Université Libre de Bruxelles

Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle

The Hidden Benefits of Limited Communication and Slow Sensing in Collective Monitoring of Dynamic Environments

T. Aust, M. S. Talamali, M. Dorigo, H. Hamann, and A. Reina

IRIDIA – Technical Report Series

Technical Report No. TR/IRIDIA/2022-005

May 2022

IRIDIA – Technical Report Series
ISSN 1781-3794

Published by:
IRIDIA, *Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle*
Université Libre de Bruxelles
Av F. D. Roosevelt 50, CP 194/6
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2022-005

The Hidden Benefits of Limited Communication and Slow Sensing in Collective Monitoring of Dynamic Environments

Till Aust[1,2][0000−0003−2863−1341], Mohamed S. Talamali[3][0000−0002−2071−4030], Marco Dorigo[1][0000−0002−3971−0507], Heiko Hamann[2][0000−0002−2458−8289], and Andreagiovanni Reina[1][0000−0003−4745−992X]

[1] IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
[2] Institute of Computer Engineering, University of Lübeck, Lübeck, Germany
[3] Sheffield Hallam University, Sheffield, UK
andreagiovanni.reina@ulb.be

Abstract. Most of our experiences and also our intuition is usually built on a linear understanding of systems and processes. Complex systems in general and more specifically swarm robotics in this context leverage non-linear effects to self-organise and to ensure that 'more is different'. In previous work the non-linear and therefore counter-intuitive effect of 'less is more' was shown for a site-selection swarm scenario. Although it seems intuitive that being able to communicate over longer distances should be beneficial, swarms were found to sometimes profit from communication limitations. Here, we built on this work and show the same effect for the collective perception scenario in a dynamic environment. We found an additional effect of 'slower is faster'. In certain situations, swarms benefit from sampling their environment less frequently. All our work is based on simulations using the ARGoS simulator extended with a simulator of the smart environment for the Kilobot robot called Kilogrid. Our findings are supported by an intensive empirical approach and a mean-field model. Both effects seem important for designing swarms.

1 Introduction

In our recent research about information spreading in groups of individuals [30], we discovered a counter-intuitive mechanism by which reducing interactions between the individuals makes the group more capable to adopt new better opinions. This effect, that we call *less is more*, manifests when groups need to make consensus decisions and individuals follow a relatively simple voting behaviour. Such conditions can be particularly relevant for the design of algorithms for swarms of minimalistic robots that make best-of-n decisions [30]. Such algorithms are based on opinion dynamics models, in which every robot has an opinion about the option it currently considers the best (among n alternatives) and sends messages to neighbouring robots to recruit them on that option [33].

In this study, we confirm the generality of our previous finding [30] by reproducing the *less is more* effect (LIME) in a different scenario: collective perception of a predominant environmental feature in a dynamic environment. Studying this new scenario allows us to control the speed through which recruitment of new robots takes place, that is the key parameter to trigger the LIME. Controlling this parameter was not possible in the previous study of Talamali et al. [30] as the recruitment speed was constrained by robot travel times to specific locations. These times also had high variance and depended, for example, on traffic congestion and robot density. This scenario allows for a much cleaner analysis. Our results confirm and clarify the mechanisms. More importantly, we found a new surprising effect that was not reported earlier in this type of systems: the *slower is faster* effect [8, 24, 27, 26, 17]. To adapt faster, recruited robots must be slower in disseminating their opinions and recruiting other robots. This is a second surprising and counter-intuitive mechanism of this simple voting system. The slower is faster effect occurs when individuals are sparsely connected and make noisy estimates—two conditions commonly found in swarm robotics [9].

With this paper, we also release open-source code [2] supporting realistic simulations of the Kilogrid platform [1] (technology for Kilobots [20] to operate in smart environments) in ARGoS [14]. This simulation code, combined with the ARGoS Kilobot plugin [13], allows the use of identical code in simulation and reality (both for Kilobots and Kilogrid). Despite the limited adoption of Kilogrid in other research labs than IRIDIA (ULB), we believe that supporting realistic physics-based simulations can help spreading the technology and encourage collaborations between laboratories with and without such equipment.

2 Collective perception in a dynamic environment

The task of the robot swarm is to make a consensus decision in favour of the predominant element of the environment [32]. We assume that the robots can individually estimate each element concentration (i.e., the relative amount of the element present in the environment) to form their opinion which they share with each other. While individual estimates are noisy, the swarm collectively filters noise and converges to an accurate collective decision [32]. Individual estimation errors can be caused, for example, by simple error-prone sensing devices (readings distant from the ground truth, e.g. [12, 10]), spatial correlations (clustered information in localised areas rather than uniformly in the environment, e.g., [3, 4, 29]), and limited sensing range. In simulations we control sources and levels of sampling errors. Our analysis allows us to disentangle the impact of sampling errors (i.e., noise) from other system dynamics of interest (e.g., recruitment times).

