

From Ants to Birds: A Novel Bio-Inspired Approach to Online Area Coverage

Luca Giuggioli, Idan Arye, Alexandro Heiblum Robles, Gal A. Kaminka

Abstract Online coverage path planning is a canonical multi-robot task, where the objective is to minimize the time it takes for robots to visit every point in an unknown area. Two general major approaches have been explored in the literature: a stigmergic approach, inspired by *ant behavior*, relies on active marking of the environment. In contrast, the *collaborative approach* relies instead on localization, memory of positions, and global communications. In this paper, we report on a new approach, inspired by territorial *bird chirping*, which borrows from both previous approaches: it relies on localization and memory, but not on global communications. We provide a detailed analytic and empirical evaluation of this model.

1 Introduction and Background

Coverage path planning is a canonical robotics task, with many applications such as environmental monitoring, surveillance, exploration, and search [5]. In *online coverage*, one or more robots is to move inside an unknown target area, such that every point in the area is visited by one or more robots, often with the secondary objective

Luca Giuggioli

Bristol Centre for Complexity Sciences, Department of Engineering Mathematics and School of Biological Sciences, University of Bristol, UK e-mail: Luca.Giuggioli@bristol.ac.uk

Idan Arye

The MAVERICK Group, Computer Science Department, Bar Ilan University, Israel e-mail: idanarye@gmail.com

Alexandro Heiblum Robles

Department of Engineering Mathematics, University of Bristol, UK e-mail: Alex.Heiblum@bristol.ac.uk

Gal Kaminka

The MAVERICK Group, Computer Science Department, Bar Ilan University, Israel e-mail: galk@cs.biu.ac.il

of minimizing the time for such full coverage [8]. For multiple robots, two major approaches emerge for online coverage: A *stigmergic* [25] *approach*, which relies on environmental marking by the robots to direct their motion towards uncovered territory, and a *collaborative approach*, which relies on global communications to have robots explicitly—and remotely—coordinate their actions.

The stigmergic approach is often thought to be inspired by *ants*, though it is used by some mammals as well [7], e.g., foxes [20]. Here, robots mark visited points as they move around, simultaneously reading previously left markings. Robots move away from points marked by others [27, 10, 17]. This causes them to divide the environment into territories, each maintained by a single robot. Clear benefits of this approach include simplicity of the control algorithm (a random walk), and the fact there is no need for localization or memory of markings; robots use the markings themselves to identify locations visited by themselves or others. Unfortunately, coverage is often redundant, and relies strongly on the duration of the markers existence; moreover, building robots with actual marker reading and writing mechanisms is quite difficult in practice [10].

The collaborative approach is often associated with artificial methods (though it could just as easily be inspired by human teamwork). Here robots communicate with each other (in most studies, regardless of their distance), and divide up the area between them, e.g., [13, 11, 1]. The coverage time can be minimized using such collaborative algorithms. However, this approach requires not only localization (to identify current position) and navigation (to move to agreed-upon new locations), but also memory (to store visited locations and future positions), communications that allow task assignment, and most importantly, a shared coordinate system that allows a shared understanding of the regions to be divided between the robots. Satisfying these requirements in practice is again difficult. While memory and localization can perhaps be fairly easily had (some commercial vacuum cleaners now employ SLAM), establishing global communications, and a shared coordinate system in an unknown area, with unknown initial poses, is a difficult challenge [24, 18].

There have been surprisingly few attempts at addressing these challenges. Rekleities et al. have worked on coverage with range-limited communications [21], yet still rely on a shared coordinate systems. Batalin et al. have replaced the need for communications and localization with the need to sense others remotely, distinguishing robots from other objects in the environment [2]. Rutishauser et al. examine collaborative algorithms that display graceful degradation (to random motions) when positional, sensory, and communication failures accumulate [22], and are therefore less reliant on explicit collaboration. Durham et al. have discussed a related approach to ours, for offline coverage, whereby robots that meet exchange information by pairwise gossip communications, to statically partition a known area between them [4]. In contrast, our online coverage approach only requires robots to detect each other, but no information exchange is necessary.

