

Random Walks in Swarm Robotics: An Experiment with Kilobots

Cristina Dimidov¹, Giuseppe Oriolo¹, and Vito Trianni²(✉)

¹ DIAG, Sapienza University of Rome, Rome, Italy
dimidovc@yahoo.com, oriolo@diag.uniroma1.it

² ISTC, National Research Council, Rome, Italy
vito.trianni@istc.cnr.it

Abstract. Random walks represent fundamental search strategies for both animal and robots, especially when there are no environmental cues that can drive motion, or when the cognitive abilities of the searching agent do not support complex localisation and mapping behaviours. In swarm robotics, random walks are basic building blocks for the individual behaviour and support the emergent collective pattern. However, there has been limited account for the correct parameterisation to be used in different search scenarios, and the relationship between search efficiency and information transfer within the swarm has been often overlooked. In this study, we analyse the efficiency of random walk patterns for a swarm of Kilobots searching a static target in two different environmental conditions entailing a bounded or an open space. We study the search efficiency and the ability to spread information within the swarm through numerical simulations and real robot experiments, and we determine what kind of random walk best fits each experimental scenario.

1 Introduction

Animal search patterns are often described as random walks to represent situations in which the location of the target is not a priori known and the individual cognitive limitations and ecological conditions do not support more complex search strategies [2, 4, 6]. Since the early studies of the botanist Robert Brown about the motion of pollen particles—usually referred to as Brownian motion—the theory about random walks has been thoroughly developed and applied in disparate contexts, from animal behaviour to human cognition [1, 10]. Among the several possible random walk variants, two classes have emerged as the most useful models for describing animal search patterns: correlated random walks (CRWs) and Lévy walks (LWs) [6, 22]. The former are characterised by some correlation between the orientation of consecutive movements, while the latter are characterised by several small displacements interleaved by long relocations. LWs, in particular, provide an optimal search behaviour in case of sparse targets, thanks to the long motion steps that allow to search in different areas without frequently passing over sites that have been already visited [20, 21]. For this reason, several experimental studies have been conducted to demonstrate Lévy

behaviour in Nature, sometimes generating strong controversies [7, 11]. Recently, even theoretical results about the superiority of Lévy strategies have been questioned, showing that the optimal strategy is highly influenced by the specific contingencies of the search problem, such as the existence of a direction bias [13].

In swarm robotics too, it is important to deliver efficient search strategies for a number of relevant tasks, from search & rescue to demining, from surveillance to space exploration. The basic assumptions of swarm robotics entail limited individual abilities (i.e., local sensing, low processing power) which do not support complex localisation and mapping. As a consequence, random walks represent fundamental building blocks for the individual behaviour, and are often used when no environmental/social cue can be exploited [17]. For instance, a biased random walk was exploited for the basic motion of robots in the context of a commercial pollination task [3]. In a resource exploitation task, robots were programmed to search for resources within a bounded space exploiting a simple random walk with obstacle avoidance [5]. In collective decision-making problems, robots used a CRW to move in the different parts of a bounded space to find available options and share information with other robots [14, 19]. In a collective foraging task, the parameter of a biased CRW have been tuned to maximise search efficiency [9]. In the context of consensus decisions, the effects of the diffusion coefficient on the collective dynamics and the robot interaction network have been investigated [18]. These are just a few recent studies that demonstrate the pervasiveness and relevance of random walks in swarm robotics research.

Despite the importance of random walks for the design of swarm robotics systems, there is no systematic study that can help in the choice of the correct type and parameterisation, to the best of our knowledge. In some cases, a LW is chosen on the basis of the theoretical results only [8, 16]. More often, the type and parameterisation of the random walk behaviour are chosen heuristically without an analysis of their effect. This may lead to lower performance as soon as the working conditions change (e.g., when the size of the working space is increased [5]). In other cases, parameters are tuned with some optimisation algorithm [9]. Additionally, the type and parameterisation of the random walk are often chosen in relation to the efficiency of the individual behaviour, and the effects they have on information sharing within the robot swarm are often overlooked [18]. Given that the way in which robots move determines also the probability of their encounters, there may be a trade-off between the ability to widely search the environment and the ability to share information with other robots, which sometime requires special reverse behaviours (i.e., homing to a central place [14]) to allow information sharing. Overall, we believe that there is a compelling need for a systematic study of random walks in swarm robotics, taking into account possible trade-offs between search efficiency and information sharing, and investigating scenarios relevant for a wide range of applications.