We conveniently model the collective perception problem similarly as done previously [32]. The to-be-estimated environmental feature is the predominant colour of the ground which is comprised of squared tiles ($5\,\mathrm{cm}^2$). We consider tiles with $n = 2$ colours: blue and yellow (see Fig. 1a)[4]. The difficulty of the perception problem $\kappa \in [0, 1]$ is determined by the ratio between the concentration of tiles in the two colours: $\kappa = q_b/q_y$ where q_b and q_y are the concentrations

[4] Our decision to pick the blue and yellow colours has not been random.

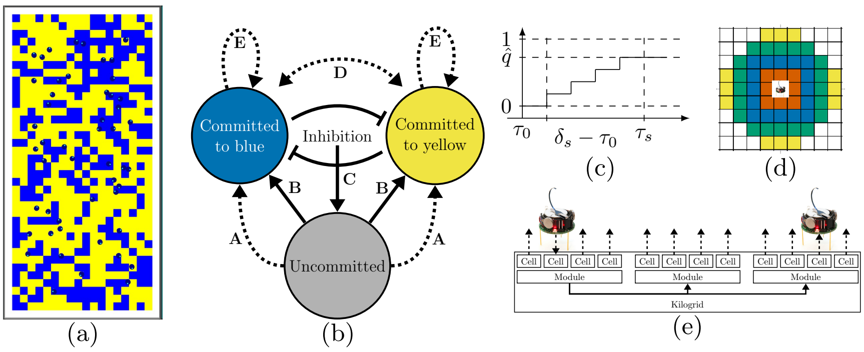

Fig. 1: (a) collective perception scenario for simulated Kilogrid in ARGoS (simulation code, see [2]), swarm of $N = 50$ Kilobots (small circles); (b) robots controlled by finite state machine with $n+1$ states (here $n = 2$), transitions: self-sourced (dashed arrows) or social (solid arrows) evidence; uncommitted→commit through (A) discovery or (B) recruitment; committed robots change state by: (C) cross-inhibition, (D) direct switching, or (E) stay; (c) once recruited, robots gradually increase com. probability w for sampling time $\tau_s = s\,\delta_s$, s samples every δ_s seconds; (d) robot-to-robot com. is virtualised using Kilogrid; (e) focal Kilogrid module receives Kilobot message and sends to all cells within com. range (prop. to param. c; com. range r_c, $c = \{2, 3, 4, 5\}$.

of blue and yellow tiles, respectively. Without loss of generality we assume that yellow is the predominant colour in all our experiments, $q_y > q_b$. The concentration of blue/yellow tiles corresponds to the number of blue/yellow tiles divided by the total number of tiles in the environment. The tiles are uniformly randomly distributed, hence reducing spatial correlations. However, spatial correlations are still present because of the unicolor tiles. Taking samples in a close proximity leads to biased measurements due to high correlations (see Sec. 5.2).

We consider a virtual dynamic scenario. We focus on a sudden change → of colours from blue to yellow that happens right at the moment when we start our simulations. The previously most frequent colour blue is reduced in concentration and yellow is the predominant colour after the change. All our experiments are initialised with the robot swarm in a state of full (100%) commitment in favour of blue, emulating a virtual history of the system, with every robot holding an estimate $q_b = 0.8$. We say that when starting the simulation a sudden environmental change has happened and the proportion of yellow tiles has just increased to $q_y > q_b$. In the next section, we describe how robots reassess the environment's state and reconsider their opinion.

3 A minimalist behaviour for a rich collective response

The robots have minimal requirements in terms of memory, computation, sensing, and communication capabilities. Compared with previous work that investigated decentralised consensus decision making in the collective perception scenario, our algorithm has the fewest requirements, in line with our quest for minimalism. Different from previous work that required the storage of all available alternatives and all received messages [22, 23, 4], here the robots only store the information about a single opinion (i.e., the colour considered predominant and the estimated concentration) the last received message from a neighbour, and a temporary variable to estimate possible environmental changes. Different from previous work requiring more advanced computation based on Bayesian inference [22, 23, 7, 6] or fusion operators from epistemic logic [4], our robot behaviour is defined by a small finite state machine with standard reactive transitions. Different from previous work that required sensors being capable of reading the predominant element at every measurement step [22, 23, 6], here the robot can only sense the presence or absence of one of the elements at a time. Different from previous work requiring maintenance of shared collective knowledge through rich inter-robot communication [28, 29], here robots send simple messages with a few bits of information, only indicating their preferred element (i.e., their chosen colour, for $n = 2$ that is one bit of information). Previous work in collective perception that is comparable to ours in its simplicity of individual robot requirements is that of Valentini et al. [32]. We extend previous analyses by considering a dynamic environment which has only been considered in few consensus decision making studies for the site selection scenario [15, 16, 25, 30], while here we consider the collective perception scenario.

Despite both, the minimalist robot control algorithm and the robots' noisy measurements, the swarm is able to collectively gather and process the data to conclude with an accurate consensus decision (picking the dominant colour). The control algorithm is based on simple reactive rules based on limited memory. The robot's algorithm can be described as four routines that are executed in parallel: motion, opinion update, sampling, and broadcasting.

