

Initial Experiments in Using Communication Swarms to Improve the Performance of Swarm Systems

Stephen M. Majercik

Bowdoin College, Brunswick ME 04011, USA

Abstract. Swarm intelligence can provide robust, adaptable, scalable solutions to difficult problems. The distributed nature of swarm activity is the basis of these desirable qualities, but it also prevents swarm-based techniques from having direct access to global knowledge that could facilitate the task at hand. Our experiments indicate that a swarm system can use an auxiliary swarm, called a *communication swarm*, to create and distribute an approximation of useful global knowledge, without sacrificing robustness, adaptability, and scalability. We describe a communication swarm and validate its effectiveness on a simple problem.

1 Introduction

Swarm intelligence is a natural phenomenon in which complex behavior emerges from the collective activities of a large number of simple individuals who interact with each other and sense the environment only in their immediate area. They have limited memory and a limited repertoire of simple, reactive behaviors. Yet, the swarm as a whole is highly organized and complex behavior emerges at the global level that goes beyond the capabilities of the individuals.

We are interested in artificial *task-swarms*, i.e. swarms that have been designed by people to accomplish a particular task. As a motivating example, consider a swarm of unmanned aerial vehicles engaged in reconnaissance. An independent swarm might be assigned the task of protecting them from attacks by hostile swarms, and it would be helpful for the members of this swarm to have information about the size of their swarm relative to that of the attacking swarm in order to adapt their behavior to their current size (dis)advantage. But, the distributed nature of swarm activity, although providing the advantages of adaptability, robustness, and scalability, also prevents swarm members from having direct access to *global knowledge*, i.e. knowledge about the state of the entire swarm and/or environment (including other swarms). Since the nature of a task-swarm precludes the possibility of its members having perfect global knowledge in real time, we will be concerned with the creation and distribution of *pseudo-global knowledge* (PGK), or knowledge about the state of the system that may be imperfect because it is only an approximation of the true state of the system, or because the state being approximated is a past state.

PGK needs to be supplied to a task-swarm in a way that does not sacrifice robustness and scalability, ruling out any approaches that depend on a small number of specific bots to create and deliver the information. Task-bots could attempt to construct PGK based on environmental information they collect and information they receive from other task-bots, but this may be difficult for four reasons. First, the bots' task may limit their contact with other bots and the fraction of the environment they visit, such that they are unable to aggregate enough information in a short enough time span for the information to be useful. Second, their task may interfere with the information collection process. Third, as the scale of aerial bots approaches the micro or even nano level, their capabilities might be too limited to allow them to do their task and engage in the communication that would be necessary to construct PGK. Finally, very small-scale bots might have only enough power for very short range communication, making the construction of PGK difficult. Stigmergic communication, in which individuals communicate indirectly by altering the environment, is an alternative means of communication. This type of communication has been implemented using digital pheromones in [2], but that approach relies on a network of grounded sensors. Lack of space prevents a detailed literature review.

We propose to address these issues by using a specialized, auxiliary swarm, called a *communication swarm*, whose members, called *comm-bots*, create PGK and distribute that knowledge to other swarms. A communication swarm, or *comm-swarm*, operates in a swarm-like, distributed fashion, and so preserves the adaptability, robustness, and scalability of the entire system.

In Section 2, we introduce communication swarms using a simple illustrative problem and present the results of some initial experiments that validate the idea of communication swarms. We discuss possible future work in Section 3.

2 The FLY-HOME Problem and Experimental Results

The operation of a comm-swarm will depend on the nature of the PGK it is providing; the FLY-HOME problem illustrates the operation of a particular comm-swarm. In this problem, the bots in each of two task-swarms need to determine whether their swarm has fewer members than the other swarm and, if so, to fly to a specified location. The specific action triggered is not critical; the essential idea is that the PGK allows the swarm to act appropriately. We note that this problem would be extremely difficult for task-bots making decisions based on information gathered locally, even over an extended time period, since they would find it impossible to determine, if surrounded by their fellow task-bots over a period of time, whether they were actually in the majority, or were in the minority but were part of a congregation around the home location. Comm-swarms provide a solution to this problem.

