Proceedings of the ICAR 2009
14th International Conference on Advanced Robotics
Munich, Germany, June 22 - 26, 2009

Self-Organised Recruitment in a Heteregeneous Swarm

Carlo Pinciroli, Rehan O'Grady, Anders Lyhne Christensen, Marco Dorigo

Abstract— **In this paper, we present a heterogeneous recruitment system which allows aerial robots to recruit groups of wheeled robots. The system is novel because it is self-organized, is based only on simple probabilistic rules and relies only on local communication. Our approach is inspired by the aggregation behavior of cockroaches. The system allows aerial robots to recruit wheeled robot groups of different sizes in parallel. Although governed by probabilistic rules, we show that our system stabilises to a steady state distribution of wheeled robots that corresponds to the desired sizes of recruited groups. The system can also handle perturbations as robots enter and leave the system during execution—robots are redistributed accordingly and the system will once again stabilise. The system also displays sensible behaviour when there are not enough wheeled robots to satisfy the recruitment quotas. We conduct experiments based on physically embodied simulation of the robots to demonstrate these properties of the system, and to explore the scalability of our system.**

I. Introduction

In this paper, we consider a heterogeneous swarm of flying and wheeled robots, called *eye-bots* and *foot-bots* respectively. When this heterogeneous swarm is deployed into an environment, the eye-bots search for tasks, and then recruit groups of foot-bots to perform the various tasks they have found. The size of the recruited groups varies depending on the size and nature of the tasks.

In this study we focus on the recruitment aspect of the above scenario. Designing a recruitment mechanism for a swarm of robots is non trivial. A simple recruitment solution based on point to point communication could be implemented, in which a recruiting eye-bot communicates in turn with one foot-bot at a time. However, using such a solution it is not clear how multiple eye-bots could recruit groups of foot-bots at the same time, at least not without introducing more explicit point to point communication between the eye-bots (that would then have to negotiate over individual foot-bots). Such a system would scale very badly, as the number of point to point communication messages would increase exponentially with the number of robots in the system.

Instead, we propose a distributed recruitment mechanism inspired by the aggregation behaviour of cockroaches under shelters (see Section II). Communication is restricted to strictly local broadcast communication. Eye-bots recruit groups of foot-bots by broadcasting join and leave probabilities to any foot-bot that is in an area directly beneath them. Unrecruited foot-bots perform a random walk. Once the system has reached its steady state distribution, groups of foot-bots can be deployed to perform their various tasks.

Our control is completely distributed, and we rely only on local communication. The use of such self-organising principles tends to lead to robust, scalable platforms [1]. However, the behaviour of such system is often hard to predict. In this paper, we conduct a series of physically embodied simulation experiments to confirm that our recruitment system displays desirable properties, and to test the system's scalability.

We experimentally demonstrate that our system displays the required property of *stabilisation*: the system reaches an equilibrium that is sufficiently stable that it can be detected by the individual eye-bots.

In addition, for our system to ever be useful in the real world, it would need to be effective in a dynamically changing environment, and therefore would need to cope with the arrival and departure of eye-bots and foot-bots while the system is running. We demonstrate experimentally that our system displays the required property of *redistribution*: the system is able to dynamically redistribute foot-bots among recruiting eye-bots, and once again arrive at a stable steady state.

Another constraint is that at any given moment in time, there may not be enough foot-bots available to fill the quota of every eye-bot (*quota* refers to the desired group size of a single eye-bot). We demonstrate experimentally that our system displays the required property of *balancing*: the system settles to a state where each eye-bot has filled the same percentage of its quota of foot-bots (independently of the quota size). In a system that did not display this property, small quotas might always fill up at the expense of larger quotas (or vice-versa)[1].

II. Related Work

Several centralized and communication-based techniques for recruitment and task allocation in multi-robot systems have been proposed, see for example [3], [7], [8]. In social insects such as ants and bees, recruitment plays an important role in, for instance, food source exploitation. The ants *Leptothorax acervorum*, *L. muscorum*, *L. nylanderi* and *Temnothorax albipennis* are known to be able to recruit nest mates using a combination of pheromone secretions and a guidance technique called *tandem running* [16], [4]. When

Carlo Pinciroli, Rehan O'Grady and Marco Dorigo are with IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium ({cpinciro,rogrady,mdorigo}@ulb.ac.be). Anders Lyhne Christensen is with the DCTI-ISCTE, Portugal (anders.christensen@iscte.pt).

[1]When there are not enough foot-bots, an additional deadlock resolving mechanism would need to be implemented in the future—one solution might be to let eye-bots that have stabilised at a group size that is below their quota probabilistically release their own recruited foot-bots to help other eye-bots fill their quotas. However, for such a mechanism to be implemented the system would require *stabilisation* and *balancing* as prerequisites.

performing tandem running, the recruiting ant periodically waits for the recruited ant which in turn touches the recruiter to indicate that it can continue. Krieger and Billeter [11] have demonstrated an approach inspired by tandem running on real robotic hardware. However, the technique only allows for one robot to be recruited at a time.

There is a large body of literature on heterogeneous robotic groups [19], [24]. Gage and Murphy [5] have, for instance, demonstrated how a single unmanned aerial vehicle (UAV) can recruit unmanned ground vehicles (UGVs) in a landmine detection task. However, to the best of our knowledge, no existing robotic study investigates group size regulation in heterogeneous robot groups. Group size regulation is important in tasks ranging from search and rescue, where a certain number of robots may be required to shift a victim or a heavy object [9], to rough terrain navigation where an appropriate number of robots need to collaborate in order to overcome certain obstacles [17], [2].