The **motion routine** is independent from the other parts of the robot's behaviour. The robot's motion is neither influenced by its opinion nor by social or environmental inputs. The motion routine is a random walk implemented as a random waypoint mobility model [5, 30]. However, it could be substituted by any other algorithm implementing random diffusion. Using the random waypoint model, robots select random positions as their destinations. Once the destination is reached, robots select the next random destination. Robots avoid collisions with surrounding walls by selecting random destinations that are at least three robot-body lengths (approximately 10 cm) away from walls. As robot's motion is subject to noise, the robot can still approach walls. Once it gets at a distance smaller than three robot-body lengths from any wall, the robot starts a wall avoidance manoeuvre by rotating away from the wall and moving straight. The robots have no proximity sensing, therefore they do not implement any obstacle avoidance to prevent collisions with each other. To avoid robots remaining stuck in traffic jams caused by groups of robots moving in opposite directions (or robots not moving due to malfunctioning motors), robots select new random destinations if the previous destination was not reached within two minutes.

The **opinion update routine** is essential to solve the collective perception task because it determines how robots change their opinions and, hence, defines the collective behaviour. We require a large majority of the swarm in agreement for the predominant colour. While we present this routine for $n = 2$ colours, it does not require any changes to scale to numbers $n > 2$. Robots update their states every $\tau_u = 2\,\mathrm{s}$ following the cross-inhibition update shown in Fig. 1b. Robots can be in $n + 1$ possible opinion states, in the investigated case of $n = 2$ colours they are: committed to blue, committed to yellow, or uncommitted. Transitions between states are triggered by new self-sourced or social evidence. Self-sourced evidence (dashed arrows in Fig. 1b) is available when after a period of length τ_u the robot completed sampling a colour different from and better than its current opinion, that is, the robot has discovered in the last τ_u seconds that there is a colour more frequent than the colour of its current opinion. Social evidence (solid arrows in Fig. 1b) is available when after a period of length τ_u the robot received a message from a neighbour committed to a different colour (if multiple messages have been received, they have been overwritten and only the most recent stays in memory). If both self-sourced and social evidence are available, the robot randomly selects one of the two, discarding the other. The new evidence triggers a state change: (a) committed robots with new social evidence become uncommitted—a *cross-inhibition* transition; (b) committed robots with new self-sourced evidence become committed to the colour corresponding to the new evidence—a *discovery* transition; (c) uncommitted robots with any type of evidence, become committed as per the new evidence—a *recruitment* transition.

The **sampling routine** controls how information about the concentration of one element is collected from the environment. The robot continuously repeats sampling in cycles of collecting s samples. Each sample is a binary value indicating presence (1) or absence (0) of the environmental element of interest. Here, robots sample whether the colour at their position is of a given colour. The concentration estimate $\hat{q}_i$ is the proportion between the number of samples s_i^+ in which the element was present and the total number of samples s: $\hat{q}_i = s_i^+/s$. A new sampling cycle starts when the previous cycle has collected s samples, or when the robot changes opinion through social evidence. In the former case and when the robot becomes uncommitted, the robot determines which colour to sample randomly. Here, the robot selects the colour of the ground beneath itself. Instead when the robot becomes committed, it starts sampling the colour it has been recruited to. The random selection of the colour to sample allows the robot either to update the colour concentration estimate when it samples its commitment colour, or to gather potential self-sourced evidence when it samples a different colour. Instead, when a robot is recruited and takes a new opinion i, it immediately starts to sample i to obtain the information needed to regulate its messaging frequency (weighted voting, as described in the broadcasting routine). This means that once a robot is recruited and takes a new opinion i, it cannot instantaneously recruit other robots but a minimum amount of time is required to gather information about i first. The mathematical analysis of [30] showed that this temporal delay between change of opinion and recruitment of other robots—the sampling time τ_s—is the key mechanism that leads to the LIME. Therefore, the sampling time τ_s is the control parameter of this study and it corresponds to $\tau_s = s\ \delta_s$, where δ_s is the time between two samples. As analysed in Sec. 5.2, the sampling parameters s and δ_s are also linked to the estimation noise and have a determining impact on the collective dynamics.

The **broadcasting routine** implements a continuous 'narrowcast' of recruitment messages, that is, a broadcast to all robots within communication range r_c (i.e., neighbours). The robot scales its frequency of communication proportionally to the estimated concentration of the environmental element. The higher the estimated concentration of i is, the more recruitment messages for colour i the robot sends. The robot sends a message on average every τ_m/w seconds where $\tau_m = 0.5\,\mathrm{s}$ is the maximum communication frequency of our robots and $w = \min(2\hat{q}_i, 1)$ is the concentration weight for colour i. We multiply by two $(2\hat{q}_i)$ because we need to find the predominant element and any concentration $> 50\%$ represents the absolute majority. For lower concentrations, w scales linearly between 0 and 1. While in the case of $n = 2$ a concentration $< 50\%$ indicates predominance of the other colour, this does not generalise to $n > 2$ and therefore we do not consider this deductive mechanism. A newly recruited robot does not have a concentration estimate yet. It gradually increments its communication frequency as it collects samples (see Fig. 1c). We compute $w = \min(2\hat{q}_i, 1)$ using $\hat{q}_i = s_i^+/s$ even if the number of samples $s_i^+ < s$. This mechanism helps avoiding situations of vocal minorities: when a large proportion of the population changes their commitment, and only a small proportion of robots communicates, while the majority are silent. In our implementation, just-recruited robots are not silent, yet less vocal. Uncommitted robots do not communicate until they get recruited or make a discovery transition.

