

# A Multi-Robot Coverage Approach based on Stigmergic Communication

Bijan Ranjbar-Sahraei<sup>1</sup>, Gerhard Weiss<sup>1</sup>, and Ali Nakisaei<sup>2</sup>

<sup>1</sup>Dept. of Knowledge Engineering, Maastricht University, The Netherlands  
<sup>2</sup>National Organization for Development of Exceptional Talents, Shiraz, Iran  
{b.ranjbarsahraei,gerhard.weiss}@maastrichtuniversity.nl,  
ali.n123@gmail.com

**Abstract.** Recent years have witnessed a rapidly growing interest in using teams of mobile robots for autonomously covering environments. In this paper a novel approach for multi-robot coverage is described which is based on the principle of pheromone-based communication. According to this approach, called **StiCo** (for “Stigmergic Coverage”), the robots communicate indirectly via depositing/detecting markers in the environment to be covered. Although the movement policies of each robot are very simple, complex and efficient coverage behavior is achieved at the team level. **StiCo** shows several desirable features such as robustness, scalability and functional extensibility. Two extensions are described, including **A-StiCo** for dealing with dynamic environments and **ID-StiCo** for handling intruder detection. These features make **StiCo** an interesting alternative to graph-based multi-robot coverage approaches which currently are dominant in the field. Moreover, because of these features **StiCo** has a broad application potential. Simulation results are shown which clearly demonstrate the strong coverage abilities of **StiCo** in different environmental settings.

## 1 Introduction

In recent years there has been a rapidly growing interest in using teams of mobile robots for covering and patrolling environments of different types and complexities. This interest is mainly motivated by the broad spectrum of potential civilian, industrial and military applications of multi-robot surveillance systems. Examples of such applications are the protection of safety-critical technical infrastructures, the safeguarding of country borders, and the monitoring of high-risk regions and danger zones which cannot be entered by humans in the case of a nuclear incident, a bio-hazard or a military conflict. Triggered by this interest, today automated coverage is a well established topic in multi-robot research which is considered to be of particular practical relevance.

Currently available theoretical and algorithmic approaches to multi-robot coverage are typically of a computational complexity which excludes their usage in non-trivial application scenarios. Moreover, many of these methods are based

on unrealistic assumptions. Examples of such assumptions are idealized sensors/actuators or sensors with infinite range (e.g. [1]), convexity and/or stationarity of the environment (e.g. [2]), the availability of unlimited communication bandwidth, and fully reliable direct communication links (e.g. [3]).

This article, which is an extension to the work reported in [4], presents a multi-robot coverage approach called **StiCo** (“Stigmergic Coverage”) that avoids such type of assumptions. Specifically, **StiCo** is of a very low computational complexity and is designed for robots with very simple low-range sensors. Moreover, this approach does not rely on direct communication among robots. Instead, the covering robots coordinate on the basis of an indirect communication principle known as stigmergy. According to this principle, which was first observed in biological systems such as ant and termite colonies, natural entities improve their collective performance by influencing one another in their individual performance through local messages they deposit in their shared environment. In computer science, and especially in the field of ant algorithms (e.g., [5]), a number of computational variants of stigmergy have been developed and it has been shown that they allow for very efficient distributed control and optimization in a variety of problem domains (e.g., [6]). In addition to efficiency and distributedness, stigmergy-based coordination has several other properties which are also essential to multi-robot covering algorithms, including robustness, scalability, adaptivity and simplicity. In particular, a main advantage of stigmergy-based communication is its suitability for applications in environments with limited or intermittent network connectivity (e.g., in a devastated area after an earthquake or a military area under attack of jammers) [7, 8]. This makes **StiCo** applicable in principle even in destructed environments where limited or no direct communication is possible. In addition to this, because robots use their environment for saving and transmitting messages no critical requirements are imposed on the storage memory of the individual robots.

The rest of this article is organized as follows. Related works is overviewed in Section 2. Section 3 gives a precise system description and problem formulation. **StiCo** is described in detail in Section 4. Simulation results are shown in Section 5 and Section 6 concludes the article.