In this paper we propose a novel online coverage approach that is inspired by the territorial behaviour of higher organisms in particular certain species of birds [12] and mammals [23]. The key idea is that when two robots meet, they detect each other (in birds, this is done by chirping a challenge which is then countered), then

remember the location of the encounter and treat it as a border landmark. Robots using this approach are assumed to have localization and memory, but do not need global communications or a shared coordinate system. Moreover, they utilize the simple-to-implement random walk algorithm for their motion. We provide a comprehensive mathematical and empirical analysis of this approach. Specifically, we analyze its characteristics and determine its efficiency in terms of coverage time as a function of its control parameters.

2 The Memory-Based Territorial Exclusion Model

We are given a set of N identical circular robots of radius r with a radially uniform detection distance $d \geq r$ that move in continuous space with speed of magnitude $|\vec{v}|$ within an arena of size A with periodic boundary conditions (toroidal domain). Every time a robot detects another, they both remember the location (in their own coordinate system), and move away from it. They mark the location in memory, and then use this to avoid the location if they run into it again. Thus upon detecting another robot or remembering a mark, the robots turn away.

The position of each robot is updated following Algorithm 1, which controls a single motion step of size d at most. First (line 2), a new motion vector $\vec{\ell}$ is generated by calling a correlated random walk procedure with parameters λ, ϑ described below (Section 2.1). Then, the robot considers whether another robot is detected along the motion vector, within the detection distance d (line 3). If so, the robot re-

members the location (line 4) and sets a revised, shortened motion vector to it, \vec{b} , reaching the detected robot or mark by not moving beyond them. Alternatively, if a previously marked location is retrieved from memory, it is used to set $\vec{\ell}$ (line 5–6). The robot then moves along $\vec{\ell}$. If it encounters a robot or remembers a mark in the new location, it repels towards the opposite direction, following the exclusion rule described in Section 2.2.

Algorithm 1 SINGLE-STEP

Require: Current position \vec{p}
 1: Randomly choose λ, ϑ
 2: $\vec{\ell} \leftarrow \text{RANDOMWALK}(\lambda, \vartheta)$
 3: **if** A robot is detected in position \vec{b} within d
 OR
 Remembered location \vec{b} is within d **then**
 4: $\vec{\ell} \leftarrow \vec{b}$
 5: Move to new position $\vec{\ell}$
 6: **If** robot or mark detected, REPEL.

2.1 Selecting a Motion Vector

The robots possess a degree of persistence in their motion and move as correlated random walkers (see e.g. [3]). At one extreme each new step is uncorrelated with the direction of the previous step and their movement is random. At the other ex-

tre the direction of movement is always identical to the one at the preceding step and their movement is ballistic. At each time step, each robot randomly selects a step size λ and an angle ϑ . The step size is sampled from an exponential distribution of mean size equal to half the size of the robot diameter. The new angular direction is selected relative to the previous movement direction. Except at the start of the simulations when the choice of angle is completely random, each robot selects from a distribution of turning angles $f(\theta)$. This distribution is a (symmetric) wrapped Cauchy distribution, that is a Cauchy distribution $\mathcal{C}(x) = \rho [\pi(\rho^2 + x^2)]^{-1}$, which is wrapped around the origin. For values between $-\pi \leq \theta \leq \pi$ [16].

The parameter ρ is called the concentration parameter and represents the mean cosine of the distribution $\langle \cos(\theta) \rangle = \int_{-\pi}^{\pi} \cos(\theta) f(\theta) d\theta = \rho$. In other words it indicates the ‘strength’ by which a robot would go forward at each step. Integration and inversion of $f(\theta)$ allows to sample from a uniform distribution U between 0 and 1 and obtain θ angles distributed according to $f(\theta)$ via $\theta = 2 \arctan \left\{ \frac{1-\rho}{1+\rho} \tan \left[\pi \left(U - \frac{1}{2} \right) \right] \right\}$.

In the limit $\rho \rightarrow 0$ the distribution $f(\theta)$ reduces to a uniform one between $-\pi$ and π , while in the limit $\rho \rightarrow 1$ the distribution $f(\theta)$ reduces to a Dirac delta distribution. In the former case, sampling turning angles from a uniform distribution means that the robot moves at random, whereas in the latter

limit the robot moves ballistically, i.e. always going straight (except upon encountering others). For intermediate values, the robot moves as a so-called correlated random walker with the degree of persistence determined by ρ . For computational simplicity the random step length for which a robot is initially selected to move is rounded down to its first integer. The call to Algorithm RANDOMWALK represents the computation of a new motion vector based on the sampled parameters λ, ϑ .

