THE JOURNAL OF NAVIGATION (2008), **61**, 9–29.
doi:10.1017/S0373463307004444 Printed in the United Kingdom

Particle Swarm Guidance System for Autonomous Unmanned Aerial Vehicles in an Air Defence Role

Alec Banks

(*Tornado In-service Software Maintenance Team*)
(Email: test-flt@tismt.raf.mod.uk)

Jonathan Vincent and Keith Phalp

(*Bournemouth University*)

This work investigates the utilisation of Particle Swarm Optimisation (PSO) for the non-deterministic navigation of Unmanned Aerial Vehicles (UAVs), allowing them to work cooperatively toward the goal of protecting a wide area against airborne attack. To negate the PSO's inherent weakness in dynamic environments, a neighbourhood scheme is proposed that not only enables the efficient interception of targets several times faster than the UAVs but also facilitates the maintenance of effective airspace coverage. Empirical results suggest that these techniques may indeed be of use in autonomous navigation systems for UAVs in air defence roles.

KEY WORDS

1. UAV. 2. Guidance system. 3. Particle Swarm Optimisation. 4. Dynamic Environment.

1. INTRODUCTION. Automated or remote controlled systems are an attractive option for military combative systems since they remove humans from dangerous environments. McCarley and Wickens (2004) indicated that remote control can be problematic for a variety of reasons such as the lack of timely sensory feedback. Further, the ability of current defence technology and strategies to deal with the possible use of Unmanned Aerial Vehicles (UAVs) in modern asymmetrical warfare is being questioned (Miasnikov 2005).

We consider the possibility that a 'swarm' of UAVs, equipped with RADAR and missile capability, could be used to defend a large airspace, autonomously, against aerial attack if strategies could be developed that allowed them to behave cooperatively. Indeed, a large amount of research has already taken place into autonomous control of UAV groups using a variety of paradigms (for example, Chandler *et al.* (2002), Alighanbari *et al.* (2003), Jin *et al.* (2003), Frew and Lawrence (2005) and Krishna *et al.* (2005)). Swarm intelligence explores the emergent properties of natural cooperative systems, such as bird flocks or ant colonies, and can be applied to search problems. From this paradigm we adopt the Particle Swarm Optimisation (PSO) approach for our investigation.

PSO uses a model of social learning that has been found to reproduce the behaviour of naturally occurring swarms. In these swarms each individual is only capable of simple behaviour but through sharing knowledge they produce complex and computationally useful emergent group behaviour that can be used to solve both continuous (Kennedy and Eberhart 1995) and discrete (Kennedy and Eberhart 1997) non-linear problems. Doctor and Venayagamoorthy (2004) indicated that this emergence could be harnessed to automate the control of UAVs to attack static targets. In this work, however, we focus on developing swarms that are able to defend fixed targets against airborne assault, by way of proof-of-concept, and as a basis for further explorations of swarm intelligence in this domain.

The remainder of this paper is organised as follows. Section 2 provides a brief overview of the PSO algorithm and its development. Section 3 describes the abstract world model, i.e. the simplifications made to the problem domain to facilitate development, visualisation and analysis. Section 4 discusses the development of the PSO algorithm for UAV guidance. Section 5 describes some of the properties of the simulation environment and section 6 details the parametric study conducted to explore the performance. Section 7 presents the results of the parametric study and section 8 comprises a general discussion on the nature of PSO and its utility in this domain. Section 9 concludes.

2. PARTICLE SWARM OPTIMISATION.

The concept of using a collection of autonomous particles that act together to produce an emergent behaviour was initially developed to solve the problem of rendering natural looking images in computer animations (Reeves 1983). For some animations it is necessary to render group behaviour (such as that of a flock of birds) with higher order dynamics than simple particles. To script the behaviour of each individual is possible, but tedious and difficult to make lifelike. Reynolds (1987) used Reeves' particle system as the basis for his higher order (in terms of objects being modelled) flocking algorithm. He took the particle movement and added orientation and inter-object communication. These additional behaviours allowed individual "Boids" (*Bird-oid* objects) to follow some simple flocking rules: Boids should avoid colliding with fellow Boids, they should attempt to match velocity vectors, and try to stay close to each other. Adopting this underlying model removed the need for the animator to specify each individual flight path.

Kennedy and Eberhart (1995) sought to extend Reynolds' model to reflect human social behaviours. More importantly, they replaced the simple roost goal developed by Heppner and Grenander (1990) with a more realistic goal – that of searching for food. It was this goal that led researchers to use non-trivial mathematical problems as the fitness function for the particles (no longer called Boids since they had become more generalised). The resulting algorithm for calculating the position update vector ($\mathbf{v}$) can be defined as:

$$v' = v_i + c_1 r_1 (p_i - x_i) + c_2 r_2 (p_g - x_i) \tag{1}$$

where c is a constant for balancing between group and individual influence, r_1 and r_2 are uniformly distributed random numbers in the range 0·0 to 1·0, p_i is the

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

previous best position of the i^{th} individual, and p_g is the previous best position of the group as a whole. For an n-dimensional space the particle velocity is calculated for each dimension and then resolved into a final vector for updating the particle's position. Thus, individual particles are accelerated toward the best solutions found, based on a combination of their personal experience and the experience of the group, hence providing a trade-off between local and global search behaviour.