4 Simulated Kilobots and Kilogrid

We conduct our experiments using Kilobots which are cheap, simple, and small robots widely employed in swarm robotics [20, 21, 19, 9]. By regulating the frequency of two vibration motors, the Kilobots move on a flat surface at speeds of about $1\,\mathrm{cm/s}$ in roughly straight motion and rotate at the spot at about $45\,^\circ/\mathrm{s}$. The Kilobot has a diameter of 3.3 cm, can display its internal state through a coloured-LED, and can communicate with other robots and other devices through a infrared (IR) transceiver. The range of communication varies depending on lighting conditions and ground material [11]; in ideal conditions $r_c \approx 10\,\mathrm{cm}$. The Kilobot's control loop is executed at approximately 32 Hz.

Given these limited robot capabilities, researchers working with Kilobots have developed systems of augmented reality to allow Kilobots to interact with virtual environments [1, 18, 31]. We employ the Kilogrid system [31], which is a lattice of squared electronic modules covered with a transparent glass. The Kilobots can move on a glass surface while communicating with static modules beneath which are equipped with the same IR transceivers as Kilobots. Each module, sized $10 \times 10\,\mathrm{cm}^2$, is composed of smaller squared cells sized $5 \times 5\,\mathrm{cm}^2$

in a 2×2 configuration. In our setup, we use a Kilogrid composed of 10×20 modules for a total of 800 cells in a rectangular environment sized $1 \times 2\,\mathrm{m}^2$.

The collective perception scenario is implemented by assigning a colour to each internal Kilogrid cell. There are a total of 684 coloured tiles in our environment and 116 non-coloured tiles at the environment's boundaries. Cells adjacent to the walls are colourless because robots do wall avoidance when under two tiles away from any wall and should rarely visit these areas. All cells continuously signal their colour to human observers using coloured-LED and to the Kilobots through IR messages. Without loss of generality or simplicity, the Kilogrid provides more information to the Kilobots. Our idea is to improve the robots' movement which is subject to noise and unreliable [13]. The cell's IR messages contain, the colour, the cell's coordinates (x, y in the 20×40 Kilogrid's plane) and a wall flag. The coordinates are used to implement the above mentioned random waypoint mobility model [5, 30] to let robots effectively diffuse in space. The 0/1 wall flag indicates a wall at distance <10 cm and triggers wall avoidance.

The Kilogrid also allows extending the robot-to-robot communication range which is otherwise physically limited to $r_c \approx 10\,\mathrm{cm}$. Our robots communicate with each other via the Kilogrid. They send their IR messages to the cell beneath them. The cell sends the message to all the cells at an Euclidean distance $< c$ resulting in an effective range of $r_c \approx 2.5 + 5(c - 1)\,\mathrm{cm}$ (see Figs. 1d, e). Hence, we can test communication ranges beyond the Kilobot's limitations.

We do experiments in simulation only. There was already a dedicated ARGoS plugin to run accurate simulations with Kilobots [13], but there was no simulation software for the Kilogrid. Therefore, we implemented a (sub-)plugin for ARGoS to simulate the Kilogrid; another contributions of this paper (open-source code available at [2]). The Kilogrid is programmed via code executed on each module. To simulate the Kilogrid, we developed an ARGoS loop function that runs the control cycle of all Kilogrid modules in each simulation step. Module-to-module communication is done by CAN bus, module-to-robot through IR messages, and modules can send data to the PC control station (e.g., log files). Following the ARGoS paradigm of using identical code for simulations and real-world experiments, we developed a simulated module interface that provides all functions available on the real Kilogrid module controller. The code for simulated and real modules has only minimal differences (documented in the code repository) that have been included to optimise simulation speed.

5 Results: less is more & slower is faster

We test the ability of the robot swarm to adapt to sudden environmental changes. All robots start committed to blue ■ (predominant colour before the change) with a high estimate $q_b = 0.8$. We assume the change ■→■ happens right at the beginning of our experiment, which is initialised with an environment with more yellow tiles ■ than blue. The swarm is expected to perceive the change, reconsider its previous decision, and converge to a large majority (consensus decision) in favour of yellow. We consider the swarm capable to adapt to the change when over a 5 minutes interval the mean of the number of robots committed to yellow is greater than 70% of the swarm size. In this way, we avoid to count short-lived random fluctuations as successful adaptations. Instead we want the swarm to reach a stable majority. The adaptation time is measured as the moment when the subpopulation committed to yellow reaches 70% at the beginning of the 5 minutes. We run 30 simulations per condition in a variety of experimental conditions where we consider different communication ranges r_c, number of ground samples s collected during a sampling cycle, time between two samples δ_s, and problem difficulty κ.

