2021/2500M-0136

ISSN 1797-3457 (verkkojulkaisu)
ISBN (PDF) 978-951-25-3360-2

SUMMARY REPORT

CounterSwarm – Turning Collective Intelligence against Hostile Drone Swarms

Fabrice Saffre (fabrice.saffre@vtt.fi, 040 187 7817),
Hannu Karvonen, Jari Laarni, Toni Lastusilta, Antti Väätänen
VTT Technical Research Centre of Finland Ltd

Abstract: The CounterSwarm project aims at testing the hypothesis that the best way to defend efficiently and affordably against a drone swarm is to play "tit for tat" and use another drone swarm. In nature, the species that most successfully resist an attack by social insects (the original inspiration for autonomous drone swarms) are other social insects. The reason is simple: both sides have the same characteristics in terms of their intrinsic ability to self-organise in response to highly dynamic circumstances, based on incomplete information and without any centralised control structure (and the vulnerability associated with it). The end goal of this project is to identify methods for adapting the collective intelligence paradigm to swarm tactics, and to develop conceptual, analytical, and numerical tools to evaluate performance in defence applications.

1. Introduction

The use of drones as convenient, remotely controlled airborne sensor platforms has become routine over the last decade in a wide range of application areas. Reconnaissance and surveillance missions are among these areas and have recently received considerable attention from users and providers of defence solutions. With experimental deployment now a fact, the focus of research and development has effectively moved to the next step, namely the incorporation of machine intelligence into drone design in the pursuit of ever-increasing autonomy and reduced reliance on direct human control.

Among the many aspects worthy of this domain's researchers' attention is the interaction between individual autonomous units operating simultaneously within each other's sensing and actuation range. Coordinated flight in a swarm has been the topic of many recent publications, which have highlighted the various opportunities afforded by collective behaviour and identified the challenges in implementing it in a decentralised, adaptive and dynamic way (as opposed to centralised choreography with predefined static trajectories).

In this context, the investigation of interaction patterns between multiple swarms that may have different, possibly mutually exclusive goals presents a substantial scientific and technical challenge. Of relevance to the defence sector are the dynamics of a confrontation between two autonomous swarms, one of which is trying to reach a predefined target (the attacker) and the other to prevent it (the defender).

Many examples of this type of swarm vs. swarm confrontations exist in nature and can be used as blueprints to define stochastic, rule-based decision systems that, when applied independently (but sometimes based on shared information) by individual autonomous units, can foster the emergence of a collective response to the threat posed by the other swarm. This hypothesis forms the basis and justification for the CounterSwarm project, of which objectives, methods, results, and conclusions are summarised in this report.

2. Research objectives and accomplishment plan

The research objectives of the project are two-fold:

- To establish a solid scientific foundation for the modelling and understanding of the dynamics of interaction between two mutually hostile swarms comprised of autonomous units. This framework (combining numerical and analytical methods) is expected to shed light on fundamental properties and be applicable to naturally occurring swarms as well as to their artificial counterparts.
- To develop a (collection of) simulation tool(s) designed to study the performance of various attack and defence tactics (specified in the form of decision rule sets and parameter values) in the specific scenario in which one swarm is attempting to inflict damage to a target while the other is trying to protect it. This second framework will also be used for demonstration and training purposes.

These two complementary objectives were aimed to be achieved by creating our own collection of simulation engines, which put the emphasis on machine intelligence aspects (decision-making and interaction dynamics) and not, for example, on accurate flight modelling (as many available drone simulators do). These engines were used to explore the parameters space, uncovering and documenting fundamental system properties (first objective), and identifying the key characteristics that a swarm-based counterswarm solution should possess in order to fulfil its mission (second objective).

3. Materials and methods

On a fundamental level, the confrontation between two antagonistic groups can be regarded as a series of encounters between individuals with various possible outcomes, which together determine the result of the overall conflict (one side's victory, draw, etc.). Even such a simplified model allows for the exploration of multiple relevant factors, such as the effect of numerical superiority, the relative strength of individuals (e.g., one side being comprised of units that are inherently more or less likely to win a duel) and the way they combine to generate final outcomes of variable statistical weight.

