Hybrid Insect-Inspired Multi-Robot Coverage in Complex Environments

Bastian Broecker[1], Ipek Caliskanelli[1], Karl Tuyls[1], Elizabeth Sklar[1,2], and Daniel Hennes[3]

[1]Department of Computer Science, University of Liverpool, UK
[2]Department of Informatics, King's College London, UK
[3]European Space Agency, NL
{bastian.broecker, ipek.caliskanelli, k.tuyls}@liverpool.ac.uk,
elizabeth.sklar@kcl.ac.uk, daniel.hennes@esa.int

Abstract. Coordination is one of the most challenging research issues in distributed multi-robot systems (MRS), aiming to improve performance, energy consumption, robustness and reliability of a robotic system in accomplishing complex tasks. Social insect-inspired coordination techniques achieve these goals by applying simple but effective heuristics from which elegant solutions emerge. In our previous research, we demonstrated the effectiveness of a hybrid ant-and-bee inspired approach, HybaCo, designed to provide coordinated multi-robot solutions to area coverage problems in simple environments. In this paper, we extend this work and illustrate the effectiveness of our hybrid ant-and-bee inspired approach (HybaCo) in complex environments with static obstacles. We evaluate both the ant-inspired (StiCo) and bee-inspired (BeePCo) approaches separately, and then compare them according to a number of performance criteria using a high-level simulator. Experimental results indicate that HybaCo improves the area coverage uniformly in complex environments as well as simple environments.

1 Introduction

Recent years have seen a rapidly growing interest in distributed multi-robot systems for automatically exploring and surveilling environments of different sizes, type and complexity. *Multi-robot systems (MRS)* consist of multiple interacting robots, each executing an application-specific control strategy, which is not centrally steered. *Coordination* is a key challenge when deploying teams of distributed multi-robot systems as these systems are often limited in resources (e.g., on-board processors, batteries, controllers). Lightweight interactions among robots (e.g., facilitated by wireless communication) are not only a desired feature for such platforms, but also necessary to overcome practical deployment issues stemming from inconsistent network connectivity. Simple yet effective heuristics that avoid complex, heavy computation and establish lightweight interactions are therefore highly desirable for MRSs. Biologically inspired, or *bio-inspired*, solutions for the challenging problem of multi-robot coordination are gaining traction.

Swarm algorithms, like ant colony optimisation [10], rely on pheromonal knowledge for communication between agents. Such insect-inspired multi-agent research has also opened the possibility of applying some of these techniques to robotic systems, i.e., swarm robotics [11, 16]. Swarm robotic systems are motivated by a wide range of application areas, such as surveillance and patrolling, where mobile guard robots are considered an alternative and improved mechanism over fixed security cameras and even humans. Other application areas include exploration and identification of hazardous environments (e.g., nuclear plants and fire detection), mobile sensor networks, wireless sensor and robot networks, and space exploration.

Previous research, *StiCo* [13], investigated ants' stigmergic behaviour in an attempt to address the coverage issues in multi-robot systems using self-organised coordination techniques. Another approach, *BeePCo* [4], used bee pheromone signalling processes to solve the same multi-robot coverage problem. Both *StiCo* and *BeePCo* rely on pheromone substances for coordination and communication between agents. Based on these approaches, a hybrid ant-and-bee inspired technique was developed, *HybaCo* [2], in an attempt to further improve multi-robot coverage. This included a comparative study of the *StiCo*, *BeePCo* and *HybaCo* algorithms and identified their strengths and weaknesses. The environment used in the experimental scenarios in [2] was a simplistic, obstacle-free square arena. In this paper, we extend this work and introduce more complex environments that contain obstacles, in order to represent more realistic scenarios. We demonstrate the effectiveness of *HybaCo* and compare it with *StiCo* and *BeePCo* on multiple complex environments.

2 Related Work

Many bio-inspired techniques have been examined to address multi-robot coverage by establishing lightweight coordination principles based on the behaviour of social insects. Ants [8,9] and bees [1,12] are the two main families of social insects that have inspired approaches within the fields of robotics and distributed systems as a means to improve self-organisation and autonomous coordination characteristics.