5.1 When recruitment is slow do not be too social (less is more)

We fix sample number $s = 15$ and time between two samples $\delta_s = 4\,\mathrm{s}$, and test different communication ranges $r_c \in [2.5, 225]\,\mathrm{cm}$ for problem difficulties $\kappa \in \{0.7, 0.8, 0.9\}$. Hence, once recruited, robots broadcast with low probability the new colour until they complete the sampling cycle which lasts $\tau_s = s\,\delta_s = 60\,\mathrm{s}$ (see broadcast frequency diagram in Fig. 1c). Because the positive feedback (i.e., recruited robots recruit more) is slow, we expect to observe similar dynamics as reported in [30]. Fig. 2a shows that also here we have the LIME, where more social interactions (large r_c) diminish the swarm's ability to adapt. Therefore, we confirm the predictions of [30] and show this is a general effect that can take place in other scenarios than collective site selection, where it was found first.

This counter-intuitive effect can be explained via the social impact of committed subpopulations of unbalanced sizes. A large majority is able to repeatedly mute minorities that make temporary discoveries of alternative options. The minority's opinion is slow to gain traction in the population as new recruits are slow in becoming vocal and are quickly reverted to the majority's opinion. When the communication range is large, or equivalently when the robot density is high, any minority is in contact with the large majority at all time. Instead, sparse connectivity, due to a small communication range or a low robot density, reduces the importance of subpopulation sizes. Interactions are sporadic (often limited to pairs) and the collective dynamics is governed by opinion quality (encoded via messaging frequency).

Unlike the site selection scenario [30], where the positive feedback delay was hard to manipulate, here the delay consists in the sampling time $\tau_s = s\,\delta_s$ and can easily be studied. We investigate how the collective performance varies for different sampling times and for different levels of robot connectivity. We study sampling time by varying values s as well as δ_s and robot connectivity by varying both communication range and robot density (proportional to swarm size as environment size is constant). Figs 2b-d show that the LIME on robot connectivity lives on parameter regions of slow recruitment (top part of plots) and gradually vanishes when recruitment is quick. This result is in agreement with theory predictions, as quick recruitment enables positive feedback cascades and allows well connected swarms to react fast to environmental changes. While Figs. 2b-d only show results for problem difficulty $\kappa = 0.9$, we observed qualitatively equivalent dynamics for any κ tested.

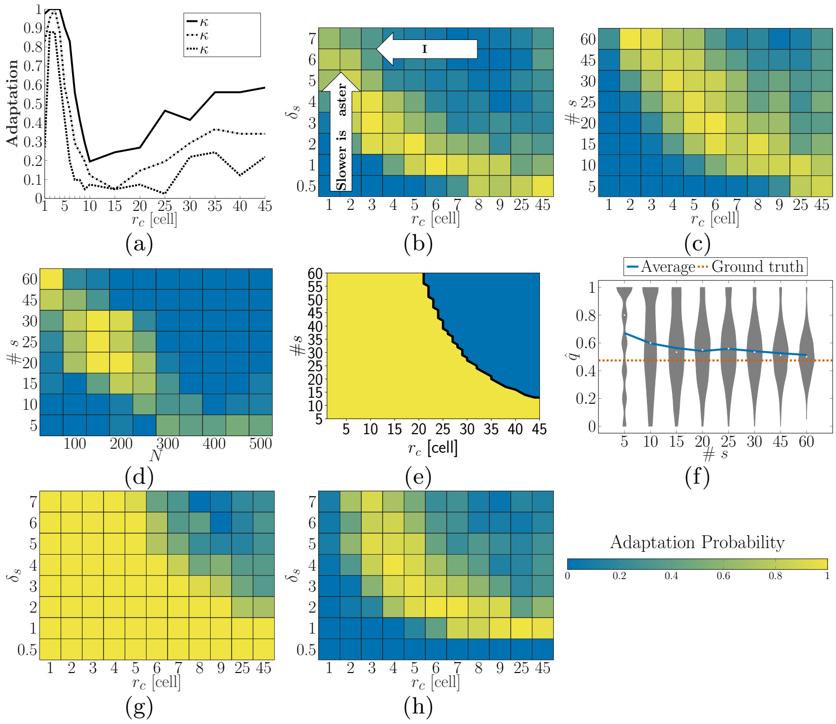

Fig. 2: (a,b,c) Less communication leads to higher adaptation probability. This effect happens only when recruitment is slow, i.e., in the top part of panels (b,c), for high sampling times $\tau_s = s\,\delta_s$, which can be caused by (b) large times between samples δ_s (with fixed $s = 15$) or (c) high numbers of samples s (with fixed $\delta_s = 1$). (b-c) When communication is limited, slower sampling time leads to higher probability of adapting. (d) The same two effects can be observed for fixed communication range ($r_c \approx 12.5\,\text{cm}$ for $c = 3$) and increasing robot density, which we control by modifying the swarm size N in an environment with a fixed size; (e) the ODE model (Eq. (1)) predicts qualitatively similar to results as seen in simulations ($N = 50$, $\kappa = 0.9$, $k = 4 \times 10^{-6}$); (f) sampling times τ_s influence the noise, because accuracy increases when robots collect more samples s, or reduces sample correlations with larger times δ_s in-between samples; (g) when estimation noise is independent of sampling and low ($\sigma = 0$) the slower-is-faster effect disappears; (h) when estimation noise is independent of sampling and high ($\sigma = 0.1$) both effects are present; if not specified, swarm size $N = 50$, difficulty $\kappa = 0.9$, 30 simulations each; colour-maps show probability to adapt.

