Minimalist Protocols for Quorum Sensing in Robot Swarms

Fabio Oddi[1,2], Andreagiovanni Reina[3,4], and Vito Trianni[2]

[1] DIAG, Sapienza University of Rome, Italy
[2] ISTC, National Research Council, Rome, Italy
[3] CASCB, Universität Konstanz, Konstanz, Germany
[4] Department of Collective Behaviour, Max Planck Institute of Animal Behavior, Konstanz, Germany
oddi@diag.uniroma1.it, andreagiovanni.reina@gmail.com, vito.trianni@istc.cnr.it,

Abstract. Quorum sensing is a key mechanism enabling coordinated behaviour in populations of autonomous agents, and is extensively studied in biological systems spanning from bacteria populations to social insects. In swarm robotics too, quorum sensing aroused much interest, but remained mostly constrained to the implementation of collective decisions. Here, we propose protocols to estimate the quorum level that are suitable for resource-constrained robots, and evaluate the precision and speed of the quorum assessment across a large spectrum of swarm state conditions. Through systematic experimentation, we evaluate the proposed protocols for different swarm densities and working area sizes. Our findings shed light on the trade-off between computational requirements and expected performance, aiding in the selection of appropriate quorum sensing protocols for future swarm robotics research.

1 Introduction

Quorum sensing (QS) is a widespread phenomenon in both biological and artificial systems, playing a pivotal role in enabling coordinated system behaviour [3,2]. It can be defined as a "consistent population-dependent modulation of discrete modes of behaviour" [21], meaning that a population of agents can coordinately switch to a different behaviour when some population-level feature (e.g., density) surpasses a given threshold (i.e., a quorum). QS is largely studied in biology and especially in bacteria [36,17], in which a density-dependent concentration threshold of signalling molecules can trigger a behaviour change (e.g., bioluminescence in *Vibrio fischeri* [13]). In social insects, QS has been studied mainly as a decentralised mechanism necessary to implement a collective decision (e.g., migrating to a new nest site [24,31,7]). Indeed, collective decision making requires two processes: first, building consensus, so that a sufficient number of individuals in the population selects the same alternative (possibly among many); second, decision implementation, which requires that individuals change their behaviour according to the selected alternative. In nest site selection, for instance, first a new nest location must be identified, and then the colony can relocate to the new nest. To minimise colony splitting, a robust QS mechanism must be in place, enabling to implement the decision only when sufficient support is available. This also leads to a speed-accuracy trade-off, because a fast-and-frugal estimation of the quorum can lead to frequent errors in the decision implementation [10].

In swarm robotics too, QS has been conceived as a mechanism for the implementation of collective decisions [20,6,21,8,16]. Indeed, to orchestrate swarm-level activities across multiple functional tasks, the robot swarm must be capable of collective decision making and task switching, whereby in face of multiple alternatives the swarm selects the most profitable or urgent task and coordinately executes it. Minimalist approaches for best-of-N decision problems have been proposed in the past [35]. Recently, it has been shown that a hierarchy of collective decisions can simplify best-of-N problems, improving both speed and accuracy [18]. Such approaches can be generalised to coordinated task switching, whereby the swarm is required to coordinately move from one task to another maintaining coherence and prioritising the most relevant task or suitably distributing among many [15,1]. In both task switching and decision sequences, QS strategies are fundamental to establish when the swarm is ready to move on. This must be optimised for speed to avoid unnecessary delays and energy consumption. Additionally, recovery mechanisms need to be developed for robots that wrongly recognise the quorum or that engage in a different task, by means of systematic coherence check within the swarm.

In this paper, we move beyond previous research on QS for robot swarms [20,6,21] by studying the system dynamics in isolation from the collective decision-making process and by comparing three alternative protocols in terms of their computational and memory requirements. We make a theoretical analysis through an urn model that informs the swarm robotics simulations. We test the ability and speed of the robot swarm to detect a quorum for a wide range of quorum levels. Moreover, we discuss how much QS is impacted by swarm density and size of the working area, which contribute to determining the effectiveness of the sampling strategy. Our results show that effective QS protocols can be implemented on resource-constrained robotic platforms.

2 Background

As mentioned above, QS finds application in swarm robotics, particularly in association with collective decision making, as a mechanism to determine when one alternative gathered sufficient support. Over the years, a few approaches have been developed to estimate the swarm state and react accordingly with a system-level response. Bacteria-inspired QS did not find much application in robotics as it requires that some chemical product is produced and diffused in the environment, and its concentration measured [25]. Much as with pheromone communication, dealing with chemicals imposes technical challenges difficult to overcome [11], or that require some special infrastructure to be simulated effectively [32,30]. Inspired by QS in social insects, approaches based on sampling the state of neighbouring individuals received more attention. Notably, two main approaches can be followed employing either anonymous or identity-aware interactions, each offering distinct advantages and disadvantages in different scenarios.