In this paper, we move a step in this direction by studying different types of random walk for two specific search problems involving a single static target in either a bounded or an open space. We develop our study with respect to a specific robotic platform—the Kilobot robot [15]—and consider its properties

and limitations in the performance analysis. We devise a random walk behaviour determined by the distribution of step lengths and of turning angles. The former is a Lévy distribution controlled by the parameter α , which allows one to obtain step-length distributions varying from Gaussian—resulting in a (correlated) random walk—to a power law—resulting in a LW. The distribution of turning angles corresponds to a wrapped Cauchy distribution controlled by the parameter ρ , which allows to obtain a distribution varying from uniform—resulting in an isotropic random walk—to a delta (i.e., ballistic motion, see Sect. 2). We implement the target search scenarios in both multiagent simulations and physical robots (see Sect. 3), and we measure the efficiency in searching the target and in sharing information about its discovery for a wide range of parameterisations and swarm sizes (see Sect. 4). We find that the search efficiency and the information transfer highly depend on the scenario and on the local robot density. These results can be exploited in future swarm robotics research to support a principled choice of the random walk behaviour to be used as building block for more complex problems (see Sect. 5).

2 Random Walk Models

In its simplest form, a random walk can be thought of as a sequence of straight movements and direction changes. For instance, the simplest isotropic random walk in one dimension has a fixed step-length δ and equal probability of moving left or right $p_l = p_r = 1/2$ at each time step τ . In these conditions, the resulting motion is diffusive with a diffusion constant $D = \delta^2/2\tau$ [6].

In two or more dimensions, changes of directions are not limited to left-right motion, but any possible orientation must be taken into account, leading to the need of introducing a probability distribution for the turning angles. An isotropic random walk in multiple dimensions is obtained with a uniform distribution, while a biased or correlated random walk is obtained if such distribution is not uniform. A biased random walk has a preferential (absolute) direction of motion that does not depend on the current orientation of the random walker, but may depend on its location in space (e.g., a drift from wind or currents). A CRW, instead, is characterised by a positive correlation between consecutive movements, that is, the walker with higher probability moves in a similar direction to the previous one. In this study, we focus on random walks in two dimensions, and the turning angle θ is chosen from a wrapped Cauchy distribution, characterised by the following PDF [2, 12]:

$$f_w(\theta; \mu, \rho) = \frac{1}{2\pi} \frac{1 - \rho^2}{1 + \rho^2 - 2\rho \cos(\theta - \mu)}, \quad 0 < \rho < 1, \quad (1)$$

where μ represents the average value of the distribution ($\mu = 0$ in this study), while ρ determines the distribution skewness: for $\rho = 0$ the distribution becomes uniform and provides no correlation between consecutive movements, while for $\rho = 1$ a Dirac distribution is obtained, corresponding to straight-line motion.

For what concerns the straight motion steps, these may be of fixed length or can vary according to a certain step-length distribution. A Brownian motion is characterised by a step-length distribution with finite second moment, so that the distribution tends to a normal distribution according to the central limit theorem. On the other hand, a LW is characterised by a step-length distribution that follows a power law:

$$P_\alpha(\delta) \sim \delta^{-(\alpha+1)}, \quad 0 < \alpha \leq 2 \quad (2)$$

so that long movements are performed with non-null probability, as the distribution is heavy-tailed [22]. The Lévy distribution can be defined in terms of its Fourier transformation:

$$F(k) = e^{-\beta|k|^\alpha}, \quad 0 < \alpha \leq 2, \quad (3)$$

where β is a scale parameter, and α is the parameter defining the power law scaling. In particular, for $\alpha = 2$ the distribution becomes Gaussian, while for $\alpha \rightarrow 0$ the random walk reduces to straight-line paths.