There is some literature on aggregation and group size regulation in homogeneous groups. Dorigo *et al.* [13] evolved two dimensional distributed aggregation in a swarm of embodied robots, Martinoli *et al.* [14] investigated the effects of probabilistic parameters on the size of object clusters collected by Khepera robots. Neither work provided an explicit group size control mechanism. Melhuish *et al.* [15] controlled group sizes in a swarm of abstracted agents using a firefly-like synchronisation mechanism. However, group size control was not fine grained to the level of individual robots, only one group was formed at a time and the physics of an embodied system was not taken into account.

Our approach is inspired by the aggregation behaviour of cockroaches. This behaviour is accurately mimicked by Jeanson's *et al.* model [10] in which cockroaches can probabilistically switch from resting to performing a random walk and back. The likelihood of resting rather than random walking increases with the number of other nearby resting cockroaches. A positive feedback mechanism then results in aggregation into groups. Other experimental evidence also shows that cockroaches prefer to aggregate in dark places [23]. When multiple dark shelters are available in the environment, the majority of cockroaches aggregate under only one shelter rather than spreading evenly (even when the shelters are identical) [12].

Garnier *et al.* [6] used Jeanson's behavioral model to show that a group of cockroach-like robots can achieve a collective choice between two different shelters in the environment through simple local interactions.

III. Robotic Platform and Simulation Environment

Our system comprises eye-bots (Figure 1(a)) and foot-bots (Figure 1(b)). Eye-bots are quad-rotor equipped aerial robots capable of flying and attaching to the ceiling. Although our overall scenario uses the flying capabilities of the eye-bots, in this study our experiments always consider the eye-bots to be attached to the ceiling. Eye-bots are equipped with a high resolution camera which allows them to monitor what happens on the ground [22]. Foot-bots are mobile robots that use a combined system of track and wheels to move. They are equipped with infrared proximity sensors, an omnidirectional camera, and an RGB LED ring that enables them to display their state to robots within visual range.

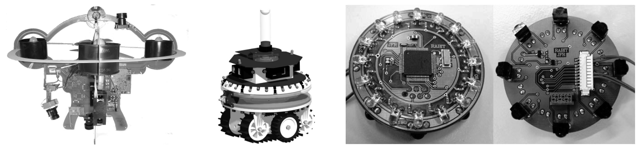

(a) Eye-bot (b) Foot-bot (c) Range and Bearing Sensor

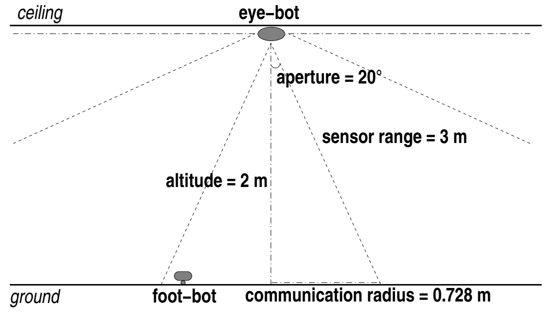

(d) Range and Bearing Communication Range

Fig. 1. Heterogeneous Robotic Platform. (a,b) The robots. (c) The range and bearing sensor. (d) Diagramatic representation of the communication range of the range and bearing sensor.

Communication between eye-bots and foot-bots is achieved through a range and bearing system [21] mounted on both robots. This system allows the robots to broadcast and receive messages either from neighbors in the same plane, or in a cone above the foot-bots or beneath the eye-bots. Furthermore, the system allows for *situated communication*, meaning that recipients of a message know both the content of the message and the spatial origin of the message (in their own frame of reference), see Figures 1(c) and 1(d).

At the time of writing, the robotic platform is still under development. The results presented in this paper have thus been obtained in simulation. A custom physics based simulator called ARGoS [20] has been developed to reproduce the dynamics of the robots' sensors and actuators with reasonable accuracy (see Figure 4(left)).

IV. Our Approach

Our approach is inspired by Jeanson's aggregation model of cockroaches under shelters [10]. Eye-bots play the role of the shelters, while foot-bots play the role of cockroaches. Previous robotic systems have mimicked the aggregation behaviour of cockroaches, thus preserving the property that the aggregate group size was passively dependent on initial conditions. Our approach is original in that we transform the stop/go probabilities into join/leave group probabilities computed and transmitted by the eye-bots, thus allowing the system to actively control the distribution of recruited foot-bot group sizes.

In our system, all recruitment takes place in a central recruitment area. Foot-bots not currently engaged in task

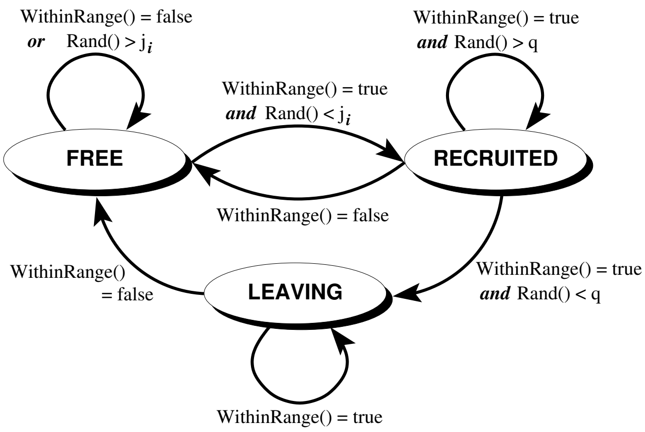

Fig. 2. State transition logic for foot-bots at each time step. `WithinRange()` is a function returning *true* when the robot is within the communication range of an eye-bot. `Rand()` is a function returning a random number in $\mathcal{U}(0,1)$. j_i is the join probability for eye-bot i, q_i is the common disaggregation probability.

execution gather in the recruitment area. Eye-bots return to the recruitment area to recruit groups of foot-bots. Foot-bots move much more slowly in the environment than the flying eye-bots. The centralised recruitment area thus is much less expensive in terms of both time and expended energy than any alternative system of allocating foot-bots to tasks that would require the foot-bots to spread in the environment.