Anonymous protocols are mostly inspired by studies performed in honeybees and ants during nest site selection [31,24]. In these studies, individuals base QS on the encounter rate with conspecifics, as the estimation takes place at the candidate site: the higher the encounter rate, the stronger the support for the site. Models and algorithms replicate QS via encounter rates, either through a leaky integrator [21] or by maintaining an anonymous buffer of messages [20]. These approaches clearly minimise memory requirements and computational overhead, and can be implemented with very simple logic.

In contrast, identity-aware protocols leverage unique identification mechanisms to recognise each agent within the swarm. These protocols integrate individual identities alongside state information, enabling more precise decision making at the cost of higher memory requirements and larger computational and communication complexity, posing challenges in resource-constrained settings. A class of approaches closely related to QS is decentralised node counting in networks of autonomous agents [12,29], as QS can also be seen as the problem of counting how many agents are in a given state. Closely related to approaches inspired by social insects, population sampling methods allow obtaining estimates of the population state [6,8,5]. In these studies, agents share their state upon encounter, and each agent stores/updates other agents' states to evaluate if the quorum has been reached using a sufficiently large sample of the population. Both majority and k-unanimity voting rules are tested [6], the former presenting faster dynamics.

The choice between anonymous and identity-aware protocols hinges on the specific requirements and constraints of the swarm robotics application at hand. While anonymous protocols offer simplicity and efficiency in memory usage, they may struggle to maintain robustness in scenarios necessitating fine-grained coordination or prolonged interactions among a subset of agents. Conversely, identity-aware protocols provide enhanced precision and adaptability but demand greater computational resources and may introduce complexities in implementation and maintenance. Thus, understanding the trade-offs between these approaches is essential in designing effective QS protocols for swarm robotics applications, also considering performance in terms of precision and speed.

3 Problem Description

We consider a QS problem in which a group of N agents must collectively recognise if a given portion of the group agrees about an opinion. To focus on the QS dynamics aside from the consensus-building process, we consider agents with two possible states, *committed* and *uncommitted*. Agents are randomly initialised in one or the other state, and never change it. As a consequence, the swarm has a constant fraction of *committed* agents, hereafter referred to as the *ground truth* $G \in [0, 1]$. We consider that the quorum is reached when a minimum fraction

$Q \in [0, 1]$ of the swarm recognises that the percentage of teammates in the *committed* state is larger than a given *threshold* $\tau \in [0, 1]$. Ideally, one would expect that $Q = 1 \iff G \geq \tau$, and conversely $Q = 0 \iff G < \tau$.

As mentioned in Section 2, agents can estimate whether or not the quorum is reached by sampling the state of other agents within the population either anonymously or by being aware of the identity of the interacting individuals. Anonymous protocols minimise assumptions about interactions among agents and reduce the memory requirements, but can be less effective in case of multiple encounters among the same individuals because these individuals will be double counted. This double-counting problem can be resolved through identity-aware protocols, which however require that identity recognition is possible and are more demanding in terms of memory. In swarm robotics, robots are often assigned a unique identification number (ID), which can be communicated to neighbours together with information about the agent's state. Identity-aware protocols can therefore be easily implemented, provided that the memory requirements are considered. In the following, we first discuss theoretical implications related to choosing anonymous or identity-aware protocols, and then we detail the minimal implementations we propose for robot swarms.

3.1 Urn Models of Quorum Sensing

We assume a well-mixed population, where the probability that two agents interact is the same for any couple of agents. Hence, the sampling process carried out by agents to estimate the quorum state can be easily modelled using simple urn models. Despite their simplicity, urn models are often exploited to provide guidance for the understanding of stochastic processes in complex systems such as robotic swarms [14]. Urn models describe probabilistic events as extractions of balls from an urn. In our case, balls represent agents and the ball's colour represents their state (i.e., committed or uncommitted). Therefore the probability of obtaining information from a neighbour agent can be modelled by the extraction of a ball from an urn. Anonymous protocols can be represented by urn sampling with replacement, because the same agent (ball) can be sampled multiple times, therefore the ball is replaced in the urn after being sampled. Instead, identity-aware protocols can be represented by urn sampling without replacement as double-counting is prevented by keeping track of the agent identities. For both cases, we provide models to compute the probability of detecting a quorum as a function of the number of samples n.

Urn sampling with replacement. Consider an urn containing N balls with W white balls and $B = N - W$ black balls (i.e., $G = \frac{W}{N}$). The probability of extracting a white ball is $P_w = G$ and of extracting a black ball is $P_b = 1 - G$. Hence, the probability of having k white balls extracted with replacement in n trials is

$$P_r(n, k) = \binom{n}{k} P_w^k P_b^{n-k}, \qquad k \leq n. \tag{1}$$

So, the probability of having at least τ percent (e.g., 80% in our experiments) of the balls extracted within n trials to be white can be written as

$$\mathcal{P}(n,\tau) = \sum_{k=\lceil n\tau \rceil}^{n} P_r(n,k), \tag{2}$$

where $\lceil . \rceil$ is the ceil operator. Note that here $\mathcal{P}(n,\tau)$ does not depend on the population size N, as a consequence of the replacement.