To summarise, a pure CRW can be obtained by controlling the distribution of turning angles through the parameter ρ , while having a Gaussian step-length distribution—corresponding to $\alpha = 2$ in (3). On the other hand, a pure LW can be obtained by controlling the step-length distribution through the parameter α , while having a uniform turning-angle distribution—corresponding to $\rho = 0$ in (1). As a special case, a Brownian motion is obtained when $\alpha = 2$ and $\rho = 0$. In this work, we study pure CRWs and LWs, and we also address a hybrid form that joins together the non-uniform distribution of turning angles and the heavy-tail distribution of step lengths [2]. In this way, both the properties of correlated movements and long relocations can be obtained at the same time, which may lead to improved search efficiency.

3 Experimental Setup

As already mentioned, the type and parameterisation of a random walk must be chosen on the basis of the search problem to be faced. In this study, we focus on the search of a single static target by a swarm of N robots, and we also allow robots to exchange information about the target discovery. The target is represented by a circle of radius $r_t = 10$ cm, and robots recognise the target when they pass over it. Robots can communicate with all neighbours within a radius $r_c = 10$ cm, and can exchange short messages to communicate whether the target was discovered (by themselves or by some teammate). Robots always move following the random walk behaviour, even after discovering the target: in this way, they can interact with other robots upon encounters.

We look for the effects of different parameterisation and swarm sizes on the search performance. We measure the efficiency of the individual search as the average *first passage time* t_f , measured as the average time taken by robots to pass over the target for the first time. The *fraction* of robots that individually discover the target is referred to as ϕ_t . We also measure the efficiency of

information sharing as the *convergence time* t_c , which is the time in which all robots have information about the target (e.g., an identification code). Robots can obtain such information directly from the target when passing over it, or indirectly from other robots that already have such information. An experiment is terminated after a fixed amount of time T or when a sufficient number of robots have discovered the target.

We consider two simple scenarios that represent two common situations in swarm robotics. The first scenario consists in a *bounded space* surrounded by walls (square area with side L) in which robots and target are uniformly distributed. Given that Kilobots are not able to perceive obstacles, they are let free to collide with walls and with each other, and eventually disentangle owing to the random turns determined by their random walk behaviour. As we will discuss, collisions with walls have a bearing on the efficiency of the random search. The second scenario consists in a *open space* with a central place (e.g., a “home” location) from which all robots start searching. The central place is globally known to the robots, thanks to some form of dead reckoning or by some globally visible cue (e.g., a light source for the Kilobots [15]). In this scenario, the target is placed at a fixed distance d_t from the central place, but robots can travel very long distances away from the central place because of the absence of boundaries. To avoid that robots overly depart from the central place, we implement a biased random walk, so that a motion step is performed towards the central place with probability $p_b = 0.2$, otherwise following the direction given by the wrapped Cauchy distribution. In this way, we provide an elegant mechanism to maintain the swarm bounded in the vicinity of the central place, with the possibility of tuning the probability p_b to vary the diffusion properties. As a final remark, note that the motion step is interrupted and a new random direction is chosen every time a robot passes over the target or the central place.

Multi-agent Simulations. In order to analyse the properties of several parameterisations for the random walk as well as several swarm sizes, we devised a simple multi-agent simulation that abstracts the physical details of the Kilobots and focuses on the collective motion pattern. This computational model has been devised to provide a link between abstract analytical models and the robotic implementation [18], and is particularly useful to test a wide range of parameters for the random walk and the swarm size (see Sect. 4). Our simulation treats agents as dimensionless particles that can move at constant speed $v = 1$ cm/s and can communicate with all neighbours within the radius r_c , much as the Kilobots can do. Differently from Kilobots, however, agents do not collide with each other and they can change direction instantaneously. Finally, agents can collide with walls—if present—so that the motion vector orthogonal to the wall is cancelled, and only the parallel component is executed. For the bounded space scenario, we use a square arena with side $L = 1$ m. For the open space scenario, we vary the target distance $d_t \in \{0.25, 0.5, 0.75, 1.0\}$ m. Each simulation is run for $T = 10^5$ steps with an integration time step $\Delta t = 0.5$ s, to grant enough time for the agents to individually find the target, hence having a sufficiently

