

A Reliable Distributed Algorithm for Group Size Estimation with Minimal Communication Requirements

Manuele Brambilla, Carlo Pinciroli, Mauro Birattari and Marco Dorigo

Abstract— This study proposes a method that lets individual robots in a group estimate the size of the group in a distributed manner. The process is loosely based on the signaling behavior of fireflies and crickets. Holland *et al.* first devised a method based on local robot signaling to estimate the group's size. Each robot emits a signal and can perceive the signals of neighboring robots in close proximity. Following the approach of Holland *et al.*, our robots count the number of emitted signals over a suitably defined period of time. Experiments show that the estimates calculated with Holland *et al.*'s method display a great deal of noise. We modify their method so as to sensibly stabilize the output. We assess the quality of our method through extensive simulation-based experiments.

I. INTRODUCTION

Swarm robotics [9] is a branch of collective robotics focused on the study of relatively large groups of robots with limited sensor and communication capabilities. Swarm robotic systems naturally display a high level of redundancy and parallelism, thus making them suitable for complex and high risk scenarios such as rescue missions [2] and space exploration [3]. For such applications, some of the robots are likely to be lost or experience failures, thus making the number of surviving robots an important piece of information for tuning the action of the swarm.

Furthermore, in many applications the performance of the swarm varies with its size and an optimal value for size exists. For example, a minimum number of robots are required to transport a heavy object, and as the number of robots assisting with the task increases, the overall performance increases superlinearly [8]. At the same time, the coordination of the swarm becomes more complex; beyond a certain number of robots, this negative effect dominates and performance decreases [6]. Therefore, we can identify an optimal number of robots corresponding to maximum performance.

Given the effect of swarm size on performance, the ability to estimate the swarm's size is frequently of significant importance. Based on this piece of information, robots may switch to the fittest operational regime and/or divide into groups of optimal sizes for the tasks to perform. In the typical usage scenario, a swarm of known size is deployed and failures are expected as the task progresses, thus lowering the group's size. When the loss of robots drives the swarm's size below a certain threshold, robots may send messages to the operator and ask for the addition of new robots.

Manuele Brambilla, Carlo Pinciroli, Mauro Birattari and Marco Dorigo are with Université Libre de Bruxelles, Brussels, Belgium ({mbrambill, cpinciro, mbiro, mdorigo}@ulb.ac.be).

In this paper, we propose a simple distributed algorithm that lets each robot in a swarm estimate individually and in a reliable manner the size of the swarm. The algorithm works assuming only minimal communication requirements, i.e. each robot can only send signals to its neighbors within a very limited range. Our method is based on the pioneering work of Holland *et al.* [5] on group size estimation, which in turn was inspired by the simple flash synchronization process of fireflies [10]. As discussed in more detail in Section III, Holland *et al.*'s method is very promising, but provides a noisy and unreliable estimate of the swarm's size for practical use. We modify this method so as to obtain more reliable estimates.

The paper is organized as follows. In Section II we detail the basic methodology we applied throughout the paper. In Section III we introduce the method of Holland *et al.* to obtain an estimate of the group size and discuss its performance. Section IV explains our modifications to improve performance. The robustness, scalability and stability of our modified method is discussed in Section V. Section VI concludes the paper and proposes future research directions.

II. METHODOLOGY

A. Objectives

The aim of our work is to design a reliable method that lets each individual robot in a swarm estimate the group's size. By *reliable*, we mean that our method should display properties which make it applicable in a real scenario.

The first desirable property such a method should display is low individual estimation error, i.e., the difference between a robot's estimate and the actual group size should be small. Secondly, each robot's estimation should be *stable*, meaning that once settled on a value, the estimate should not fluctuate significantly. Finally, another important property is a high degree of *agreement* throughout the swarm, i.e., the estimates of the individuals in the swarm should not be too different among each other.

B. Quality Assessment

To compare the original method of Holland *et al.* with our modified one, we defined suitable measures of the desired properties discussed above.