Urn sampling without replacement. Consider an urn containing N balls with W white balls and $B = N - W$ black balls (i.e., $G = \frac{W}{N}$). If there is no replacement, the probability of having $k \leq W$ white balls within n extractions can be written as

$$P_n(n,k) = \frac{\binom{W}{k}\binom{N-W}{n-k}}{\binom{N}{n}}, \qquad k \leq n. \tag{3}$$

Here, the probability of having at least τ percent of the balls being white is

$$\mathcal{P}(n,\tau) = \sum_{k=\lceil n\tau \rceil}^{n} P_n(n,k). \tag{4}$$

3.2 Implementation of Quorum Sensing in Robot Swarms

We consider a swarm of N robots randomly moving in a square arena (side length L). A total of $\lceil GN \rceil$ robots are randomly chosen and initialised in the committed state, the rest are set to the uncommitted state. Robots can communicate with neighbours within a radius r, sharing their unique ID and a bit b_c indicating their state ($b_c = 1$ for committed, $b_c = 0$ for uncommitted). Messages are broadcast every t_c seconds, and upon reception they are stored and processed to evaluate the existence of the quorum. Owing to communication, sampling of the swarm state can be performed.

To estimate the quorum level, each robot maintains a buffer $\mathcal{B}$ of received messages and, at each buffer update, computes the proportion of neighbours in each state and compares it to the threshold τ. To have a large enough sample over which to compute the qualified majority, we impose a minimum buffer dimension B_m before considering the quorum assessment:

$$b_q = |\mathcal{B}| \geq B_m \wedge \sum_{m \in \mathcal{B}} b_c(m) \geq \tau|\mathcal{B}|. \tag{5}$$

Here, b_q is a bit representing the quorum detection state of the agent, and $b_c(m)$ is the commitment bit stored in message m. Note that the quorum detection state can transition to 0 if the conditions on the buffer $\mathcal{B}$ do not hold anymore. We propose three different approaches over the same problem formulation to implement both anonymous and identity-aware protocols.

Anonymous Protocol. To implement an anonymous protocol, the robot ID contained in a received message m is ignored and only $b_c(m)$ is stored in $\mathcal{B}$. The buffer $\mathcal{B}$ implements a FIFO method for memory management with fixed size $B_M \leq N$, closely following the implementation in [20]. Hence, the maximum memory requirement is B_M bits. Recall that the commitment state of the same robot can appear multiple times within the buffer, especially if robots remain in mutual proximity for a sufficient time. This may bias the estimation of the quorum level and ultimately the accuracy of the algorithm.

Identity-Aware Protocol with Message Broadcasting (ID+B). In this case, a qualified majority is computed only with information coming from different robots, similarly to what is proposed in [6]. The buffer $\mathcal{B}$ stores any received message m in a list of tuples $\langle k(m), b_c(m), t\rangle$, where $k(m)$ is the robot ID, $b_c(m)$ the corresponding commitment state, and t is a timeout for storing a message drawn from an exponential distribution with average T_m. Through this timeout, old messages are removed from the buffer, forcing the robot to make QS estimates on fresh information. If a new message m' is received from robot $k(m')$ and there is already a tuple from the same robot in the buffer, then the new message is discarded only if the status bit is unchanged, otherwise the old message is replaced with the new one and a new timeout T'_m is computed. This approach allows mitigating the effects of repeated encounters among the same robots, while keeping track of changes in the neighbour status. Overall, if we use B_k bits to store the robots' IDs and B_t bits for timeouts, this protocol requires at most $N(B_k+B_t+1)$ bits to store the message buffer. As we will see in Section 4, the buffer size can be limited by reducing the average timeout T_m.

Identity-Aware Protocol with Message Re-broadcasting (ID+R). One limitation of the above protocols is that messages are received only within the communication range r. To address this limitation, we propose a simple rebroadcasting protocol, which allows to diffuse information widely within the swarm by forwarding the same message multiple times. Here, the buffer $\mathcal{B}$ stores any received message m in a list of tuples $\langle k(m), b_c(m), t, b_r\rangle$, where the additional bit b_r indicates whether or not the information has been rebroadcast. The rebroadcast approach implements a FIFO strategy: every time the robot can communicate (i.e., every t_c seconds), it rebroadcasts the oldest message that has not been rebroadcast yet and the corresponding bit is set in $\mathcal{B}$. When there is no message to rebroadcast, the robot shares its own state. Note that the choice of a FIFO may delay the broadcast of one agent's state, as the latter is shared only when the FIFO is empty. However, preliminary tests without a FIFO revealed that such a delay has a negligible effect in practice for the studied settings (data not shown). Concerning memory requirements, if we do not consider a rebroadcasting FIFO buffer (which can be implemented within $\mathcal{B}$ at the expense of some additional computation), this protocol requires at most $N(B_k + B_t + 2)$ bits.