large sample to compute the statistics of the first passage time t_f . Given the discrete-time simulation and the finite velocity of the agents, the random walk is implemented by letting agents move straight for a number of simulation steps drawn from the Lévy distribution of Eq. (3), with $\beta = 5$. Turning is instantaneous and the turning angle is directly sampled from the wrapped Cauchy distribution of Eq. (1), with $\mu = 0$.

Experiments with Kilobots. The experiments performed with the Kilobots were made only in the bounded space scenario.¹ We built a square arena with $L = 90$ cm and distributed the robots and the target uniformly. The target here is represented by an immobile Kilobot programmed to continuously and frequently broadcast a target identification code id_t . As soon as a robot enters the communication range of the target Kilobot and receives the id_t code for the first time, it stores a time stamp to be used for the average first passage time statistics. Additionally, robots communicate with each other and share information about the target discovery by broadcasting a discovery identification code id_d , either when a robot has discovered the target on its own or when a robot has received the information from another robot. In the latter case, robots just act as relay. Upon reception of id_d for the first time, a robot stores a time stamp to be used to compute the convergence time t_c . The experimental run is terminated when a sufficient number of robots has discovered the target (i.e., $\phi_t > 0.7$) or after one hour of experimentation (i.e., $T = 3600$ s).

The implementation of the random walk behaviour for the Kilobot is rather straightforward, with the main differences with respect to simulation given by collisions and finite angular velocity. The distribution function for CRW and LW have been implemented as custom library functions in C++ and then ported to AVR C—the Kilobot programming language. Given the discrete control step of the Kilobots, both straight motion and turning are obtained through counters drawn from the Lévy and wrapped Cauchy distribution, respectively.

4 Results

Simulations of the Bounded Space Scenario. We have performed a wide simulation study to analyse different random walk behaviours by varying the parameters $\rho \in [0, 0.9]$ and $\alpha \in [1.0, 2.0]$, and by testing the search efficiency for varying swarm size $N \in [10, 100]$. The parameter ρ spans the whole range of possible values for the wrapped Cauchy distribution, excluding only very high values $\rho > 0.9$. The parameter α varies in the range in which the Lévy distribution has a finite mean. For each experimental condition, we performed 100 independent runs. The statistics for the individual search efficiency are shown in Fig. 1A, D. Here, the heat map shows average statistics over different runs and

¹ Given the size of the robots and characteristics of the random walks, we have found particularly impractical to run experiments in the unbounded arena scenario due to the small robot arena available for experimentation with Kilobots.

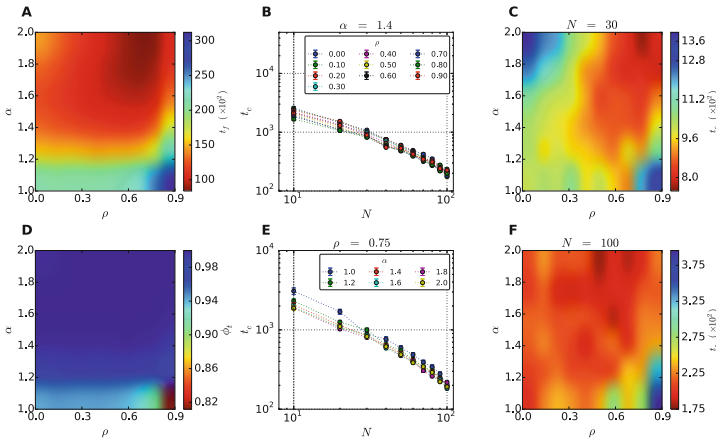


Fig. 1. Bounded space scenario. The search efficiency is represented by the average first passage time t_f (A) and the fraction of agents that discovered the target ϕ_t (D). The efficiency of information diffusion is represented by the convergence time t_c . The scaling for some fixed parameterisation is shown in panel B for $\alpha = 1.4$ and panel E for $\rho = 0.75$. The joint dependency on α and ρ is shown in panel C and F respectively for $N = 30$ and $N = 100$.