We proceeded by first running a set of experiments to find, for each method, the parameter choice that minimized the mean relative estimation error across the swarm. The individual mean relative error of robot i is defined as

$$E^i = \frac{1}{t_2 - t_1} \sum_{t=t_1}^{t_2} \left| \frac{N - \hat{n}_t^i}{N} \right|$$

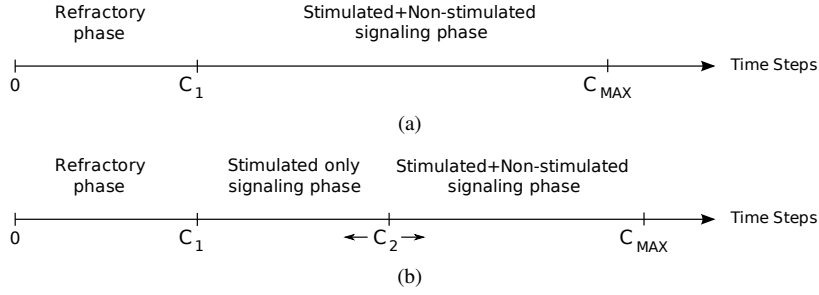


Fig. 1. The phase cycles of the HMH method (a) and of the proposed method (b).

where N is the actual number of robots in the swarm and \hat{n}_t^i is robot i 's t -th estimate. In our experiments, to ensure fair comparisons, the observation period $[t_1, t_2]$ during which estimates are taken begins after the initial stabilization phase is through. The parameters we selected are those that minimize the average of the individual errors over the swarm, defined as

$$E = \frac{1}{N} \sum_{i=1}^N E^i.$$

Furthermore, to ensure fairness, we searched for the best parameter choice of each method evaluating the same number of setups.

All our results have been obtained through extensive simulations. Implementation on real robots¹ is currently under development.

III. THE HOLLAND, MELHUISH AND HODDELL METHOD

A. Basics

Holland *et al.* proposed a modification of Mirolo and Strogatz' model in which fireflies behave like coupled oscillators [7] to allow robots to estimate the swarm's size. Each robot possesses an internal counter c ranging from 0 to C_{MAX} which is increased at each time step by a fixed quantity $\delta = 1$. When $c > C_{MAX}$ the robot emits a signal and sets $c = 0$. Depending on the value of c , a robot can be in one of the two phases of the algorithm: it is either in the *refractory* phase when $c \in [0, C_1]$ or in the *stimulated+non-stimulated* signaling phase when $c \in (C_1, C_{MAX}]$ (see Figure 1(a)).

During the *stimulated+non-stimulated* signaling phase, at each time step a robot can emit a signal with probability p . When this happens, the robot has emitted a *non-stimulated* signal. After signaling, the robot resets c to 0. A neighboring robot that perceives the signal resets c to 0 and emits a signal too. In the latter case the robot has emitted a *stimulated* signal. This causes a cascade of stimulated signals across the swarm. To avoid an infinite sequence of signal waves traversing the swarm back and forth, robots that have signaled enter into the *refractory* phase which prevents them from signaling even if a neighbor does. Therefore, robots

cycle between two successive phases: the *stimulated+non-stimulated* signaling phase, in which a robot can signal or be stimulated to signal; and the *refractory* phase, in which a robot ignores its neighbors' signals.

With this method, in a swarm of N robots, a robot should emit a non-stimulated signal every $N - 1$ stimulated signals. Each robot counts the number s_t^i of stimulated signals it has emitted since the last non-stimulated one. Therefore, $\hat{n}_t^i = s_t^i + 1$ can be used as an estimate of the group's size.