In [13, 14], Ranjbar-Sahraei et al. developed their first *stigmergic* approach using a simulation environment to address coordination issues in multi-robot systems. Extensive experimental results on both simple (obstacle free) and complex (with static and dynamic obstacles) environments showed that the *stigmergic* approach (StiCo) applies (almost) uniform coverage. Later in [17], Ranjbar-Sahraei et al. extended their research deployed on a physical robot swarm, which validated the correctness of the simulated results. The mathematical model of this approach is shown in [15] that derives a probabilistic macroscopic model for ant-inspired coverage by multi-robot platform.

In previous work on *pheromone signalling (PS)* based load-balancing, [5, 7] presented a dynamic technique for *wireless sensor networks (WSNs)* that is applied at run time in the application layer of a communication network.

PS is inspired from the pheromone signalling mechanism found in bees and provides distributed WSN control that uses local information only. [6] extend the initial *PS* technique by introducing additional network elements in the form of robotic vehicles for *wireless sensor and robot networks (WSRNs)*. Different subclasses of cyber-physical systems are merged together to increase the area coverage *effectively*, which directly increases the service availability and extend the network lifetime by benefiting from their heterogeneity. The same pheromone signalling principle is applied to multi-robot systems in [2, 3] and explained in detail in the next sections.

3 Background: Comparison Between StiCo and BeePCo

This section provides some background information about the *StiCo* and *BeePCo* techniques separately, which is essential for later describing our hybrid *HybaCo* in Section 4.

3.1 StiCo Principle

The *StiCo* approach follows the principle of indirect, stigmergic coordination to establish efficient coverage of an environment by simple means. Classic stigmergic coordination in an *ant system (AS)* is characterised by two properties: (1) agents have a tendency to move straight with minor deviations, and (2) traces (signal trails) act as sources of attraction. In contrast, StiCo robots orbit in circles (instead of moving straight), and their traces have repulsion characteristics (instead of attraction). These two key differences transform the path-finding characteristic of an AS into the efficient area coverage provided by StiCo.

In StiCo, robots are equipped with two simple sensors (pointing in the front-left and front-right directions like ant antennae), capable of detecting immediate traces. Each robot rotates in a circle with a predetermined radius. Based on the circling direction (clockwise, CW, or counter-clockwise, CCW), one sensor is considered as the interior sensor and the other as the exterior one. When the interior sensor detects *pheromone* (virtual substance marking the trace of a robot), the robot changes its circling direction immediately. Otherwise, if the exterior sensor detects pheromone, the robot continues rotating in the same direction until it no longer detects any pheromone. For a more detailed description we refer to [13].

3.2 BeePCo Principle

The *BeePCo* approach follows the principle of pheromone-signalled coordination found in bees to establish direct, lightweight communication between agents and can address multi-robot coverage problems. The *BeePCo* algorithm consists of four parts, which are executed on every robot in the MRS: two parts are time-triggered (differentiation cycle and decay of pheromone), whereas the other two (propagation of pheromone and motion direction and magnitude) occur in parallel invoked by one event-triggered process. During the propagation, robots send pheromone to their neighbours within direct communication range. If a robot receives pheromone, it makes a decision to move and selects a target destination in the opposite direction of the pheromone received. The movement decision is based on vector addition; further description can be found in [4].

3.3 Comparison Between StiCo and BeePCo

Here we characterise the strengths and weaknesses of *StiCo* and *BeePCo*. Three key differences between these two techniques are illustrated in Table 1. The first

Table 1: Differences between *StiCo* and *BeePCo*

property	*StiCo*	*BeePCo*
communication	Indirect	Direct
movement	Circular	Vector-based
speed to converge	Normal	Fast

difference between *StiCo* and *BeePCo* lies in how pheromone is used for *communication*. In the *StiCo* approach, communication between agents is implemented using *indirect* pheromone trails, where pheromone signals are deposited in the robots' environment without knowing which or whether another agent will receive the signal. In contrast, in the *BeePCo* approach, pheromone signalling is implemented by *directly* sending signals to robots within a specific range.