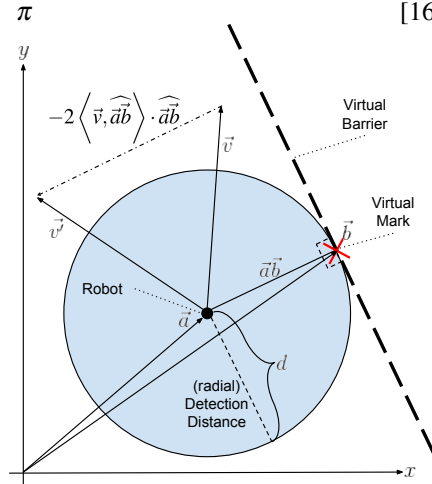


Fig. 1: Representation of the robot exclusion (repel) rule whereby one agent changes its direction upon encountering a virtual barrier.

2.2 Repelling Away from a Collision or a Remembered Mark

Two robots detect each other whenever their distance becomes equal to d . Upon detection the two robots will mark the positions where such detection has occurred and retreat from each other. The retreat represents an exclusion interaction at the

moment of detection, but the detection position is also a virtual mark is remembered by the robots for the duration T (called memory time). Neither of the two robots will cross the virtual mark. As each robot may interact with any neighbour that comes within a detection distance, the number of active marks that a robot has with any other members of the population fluctuates over time and depends on the various parameters of the model.

Consider a robot R_1 executing Algorithm 1. Assume that it is in collision course with another robot, robot R_j , which is considered static initially. At the predicted location of collision, robot R_1 marks the terrain at location \vec{b} . Robot R_1 's location now gets updated by accounting for the presence of the virtual mark at \vec{b} with robot j fixed. We assume symmetrical detection, i.e., Robot R_j now also remember a virtual mark at \vec{b} , though in its own coordinate system.

The geometry of the collision is as follows. After the mark location at \vec{b} is established, robot 1 moves up to \vec{a} when it collides with the virtual mark at position \vec{b} . The mark now acts as a virtual barrier since robot 1's velocity gets reflected as if a barrier tangent to the robot's detection circle was present at \vec{b} . Formally the reflected velocity vector \vec{v}' changes from the old velocity \vec{v} through the relation

$$\vec{v}' = \vec{v} - 2 \left\langle \vec{v}, \widehat{\vec{a}\vec{b}} \right\rangle \cdot \widehat{\vec{a}\vec{b}}, \quad (1)$$

where $\widehat{\vec{a}\vec{b}}$ is the normalized vector from \vec{a} to \vec{b} and $\langle \vec{z}, \vec{z}' \rangle$ represents the scalar product between \vec{z} and \vec{z}' , that is the projection of \vec{z} along \vec{z}' . As the reflection of the velocity is performed only when the robot is moving towards the virtual barrier, the exclusion rule applies only when $\left\langle \vec{v}, \widehat{\vec{a}\vec{b}} \right\rangle > 0$. All of this computation is carried out by the REPEL algorithm.

3 Simulating the Behavior of Robotic Birds

The simulator was developed in Java using the MASON simulation framework (<http://cs.gmu.edu/eclab/projects/mason/>). The simulator uses MASON's visualization facilities to draw the moving robots. It portrays each robot as a small disk of radius r . The radius is drawn in a different color to indicate the current movement direction from the robot's centre to the radius tip. A larger circle around the robot's disk indicates its detection circle. For ease of visualization the detection circle's radius is half its true value, so that the touching of two circles indicates that a collision between two robots has occurred.

There are in total eight parameters in the model. Those that we have kept fixed are, in arbitrary units, the size of the toroidal arena $A = 100 \times 100$, the magnitude of the robots' speed $|\mathbf{v}| = 1$, their size (robot radius $r = 1$) and the mean of the distribution of step length $\lambda = 1$. We have changed the remaining four parameters