B. Discussion

Holland, Melhuish and Hoddell's method (HMH) has a strongly probabilistic nature and so the values of s_t^i over time display dramatic fluctuations. The signaling probability p plays a key role in shaping such fluctuations: for a high value of p , robots signal too frequently in a non-stimulated manner. In this way, signals are likely to overlap, thus making the estimates \hat{n}_t^i significantly lower than N . For this reason, small values of p are preferable. In their method, Holland *et al.* dampen fluctuations with a weighted average of \hat{n}_t^i over time:

$$\hat{n}_t^i = \alpha(s_t^i + 1) + (1 - \alpha)\hat{n}_{t-1}^i.$$

In their paper, Holland *et al.* suggest $\alpha = 0.85$. As reported in Table I, the minimum error E obtained with this method is more than 50%. This is the key drawback of the method, because it entails two undesired consequences. First, for an individual robot it is not possible to ultimately decide about the size of the group, thus making the estimate basically useless. Second, the swarm as a whole displays a very low degree of agreement – even averaging an individual's estimate with its neighbours' is not likely to improve the method's performance sensibly.

Furthermore, our experiments showed that the length of the *refractory* phase C_1 has no effect the error of the system, but affects its stabilization speed. High values for C_1 , in fact, oblige the robots to wait for a long time before exiting the refractory period, thus slowing down the entire process. In our experiments we set the length of this phase to the same value used in Holland *et al.* experiments: $C_1 = 20$.

Finally, the choice of a value for C_{MAX} derives from a trade off between scalability and speed. In fact, if C_{MAX} is too low, there could be not enough time for all the robots to signal, thus making the method not scalable with N . The value of C_{MAX} depends also on the probability of emitting

¹The platform we are porting our software to is the *e-puck* robot (<http://www.e-puck-robot.org>).

TABLE I

MEAN RELATIVE ERROR OF THE DISCUSSED METHODS WHEN PARAMETERS ARE SET TO OPTIMAL VALUES.

Method [Averaging Method]	C_1	C_{MAX}	p	ϵ	α	b	w	Error [%]
HMH [Time-Weighted Average]	20	260	0.01	-	0.85	-	-	58.24
Modified Method [No Average]	20	400	0.12	10	-	-	-	1.2
Modified Method [Time-Weighted Average]	20	400	0.12	10	0.85	-	-	9.60
Modified Method [Moving Average]	20	400	0.12	10	-	-	9	3.68
Modified Method [Mixed Average]	20	400	0.12	10	-	2	9	4.29

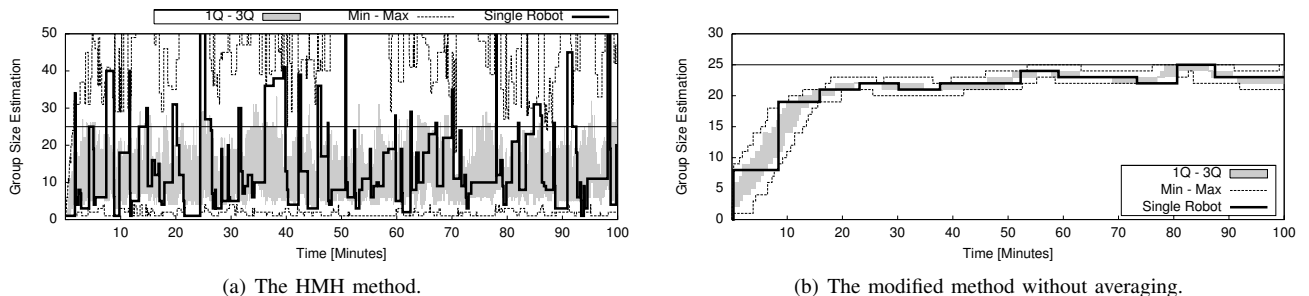


Fig. 2. Results of test runs of the method with 25 simulated robots. Each graph shows the distribution of the estimates of the group's size for each robot over time. The thin dotted line delimits the maximum and minimum estimates in the swarm; the gray area corresponds to the estimations falling into the first and third quartiles; the black solid line is an example of estimate of a single robot. A qualitative measure of agreement across the swarm is provided by the thickness of the max-min span or of the quartile span. In these graphs and in Figures 4– 6, the simulation time step lasts 100 ms.

a non-stimulated signal p , so that a high C_{MAX} value with a low probability p increases the time needed by the robots to obtain the estimate of the group's size.