The second difference between *StiCo* and *BeePCo* lies in the type of *movement* effected by the robots. When robots run *StiCo*, their motion is applied in a circular fashion and only the direction of circling changes. When robots run *BeePCo*, their motion is guided by vectors which influence the straight-line direction and distance for each move.

The difference between indirect and direct communication methods, taken in combination with the different motion methods, have significant impact on the *speed* with which these algorithms converge. The duration from the moment of deployment until the robots reach a stable configuration and until their energy is depleted varies significantly. *BeePCo* produces faster convergence in comparison to *StiCo*, and this is one of the strengths of *BeePCo*. However, because robots in *BeePCo* use direct communication based on transmission range, the robots stop moving once they are not in each other's communication range; and this is the main weakness of this approach.

4 HybaCo: Hybrid Bee-Ant Coverage Algorithm

In [2], we proposed *HybaCo*, which combines the effectiveness of the two pheromone-based approaches (*StiCo* and *BeePCo*) detailed above while overcoming the major weaknesses of each approach taken alone. In this section, we describe *HybaCo* briefly once again before we evaluate it in complex environments. The most important performance bottleneck of the *BeePCo* occurs when the robots move far apart from each other and lose their communication network. This prevents pheromone exchange and, as a result, robots do not move any more. The biggest problems with the *StiCo* approach are the extended time to converge and the lack of coverage redundancy—which is a desirable feature when considering practical deployment. In order to solve these issues, our hybrid approach begins with *BeePCo* but changes dynamically to *StiCo* when the communication network between the robots is lost [2]. Robots apply the *BeePCo* technique when the communication network is still active; but as the robots move further apart from each other, the communication network dies. When the robots lose connectivity (communication links) with all of their neighbours (i.e., others within transmission range), they assume that the *BeePCo* technique is no longer effective and so they switch to *StiCo*. After some time using *StiCo*, the robots will again get close enough to transmit pheromones to each other, at which point they will switch back to *BeePCo*. It is important to underline that the ANT and BEE pheromones are *different*, and are thus declared separately in the algorithm.

5 Experimental Evaluation

We evaluated three main algorithms (*StiCo*, *BeePCo* and *HybaCo*) using custom-built abstract simulators. The original *StiCo* [13, 14] was developed in C++, which is extended for the *HybaCo* experiments described here. The set of experiments presented in this section compare three important evaluation metrics:

1. the *area covered* by the robots;
2. the *distribution* of robots in their environment; and
3. the *time* it takes to converge (or stabilise).

In this study, *area coverage* is defined as the maximum of the total non-overlapping area covered by the sensors of the involved robot(s), as defined in [17]. The *distribution* of robots in the arena shows that the robots are moving around without leaving unattended gaps in the environment. This is measured by overlaying a grid of $160,000$ cells ($1cm \times 1cm$) on the environment and calculating how many cells are unattended at any moment in time. For the entire environment, all cells that are "attended" (or covered) by one or more robots are summed and divided by the total number of cells, resulting in the *percentage of coverage*. This metric indicates how evenly the robots are distributed in the arena and what the level of redundancy is for the different algorithms. The result is illustrated in the heatmaps shown in Figures 2, 4 and 6. The lighter the colour of the area's in the arena, the higher is the percentage of the area being covered over the total time of the experiment. The more evenly the total area is coloured, the more uniform is the distribution of the robots' positions over time.

Finally, the time it takes to converge is the duration from the moment of deployment of the robots until they achieve a stable configuration (i.e., the moment the algorithm performs nearly optimally before resources are depleted). The experiments were carried out with sets robots which allow maximal coverage of $\approx 75\%$ (depending on the environment this number is between 56 to 60 robots), each robot having a sensing and communication radius of $25cm$. The robots' environment (arena) size is $400cm \times 400cm$. Initially, the robots are deployed randomly in the central square region of size $5cm \times 5cm$.