For instance, it is possible to investigate quantitatively how greater numbers can offset individual weakness (low probability of winning a one-to-one confrontation), which is of obvious interest to understand and evaluate the performance of swarm tactics. In the natural world, one can think of the way in which bees respond to a hornet attack, the interplay between the probability of individual success and sheer numbers conveniently abstracting from the specific "weapon" used by either species. In the technology realm, an equivalent scenario would be a swarm of lightweight unmanned aerial vehicles (UAVs) with limited capabilities confronting a small number of heavier, more advanced units.

This basic quantitative model was designed for simplicity and tractability, with numerical methods used to supplement mathematical analysis. It features only three independent parameters:

1. SA: the initial size/strength of swarm A
2. SB: the initial size/strength of swarm B
3. WA: the probability that A wins over B in a duel (it is assumed that every encounter ends up with one individual being eliminated, so the probability of B winning over A is simply 1 – WA)

However, this is a completely non-spatial model, which effectively hypothesises that both swarms occupy the same area and that every unit on either side is susceptible to engage every unit on the opposing side. Although this may be a suitable description for some confrontations, it does not adequately capture more complex scenarios. For instance, an attacker may be probing enemy defences in search of a weak point, which implies that there are multiple sites and vectors and that opens the possibility of more advanced manoeuvres, such as flanking. It is worth observing that such "tactics" need not be intentional or result from explicit planning but can emerge from the combination of positive and negative feedbacks known to play a critical part in the organisation of many insect societies.

To capture the spatial aspects of such a confrontation between two colonies of insects or swarms of autonomous devices governed by collective intelligence principles, another method is required, which is why a model for perimeter defence was created. This model is one-dimensional with periodic boundary conditions (i.e., a ring) and treats the confrontation as a set of local engagements between sub-populations of attackers and defenders that have the ability to move between multiple "fronts".

In many conflict situations, the objective of the attacker is to achieve a decisive advantage in one location at one point in time to break through enemy lines. The objective of the defender is to prevent this from happening by building up defences at the point of maximum pressure. This is clearly a dynamic problem for both sides, as the initial distribution of forces might be such that neither can achieve their goal (e.g., the attacker has a slight advantage in one location but not strong enough for a clear breakthrough, which is ideal for neither side). How the confrontation develops depends on the way in which each side responds to the other's actions by redeploying their forces. This model was used extensively to document the fundamental dynamics of swarm vs. swarm confrontation.

A third and last formal model was also developed that introduced some restrictions to the way in which individuals move along the perimeter. This allowed for the study of the influence of mobility (at what point and in which circumstances does an ability to redeploy faster than the opponent turn into a clear advantage) and of potential asymmetries between attacking and defending swarms (the former having to potentially travel longer distances around the perimeter).

The implemented numerical tool (i.e., the "CounterSwarm simulator") used to achieve the second research objective (to evaluate the performance of a variety of swarm tactics in a scenario involving an attack on a designated target) followed an iterative development cycle that took into account comments and suggestions from the Finnish Defence Forces (FDF) about an early prototype. The final version of the tool consisted of a Monte Carlo simulation running in a discretised 3D space one square kilometre in area and 100 m in height, comprised of one hundred thousand 10 x 10 x 10 m cubic cells. As requested by the FDF, it featured a complex geometry in which buildings of variable shapes and heights obstructed certain paths to the target (located at the centre of the modelled area, see Figure 1), resulting in "channelling" effects that could impact the relative performance of attack and defence tactics.