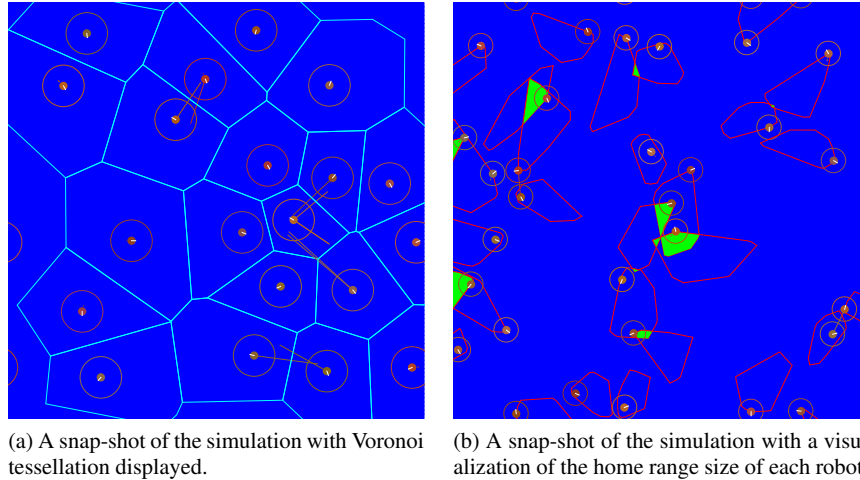


Fig. 2: Visualization of the simulator with Voronoi partitioning in panel (a) and home range size in panel (b). The circles around each robot represent half the size of the robot’s detection distance, while the long lines centered on a robot indicate the locations of the virtual marks. The direction of motion of each robot is shown by a small radial vector pointing outward from each robot. Although home ranges and Voronoi tessellation are correctly computed also for the agents that are close to the ‘edges’ of the toroidal domain, they are not visualized here. The home range size is computed with a minimum convex polygon estimate over 50 time steps. Overlaps between the estimated home ranges are colored in green.

consisting of the number of robots $1 \leq N \leq 100$, the detection distance $2 \leq d \leq 30$ (arb. units), the memory time $0 \leq T \leq 100$ (arb. units), and the random walk persistence $0 \leq \rho \leq 1$.

Additionally, the simulation draws lines from each robot to the virtual marks that were generated when other robots were detected. These lines remain visible for an amount of time T and show the mark locations where the robot may collide with. To get a perception of the global patterns we display at each time step the Voronoi tessellation [19] of the arena and the size and locations of the boundaries of the so-called agent home range (see e.g. [6]) that gives an indication of where a robot has been over a prescribed amount of time. Fig. 2 shows the Voronoi partitioning and the outer boundaries of the home ranges measured by computing the minimum convex polygon (MCP) (see e.g. [28]) of a robot’s localizations. Although other more sophisticated procedures exist (see e.g. [14]), MCP is sufficient for our purpose.

To mimic experimental conditions where no global clock exists that causes all the elements of the system to update their state at the same time [9], we use the following asynchronous update scheme. At the start of a simulation the N robots are numbered. At each time step robots are displaced sequentially from 1 to N . Within one simulation step, say robot R_1 is selected, then its position is updated

considering all other robots as static. When R_2 is then selected, R_1 's updated position is accounted for in moving R_2 , and so on with all other robots. A discrete simulation step is considered elapsed once all the N robot positions have been updated. This sequential scheme has also the advantage of reducing the computational costs of dealing with multi-agent detections if synchronous updating were to be used.

3.1 Measured outputs

For each parameter set of the model the measured outputs are obtained by running 1000 simulations starting from the same initial condition (the number of simulations for each figure is specified in the caption). The outputs that aim at giving an indication of the degree of spatial heterogeneity of the system are of two kinds. One kind is instantaneous and is obtained by observing the locations of each robot at a given time and averaging over simulation runs. The instantaneous measurements include the size of each Voronoi cell, the distance with respect to Voronoi neighbours, the number of active virtual marks both within each Voronoi cell and within Voronoi neighbours. The time-integrated measurements require accumulation of the data over each simulation run for a certain period of time and include the robot home range size, the spatial overlap and the coverage time. Outputs of the simulations are represented via the values of the mean, the standard deviation and the coefficient of variation (CV), that is the ratio of the standard deviation divided by the mean.

Whenever a simulation run starts, the initial directions of the robots are randomized and the robots are placed in an hexagonal pattern. To do so it calculates the maximal circle radius for packing N circle in the toroidal arena of size A , and chooses initial placements for the robots as if they were circles of that radius. After the initial placement, to 'thermalize' the initial configuration the simulation is run for a burning time $t_b = 100T$.