different swarm sizes. Indeed, the search efficiency is an individual measure that does not depend on the swarm size and robot density. By looking at the average first passage time t_f in Fig. 1A, it is possible to note that the search efficiency is higher when $\alpha \geq 1.8$ and $\rho = 0.7$, which corresponds to the areas in which t_f is minimised. This means that a pure CRW is sufficient to maximise individual search efficiency, and that the correlation coefficient ρ needs to be adequately tuned because an excessive persistence in maintaining a certain direction of motion is deleterious. Indeed, due to the absence of a collision avoidance strategy, agents that persist in moving in the same direction remain blocked against walls or corners for long time. For the same reason, the LW does not provide any advantage here, as the long relocations lead the agents to collide with walls for long time without any possibility to escape until a new direction away from the wall is chosen. The fraction of agents visiting the target ϕ_t is near 1 for much of the parameter space (see Fig. 1D). In particular we can see a slight influence on the results when $\alpha = 1$ confirming that for small Lévy exponents the agents collide with walls too often due to excessively long relocation steps.

The above measures are mainly related to the individual search behaviour, and do not depend on the swarm size. The situation is different for the diffusion of information, which instead can be influenced by the number of agents in the swarm. Within a bounded space, the swarm size determines also the agent density, so that higher densities correspond to more chances of agents to interact with each other. This is reflected in the convergence time, which decreases with increasing swarm size, as shown in Fig. 1B, E where we plot the convergence time for some specific values of α and ρ . Both parameters have a minor impact on the convergence time, which is mainly dependent on the swarm size N .

To appreciate the effect of the random walk parameters on the convergence time, we have plotted t_c for specific values of the swarm size (see Fig. 1C, F). It is evident that the convergence time is higher either when the agents perform a very localised random walk (i.e., a Brownian motion), especially for small swarm size, as shown in the top left corner of Fig. 1C or when the agents perform almost a ballistic motion (bottom-right corner of Fig. 1C, F). Instead, the smallest values for t_f coincide with the region of highest efficiency, revealing that a good search behaviour also corresponds to a good behaviour for information diffusion, as the agents meet more often if they efficiently move within the bounded arena.

Simulations of the Open Space Scenario. In an open space, agents can move large distances away from the central place. The biased random walk ensures that the search remains somehow constrained around the central place. Nevertheless, it can happen that the target is never found by any agent, or that the system does not converge. To evaluate the performance of the system, we perform 100 runs for each experimental condition obtained by varying the parameters as for the bounded space scenario, and we report the statistics for those parameterisation in which at least 75 % of the runs were successful. We report here a thorough analysis for a fixed target at distance $d_t = 0.5$ m, while other distances are discussed in less detail later on. Looking at Fig. 2A, we note that the parameter α controlling the Lévy distribution has the most influence on the search efficiency, while ρ as a smaller effect. The mean first passage time t_f is minimised for $\alpha \approx 1.4$ and $\rho \geq 0.75$. This parameterisation corresponds to a random walk characterised by both highly correlated movements and long relocations, which proves useful to explore widely the space around the central place. The fraction of agents visiting the target is highest for similar parameterisation (see Fig. 2D).