We consider the following five algorithmic variations of *StiCo* and *BeePCo* in our comparisons:

– ***StiCo***: The robots execute the stigmergic principle for coordination [14], as described in Section 3.1.
– ***BeePCo***: The robots execute the bee-pheromone signalling principle [4] for coordination, as described in Section 3.2.
– ***BeePCo-with-rotation***: The *BeePCo* algorithm is extended with a rotational move. This is an intermediate algorithm between *BeePCo* and *HybaCo* in which the robots execute the *BeePCo* approach until they lose communication links with each other. Once communication is lost, the robots apply the rotational move described in [2].
– ***HybaCo***: The robots execute the bee-and-ant inspired coordination principle as introduced in [2] and described in Section 4.
– ***MaxCo***: This represents the optimal case where the robots' transmission range does not intersect with each other. This scenario is a benchmark for the maximum possible coverage of deployed robots with zero surveillance area overlap. This can also be referred to as *potential* coverage.

These five different algorithms are evaluated over three different environmental setups (topdown, L-shaped, floor plan) as Fig. 1 illustrates. The experimental results presented in this paper is average of five individual runs over each environmental setup for all five algorithms.

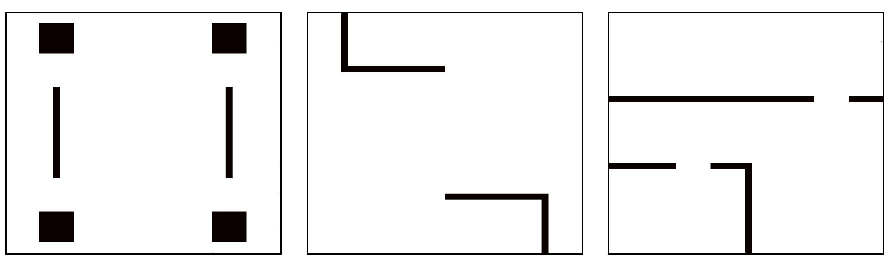

(a) Top down (b) L-shaped (c) Floor plan

Fig. 1: Three environment setups used to evaluate the performance.

Figures 2 and 3 illustrate the experimental results of an MRS of 56 robots in an arena containing four blocks of square obstacles in the corners and two walls comparing the performance of the *StiCo, BeePCo, BeePCo-with-rotation* and *HybaCo* approaches against each other. Fig. 2 shows the heatmap images of this arena containing a number of static obstacles and our observations are as follows. *StiCo* robots had difficulties passing the walls and the square obstacles, therefore, focused on the middle of the arena, as shown in Fig. 2a with darker corners. In *BeePCo*, the robots stop spreading after communication links with the other robots are broken, because they are outside of the inter-robot transmission range.

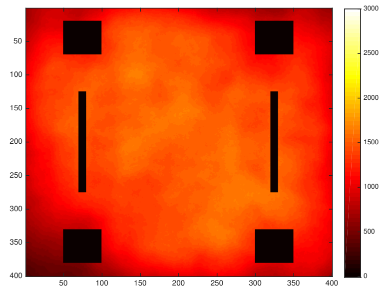

(a) StiCo

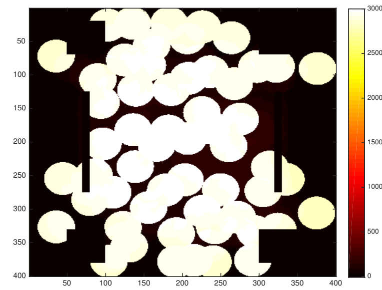

(b) BeePCo

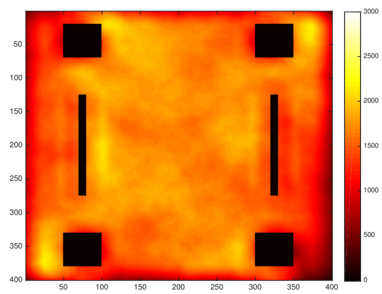

(c) BeePCo with Rotation

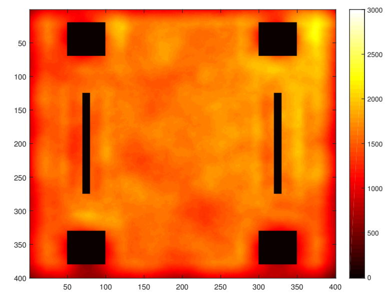

(d) HybaCo

Fig. 2: The distribution of robots in the arena using an MRS of 56 robots on BeePCo, StiCo, BeePCo-with-rotation and HybaCo.