## 2 Related Work

Our work is built on the notion of stigmergic communication introduced by Marco Dorigo [5]. The basic idea underlying this form of communication is that pheromones are used as a medium for transmitting messages among artificial ants. During the last years, computational variants of Dorigo’s method (also known as ACO) have been developed and it has been shown that it allows for very efficient distributed control and optimization in a variety of problem domains [6]. Wagner et al. [9] were the first who invested stigmergic multi-robot coordination for covering/patrolling the environment. In their approach a group of robots were assumed which are able to (1) deposit chemical odor traces and (2) evaluate the strength of smell at every point they reach. Based on these

assumptions, they used robots to model an un-mapped environment as a graph and they proposed basic graph search algorithms (such as Depth-First-Search and Breadth-First-Search) for solving robotic coverage problems. Many other researchers used this graph-based modeling scheme in order to design solutions for multi-robot patrolling/covering problems [10–14]. For example, in [10] Elor and Bruckstein mixed *cycle finding* algorithm with *spreading* algorithm in order to provide a finite-time cycle-based patrolling approach.

In contrast to all of the mentioned graph-based techniques, we use a geometrical framework which does not require to model the whole environment as a graph. Specifically, our geometrical approach is similar to Voronoi-based techniques that have recently been introduced for solving robot coverage problems (e.g., see Cortes et al. [3, 15] and Schwager et al. [2, 16]). These Voronoi-based techniques aim at devising coverage algorithms which work according to the following basic rule: *Each vehicle moves toward the center of its Voronoi region*. Based on this rule many researchers have proposed modified covering approaches which are adaptable to changes in the environment and are provably convergent (e.g., [2, 17]). However, all these geometrical algorithms require a group of robots with the capability of direct communication and in most of the cases also need very complex mathematical computations (e.g., calculating margins and center of mass for an individual Voronoi-region) which limits their potential real-world usage.

Another related research topic is focused on the “real” implementation of stigmergic communication in real world experiments. For example, chemical substances such as ethanol ( $C_2H_5OH$ ) are already used instead of natural pheromones [18]. However, with recent developments in communication technology, electrical devices such as Radio Frequency Identification Devices (RFIDs) have gained much interest for such applications. In [7, 8] RFIDs are used for map building and simple pheromone-based explorations. Moreover, in [19] coordinated exploration and multi-robot SLAM for large teams of rescue robots is tackled by using RFIDs as environment features, which are detectable via UHF antennas. Based on characteristics of **StiCo**, it can be implemented on real robots with both chemical pheromones and digital markers.

### 3 Problem Formulation

The basic intention behind the work described here is to design a *motion policy* which enables a group of *robots*, each equipped only with simple *sensors*, to efficiently *cover* a possibly complex *environment*. Moreover, the basic idea pursued is to utilize the principle of pheromone-based coordination and to let each robot deposit *pheromones* on boundaries of its *territory* to inform the others about the already covered areas. This section defines and clarifies some key terms which are relevant to this intention and idea and will be used throughout this article.

- **Environment:**  $Q \subset \mathbb{R}^2$  is an *allowable environment* with area  $A$ , where “allowable environment” is defined as a closed and simply connected set which has a finite number of strict concavities [20].

- **Robot:** A *Dubins* vehicles [21] described by the dynamical system

$$\dot{x} = v \cos \theta, \quad \dot{y} = v \sin \theta, \quad \dot{\theta} = \omega, \quad (1)$$

where  $x, y \in \mathbb{R}$  denote the vehicle position and  $\theta \in \mathbb{S}^1$  denotes its orientation. The control inputs  $v$  and  $\omega$ , describe the forward linear velocity and the angular velocity of the vehicle respectively, while  $v$  is set equal to  $v_0$  (i.e. the nonholonomic vehicle is constrained to move at a constant linear speed) and the control input  $\omega$  takes value in  $[-1/\rho, 1/\rho]$ ;  $1/\rho$  being the maximum curvature.

- **Sensor:** Each robot is equipped with two ant-antenna like sensors, placed on the front-right and front-left corners. These sensors have the ability to detect presence of pheromones from a predetermined distance called  $R_d$ , where  $R_d$  is considered to be very small.
- **Pheromone:** A chemical substance or an electrical marker placed at an arbitrary position  $(x_p, y_p)$ . The pheromone is fully evaporated (naturally or artificially) after time  $T_e$ .
- **Territory:** Inspired by real ants, each robot considers a circular environment of area  $A_T$  as its territory and circles around this area persistently. The area of territory is related to angular and linear velocity of robot as:  $A_T = \pi(v/\omega)^2$ .
- **Motion Policy:** The motion policy tells a robot what to do at each iteration of time. Therefore, when a robot detects pheromone, it decides based on this policy what to do next.
- **Coverage:** We consider an environment to be covered, as a condition that no two robot territories share a common area of the environment. Therefore, the motion policy should guide the robots in a way that their territory intersections decrease as time passes. When the full coverage is achieved (i.e. no territories have intersection), each robot patrols its territory by moving on the territory border, persistently.