IV. REDUCING THE ERROR

The HMH method provides a distributed and simple method to estimate group size, but for real applications the results are too noisy. As the above discussion demonstrated, this is due to the low degree of stability of the individual estimates over time caused by the fact that when a robot emits a non-stimulated signal, after the *refractory* phase it has the same probability as any other robot to emit another non-stimulated signal.

Instead, it would be desirable if those robots that have recently emitted a non-stimulated signal were less likely to emit another one, while robots that did not have the occasion to emit a non-stimulated signal were more likely to emit one. This way, the estimates would be less noisy, because the probability of having exactly $N - 1$ stimulated signals between two non-stimulated ones would be higher.

What would be needed, in other words, is an ordering mechanism that lets each robot emit a non-stimulated signal every $N - 1$ stimulated emissions. However, forcing a strict signaling order requires a distributed agreement protocol and more powerful communication abilities than those assumed for our robots [4]. With the simple signals our method can afford, we can obtain an approximated but satisfactory ordering with a simple modification to the original method.

The key idea is the following: between the *refractory* phase and the *stimulated+non-stimulated* signaling phase, we introduce a *stimulated only* signaling phase in which a robot can emit a stimulated signal, but has zero probability of emitting a non-stimulated one. Therefore, as shown in Figure 1(b), a robot is in the *refractory* phase when

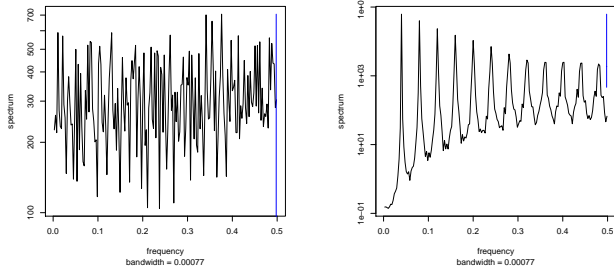
$c \in [0, C_1]$, in the *stimulated only* signaling phase when $c \in (C_1, C_2]$ and in the *stimulated+non-stimulated* signaling phase when $c \in (C_2, C_{MAX}]$. Whenever a robot emits a signal (stimulated or not), it resets c to 0.

Furthermore, while C_1 remains constant over time and is set to the same value for all robots, the value of C_2 is allowed to vary. The value of C_2 is maximum and equal to C_{MAX} for a robot that emitted a non-stimulated signal in the previous time step. This means that a robot that has just signaled does not pass through the *stimulated+non-stimulated* signaling phase. After each complete cycle of c , C_2 is gradually decreased subtracting a quantity ϵ . Clearly, at the same time, the length of the *stimulated+non-stimulated* signaling phase increases by the same quantity. As a result, every robot has a different length of the *stimulated only* signaling phase and the *stimulated+non-stimulated* signaling phase, proportional to the time elapsed since the last non-stimulated signal. The decrease of C_2 continues until it reaches C_1 : in this case, the *stimulated only* signaling phase disappears and the phase cycle that results is the same as in the HMH algorithm. The swarm is initialized with $C_2 = C_1$ and $c = 0$.

In this mechanism, the choice of ϵ , C_1 and C_{MAX} are critical. Together, they characterize the number of phase cycles τ needed for the *stimulated only* signaling phase to disappear:

$$\tau = \frac{C_{MAX} - C_1}{\epsilon}$$

With a value of τ that is too low, the *stimulated only* signaling phase disappears quickly and its benefits vanish – the method behaves similarly to the HMH's and suffers the same drawbacks. Hence, a slow decrease is preferable, as it lets room for each robot to signal in a non-stimulated manner. In general, the larger the swarm size, the larger τ should be to obtain good estimates \hat{n}_t^i . As discussed in Section I, in



(a) Spectrum of the HMH method. (b) Spectrum of the modified method without averaging.

