
Eliyahu Osherovich
Vladimir Yanovski

The Computer Science Department,
Technion-Israel Institute of Technology,
Haifa 32000,
Israel
oeli@cs.technion.ac.il

Israel A. Wagner

IBM Haifa Research Lab,
Haifa, 31905,
Israel

Alfred M. Bruckstein

The Computer Science Department,
Technion-Israel Institute of Technology,
Haifa 32000,
Israel

Robust and Efficient Covering of Unknown Continuous Domains with Simple, Ant-Like A(ge)nts

Abstract

We propose a new 'Mark-Ant-Walk' algorithm for robust and efficient covering of continuous domains by ant-like robots with very limited capabilities. The robots can mark places visited with pheromone marks and sense the level of the pheromone in their local neighborhood. In the case of multiple robots, these pheromone marks can be sensed by all robots and provide the only way of (indirect) communication between the robots. The robots are assumed to be memoryless, and to have no information besides that mentioned above. Despite the robots' simplicity we show that they are able, by running a very simple rule of behavior, to ensure efficient covering of arbitrary connected domains, including non-planar and multidimensional domains. The novelty of our algorithm lies in the fact that, unlike previously proposed methods, our algorithm works on continuous domains without relying on some induced underlying graph effectively reducing the problem to a discrete case of graph covering. In addition, we demonstrate various benefits of our algorithm such as ability to cope with arbitrary initial pheromone profile and a bounded constant time between two successive visits of the robot at the same location, making it suitable for patrolling. Finally, we provide a new theoretical bound on covering time of a wide family of such mark and walk algorithms.

KEY WORDS—ant robotics, covering, domain coverage, swarm intelligence, distributed robots

The International Journal of Robotics Research
Vol. 27, No. 7, July 2008, pp. 815–831
DOI: 10.1177/0278364908092465
©SAGE Publications 2008 Los Angeles, London, New Delhi and Singapore
Figures 4, 8, 9, 11–14 appear in color online: <http://ijr.sagepub.com>

1. Introduction

Suppose we want to cover (or clean or search or paint) a connected domain in \mathbb{R}^2 using one or more simple robots having effectors (or arms) that can sweep a well-defined neighborhood of the robots when they are stationary. We say that a domain was *covered* by the (team of) robots if each and every point of the domain was swept by a robot 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 on, we face the problem of complete covering of continuous domains by such devices.

1.1. Problem Constraints

The approach to solving covering problems depends, of course, on the capabilities we assume for our robots and on various environmental constraints. Many algorithms can be, and actually have been, developed to accommodate constraints and assumptions on the robots used for the covering problem. The various considerations include:

1. domain type (e.g. discrete versus continuous, simply-connected or multiply-connected region, a general graph or a grid, etc.);
2. robot capabilities (their communication capabilities, the amount of on-board memory, the size of their footprint

and the areas swept by the effectors and range of robots' sensors);

3. the type of information the robots are assumed to be able to get or gather via their sensors either globally or locally, off-line or on-line;
4. local behavior and interaction model in case of multiple robots (such as synchronous or asynchronous operation, centralized or distributed control, etc.).

In this paper we adopt the model proposed in Yanovski et al. (2001), which assumes that the robots are anonymous (i.e. all robots are identical), memoryless (i.e. have no ability to 'remember' anything from the past) and have no means of direct communication which means there is no direct exchange of information between the robots. The only means of (indirect) communication is via some marks they leave in their environment. This model was originally inspired by ants and other insects using chemicals called *pheromones* that are left on the ground as a means of achieving indirect communication and coordination. (Note that recent publications show that ants are not memoryless, e.g. Harris et al. (2005). However, it is certainly true that for certain ant species these individual capabilities play a limited role in navigation and trail-laying or trail-following mechanisms.)

Ant colonies, despite the simplicity of single ants, demonstrate surprisingly good results in global problem solving and pattern formation (Schöne 1984; Hölldobler and Wilson 1990; Bruckstein 1993; Dorigo et al. 1996, 1999). Consequently, ideas borrowed from insects (and especially ants) behavior research are becoming increasingly popular in robotics and distributed systems (Dorigo et al. 1996, 1999; Russell 1999; Bonabeau and Théraulaz 2000; Koenig and Liu 2001; Wagner and Bruckstein 2001). Simple robots were found capable of performing quite complex distributed tasks while enjoying benefits of being small, cheap, easy to produce and easy to maintain.

This paper is organized as follows. Our formal robot model is presented in Section 2.1. In Section 2.2 we define the Mark-And-Walk (MAW) covering algorithm, followed by a short survey of previously proposed covering algorithms in Section 3. As mentioned earlier the number of such algorithms is fairly large; we therefore limited our survey to those that share some common principles with the algorithm suggested in this paper. Formal proofs of complete coverage and efficiency analysis are given in Sections 4.1 and 4.2, respectively. In Section 5 we provide various extensions of the basic MAW algorithm including its multi-agent or swarm-system performance and its behavior when the environment contains some initial pheromone profile (false pheromone marks). Results of simulations are presented in Section 6. Section 8 provides a summary of our results and a discussion of possible extensions and implementation details.




Fig. 1. Robot's sensing (dark gray) and marking (light gray) areas.

2. The Basic Mark-Ant-Walk (MAW) Covering Algorithm

2.1. Robot Model

We define the mathematical problem of robot covering along with the model for the robots that we use throughout this paper.

The domain to be covered is denoted Ω . We currently only consider two-dimensional (2D) domains (however, extensions will be given in Section 5.4). Given any two points $a, b \in \Omega$, the *geodesic distance* between a and b is denoted $\|a - b\|$, i.e. the length of the shortest path restricted to lie entirely in the domain Ω that connects a and b . For the sake of brevity we shall omit the word 'geodesic' and simply use 'distance'. At the moment, we assume that this length is measured as a common Euclidean length in 2D space; extensions to other measures are discussed in Section 5.4. We shall then assume that the robot is able to sense the pheromone level at its current position p and in a closed 'geodesic' ring $R(r, 2r, p)$ lying between the internal radius r and the external radius $2r$ around p (Figure 1). R is formally defined as

$$R(r_1, r_2, c) = \left\{ a \in \Omega \mid r_1 \leq \|a - c\| \leq r_2 \right\}, \quad (1)$$

and r is considered to be an intrinsic parameter of the robot.

Additionally, our robot is able to set a constant arbitrary pheromone level in the area swept by its effector which is, we assume, an open disk $D(r, p)$ of radius r around its current location p . The formal definition of D is:

$$D(r, c) = \left\{ a \in \Omega \mid \|a - c\| < r \right\}. \quad (2)$$

Note that, as mentioned before, all distances are measured as geodesic distances. Hence, for example, only area \mathcal{A} in Figure 2 will be available to the robot. Area \mathcal{B} , on the other hand, will not be 'visible' to the robot since the distance from the robot location p to any point in \mathcal{B} is greater than $2r$.

Furthermore, we shall assume that our time steps are discrete. We denote the pheromone level of point $a \in \Omega$ at time instance t ($t = 0, 1, 2, \dots$) by $\sigma(a, t)$.




Fig. 2. Geodesic distance example.