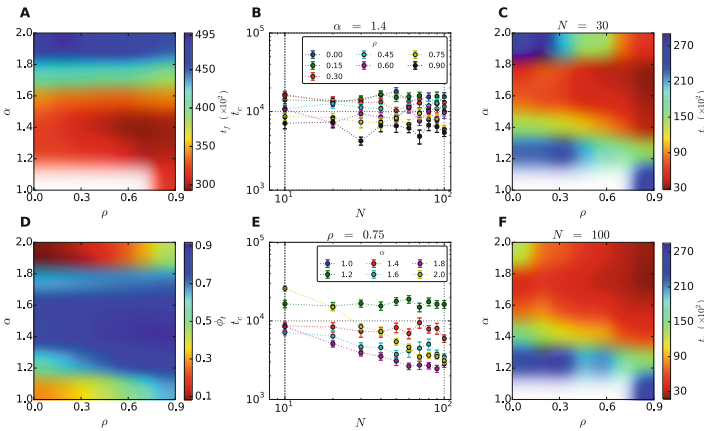


Fig. 2. Open space scenario. Results for a target at distance $d_t = 0.5$ m. See the caption of Fig. 1 for a description of the different panels.

We have also analysed the effects of the different parameterisations on the convergence time t_c . By varying the Lévy exponent α and fixing the correlation to the best-efficiency value ($\rho = 0.75$), we can appreciate how the long relocations of the LW affect the convergence time (see Fig. 2B). When the behaviour is close to a CRW ($\alpha \geq 1.6$), the convergence time decreases with the swarm size N . Indeed, without long relocations the agents remain close to the central place, hence the convergence time is mainly determined by the local density of agents, which in turn depends on N . On the other hand, when the random walk has the typical characteristics of a LW ($\alpha < 1.6$), the influence of the swarm size N on the convergence time vanishes. In this case, the long relocations coupled with the biased random walk allow agents to search widely the arena and to interact with each other in a way that is not dependent on the local density, leading to a very scalable behaviour. This is confirmed by the analysis of t_c for a fixed Lévy parameter $\alpha = 1.4$ (see Fig. 2E), which highlights that this scalability of the convergence time is determined by the feature of the step-length distribution, and not by the correlated movements, because for any value of ρ the convergence time appears to be independent from the swarm size N . If we look at the combined effects of α and ρ on t_c for a fixed swarm size (Fig. 2C, F), we can see a faster convergence for high values of both α and ρ . This means that the scalable behaviour offered by the LW does not minimise convergence time, but a good trade-off can be found to obtain at the same time scalability and fast convergence (e.g., $\alpha = 1.4$ and $\rho = 0.9$).

The other cases analysed in the open space scenario validate the analysis exposed above. When the distance of the target from the central place is small ($d_t = 0.25$ m, Fig. 3A), the search efficiency is maximised by a random walk with few long relocations (high values for α) and high persistence of movements

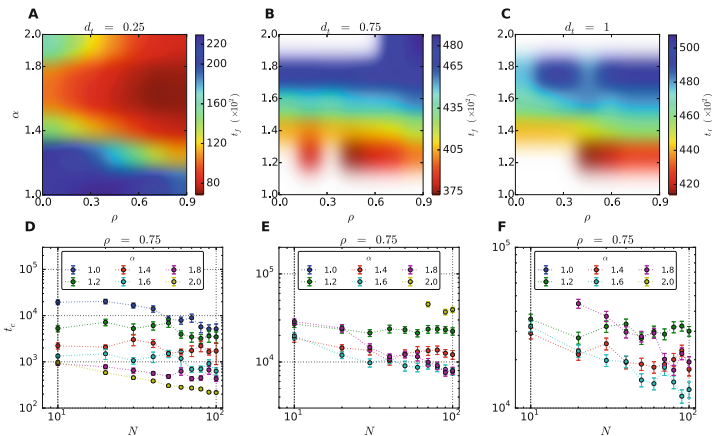


Fig. 3. Average first passage time t_f (top) and convergence time t_c (bottom) in the open space scenario for varying target distance (A, D: $d_t = 0.25$ m; B, E: $d_t = 0.75$ m; C, F: $d_t = 1.0$ m).

(high values for ρ). With a very close target, a localised random walk is more efficient because long relocations are not necessary, similarly to the bounded space scenario. With increasing distance (see Fig. 3B, C), ρ loses its influence and the search efficiency is mainly determined by the extent to which agents are able to perform long relocations.