Fig. 3. Spectral analysis of the sequence of robot ID signaling in a non-stimulated way. Since the robot ID is immaterial, the sequences are analyzed as symbolic time series. The total Fourier spectrum of the symbolic sequence is defined as the sum of the squared modulus of the individual indicator sequence spectra [1].

the typical usage scenario we assume that the initial size of the group is known, so τ can be adjusted to a reasonable value. Furthermore, as casualties due to failures reduce the group's size, τ increases, thus improving the quality of the estimates.

The insertion of the *stimulated only* signaling phase renders the choice of p noncritical. As discussed with more detail in Section V-C, although an optimal value for p exists, even suboptimal values do not result in a significant degradation of performance.

Experiments show that the proposed modification alone, without averaging, significantly improves the quality of the estimation. Table I reports the mean relative error E obtained with the best parameter values we found and Figure 2(a) shows the results obtained with the HMH method, while the proposed modification is depicted in Figure 2(b). There is also another interesting way to analyze the effect of our modification. Consider the series $S(k)$ of the IDs of the robots that (over time) emit a non-stimulated signal (i.e., $S(k) = i$ if the swarm's k -th non-stimulated signal has been emitted by robot i). Our modification imposes on $S(k)$ a periodicity. The spectrum in Figure 3(b) shows a large peak for frequency $1/N = 1/25 = 0.04$ and subsequent smaller peaks for frequencies $k/N, k \in \mathbb{N}, k > 1$. Each peak corresponds to an increasing part of the series $S(k)$ overlapping with itself. The spectrum corresponding to the HMH method (Figure 3(a)) does not display that regularity.

To provide a qualitative comparison of individual stability between the HMH method and our modified version, in Figure 2 we show the behavior of a single individual picked at random. In an analogous way, agreement is visualized plotting the maximum and minimum values of the group size estimate in the swarm, and the first and third quartiles. The two graphs show that stability and agreement in our modified version are significantly better than those of the HMH method. However, in the modified method, individual estimates still display slight fluctuations around the equilibrium value. Fluctuations, albeit small, are undesirable because they prevent the robot from ultimately deciding

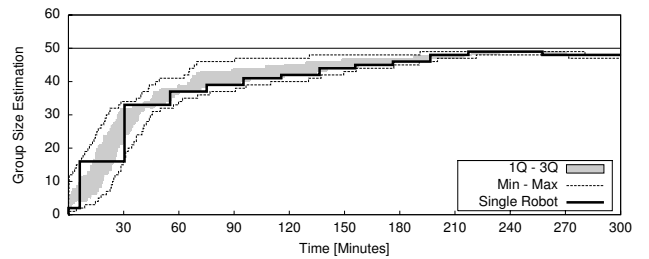


Fig. 4. Result of a test experiment with 50 simulated robots.

about the swarm's size. For a robot to take a decision based on the swarm's size, the estimation needs to stay constant for a sufficient number of successive observations, beyond which the robot considers the estimate definitive. As discussed in Section V-D, averaging can help reduce fluctuations, thus making the estimates more stable.

V. PROPERTIES

A. Scalability

In the simulated experiments shown in Figure 2, 25 simulated robots are used. We are currently studying the scalability properties of the modified method and the preliminary results we obtained running our method with an increasing number of robots are very promising. As an example, in Figure 4 we report the results of an experiment in which 50 robots are used in simulation. The results are obtained with $C_1 = 20$, $C_{MAX} = 700$, $p = 0.12$ and $\epsilon = 10$. The swarm stabilizes around the correct value with very high levels of stability and agreement across the swarm.

B. Adaptivity

The modified method proves also to be adaptive to changes in the swarm size. Figure 5 shows the results of a test experiment where initially 25 robots are used. After the estimates across the swarm successfully stabilize around the right value, a second phase starts in which 5 robots are added (Figure 5(a)) and 10 robots are removed (Figure 5(b)). These results are obtained using the same setup as in Section V-A. In the first case (group size is increased), the system behaves slightly worse than in the second case (when group size is decreased). This is due to the fact that increasing the group size lowers τ , and vice versa. As explained in Section IV, a higher τ highlights the benefits of the *stimulated only* signaling phase, thus ensuring better adaptivity. Nevertheless, even without further parameter tuning, the method shows a good adaptivity to changes in group size.