Foot-bots can be in one of three states:

- state FREE: The foot-bot is not recruited. The foot-bot performs random walk with obstacle avoidance. This state is signaled with the foot-bot's LEDs lit up in green.
- state RECRUITED: The foot-bot has been recruited by an eye-bot. This state is signaled with the foot-bot's LEDs lit up in red.
- state LEAVING: The foot-bot is leaving a group of recruited foot-bots to which it previously belonged. This state is signaled with the foot-bot's LEDs lit up in blue.

Eye-bots influence the size of their recruited group of foot-bots by actively changing the join and leave probabilities that they broadcast. Each eye-bot has a *quota* (i.e., a desired group size) of foot-bots that it is trying to recruit. The join probability j_i for each eye-bot is set proportional to the size of its quota. Once its quota is filled, the eye-bot prevents any more foot-bots from joining its aggregated group by setting its join probability to zero ($j_i = 0$). Figure 2 shows the state transition logic of foot-bots upon receipt of a message from an eye-bot.

Eye-bots obtain the number of foot-bots aggregated beneath them by using their cameras to count the number of foot-bots in state RECRUITED (i.e., lit up in red).

Through mathematical analysis and experimentation conducted in a previous study [18], we established an appropriate way to set the leave probabilities q_i for each eye-bot. We found that high values of q_i led to a highly flexible system, but discouraged system stability. On the other hand, low values for q_i led to stable systems, but discouraged *redistribution*: the redistribution of foot-bots through the system was difficult as eye-bots rarely released recruited foot-bots. Our solution is to let the eye-bots vary q_i between two values while the system is running—a high value ($q_{max} = 0.05$) that encourages *redistribution*, and a low value ($q_{min} = 0.00001$) that encourages *stability*. Parameter q_i is 'spiked' to the high value when the system starts, and whenever the system is perturbed by the arrival or departure of an eye-bot. After an eye-bot spikes its q_i value, the q_i value is then linearly decayed over T_{dec} time steps until it arrives at q_{min}, at which it remains until the next spike. Modifying q_i dynamically in this way allows the system to redistribute effectively at moments when the fast transfer of foot-bots is desirable, but nonetheless to stabilise as q_i decreases.

To detect stabilisation, each eye-bot continually checks if the fluctuations in the number of its recruited foot-bots have stayed within some tolerance boundary for a given length of time. A more detailed discussion of the parameter settings we have chosen, an analysis of the effects of the various parameters and the definition of the stabilisation boundary condition can be found in [18].

V. Two Robot Experiments

We conducted two sets of experiments in a rectangular arena with 30 foot-bots and 3 eye-bots (see Figure 3). In the first set of experiments we used eye-bot quotas of 15, 10 and 20. In the second set of experiments, we used eye-bot quotas of 12, 12 and 12. There are three phases in each experiment. In the first phase, only eye-bot 1 and eye-bot 2 are active. This first phase allows us to test whether the eye-bots recruit the correct number of foot-bots in the simple case where there are enough foot-bots to satisfy all eye-bot quotas. The second phase starts at time T_1 with the activation of eye-bot 3. Eye-bot 3 now has a quota equal to or greater than the quotas of the other eye-bots in the system, thus providing a good test for the *redistribution* capability of our system. In addition, with the addition of eye-bot 3, the sum of the quotas becomes greater than the total number of foot-bots in the arena. We can therefore also test the *balancing* property of the system. The third and final phase is initiated at time T_2 by the disactivation of eye-bot 2, and again tests the *redistribution* and *balancing* properties of the system, this time in response to the addition of more free foot-bots into the arena. *Stabilisation* can be tested in each of the three phases of the experiment.

In both quota sets, eye-bots 1 and 2 successfully stabilise to the correct number of foot-bots between time step 0 and time step T_1. After the introduction of eye-bot 3, the system on average quickly settles to an equilibrium where all three eye-bots have filled the same percentage of their quota, showing that the system also displays the *balancing* property.

VI. Scalability Experiments

A. Experimental Setup

To test the scalability of our system, we set up experiments with larger numbers of eye-bots. For simplicity, we consider