3.3 Implementation with Kilobots

In this study, QS has been tested with Kilobot robots [27,28] simulated through ARGoS [23,22]. To support the experimentation, we use the ARK system that significantly enhances the experimentation possibilities with Kilobots [26,9]. The motion pattern implemented with the Kilobots is a random waypoint model [4], which corresponds to a directed movement towards a position randomly chosen within the working area. This motion pattern keeps the swarm constrained within the predefined working area without the need for collision avoidance—which cannot be performed by Kilobots—and can be easily implemented employing ARK as a global positioning system [33].

Kilobots communicate at a maximum rate of about 2 Hz, that is, $t_c = 0.5$ s. Messages can be effectively received within a radius of $r = 0.1$ m. We have implemented the buffer $\mathcal{B}$ as a doubly linked list to optimise traversal, insertion and deletion of messages. A static vector of indices is also used to quickly check if messages are already present in the buffer, and access them. These structures come with a negligible overhead in memory requirements and can be easily implemented on the Kilobots (see the open-source code [19]).

4 Results

Starting from the probabilistic urn models, we first show how the probability of sampling a qualified majority of white balls changes with and without replacement, mirroring the usage of anonymous and identity-aware protocols in a well-mixed population. To this end, we compute the maximum threshold $\tau_M \leq G$ that ensures a high probability—higher or equal than 80%—of detecting a quorum. In other words, we determine how precise the threshold τ should be (i.e., how close it should be to the ground truth G) to detect the quorum with high probability $\mathcal{P}(n, \tau) \geq 0.8$. Figure 1 shows the value of τ_M normalised on G, for varying sample size $n \in [1, 25]$ and ground truth $G = i/N, i \in \{\lceil N/2 \rceil, \dots, N\}$. When $N = 25$, the urn sampling without replacement (NR25) predicts higher precision in detecting the quorum than the case with replacement (R25). On the other hand, when the population size is significantly larger than the number of samples ($N = 100$ in Figure 1), differences fade away, as it is possible to no-

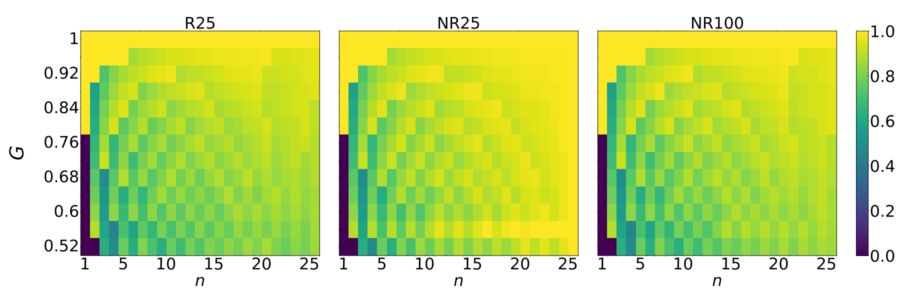

Fig. 1. Prediction of the urn models for sampling with and without replacement. The heatmap shows—for each pair n, G—the normalised maximum threshold τ_M/G that ensures a probability $\mathcal{P}(n, \tau_M) \geq 0.8$. Left: results with replacement and $N = 25$ balls (R25). Centre: results without replacement and $N = 25$ balls (NR25). Right: results without replacement and $N = 100$ balls (NR100).

tice comparing NR100 with R25.[5] Indeed, if the well-mixed assumption holds, an anonymous protocol should be as good as the identity-aware protocol when $N \gg n$. As an additional result, we can determine what is the minimum sampling size B_m to obtain a good precision. Looking at Figure 1, we can see that for $n < 5$ the precision is not very good for all sampling strategies, but for larger values of n the value of τ_M need not be too distant from G. We therefore fix $B_m = 5$ for all swarm robotics experiments, considering that this is a minimum requirement and that the size of the buffer $\mathcal{B}$ can grow larger than B_m.