We model the collective adaptation dynamics using a mean-field model built as a system of ODEs that describes the proportion of robots in each opinion state. Let x_i be the proportion of robots committed to the environmental element i and let x_u be the proportion of uncommitted robots, with $x_u + \sum_i x_i = 1$. The opinion dynamics model reads as

$$x_i = \underbrace{\frac{q_i}{\tau_s} x_p}_{A} + \underbrace{\frac{1}{\tau_s} \frac{k\, r_c^2\, N\, x_p}{1 + k\, r_c^2\, N\, x_p} q_i\, x_i}_{B} - \underbrace{\frac{k\, r_c^2\, N\, x_i}{1 + k\, r_c^2\, N\, x_i} \sum_{j \neq i} q_j\, x_j}_{C}$$
$$+ \underbrace{\frac{q_i}{\tau_s} \sum_{j \neq i} [x_j\, \mathcal{H}(\hat{q}_i - \hat{q}_j)]}_{D_{ji}} - \underbrace{\frac{x_i}{\tau_s} \sum_{j \neq i} [q_j\, \mathcal{H}(\hat{q}_j - \hat{q}_i)]}_{D_{ij}}\,, \tag{1}$$

where $\mathcal{H}$ is the unit step function, and k is a proportionality factor to fit the ODE system to the observed dynamics of the simulated swarm robotics systems (e.g., speed of robots, communication frequency, robots' opinion update time). The five terms on the rhs of Eq. (1) model the discovery, recruitment, cross-inhibition, and direct switching transitions, respectively (the capital letters below each term correspond to the transitions depicted in Fig. 1b).

The model of Eq. (1), as previously published [30], describes 'slow' (i.e., not instantaneous) recruitment through the Holling function type 2. As the sampling time τ_s is decreased, the recruitment becomes quicker and the effect of the Holling function reduces. As a result the recruitment rate becomes approximately linear on neighbourhood size. Through bifurcation analysis in the case of $n = 2$, we identify two states of the system with varied communication range (in Fig. 2e), or equivalently the swarm density (not shown). Prior to the subcritical bifurcation (low r_c or N), the system has a single stable equilibrium that represents a consensus decision for the colour with the highest concentration, therefore, in this parameter range adaptation is guaranteed. After the bifurcation (high r_c or N), a second stable equilibrium appears representing a consensus decision for the inferior alternative. In this parameter range, the swarm when initialised in the equilibrium for the inferior colour can only switch between the coexisting attractors through high random fluctuations and the swarm may take longer to adapt. This bifurcation analysis for sampling time τ_s and communication range r_c shows results that are qualitatively equivalent to the dynamics observed in simulations (see inset of Fig. 2e).

5.2 With noisy estimates and few neighbours, slower is faster

The results of Figs. 2b-d also show new interesting dynamics that were not found in the previous study [30]. When robot connectivity is low (i.e., sporadic social interactions) the swarm is only able to adapt when the sampling time is high (either the number of samples s or the time between readings δ_s are large). This mechanism corresponds to the slower is faster effect [8, 24, 27, 26, 17] by which the swarm is able to adapt at a quicker speed (i.e., within 40 minutes) when the robots perform their task of estimating the environmental element concentration at a slower pace. To study this phenomenon, we did further simulations. Slowing down the sampling process impairs the impact of slower recruitment and smaller errors on the colour concentration. Increasing either sample number s or time between samples δ_s has similar effects as reducing the estimation noise (e.g., see Fig. 2f). Therefore, we investigate whether adaptation of sparse swarms is limited by high noise or quick recruitment.

In additional simulations, robots sample colour concentration estimates ($\hat{q}_b$ or $\hat{q}_y$) from a normal distribution rather than observing tile colours. The normal distribution has the correct concentration of the colour in the environment as mean and we test various standard deviations σ. We disentangle noise from sampling time and study their impact separately. Without noise ($\sigma = 0$), swarms with low communication range are able to adapt to changes, hence sampling times have no impact on the collective ability to adapt (see Fig. 2g). Interestingly, when noise is high ($\sigma \in \{0.1, 0.2\}$, Fig. 2h), swarms are only able to adapt when robots take a long time to estimate (i.e., recruited robots are slow in becoming recruiters themselves). The slower a robot starts disseminating its opinion, the faster this opinion spreads throughout the swarm.