C. Robustness

As the aim of this work is to provide a reliable method for group size estimation, we studied robustness to parameter variation. We studied the effect of perturbing the probability to signal p around its best found value 0.121 and the decrease quantity ϵ around its best found value 10. Table II summarizes the mean relative error obtained in each of these setups. The results show that the error does not change considerably with slight modifications of the parameters,

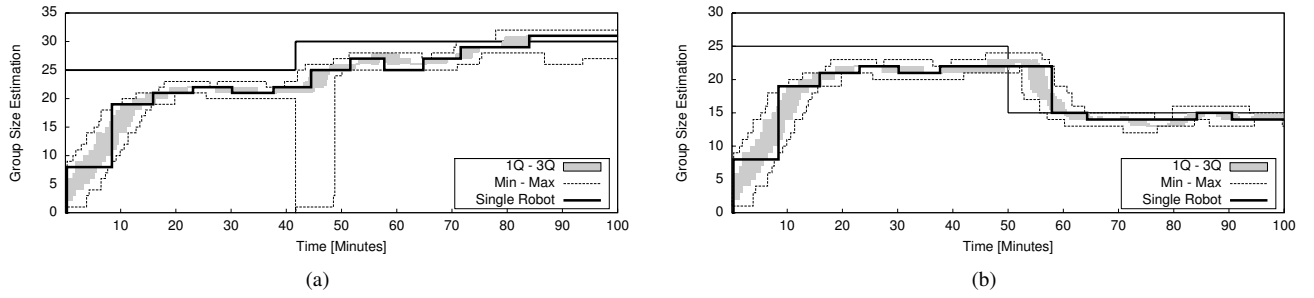


Fig. 5. Adaptivity tests. In these two-phase experiments, the swarm size is 25 during the first phase. In the second phase, the swarm size is changed to (a) 30 and (b) 15.

TABLE II

ROBUSTNESS – MEAN RELATIVE ERROR OF THE MODIFIED METHOD FOR DIFFERENT VALUES OF p , THE PROBABILITY TO SIGNAL, AND OF ϵ , THE DECREASE RATE OF C_2 PER CYCLE.

		0.115	0.116	0.117	0.118	0.119	p					
							0.120	0.121	0.122	0.123	0.124	0.125
	8	2.52	6.04	5.68	3.84	8.8	6.12	8.12	6.28	4.56	9.92	2.96
	9	1.92	2.84	4.56	6.2	1.96	6.88	1.96	1.56	2.96	4.28	2.04
ϵ	10	2.76	3.92	3.52	1.8	1.36	1.2	0.92	3.24	1.16	3.16	4.64
	11	2.48	4.72	3.76	3.96	1.56	3.12	3.44	3.12	2.28	1.08	3.2
	12	2.24	2.76	3.56	3.76	1.12	3.52	4.16	4.28	3.96	4.6	3.8

thus giving good performance even in case of suboptimal parameter settings.

D. Stability

To improve the stability of individual estimates \hat{n}_t^i , we have tried several averaging strategies. Table I reports the mean relative errors obtained with the three different averaging strategies we studied. As can be seen, the price to pay to improve stability is a slightly increased error and less adaptivity (i.e., the system is slower to react to changes of group size). The choice of whether or not to privilege precision over stability depends on the situation in which the method is employed.

1) *Time Weighted Average*: The average used by Holland *et al.* is a time-weighted average process based on the past history of the system. At every new estimation, the averaged output is computed adding the most recent values (multiplied by a constant α) and the previous average (multiplied by a constant $1 - \alpha$). Parameter α influences the trade off between adaptivity and stability. The lower the parameter, the more stable the output. The price to pay for stable estimates is that if the group size changes (e.g., due to malfunctioning robots) longer time is needed to stabilize to a new value.