Mark-Ant-Walk step rule (current time is t and robot location is p)	
(1)	Find $x :=$ a point in $R(r, 2r, p)$ with <i>minimal</i> value of $\sigma(x, t)$ (In case of a tie, i.e. when the minimal value is achieved at several places, make an arbitrary decision) <i>/* note that $\ p - x\ \geq r$ */</i>
(2)	If $\sigma(p) \leq \sigma(x)$ then $\forall u \in D(r, p)$ set $\sigma(u) = \sigma(x) + 1$ <i>/* we mark open disk of radius r around current location */</i>
(3)	$t := t + 1$
(4)	move to x

Fig. 3. MAW step rule.

2.2. The Mark-Ant-Walk (MAW) Algorithm

Initially, we consider the case where no point is marked with the pheromone, thus all σ values are assumed to be equal to zero:

$$\sigma(a, 0) = 0; \forall a \in \Omega. \tag{3}$$

Suppose further that a starting point is (randomly) chosen for a (single) robot and then a MAW step rule (as described in Figure 3) is applied repeatedly. Note that there is no explicit stopping condition for this algorithm. Nevertheless, one can use an upper bound (that will be provided later in this paper) on the cover time in order to stop robots after a sufficient time period has elapsed, effectively guaranteeing complete covering.

It is easy to see that the robot markings create some kind of *potential field*, where high values of the pheromone level

roughly indicate areas that have been visited many times and lower values of the pheromone level correspond to a smaller number of robot visits. According to the MAW rule, the robot tries to avoid areas with high pheromone values by moving toward lower levels, striving to reach areas yet uncovered.

3. Related Work

A first step toward a ‘pheromone-marking’-oriented model was taken by Blum and Sakoda (1977) and Blum and Kozen (1978) where pebbles were used to assist the search. Pebbles are tokens that can be placed on the ground and later removed. The idea of using pebbles for unknown graph exploration and mapping was further developed in Bender et al. (1998).

3.1. Discrete Domain Covering Problems

The covering of discrete domains (graphs) is an important problem in theoretical computer science and thus a number of solutions have been proposed and comprehensively analyzed by researchers. Well-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. Formal proofs of complete coverage and efficiency analysis can be found in texts on discrete mathematics and algorithms (e.g. Rivest and Leiserson 1990). Note that the DFS algorithm can be readily adapted to fit our robot model as opposed to the BFS algorithm, which requires additional on-board memory in order to maintain a queue of already discovered but yet unvisited vertices. Moreover, the amount of on-board memory required depends on the graph to be explored.

Two further algorithms that fit our paradigm entirely, i.e. they rely on a group of identical autonomous robots that mark the ground with pheromones, were suggested for efficient and robust graph covering. The Edge-Ant-Walk marks the graph edges (Wagner et al. 1996). The Vertex-Ant-Walk leaves marks on graph vertices instead (Yanovski et al. 2001; Wagner et al. 1998). Both algorithms provided significant improvement over DFS in terms of robustness along with quite efficient cover times. Both are capable of completing the traversal of the graph with multi-robot teams even in the case when almost all robots die and/or the environment graph changes (edges/vertices are added or deleted) during the execution.

Different research related to our subject work was published in de Silva and Ghrist (2006) where the authors deal with coverage in coordinate-free sensor networks via homology.

3.2. Continuous Domain Covering Problems

Covering continuous domains is a relatively new problem. Several algorithms addressing this problem are summarized in a good survey paper by Choset (2001). We shall provide a short review of algorithms developed for robot models that are close to that used in this paper.

3.2.1. Random Walk and Probabilistic Covering

Random walks were defined for both discrete and continuous domains and enjoy unrivaled robustness and scalability (i.e. adding more robots does not require any change in the algorithm and can only improve the coverage time). However, they cannot guarantee complete coverage and can only be analyzed in terms of expected coverage time. Although some researchers have demonstrated a way to estimate this expected coverage time, i.e. in terms of electrical resistance of the domain (Chandra et al. 1989; Wagner et al. 2000), in this work we would like to have solutions that guarantee complete coverage after a deterministic and bounded time period.

3.2.2. Motion Planning Guided by a Potential Field

A popular approach is to introduce an *artificial potential field* concept in order to accomplish the robot motion planning task (e.g. Khatib 1986; Zelinsky et al. 1993). In Zelinsky et al. (1993), for example, the authors used the distance transform as the potential field. Such an approach could also be adopted by our robots, the potential being represented by the odor level. However, this type of work assumes that the potential field can be constructed prior to the start of the robot's motion and thus requires a global knowledge of the domain boundaries and obstacles. Such knowledge is not available in our model.

3.2.3. Trail-laying Algorithms for Continuous Domains

Several authors proposed the use of trails that mark the path traveled by the robot so far, and proposed local behavior for the robots resulting in some kind of peeling/milling of the domain. Two major approaches for organizing the motion exist: contour-parallel and direction-parallel. In the former approach, the robot moves along the boundaries of the domain. In the latter approach, the robot moves back and forth in some pre-defined direction. These approaches often fail for non-convex domains and thus the whole domain must be approximated as a union of convex non-overlapping cells which, in turn, can be covered with the assumed type of motion (Choset and Pignon 1997; Butler 2000; Acar et al. 2003).

Another representative of trail-based algorithms is the Mark-And-Cover (MAC) algorithm presented in Wagner et al. (2000), which is actually an adaptation of the DFS algorithm to continuous domains. This algorithm provides efficient and effective coverage with a provable bound on cover time. Additionally, the robot model used in the paper fits our paradigm very well. Nevertheless, the drawback of the MAC algorithm and, in fact, of all trail-based algorithms, is their sensitivity to false trails and robot failure. Moreover, in the multi-robot case trails of the robots interfere with the motion of the others and may hamper their efficiency. Another shortcoming of these algorithms is seen in the situation where the domain needs to be covered repeatedly, e.g. in surveillance tasks or in the scenario described in Gage (1994) where autonomous robots are used to de-mine minefields using imperfect sensors (in the sense that 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 successive visits at any point can be bounded in terms of the (unknown) size of the problem (see Section 5.1).

3.2.4. Tessellating Algorithms for Continuous Domains

Another possible approach is to split the domain into tiles such that each tile is easily covered by the robots (e.g. convex tiles

are suitable for ‘onion peeling’ or ‘back and forth’ algorithms). After a particular tile is covered, the robot goes to a new one that is a neighbor of the current tile (Choset and Pignon 1997; Koenig and Liu 2001; Koenig et al. 2001). This approach, in fact, takes us back to a graph covering of an underlying graph whose vertices are associated with the tiles, the edges between vertices being defined according to the inter-tile connectivity.

4. The MAW Algorithm: Formal Proof and Efficiency Analysis

4.1. MAW: Proof of Complete Coverage

Let us first 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 any two points are close enough at any time, i.e. their distance from each other is $\leq r$ and they must have pheromone levels that differ by one at most. This property closely resembles the Lipschitz continuity for functions $f(\cdot)$:

$$|f(p_1) - f(p_2)| \leq \kappa \|p_1 - p_2\|, \quad (4)$$

for some constant κ . However, our σ function measuring the level of the pheromone at each location is, by definition, not continuous and thus cannot be Lipschitz continuous. Instead, σ obeys the following inequality:

$$|\sigma(p_1) - \sigma(p_2)| \leq \kappa \|p_1 - p_2\| + 1. \quad (5)$$

We call this the *proximity principle*, and it has also been used in several previous papers (Wagner et al. 1996, 1998; Yanovski et al. 2001).