Unfortunately, we cannot provide an explanation of this effect by theoretical analysis as done for the LIME. The mean-field model of Eq. (1) describes a noiseless system and cannot model the slower is faster effect. In future work we will test stochastic models to study this phenomenon.

6 Discussion

We have shown that our previous results [30] generalise to a different scenario: collective perception of dynamic environmental features. This scenario allows for a more in-depth analysis not possible in the previous scenario. We have clarified the relationship between recruitment speed and ability to collectively adapt to environmental changes. The collective task that we study here is equivalent to enabling the swarm to revise an incorrect collective decision that led the swarm to reach a consensus for the inferior alternative and avoids lock-in states.

Our results explain the importance of considering the interplay between sampling time and the communication range when designing the robot behaviour as it can have a paramount effect on the collective dynamics. Through rigorous mathematical and computational analysis, we explain the mechanisms that cause the LIME, which is triggered by slow recruitment. During our investigations, we also stumbled upon a new effect: slower individual dissemination enables faster global agreement. We are unable, for the moment, to explain mathematically the slower is faster effect. However, our computational analysis confirms that the results on speed are not impaired with estimation noise. Our future research will investigate the mechanisms causing such unexpected dynamics which are highly relevant for swarm robotics studies as they manifest when swarm connectivity is sparse and robots follow a simple behaviour subject to high levels of noise.

Acknowledgements The authors thank Anthony Antoun, Marco Trabattoni, and Jonas Kuckling for technical support concerning Kilogrid and simulations on HPC. MD and AR acknowledge financial support from the Belgian F.R.S.-FNRS, of which they are Research Director and Chargé de Recherches, respectively.

References

1. Antoun, A., Valentini, G., Hocquard, E., Wiandt, B., Trianni, V., Dorigo, M.: Kilogrid: A modular virtualization environment for the kilobot robot. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3809–3814. IEEE (2016)
2. Aust, T., Reina, A.: Open-source code for simulating the Kilogrid in ARGoS, https://github.com/tilly111/adaptive_symmetry_breaking
3. Bartashevich, P., Mostaghim, S.: Benchmarking collective perception: New task difficulty metrics for collective decision-making. In: Moura Oliveira, P., Novais, P., Reis, L.P. (eds.) Progress in Artificial Intelligence. pp. 699–711. Springer International Publishing, Cham (2019). https://doi.org/https://doi.org/10.1007/978-3-030-30241-2˙58
4. Bartashevich, P., Mostaghim, S.: Multi-featured collective perception with evidence theory: tackling spatial correlations. Swarm Intelligence **15**(1), 83–110 (2021). https://doi.org/10.1007/s11721-021-00192-8
5. Bettstetter, C., Hartenstein, H., Pérez-Costa, X.: Stochastic properties of the random waypoint mobility model. Wireless Networks **10**(5), 555–567 (sep 2004). https://doi.org/10.1023/B:WINE.0000036458.88990.e5
6. Ebert, J.T., Gauci, M., Mallmann-Trenn, F., Nagpal, R.: Bayes bots: Collective bayesian decision-making in decentralized robot swarms. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 7186–7192 (2020). https://doi.org/10.1109/ICRA40945.2020.9196584
7. Ebert, J.T., Gauci, M., Nagpal, R.: Multi-feature collective decision making in robot swarms. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems. pp. 1711–1719 (2018)
8. Gershenson, C., Helbing, D.: When slower is faster. Complexity **21**(2), 9–15 (2015). https://doi.org/10.1002/cplx.21736
9. Hamann, H.: Swarm Robotics: A Formal Approach. Springer International Publishing, Cham (2018). https://doi.org/10.1007/978-3-319-74528-2
10. Lee, C., Lawry, J., Winfield, A.F.: Negative updating applied to the best-of-n problem with noisy qualities. Swarm Intelligence **15**(1), 111–143 (2021)
11. Nikolaidis, E., Sabo, C., Marshal, J.A.R., Reina, A.: Characterisation and upgrade of the communication between overhead controllers and Kilobots. Tech. rep., White Rose Research Online (April 2017)
12. Parker, C.A.C., Zhang, H.: Biologically inspired decision making for collective robotic systems. In: 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566). vol. 1, pp. 375–380. IEEE (2004). https://doi.org/10.1109/IROS.2004.1389381, http://ieeexplore.ieee.org/document/1389381/
13. Pinciroli, C., Talamali, M.S., Reina, A., Marshall, J.A.R., Trianni, V.: Simulating Kilobots within ARGoS: models and experimental validation. In: M. Dorigo et al. (ed.) Swarm Intelligence (ANTS 2018), LNCS, vol. 11172, pp. 176–187. Springer, Cham (2018). https://doi.org/https://doi.org/10.1007/978-3-030-00533-7˙14
14. Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., et al.: Argos: a modular, parallel, multi-engine simulator for multi-robot systems. Swarm intelligence **6**(4), 271–295 (2012)
15. Prasetyo, J., De Masi, G., Ferrante, E.: Collective decision making in dynamic environments. Swarm Intelligence **13**(3), 217–243 (2019). https://doi.org/10.1007/s11721-019-00169-8