Swarm robotics experiments are performed in simulation, testing different swarm densities and arena sizes. Specifically, we consider a low density case (LD25) with $N = 25$ robots in a large square arena ($L = 1\,\text{m}$), a high density case (HD25) with $N = 25$ robots in a small square arena ($L = 0.5\,\text{m}$), and a high density case (HD100) with $N = 100$ robots in a large square arena ($L = 1\,\text{m}$). To give an idea of the consequences of the robot density, consider a random geometric network induced by the robot interactions with average degree $\langle k \rangle = \pi N r^2 / L^2$ [34]. Hence, for LD25 we have less than one neighbour per robot on average ($\langle k \rangle = \pi/4$), while for HD25 and HD100 we have more than three neighbours on average ($\langle k \rangle = \pi$), and a value closer to the percolation threshold $\langle k_c \rangle \approx 4.51$. We study the anonymous and the two identity-aware protocols (ID+B and ID+R) varying the memory requirements, which are determined by the maximum buffer length B_M for the former and by the average timeout T_m for the latter two. Figure 2 shows how different settings influence the buffer length. Given that the anonymous protocol has a fixed buffer length B_M, we plot the number M of unique messages in $\mathcal{B}$, hence excluding double counting. It is possible to notice that the anonymous protocol accumulates a lower amount of information about the swarm than the identity-aware protocols, mainly due to double counting. Additionally, the re-broadcasting protocol gives a significant speed advantage, converging to the stationary value earlier than what simple broadcasting can achieve. High-density conditions ensure good interaction rates among robots and efficient sampling. Low-density conditions suffer from inefficient communication, and the difference between anonymous and identity-aware protocols is reduced.

To evaluate the quality of the QS protocols, we compute the fraction $Q(t)$ of robots that at time t recognise the quorum given a ground truth $G = i/N, i \in \{\lceil N/2 \rceil, \dots, N\}$ and the threshold $\tau \in [0.5, 1]$ in all the mentioned experimental conditions. We then compute the average quorum detection $\hat{Q}(G, \tau)$ as the average value of $Q(t)$ over $T = 900\,\text{s}$ and 100 independent runs. Figure 3 shows the isolines for $\hat{Q} = 0.8$ and $\hat{Q} = 0.2$, which represent boundaries of regions in which the swarm consistently recognises a quorum ($\hat{Q} \geq 0.8$) or reject it ($\hat{Q} < 0.2$). As mentioned in Section 3, ideally the region $\mathcal{R}$ between these isolines should be minimised, as this region corresponds to an undecided state, possibly leading to QS errors. It is possible to notice that in all settings the anonymous protocol leads to a wider region $\mathcal{R}$. Among the two identity-aware protocols, ID+R has a slight advantage consistently across experimental setups. When increasing the

[5] Recall that the urn model with replacement does not depend on system size N.

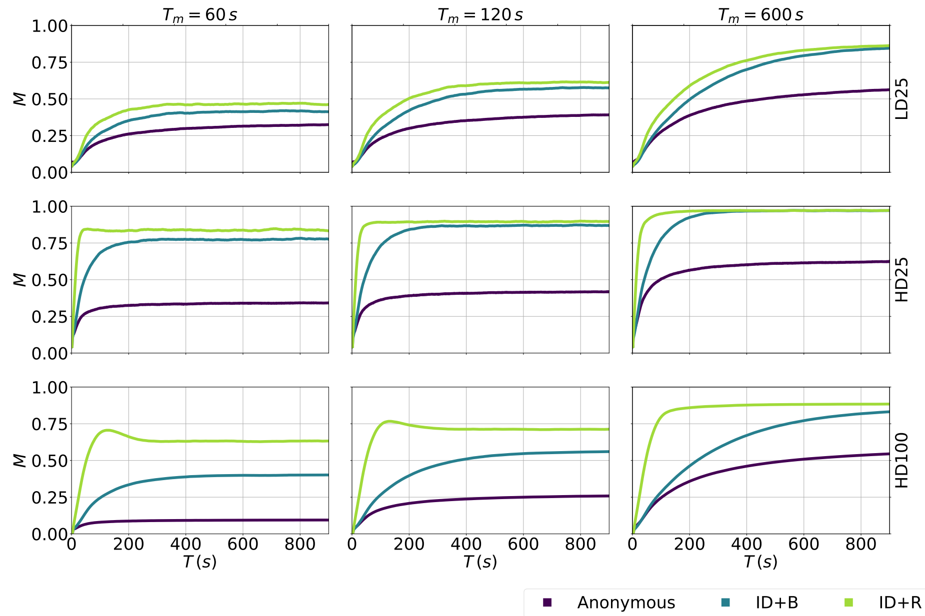

Fig. 2. Evolution over time of the average number M of unique messages in the buffer normalised over the maximum allowable size (i.e., $N-1$). From top to bottom, the density-size scenario is changed, respectively LD25, HD25 and HD100. From left to right, the average timeout T_m and the maximum buffer length B_M values increase. We set $B_m \in \{10, 13, 24\}$ for $N = 25$, and $B_m \in \{10, 32, 99\}$ for $N = 100$. Data are averaged over all robots across $R = 100$ simulation runs.

timeout T_m and maximum buffer size B_M, the region $\mathcal{R}$ always gets smaller as the qualified majority is evaluated on a larger sample. Similarly, moving from low to high densities and larger numbers of robots increases the sampling size and in turn the quality of the estimation.