Second, we look at the diameter d of the domain, defined as the length of the longest geodesic path that can be embedded in the domain: consider a pair of points $a, b \in \Omega$. Since Ω is connected, there is at least one path that connects a and b and lies entirely in the domain Ω . Among all such paths (connecting a and b , that are restricted to lie entirely in the domain Ω), there is at least one that is the shortest. We call the length of this shortest path the *distance* between a and b and denote it by $\|a - b\|$. Among all possible pairs of points $a, b \in \Omega$, there exists a pair (a_0, b_0) that has the greatest possible distance between them, i.e.

$$\|a_0 - b_0\| \geq \|a - b\| \quad \forall a, b \in \Omega. \quad (6)$$

We call the length of this longest geodesic path the *diameter* of the domain and denote it by d . In other words:

$$d \triangleq \sup_{(a,b) \in \Omega} \|a - b\|. \quad (7)$$

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 upper bounded by $\lceil d/r \rceil$. This in turn means that once a value of $\lceil d/r \rceil + 1$ is reached by $\sigma(\cdot)$ at any point, no unmarked point may exist and thus the whole domain has been covered by the robot. Finally, we show that the maximal pheromone marker value goes to infinity as time goes to infinity; hence we shall surely, at some time, reach the value of $\lceil d/r \rceil + 1$. Formal proofs are given below.

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

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

Proof. Note that the distance $\|a - b\|$ is the length of the shortest path between these two points, which is restricted to lie entirely in the domain Ω . We shall prove the lemma by mathematical induction on the step number. The lemma is clearly true at $t = 0$ when all the marks are assumed to be zero. Assuming it is also true at time t , we shall show it remains true at time $t + 1$. Let us look at two arbitrary points $a, b \in \Omega$, such that $\|a - b\| \leq r$. In the trivial case neither a nor b change their marker values at the $(n + 1)$ th step; therefore, the lemma continues to hold. 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, while the other one remains unchanged. Assuming the current robot’s location is p_t we conclude that a belongs to $D(r, p_t)$, otherwise it could not change its marker value. Therefore, $\|a - p_t\| < r$. However, b does not change its marker value and thus $\|b - p_t\| \geq r$. Combining these constraints, we obtain:

$$\begin{cases} \|a - b\| \leq r \\ \|a - p_t\| \leq r \Rightarrow r \leq \|b - p_t\| \leq 2r, \\ \|b - p_t\| \geq r \end{cases} \quad (8)$$

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 the pheromone level at the current robot’s location is smaller than or equal to that at x : $\sigma(p_t, t) \leq \sigma(x, t)$. In this case we have:

$$\sigma(a, t + 1) = \sigma(x, t) + 1. \quad (9)$$

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

$$\sigma(b, t) \geq \sigma(x, t), \quad (10)$$




Fig. 4. The difference between marker values of any two points is bounded.

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

$$\begin{cases} |\sigma(a, t) - \sigma(b, t)| \leq 1 \\ \sigma(b, t) \geq \sigma(x, t) \end{cases} \Rightarrow \sigma(a, t) \geq \sigma(x, t) - 1; \quad (11)$$

and on the other hand:

$$\begin{cases} |\sigma(a, t) - \sigma(p_t, t)| \leq 1 \\ \sigma(p_t, t) \leq \sigma(x, t) \end{cases} \Rightarrow \sigma(a, t) \leq \sigma(x, t) + 1. \quad (12)$$

Combining inequalities (11) and (12), we obtain

$$|\sigma(a, t) - \sigma(x, t)| \leq 1. \quad (13)$$

Using the system of inequalities (11), we conclude that

$$0 \leq \sigma(b, t) - \sigma(x, t) \leq 2. \quad (14)$$

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 derive the desired result:

$$|\sigma(a, t+1) - \sigma(b, t+1)| \leq 1. \quad (15)$$

Thus the lemma is proved.

Lemma 2 *The difference between marker values of any two points is bounded at any time instance, i.e.*

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

where d is the diameter of Ω .

Proof. Let us consider a path connecting the points a and b . We can always split the path into sub-paths of length r , as depicted in Figure 4. According to Lemma 1 the difference between the marker values at the endpoints of every such sub-path is limited by 1; we can therefore conclude that

$$|\sigma(a, t) - \sigma(b, t)| \leq \left\lceil \frac{l}{r} \right\rceil, \quad (16)$$

where l represents the length of the path. Obviously, among all paths connecting a and b the shortest path will provide best upper bound on the difference between the pheromone levels at a and at b . Since the longest geodesic path in Ω is limited by d we obtain the desired result:

$$|\sigma(a, t) - \sigma(b, t)| \leq \left\lceil \frac{d}{r} \right\rceil. \quad (17)$$

It is implied by Lemma 2 that the difference between marker values is bounded. Our next step will be to show that the maximal marker value tends to infinity as t goes to infinity. First, we prove that marker values can only grow and never decrease.

Lemma 3 *Pheromone level values at any point form a non-decreasing sequence, i.e.*

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

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, i.e.

$$\sigma(u, t + 1) < \sigma(u, t). \quad (18)$$

Let us now look at point p_t , the location of the robot at time t . Obviously u belongs to $D(r, p_t)$ (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

$$\begin{cases} \sigma(x, t) + 1 = \sigma(u, t + 1) < \sigma(u, t) \\ \sigma(p_t, t) \leq \sigma(x, t) \\ \|u - p_t\| < r. \end{cases} \quad (19)$$

This implies

$$\begin{cases} |\sigma(u, t) - \sigma(p_t, t)| > 1 \\ \|u - p_t\| < r, \end{cases} \quad (20)$$

which contradicts Lemma 1.

Next, we show that the maximal value of the pheromone level grows together with the step number t .

Lemma 4 *At any time instance t maximal pheromone level in Ω is bounded from below by t/n for some constant n .*

Proof. Imagine that the domain Ω is tessellated into n cells so that the diameter of every such cell is less than r (for convex cells we can, alternatively, require that the diameter of the circumscribing circle be less than r). Let us examine the following expression:

$$S_t = \sum_{i=1}^n m_t^i - \sigma(p_t, t), \quad (21)$$

where m_t^i is the minimal marker value over the i th cell at time t and $\sigma(p_t, t)$ is the marker value at the robot's location p_t at time instance t . It was shown in Lemma 3 that marker values of any point inside Ω form a non-decreasing sequence. Hence we claim that

$$S_{t+1} > S_t. \quad (22)$$

Indeed, for the non-marking step, the sum of the minima does not change: $\sum_{i=1}^n m_t^i = \sum_{i=1}^n m_{t+1}^i$. However, $\sigma(p_t, t) > \sigma(p_{t+1}, t + 1)$ and therefore $S_{t+1} > S_t$. For the marking step, assuming that the robot goes from cell k into cell l , we have $\sigma(p_t, t) \leq \sigma(p_{t+1}, t + 1)$ and therefore $m_t^k \leq \sigma(p_{t+1}, t + 1)$. Additionally, the whole cell k was marked during this step and thus $m_{t+1}^k = \sigma(p_{t+1}, t + 1) + 1$. Hence we have

$$\begin{cases} m_t^k - \sigma(p_t, t) \leq 0 \\ m_{t+1}^k - \sigma(p_{t+1}, t + 1) = 1. \end{cases} \quad (23)$$

Since the sum of the other minima cannot decrease as was shown in Lemma 3, we conclude again that $S_{t+1} > S_t$. Given that $S_0 = 0$, we readily conclude that

$$S_t \geq t \Rightarrow \sum_{i=1}^n m_t^i \geq t \quad \forall t, \quad (24)$$

which leads us to the conclusion that there exists $k \in 1, 2, \dots, n$ such that

$$m_t^k \geq \frac{t}{n}. \quad (25)$$

Hence the lemma is proved.