The algorithm governing task-bot motion is similar to Reynolds' boids algorithm [3]. Each bot tries to maintain a specified speed, subject to a maximum, while staying close to its neighbors (cohesion), matching their average velocity (alignment), and keeping a distance from them (separation). Neighbors are those

bots that are within a circle of a specified radius. Weights specify the strength of separation (0.0–100.0), cohesion (0.0–1.0), and alignment (0.0–1.0). Time proceeds in discrete steps and, at each step, the factors just described are used to update a bot’s velocity, which is then applied to calculate its new position.

We give an overview of our comm-swarm based algorithm for the FLY-HOME problem; lack of space prevents details. Both task-bots and comm-bots maintain cumulative counts of the number of bots they have encountered in each task-swarm. At each time step: 1) each comm-bot increases its cumulative counts by the number of task-bots from each swarm in its neighborhood, 2) each comm-bot’s cumulative counts are discounted by a factor of 0.95 to allow it to gradually “forget” past counts and adapt to changes in swarm sizes more quickly, 3) each task-bot increases its cumulative counts by the cumulative counts of each comm-bot in its neighborhood, and 4) each task-bot decides whether to fly home based on its counts. If a task-bot’s count of bots in its swarm is less than its count of bots in the other swarm, it flies home *at that time step only* by 1) changing its motion parameters such that it coheres and aligns more strongly with bots from its swarm and enforces less of a separation from them, making it possible for the task-bots flying home to congregate closely around the home area, and 2) adjusting its velocity to include a component toward the home position. If a task-bot’s count of bots in its swarm is greater than or equal to its count of bots in the other swarm, it uses its original behavioral parameters to update its velocity and position at that time step, possibly interrupting its flight home.

In our experiments, Task-Swarms 1 and 2 were initialized with 500 and 1000 bots, respectively, and the comm-swarm was initialized to the comm-swarm size being tested, all bots randomly distributed in a 2500 pixel \times 1350 pixel non-toroidal environment. (Although aerial swarms would operate in three dimensions, we used a 2-dimensional version of the problem to test the comm-swarm idea.) The home location was the center of the environment, and a swarm was defined to be at home if at least 80% of its bots were in a 400 pixel \times 400 pixel area centered on the home location, and not at home if less than 50% of its bots were in that area. All swarms were given 50 time steps to allow the behavior prescribed by their parameters to emerge; then each task-bot began to run the FLY-HOME algorithm described above. At time step 150, the size of Task-Swarm 2 was reduced to 250 by randomly removing 750 bots, making it half the size of Task-Swarm 1. Thus, at time step 50, the bots in Task-Swarm 1 should fly home, but at time step 150, the bots in Task-Swarm 1 should leave home, while the bots in Task-Swarm 2 should fly home. The success of a comm-swarm was measured by the total time steps of delay (TSD) between the trigger of each arrive-home or leave-home event and the accomplishment of the event. The TSD is defined as: $(t_{AH-1} - 50) + (t_{LH-1} - 150) + (t_{AH-2} - 150)$, where t_{AH-1} is the time step at which Task-Swarm 1 arrives home, t_{LH-1} is the time step at which Task-swarm 1 leaves home, and t_{AH-2} is the time step at which Task-Swarm 2 arrives home.

We designed three swarms for Task-Swarm 1 (a, b, and c) and three swarms for Task-Swarm 2 (a, b, and c), such that each had a qualitatively distinct behavior in terms of the size of the clusters formed by the bots and the dynamics

of those clusters. We tested a sample of possible comm-swarms on the nine possible task-swarm pairs. We tested both random comm-swarms, in which there was no explicit separation, cohesion, or alignment specified and the velocity at each step was randomly generated, and non-random comm-swarms that included separation, cohesion, and alignment factors. Our hypothesis was that the more complex dynamics of the non-random swarms (e.g. interacting clusters of bots) could yield more efficient information propagation.