## 4 Design of the StiCo Approach

The basic notion underlying **StiCo** is to partition the environment into equal circular regions (also called territories) where each robot takes responsibility to guard one of these regions. The robots need not communicate directly, but deposit pheromones on the borders of their territory for instructing other robots to not enter it. In this way **StiCo** answers the core question “*How should robots move in order to decrease the intersections of their territories*”.

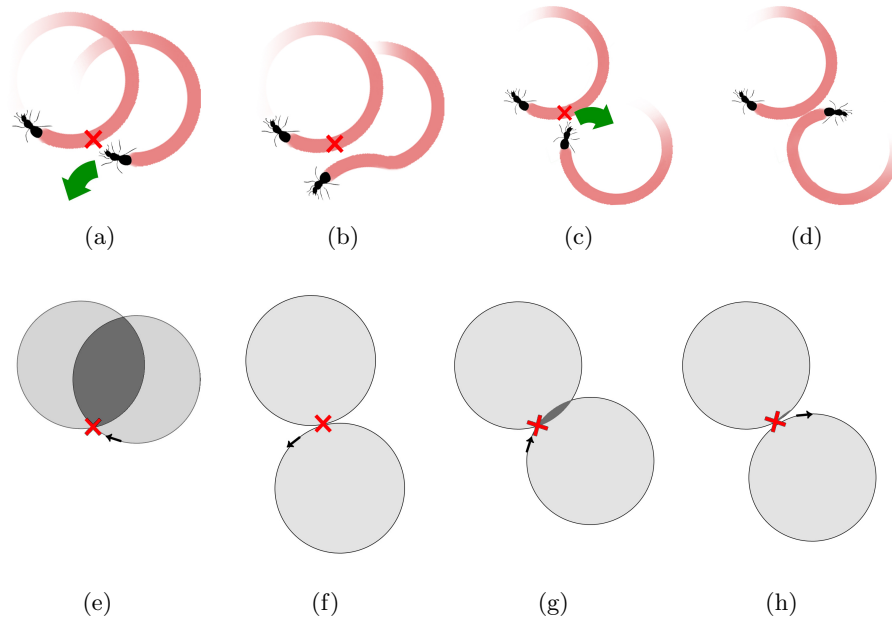
### 4.1 Basic StiCo

In **StiCo**, each robot starts to move with a constant forward linear velocity  $v_0$ , and a constant angular velocity  $w_0$ , which results in a circular motion on the border of a territory with radius  $v_0/w_0$ . The forward linear velocity remains

constant during the whole mission. However, in different situations the angular velocity might increase or decrease based on the motion policy.

In order to adjust the angular velocity, based on the circling direction (CW or CCW), one sensor serves as the interior sensor (the one nearer to the center of territory) and the other one as the exterior sensor.

When the interior sensor detects a pheromone (Figure 1a), it indicates to the robot that it is about entering another territory, and therefore the robot changes its circling direction immediately (Figure 1b). In this way, the robot establishes its territory in a new region without any intersection with the other territory. Otherwise, if exterior sensor detects a pheromone (Figure 1c), this tells the robot that it is passing near another territory (however, not completely entering it as in Figure 1a). In this case the robot rotates (i.e., magnitude of  $w_0$  is increased up to  $1/\rho$ ) until it does not detect pheromone any more and then circles in the same direction with the constant angular velocity  $w_0$  (Figure 1d). Therefore, the intersection between two territories is fully eliminated with a small displacement of territory.



**Fig. 1.** StiCo coordination principle: (a)-(b) before and after pheromone detection by internal sensor. (c)-(d) before and after pheromone detection by external sensor. (e)-(f) covered area before and after pheromone detection by internal sensor. (g)-(h) covered area before and after pheromone detection by external sensor.