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. According to Lemma 4, after $n\lceil d/r \rceil + 1$ steps at least one of the m^i values will be greater than $\lceil d/r \rceil$ and thus the whole domain will be covered.

4.2. MAW: Efficiency Analysis

As proved in Theorem 1 the domain Ω will be covered by a single robot after $n\lceil d/r \rceil + 1$ steps where d is the diameter of the domain, r is the covering radius of the robot's effector and n is the number of cells in some tessellation of Ω (see Lemma 4 for the definition of n). We shall now analyze how good this upper bound is. In order to make such a comparison we shall find an approximation to n to find out what is the best upper bound possible.

Let us denote by N_r the minimal number of steps required by the robot to cover the whole domain Ω . Here we only assume that the robot can cover an open disk of radius r at every step. However, no assumption is made regarding the algorithm governing the robot's behavior implying, among other things, that the robot is allowed to make steps of arbitrary length.

Clearly $N_r \geq A_\Omega/a_r$, where A_Ω and a_r are the areas of the domain Ω and the robot's effector, respectively. Obviously, this is an ultimate lower bound and no algorithm can beat it. However, this bound fails to be tight enough for domains whose shape factor (ratio of the squared domain perimeter to its area multiplied by 4π) is far from 1 (not round). For example, domains comprising a finite number of curves and line segments would have zero area yielding a lower bound of zero, which is of course meaningless. A possible solution is to use the area of the 'augmented' domain $\bar{\Omega}$ which results from an inflation or expansion Ω in all directions by r , i.e. Ω has undergone a morphological dilation with a disk of radius r (Wagner et al. 2000). In this case we have:

$$A_{\bar{\Omega}} \leq A_\Omega + rP_\Omega + \pi r^2, \quad (26)$$

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

We suggest another way of performance assessment which is correct by construction and does not depend on geometric properties such as area or perimeter. Consider the best possible algorithm that covers an open disk of radius $r/2$ in each step. Note that here we consider any algorithm with the only limitation being the coverage radius; we do not limit the algorithm in any other way. Assume that the domain Ω can be covered by such an algorithm, i.e. there exists a finite sequence of points P_1, P_2, \dots, P_N of robots' locations that results in complete coverage, that is:

$$\Omega \subset \bigcup_{i=1}^N D\left(\frac{r}{2}, P_i\right). \tag{27}$$

Alternatively, we can say that for every point a in Ω there exists a number k ($1 \leq k \leq N$) such that $a \in D(r/2, P_k)$. If we consider the best possible algorithm, denoting its coverage time by $N_{r/2}$ as before, we easily conclude that the upper bound on the number of steps required by MAW algorithm is:

$$N_{r/2} \left\lceil \frac{d}{r} \right\rceil + 1. \tag{28}$$

Indeed, if we consider a particular coverage path or the robot described by the sequence of its successive positions: $P_1, P_2, \dots, P_{N_{r/2}}$ we can always perform Voronoi tessellation around these points. Each cell in this tessellation will have diameter smaller than r and thus this tessellation will be like the one we used in Lemma 4.

We have only to estimate the d/r fraction. Obviously, $d/r \leq N_{r/2}$ on one hand and on the other hand we can estimate the lower bound $d/r \geq \pi/4 \sqrt{N_{r/2}}$ (for domains that have a shape close to a circle). Hence we have shown that the upper bound time is polynomial with respect to best possible time of algorithm whose covering radius is $r/2$, i.e.

$$t_{\text{coverage}} \leq N_{r/2}^x, \tag{29}$$

where $0.5 \leq x \leq 2$.

The main question here is whether we can conclude that

$$t_{\text{coverage}} \leq N_r^x, \tag{30}$$

i.e. whether our coverage time is bounded by a polynomial function of N_r , which is the best possible coverage time among all algorithms whose covering radius is r . The general answer is 'no'. A general theorem regarding limitations of such a type of algorithms is as follows.

Theorem 2 *Given an algorithm whose marking area is an open disk of radius r and step size is $\geq r$, the time for complete coverage t_{coverage} of a continuous domain is (tightly) limited from below by $N_{r/2}$. Moreover, t_{coverage} cannot be expressed as a bounded function of N_x for any $x > r/2$.*




Fig. 5. A 'pathological' star-shaped domain.

Proof. The proof is by example of such a domain. We consider a domain comprising n line segments emanating from common origin O , as shown in Figure 5. Assuming that the length of each line segment is $r/2$ we easily verify that $N_{r/2} = n$ and that $t_{\text{coverage}} = n$ (assuming that the robot's initial position was an end point of any line segment). Hence t_{coverage} is (tightly) bounded from below by $N_{r/2}$. If we assume that n is infinite we easily conclude that such a domain cannot be covered in a finite number of steps by our algorithm (again, provided that the initial robot location was at an end point of any line segment), while this domain can be covered in two steps by an optimal algorithm whose covering radius is $> r/2$. We simply go to the origin O in the very first step and the whole domain will be covered in the next step. Hence, the theorem is proved.

The above theorem is of course quite general and, for a particular domain, one can have

$$N_r \propto N_{r/2}. \tag{31}$$

The best possible time would therefore be linear in terms of N_r and MAW's coverage time would be limited from above by N_r^2 multiplied by some constant.

Let us elaborate more about the relationship between N_r and $N_{r/2}$. Consider a domain of area A and perimeter P . Consider also an optimal coverage with radius r . According to our definitions, this coverage requires exactly N_r steps. Let us now look at the Voronoi tessellation around corresponding robots' locations; there are N_r cells in this tessellation. Cells that do not include the boundary of the domain are convex, those that do include domain boundaries may not be convex. Each convex cell in this tessellation can be covered by a finite (and well-defined) number of disks of radius $r/2$; for these cells we can therefore conclude that $N_r \propto N_{r/2}$. For non-convex cells this

claim may not be correct; for domains whose tessellation consists mainly of convex cells we therefore have approximately $N_r \propto N_{r/2}$. Note that a similar analysis was carried out on the MAC algorithm (Wagner et al. 2000) where the authors claim that MAC algorithm is asymptotically linear for domains that obey $A \gg Pr$. This is, in fact, equivalent to saying that the number of cells in the Voronoi tessellation described above which do not contain the domain boundary is large compared to the number of the cells that do.