16. Prasetyo, J., De Masi, G., Ranjan, P., Ferrante, E.: The best-of-n problem with dynamic site qualities: Achieving adaptability with stubborn individuals. In: Swarm Intelligence (ANTS 2018), LNCS, vol. 11172, pp. 239–251. Springer, Cham (2018)
17. Rahmani, P., Peruani, F., Romanczuk, P.: Flocking in complex environments—attention trade-offs in collective information processing. PLoS Computational Biology **16**(4), 1–18 (2020). https://doi.org/10.1371/journal.pcbi.1007697, http://dx.doi.org/10.1371/journal.pcbi.1007697
18. Reina, A., Cope, A.J., Nikolaidis, E., Marshall, J.A.R., Sabo, C.: ARK: Augmented Reality for Kilobots. IEEE Robotics and Automation Letters **2**(3), 1755–1761 (2017). https://doi.org/10.1109/LRA.2017.2700059
19. Reina, A., Ioannou, V., Chen, J., Lu, L., Kent, C., Marshall, J.A.: Robots as actors in a film: No war, a robot story. arXiv preprint arXiv:1910.12294 (2019)
20. Rubenstein, M., Ahler, C., Hoff, N., Cabrera, A., Nagpal, R.: Kilobot: A low cost robot with scalable operations designed for collective behaviors. Robotics and Autonomous Systems **62**(7), 966–975 (2014). https://doi.org/10.1016/j.robot.2013.08.006
21. Rubenstein, M., Cornejo, A., Nagpal, R.: Programmable self-assembly in a thousand-robot swarm. Science **345**(6198), 795–799 (aug 2014). https://doi.org/10.1126/science.1254295
22. Shan, Q., Mostaghim, S.: Collective decision making in swarm robotics with distributed bayesian hypothesis testing. In: International Conference on Swarm Intelligence. pp. 55–67. Springer (2020). https://doi.org/10.1007/978-3-030-60376-2˙5
23. Shan, Q., Mostaghim, S.: Discrete collective estimation in swarm robotics with distributed bayesian belief sharing. Swarm Intelligence **15**(4), 377–402 (2021). https://doi.org/10.1007/s11721-021-00201-w
24. Slobodkin, L.B.: Growth and Regulation of Animal Populations. Holt, Rinehart and Winston, New York (1961)
25. Soorati, M.D., Krome, M., Mora-Mendoza, M., Ghofrani, J., Hamann, H.: Plasticity in collective decision-making for robots: Creating global reference frames, detecting dynamic environments, and preventing lock-ins. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4100–4105 (2019)
26. Stark, H.u., Tessone, C.J., Schweitzer, F.: Decelerating microdynamics can accelerate macrodynamics in the voter model. Physical Review Letters **101**(1), 018701 (jun 2008). https://doi.org/10.1103/PhysRevLett.101.018701, https://link.aps.org/doi/10.1103/PhysRevLett.101.018701
27. Stark, H.U., Tessone, C.J., Schweitzer, F.: Slower is faster: fostering consensus formation by heterogeneous inertia. Advances in Complex Systems **11**(04), 551–563 (2008). https://doi.org/10.1142/S0219525908001805, https://www.worldscientific.com/doi/abs/10.1142/S0219525908001805
28. Strobel, V., Castelló Ferrer, E., Dorigo, M.: Managing byzantine robots via blockchain technology in a swarm robotics collective decision making scenario. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems. p. 541–549. AAMAS '18, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2018)
29. Strobel, V., Castelló Ferrer, E., Dorigo, M.: Blockchain technology secures robot swarms: A comparison of consensus protocols and their resilience to Byzantine robots. Frontiers in Robotics and AI **7**, 54 (2020). https://doi.org/10.3389/frobt.2020.00054

30. Talamali, M.S., Saha, A., Marshall, J.A.R., Reina, A.: When less is more: Robot swarms adapt better to changes with constrained communication. Science Robotics **6**(56), eabf1416 (2021). https://doi.org/10.1126/scirobotics.abf1416
31. Valentini, G., Antoun, A., Trabattoni, M., Wiandt, B., Tamura, Y., Hocquard, E., Trianni, V., Dorigo, M.: Kilogrid: a novel experimental environment for the kilobot robot. Swarm Intelligence **12**(3), 245–266 (2018)
32. Valentini, G., Brambilla, D., Hamann, H., Dorigo, M.: Collective perception of environmental features in a robot swarm. In: Dorigo, M., Birattari, M., Li, X., López-Ibáñez, M., Ohkura, K., Pinciroli, C., Stützle, T. (eds.) Swarm Intelligence. pp. 65–76. Springer International Publishing, Cham (2016). https://doi.org/10.1007/978-3-319-44427-7˙6
33. Valentini, G., Ferrante, E., Dorigo, M.: The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. Frontiers in Robotics and AI **4**, 9 (2017). https://doi.org/10.3389/frobt.2017.00009