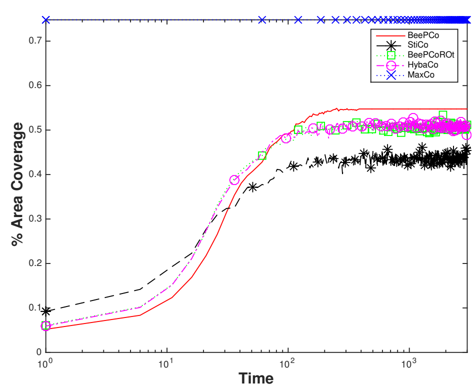

Fig. 3: The percentage of area coverage using MRS of 56 robots on different techniques.

Therefore, in Fig. 2b, a white circle represents an area covered by a robot after it stopped moving and did not receive further pheromone from other robots. This results in entirely uncovered areas in the arena, because the robots are not able to move. The *BeePCo-with-rotation* and *HybaCo* approaches provide more even distribution than *BeePCo*, where *BeePCo-with-rotation* distinctively focuses more around the obstacles. This behaviour is less obvious in *HybaCo* and thus makes *HybaCo* more uniformly distributed than *BeePCo-with-rotation*.

Figure 3 shows the percentage of the area coverage with all four techniques. The maximum possible area coverage is shown by *MaxCo*. Among the four techniques, *StiCo* provides the lowest percentage of covered area, whereas *BeePCo* outperforms and achieves the highest percentage of covered area. The difference between *BeePCo-with-rotation* and *HybaCo* is not significant, although *HybaCo* distributes robots more uniformly.

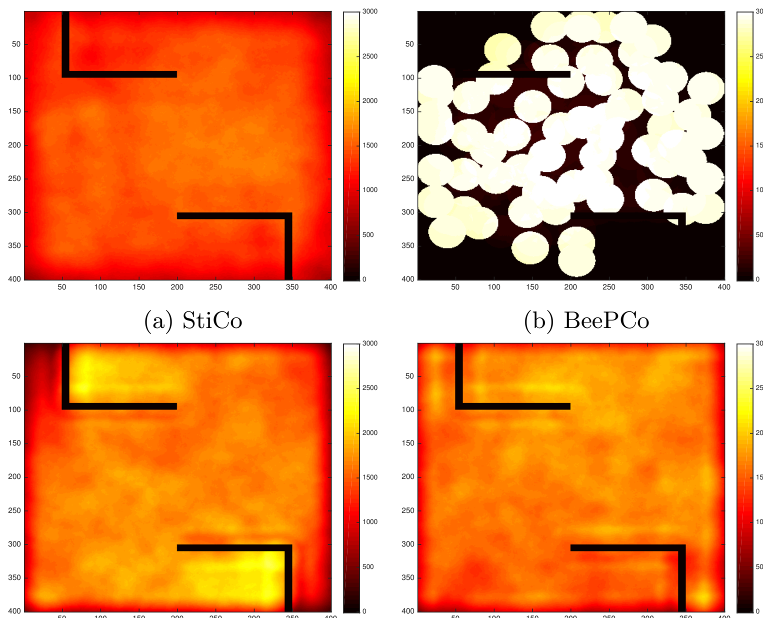

(a) StiCo (b) BeePCo

(c) BeePCo-with-rotation (d) HybaCo

Fig. 4: The distribution of robots in the arena using an MRS of 60 robots on BeePCo, StiCo, BeePCo-with-rotation and HybaCo.