Figures 1e-1h, illustrate how **StiCo** works. As shown in Figures 1e and 1g, before pheromone detection there is an intersection between two territories. When one robot detects pheromone and changes its territory area the intersection is fully eliminated, as shown in Figures 1f and 1h. As a result, the robots keep their territories disjoint. In the case of a large swarm of robots, eliminating one intersection may cause the emergence of other intersections. However, as the simulation results of this article show, the statistical chance for decrement of intersections is significantly higher than the chance for its increment – and this makes **StiCo** a very efficient coverage approach.

**StiCo** is further detailed in Algorithm 1.

---

#### Algorithm 1 **StiCo**

---

**Require:** Each robot can deposit/detect pheromone trails  
Initialize: Choose circling direction (CW/CCW)  
**loop**  
    **while** (no pheromone is detected) **do**  
        Circle around  
        deposit pheromone  
    **end while**  
    **if** (interior sensor detects pheromone) **then**  
        Reverse the circling direction  
    **else**  
        **while** (pheromone is detected) **do**  
            Increase the magnitude of angular velocity (Rotate)  
        **end while**  
    **end if**  
**end loop**

---

## 4.2 **StiCo** Extensions

By applying **StiCo** on a swarm of robots, complex behavior emerges and robots disperse in the environment homogeneously to cover the maximal possible area. Although this novel coverage approach generates very efficient coverage results based on relatively simple motion rules, it can be extended in two important ways. First, toward dealing with dynamically changing environments. In such environments it is difficult how to choose certain parameters of the multi-robot system such as motion speed of the individual robots. As a solution to this, **StiCo** can be extended by treating the territory area of each robot as an adaptable term: a robot can increase the size of its territory (by decreasing the magnitude of angular speed) when detecting large uncovered areas near its territory borders. We call this motion policy **A-StiCo** (for "Adaptive **StiCo**"). Second, **StiCo** allows to easily add intruder detection behavior: robots decrease their respective territories as soon as they detect the presence of an intruder. By adding this behavior to **A-StiCo**, the density of robots near the intruder increases automatically. The

effective behavior of **ID-StiCo** (for "Intruder Detection **StiCo**") suggests to use this coverage approach in surveillance missions of unknown environments.

## 5 Simulation Results

In this section, we demonstrate the evolution of **StiCo** on four simulation scenarios. In the first scenario robots are initialized in the center of an obstacle-free environment and disperse in it homogeneously in order to partition the environment into circular regions. In this simulation, the scalability of **StiCo** is demonstrated by using a unique motion policy for robotic swarms of different sizes. In the second scenario, obstacles are used to generate a non-convex coverage problem. The main goal of this simulation scenario is to demonstrate the robustness of **StiCo** in complex environments. Then, two possible extensions on **StiCo**, **A-StiCo** and **ID-StiCo**, are discussed and related simulations results are illustrated. A video presentation of different characteristics of **StiCo** is available at: <http://youtu.be/DOlyqDN2a9o>.

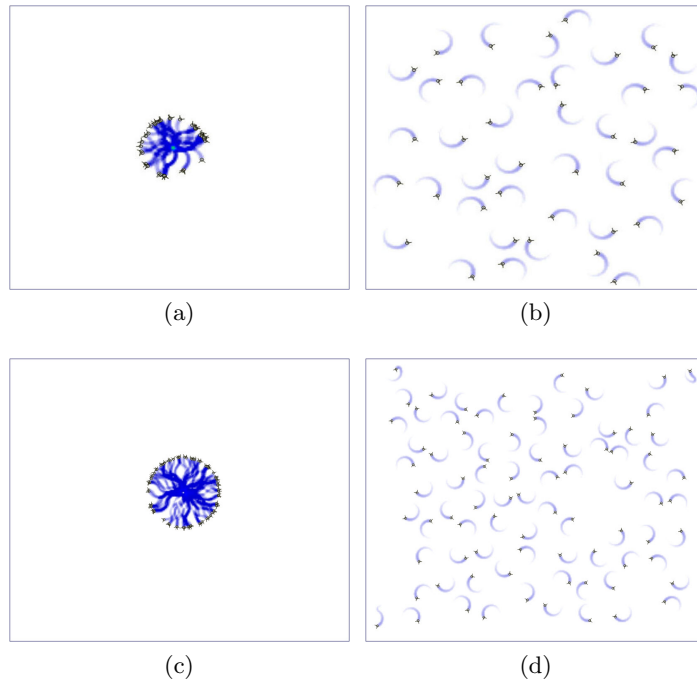
All of the simulations are implemented on a robotic swarm of identical members initialized in the center of a  $40m \times 40m$  field. The pheromones are simulated with a high resolution, equal to  $300 \times 300$  and the evaporation time is  $T_e = 1.5s$ . Moreover, we pay careful attention to numerical accuracy and optimization issues in the pheromones update policy.