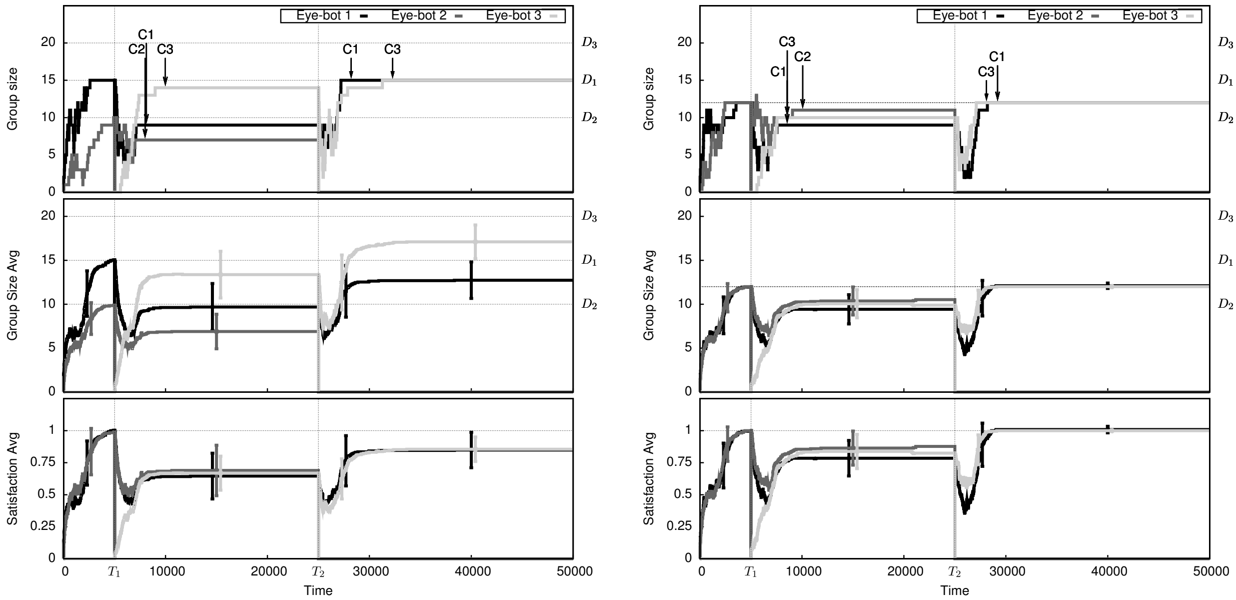

(a) Quotas: 15,10,20

(b) Quotas: 12,12,12

Fig. 3. Results of two-robot experiments. Both groups of 3 plots (a,b) describe the results of a set of 20 experiments. The top graph represents a single representative experiment that we selected from the set of 20, the middle graph shows the average results of all experiments (bars at selected time steps indicate the standard deviation), the bottom graph shows the percentage of each eye-bot's quota that is filled (here referred to as *satisfaction avg*). In our simulations, a single time step corresponds to 100 ms of real time. In every experiment, eye-bots 1 and 2 are introduced at time step 0, eye-bot 3 is introduced at time step $T_1 = 5,000$, eye-bot 1 leaves the experiment at time step $T_2 = 25,000$. Eye-bot quotas are indicated by symbols D_1, D_2, D_3. Symbols $C1, C2, C3$ indicate the moments at which the correspondingly numbered eye-bot detected system stabilisation. The decaying period of q_i is set to $T_{dec} = 2,000$.

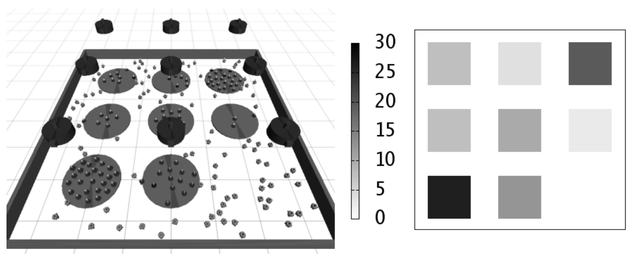

Fig. 4. Snapshot from scalability experiments. Left: Simulation snapshot. Right: Abstracted representation of this simulation snapshot—the grey intensity level of each square is proportional to the recruited group size of the correspondingly positioned eye-bot (i.e., to the number of foot-bots recruited by that eye-bot).

the centralised recruiting area to consist of a square grid of eye-bots (we use varying numbers of eye-bots in our different scalability experiments). A snapshot from one of our experiments is shown in Figure 4.

In all of the experiments in this section, every eye-bot has a recruitment quota of 25 foot-bots to fulfil. Although this quota parity would not be very likely in a real deployment scenario, this simplification allows us to concentrate our analysis on other properties of the system, without being distracted by the role of different quota sizes on our results.

B. Stabilisation and Balancing

In this section, we describe a series of experiments that we ran to test the *stabilisation* and *balancing* properties of our system with larger numbers of eye-bots. For each set of experiments in this section, the number of foot-bots in the system is set to 20 × the number of eye-bots. Therefore, when foot-bots are distributed in a balanced way, we would expect the system to stabilise at group sizes of around 20 foot-bots per eye-bot. Each experiment is run for a total of 7,500 time steps, each step being 100 ms long, for a total real time of 750 s.

The results for 16 eye-bots and 25 eye-bots can be seen in Figure 5. In the sample runs (top plots) for both the 16 eye-bot and 25 eye-bot experiments, the snapshots (grids of grey squares) show that the system is growing in a balanced way. The grey intensities for all of the squares in any particular snapshot are quite homogeneous. The grey intensities get darker as the experiment continues, corresponding to the growing recruited group sizes.

In these experiments, T_{dec} is set to 3,500 time steps (350 s). This explains why during the first 4,000 steps (400 s) of the experiment, the size of the recruited groups fluctuate considerably. This is evident in the central plots of Figure 5 which show a sample experiment taken at random among the 80 that have been run. Despite the initial fluctuations due to the system's initialisation of $q_i = q_{max}$ the system correctly stabilises, and the last 2,500 steps of the simulations are smooth in the individual experiment as well as in the averaged results.

C. Redistribution

In this section, we describe a series of experiments that we ran to test the *redistribution* property of our system with