Figures 4 and 5 show experimental results on an MRS of 60 robots in the arena which contains L-shaped obstacles (e.g., separators or boxes in a room). Figure 4 illustrates the distribution of the robots in this environment on *StiCo, BeePCo, BeePCo-with-rotation* and *HybaCo*. *StiCo* performs uniformly distributed coverage, although the corners of the arena are slightly less covered, whereas the robots had no issues getting around the L-shaped obstacles, nor providing coverage around the obstacles. In *BeePCo* approach, the robots stop spreading after communication links with the other robots are broken because they are outside of the inter-robot transmission range. The white circles in Fig. 4b show the robots' non-overlapping coverate within the experimental arena. Because the robots do not move for a long time, the percentage of area coverage in

BeePCo is higher and more unevenly distributed, as opposed to the other three techniques. Figure 4c and 4d illustrate that *BeePCo-with-rotation* and *HybaCo* have a more uniform robot distribution over the environment in comparison to *StiCo* and *BeePCo*. The distribution of the robots is remarkably high around the obstacles in *BeePCo-with-rotation* which shows that robots struggle to get away from the obstacles. This situation is not applicable in *HybaCo*, therefore, *HybaCo* shows more uniformly distributed robots, as Fig. 4d illustrates.

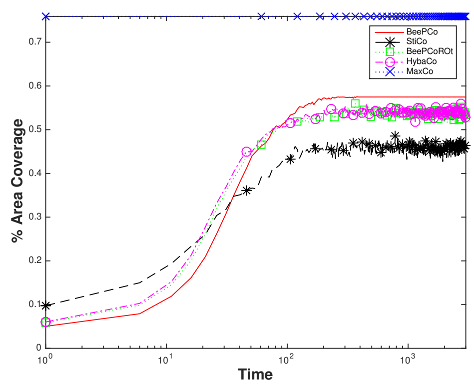

Fig. 5: The percentage of area coverage using MRS of 60 robots on different techniques.

Furthermore, Fig. 5 shows the percentage of the covered area in the environment containing L-shaped obstacles. The maximum possible area coverage of 60 robots is represented by *MaxCo*, which is $\approx 75\%$. *StiCo* achieves the least area coverage, whilst *BeePCo* gives the most among these four techniques, although it does not allow robots to move any further once the transmission connection between robots has stopped. The difference in the percentage of the covered area is very small between *BeePCo-with-rotation* and *HybaCo*, as the figure shows. In terms of the time the algorithms take to distribute robots in the period between 10^0 and 10^1 seconds, *StiCo* initially scatters the robots faster than *BeePCo* and converges faster, unlike our expectations, because *StiCo* has a more gradual manner of moving outwards, i.e., circling, whereas in *BeePCo* robots move outwards in a direct line. Later, in the time period between 10^1 and 10^2 seconds, *BeePCo* robots benefit from the direct communication exchange and spread out much faster than *StiCo*, *BeePCo-with-rotation* or *HybaCo*.

Figures 6 and 7 illustrate an MRS of 60 robots on an arena representing floor plans. In Fig. 6, heatmap images represent the distribution of the robots using *StiCo, BeePCo, BeePCo-with-rotation* and *HybaCo* approaches individually. The bottom left corner represents a room where all four sides of the room are surrounded by walls, apart from a small doorway gap close to the middle of the arena. Although *StiCo* and *BeePCo* have some coverage in this room, both *BeePCo-with-rotation* and *HybaCo* provide better coverage and improve the robot distribution as well. The top left corner in the arena is the least cov-