4 Results

Given our interest in proposing a new coverage algorithm, we have focused on the analysis of the memory-based territorial exclusion model rather than on a comparison between the various algorithms employing stigmergic or collaborative approaches. We do so by characterizing the heterogeneity of the spatial arrangement of the robots and computing coverage times.

The degree of spatial heterogeneity of the emerging spatial segregation patterns is obtained by studying the variability of the Voronoi partitioning. In Fig. 3 we display CV of the Voronoi tile size at time t_b as a function of detection radius and memory time. The CV of Voronoi partitioning is a measure of the relative strength of the fluctuations in the Voronoi cells among each robot and is mainly dependent on the detection distance d . For a given T the robots are allowed to wander more

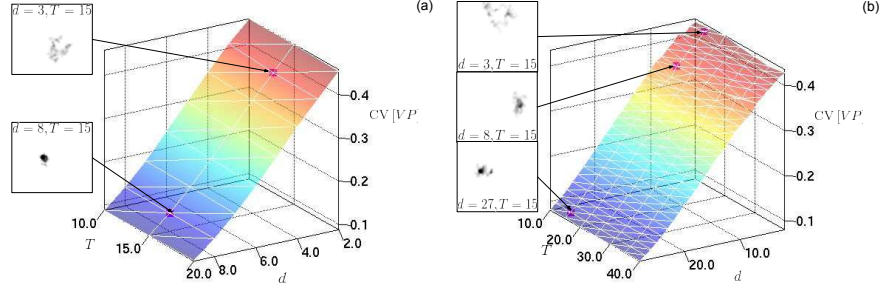


Fig. 3: Coefficient of variation of Voronoi cell size as a function of the memory time and detection distance for 100 robots in panel (a) and 10 robots animals in panel (b). Variability is evaluated through 1000 simulation runs for each parameter combination. Heat maps of one simulation run that indicate where one robot (selected at random) has been over 1000 time steps are also plotted for different combinations of values of d and T . The concentration parameter ρ is set to 0.

throughout space the smaller is d . Voronoi tiles with varying shape and size appear more frequently the smaller is the detection distance as indicated by the increase in CV while d is decreased in panel (a) and (b). The dependence of the spatial hetero-

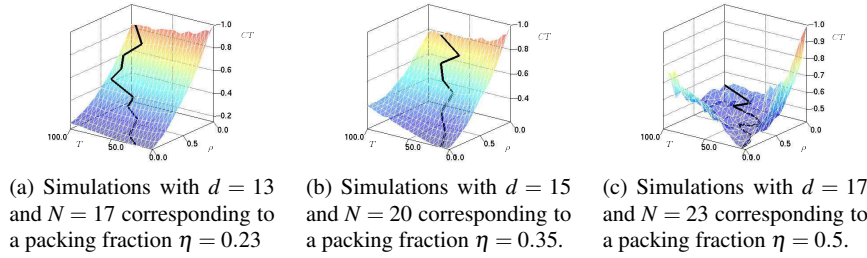


Fig. 4: Coverage time as a function of the concentration parameter ρ and memory time T for different choices of packing fraction η averaged over 1000 simulations. The black line in each panel tracks the minimum as a function of T for each different value of ρ . In computing the coverage time a grid of 10^4 rectangles is superimposed on the spatially continuous domain contained in A . Whenever the centroid of a robot is within a distance d from any point of a rectangle, that rectangle is considered covered. Once all rectangles have been covered the simulation runs are halted and the number of time steps are recorded.

genity of the robot position on T for a fixed d is in general very minor, as shown in panel (a) with 100 robots. However, this is not the case anymore when one uses a small number of robots for which a decrease in T reduces the movement constraints and allows for more variability between the positions of the robots (and the result-

ing Voronoi tiles) as displayed in the top right part of panel (b). The associated heat maps with $(d, T) = (3, 15)$ and $(d, T) = (8, 15)$ in panels (a) and (b), which represent the spatial occupation probability where a single agent have been, also gives some indication of a wider range and variability in the sizes of the Voronoi cells when a smaller number of robots is being deployed.