5. Extensions

5.1. Repetitive Coverage

In some scenarios we might be interested in repetitive coverage of the domain. For example, repetitive coverage is necessary in the aforementioned scenario when robots perform minefield de-mining and their mine detection is not perfect, i.e. the probability of detecting a mine when the robot's sensors sweep above it is less than one. Repetitive coverage is therefore required to minimize the probability of leaving any mines undetected. In this case, we have to give an upper bound on the time between two successive visits of the robot in order to guarantee an improvement in detection probability. This requirement also arises naturally in tasks such as surveillance or patrolling where robots are required to visit every point over and over and the time between two successive visits must be limited by a constant. We shall now show that our algorithm has the property of patrolling. Before we can provide this result, we need the following lemma.

Lemma 5 *For any two time instances t_1 and t_2 , if $t_2 > t_1$ then the following inequality must hold:*

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

Proof. Let us write $t_2 = t_1 + n$ for some natural n and prove the lemma by mathematical induction on n . For $n = 1$, the lemma holds due to Equation (22). Assuming that the lemma holds for some n , we can easily conclude that the lemma also holds for $n + 1$.

Theorem 3 *For any point $a \in \Omega$, the time period between two successive visits of the robot is bounded from above by*

$$2n \left(\left\lceil \frac{d}{r} \right\rceil + 1 \right).$$

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 at time t_s , i.e.

$$\sigma_{\min}(t_e) > \sigma_{\max}(t_s), \quad (32)$$

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 Equation (21). On the one hand:

$$\begin{aligned} S_{t_s} &= \sum_i^n m_{t_s}^i - \sigma(p_{t_s}, t_s) \geq \sum_i^n m_{t_s}^i - \sigma_{\max}(t_s) \\ &\geq \sum_i^n \sigma_{\min}(t_s) - \sigma_{\max}(t_s) = n \sigma_{\min}(t_s) \\ &\quad - \sigma_{\max}(t_s) \end{aligned} \quad (33)$$

and on the other hand:

$$\begin{aligned} S_{t_e} &= \sum_i^n m_{t_e}^i - \sigma(p_{t_e}, t_e) \leq \sum_i^n m_{t_e}^i - \sigma_{\min}(t_e) \\ &\leq \sum_i^n \sigma_{\max}(t_e) - \sigma_{\min}(t_e) = n \sigma_{\max}(t_e) \\ &\quad - \sigma_{\min}(t_e). \end{aligned} \quad (34)$$

According to Lemma 5,

$$\begin{aligned} t_e - t_s &\leq S_{t_e} - S_{t_s} \\ &\leq n \sigma_{\max}(t_e) - \sigma_{\min}(t_e) - n \sigma_{\min}(t_s) + \sigma_{\max}(t_s) \\ &= n (\sigma_{\max}(t_e) - \sigma_{\min}(t_s)) + (\sigma_{\max}(t_s) \\ &\quad - \sigma_{\min}(t_e)). \end{aligned} \quad (35)$$

From Equation (32), we obtain

$$t_e - t_s < n (\sigma_{\max}(t_e) - \sigma_{\min}(t_s)). \quad (36)$$

According to Lemma 2,

$$\sigma_{\min}(t) \geq \sigma_{\max}(t) - \left\lceil \frac{d}{r} \right\rceil. \quad (37)$$

Therefore, from Equation (36) we obtain

$$\begin{aligned} t_e - t_s &\leq n (\sigma_{\max}(t_e) - \sigma_{\min}(t_s)) \\ &\leq n \left(\left(\sigma_{\min}(t_e) + \left\lceil \frac{d}{r} \right\rceil \right) + \left(\left\lceil \frac{d}{r} \right\rceil - \sigma_{\max}(t_s) \right) \right) \\ &= n \left(2 \left\lceil \frac{d}{r} \right\rceil + \sigma_{\min}(t_e) - \sigma_{\max}(t_s) \right). \end{aligned} \quad (38)$$




Fig. 6. Repetitive coverage illustration: (a) pheromone level along some path in Ω and (b) the worst case that guarantees that the whole domain was re-covered.

Since t_e is the first time instance such that $\sigma_{\min}(t_e) > \sigma_{\max}(t_s)$ and the pheromone levels can change by at most 2, we obtain

$$\sigma_{\min}(t_e) - \sigma_{\max}(t_s) \leq 2. \tag{39}$$

Substituting the above inequality into Equation (38) yields

$$t_e - t_s \leq S_e - S_s \leq 2n \left(\left\lceil \frac{d}{r} \right\rceil + 1 \right), \tag{40}$$

which completes the proof.

We now make two important observations. First, the upper time bound between two successive visits by the robot does not depend on the current pheromone level distribution. This is determined completely by the geometric parameters of the problem: r , n and d . Second, we observe that this time limit is twice as long as the time period needed for complete coverage. This situation is quite intuitive. Indeed, observe the pheromone level along some path in Ω , such as that shown in Figure 6a. In this case, the robot may start by ‘filling’ the hollow area on the right until it becomes a hill and only then covers the left-hand part, which used to be a summit point and has now become the lowest point in the pheromone level profile as shown in Figure 6b.

5.2. Initial Pheromone Profile

So far, we have always assumed that the robot starts with a domain that does not contain any pheromone marks. Unfortunately, such a clean environment is not always available in

reality. For example, spurious pheromone marks may arise as a result of previous attempts to explore the domain by similar robots, which might have used other algorithms and thus the initial pheromone level distribution does not necessarily obey the proximity principle. In general, we assume that the initial pheromone level distribution is given by some function

$$N : \Omega \mapsto \mathbb{Z}^*, \tag{41}$$

where \mathbb{Z}^* denotes the set of non-negative integers. As a short digression we have to note that such initial pheromone marks pose a severe problem to all trail-based algorithms. The reason is that such algorithms, for the sake of efficiency, do not get close to their own trails and thus any initial pheromone marks would be interpreted as trails, resulting in uncovered areas around such marks. The result may be even worse if such false trails split the domain into several disconnected parts, in this case the robot will not be able to exit the part where it was located initially. Our algorithm, on the contrary, can easily overcome this problem as we prove below. Actual covering times in the presence of an arbitrary initial pheromone profile are demonstrated in Section 6.3. Let us start with several lemmas.

Lemma 6 *Immediately after point $a \in \Omega$ has been marked by the robot for any point $b \in \Omega$ such that $\|a - b\| \leq r$, we have*

$$\sigma(b, t + 1) \geq \sigma(a, t + 1) - 1, \tag{42}$$

where t denotes the time instance when the new pheromone level was assigned to a .

Proof. Since the pheromone value of a changes during the t th step we conclude that $a \in D(r, p_t)$, where p_t denotes the robot’s location at time t . We also know that $\|a - b\| \leq r$ and thus b either belongs to $D(r, p_t)$ or to $R(r, 2r, p_t)$. Hence there are two possible scenarios: either b belongs to $D(r, p_t)$ or b belongs to $R(r, 2r, p_t)$. In the former case, $\sigma(b, t + 1) = \sigma(a, t + 1)$ since the algorithm assigns the same value to all points in $D(r, p_t)$ and the lemma clearly holds. In the latter case, $b \in R(r, 2r, p_t)$, we recall that the algorithm seeks for the minimal pheromone value inside $R(r, 2r, p_t)$, say attained at some point x and sets a new pheromone level inside $D(r, p_t)$ to be equal to $\sigma(x, t) + 1$. Hence we have

$$\begin{cases} \sigma(b, t + 1) = \sigma(b, t) \geq \sigma(x, t) \\ \sigma(a, t + 1) = \sigma(x, t) + 1. \end{cases} \tag{43}$$

This leads us again to the conclusion

$$\sigma(b, t + 1) \geq \sigma(a, t + 1) - 1. \tag{44}$$

Hence the lemma is proved.