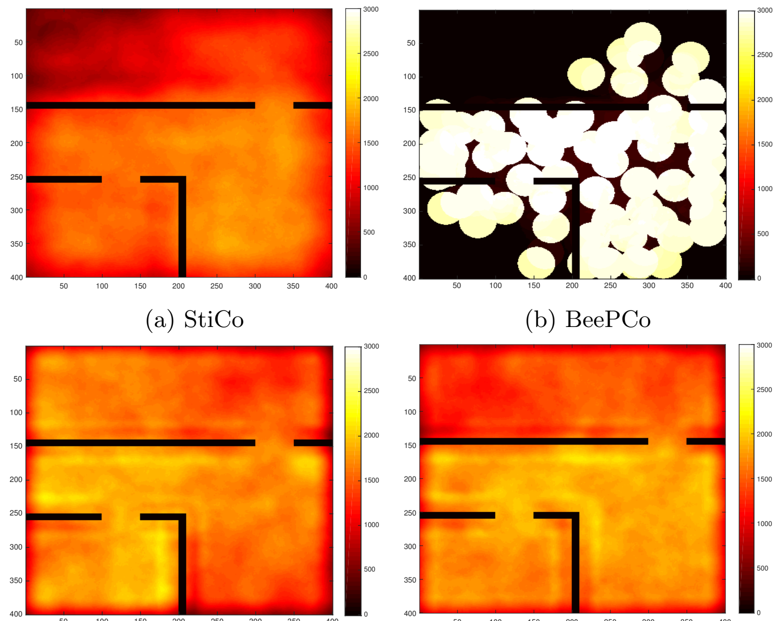

(a) StiCo (b) BeePCo

(c) BeePCo-with-rotation (d) HybaCo

Fig. 6: The distribution of robots in the arena using an MRS of 60 robots on BeePCo, StiCo, BeePCo-with-rotation and HybaCo.

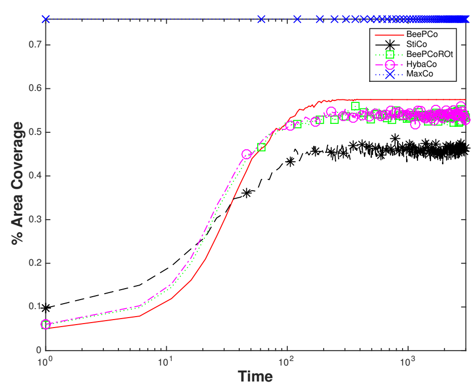

Fig. 7: The percentage of area coverage using MRS of 60 robots on different techniques.

ered area for all four approaches due to the long wall stretching across the arena from side-to-side, apart from the little doorway gap. Because the doorway is small, robots struggle to find the gap and pass through the wall, especially covering the upper left corner. *BeePCo* has no coverage on the left corner above the wall, whereas *StiCo* has brief coverage of the same corner. *BeePCo-with-*

rotation and *HybaCo* visibly improve coverage of the same corner as a result of the advantages of the combined approaches. Similar to Figs. 2 and 4, *BeePCo-with-rotation* has more coverage around the walls, whereas *HybaCo* uniformly improves the coverage of the arena, mostly beneath the long wall.

In terms of the percentage of the area coverage, as shown in Fig. 7, *BeePCo* outperforms in this set of experiments too. The maximum possible area coverage of 60 robots is represented by *MaxCo*, which is $\approx 75\%$, similar to the L-shaped arena. The difference in the percentage of the area coverage in *BeePCo-with-rotation* and *HybaCo* is very small, whereas *StiCo* is considerably lower than the other three techniques. In terms of the time for the robots to converge, *StiCo* takes the longest to stabilise, whereas *BeePCo* has the steepest hill in the time period between 10^1 and 10^2 seconds and stabilises fastest among the four techniques compared in this study. A significant performance improvement can be observed in *BeePCo-with-rotation* and *HybaCo* in comparison to *StiCo* and *BeePCo*. Specifically, the experiments show that merging the strengths of both *StiCo* and *BeePCo* leads to superior results with respect to uniform distribution of the robots in the arena.

6 Conclusions

This paper compares four social insect inspired multi-robot coverage approaches, namely stigmergic behaviour of ants (*StiCo*), the pheromone signalling process of bees (*BeePCo*), a derived method based on *BeePCo* (*BeePCo* rotation), and an ant-and-bee inspired hybrid approach (*HybaCo*) in realistic, complex environments. We have shown the performance of all four approaches with respect to a number of criteria, including area coverage, uniformity of distribution and speed of convergence, with a particular focus on our hybrid bee-and-ant inspired approach that merges the strengths of *StiCo* and *BeePCo* into one algorithm. The advantages and disadvantages of these two techniques have been highlighted. In the experimental analysis, we evaluated the effectiveness of the proposed hybrid bee-and-ant inspired approach, i.e. *HybaCo* in MRSs with 56 and 60 robots in a number of different complex environment each containing numerous static objects and reported our observations.

StiCo moves at all times and applies (almost) uniform coverage over the arena. *BeePCo* achieves a higher percentage of area coverage in comparison to *StiCo*; however, it produces non-uniform coverage because the robots stop moving when they step outside of each others' transmission range. The experimental results show that *BeePCo-with-rotation*, which is an extension of *BeePCo*, improves the distribution of the robots, but does not provide the same percentage of area coverage as BeePCo. Finally, our experiments confirm our earlier results [2] and show that *HybaCo* merges the strengths of the *StiCo* and *BeePCo* algorithms.