A way to measure the efficiency of the swarm algorithm is to estimate the coverage time (CT) of the domain A for different choice of detection distance d and number of robots N as a function of the concentration parameter ρ and the memory time T . We do so in Fig. 4 and for a better appreciation of the robot density we actually use the packing fraction η in place of N . As the robot collision distance is $d/2$ we define $\eta = \pi N(d/2)^2/A$, i.e. as the ratio between the maximum total area robots may occupy as they collide with each other, $\pi N(d/2)^2$, and the domain size A . For low packing fraction we notice a general decrease of coverage time towards higher persistent walk and lower memory time. We also notice that for a given fixed concentration ρ , the value of memory time T that minimizes the coverage time, drawn as a continuous black curve in all three panels, is neither random nor ballistic, but intermediate between the two, with larger values the smaller the memory time. Eventually for higher packing fraction, the coverage time develops a region in the $T - \rho$ parameter domain where minimization of the CT is possible. The transition to the appearance of a region of global minimization of the coverage time is smooth and is noticeable when $\eta \gtrsim 0.49$.

It is also of interest to know the dependence of the coverage time for a *given* memory time T as a function of ρ and the number of robots, which can also be evinced from Fig. 4. As shown in panel (a) and (b), when packing fraction is sufficiently low, the minimum coverage time is obtained with a ballistic walk. To show clearly this effect we plot in Fig. 5 the coverage time as a function N for different concentration ρ and we compare to the ‘perfect’ coverage algorithm whereby each location is visited only once by one robot and based on an initial configuration of the robots that gives the lowest possible coverage time. Considering the torus we con-




Fig. 5: Coverage time as a function of the number of robots averaged over 1000 simulations. From top to bottom the first ten curves have been drawn with values of ρ increasing by 0.1 starting from the random case at $\rho = 0$ to the ballistic case at $\rho = 1$. The detection distance is $d = 15$ and the memory time is $T = 20$. The bottom curve is the perfect algorithm. Due to optimal initial placement the perfect algorithm possess a zero coverage time beyond 14 robots.

struct the perfect algorithm by computing $CT = \sqrt{A + (2d)^2} \sqrt{A} / (2dN) - 2d$ where $\sqrt{A + (2d)^2}$ is the length that a robot needs to travel to wrap around the torus' edge before its original place and $\sqrt{A} / (2d)$ is the number of required trips to go to cover the available area for each robot. The final subtraction is included because robots stop before reaching their original location.

To understand better the role of the walk persistence and spatial constraints due to the collisions with other robots, in Fig. 6 we plot CT of a single robot with zero detection distance and zero radius in a circular geometry of radius R . The robot follows the movement rules as in the swarm and reflects the normal component of the velocity vector by colliding with the circular confining wall in the same way robots in the swarm get reflected upon encountering a virtual barrier. We study how the coverage time changes as a function of the robot directional persistence. We do so by plotting CT versus the dimensionless ratio $\zeta = -\lambda / [R \ln(\rho)]$, which represents an effective persistence parameter being the ratio between $-\lambda / \ln(\rho)$, the average distance a correlated random walker would move without turning [26], and the size of the confining domain R .

To measure the coverage time we partition the circular arena with a rectangular grid of $L \times L$ cells, and only the n cells whose centers are inside the arena are taken into account. The coverage time $C(n)$ is then defined as the time that takes a walker to visit all these n sites. We make $C(n)$ dimensionless by considering the ratio $\phi = C(n) / \tau_n$, where τ_n is the time it takes a random walker with no correlation ($\rho = 0$) to cover n distinct sites in open space.

5 Discussions

We have proposed a bio-inspired distributed spatial coverage algorithm that does not require robots to deposit 'marks' on the terrain. Rather than exploiting the stigmergic nature of scent-mediated territorial exclusion, we mimic a form of cognitive




Fig. 6: Dimensionless coverage time as a function of the inverse normalized persistence parameter ζ for a single robot in a circular arena averaged over 1000 simulations following the same movement and interaction (reflection) rules as those in the swarm. Random movement corresponds to $\zeta = 0$ ($\rho = 0$) and ballistic movement corresponds to $\zeta = \infty$ ($\rho = 1$). The circular arena and the speed of the agents were set at 1. From top to bottom the curves correspond to a choice of a square grid superimposed on the arena, respectively, of size $L = 1000, 316, 100, 31, \text{ and } 10$.

territorial behaviour whereby animals remember the locations where they exchange visual or auditory signals with neighbouring individuals to make each others' presence conspicuous. A robotic implementation of this behaviour consists of making each agent consider, for a finite amount of time, the locations of direct collision with or proximity to other robots as virtual barriers that cannot be encroached. This form of interaction generates a dynamical segregation that reduces spatial oversampling between the robots.