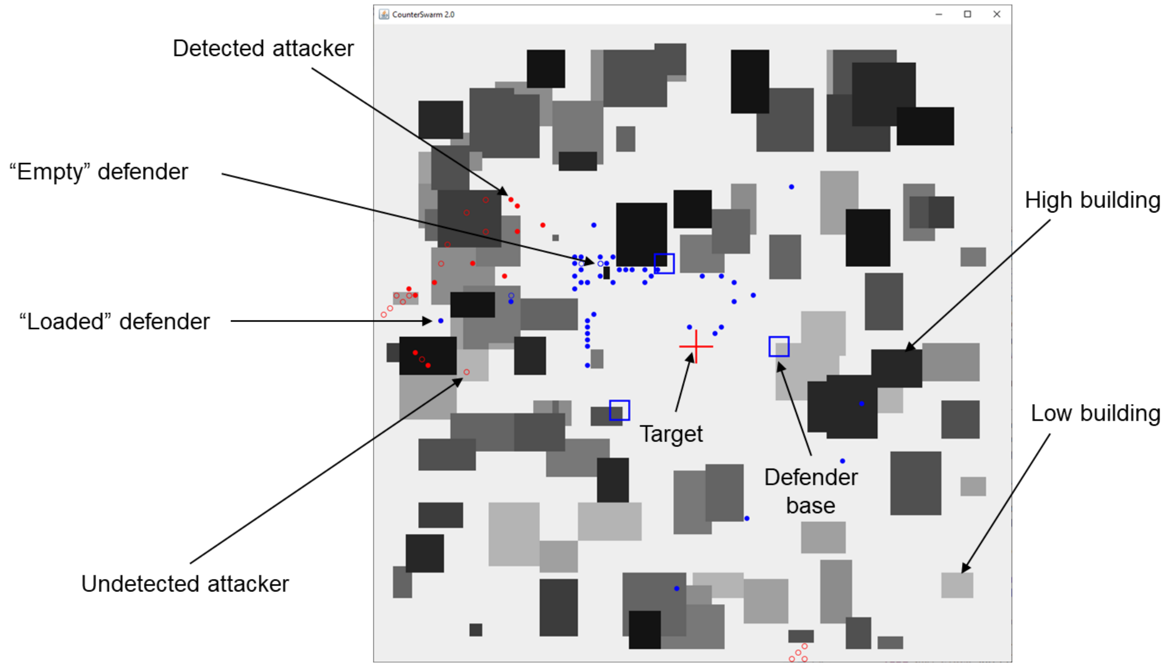

Figure 1: CounterSwarm simulator's situational picture with legends of the used symbols

Additionally, six military experts were interviewed between 13/8/2021 and 7/9/2021. These experts represented the Headquarters of the Defence Forces, the Headquarters of the Land Forces, the Naval Academy, and The Finnish Defence Research Agency. Three VTT experts conducted the interviews remotely through Microsoft Teams. Interviews were audio-recorded if the interviewee consented to audio recording. If the recording consent was not given, detailed notes were taken during the interview.

Each of the interviews lasted for about two hours, and they were divided into two main parts. First, there was a general discussion about the use of robotic swarms in military missions; and after that, there was a detailed discussion about the CounterSwarm simulator and the ways it could be developed further.

4. Results and discussion

This chapter discusses the results of the CounterSwarm project and is divided into two subsections: the quantitative model and the interview results. In addition to these summarised results, also a small-scale literature review in the beginning of the project, the developed CounterSwarm simulator (and its different scenarios), and an introductory video to the CounterSwarm simulator produced in the end of the project can be counted as the project's results.

4.1 The quantitative model

The key result from the basic quantitative model is best summarised by the equation giving the probability of swarm A winning the confrontation:

$$P(A,0) = \frac{(S_B - 1 + (S_A - A))!}{(S_B - 1)!\,(S_A - A)!} W_A{}^{S_B} (1 - W_A)^{S_A - A}$$

This equation allows the prediction of the most likely outcome based on the relative individual and collective strengths of the two opposing swarms before the encounter.

The perimeter defence model exposed the very complex behaviour of even a relatively simple approximation of collective decision-based resource allocation in a very basic set-up (two fronts). It confirmed that seemingly "advanced" tactics resembling decoy and flanking manoeuvres could emerge from simple amplification mechanisms, without central planning. An in-depth exploration of the four-dimensional parameters space revealed the presence of stable and unstable equilibria, with time-series featuring amplified and dampened oscillations depending on parameter values and initial conditions.

