

Evolved swarming without positioning information: an application in aerial communication relay

Sabine Hauert · Jean-Christophe Zufferey · Dario Floreano

Received: date / Accepted: date

Abstract In most swarm systems, agents are either aware of the position of their direct neighbors or they possess a substrate on which they can deposit information (stigmergy). However, such resources are not always obtainable in real-world applications because of hardware and environmental constraints. In this paper we study in 2D simulation the design of a swarm system which does not make use of positioning information or stigmergy.

This endeavor is motivated by an application whereby a large number of Swarming Micro Air Vehicles (SMAVs), of fixed-wing configuration, must organize autonomously to establish a wireless communication network (SMAVNET) between users located on ground. Rather than relative or absolute positioning, agents must rely only on their own heading measurements and local communication with neighbors.

Designing local interactions responsible for the emergence of the SMAVNET deployment and maintenance is a challenging task. For this reason, artificial evolution is used to automatically develop neuronal controllers for the swarm of homogenous agents. This approach has the advantage of yielding original and efficient swarming strategies. A detailed behavioral analysis is then performed on the fittest swarm to gain insight as to the behavior of the individual agents.

Keywords swarm intelligence · swarming without positioning · Micro Air Vehicles (MAVs) · communication relay · artificial evolution · situated communication · SMAVNET

Laboratory of Intelligent Systems
Ecole Polytechnique Fédérale de Lausanne (EPFL)
Tel.: +41 21 693 77 59
Fax: +41 21 693 58 59
E-mail: sabine.hauert@epfl.ch

1 Introduction

Existing aerial robotic swarms use relative or global positioning information to navigate in their environment using either map-based strategies (Kuiper and Nadjm-Tehrani, 2006; Parunak et al, 2005; Sauter et al, 2005; Elston and Frew, 2008; Flint et al, 2002; Lawrence et al, 2004; Pack and York, 2005; Yang et al, 2005), Reynolds' Flocking (Reynolds, 1987) or Artificial Physics (Spears et al, 2005) approaches (Basu et al, 2004; De Nardi and Holland, 2007; Holland et al, 2005; Kadrovach and Lamont, 2001; Merino et al, 2006), or predefined swarm formations (Vincent and Rubin, 2004). Other researchers have explored the use of artificial evolution to automatically determine position-aware swarm controllers (Gaudio et al, 2005; Lin et al, 2004; Richards et al, 2005; Soto and Lin, 2005; Wu et al, 1999). While these approaches have the potential to achieve tasks such as surveillance, sensing, area coverage, target tracking and communication relay, they all rely heavily on local or global positioning information.

Swarming without positioning is interesting because it alleviates the need for information which is often unavailable in real-life applications because of hardware and environmental constraints. Off-the-shelf sensors such as cameras, laser range finders, radars, ultrasound and infrared sensors are capable of providing relative positioning. However, interesting usability ranges for aerial swarm deployments (>100 m) typically entail expensive hardware, in terms of energy, cost and weight, which is incompatible with the scalable nature of swarms composed of large numbers of simple and inexpensive aerial robots. Alternatively, wireless technologies can be used to estimate the range or bearing between agents of the swarm and position beacons using time of arrival (TOA), time difference of arrival

(TDOA), angle of arrival (AOA) or the received signal strength indicator (RSSI). However, depositing beacons in the environment is generally not practical for the rapid deployment of swarm systems in unknown environments (Hu and Evans, 2004). Finally, global positioning system (GPS) technologies are unreliable or impossible in cluttered areas where line-of-sight with the necessary satellites is not available (Siegwart and Nourbakhsh, 2004), thus limiting the deployability of the system in certain environments. Taking these limitations into account, we aim at a system which does not depend on such positioning technologies.

Very few examples of swarms that do not use stigmergy or positioning are known. For example, in work by Nembrini et al (2002), a swarm of ground robots is capable of generating a coherent movement towards a light beacon while avoiding obstacles and maintaining global shape in simulation. Very simple rules based on local wireless communication are responsible for the behavior of the swarm, and no positioning information is used. However, the resulting rules are adapted to wheeled robots with very simple dynamics, whereas fixed-wing MAVs need to maintain forward motion in order not to stall. Finally, the developed local interactions are not suitable for the generation of aerial communication relay and no methodology for designing the local interactions is given.

Our target application scenario consists in implementing a rapidly-deployable aerial communication network between users located on ground. Such emergency systems could replace damaged, inexistent or congested networks and can play an important role in disaster mitigation (Oh, 2003). The long-term objectives include relaying multiple users as can be seen in Fig. 1. The aerial nature of the system is interesting in that it allows for line-of-sight transmissions between MAVs, which is more energy efficient than communication through obstacles at ground level. Furthermore, MAVs can fly over difficult terrain such as flooded areas or debris.

It was chosen to approach the problem in 2D simulation using a homogenous swarm of twenty agents displaying the paradigms of swarm robotics (Şahin, 2005). Rather than positioning sensors, agents are only equipped with local communication capabilities (typically a WiFi module) and a heading sensor. As a first step, a simplified scenario is considered in which a base station (BS) and a user station (US) located on the ground must communicate via the SMAVNET.

Determining the atomic rules responsible for the deployment and maintenance of the SMAVNET is not straightforward and no methodology currently exists to acquire such knowledge. The task is especially difficult given the lack of positioning information in the system,




Fig. 1 Artistic view of the use of a swarm of MAVs for establishing communication networks between users stations located on ground.

which prevents from applying existing algorithms described previously. To overcome these difficulties, we use genetic algorithms to evolve neural controllers for the swarm of agents (Nolfi and Floreano, 2000). Evolutionary approaches have the potential to find creative swarm strategies which might otherwise not have been considered.

A good understanding of the evolved inter-agent interactions can provide the basis for hand-designing controllers in the future. The importance of being able to understand the behavior of the agents is increased in real-life applications because it allows for rapid modification of the swarm behavior (i.e. no need to re-evolve a controller). Also, basic understanding of the underlying rules is necessary if we are to perform a safeness and liveness analysis of the system following swarm engineering paradigms (Winfield et al, 2005a,b). For these reasons, a qualitative behavioral analysis is conducted on the best evolved controller.

The following section describes the experimental setup in terms of scenario, simulation setup and evolutionary methodology. Section 3 presents the performance of the best evolved agent controller and the local interactions responsible for the deployment and maintenance of the SMAVNET. The robustness of the system to variations in the setup parameters and MAV failures is assessed. Finally the conclusion section presents the current status and future steps within the scope of the SMAVNET project.