The convergence time presents similar dynamics to the case with target distance $d_t = 0.5$ m (see Fig. 3D, F). When the target is very close, there is a generalised tendency to have slightly lower values when N increases. This is because the target is quickly found and hence the information spreads faster for larger swarms as the agents return to the central place. We note also that the parameter α has a clear influence on the convergence time, with higher values corresponding to faster convergence. When $d_t \geq 0.75$ instead, high values of α lead to a low success rate (hence data are not plotted). Here, too, smaller values of α lead to a scalable behaviour with respect to the swarm size N .

Results with Kilobots. The experiments with Kilobots have been performed with groups of 10 and 30 robots. Following the simulation results, the highest search efficiency in the bounded space scenario was obtained with a pure CRW. Consequently, we have decided to examine the cases in which $\alpha = 2$ and $\rho \in \{0.1, 0.5, 0.9\}$ to evaluate the effects of different CRWs on the search efficiency.

We perform 20 runs for each experimental condition, and we evaluate the performance of the system in case at least 75 % of the experiments were successful. The results obtained are in line with those obtained in simulation, with smaller t_f for larger ρ values (see Fig. 4). Differently from what observed in simulation, the high persistence provided by $\rho = 0.9$ does not lead to a lower efficiency. We argue that this is due to the effects of collisions and the time taken by Kilobots to turn on the spot, which result in a lower diffusion coefficient. This means that a higher persistence may be required to obtain a good search efficiency. Similarly for what concerns the convergence time t_c we observe a decrease with the swarm size N , with a lesser extent for $\rho = 0.9$. In this case, Kilobots perform a CRW with long straight paths, and consequently it takes time to get out from a collision condition among robots. This phenomenon has a strong influence on the convergence time t_c , as shown in Fig. 4.




Fig. 4. Average first passage time t_f (left) and convergence time t_c (right) in the bounded space scenario in the experiments with kilobots. Boxes represent the inter-quartile range from the median. Whiskers extend 1.5 times beyond the inter-quartile range. Empty circles mark the outliers.

5 Conclusions

The analysis of different random walks in a bounded space shows that a CRW with a relatively high persistence is the best strategy to adopt. The LW has worse performance due to the effect of collisions with walls (or other robots) that result in a poor overall performance. In the unbounded space scenario, instead, the best strategy is the LW, although some correlation in the movement provide some advantage as well. The distance between the target and the central place can have a high influence on the performance of the system. Additionally, the LW provides scalability to the system in terms of information diffusion, as the convergence time becomes independent of the swarm size, especially for the values of the Lévy exponent α that maximise efficiency. Hence, the Lévy strategies are advantageous because they provide a very scalable behaviour, even if at the cost of some performance. Notwithstanding the specificities of the Kilobot platform, these results can help in selecting the type of random walk to use in a swarm robotics context, on the basis of the particular experimental scenario to be tackled.

In future work, we plan to analyse other free parameters that have not been tested in this study (e.g., the amount of bias). Additionally, we plan to perform a systematic study with varying density of targets in order to analyse the system efficiency and the diffusion of information when there are more than one possible outcomes, in particular for what concerns the unbounded arena scenario. Finally, robots could exploit their communication abilities to provide locally some bias to the random walk performed by the neighbours. In this way, the transmission of information between robots could alter the random walk pattern and improve the search efficiency beyond the individual capabilities.

Acknowledgments. Vito Trianni acknowledges support from the project DICE (FP7 Marie Curie Career Integration Grant, ID: 631297).