Although this lemma resembles Lemma 1, it does not guarantee that the proximity principle is obeyed in the case of an environment with arbitrary initial pheromone level. We demonstrate a stronger result in Lemma 8. We shall first prove that the pheromone level at marked points, i.e. pheromone left by the robot, never decreases.

Lemma 7 *Pheromone level values at any marked point form a non-decreasing sequence; that is*

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

given that u was marked by the robot at the time prior to t .

Proof. Let us assume the contrary, i.e. for some time instance t and for some point $u \in \Omega$ we have:

$$\sigma(u, t) < \sigma(u, t + 1). \quad (45)$$

Let t be the first such time instance. As usual we denote by p_t the robot's location at time t . According to the MAW algorithm, the robot seeks the minimal pheromone level in $R(r, 2r, p_t)$. Say this minimal level is attained at some point $x \in R(r, 2r, p_t)$. Since u changes its value during the t th step we conclude that

$$\sigma(p_t, t) \leq \sigma(x, t). \quad (46)$$

Otherwise no change occurs, according to the MAW algorithm. Since all points in $D(r, p_t)$ get the same value in the marking step, we conclude

$$\sigma(u, t + 1) = \sigma(p_t, t + 1) = \sigma(x, t) + 1 \geq \sigma(p_t, t) + 1. \quad (47)$$

Moreover, according to our assumption,

$$\sigma(u, t) > \sigma(u, t + 1) \Rightarrow \sigma(u, t) \geq \sigma(p_t) + 2. \quad (48)$$

Thus, if we assume that the pheromone level at some point u decreases at time instance t , Equation (48) must hold. Showing that this inequality is wrong we actually obtain a contradiction to the assumption and thus prove the lemma. Let us look at the time instance t_u when the current pheromone level of u ($\sigma(u, t)$) was set. According to Lemma 6,

$$\sigma(p_t, t_u + 1) \geq \sigma(u, t_u + 1) - 1. \quad (49)$$

Since $t_u < t$ and t was chosen to be the first time when the pheromone level at any point in Ω decreases, we conclude that

$$\sigma(p_t, t) \geq \sigma(p_t, t_u + 1). \quad (50)$$

Substituting it into Equation (49) we obtain

$$\sigma(p_t, t) \geq \sigma(u, t) - 1. \quad (51)$$

The inequality in Equation (48) therefore does not hold. This contradiction completes the proof of Lemma 7.

Lemma 8 *If both a and b had been marked by the robot at some time instance t , their pheromone levels obey the proximity principle, i.e.*

$$\text{if } |a - b| \leq r, \text{ then } |\sigma(a, t) - \sigma(b, t)| \leq 1. \quad (52)$$

Proof. Since both a and b had been marked prior to time instance t there exist time instances t_a and t_b when a and b obtained their current pheromone levels accordingly. Applying Lemma 6, we get

$$\begin{cases} \sigma(b, t_a + 1) \geq \sigma(a, t_a + 1) - 1 \\ \sigma(a, t_b + 1) \geq \sigma(b, t_b + 1) - 1. \end{cases} \quad (53)$$

Or, substituting $\sigma(b, t_b + 1) = \sigma(b, t)$ and $\sigma(a, t_a + 1) = \sigma(a, t)$,

$$\begin{cases} \sigma(b, t_a + 1) \geq \sigma(a, t) - 1 \\ \sigma(a, t_b + 1) \geq \sigma(b, t) - 1. \end{cases} \quad (54)$$

Since $t > t_a$ and $t > t_b$ we can apply Lemma 7, i.e.

$$\begin{cases} \sigma(b, t) \geq \sigma(b, t_a + 1) \\ \sigma(a, t) \geq \sigma(a, t_b + 1). \end{cases} \quad (55)$$

Substituting it into Equation 54, we get

$$\begin{cases} \sigma(b, t) \geq \sigma(a, t) - 1 \\ \sigma(a, t) \geq \sigma(b, t) - 1, \end{cases} \quad (56)$$

which means

$$|\sigma(a, t) - \sigma(b, t)| \leq 1. \quad (57)$$

Hence the lemma is proved.

Lemma 9 *The maximal pheromone level tends to infinity as t goes to infinity.*

Proof. The proof is identical to that of Lemma 4. We again introduce a virtual tessellation of domain Ω into n cells so that every such cell can be inscribed into a circle of diameter less than r . As before, we look at the sum

$$S_t = \sum_{i=1}^n m_t^i - \sigma(p_t, t). \quad (58)$$

The only difference is that here m_t^i denotes the minimal marker value that was set by the robot and not as a result of the initial pheromone profile. As in Lemma 1, we get

$$S_t \geq t \forall t. \quad (59)$$

Thus the lemma is proved.

Theorem 4 For any initial pheromone profile $N : \Omega \mapsto Z^*$, the domain Ω will be covered after $n(M_N - m_N + \lceil n/r \rceil) + 1$ steps, where M_N and m_N denote the maximal and the minimal pheromone levels at time $t = 0$, respectively.

Proof. Let us denote by M_N and by m_n the maximal and minimal values of the initial pheromone level given by N , respectively. According to Lemma 9, the maximal pheromone level in Ω grows and will eventually reach the value of $M_N + \lceil n/r \rceil + 1$. We claim that at this moment the whole domain is covered by the robot. Indeed, let us look at some point $a \in \Omega$ that obtained this pheromone level at step t . According to Lemma 6, for any point $b \in \Omega$ such that $\|a - b\| \leq r$ we have

$$\sigma(b, t + 1) \geq \sigma(a, t + 1) - 1 = M_N + \left\lceil \frac{n}{r} \right\rceil. \quad (60)$$

Since $M_N + \lceil n/r \rceil > M_N$ we conclude that all such points are covered. In the same manner, we determine that all points with distance from a less than or equal to $2r$ are also covered. The maximal distance between points in Ω is bounded by d ; if any one point reaches the pheromone level of $M_N + \lceil n/r \rceil + 1$ we are therefore assured that the minimal possible pheromone level is $M_N + 1$, meaning that the whole domain Ω has been covered. The time needed to cover the domain is

$$t_{\text{cover}} = n \left(M_N - m_N + \left\lceil \frac{n}{r} \right\rceil \right) + 1, \quad (61)$$

where m_N denotes the minimal pheromone level at time $t = 0$. Since the whole domain is covered we can guarantee, by Lemma 8, the proximity principle is obeyed by any two points in Ω and further repetitive coverage is governed by Theorem 3.

5.3. 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 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 shall assume that the clock phases of all robots are slightly different so that no two robots are active at the same time. Each robot therefore sees other robots as regular stationary obstacles and acts accordingly. With this approach, we avoid the problem of simultaneous attempts to set (probably different) pheromone levels at a particular location by multiple robots, since only one robot is active at any given time.