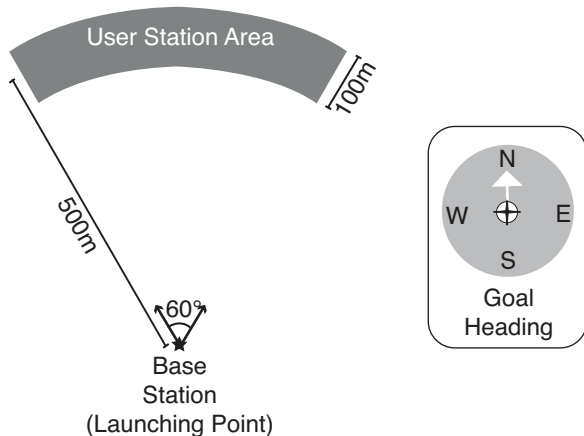


Fig. 2 Scenario environment as defined in our application. The swarm should be capable of finding any user station placed within a $\pm 30^\circ$ angle of the desired goal heading and within 500 ± 50 m of the base station.

2 Experiment Setup

2.1 Scenario

Our preliminary test case scenario consists in having a swarm of MAVs search for a single user station positioned on ground while maintaining a connection to the base station from which they are launched. MAVs must remain coherent (i.e. interconnected) and directly or indirectly connected to the base station so as not to get lost. A communication link between the base station and the user station is then to be maintained for the remaining of the 30 min trial duration (i.e. the current battery life-time expectation of the physical platform under development (Leven et al, 2007)). Experiments are done in 2D, meaning all MAVs are assumed to fly at the same altitude. At this stage, collision avoidance has not been taken into account and will later be implemented by altitude differentiation. MAVs are launched every 15 ± 7.5 seconds within a 50 m radius from the base station to model the fact that MAVs will be launched by hand by a single operator. The user station is positioned in the area 500 ± 50 m away from the base station and within a $\pm 30^\circ$ angle of a given goal orientation reflecting the user’s approximate knowledge of the direction of the user station (Fig. 2).

2.2 MAV model

The MAV dynamic model is implemented in 2D based on a first order flight model which was chosen to reproduce trajectories obtained with our current simple and low-cost fixed-wing airframe (Leven et al, 2007).

The resulting agents fly at a speed of 14 m/s affected by $\pm 5\%$ uniform noise and are unable to hover or make turns sharper than 18 m in radius. Uniform noise in the range $[-5, 5]^\circ/s$ is added to the turn rate of the MAV. A smoothing function ensures that the turn rate can not be modified abruptly (the maximum change in turn rate is of $100^\circ/s^2$). Such physical constraints enforce a more complex controller with respect to ground robots or hovering platforms. At this stage, the only internal sensor used for swarming is a heading measurement sensor affected by Gaussian noise with a standard deviation of 5° .

MAVs are capable of wireless communication with other MAVs and the base station or the user station. The communication model assumes that communication between two agents is perfect if the individuals are less than 90 m apart, noisy from 90 m to 100 m and inexistent when separated by 100 m or more. The probability of entirely dropping a message increases linearly between 90 m and 100 m from 0 to 1. Similar disc models have been used in robotic swarm simulators (Nembrini et al, 2002) based on assumptions introduced in work by Winfield (2000).

Agents can send two types of messages, control messages and data messages. Control messages are only used for the coordination of the swarm and are broadcasted by each agent every 50 ms. Data messages are related to the application of the swarm (e.g. video relay, voice relay, etc.), and are sent between the base station to the user station every 50 ms (the direction of the message flow is non-relevant). A data message is assumed to have reached its destination if there exists, at that given time-step, at least one communication pathway between the base station and the user station.

2.3 Neural Controller

MAVs are controlled by means of a neural network (Fig. 3) which outputs the turn rate to adopt. The speed of the MAV is assumed to remain constant. Since no positioning information is available, the inputs are exclusively derived from an absolute heading sensor and the control messages received from neighboring MAVs (situated communication (Støy, 2001)). Inputs to the neural controller are defined as follows:

- *Heading* Heading of the MAV as computed using a magnetic compass.
- *BSHopCount* Minimum number of network hops that separate the base station from the MAV. MAVs which are disconnected from the base station are assigned a *BSHopCount* value of N where, N is the maximum *BSHopCount* which can be obtained in

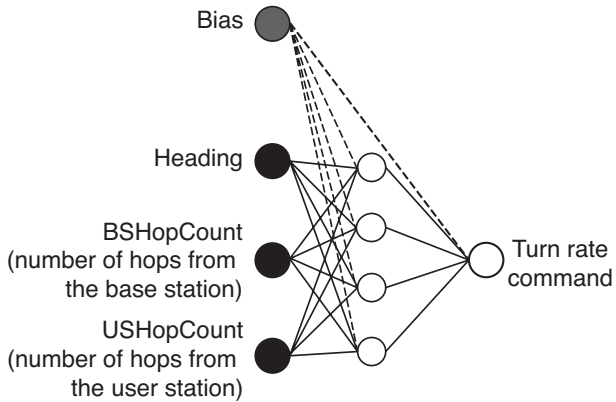


Fig. 3 Artificial neural network architecture used for the control of the MAVs.

a given network. For our application, N corresponds to the number of MAVs in the swarm ($N=20$). Hop information is transmitted in a decentralized and scalable manner at a rate of 1 hop/50 ms as described in Appendix A.

- *USHopCount* Minimum number of network hops that separate the user station from the MAV. When disconnected from the user station MAVs are assigned a value of N .

Inputs were scaled to fit the range $[-1,+1]$ and the neuronal transfer function was chosen to be a hyperbolic tangent (\tanh). The synaptic weights were in the range $[-4,+4]$ and coded on 8 bits. Neural architectures possessed 4 hidden neurons.¹

2.4 Genetic Algorithm

For our application, 15 independent evolutions were conducted using a genetic algorithm (Nolfi and Floreano, 2000). Each MAV in the swarm is equipped with identical neural controllers (homogenous swarm) to allow for scalable systems composed of interchangeable agents, an advantage in real-life applications. Colony level selection was used to favor inter-agent cooperation (Waibel et al, submitted). Each neural controller is genetically represented by a binary string composed of a series of 8-bit blocks representing the synaptic weights of the neural controller. Populations of 100 individuals were evolved using a rank-based truncation selec-

¹ We sequentially tested neural architectures with zero to four hidden neurons. Architectures with four hidden neurons were found to yield swarm controllers with the highest fitness.

tion, one-point crossover, bit mutation and elitism. The genomes of the first generation were initialized randomly. Each individual represents an MAV controller which is copied to all members of the swarm. The mean fitness of each swarm (Sec. 2.5) was evaluated for 10 user stations positioned within the user station area (Fig. 2)².

After ranking the individuals according to their measured fitness values, the top 10 were copied to the new population (elitism). The remaining population was generated from the crossover of two randomly paired individuals within the 30 first ranks. One point crossover was applied to each pair with a probability 0.2 and each individual was then mutated with a probability 0.01 per bit.