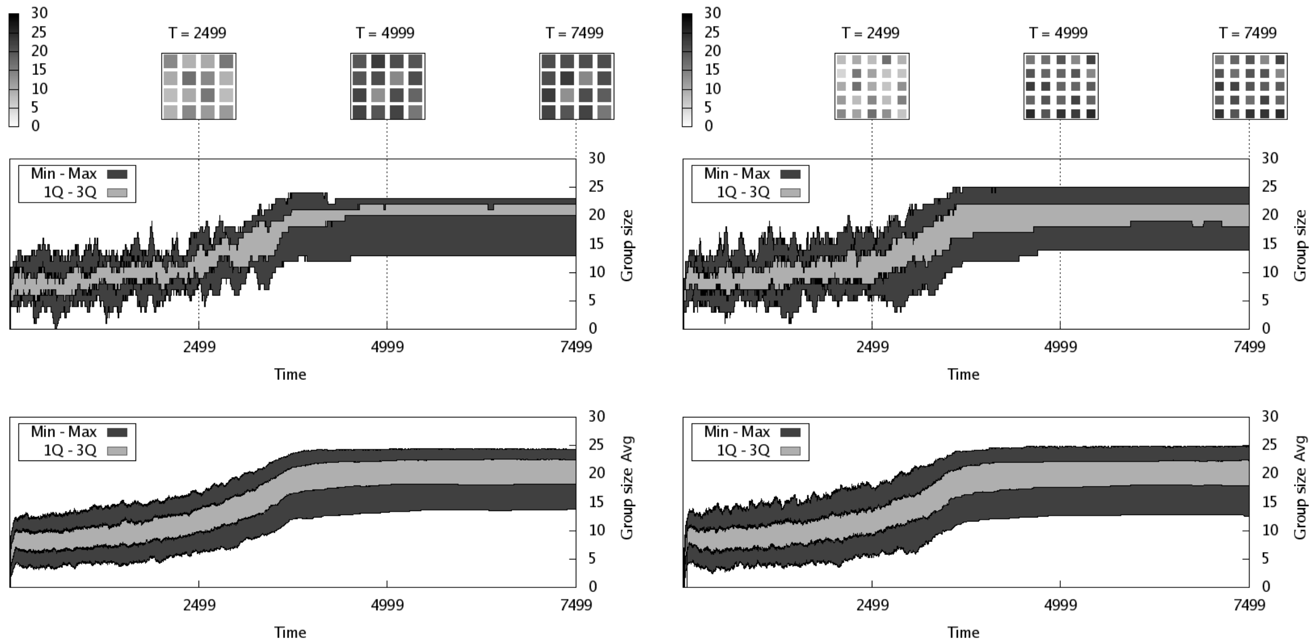

(a) 16 eye-bots in a 4x4 recruiting area. 320 foot-bots.

(b) 25 eye-bots in a 5x5 recruiting area. 500 foot-bots.

Fig. 5. Scalability experiments testing the *stabilisation* and *balancing* properties of the system. Results are shown for two sets of experiments with 16 eye-bots (a) and 25 eye-bots (b). 80 experimental runs per set of experiments. The top plots show the behaviour of the system in a single sample experiment that we have selected. The grids of squares represent snapshots of the state of the system at given moments in time during this sample experiment. The grey intensity of each individual square corresponds to the number of foot-bots recruited at that time by a single eye-bot. The min and max lines show the size of the largest recruited group of foot-bots and the size of the smallest recruited group of foot-bots at any given moment. The 1Q–3Q area shows the inter-quartile range of the distribution of recruited group sizes among the eye-bots. The 1Q is the first quartile and shows the minimum recruited group size once we discard the lowest 25% of groups. The 3Q line is the third quartile, and shows the maximum recruited group size once we discard the highest 25% of the data. The bottom plot shows the same data averaged over all 80 runs.

larger numbers of eye-bots. To do this, we run experiments in which we perturb the system while it is running and see how well the system succeeds in redistributing foot-bots in response to the perturbation.

As these experiments are longer and more computationally expensive, we use a 3x3 formation of 9 eye-bots. We perturb the system by either activating or disactivating an eye-bot at time step 7,500 (after 750 s). The activation of an eye-bot corresponds in a dynamic real world scenario to the arrival of a new recruiting eye-bot in the recruitment area. The disactivation of an eye-bot results in the release of more free foot-bots into the system, and thus corresponds in a dynamic real world scenario to the return of foot-bots that have successfully completed a task to the recruitment area. We always perturb either the central eye-bot in the 3x3 formation or the eye-bot in the lower right hand corner.

In these experiments, we assume that at the moment an eye-bot either activates or disactivates, it sends a signal to its immediate neighbours (those eye-bots within a radius of 2.828 m). This signal induces the neighbouring eye-bots to spike their q_i. These immediate neighbours also propagate the signal, and in this way, the entire swarm spikes q_i to q_{max}. The perturbation effects of an eye-bot activation or disactivation event are thus applied throughout the whole swarm of eye-bots.

The results of these experiments can be seen in Figure 6. In each of the four experimental sets, both the snapshots from the sample run and the average dynamics show that the system stabilises before the perturbation event. More interestingly, the final snapshots in each of the four experimental sets are relatively homogeneous in their grey intensity levels, showing that in the sample runs the system succeeded in redistributing foot-bots. The relatively narrow max-min and inter-quartile ranges in all four experiments indicate that this was also true on average.

Looking at the averaged data (4 bottom plots), we can see that the type of perturbation event seems to play a more important role in the system's ability to redistribute foot-bots than the location of the eye-bot that caused the perturbation event. In particular, the max-min and inter-quartile ranges for the experiments in which an eye-bot disactivated (redistribution of foot-bots released from disactivated eye-bot to the 8 other eye-bots) are much closer together than the corresponding ranges for the experiments in which an eye-bot activated (redistribution of foot-bots from 8 already active eye-bots to newly activated eye-bot). In contrast, for the same type of perturbation event (activation or disactivation) there seems to be little difference between the average max-min and inter-quartile ranges for the experiments in which a corner eye-bot was perturbed or the center eye-bot was perturbed.