Let us find the upper bound for complete coverage provided we have n robots. Using the same notation as in Equation (21), we have

$$S_t = \sum_{i=1}^n m_t^i - \sum_{j=1}^n \sigma(p_t^j, t), \quad (62)$$




Fig. 7. Use of different metrics: (a) L_1 metric; (b) L_2 metric; and (c) L_∞ metric.

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

$$S_{t+1} > S_t \quad (63)$$

and consequently

$$S_t \geq t, \quad (64)$$

which leads us to the same upper bound as for a single robot. Hence, adding more robots does not necessarily guarantee better coverage time. However, simulations (see Section 6) demonstrate that there is a substantial improvement when we use more robots.

5.4. Generalization for Other Metrics

Until now, we have always assumed that the domain Ω is a flat 2D domain and the usual Euclidean notion of the distance was used. Nevertheless, it is easy to verify that all the proofs remain valid if we change the Euclidean (often referred to as L_2) distance to another valid metric. For example, we could use L_1 distance or, alternatively, the L_∞ distance which is particularly suitable for computer simulations. Of course, each choice of the metric changes the form of the robot's effector. Three different forms, shown in Figures 7a–c, correspond to the L_1 , L_2 and L_∞ metrics accordingly.

We can therefore use the algorithm in any metric space, e.g. we are not limited to 2D spaces as the results remain valid for higher dimensions. For example, we can use the same algorithm for covering 3D volumes assuming the robot's effector is a ball of radius r or, more probably, a regular octahedron or a cube if we choose to work with L_1 or L_∞ metrics, respectively.

6. Simulations and Experiments

6.1. General Notes

We used L_∞ metric in our experiments as the square shape of the effector and the sensing area that correspond to this metric is particularly suitable for computer simulations. In experiments with an initial pheromone profile, the robots were forced




Fig. 8. Simulation domains: (a) A; (b) B; and (c) C.

to start at an unaffected location, i.e. at locations with minimal pheromone level at time $t = 0$. Additionally, 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. Experiments were conducted on the domains shown in Figure 8.

All domains are of size 100×100 pixels and the marking radius in all experiments was set to 3, i.e. each step robot marks a square of 5×5 pixels.

6.2. Comparing MAW to other algorithms

In this experiment we studied the performance of three different algorithms: MAW, MAC (Wagner et al. 2000) and Random Walk. All algorithms used the same square effector of size 5×5 pixels; additionally, the steps of the Random Walk algorithm were restricted to the interval $[r, 2r]$ as in the steps in the MAW algorithm.

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 generally performs better than the theoretical upper bound obtained in Section 4.2. (Recall that the bound is quadratic and results were better than that of MAC, which is asymptotically linear). Cover time of the Random Walk was omitted from Figures 9b and 9c as the values were so big that the difference between the MAC and the MAW algorithms became invisible on this scale.

Note that our upper bound on coverage time is quadratic, while the above results suggest that the actual coverage time is linear. We can demonstrate the predicted quadratic coverage time by tailoring specific tie-breaking rules for specific domains. For example, we can use a domain comprising n loops as shown in Figure 10.

For this domain and specific tie-breaking rules we can obtain quadratic coverage time for the MAW algorithm, while the MAC algorithm still demonstrates linear coverage time. These results are shown in Figure 11.

6.3. MAW in an Environment with an Initial Pheromone Profile

In this section we present the results of our simulation of the MAW algorithm in the presence of an initial pheromone profile. In the first scenario, we ran one robot on Domain A, each time changing the amount of affected pixels (having non-zero pheromone level at $t = 0$). The initial pheromone level values are uniformly distributed in interval $[1, 10]$, i.e. given that 60% of the pixels are affected there are about 6% that obtained a value of 1 and 6% that obtained a value of 2, etc. An example of such a profile with 60% of affected pixels is shown in Figure 12a.

In Figure 12b we demonstrate the cover time as a function of the amount of affected pixels in this scenario.

In another scenario, we chose to explore the influence of the pheromone profile with constant values on the performance of the algorithm. This time, the initial pheromone values in each experiment were constant and again randomly distributed in the space. An example of such a pheromone profile for a constant value of 10% and 30% affected pixels is shown in Figure 13a.

We ran a series of experiments with initial pheromone profile values of 10, 20, 30, 40 and 50. The results of cover time versus the percentage of affected pixels is shown in Figure 13b.

As we can see, the value of the initial pheromone profile does not play a significant role in this scenario, at least in the limits from 10 to 50 used here. This is probably due to the




Fig. 9. Cover time: (a) domain A; (b) domain B; and (c) domain C.




Fig. 10. Domain comprising n loops.

nature of the algorithm that discards high pheromone values in the presence of lower values. To check this we conducted another experiment where the initial pheromone profile occupies a compact space in the domain, i.e. given 40% of affected pixels we form a plateau of the pheromone level as shown in Figure 14a. In this case, the influence of the initial pheromone level value is quite pronounced as expected. See Figure 14b for covering time versus affected pixels percent-

age for the initial pheromone level values of 10, 20, 30, 40 and 50.

As our experiments demonstrate, the MAW algorithm has little sensitivity to ‘non-compact’ distribution of initial pheromone profile. Note that the initial pheromone profile does not affect the Random Walk but completely destroys the MAC algorithm, making it unable to cover the domain completely.

7. Practical Implementation

We would like to address several issues that may arise in a hardware implementation of the algorithm. First, using pheromones as marking is a widespread phenomena in nature.




Fig. 11. Worst-case coverage time of specific domain.




Fig. 12. MAW in presence of initial pheromone profile 1: (a) example of initial pheromone profile 1 and (b) cover time.

However, in robotics such an approach is rare due to complexity of ventilation systems and odor sensors. A much better alternative would be the use of color or electromagnetic markings.

Another practical problem may be associated with the requirement to mark an open disk. This requirement can be eas-

ily eliminated, however, if we extend the robot’s sensing range to $3r$. In this case the robot does not need to mark an open disk; instead, a single point at its current location can be marked.

Note also that current analysis does not address some issues of a practical implementation such as imperfect sensors and/or actuation. These issues will be addressed in further research.




Fig. 13. MAW time in presence of initial pheromone profile 2: (a) example of initial pheromone profile 2 and (b) cover time.