2.5 Fitness

We aim at favoring systems which can rapidly establish connections in a coherent manner and maintain them throughout the entire trial duration. To do so, the fitness F of each swarm is measured by the mean connectivity between the base station and the user station over the entire trial duration ($T=30$ min):

$$F = \frac{1}{T} \sum_{t=0}^T c(t) \quad (1)$$

The connectivity $c(t)$ of the network corresponds to the number of disjoint paths between the base station and the user station, which also corresponds to the minimum number of MAVs that need to be removed from the network for the communication link to break. In order to put an additional pressure against the loss of MAVs, trials during which individuals were disconnected from the base station for more than 30 s were assigned a fitness of 0.

3 Results

The evolutionary run during which the controller with the highest fitness was found can be seen in Fig. 4. Of the remaining 14 evolutionary runs, 2 were capable of finding controllers displaying similar strategies as the best individual while the remaining 12 evolutionary runs converged to a local optimum. This shows that the evolutionary process is not straightforward. In

² Five user stations were positioned 550 m away from the base station at 30° , 15° , 0° , -15° , -30° angles from the predefined goal heading (Fig. 2), the same angles were applied to the five user stations placed at a distance of 450 m.

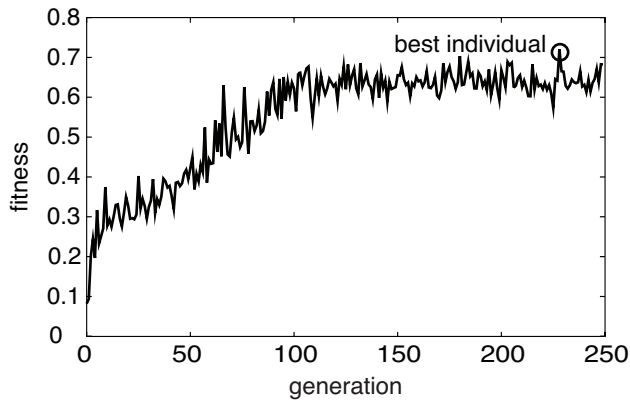


Fig. 4 Maximum individual fitnesses for each generation of the best evolutionary run. The individual with the highest fitness selected for the behavioral analysis was found at generation 228.

this paper, we focus on the performance and behavior of the best evolved controller in order to gain insight as to how swarming without positioning can be achieved and what underlying mechanisms are used. We then assess the robustness of the evolved approach to variations on the parameters of the experimental setup and MAV failures.

3.1 Performance

After evolution, the best controller was tested 1000 times on randomly positioned user stations within the area described in Fig. 2. In the end, 97.5% of the user stations were found and only 2.24% of the 20'000 deployed MAVs (1000 runs with 20 MAVS) were lost. The connectivity measures of the networks over the 1000 trials are shown in Fig. 5 (left). The connectivity measures equal 0 during the first couple of minutes of a deployment because few MAVs have been launched, and the swarm has not traveled far enough to find a user station. However, once the connection between the base station and the user station is established, it is maintained in a robust manner, this can be seen by the fact that the connectivity measures remain stable to the end of the trials. After 30 min, 22.4% of the 1000 trials had a connectivity of zero, 74% had a connectivity of one and 3.6% had connectivity of two. Over the 1000 trials, only those where the user station was not found (2.5%) displayed a constant connectivity of zero. The remaining trials maintained at worst intermittent connections, displaying varying connectivities between zero, one and two. A perfect (uninterrupted) connection between the base station and the user station is not required as long as the swarm remains coherent. Fig. 5 (right) shows statistics on the mean connectivity over 30 min trials

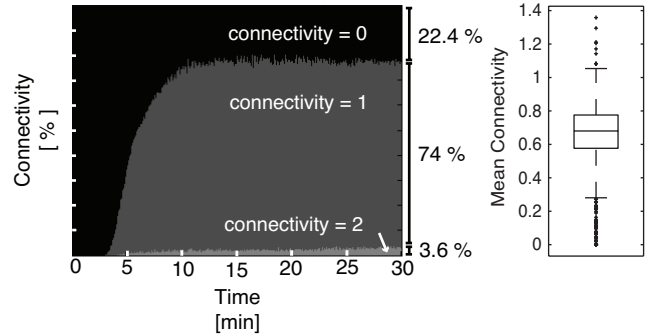


Fig. 5 Left: Connectivity of the best evolved controller when tested one thousand times on randomly positioned user stations within the user station area (Fig. 2). At each time-step, the proportion of networks (out of the 1000 trials) having connectivities of 0, 1 and 2 is shown. Connectivity values above 2 were not encountered. Right: Statistics on the mean connectivity over 1000 trials of 30 min. Each box has lines at the lower quartile, median, and upper quartile values. The whiskers extend to the farthest data points that are within 1.5 times the interquartile range. Outliers are shown with a + sign.

as described in Eq. 1. The median fitness over 1000 trials is of 0.68 which is coherent with the performance of the best individual found through evolution, as shown in Fig. 4.

3.2 Behavioral Analysis

Through a behavioral analysis we aim at extracting the underlying principles responsible for the emergence of the evolved swarm strategy. These principles could prove useful for hand designing swarm strategies for MAV swarms in the future, across additional scenarios than the one presented here. Furthermore, a good understanding of the behavior of a swarm system is essential if it is to be deployed in real-life applications. The methodology consists in looking at the global strategy displayed by the swarm, followed by an analysis on the effect of each input to the MAV neural controllers on the individual behaviors of the MAVs.

In the evolved strategy, MAVs form a tight chain which coherently moves from one side of the search area to the other as can be seen in Fig. 6. As soon as the user station is found, the MAVs reorganize to maintain the communication link active. A stable communication link is maintained by all MAVs performing the smallest possible circular trajectory given their dynamics. Notice that the area up to 450 m away from the base station is also covered although the controllers are not explicitly evolved for finding user stations in this zone.³

³ Videos showing typical swarming behaviors can be found at <http://lis.epfl.ch/smavs>