Finally, we evaluate how fast the different protocols can lead to a reliable recognition of the quorum. To this end, for each value of $\tau \in [0.5, 1]$, we consider the smallest value $\hat{G}$ that results in a reliable estimation:

$$\hat{G} = \arg\min_{G} \left(\hat{Q}(G, \tau) \geq 0.8 \right), G = i/N, i \in \{\lceil N/2 \rceil, \ldots, N\}. \tag{6}$$

For such values of τ and $\hat{G}$, we record the time at which $Q(t)$ exceeds the 0.8 threshold. Figure 4 shows the median quorum recognition time T_c across 100 runs. It is possible to notice that the anonymous protocol is rather fast, but as we have seen in Figure 3 it trades off precision for speed. Conversely, the identity-aware protocols present similar behaviour with low density. However, ID+R is much faster than the other approaches with high density, because rebroadcasting leads to faster diffusion of information and an improved sampling (as also seen with the growth rate of the message buffer $\mathcal{B}$ in Figure 2).

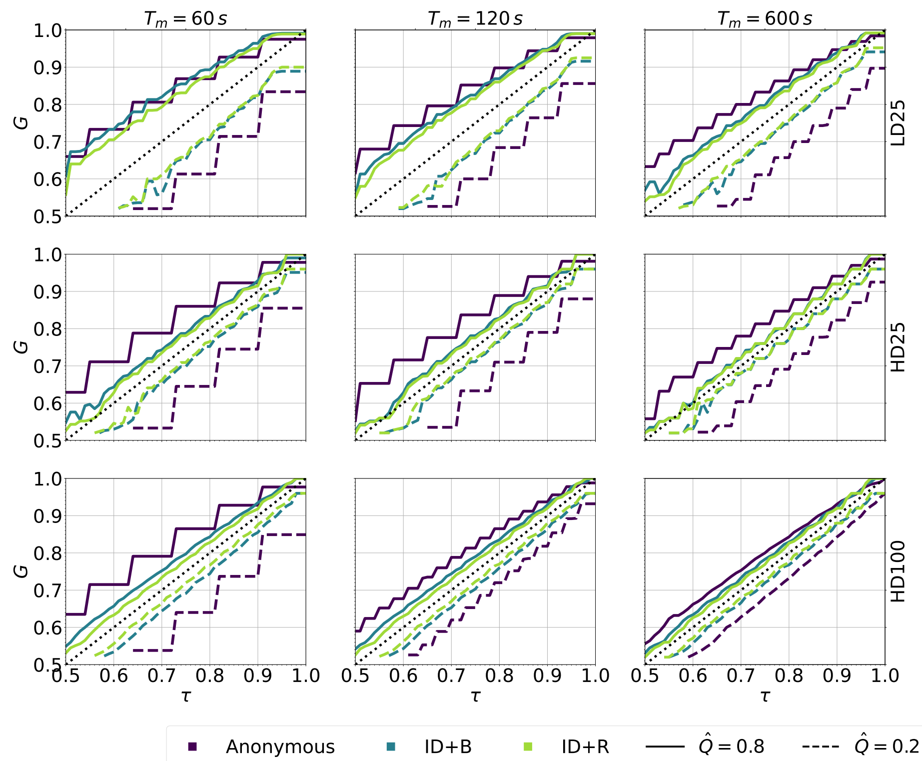

Fig. 3. Isolines of the average quorum detection for $\hat{Q} = 0.8$ (solid lines) and $\hat{Q} = 0.2$ (dashed lines). The isolines are computed through linear interpolation between the grid points used for computing the average quorum detection for all values of G and τ.

5 Conclusions and Future Work

In this paper, we explored minimalist approaches to quorum sensing (QS) based on anonymous or identity-aware sampling protocols. Our results demonstrate that anonymous protocols suffer from the double counting problem, which can be mitigated only if the robot swarm population is sufficiently well-mixed. However, robotics settings typically have strong spatial and communication correlations, hence the system is often far from being well-mixed. Identity-aware approaches are effective in both low and high densities, and message rebroadcasting can boost the QS process both in accuracy and speed. We hypothesise that, whenever the density is sufficiently high to enable information sharing, ID+R can be an effective solution to QS in particularly challenging scenarios characterised by large spatial heterogeneities in the robot distribution. Finally, ID+R should also adapt well to dynamic conditions where the ground truth changes over time, owing to the speed of diffusing new information. These conditions will also be the subject of future investigation.