An undesired effect of this averaging strategy is that the robots underestimate the real group size. This is due to the fact that this way of averaging takes into account the complete history of the estimates, and that at the beginning of the process, all robots' estimates are initialized to 1, i.e.,

$$\forall i \quad \hat{n}_0^i = 1.$$

Over time, the estimates progressively grow up to the stable value. Given that the past history is always considered during the process, future estimates are always affected by the initial low values, thus causing the system to underestimate the

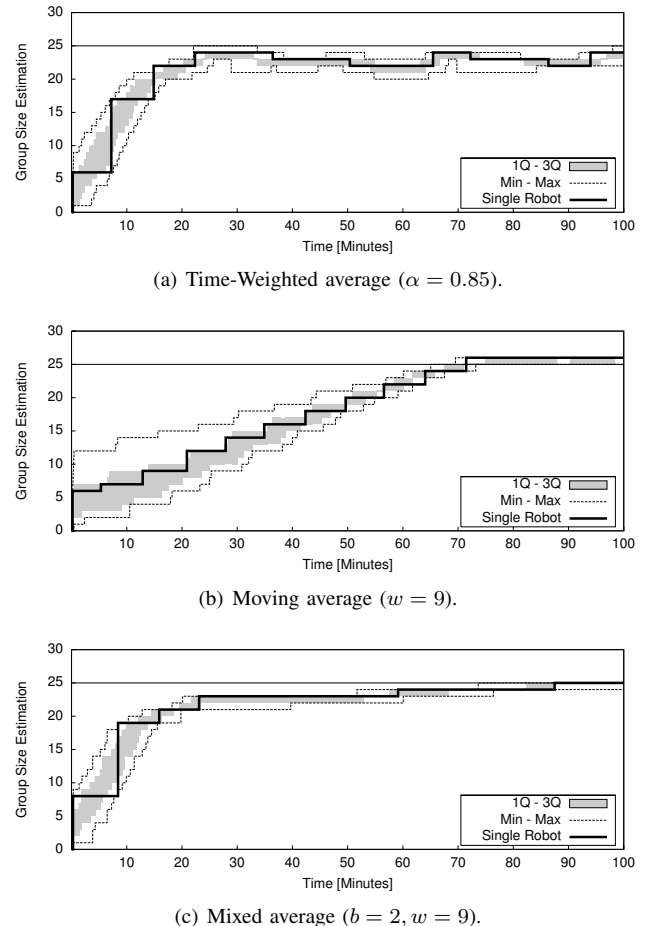


Fig. 6. Comparison of different averaging strategies for the modified method.

group size. A test experiment with this kind of averaging is reported in Figure 6(a).

2) *Moving Average*: To eliminate the underestimation effect caused by the time-weighted average, a moving average is a possible solution. With this strategy, the next estimate of a robot is calculated averaging its last w observations s_j^i :

$$\hat{n}_t^i = \frac{1}{w} \sum_{j=t-1}^{t-w} (s_j^i + 1)$$

This average solves the underestimation problem, and stability and agreement improve as w increases. On the other hand, increasing w makes stabilization time increase as well, because w estimates are needed to fill the averaging window. Hence, the moving average introduces a trade off between stability and stabilization speed. In our experiments we obtained satisfactory results with the averaging window $w = 9$. Results are shown in Figure 6(b).

We also tested a weighted moving average, in which weights are higher for recent observations and lower for older ones. We have computed the average using both linear and exponential weights, noticing no significant improvement.

3) *Mixed Average*: The trade off between stability and stabilization speed introduced by the moving average can be solved with a mixed strategy. Figure 6(c) shows that, at the beginning of the process, estimates grow fast towards the stable value. It is only after this phase that averaging is needed. Introducing the moving average at this point allows the system to settle fast around a value and dampen fluctuation from then on, thus retaining the good levels of stability and agreement of the pure moving average.