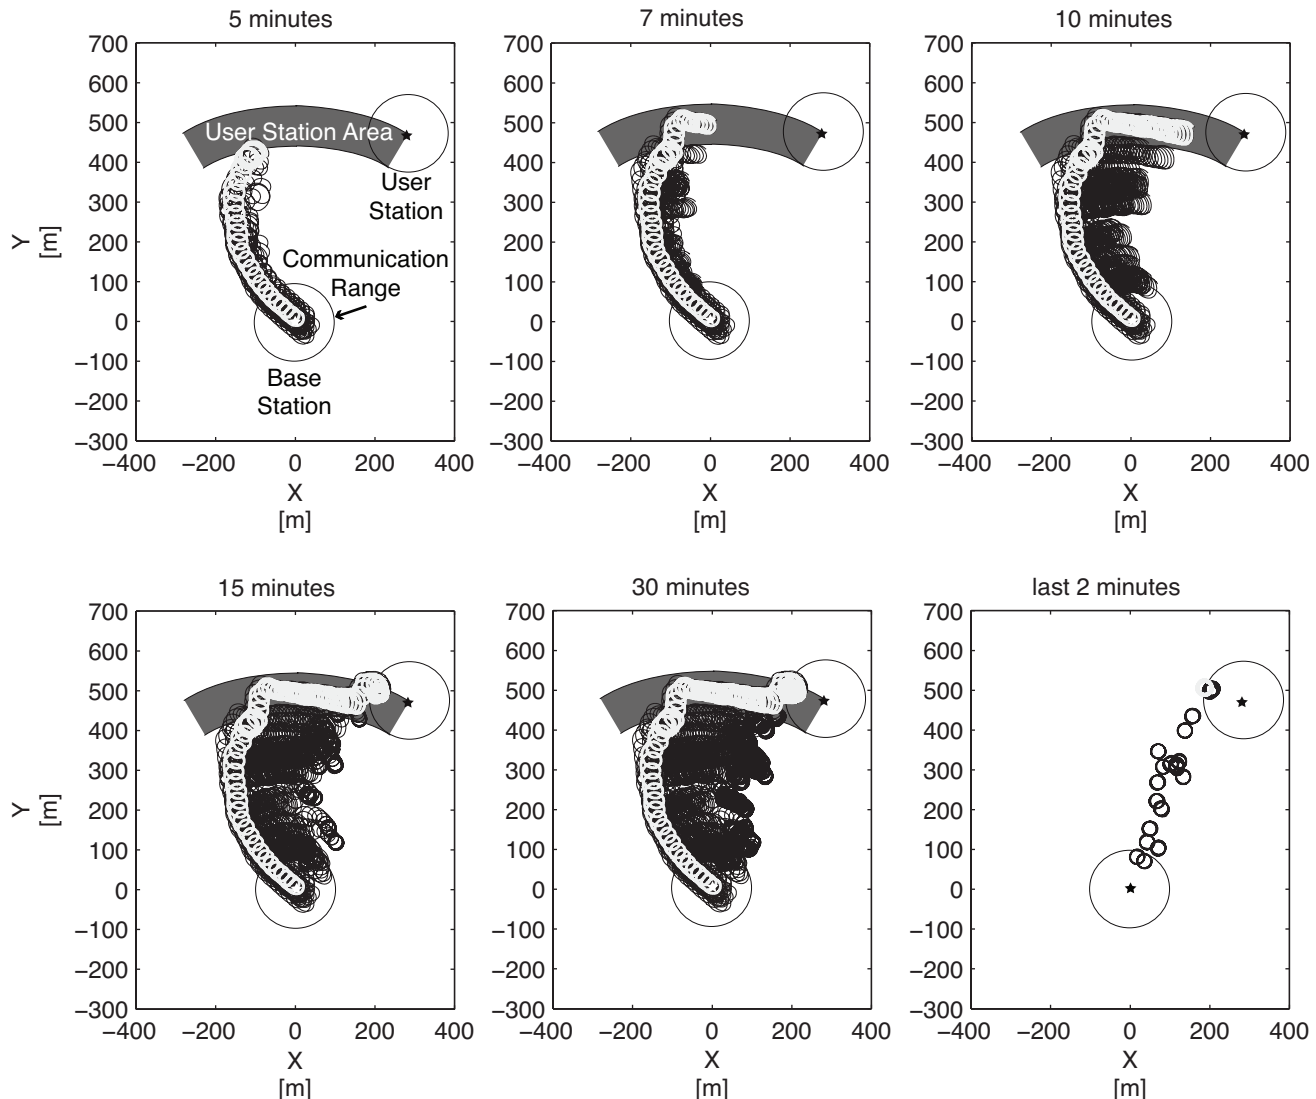


Fig. 6 Trajectories of all the MAVs during a 30 min trial of the best individual when the user station is located in the North-East corner of the search area. The trajectory of the first launched MAV is shown by a light grey line. The last two minutes of the trial show the behavior of the MAVs once the communication link between the base station and the user station is stabilized.

MAVs do not move straight, rather they adopt near constant turn rates, yielding circular trajectories. By periodically modulating their turning rate, they can describe circular trajectories which can be characterized by a velocity vector of global direction and speed. The final trajectory resembles a prolate cycloid. The periodic modulation in turn rate for a fixed $BSHopCount$ and $USHopCount$ is controlled by the MAV's heading input (Fig. 7).

The global orientation and progression speed of a given trajectory depends on how far an MAV is from the base station (Fig. 8). A low $BSHopCount$ results in trajectories directed to the West of the search area (i.e. 30° to 0° angle from the goal heading) and advancing

at low speeds. The slow overall speed of the MAVs allows them to “wait” for new launches. MAVs which are further away from the base station (more than 5 hops) adopt trajectories directed at negative angles from the goal heading, this is why the swarm chain is slightly curved (Fig. 6). When disconnected from the base station ($BSHopCount = N$), MAVs backtrack in order to maintain a coherent swarm and remain connected to the base-station.

After all MAVs have been launched, and the entire swarm disconnects from the base station, all MAVs perform the backtracking trajectory to reconnect to it. By alternating these connection and disconnection phases, the swarm is able to synchronize by having all MAVs




Fig. 8 Effect of the $BSHopCount$ on the velocity vector of the directed circular trajectories described by the best evolved MAV controllers when plotted over 30 s. The $USHopCount$ is set to N.

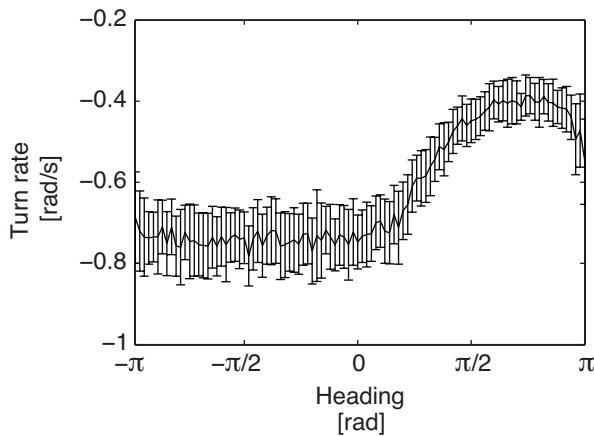


Fig. 7 Effect of the $Heading$ input on the turn rate of an MAV when the $BSHopCount$ and $USHopCount$ are maintained constant to 5 and 20 respectively.

display the same heading at a given time (Fig. 9). The connection and disconnection of the entire swarm generates a periodic signal on which MAVs gradually align. Examples of such spontaneous synchronization can be found in nature in the firefly flashing patterns for example (Camazine et al, 2001). The alternation of connection and disconnection periods also allows the MAVs to move in a coherent manner to the East of the search area. This sweeping behavior is due to the difference in MAV turn rate when connected or disconnected from the base station (Fig. 10).

Finally, MAVs that are connected to both a base station and a user station ($BSHopCount < N$ and $USHopCount < N$) perform the smallest possible circular trajectories. This ensures the maintenance of the communication link between the base station and the user station (Fig. 11).