For random comm-swarms, we sampled all possible combinations of three swarm sizes (200, 600, and 1000 bots), three speeds (50, 100, and 200 pixels per time step), and three neighborhood radii (25, 50, and 100 pixels). These values were chosen based on exploratory tests that indicated varied behaviors over these sets of values. These tests also indicated that performance could sometimes be improved if the comm-bots could 1) possibly separate from the task-bots, and 2) possibly cohere and align with the task-bots. Each possible combination of population size, speed, and radius was tried with each of the four possible combinations of these two factors, for a total of 108 possible comm-swarms.

We limited the number of non-random comm-swarms tested by fixing the speed and neighborhood radius values to 200 and 100, respectively, both because these were the optimal values for these two parameters in our tests of random comm-swarms, and because it seemed likely that the largest values for these two parameters would yield better performance. We sampled all possible combinations of three swarm sizes (200, 600, and 1000 bots), three separation strengths (20.0, 60.0, 100.0), three cohesion strengths (0.2, 0.6, 1.0), and three alignment strengths (0.2, 0.6, 1.0). As was the case with random comm-bots, we tested each of these parameter settings with each of four other scenarios: the comm-bots separating (or not) from the task-bots, and 2) cohering and aligning (or not) with the task-bots. This led to a total of 324 possible comm-swarms.

The emergent quality of swarm behavior produces a high variance in observed behavior. This made it difficult to designate any one swarm as the “best” for any given task-swarm pair. For each task-swarm pair, we determined the five random comm-swarms with the lowest average TSDs. There were six comm-swarms that were in this group for five task-swarm pairs; the next best comm-swarms were in that group for only two task-swarm pairs. Furthermore, all six of these better comm-swarms had the same speed (200), neighborhood radius (100), and cohered and aligned with task-bots. The number of comm-bots did not appear to make a difference (we conjecture that an increase in the number of bots results in higher bot counts, which make it more difficult for the relative counts to be reversed when the majorities are reversed), and whether the comm-bots separated from the task-bots was not important (perhaps because the cohesion and alignment with task-bots that are, themselves, separating from each other is sufficient). Further investigation is needed on both issues. We chose to test further the swarm from this group of six that had the lowest TSD in the most swarm pairs (three swarm pairs, compared to one or none for each of the other five). This was the 600 bot comm-swarm that did not separate from task-bots. See Table 1 for results.

No non-random comm-swarm displayed superior performance, even in the weak sense described above for random comm-swarms, but swarms with at least 600 bots appeared to perform somewhat better. For the sake of comparison with the random comm-swarm shown in Table 1, we chose to further investigate non-random comm-swarms of that size and with the same characteristics as that random comm-swarm, but with non-zero separation, cohesion, and alignment strengths. These tests indicated that swarms with higher separation strengths and cohesion and alignment strengths of 0.2 were more effective, leading us to do more extensive testing of the swarm that had separation, cohesion, and alignment strengths of 100.0, 0.2, and 0.2, respectively. See Table 1 for results. The non-random comm-swarm outperforms (in boldface) the random comm-swarm in eight of the nine cases, reducing the TSD by an average of 21.4%. These eight reductions are at a significance level of 0.05 or less (0.0001 in five cases). Observations of a graphical representation of the comm-swarms suggests that the high separation factor between comm-bots serves to amplify the “follow-the-leader” fluctuations induced by the cohesion and alignment between comm-bots, producing a single moving cluster of comm-bots that covers the area quickly and repeatedly, providing better coverage than the more homogeneous coverage of random movement. In both types of comm-swarms, a degree of cohesion and alignment with task-bots appears to be critical; we conjecture that this is necessary to ensure the effective distribution of information.