The benefit of this averaging approach depends on the length of the initial phase when averaging is off. Since robots perform their estimations \hat{n}_t^i only when they signal in a non-stimulated manner, the length of the initial phase can be expressed by the number of non-averaged estimates b before averaging is turned on. The value of b must be high enough to allow the robots to settle around a good estimate. In our experiments we found that $b = 2$ is a reasonable value.

For what concerns the averaging period w , the discussion in Section V-D.2 still holds. Therefore, $w = 9$ is a reasonable value in this case too. The results obtained with this average are shown in Figure 6(c).

VI. CONCLUSIONS AND FUTURE WORK

In swarm robotics, and more generally in distributed systems, the ability of group members to estimate the group's size is often useful. In this work, we propose a method that lets each robot in a swarm estimate the group's size in a distributed manner by exploiting only local information. We base our work on the assumption that the robots are only able to send and receive signals within a limited range.

Our method is a modification of the one proposed by Holland *et al.*, who were the first to propose to use local robot signaling for estimating the group's size from signaling frequency. The HMM method is very interesting, but its mean relative error is considerable. We identified the reason for the

problem as being the fact that the robots do not signal in an ordered way, thus causing great deals of noise.

We modified the HMM method by imposing a partial ordering on robot signaling, while retaining the distributed nature of the original method. Our method outperforms the original one, both in terms of precision and in terms of overall stability of the estimate and agreement across the swarm.

We are currently deepening our understanding of the proposed method. Preliminary results show that the method is adaptive to changes of group size, scalable when applied to larger swarms and robust, i.e., the performance degrades gracefully for small perturbations of the parameter values.

We are planning to test our method on *e-puck* robots. At present, in our simulations we do not model message loss due to noise or interference, and we believe that the next major point will be coping with this issue.

Acknowledgments. This research was carried out in the framework of Swarmanoid, a project funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission under grant IST-022888. Mauro Birattari and Marco Dorigo acknowledge support from the Belgian F.R.S.-FNRS, of which they are Research Associate and Research Director, respectively.

REFERENCES

- [1] V. Afreixo, P.J.S.G. Ferreira, and D. Santos. Fourier analysis of symbolic data: A brief review. *Digital Signal Processing*, 14:523–530, 2004.
- [2] J. Casper, M. Micire, and R.R. Murphy. Issues in intelligent robots for search and rescue. In *Unmanned ground vehicle technology II*, pages 292–302. SPIE, Bellingham WA, USA, 2000.
- [3] S. Curtis, J. Mica, J. Nuth, G. Marr, M. Rilee, and M. Bhat. ANTS (Autonomous Nano-Technology Swarm): An artificial intelligence approach to asteroid belt resource exploration. International Astronautical Federation, 51th Congress, 2000.
- [4] X. Défago, A. Schiper, and P. Urbán. Total order broadcast and multicast algorithms: Taxonomy and survey. *ACM Computing Surveys*, 36(4):372–421, 2004.
- [5] O. Holland, C. Melhuish, and S. E. J. Hoddell. Convoying: using choring for the formation of travelling groups of minimal agents. *Robotics and Autonomous Systems*, 28:207–216, 1999.
- [6] K. Lerman and A. Galstyan. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, 13(2):127–141, 2002.
- [7] R. E. Mirollo and S. H. Strogatz. Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6):1645–1662, 1990.
- [8] F. Mondada, M. Bonani, A. Guignard, S. Magnenat, C. Studer, and D. Floreano. Superlinear physical performances in a SWARM-BOT. In *Proceedings of the VIIIth European Conference on Artificial Life*, pages 282–291. Springer, Berlin, Germany, 2005.
- [9] E. Sahin. Swarm robotics: From sources of inspiration to domains of application. In Erol Sahin and William Spears, editors, *Swarm Robotics Workshop: State-of-the-art Survey*, number 3342 in Lecture Notes in Computer Science, pages 10–20, Berlin, 2005. Springer.
- [10] H.M. Smith. Synchronous flashing of fireflies. *Science*, 82(2120):151–152, 1935.