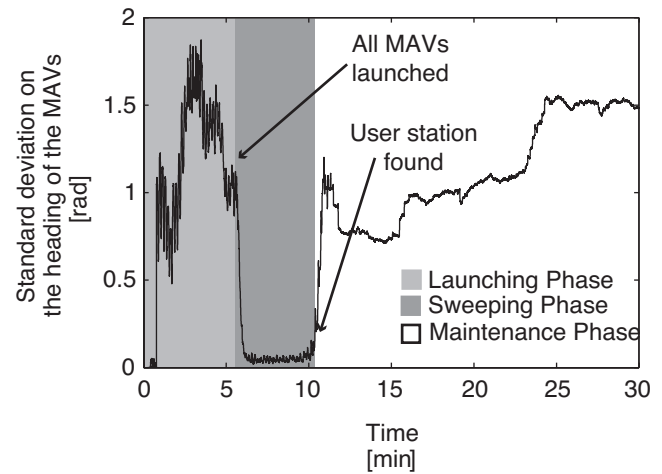


Fig. 9 Standard deviation on the heading of all MAVs during a single run with the user station placed in the North-East corner of the user station area. Three phases displaying different levels of synchrony can be observed. The first phase corresponds to the deployment in which the MAVs are being launched. During this phase, the MAVs are not synchronized. The second phase represents the sweeping phase which is highly synchronized. The third phase corresponds to the maintenance of the communication network. During this phase, the MAVs are not synchronized.

3.3 Robustness

The fact that the swarm can support significant parameter variations with little or graceful degradation, further supports the usability of the evolved strategy in real-life scenarios. For this reason, we test the best evolved controller against variations on the swarm sizes, communication ranges, MAV launch intervals and MAV speeds, as well as MAV failures. Fig. 12 shows the effect of each parameter variation on the mean connectivity over 30 min trials. Statistics are produced over 1000 tri-

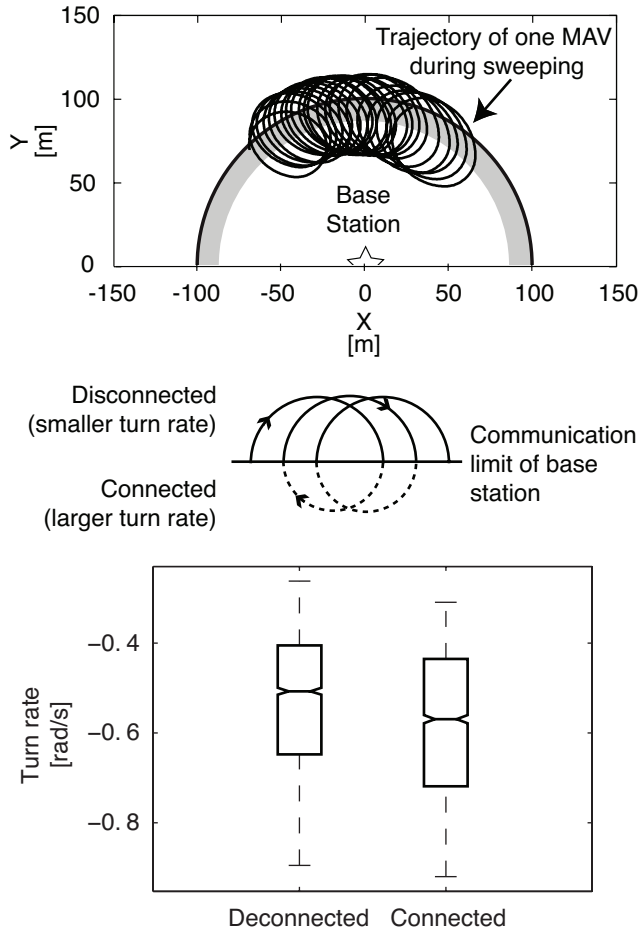


Fig. 10 The top figure shows the trajectory of an MAV along the communication limit of the base station (the grey area shows the noisy communication zone). This sweeping behavior, at the frontier of the communication range, is induced by the alternation of turn rates when connected and disconnected from the base station. The figure in the center shows an abstraction of this mechanism. A one-way ANOVA plot in the bottom figure shows that the turn rates when connected and disconnected are significantly different ($p < 0.01$). The boxes have lines at the lower quartile, median, and upper quartile values. The whiskers are lines extending from each end of the box to show the extent of the rest of the data. The notches represent a robust estimate of the uncertainty about the medians for box-to-box comparison. The notches permit the assessment of the significance of the differences of the medians. When the notches of two boxes do not overlap, the corresponding medians are significantly different at (approximately) the 95% confidence level (McGill et al, 1978).

als with user stations randomly positioned within the user station area described in Fig. 2. For each set of experiments, only one parameter was changed, the remaining parameters being set to the original values described in the experimental setup (Sec. 2).

Results show a decrease of performance by at most 10% with respect to the original median performance of 0.68 (Fig. 5, right) for any of the following variations in setup parameters:

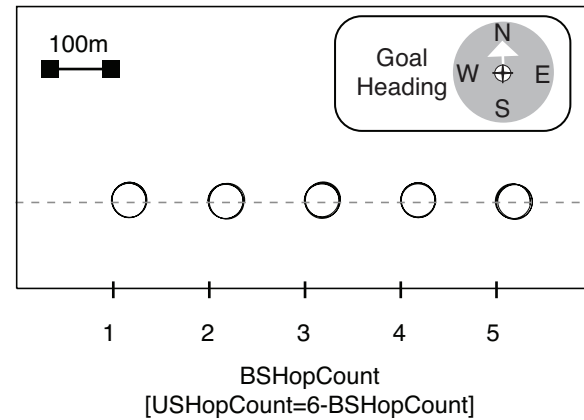


Fig. 11 Effect of the $BSHopCount$ and $USHopCount$ on the trajectory of the best evolved MAV controllers when plotted over 30s. Here the MAV is assumed to be connected to the base station which is $BSHopCount$ away and to the user station which is $USHopCount = 6 - BSHopCount$ away, assuming a minimum path between the base and user station of length 6.

- variations in swarm size from 17 to 30 (30 being the maximum tested)
- variations in communication range from 100 m to 200 m (200 m being the maximum tested)
- variations in launch interval from 10 s to 20 s
- variations in MAV speed from 10 m/s to 16 m/s

Performances decreased by more than 10% are typically due to the fact that MAVs are not spaced out enough within the MAV chain (small launch intervals, low MAV speeds, short communication ranges). In such deployments, MAVs are launched before the MAV chain is able to sufficiently distance itself. These swarms are not able to form MAV chains which reach out to far away user stations. Launching small MAV swarms has a similar effect (i.e. the formation of short chains).