### 5.1 Scenario 1: Convex Environment

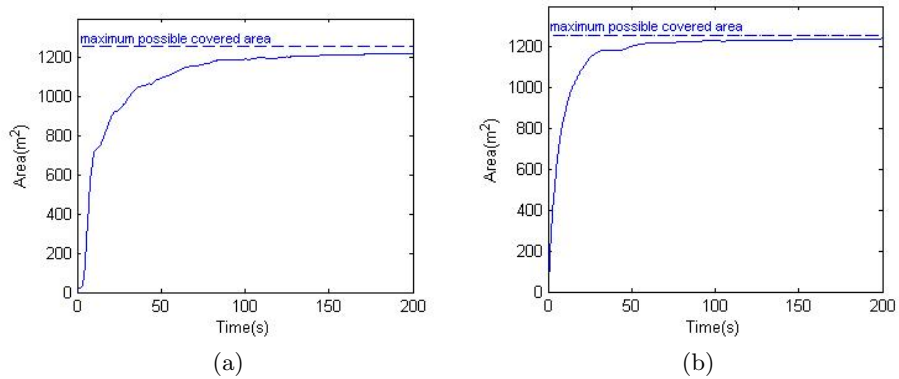
Scenario 1 consists of a convex environment (a square of size  $1600m^2$  shown in Figure 2). All of the robots are initialized in the center of this environment with different initial angles. The execution of **StiCo** on a group of 40 robots is illustrated in Figures 2a and 2b and its execution on a group of 80 robots is illustrated in Figures 2c and 2d.

The snapshots shown in Figure 2, confirm our predictions in Section 4 that the intersected area between territories is completely eliminated after a while and robots are dispersed in the environment homogeneously. It is intuitively clear that when robots are placed in a configuration that no two territories have intersection, then the whole configuration remains stable and the robots move on the borders of their territories, persistently.

In order to depict the coverage performance of this algorithm in respect of time, we run **StiCo** for 30 times with different initial positions. Then by averaging the overall covered area in each iteration over different simulation runs, we can compute the *estimated covered environment*. Moreover, based on basic geometry, the maximum possible fraction of a square which can be covered by a set of disjoint identical circles is 78.5%. Therefore, in the best case  $1256m^2$  of the considered environment can be covered. Figure 3 shows a comparison of the estimated covered environment with this maximum possible coverage for both groups of 40 and 80 robots. As can be seen, **StiCo** converges to this maximum in both cases.



**Fig. 2.** The evolution of coverage achieved by **StiCo**: (a),(b) Initial and final position of the 40-robot group. (c),(d) Initial and final position of the 80-robot group.

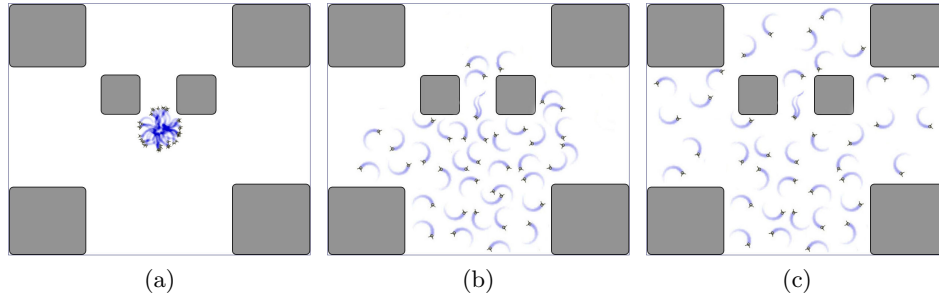


**Fig. 3.** The estimated covered environment: (a) 40-robot group. (b) 80-robot group.

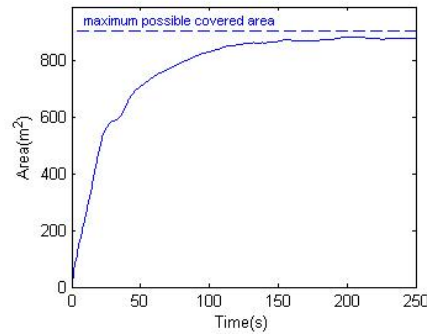
## 5.2 Scenario 2: Non-Convex Environment

In order to demonstrate potential capabilities of the **StiCo** approach, we consider as a second scenario a non-convex environment as shown in Figure 4a. This





**Fig. 4.** Evolution of coverage of a non-convex environment achieved through **StiCo**: (a) Initial snapshot. (b) Intermediate snapshot. (c) Final snapshot.