The third model (including variable mobility considerations and a more accurate representation of spatial constraints) is still work in progress at the time of writing. Preliminary results covering only a small region of the parameters space seem to indicate a stabilising effect, with amplification of initial heterogeneities far more common than oscillations.

The results of the evaluation of swarm tactics in the more specific scenario (i.e., complex 3D environment) take the form of a set of rules and parameter values that provide the best chance of protecting a designated target against an attacking swarm. These emphasise a combination of:

- Quick response (i.e., fly toward the attacker to intercept it some distance from the target)
- Limited recruitment capability (i.e., allow defenders to respond to a threat detected by others, but only when already close to the event)
- Restricted perimeter (i.e., chase a retreating attacker only a short distance away from the target, to avoid falling for decoy manoeuvres)

4.2 The interview results

This section's interview results summary is based on the opinions of individual experts, and thus does not present a consensus among all interviewees. It is divided into two parts: the general-level results regarding swarming and the results regarding the developed simulator in the project.

4.2.1 The general-level results regarding swarming

In general, the interviewees thought that autonomous robotic swarms can be seen as a potential game changer in how modern warfare is conducted. For example, the boundaries between different military branches may become blurry if all types of robotic swarms can be used in all branches. Robotic swarms can be used on land, at sea and in the air, and swarming can also be applied in cyberspace. A war between drone swarms is technically possible in the near future (circa 5–10 years), but poor weather and environmental conditions may compromise their effective use.

In general, ethical and judicial restrictions will prevent the use of swarms of autonomous vehicles in direct military offences. The role of humans is to set limits to the warfare between autonomous systems and prevent further escalation of the situation. However, this aim becomes more difficult, as the tempo of the warfare increases in the future.

Encirclement and simultaneity are key features of military swarming: swarming makes it possible to encircle the enemy drones and attack them simultaneously from multiple directions. A surrounding swarm can also conduct pulsing hit-and-run attacks by appearing and disappearing repeatedly. In order to achieve these positive impacts reliably a high-level of autonomy is required so that the members of the swarm can decide on how to act. Decision-making is decentralized and conducted where the operation is carried out.

Swarming also promotes flexibility: the mission can be continued, even though a large part of the drones has been destroyed. Disposability and cost-effectiveness are key indicators here: if the swarm is composed of inexpensive drones, the whole swarm can be sacrificed if needed. It is not necessarily feasible to construct swarms that include both manned and unmanned units, because they may restrict each other's abilities.

Swarming provides new opportunities to deceive the enemy by saturating the airspace or by leading new swarms periodically to the airspace. From the defender's perspective, it is very difficult to recognize armed drones carrying explosives or weapons from harmless decoy units if the airspace is saturated with these swarms of drones.

In the first phase, autonomous swarms can be used in surveillance/reconnaissance operations and in area monitoring. For example, a drone swarm could conduct long-term patrolling in a military area, detect and recognize possible unknown objects and react to them quickly. Since it is difficult to detect an attacking swarm of drones if it does not emit anything, a layered system is required for the detection and identification of the swarm. The layered system should comprise of various detection systems (e.g., sensors, radars, and lasers). It was especially emphasized by the interviewees that it may be difficult to detect and identify the swarm using only radars.

An important defensive manoeuvre against robotic swarms is interfering and/or preventing communication. One method to neutralize multiple attacking drones simultaneously is to break the electronics of a drone with an electro-magnetic pulse. The drawback of this method is that this pulse can easily destroy one's own devices at the same time. Physical mitigators, such as projectiles and collision drones, were seen especially suitable for counter-swarming.

4.2.2 Interview results regarding the developed simulator

During the latter part of the interview, we demonstrated some CounterSwarm simulator runs and discussed the key features of the simulator. In this part, one of the interviewers presented the CounterSwarm simulator to the interviewees. The interviewer noted that the purpose of discussing about the simulator is to support the development. Simulator features like swarm size, assault distance, evasion distance, guarding perimeter, counter-attack distance and shooting distance was clarified. The interviewer showed that the impact of presented control parameters enables the implementation of various tactics in the form of combinations of parameter values for the functions governing autonomous decision making by individual units, for example, pursuit, encirclement, and perimeter defence.