The graceful degradation of the performance for launch intervals larger than 20 s is due to the fact that the swarm performs the sweeping behavior to the right prior to the formation of the entire MAV chain, resulting in certain user stations in the North-West corner of the user station area (Fig. 2) to not be found. The sweeping behavior is activated by continuous connections and disconnections of the entire swarm from the base station (Sec. 3.2). In the case of long launch intervals, the swarm enters this behavior in between each new launch, for a short period of time, resulting in a slight shift of the MAV chain. Large MAV speeds are also challenging for the swarm for identical reasons.

In addition, randomly removing 1 MAV at random times during the trials (MAV failure) does not significantly decrease the performance (Fig. 13). Further-

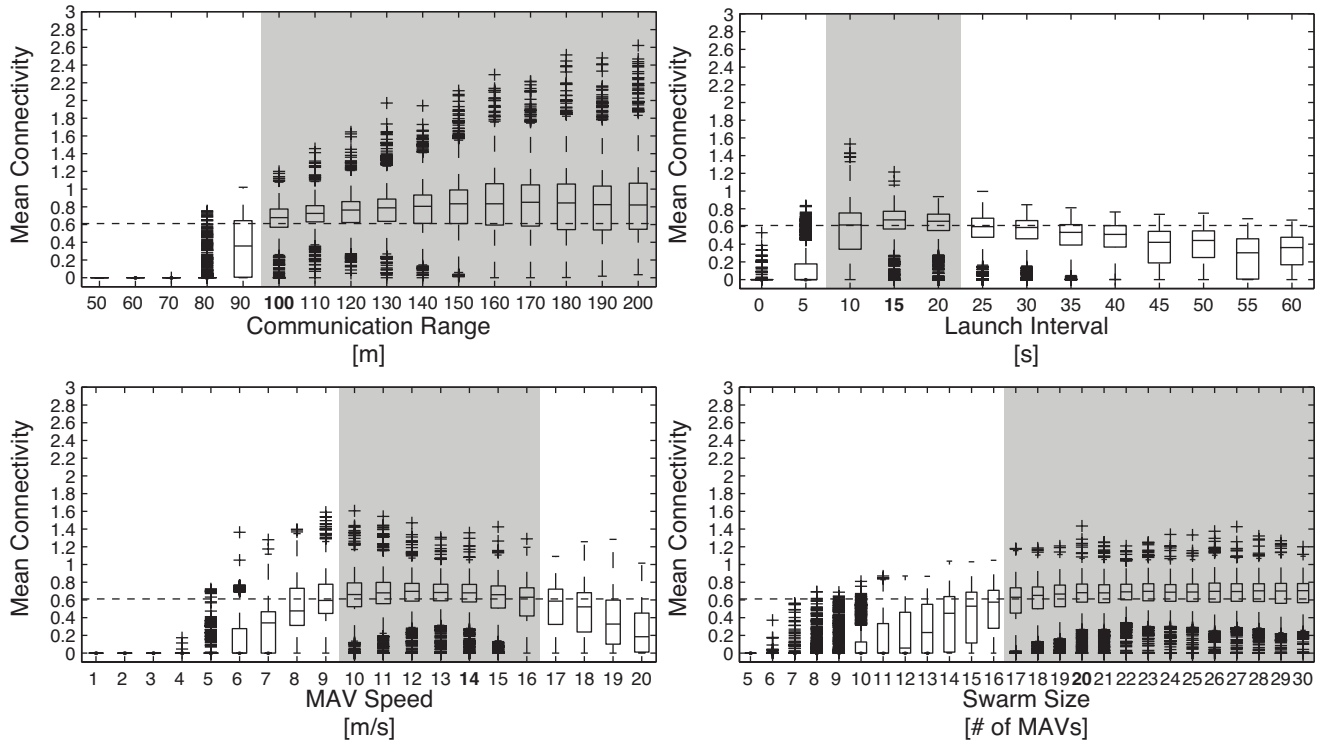


Fig. 12 Effect of varying a parameter of the experimental setup on the mean connectivity over 1000 trials of 30 min. The dashed line represents a 10% decrease in performance with respect to the original parameters shown in bold. Sectors in grey include parameter values which perform above this limit. Each box has lines at the lower quartile, median, and upper quartile values. The whiskers extend to the farthest data points that are within 1.5 times the interquartile range. Outliers are shown with a + sign.

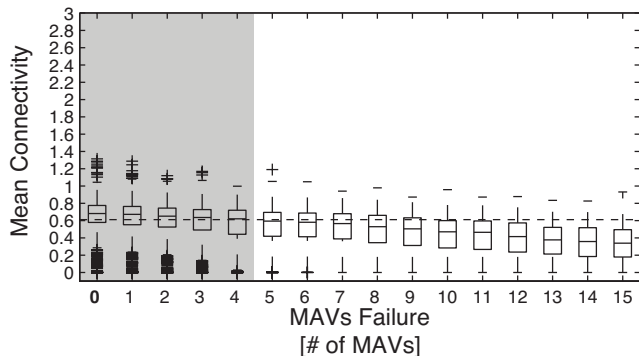


Fig. 13 Effect of MAV failures on the mean connectivity over 1000 trials of 30 min. At each trial, a fixed number of randomly chosen MAVs are removed at random times for the remaining of the 30 min trial. The dashed line represents a 10% decrease in performance with respect to the original parameter shown in bold. Each box has lines at the lower quartile, median, and upper quartile values. The whiskers extend to the farthest data points that are within 1.5 times the interquartile range. Outliers are shown with a + sign.

more, the failure of 4 MAVs or less decreases the median connectivity mean over 1000 trials by less than 10% .

Finally, while the controller was evolved in the scope of a specific experimental setup, it is robust to limited variations on the swarm sizes, communication ranges,

MAV launch intervals and MAV speeds, as well as to MAV failures.

4 Conclusion

To this day, application oriented swarm algorithms have mostly assumed that relative or global positioning concerning themselves and their neighbors was available. Instead, our agents rely on internal sensors and information derived from local wireless communication with neighboring agents. Our approach has been demonstrated in 2D simulation on an application whereby 20 MAVs are deployed to establish and maintain a multi-hop communication network between a base station and a user station located on the ground. Such a rapidly deployable communication network for information relay could play an important role in disaster mitigation.

Artificial evolution is used as an exploration tool to automatically determine novel swarm controllers. The evolved swarm is able to find more than 97% of the user stations placed within 550 m from the base station and between -30° and 30° of the given goal heading in a coherent manner. The system is then able to maintain the communication link active throughout the remaining trial durations.

Through a behavioral analysis of the best evolved controller, we were able to identify the following principles for the design of swarms of MAVs for communication relay, which do not use positioning information:

- directed circular trajectories prove to be an efficient approach for designing fixed wing aerial platform navigation by allowing variation of direction of movement and global advancement speed by only varying the turn rate with respect to the absolute heading of the MAV.
- MAVs which are not connected (even indirectly) to the base station should immediately backtrack using their heading sensor in order to attempt a reconnection.
- Alternation of connection and disconnection phases from the base station provide the basis for a synchronized sweeping movement of the entire swarm.
- MAVs which are connected to both the base station and the user station should perform the smallest possible circular trajectory in order to maintain the communication link between the two.