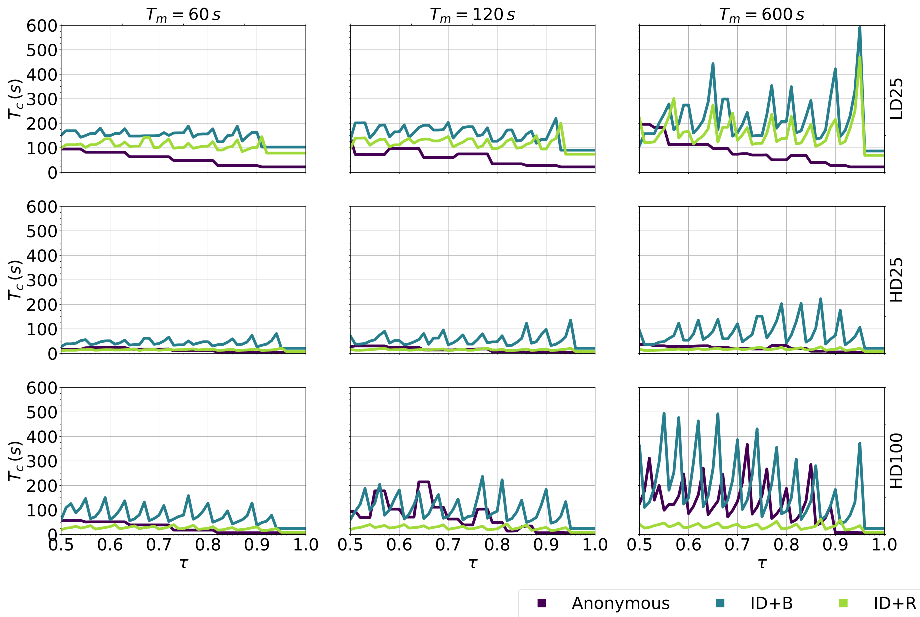

Fig. 4. Median times for the swarm to reliably detect the quorum across 100 different runs. We can interpret the oscillations as an artefact determined by the way in which we select $\hat{G}$, which results in more or less close values of $P(\hat{G}, \tau)$ to the 0.8 threshold: the closer it is, the slower is the quorum detection process.

The systematic experimentation we have conducted should provide valuable information to swarm robotics practitioners who need to define how to estimate the state of the swarm in a fast and reliable way. Further characterisation of the mechanisms presented here will improve the ability to make informed choices about the QS protocol, especially in relation to dynamic settings, different motion patterns and varying task requirements. For instance, when robot swarms perform specific tasks (e.g., foraging), their interaction topology stops approximating a well-mixed system, and the performance of the different QS protocols can be affected. Finally, recovery from errors should be explicitly accounted for. While the presented protocols can correct from errors as long as additional samples are collected, the speed and reliability of error recovery needs to be explicitly assessed and possibly adaptive mechanisms should be put in place to enable quick reaction both locally and globally.

Acknowledgements. A.R. acknowledges support from DFG under Germany's Excellence Strategy - EXC 2117 - 422037984. V.T. acknowledges support from the project TAILOR (H2020-ICT-48 GA: 952215) and from the PNRR MUR project PE0000013-FAIR.

References

1. Albani, D., Hönig, W., Nardi, D., Ayanian, N., Trianni, V.: Hierarchical task assignment and path finding with limited communication for robot swarms. Applied Sciences **11**(7), 3115 (2021)
2. Amir, Y., Abu-Horowitz, A., Werfel, J., Bachelet, I.: Nanoscale robots exhibiting quorum sensing. Artificial Life **25**(3), 227–231 (2019)
3. Atkinson, S., Williams, P.: Quorum sensing and social networking in the microbial world. Journal of The Royal Society Interface **6**(40), 959–978 (2009)
4. Bettstetter, C., Hartenstein, H., Pérez-Costa, X.: Stochastic properties of the random waypoint mobility model. Wireless Networks **10**(5), 555–567 (2004)
5. Cai, G., Sofge, D.: An urgency-dependent quorum sensing algorithm for n-site selection in autonomous swarms. In: AAMAS 2019: Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems. pp. 1853–1855. IFAAMAS, Richland, SC (2019)
6. Cody, J.R., Adams, J.A.: An evaluation of quorum sensing mechanisms in collective value-sensitive site selection. In: 2017 International Symposium on Multi-Robot and Multi-Agent Systems (MRS). pp. 40–47. IEEE (2017)
7. Cronin, A.L.: Ratio-dependent quantity discrimination in quorum sensing ants. Animal Cognition **17**(6), 1261–1268 (2014)
8. Ebert, J.T., Gauci, M., Nagpal, R.: Multi-feature collective decision making in robot swarms. In: AAMAS 2018: Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems. pp. 1711—1719. IFAAMAS, Richland, SC (2018)
9. Feola, L., Reina, A., Talamali, M.S., Trianni, V.: Multi-swarm interaction through augmented reality for Kilobots. IEEE Robotics and Automation Letters **8**(11), 6907–6914 (2023)
10. Franks, N.R., Dechaume-Moncharmont, F.X., Hanmore, E., Reynolds, J.K.: Speed versus accuracy in decision-making ants: expediting politics and policy implementation. Philosophical Transactions of the Royal Society B: Biological Sciences **364**(1518), 845–852 (2009)
11. Fujisawa, R., Dobata, S., Sugawara, K., Matsuno, F.: Designing pheromone communication in swarm robotics: Group foraging behavior mediated by chemical substance. Swarm Intelligence **8**(3), 227–246 (2014)
12. Ganesh, A.J., Kermarrec, A.M., Merrer, E.L., Massoulié, L.: Peer counting and sampling in overlay networks based on random walks. Distributed Computing **20**(4), 267–278 (2007)
13. Girard, L.: Quorum sensing in *Vibrio* spp.: the complexity of multiple signalling molecules in marine and aquatic environments. Critical Reviews in Microbiology **45**(4), 451–471 (2019)
14. Hamann, H.: Towards swarm calculus: urn models of collective decisions and universal properties of swarm performance. Swarm Intelligence **7**(2), 145–172 (2013)
15. Hsieh, M.A., Halász, Á., Berman, S., Kumar, V.: Biologically inspired redistribution of a swarm of robots among multiple sites. Swarm Intelligence **2**(2), 121–141 (2008)
16. Leaf, J., Adams, J.: The effect of uneven and obstructed site layouts in best-of-N. Swarm Intelligence (2024)
17. Miller, M.B., Bassler, B.L.: Quorum sensing in bacteria. Annual Review of Microbiology **55**(1), 165–199 (2001)

