–161
12. Blum, M., Kozen, D.: On the power of the compass. In: Proc. 19th Ann. Symp. on Foundations in Computer Science. (1978) 132–142
13. Bender, M.A., Fernández, A., Ron, D., Sahai, A., Vadhan, S.: The power of a pebble: exploring and mapping directed graphs. In: STOC '98: Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing, New York, NY, USA, ACM Press (1998) 269–278
14. Wagner, I.A., Lindenbaum, M., Bruckstein, A.M.: Smell as a computational resource — A lesson we can learn from the ant. In: Proceedings of the 4th Israel Symposium on Theory of Computing and Systems, ISTCS'96 (Jerusalem, Israel, June 10-12, 1996), Los Alamitos-Washington-Brussels-Tokyo, IEEE Computer Society Press (1996) 219–230
15. Wagner, I.A., Lindenbaum, M., Bruckstein, A.M.: Efficiently searching a graph by a smell-oriented vertex process. Annals of Mathematics and Artificial Intelligence **24**(1-4) (1998) 211–223
16. Khatib, O.: Real-time obstacle avoidance for manipulators and mobile robots. The International Journal of Robotics Research **5**(1) (1986) 90–98
17. Zelinsky, A., Byrne, J.C., Jarvis, R.A.: Planning paths of complete coverage of an unstructured environment by a mobile robot. In: International Conference on Advanced Robotics (ICAR). (1993)
18. Choset, H., Pignon, P.: Coverage path planning: The boustrophedon decomposition. In: International Conference on Field and Service Robotics. (1997)
19. Butler, Z.J.: Distributed coverage of rectilinear environments. PhD thesis, Carnegie Mellon University (2000)
20. Acar, E.U., Choset, H., Zhang, Y., Schervish, M.J.: Path planning for robotic demining: Robust sensor-based coverage of unstructured environments and probabilistic methods. I. J. Robotic Res **22**(7-8) (2003) 441–466
21. Wagner, Lindenbaum, Bruckstein: MAC vs. PC: Determinism and randomness as complementary approaches to robotic exploration of continuous unknown domains. ROBRES: The International Journal of Robotics Research **19** (2000)
22. Gage, D.W.: Randomized search strategies with imperfect sensors. In Chun, W.H., Wolfe, W.J., eds.: Proc. SPIE. Volume 2058. (1994) 270–279