**Fig. 5.** The estimated covered environment for a 40-robot group in a non-convex environment.

environment can represent, for instance, a devastated area after an earthquake, or a street map in an emergency condition.

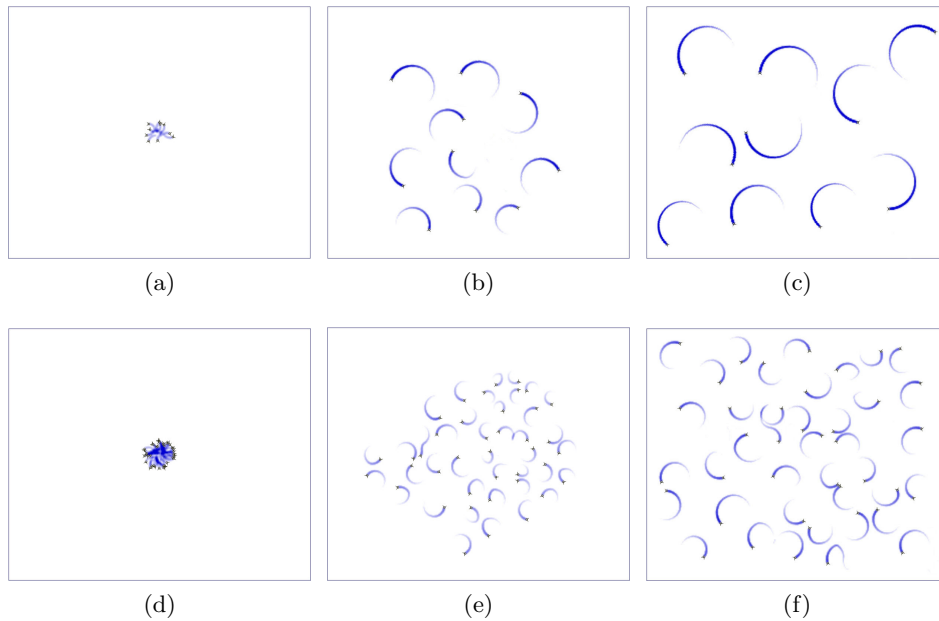
For coverage of this environment, a group of 40 robots are initiated at the center of the environment with different initial angles. **StiCo** is executed on this group and snapshots of this simulation are shown in Figure 4(a-c). (In this simulation, artificial pheromones are deposited on the borders of obstacles to make them detectable for robots).

As can be seen in this figure, the **StiCo** approach is robust to environmental complexities. Although the robots are not equipped with any path planning system, they are able to disperse homogeneously in the environment independent of where obstacles are placed.

The obstacle-free area of this environment is equal to  $1150m^2$  and, as mentioned in the preceding subsection, in the best case 78.5% of this area can be covered by a set of disjoint identical circles. Figure 5 compares this maximum with the estimated covered environment achieved by **StiCo**. As can be seen again, **StiCo** is able to reach the maximum possible coverage.