In their feedback, the interviewees thought that the simulator is useful in promoting tactical thinking about counter-swarming. Several parameters were recognized as important for controlling the behaviour of attacking and defending swarms, such as assault and counterattack distance and the probability of change in direction. The interviewees thought that it is important to continue to study the impacts of various control parameters on swarm behaviour in the counter-swarming context. Several improvements were suggested to the tactical control of robotic swarms, such as the possibility to divide the swarm into smaller groups, introduction of sub-tasks, optimisation of the counterattack and introduction of new scenarios (e.g., reconnaissance). Based on the interview results, the preliminary summarized structure of a swarm intelligence system to be developed with relevant topics in further research is outlined in the following:

1. Swarm intelligence: a graduated structure from autonomous units to coordinated swarm behaviour is needed
2. Swarm behaviour: dynamic grouping and sub-tasks are needed
3. Drone features: for example, speed, agility, and payload should be considered
4. Environment: buildings, terrain, fixed sensors and weaponry need consideration
5. Swarm control: initial parameters and autonomous parameter adjustments based on situation evolvement and mission goal is needed;
6. Analysis system for simulator runs: a decision-support system for making conclusions from numerous simulations needs to be developed.

5. Conclusions

In this summary report, we have presented the research objectives, methods, and results of the CounterSwarm project. The end goal of this project has been to identify methods for adapting the collective intelligence paradigm to swarm tactics, and to develop conceptual, analytical, and numerical tools to evaluate performance in defence applications. We have presented 1) a basic quantitative model, which was used extensively in the documentation of the fundamental dynamics of swarm vs. swarm confrontation, 2) an implemented numerical simulation tool (i.e., the CounterSwarm simulator), and discussed the results of expert interviews related to swarming and our CounterSwarm simulator. Next, we draw conclusions from each of these, examine critically the obtained research results, and discuss further work.

The control algorithms used to control unmanned military machines will play a crucial role for the outcome and a competition in algorithm development can be foreseen. In the basic quantitative model created in this CounterSwarm project, we (1) ignore orchestrated simultaneous actions and consider each encounter to be a "duel", always fought between exactly two opposing units and (2) simplify the outcome of an engagement to be either a win or a loss (so we are not addressing damaged devices continuing to participate in the conflict). We further (3) assume that the probability for the outcome of any such duel can be estimated using experience and historical data, and that the outcome entirely depends on which types (model / make) of the units are duelling. The created simple quantitative model allows us to determine the (probability of the) outcome, assuming we know both population sizes SA, SB as well as the probability WA of A to win over B in a duel. This in turn enables us to estimate the number of defenders needed to successfully "*hold the line*" (defeat the attackers) with a given probability P, cf. equation in section 4.1.

The introduced model allows the adaptive control over the allocation of units. In line with the expectation that semi-autonomous drone swarms will soon be a reality, we propose the presented approach as the basis for designing an adaptive control element to drive the deployment and redeployment of the members of a swarm based on the perceived strength of the enemy in the various locations (and on the expected chance of winning a duel).

The practical considerations regarding the quantitative model are namely the following: (a) when there are many locations on the frontline, then the travel distance between them will differ, and that (b) the shape of the frontline itself will affect the attacker and the defender differently (unless it is a straight line). Amending the presented equation correspondingly will break the symmetry, but we show that to address that (a) we can restrict redeployment to adjacent locations and argue with regard to (b) that this could be included by defining the adjacency of locations differently for the attacker and the defender.

The CounterSwarm simulator runs revealed that there is a great deal of swarm control parameters susceptible to affect the result of a swarm vs. swarm confrontation. In this project, only a few fundamental parameters (e.g., defender movement radius) could be thoroughly investigated and their impact on the final outcome of the confrontation in different battle scenarios properly quantified. Furthermore, based on our simulations we identified a set of individual rules (i.e., "quick response", "limited recruitment capability", and "restricted perimeter") that could give an autonomous drone swarm the best possible chance of repelling an autonomous drone swarm attack.