In the scenario described in Section 1, a comm-swarm would be targeted by the attacking swarm and gradually reduced in size. Thus, it is critical that the comm-swarm be able to function, albeit possibly with reduced effectiveness, with a smaller number of comm-bots. We tested the effectiveness of the two comm-swarms described above with only 60 bots, a 90% reduction in size (see Table 1), and found that the TSDs for these swarms were only, on average, 11.6% larger than their 600-comm-bot counterparts for random comm-swarms, and 18.8% larger than their 600-comm-bot counterparts for non-random comm-swarms, suggesting that these comm-swarms are robust to significant losses.

Table 1. TSD scores for all task-swarm pairs, mean and standard deviation of 20 runs

Swarm Pair	R-CS, 600 bots		NR-CS, 600 bots		R-CS, 60 bots		NR-CS, 60 bots	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
1a-2a	576.2	52.8	391.7	91.3	633.9	67.3	469.3	76.2
1a-2b	445.2	33.0	322.2	35.1	595.8	69.7	482.5	75.5
1a-2c	1204.0	147.1	894.9	142.3	1162.3	181.6	1011.3	153.0
1b-2a	553.8	98.4	442.6	73.3	560.8	81.0	488.5	69.8
1b-2b	384.4	34.2	274.6	38.3	493.4	59.4	372.7	58.4
1b-2c	1107.1	97.4	965.4	111.9	1127.0	217.1	1132.4	271.9
1c-2a	540.1	71.3	479.8	111.7	568.9	87.8	511.4	95.3
1c-2b	409.3	60.1	355.8	49.4	493.0	79.2	446.7	97.5
1c-2c	893.5	103.5	1040.9	144.1	951.3	112.8	950.3	170.1

R(NR)-CS = Random(Non-Random) Comm-Swarm, SD = Standard Deviation

3 Future Work

Given the preliminary nature of these experiments, there is a great deal of work to be done to refine and explore the capabilities of comm-swarms. For example, preliminary tests suggest that the performance of comm-swarms can be improved by introducing alternating information collection and information distribution phases, each with a separate set of parameters that tune the swarm to that activity. We are investigating this possibility further.

More importantly, we view our current work as the first step in a program to develop a general communication mechanism for cooperating swarms. Given the increasing miniaturization of actual bots, one might design a system of multiple specialized swarms that work together to accomplish a task. One of the challenges in designing such a system would be to provide a mechanism that facilitates information transfer among these swarms. We are currently developing a *communication-link swarm* that will allow information transfer between multiple, mobile task-swarms.

A general measure of the efficiency of information circulation in such a system would be important, allowing us to measure the effectiveness of a comm-swarm in a non-task-specific way. We have begun to develop a measure of the communication efficiency of comm-swarms that is based on the age of the information being distributed, and we are investigating the relationship between this measure and measures developed by others that might be useful in characterizing the effectiveness of comm-swarms: mixing measures [1], measures of information storage and transfer [5], and the moving average Laplacian of [4].

In a different arena, comm-swarms might be useful for particle swarm optimization (PSO), a flocking-inspired optimization technique in which virtual particles search the solution space guided by high-quality solutions found by themselves and by other particles. Communication among particles is critical to the success of the algorithm and we are currently developing a PSO variant that uses comm-swarms to provide an effective communication mechanism.

References

1. Finn, M.D., Cox, S.M., Byrne, H.M.: Mixing measures for two-dimensional chaotic Stokes flow. *Journal of Engineering Mathematics* 48, 129–155 (2004)
2. Parunak, H.V.D., Purcell, M., O’Connell, R.: Digital pheromones for autonomous coordination of swarming UAVs. *American Institute of Aeronautics and Astronautics* (2002)
3. Reynolds, C.W.: Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Comput. Graph.* 21, 25–34 (1987)
4. Skufca, J.D., Bollt, E.M.: Communication and synchronization in disconnected networks with dynamic topology: Moving neighborhood networks. *Mathematical Biosciences and Engineering* 1(2), 1–13 (2004)
5. Wang, X.R., Miller, J.M., Lizier, J.T., Prokopenko, M., Rossi, L.F.: Measuring information storage and transfer in swarms. In: *Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems*, pp. 838–845. Massachusetts Institute of Technology (2011)