Finally, we show that the evolved strategy is robust to parameter variations in the experimental setup and MAV failures.

In the future, we plan on applying the principles found through evolution to different scenarios than the one presented in this paper. Possible scenarios could feature multiple dynamic user stations and base stations. In order to increase the realism of the simulations we are currently developing a 3D simulator which implements communication models (for inter-MAV and base station to user station communication) based on 802.11b wireless specifications and physics-based wave propagation. Preliminary results, yet unpublished, tend to back the rapid implementation of behaviors inspired by evolution, such as the chain formation, synchronization and sweeping mechanisms described in Sec. 3.2, for novel applications within more realistic simulators.

Efforts will also be put forward to mitigate the effect of wind based on low-level reactions to wind sensor readings (Rodriguez et al, 2007), or high-level behaviors at the individual and swarm level.

Finally, we aim at transferring simulated controllers to a swarm of MAV platforms currently developed for real-life experimentation.

Acknowledgements The authors are grateful to Severin Leven who is currently developing the platform and has provided his insight on all related aspects. The authors are also in debt to the reviewers, whose constructive comments have allowed a significant improvement of this paper. This work is supported by armasuisse, competence sector Science + Technology for the Swiss Federal Department of Defense, Civil Protection and Sports.




Fig. 14 Topology of the swarm determined using local communication. The base station and the user station are represented by black circles, the MAVs by white ones and the local communication links by the lines connecting them. Tags above each MAV and the base station represent the number of hops from the base station and the user station respectively. MAVs that are isolated from the swarm receive the default values $N = 20$.

A Communication

Information about the topology of the swarm or the connection status of an MAV to the base station or the user station can be gained using local wireless communication.

MAVs are capable of determining the minimum number of network hops needed for a message to go from the base station (*BSHopCount*) and the user station (*USHopCount*) to themselves using only local communication (Fig. 14).

More generally, any MAV can approximate its hop count (h) to any ground station using equation 2.

$$h(s, i, t) = \begin{cases} N & t = 0 \\ N & n(i) = \emptyset \\ 1 & s \in g(i) \\ \min(\operatorname{argmin}_{n_i \in n(i)}(h(s, n_i, t-1)) + 1, N) & \text{otherwise} \end{cases}$$

s	= ground station index
i	= MAV index
t	= timestep of 50 ms duration
$n(i)$	= MAVs in the neighborhood of i
$g(i)$	= base and user stations in the neighborhood of i
N	= maximum number of hops between an MAV and the ground stations (N=20 in a scenario with 20 MAVs)

(2)

While this system tends towards the correct topology of the network when the system is stable, it can obviously be momentarily locally inexact because of the dynamism of the network and because hop information needs to propagate throughout the network at a speed of 1 hop per 50 m/s. Another possibility would be to have the base station or user station broadcast a notification message which would then flood the network and almost instantaneously update the hop count of all MAVs with respect to the initiating user or base stations on ground. We believe however that our solution has several advantages over this solution. First of all, it is scalable in the number of user stations on ground (additional ground stations increase slightly the size of messages sent between MAVs but do not increase the number of messages which need to be sent). Also, our approach is truly decentralized and dynamic in that it does not rely on the base stations or user stations to initiate a flooding mechanism but only on local communication between MAVs.

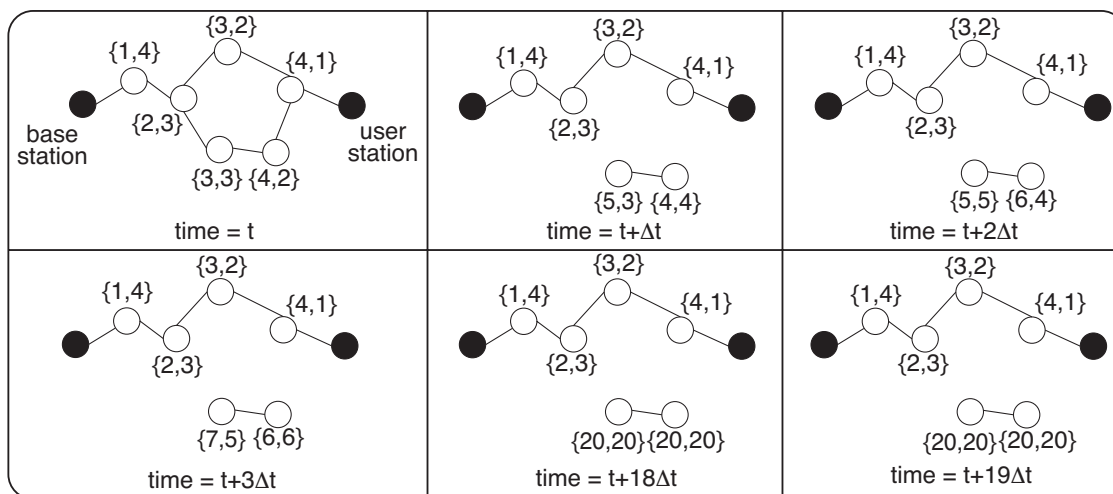


Fig. 15 Effect of the disconnection of a group of MAVs from the swarm. The base and user stations are represented by black circles, the MAVs by white circles and the local communication links by the lines connecting them. Tags above each MAV and the base station represent the number of hops from the base station and the user station respectively. Messages are sent every time-step ($\Delta t=50$ ms).

Finally, MAVs can detect if they are connected to the base station or user station either directly or indirectly thanks to the mechanism described in Fig. 15. This mechanism is derived from equation 2, which pushes the hop information of MAVs disconnected from the base or the user station to increment until the cutoff value N . MAVs that have reached the cutoff value are assumed disconnected.