Two extreme regimes of spatial heterogeneity of the swarm are present depending on the number and detection distance of the robots. At one extreme, when packing fraction is sufficiently small, robots wander over all space with rare encounter occurrences and the system resembles a fluid-like material with homogeneous mixing. At the other extreme, when packing fraction is relatively high, robots are nearly jammed and the system may be highly heterogeneous. In some areas robots may move very little around their initial placements, while in others robots move very little around their initial placement.

While the analysis of a single robot moving within a confined arena showed that an intermediate degree of correlation minimizes coverage time, we found no evidence of such minimization in our robotic swarm except by increasing packing fraction (black line Fig. 4). For lower packing coverage minimization was achieved instead when robots moved ballistically with increasingly poorer performance as persistence was reduced (Fig. 5). We ascribe this difference to the fact that virtual barriers, when robot encounters are rare, do provide a degree of confinement but only partially. Although persistence reduces spatial oversampling in a single robot, it also diminishes the chance to revisit or return towards the area where a barrier was created thus preventing the robots to generate collectively long-lasting space partitioning.

As robot confinement is only partial, a natural measure to determine the size of the confining domains is necessary. Home range estimates, being based on an arbitrary integration time, do not offer a proper representation of the spatial confinement. As a consequence it is not evident how to estimate the ζ parameter regime in which robots operate. To explain the shift in optimal coverage between the ballistic and the correlated regime between low and high packing fractions future work should address the lack of a proper tool to estimate the size of the partial confinement.

Promising directions to improve the spatial coverage in the proposed algorithm include the choice of more informed alternatives for the movement paths of the robots, e.g. by avoiding recently visited locations, a variant of the so-called self-avoiding walk [15], or by systematic plowing of the emerging territory of each robot [1].

Finally, we would like to add that although we have left unexplored the impact of robot failures, the presence of obstacles and the confining geometry of a real arena, we do not expect qualitative changes in our findings because the movement response upon encountering an immobile robot, an obstacle or a reflecting wall would follow the same interaction mechanism that a robot undergoes when encountering a virtual mark. On the other hand, the same cannot be said about perception errors as these would affect the degree of confinement of the robots and their effects on spa-

tial coverage would need to be studied. For an empirical test careful considerations should also be made on the robot sensing mechanisms and how malfunctioning of the detection or recording systems would change the efficiency of the algorithm.

Acknowledgements LG thanks the support of the Bar Ilan Robotics Consortium (BIRC) and the office of the vice president during his stay at Bar Ilan University and acknowledges discussions with Adham Sabra and Alan Winfield. The research was supported in part by ISF grant #1511/12 and EPSRC grant EP/I013717/1. As always, thanks to K. Ushi.

Data Access Statement The Java code to run the stochastic simulations is openly available in the data.bris University of Bristol repository under DOI:10.5523/bris.i1rl4k2boj410h6ui4cpblfh.

References

- [1] Agmon N, Hazon N, Kaminka GA (2008) The giving tree: Constructing trees for efficient offline and online multi-robot coverage. *Annals of Mathematics and Artificial Intelligence* 52(2–4):143–168
- [2] Batalin MA, Sukhatme GS (2002) Spreading out: A local approach to multi-robot coverage. In: Asama H, Arai T, Fukuda T, Hasegawa T (eds) *Distributed Autonomous Robotic Systems 5*, Springer Japan, Tokyo, pp 373–382
- [3] Codling EA, Planck MJ, Benhamou S (2008) Random walk models in biology. *J Roy Soc Interface* 95(5):813–834
- [4] Durham JW, Carli R, Frasca P, Bullo F (2012) Discrete partitioning and coverage control for gossiping robots. *IEEE Transactions on Robotics* 28(2):364–378
- [5] Galceran E, Carreras M (2013) A survey on coverage path planning for robotics. *Robotics and Autonomous Systems* 61(12):1258 – 1276, DOI <http://dx.doi.org/10.1016/j.robot.2013.09.004>
- [6] Giuggioli L, Kenkre VM (2014) Consequences of animal interactions on their dynamics: emergence of home ranges and territoriality. *Mov Ecol* 2:1–20
- [7] Giuggioli L, Potts JR, Rubenstein DI, Levin SA (2013) Stigmergy, collective actions, and animal social spacing. *Proc Natl Acad Sci USA* 110(42):16,904–16,909
- [8] Hazon N, Kaminka G (2008) On redundancy, efficiency, and robustness in coverage for multiple robots. *Robotics and Autonomous Systems* 56(12):1102–1114
- [9] Huberman BA, Glance NS (1993) Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences* 90(16):7716–7718
- [10] Koenig S, Liu Y (2001) Terrain coverage with ant robots: a simulation study. In: *Autonomous Agents*, ACM, pp 600–607, DOI <http://doi.acm.org/10.1145/375735.376463>
- [11] Kong CS, Peng NA, Rekleitis I (2006) Distributed coverage with multi-robot system. In: *Proceedings of the 2006 IEEE International Conference on Robotics and Automation*