References

1. Baronchelli, A., Radicchi, F.: Lévy flights in human behaviour and cognition. *Chaos Solitons Fractals* **56**(C), 101–105 (2013)
2. Bartumeus, F., da Luz, M., Viswanathan, G., Catalan, J.: Animal search strategies: a quantitative random-walk analysis. *Ecology* **86**(11), 3078–3087 (2005)
3. Berman, S., Kumar, V., Nagpal, R.: Design of control policies for spatially inhomogeneous robot swarms with application to commercial pollination. In: *Proceedings of the 2011 IEEE International Conference on Robotics and Automation (ICRA 2011)*, pp. 378–385. IEEE (2011)
4. Blackwell, P.: Random diffusion models for animal movement. *Ecol. Model.* **100**(1–3), 87–102 (1997)
5. Campo, A., Garnier, S., Dédriche, O., Zekkri, M., Dorigo, M.: Self-organized discrimination of resources. *PLoS ONE* **6**(5), e19888 (2011)
6. Codling, E., Plank, M., Benhamou, S.: Random walk models in biology. *J. Roy. Soc. Interface* **5**(25), 813–834 (2008)
7. Edwards, A., Phillips, R., Watkins, N., Freeman, M., Murphy, E., et al.: Revisiting Lévy flight search patterns of wandering albatrosses, bumblebees and deer. *Nature* **449**(7165), 1044–1048 (2007)

8. Fujisawa, R., Dobata, S.: Lévy walk enhances efficiency of group foraging in pheromone-communicating swarm robots. In: 2013 IEEE/SICE International Symposium on System Integration (SII), pp. 808–813. IEEE (2013)
9. Hecker, J., Moses, M.: Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intell.* **9**(1), 1–28 (2015)
10. Hills, T., Todd, P., Lazer, D., Redish, A., Couzin, I.: The cognitive search research group: exploration versus exploitation in space, mind, and society. *Trends Cogn. Sci.* **19**(1), 46–54 (2015)
11. Humphries, N., Queiroz, N., Dyer, J., Pade, N., Musyl, M., et al.: Environmental context explains Lévy and brownian movement patterns of marine predators. *Nature* **465**(7301), 1066–1069 (2010)
12. Kato, S., Jones, M.: An extended family of circular distributions related to wrapped Cauchy distributions via Brownian motion. *Bernoulli* **19**(1), 154–171 (2013)
13. Palyulin, V., Chechkin, A., Metzler, R.: Lévy flights do not always optimize random blind search for sparse targets. *Proc. Natl. Acad. Sci. U S A* **111**(8), 2931–2936 (2014)
14. Reina, A., Miletitch, R., Dorigo, M., Trianni, V.: A quantitative micro–macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intell.* **9**(2–3), 75–102 (2015)
15. Rubenstein, M., Ahler, C., Hoff, N., Cabrera, A., Nagpal, R.: Kilobot: a low cost robot with scalable operations designed for collective behaviors. *Robot. Auton. Syst.* **62**(7), 966–975 (2014)
16. Sutantyo, D., Kernbach, S., Levi, P., Nepomnyashchikh, V.: Multi-robot searching algorithm using Lévy flight and artificial potential field. In: IEEE International Workshop on Safety Security and Rescue Robotics (SSRR), pp. 1–6. IEEE (2010)
17. Trianni, V., Campo, A.: Fundamental collective behaviors in swarm robotics. In: Kacprzyk, J., Pedrycz, W. (eds.) *Springer Handbook of Computational Intelligence*, pp. 1377–1394. Springer, Berlin, Germany (2015)
18. Trianni, V., De Simone, D., Reina, A., Baronchelli, A.: Emergence of consensus in a multi-robot network: from abstract models to empirical validation. *IEEE Robot. Autom. Lett.* **1**(1), 348–353 (2016)
19. Valentini, G., Ferrante, E., Hamann, H., Dorigo, M.: Collective decision with 100 kilobots: speed versus accuracy in binary discrimination problems. *Agents Multi-Agent Syst.* **30**, 1–28 (2015)
20. Viswanathan, G., Afanasyev, V., Buldyrev, S., Havlin, S., da Luz, M., et al.: Lévy flights in random searches. *Phys. A: Stat. Mech. Appl.* **282**(1–2), 1–12 (2000)
21. Viswanathan, G., Raposo, E., da Luz, M.: Lévy flights and superdiffusion in the context of biological encounters and random searches. *Phys. Life Rev.* **5**(3), 133–150 (2008)
22. Zaburdaev, V., Denisov, S., Klafter, J.: Lévy walks. *Rev. Mod. Phys.* **87**(2), 483–530 (2015)