### 5.3 Extension 1: A-StiCo

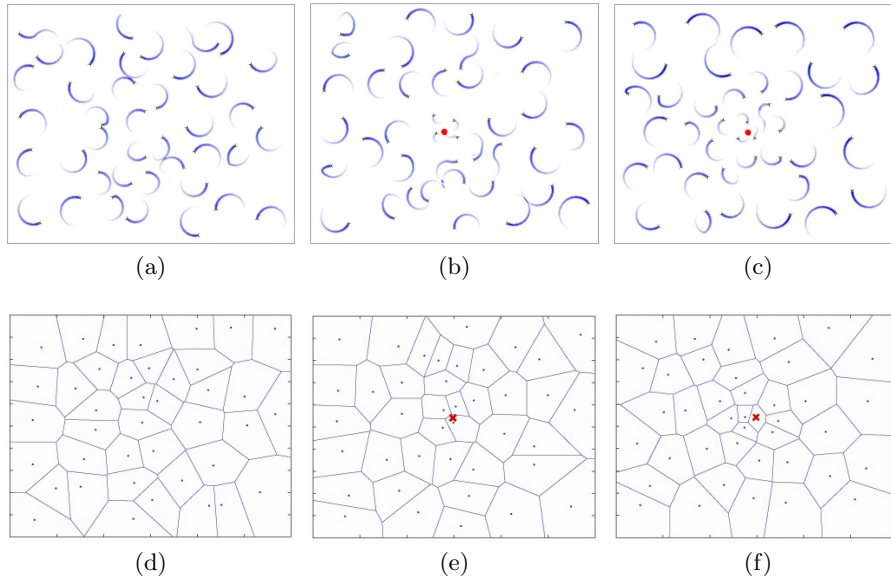
In this simulation we show that by adding adaptive behaviors to the **StiCo** approach, even more efficient coverage results can be achieved. In **A-StiCo**, when a robot does not detect pheromones for a while, it decreases its angular velocity ( $w_0$ ). Consequently, the territory area is expanded and the robot guards a larger region. Otherwise, when a robot detects pheromones very often (which means that many robots have been moving nearby), it increases its angular velocity. Consequently, the robot guards a smaller region. By adding this simple adaptive behavior to the **StiCo** approach, robots are able to cover the environment more adaptively. Figure 6 depicts the evolution of coverage for two swarms of 10 and 40 robots. In both simulations, robots start from the same initial conditions.



**Fig. 6.** The evolution of **A-StiCo**: (a)-(c) Initial, intermediate, and final snapshots after 250s, for 10 robots. (d)-(f) Initial, intermediate, and final snapshots after 250s, for 40 robots.

### 5.4 Extension 2: ID-StiCo

In Subsection 4.2 we suggested an extended form of **StiCo** called as **ID-StiCo** which realizes intruder detection behavior. According to **ID-StiCo** a robot decreases its territory area as soon as it senses an intruder. Integrating this behavior with the **A-StiCo** approach results in a very effective surveillance characteristic.



**Fig. 7.** The evolution of **ID-StiCo**: (a) Initial homogeneous configuration. (b) Intruder entrance. (c) Final configuration after 200s. (d) Voronoi diagram of initial configuration. (e) Voronoi diagram after intruder entrance. (f) Voronoi diagram of the final configuration.

Figures 7a-7c illustrates the evolution of the intruder detection behavior of this new approach for a group of 40 robots: Figure 7a shows an already achieved coverage of the environment, in Figure 7b a stationary intruder is added in the center of the environment which has resulted in an immediate reaction of nearby robots, and in Figure 7c the robots achieve a final stable configuration.

We use Voronoi diagrams to illustrate and analyze how robots move in their environment for achieving a better coverage of the regions which are closer to the intruder. Figures 7d-7f depict the Voronoi diagram for each snapshot shown in Figures 7a-7c, respectively. As can be seen in Figure 7f, the Voronoi regions close to the intruder are smaller and more concentrated than the regions distant from it – this is exactly the kind of behavior expected from an effective intruder detection approach.

## 6 Conclusion

This article addressed the multi-robot coverage problem in environments of different complexity and presented a new approach called **StiCo** which is based on indirect, stigmergic communication. **StiCo** is a fully distributed motion policy which allows for a very effective and efficient coverage performance. Compared to existing coverage approaches, **StiCo** shows several important advantages, in-

cluding scalability, robustness, very low computational complexity and memory requirements, and easy functional extensibility (as shown with two extensions – **A-StiCo** and **ID-StiCo** – of broad practical relevance). This makes **StiCo** distinct from all other currently available multi-robot coverage approaches.

We think the experimental results justify to invest further research in **StiCo**. **StiCo** opens a promising new research avenue: the comparison of multi-robot stigmergy-based coverage and graph-based environmental coverage. Currently we are working on a mathematical framework for a formal analysis of **StiCo**, and we hope this framework will also contribute to a deeper understanding of how these two types of coverage approaches compare to each other in general. Moreover, we are currently working on an implementation of **StiCo** on a group of 30 e-puck robots in our SwarmLab (<http://swarmlab.unimaas.nl/>).