References

1. Alers, S., Hu, J.: Admoveo: A robotic platform for teaching creative programming to designers. In: Proc. of the 4th Int. Conf. on E-Learning and Games: Edutainment. pp. 410–421. Edutainment '09, Springer-Verlag, Berlin, Heidelberg (2009)
2. Broecker, B., Caliskanelli, I., Tuyls, K., Sklar, E., Hennes, D.: Social insect-inspired multi-robot coverage. In: Proc. of Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS) (2015)
3. Broecker, B., Caliskanelli, I., Tuyls, K., Sklar, E., Hennes, D.: Social insect-inspired multi-robot coverage. In: Proc. of the Workshop on Autonomous Robots and Multirobot Systems (ARMS) at AAMAS (2015)
4. Caliskanelli, I., Broecker, B., Tuyls, K.: Multi-robot coverage: A bee pheromone signalling approach. In: Int. Conf. on Artificial Life and Intelligent Agents (2014)
5. Caliskanelli, I., Harbin, J., Indrusiak, L. Polack, F., Mitchell, P., Chesmore, D.: Runtime optimisation in wsns for load balancing using pheromone signalling. In: NESEA, 2012. 3rd IEEE Int. Conf. on (Dec 2012)
6. Caliskanelli, I., Indrusiak, L.: Using mobile robotic agents to increase service availability and extend network lifetime on wireless sensor and robot networks. In: INDIN, 2014. 12th IEEE Int. Conf. on (July 2014)
7. Caliskanelli, I., Harbin, J., Soares Indrusiak, L., Mitchell, P., Polack, F., Chesmore, D.: Bio-inspired load balancing in large-scale wsns using pheromone signalling. Int. Journal of Distributed Sensor Networks (2013)
8. Di Caro, G., Dorigo, M.: Antnet: distributed stigmergetic control for communications networks. J. Artif. Int. Res. 9(1), 317–365 (Dec 1998)
9. Dorigo, M.: Optimization, Learning and Natural Algorithms. Thesis report, Politecnico di Milano, Italy (1992)
10. Dorigo, M., Birattari, M., Stutzle, T.: Ant colony optimization: Artificial ants as a computational intelligence technique. Computational Intelligence Magazine, IEEE 1(4), 28–39 (2006)
11. Dorigo, M., Roosevelt, A.F.: Swarm robotics. In: Special Issue", Autonomous Robots. Citeseer (2004)
12. Lemmens, N., De Jong, S., Tuyls, K., Nowé, A.: Bee behaviour in multi-agent systems. In: Adaptive Agents and Multi-Agent Systems III. Adaptation and Multi-Agent Learning, pp. 145–156. Springer (2008)
13. Ranjbar-Sahraei, B., Weiss, G., Nakisaee, A.: A multi-robot coverage approach based on stigmergic communication. In: Multiagent System Technologies, Lecture Notes in Computer Science, vol. 7598, pp. 126–138. Springer (2012)
14. Ranjbar-Sahraei, B., Weiss, G., Nakisaee, A.: Stigmergic coverage algorithm for multi-robot systems (demonstration). In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3. pp. 1497–1498. International Foundation for Autonomous Agents and (2012)
15. Ranjbar-Sahraei, B., Weiss, G., Tuyls, K.: A macroscopic model for multi-robot stigmergic coverage. In: Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems. pp. 1233–1234. International Foundation for Autonomous Agents and Multiagent Systems (2013)
16. Şahin, E.: Swarm robotics: From sources of inspiration to domains of application. In: Swarm robotics, pp. 10–20. Springer (2005)
17. Sahraei, B.R., Alers, S., Tuyls, K., Weiss, G.: Stico in action. In: Int.l Conf. on Autonomous Agents and Multi-Agent Systems, AAMAS '13, Saint Paul, MN, USA, May 6-10, 2013. pp. 1403–1404 (2013)