References

- Basu P, Redi J, Shurbanov V (2004) Coordinated flocking of UAVs for improved connectivity of mobile ground nodes. In: Proceedings of the IEEE Military Communications Conference, IEEE Press, Piscataway, NJ, vol 3, pp 1628–1634
- Camazine S, Deneubourg JL, Franks NR, Sneyd J, Theraulaz G, Bonabeau E (2001) Self-organization in biological systems. Princeton University Press, Princeton, NJ
- De Nardi R, Holland O (2007) UltraSwarm: a further step towards a flock of miniature helicopters. In: Swarm Robotics, Lecture Notes in Computer Science, vol 4433, Springer, Berlin, pp 116–128
- Elston J, Frew EW (2008) Hierarchical distributed control for search and tracking by heterogeneous aerial robot networks. In: Proceedings of the IEEE International Conference on Robotics and Automation, IEEE Press, Piscataway, NJ: to appear
- Flint M, Polycarpou M, Fernández-Gaucherand E (2002) Cooperative control for multiple autonomous UAVs searching for targets. In: Proceedings of the 41st IEEE Conference on Decision and Control, IEEE Press, Piscataway, NJ, vol 3, pp 2823–2828
- Gaudio P, Bonabeau E, Shargel B (2005) Evolving behaviors for a swarm of unmanned air vehicles. In: Proceedings of the IEEE Swarm Intelligence Symposium, IEEE Press, Piscataway, NJ, pp 317–324
- Holland O, Woods J, De Nardi R, Clark A (2005) Beyond swarm intelligence: the UltraSwarm. In: Proceedings of the IEEE Swarm Intelligence Symposium, IEEE Press, Piscataway, NJ, pp 217–224
- Hu L, Evans D (2004) Localization for mobile sensor networks. In: Proceedings of the 10th Annual International Conference on Mobile Computing and Networking, ACM Press, New York, pp 45–57
- Kadrovach BA, Lamont GB (2001) Design and analysis of swarm-based sensor systems. In: Proceedings of the 44th IEEE Midwest Symposium on Circuits and Systems, IEEE Press, Piscataway, NJ, vol 1, pp 487–490
- Kuiper E, Nadjm-Tehrani S (2006) Mobility models for UAV group reconnaissance applications. In: Proceedings of the IEEE International Conference on Wireless and Mobile Communications, IEEE Press, Piscataway, NJ, DOI 10.1109/ICWMC.2006.63
- Lawrence D, Donahue R, Mohseni K, Han R (2004) Information energy for sensor-reactive UAV flock control. In: Proceedings of the AIAA 3rd “Unmanned Unlimited” Technical Conference, AIAA Press, Reston, VA, AIAA paper 2004-6530
- Leven S, Zufferey JC, Floreano D (2007) A simple and robust fixed-wing platform for outdoor flying robot experiments. In: Flying Insects and Robots Symposium, p 69
- Lin K, Huang K, Li G, Qui XG (2004) Control of swarming UAVs in collaborative missions. In: Proceedings of the World Automation Congress, TSI Press Series, Albuquerque, NM, USA, vol 17, pp 25–30
- McGill R, Tukey JW, Larsen WA (1978) Variations of box plots. *The American Statistician* 32(1):12–116
- Merino L, Caballero F, Martínez-de Dios JR, Ferruz J, Ollero A (2006) A cooperative perception system for multiple UAVs: application to automatic detection of forest fires. *Journal of Field Robotics* 23:165–184
- Nembrini J, Winfield A, Melhuish C (2002) Minimalist coherent swarming of wireless networked autonomous mobile robots. In: From Animals to Animats 7, Proceedings of the 7th International Conference on Simulation of Adaptive Behavior, MIT Press, Cambridge, MA, pp 273–382
- Nolfi S, Floreano D (2000) Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines. MIT Press, Cambridge, MA
- Oh ES (2003) Information and communication technology in the service of disaster mitigation and humanitarian relief. In: Proceedings of the IEEE 9th Asia-Pacific Conference on Communications, IEEE Press, Piscataway, NJ, vol 2, pp 730–734
- Pack DJ, York GWP (2005) Developing a control architecture for multiple unmanned aerial vehicles to search and localize RF time-varying mobile targets: part I. In: Proceedings of the IEEE International Conference on Robotics and Automation,

- IEEE Press, Piscataway, NJ, pp 3954–3959
- Parunak HVD, Brueckner SA, Sauter J (2005) Digital pheromones for coordination of unmanned vehicles. In: *Environments for Multi-Agent Systems*, Lecture Notes in Computer Science, vol 3374, Springer, Berlin, pp 246–263
- Reynolds CW (1987) Flocks, herds and schools: a distributed behavioral model. In: *SIGGRAPH Computer Graphics*, ACM Press, New York, vol 21, pp 25–34
- Richards MD, Whitley D, Beveridge JR (2005) Evolving cooperative strategies for UAV teams. In: *Proceedings of the Genetic And Evolutionary Computation Conference*, ACM Press, New York, vol 2, pp 1721–1728
- Rodriguez A, Andersen E, Bradley J, Taylor C (2007) Wind estimation using an optical flow sensor on a miniature air vehicle. In: *Proceedings of the AIAA Conference on Guidance, Navigation, and Control*, AIAA Press, Reston, VA, AIAA paper 2007-6614
- Şahin E (2005) Swarm robotics: from sources of inspiration to domains of application. In: *Swarm Robotics*, Springer, Berlin, Lecture Notes in Computer Science, vol 3342, pp 10–20
- Sauter JA, Matthews R, Parunak HVD, Brueckner SA (2005) Performance of digital pheromones for swarming vehicle control. In: *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, ACM Press, New York, pp 903–910
- Siegwart R, Nourbakhsh IR (2004) *Introduction to autonomous mobile robots*. Bradford Book, MIT Press, Cambridge, MA
- Soto J, Lin KC (2005) Using genetic algorithms to evolve the control rules of a swarm of UAVs. In: *Proceedings of the IEEE International Symposium on Collaborative Technologies and Systems*, IEEE Press, Piscataway, NJ, pp 359–365
- Spears WM, Spears DF, Heil R, Kerr W, Hettiarachchi S (2005) An overview of physicomimetics. In: *Simulation of Adaptive Behaviour, Workshop on Swarm Robotics*, Springer, Berlin, Lecture Notes in Computer Science, vol 3342, pp 84–97
- Støy K (2001) Using situated communication in distributed autonomous mobile robots. In: *Proceedings of the 7th Scandinavian Conference on Artificial Intelligence*, IOS Press, Amsterdam, NL, pp 44–52
- Vincent P, Rubin I (2004) A framework and analysis for cooperative search using UAV swarms. In: *Proceedings of the ACM Symposium on Applied computing*, ACM Press, New York, pp 79–86
- Winfield A (2000) Distributed sensing and data collection via broken ad hoc wireless connected networks of mobile robots. In: *Proceedings of Distributed Autonomous Systems 4*, Springer, Berlin, pp 273–282
- Winfield AFT, Harper CJ, Nembrini J (2005a) Towards dependable swarms and a new discipline of swarm engineering. In: *Swarm robotics*, Lecture Notes in Computer Science, vol 4433, Springer, Berlin, pp 126–142
- Winfield AFT, Sa J, Fernández-Gago MC, Dixon C, Fisher M (2005b) On formal specification of emergent behaviours in swarm robotic systems. *International Journal of Advanced Robotic Systems* 2(4):363–370
- Wu AS, Schultz AC, Agah A (1999) Evolving control for distributed micro air vehicles. In: *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation*, IEEE Press, Piscataway, NJ, pp 174–179
- Yang Y, Minai AA, Polycarpou MM (2005) Evidential map-building approaches for multi-UAV cooperative search. In: *Proceedings of the IEEE American Control Conference*, IEEE Press, Piscataway, NJ, pp 116–121