## References

1. I. Roman-Ballesteros and C.F. Pfeiffer. A framework for cooperative multi-robot surveillance tasks. In *Electronics, Robotics and Automotive Mechanics Conference, 2006*, volume 2, pages 163–170, sept. 2006.
2. M. Schwager, D. Rus, and J. J. Slotine. Decentralized, adaptive coverage control for networked robots. *International Journal of Robotics Research*, 28(3):357–375, 2009.
3. J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *Robotics and Automation, IEEE Transactions on*, 20(2):243–255, april 2004.
4. Bijan Ranjbar-Sahraei, Gerhard Weiss, and Ali Nakisae. Stigmergic coverage algorithm for multi-robot systems (demonstration). In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.
5. M. Dorigo. *Optimization, Learning and Natural Algorithms*. Thesis report, Politecnico di Milano, Italy, 1992.
6. M. Dorigo, M. Birattari, and T. Stutzle. Ant colony optimization. *Computational Intelligence Magazine, IEEE*, 1(4):28–39, nov. 2006.
7. R. Johansson and A. Saffiotti. Navigating by stigmergy: A realization on an rfid floor for minimalistic robots. In *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, pages 245–252, may 2009.
8. Herianto, Toshiki Sakakibara, and Daisuke Kurabayashi. Artificial pheromone system using rfid for navigation of autonomous robots. *Journal of Bionic Engineering*, 4(4):245–253, 2007.
9. Israel A. Wagner, Michael Lindenbaum, and Alfred M. Bruckstein. Distributed covering by ant-robots using evaporating traces. *IEEE Transactions on Robotics and Automation*, 15(5):918–933, 1999.
10. Yotam Elor and Alfred Bruckstein. Autonomous multi-agent cycle based patrolling. In *Swarm Intelligence*, volume 6234 of *Lecture Notes in Computer Science*, pages 119–130. Springer Berlin / Heidelberg, 2010.
11. Yotam Elor and Alfred M. Bruckstein. Multi-a(ge)nt graph patrolling and partitioning. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 02, WI-IAT '09*, pages 52–57, Washington, DC, USA, 2009. IEEE Computer Society.

12. A. Glad, O. Simonin, O. Buffet, and F. Charpillat. Influence of different execution models on patrolling ant behaviors: from agents to robots. In *Proceedings of the Ninth International Conference on Autonomous Agents and MultiAgent Systems (AAMAS'10)*, 2010.
13. Arnaud Glad, Olivier Simonin, Olivier Buffet, and François Charpillat. Theoretical study of ant-based algorithms for multi-agent patrolling. In *Proceeding of the 2008 conference on ECAI 2008: 18th European Conference on Artificial Intelligence*, pages 626–630, Amsterdam, The Netherlands, The Netherlands, 2008. IOS Press.
14. Vladimir Yanovski, Israel A. Wagner, and Alfred M. Bruckstein. A distributed ant algorithm for efficiently patrolling a network. *Algorithmica*, 37:165–186, 2003.
15. J. Cortes, S. Martinez, and F. Bullo. Spatially-distributed coverage optimization and control with limited-range interactions. *ESAIM: Control, Optimisation and Calculus of Variations*, 11:691–719, 2005.
16. M. Schwager, D. Rus, and J. J. Slotine. Unifying geometric, probabilistic, and potential field approaches to multi-robot deployment. *International Journal of Robotics Research*, 30(3):371–383, 2011.
17. A. Breitenmoser, M. Schwager, J. C. Metzger, R. Siegwart, and D. Rus. Voronoi coverage of non-convex environments with a group of networked robots. In *Proc. of the International Conference on Robotics and Automation (ICRA 10)*, pages 4982–4989, May 2010.
18. R. Fujisawa, H. Imamura, T. Hashimoto, and F. Matsuno. Communication using pheromone field for multiple robots. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 1391–1396, sept. 2008.
19. V.A. Ziparo, A. Kleiner, L. Marchetti, A. Farinelli, and D. Nardi. Cooperative exploration for USAR robots with indirect communication. In *Proc. of 6th IFAC Symposium on Intelligent Autonomous Vehicles (IAV'07)*, 2007.
20. F. Bullo, J. Cortes, and S. Martinez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. 2009. Available at <http://www.coordinationbook.info>.
21. L. E. Dubins. On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79:497–516, 1957.