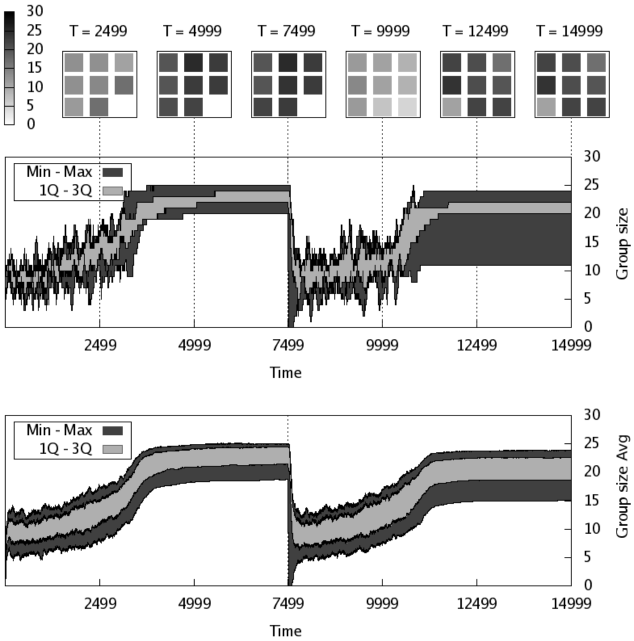

(a) Corner Robot Activates at Time step 7,500

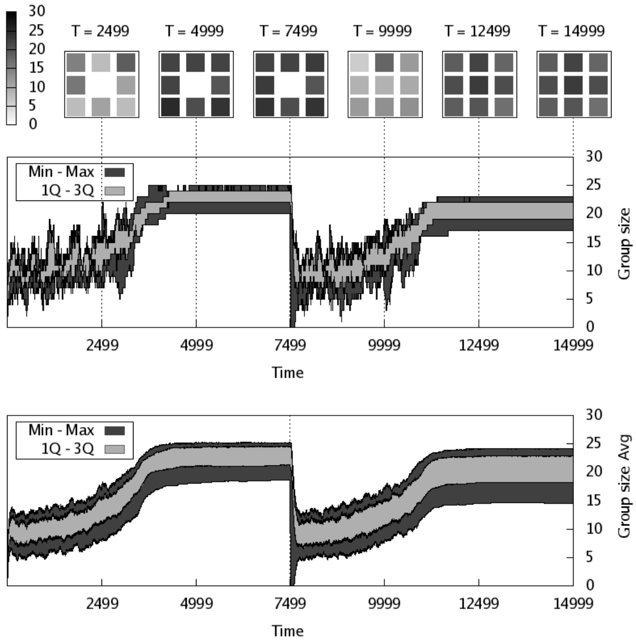

(b) Centre Robot Activates at Time step 7,500

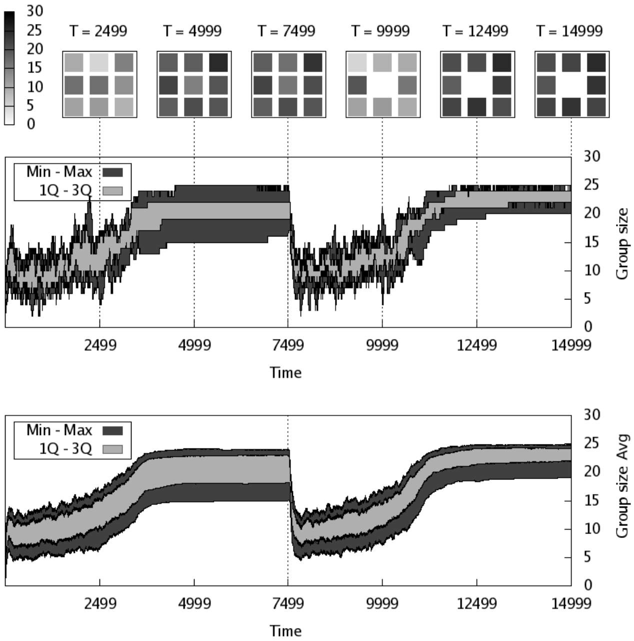

(c) Centre Robot Disactivates at Time step 7,500

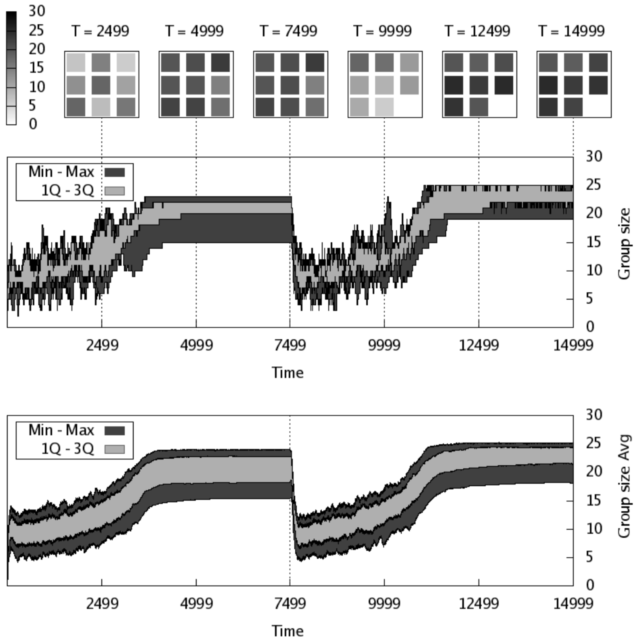

(d) Corner Robot Disactivates at Time step 7,500

Fig. 6. Set of experiments testing the *redistribution* property of the system. All experiments run with 9 eye-bots and 180 foot-bots in a recruitment area consisting of a 3x3 eye-bot formation. Results are shown for four sets of experiments (a,b,c,d). Each experimental run lasts for 15,000 time steps (1,500 s). 20 experimental runs were conducted for each set of experiments. Top plots in each set represent selected sample runs, while bottom plots represent data averaged over all 20 runs. For a more detailed explanation of the plots see previous caption from Figure 5.