18. Oddi, F., Cristofaro, A., Trianni, V.: Best-of-N collective decisions on a hierarchy. In: Swarm Intelligence. 13th International Conference, ANTS 2022. Lecture Notes in Computer Science, vol. 13491, pp. 66–78. Springer, Cham (2022)
19. Oddi, F., Reina, A., Trianni, V.: Code associated to "Minimalist Protocols for Quorum Sensing in Robot Swarms". https://github.com/Fabio930/argos3-kilobot.git (2024)
20. Parker, C.A., Zhang, H.: Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. Swarm Intelligence **4**, 199–220 (2010)
21. Pavlic, T.P., Hanson, J., Valentini, G., Walker, S.I., Pratt, S.C.: Quorum sensing without deliberation: biological inspiration for externalizing computation to physical spaces in multi-robot systems. Swarm Intelligence **15**(1-2), 171–203 (2021)
22. Pinciroli, C., Talamali, M.S., Reina, A., Marshall, J.A.R., Trianni, V.: Simulating Kilobots within ARGoS: Models and experimental validation. In: Swarm Intelligence: 11th International Conference, ANTS 2018. LNCS, vol. 11172, pp. 176–187. Springer, Cham (2018)
23. Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Caro, G.D., Ducatelle, F., Birattari, M., Gambardella, L.M., Dorigo, M.: ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. Swarm Intelligence **6**(4), 271–295 (2012)
24. Pratt, S.C.: Quorum sensing by encounter rates in the ant *Temnothorax albipennis*. Behavioral Ecology **16**(2), 488–496 (2005)
25. Pérez-Velázquez, J., Gölgeli, M., García-Contreras, R.: Mathematical modelling of bacterial quorum sensing: A review. Bulletin of Mathematical Biology **78**(8), 1585–1639 (2016)
26. 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)
27. 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)
28. Rubenstein, M., Cornejo, A., Nagpal, R.: Programmable self-assembly in a thousand-robot swarm. Science **345**(6198), 795–799 (2014)
29. Saha, A., Marshall, J.A.R., Reina, A.: Memory and communication efficient algorithm for decentralized counting of nodes in networks. PLoS ONE **16**(11), e0259736 (2021)
30. Salman, M., Ramos, D.G., Hasselmann, K., Birattari, M.: Phormica: Photochromic pheromone release and detection system for stigmergic coordination in robot swarms. Frontiers in Robotics and AI **7**, 591402 (2020)
31. Seeley, T., Visscher, P.: Quorum sensing during nest-site selection by honeybee swarms. Behavioral Ecology and Sociobiology **56**, 594—-601 (2004)
32. Talamali, M.S., Bose, T., Haire, M., Xu, X., Marshall, J.A.R., Reina, A.: Sophisticated collective foraging with minimalist agents: a swarm robotics test. Swarm Intelligence **14**(1), 25–56 (2020)
33. 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) (2021)
34. Trianni, V., De Simone, D., Reina, A., Baronchelli, A.: Emergence of consensus in a multi-robot network: From abstract models to empirical validation. IEEE Robotics and Automation Letters **1**(1), 348–353 (2016)

35. 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**, 1–43 (2017)
36. Waters, C.M., Bassler, B.L.: Quorum sensing: Communication in bacteria. Annual Reviews in Cell Development Biology **21**(1), 319–346 (2005)