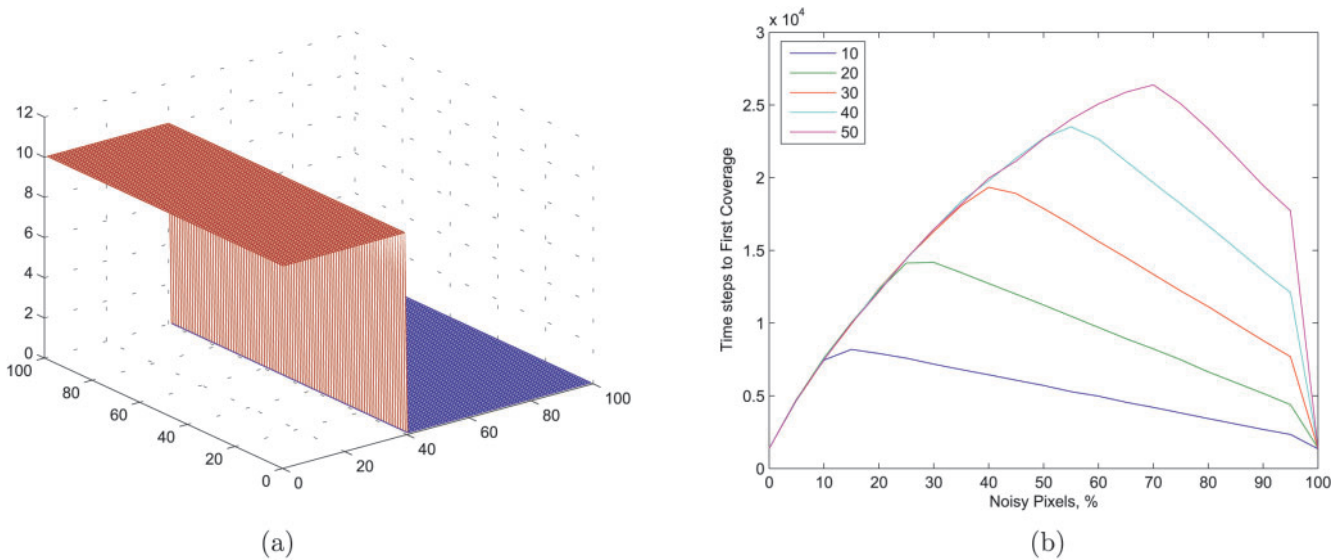


Fig. 14. MAW time in presence of initial pheromone 3: (a) example of initial pheromone profile 3 and (b) cover time.

8. Conclusions

This work has two major contributions. First, we presented a new ant-inspired algorithm for continuous domain covering. We provided a formal proof of complete coverage and upper time bounds for complete coverage and the time interval between two successive visits of the robot. We also proved that the algorithm is capable of coping with arbitrary initial pheromone profile. A formal proof was provided for multi-

agent or multi-robot coverage under the assumption that the robots have different clock phases. Second, a new method of performance analysis was suggested which provides some bounds on possible coverage time of any such mark-and-walk algorithm.

References

- Acar, E. U., Choset, H., Zhang, Y., and Schervish, M. J. (2003). Path planning for robotic demining: Robust sensor-based coverage of unstructured environments and probabilistic methods. *International Journal of Robotic Research*, 22(7–8): 441–466.
- Bender, M. A., Fernández, A., Ron, D., Sahai, A., and Vadhan, S. (1998). 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, pp. 269–278. ACM Press.
- Blum, M. and Kozen, D. (1978). On the power of the compass. In *Proceedings of the 19th Annual Symposium on Foundations in Computer Science*, pp. 132–142.
- Blum, M. and Sakoda, W. (1977). On the capability of finite automata in 2 and 3 dimensional space. In *Proceedings of Annual Symposium on Foundations in Computer Science*, pp. 147–161.
- Bonabeau, E. and Théraulaz, G. (2000). Swarm smarts. *Scientific American*, 282(3): 72–79.
- Bruckstein, A. M. (1993). Why the ant trails look so straight and nice. *The Mathematical Intelligencer*, 15(2): 58–62.
- Butler, Z. J. (2000). *Distributed coverage of rectilinear environments*. Ph. D. thesis, Carnegie Mellon University.
- Chandra, A. K., Raghavan, P., Ruzzo, W. L., and Smolensky, R. (1989). The electrical resistance of a graph captures its commute and cover times. In *STOC '89: Proceedings of the twenty-first annual ACM symposium on Theory of computing*, New York, NY, USA, pp. 574–586. ACM Press.
- Choset, H. (2001). Coverage for robotics: A survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31(1–4): 113–126.
- Choset, H. and Pignon, P. (1997). Coverage path planning: The boustrophedon decomposition. In *Proceedings of International Conference on Field and Service Robotics*, Canberra, Australia, December 1997.
- de Silva, V. and Ghrist, R. (2006). Coordinate-free Coverage in Sensor Networks with Controlled Boundaries via Homology. *The International Journal of Robotics Research*, 25(12): 1205–1222.
- Dorigo, M., Caro, G. D., and Gambardella, L. M. (1999). Ant algorithms for discrete optimization. *Artificial Life*, 5(2): 137–172.
- Dorigo, M., Maniezzo, V., and Colomi, A. (1996). Ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics—Part B*, 26(1): 29–41.
- Gage, D. W. (1994). Randomized search strategies with imperfect sensors. In W. H. Chun and W. J. Wolfe (Eds.), *Proceedings of International Society of Optical Engineering (SPIE)*, Volume 2058, pp. 270–279.
- Harris, R. A., de Ibarra, N. H., Graham, P., and Collett, T. S. (2005). Ant navigation: Priming of visual route memories. *Nature*, 438: 302.
- Hölldobler, B. and Wilson, E. O. (1990). *The Ants!* Harvard University Press, Cambridge, MA.
- Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1): 90–98.
- Koenig, S. and Liu, Y. (2001). Terrain coverage with ant robots: a simulation study. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, New York, NY, USA, pp. 600–607. ACM Press.
- Koenig, S., Szymanski, B. K., and Liu, Y. (2001). Efficient and inefficient ant coverage methods. *Annals of Mathematics and Artificial Intelligence*, 31(1–4): 41–76.
- Rivest, R. L. and Leiserson, C. E. (1990). *Introduction to Algorithms*. New York, NY, USA: McGraw-Hill, Inc.
- Russell, R. A. (1999). Ant trails - an example for robots to follow? In *Proceedings 1999 IEEE International Conference on Robotics and Automation*, pp. 2698–2703, 1999.
- Schöne, H. (1984). *Spatial orientation: The spatial control of behavior in animals and man*. Princeton, N.J.: Princeton University Press.
- Wagner, I., Lindenbaum, M., and Bruckstein, A. M. (1998). Efficiently searching a graph by a smell-oriented vertex process. *Annals of Mathematics and Artificial Intelligence*, 24(1–4): 211–223.
- Wagner, I., Lindenbaum, M., and Bruckstein, A. M. (2000). MAC vs. PC: Determinism and randomness as complementary approaches to robotic exploration of continuous unknown domains. *ROBRES: The International Journal of Robotics Research*, 19: 12–31.
- Wagner, I. A. and Bruckstein, A. M. (2001). From ants to a(ge)nts: A special issue on ant-robotics. *Annals of Mathematics and Artificial Intelligence*, 31(1–4): 1–5.
- Wagner, I. A., Lindenbaum, M., and Bruckstein, A. M. (1996). 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, pp. 219–230. IEEE Computer Society Press.
- Yanovski, V., Wagner, I. A., and Bruckstein, A. M. (2001). Vertex-ant-walk: A robust method for efficient exploration of faulty graphs. *Annals of Mathematics and Artificial Intelligence*, 31(1–4): 99–112.
- Zelinsky, A., Byrne, J. C., and Jarvis, R. A. (1993). Planning paths of complete coverage of an unstructured environment by a mobile robot. In *Proceedings of International Conference on Advanced Robotics (ICAR)*, pp. 533–538, Tokyo, Japan, November 1999.