The conducted interviews in the project also revealed the domain experts' views regarding swarming in general, our CounterSwarm simulator, and allowed us to develop a structure for the swarm intelligence system for further development. Based on our results, substantial work yet remains to be done in order to fully understand the potential and threats of drone swarms in the defence domain and the dynamics of interaction between mutually hostile drone swarms in detail.

5.1 Critical examination of the research results and further work

A critical examination of the research results reveals that the simulations and models created in this project are only the first step. For example, the current version of the developed simulator is not advanced and realistic enough to be implemented into practical military applications regarding drone swarm control. Our reasoning behind using more of a generic approach, focusing on this starting phase first on fundamentals, is that this allows us to build a more flexible solution for future development.

We can approach further work both on the 1) scientific and 2) practical level. On the scientific level, we aim to continue the work towards a unified theory of swarm interaction in which we would be able to identify fundamental properties of swarm behaviour. If we are able to develop the existing mathematical framework further, knowing the ratio of forces between the attacker and defender would allow us to estimate the probability to win a duel. In the future, this approach would allow developing a predictive tool for this purpose.

On the practical level, the developed simulator tool and the received results of the project can be used for the training of defence experts. This would allow them to be trained, for example, in tactical thinking related to autonomous swarms and also regarding how autonomous systems can change warfare in general. In addition, the gamification of the simulator and running a competition to find out best strategies with the defence experts could be a way to enhance their motivation and interest towards the developed system. Finally, developing the CounterSwarm simulator further to consider practical military scenarios and field situations would allow it to be more directly applicable to actual physical drone swarms.

6. Scientific publishing and other reports produced by the research project

1. Saffre, F., Hildmann, H., Karvonen, H., Laarni, J., Lastusilta, T., Väätänen, A. & Deneubourg, J.-L (to be submitted to Swarm Intelligence Journal). The Dynamics of Swarm vs Swarm Confrontation. Springer.

A paper about the fundamental dynamics of interaction between mutually hostile biological or artificial swarms. The influence of key parameters is documented systematically and exhaustively through a combination of mathematical and numerical methods. Implications for the design of efficient tactics for drone swarms are discussed.

2. Laarni, J., Väätänen, A., Karvonen, H., Lastusilta, T. & Saffre, F. (under review). Development of a Concept of Operations for a Counter-Swarm Scenario. 19th International Conference on Engineering Psychology and Cognitive Ergonomics. Springer, Cham.

A paper about the development of Concept of Operations (ConOps) for a counter-swarm scenario in which the defender side uses a swarm of drones to defend a target against an attacking drone swarm. The main characteristics of the ConOps are presented as well as a summary of expert interviews.

3. Saffre, F., Hildmann, H., Karvonen, H., & Lind, T. (2022). Self-Swarming for Multi-Robot Systems Deployed for Situational Awareness. In New Developments and Environmental Applications of Drones (pp. 51-72). Springer, Cham.

A book chapter describing a decentralised framework for the surveillance of an area of arbitrary size and shape by a swarm of autonomous drones, on a scale vastly exceeding the range and autonomy of individual units. The proposed solution uses a digital pheromone for coordination and features emergent division of labour. Performance is evaluated across multiple scenarios.

4. Saffre, F., Hildmann, H., & Karvonen, H. (2021). The Design Challenges of Drone Swarm Control. In International Conference on Human-Computer Interaction (pp. 408-426). Springer, Cham.

A conference paper discussing the design of generic and reusable control strategies for drone swarms. A distinction is made between direct and indirect control methods. The former is defined as requiring at least one unit being remotely operated by a human pilot. The latter involves specifying high level objectives which are then achieved autonomously, through collective intelligence.

Additionally, the following journal publication (under review at the moment) from an earlier MATINE project RoboConOps is of relevance to this topic:
Laarni, J., Koskinen, H., & Väätänen, A. (under review in Journal of Robotics and Control). Concept of Operations as a boundary object for knowledge sharing in the design of robotic swarms.