- [12] Krebs JR (1977) Song and territory in the great tit *parus major*. In: Stonehouse B, Perrins C (eds) *Evolutionary Ecology*, Macmillan Education UK, London, pp 47–62
- [13] Luo C, Yang SX, Stacey DA (2003) Real-time path planning with deadlock avoidance of multiple cleaning robots. In: *Robotics and Automation, 2003. Proceedings. ICRA '03. IEEE International Conference on*, vol 3, pp 4080–4085, DOI 10.1109/ROBOT.2003.1242224
- [14] Lyons AJ, Turner WC, Getz WM (2013) Home range plus: a space-time characterization of movement over real landscapes. *Movement Ecology* 1(1):1
- [15] Madras N, Slade G (2013) *The self-avoiding walk*. Springer Science & Business Media
- [16] Mardia KV, Jupp PE (2000) *Directional Statistics*. John Wiley, Chichester
- [17] Menezes R, Martins F, Vieira FE, Silva R, Braga M (2007) A model for terrain coverage inspired by ant's alarm pheromones. In: *Proceedings of the 2007 ACM Symposium on Applied Computing, SAC '07*, pp 728–732, DOI 10.1145/1244002.1244164
- [18] Nagavalli S, Lybarger A, Luo L, Chakraborty N, Sycara K (2014) Aligning coordinate frames in multi-robot systems with relative sensing information. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, pp 388–395
- [19] Okabe A, Boots B, Sugihara K, Chiu SN (2000) *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, 2nd edn. John Wiley
- [20] Potts JR, Harris S, Giuggioli L (2013) Quantifying behavioral changes in territorial animals caused by sudden population declines. *Am Nat* 182(3):E73–82
- [21] Rekleitis I, New AP, Rankin ES, Choset H (2008) Efficient boustrophedon multi-robot coverage: an algorithmic approach. *Annals of Mathematics and Artificial Intelligence* 52(2):109–142, DOI 10.1007/s10472-009-9120-2
- [22] Rutishauser S, Correll N, Martinoli A (2009) Collaborative coverage using a swarm of networked miniature robots. *Robotics and Autonomous Systems* 57(5):517–525, URL <http://dx.doi.org/10.1016/j.robot.2008.10.023>
- [23] Shettleworth SJ (2010) *Cognition, Evolution, and Behavior*, 2nd edn. Oxford University Press
- [24] Suzuki I, Yamashita M (1997) Agreement on a common x-y coordinate system by a group of mobile robots. In: *In proceedings of the 1996 Dagstuhl Workshop on Intelligent Robots: Sensing, Modeling and Planning*, World Scientific Press, pp 305–321
- [25] Theraulaz G, Bonabeau E (1999) A brief history of stigmergy. *Artificial Life* 5(2):97–116
- [26] Viswanathan GM, Raposo EP, Bartumeus F, Catalan J, da Luz MGE (2005) Necessary criterion for distinguishing true superdiffusion from correlated random walk processes. *Phys Rev E* 72(1):011 111
- [27] Wagner IA, Lindenbaum M, Bruckstein AM (1999) Distributed covering by ant-robots using evaporating traces. *IEEE Transactions on Robotics and Automation* 15(5):918–933, DOI 10.1109/70.795795

- [28] White GC, Garrott RA (1990) Analysis of Wildlife Radio-Tracking Data. Academic Press, Sand Diego