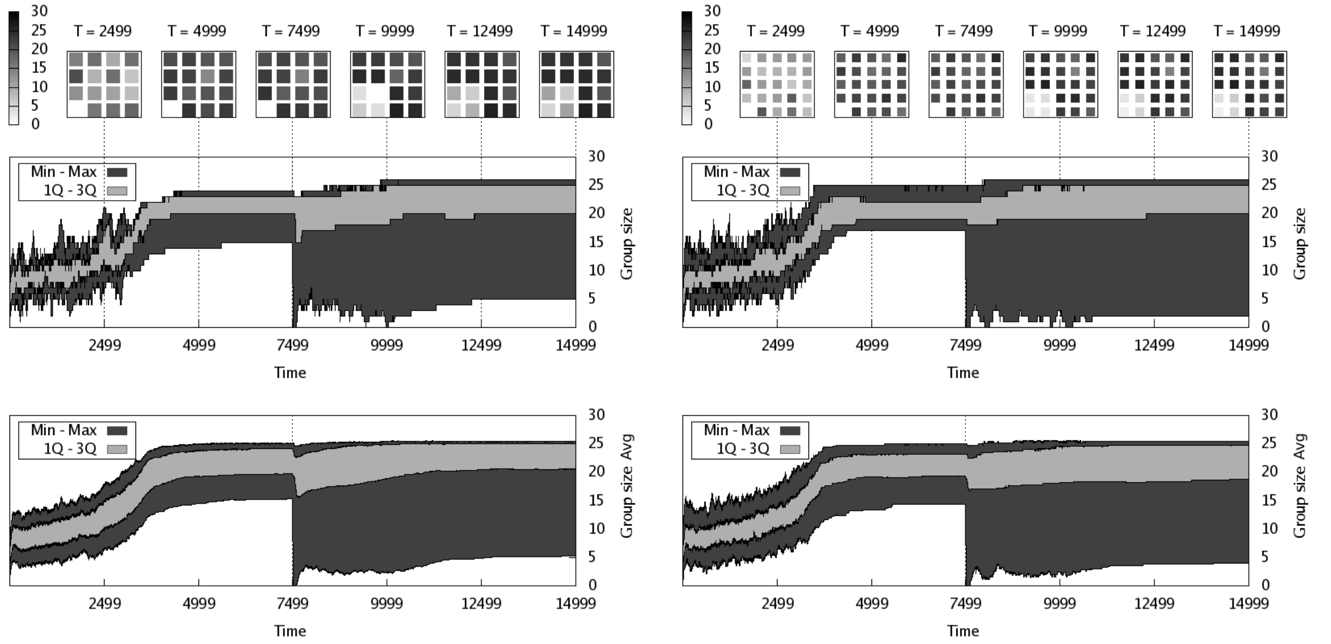

(a) (b)

Fig. 7. Set of experiments on local perturbation. In these experiments, the propagation of the q-spiking signal is limited to the direct neighbours. Results are shown for two sets of experiments (a,b). 20 experimental runs were conducted for each set of experiments. Each experimental run lasts for 15,000 time steps (1,500 s). Top plots in each set represent selected sample runs, while bottom plots represent data averaged over all 20 runs. For a more detailed explanation of the plots see previous caption from Figure 5.

D. Ongoing Work

In the previous section, knowledge of a perturbation event (i.e., an eye-bot activating or disactivating) propagated through the whole system. As a result, every eye-bot spiked its q_i probability at almost the same moment.

This type of global propagation, however, only allows for a minimum interval between new perturbation events into the system that is equal to or greater than T_{dec}. (Perturbation events include the arrival of new recruiting eye-bots in the recruitment area and the return of foot-bots to the recruitment area that have completed their task.) Ignoring this constraint on the maximum event rate would mean that the system would never stabilise, as all of the q_i rates would spike again before q_{min} was reached.

In larger dynamic environments with many tasks, however, this constraint might be very limiting. One solution might be to implement some kind of perturbation event queueing system, but this would just shift the problem, as the queues would grow longer and longer if there were too many tasks in the environment.

We are currently working on ways of solving this problem. Our first attempt involves limiting the propagation of the q-spiking signals to a locally limited neighbourhood of eye-bots centered around the eye-bot that created the perturbation event. In this way, only a small portion of the eye-bots in the recruitment area would be perturbed by incoming requests, thus letting the rest of the system stabilise undisturbed.

Preliminary experiments with 4x4 and 5x5 grids show that this method effectively lets a small portion of the system balance, as shown in Figure 7. In these experiments, the corner eye-bot is initially deactivated. Similarly to the experiments in the previous section, after 7,500 time steps the eye-bot switches to the active state and sends a q-spiking signal to its neighbours. However, in these new experiments, eye-bots that receive the signal do not propagate the signal to their neighbours. The only eye-bots that spike their q_i are therefore the immediate neighbours of the eye-bot that perturbs the system when it activates.

As the results show, at the final steady state the distribution of foot-bots is clearly divided in two distinct zones: one zone around the eye-bot that perturbs the system, and a second zone that includes all the other eye-bots (i.e., those that did not spike their q_i). The foot-bot distribution in the second (undisturbed) zone is practically identical at time steps 7,499 (steady state just before eye-bot activation) and at time step 14,999, showing that the effects of the perturbation event have remained largely local. In contrast, the homogeneous grey levels in the first zone show that local redistribution has been effective.

VII. Conclusions and Future Work

In this paper we proposed a novel self-organised recruitment mechanism for heterogeneous aerial and wheeled robots. Our recruitment system is self-organised and relies only on local communication.

We ran experiments with simulated physically embodied agents. We showed that our system is stable, and can redistribute robots when perturbed. The system also responds sensibly in the case when there are not enough robots to fulfil all of the recruitment quotas. We ran scalability experiments with up to twenty-five aerial robots and five-hundred wheeled robots and demonstrated that our system continued to perform as we increased the number of robots.

We believe that the underlying dynamics of our system are sufficiently simple that they could be implemented in other heterogeneous robotic platforms. To this end, we think it would be interesting to try other less explicit communication modalities (e.g., light intensity) as a means of transmitting probabilities.

Longer term goals involve embedding our system as part of a more complete task execution scenario, of the type outlined in the introduction.

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. Marco Dorigo acknowledges support from the Belgian F.R.S.-FNRS, of which he is Research Director.

REFERENCES

[1] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.

[2] A. L. Christensen, R. O'Grady, and M. Dorigo. Parallel task execution, morphology control and scalability in a swarm of self-assembling robots. In *Proceedings of the 9th Conference on Autonomous Robots and Competitions*. IEEE Press, Piscataway, NJ, 2009. In press.

[3] M. B. Dias, M. B. Zinck, R. M. Zlot, and A. Stentz. Robust multirobot coordination in dynamic environments. In *Proceedings of IEEE Conference on Robotics and Automation, ICRA'04*, volume 4, pages 3435 – 3442. IEEE Press, Piscataway, NJ, 2004.

[4] N.R. Franks and T. Richardson. Teaching in tandem-running ants. *Nature*, 439:153, 2006.

[5] A. Gage and R. R. Murphy. Affective recruitment of distributed heterogeneous agents. In *Proceedings of AAAI'04,*, pages 14–19. AAAI Press, Menlo Park, CA; MIT Press, Cambridge, MA, 2004.

[6] S. Garnier, C. Jost, R. Jeanson, J. Gautrais, M. Asadpour, G. Caprari, and G. Theraulaz. Aggregation behaviour as a source of collective decision in a group of cockroach-like robots. *Advances in Artificial Life*, 3630:169–178, 2005.

[7] B. P. Gerkey and M. J. Matarić. Pusher-watcher: An approach to fault-tolerant tightly-coupled robot coordination. In *Proceedings of IEEE International Conference on Robotics and Automation, ICRA'02*, pages 464 – 469. IEEE Press, Piscataway, NJ, 2002.

[8] B. P. Gerkey and M. J. Matarić. Sold!: auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation*, 18(5):758–768, 2002.

[9] R. Groß, E. Tuci, M. Dorigo, M. Bonani, and F. Mondada. Object transport by modular robots that self-assemble. In *Proceedings of the 2006 IEEE International Conferance on Robotics and Automation*, pages 2558–2564. IEEE Computer Society Press, Los Alamitos, CA, 2006.

[10] R. Jeanson, C. Rivault, Jean-Luis Denebourg, S. Blanco, R. Fournier, C. Jost, and Guy Theraulaz. Self-organized aggregation in cockroaches. *Animal Behavior*, 69:169–180, 2004.

[11] M. J. B. Krieger and J.-B. Billeter. The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, 30(1-2):65–84, 2000.

[12] A. Ledoux. Étude experimentale du grégarisme et de l'interattraction sociale chez les Blattidés. *Annales des Sciences Naturelles Zoologie et Biologie Animale*, 7:76–103, 1945.

[13] Dorigo M., Trianni V., Şahin E., Gross R., Labella T. H., Baldassarre G., Nolfi S., Deneubourg J.-L., Mondada F., Floreano D., and Gambardella L. M. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2-3):223–245, 2004.

[14] A. Martinoli, A.-J. Ijspeert, and F. Mondada. Understanding collective aggregation mechanisms: from probabilistic modelling to experiments with real robots. *Robotics and Autonomous Systems*, 29:51–63, 1999.

[15] C. Melhuish, O. Holland, and S. Hoddell. Convoying: Using chorusing to form travelling groups of minimal agents. *Robotics and Autonomous Systems*, 28:207–216, 1999.

[16] M. Moglich, U. Maschwitz, and B. Holldobler. Tandem calling: a new kind of signal in ant communication. *Science*, 186(4168):1046–1047, 1974.

[17] R. O'Grady, R. Groß, F. Mondada, M. Bonani, and M. Dorigo. Self-assembly on demand in a group of physical autonomous mobile robots navigating rough terrain. In *Proceedings of ECAL 2005*, volume 3630 of *LNAI*, pages 272–281. Springer Verlag, Berlin, Germany, 2005.

[18] R. O'Grady, C. Pinciroli, A. L. Christensen, and M. Dorigo. Supervised group size regulation in a heterogeneous robotic swarm. In *Proceedings of ROBOTICA 2009 - 9th International Conference on Autonomous Robot Systems and Competitions*, pages 113–119. IPCB, Castelo Branco, Portugal, 2009.

[19] L. E. Parker. ALLIANCE: an architecture for fault tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.

[20] C. Pinciroli. Object retrieval by a swarm of ground based robots driven by aerial robots. Mémoire de DEA, Universite Libre de Bruxelles, Bruxelles, Belgium, 2007.

[21] J. Pugh and A. Martinoli. Relative localization and communication module for small-scale multi-robot systems. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA2006)*, pages 188–193. IEEE Press, Piscataway, NJ, 2006.

[22] J. Roberts, T. Stirling, J. Zufferey, and D. Floreano. Quadrotor using minimal sensing for autonomous indoor flight. In *Proceedings of the European Micro Air Vehicle Conference and Flight Competition (EMAV2007)*, 2007.

[23] M.K. Rust, J.M. Owens, and D.A. Reierson. *Understanding and controlling the german cockroach*. Oxford University Press, UK, 1995.

[24] G.S. Sukhatme, J.F. Montgomery, and R.T. Vaughan. Experiments with aerial-ground robots. In T. Balch and L.E. Parker, editors, *Robot Teams: From Diversity to Polymorphism*, pages 345–367. AK Peters, Wellesley, MA, 2001.