Shi and Eberhart (1998) noted that without the velocity memory, the first part of equation 1, the swarm would simply contract to the global best solution found within the initial swarm boundary (providing a local search). Conversely with the velocity memory, the swarm will behave in the opposite sense, expanding to provide a global search. To assist with the balance between exploration and exploitation a modified formulation using an inertia term (ω) was introduced:

$$v' = \omega\, v_i + c_1 r_1 (p_i - x_i) + c_2 r_2 (p_g - x_i) \tag{2}$$

A further approach to controlling the behaviour of the swarm was introduced by Clerc and Kennedy (developed 1999, published 2002), where, rather than applying inertia to the velocity memory, they applied a constriction factor, χ, thus:

$$v' = \chi\,(v_i + c_1 r_1 (p_i - x_i) + c_2 r_2 (p_g - x_i)) \tag{3}$$

$$\chi = \frac{2}{\left| 2 - \phi - \sqrt{\phi^2 - 4\phi} \right|} \quad \textit{where } \phi = c_1 r_1 + c_2 r_2, \phi > 4 \tag{4}$$

In conjunction with this work, Eberhart and Shi (2000) showed that combining inertia and constriction by setting the inertia weight (ω) to χ improved performance across a wide range of problems.

It is not just particle trajectory that is important to swarm behaviour; the method by which P_g is calculated affects how particles are allowed to explore areas away from the current swarm best. Through investigations into the impact of a variety of social networks, Kennedy (1999) focussed on four basic topologies: rings, wheels, stars and random edges; sociometric network shortcuts for each were also tested. The research concluded that the topology does affect the swarm's performance but which is best is objective function dependent; the shortcuts were found to have very unpredictable effects. An inferred finding from the work was that networks that slow down communication could help prevent premature convergence in multimodal landscapes. Kennedy and Mendes (2002) supported this, concluding that the best, on average, configuration is a von Neumann topology. This is logical since it has connectivity somewhere between that of a ring (using the local neighbourhood best, which is slow and better in multimodal landscapes) and totally connected (using whole group best, which is fast and suited to simple landscapes).

3. ABSTRACT WORLD MODEL.

The purpose of this work is to investigate the potential of PSO as a guidance strategy for UAVs deployed in an air defence role. This is a hugely complex domain, and for initial proof of concept, it was therefore necessary to develop a much simplified world model to facilitate experimentation. The model is based on the following abstractions:

- All aircraft movement is restricted to a two dimensional world.
- Scaled aircraft speeds are increased to make the study feasible. Without this each simulation would take several hours. The attacking aircraft are given constant speeds relative to the defending UAV speeds, thus a relative speed of 2 allows them to fly twice as fast as the defending UAVs.
- Attacking aircraft do not fight back.
- All UAV missile launches are deemed to be successful.
- UAVs are given unlimited missiles, each with a range of 5 nautical miles.
- UAV RADAR is given a range of 40 nautical miles and a field of view of 40°. This is much narrower that would be expected from standard aircraft radar systems but made observational study easier.
- There is no ground based RADAR system.
- Where enemy aircraft are forced out of the search space it is assumed that their mission would be aborted and a score is given as if the aircraft had been destroyed. This was implemented because airspace denial was the prime objective of the UAVs in this simulation.
- To provide a benchmark against which the PSO based system could be judged a simple deterministic search pattern was implemented, loosely modelling that of a human controlled UAV group without ground-based or airborne early warning.

The nature of this work is a comparative examination of the performance improvements possible when using a PSO guidance strategy rather than a deterministic search pattern. Therefore, although some aspects of the abstract model may give rise to perceived advantages to the defending force, the results hold – i.e. that PSO is useful in this context since it performs better than a deterministic search pattern.

4. ALGORITHM DEVELOPMENT.

This section provides an insight into how the swarm algorithm was developed and some of the design decisions that were made to improve its performance. After providing the UAVs with the basic ability to move about the airspace and communicate with each other, the swarm algorithm could be developed. The constricted algorithm (equations 3 and 4) proposed by Clerc and Kennedy (2002) was used because the increased particle stability it offers is more suitable for aircraft flight, where constant changes of direction are inefficient. In Eberhart and Shi (2000) the optimal value for ϕ is given as 4·1, yielding a value for χ of 0·729. This is functionally equivalent as using the modified PSO algorithm and setting $\omega = 0{\cdot}729$ and $c_1 = c_2 = 1{\cdot}49445$. To simplify the implementation this equivalent formulation was used.

The swarm stability provided by the use of this variant of the PSO algorithm was augmented by the addition of a flight path stabiliser algorithm that forces a UAV to maintain a flight trajectory for a randomised period of time (*i.e.* it will not try to change direction on every iteration of the algorithm); the periodicity of this ranged from 1 to 20 UAV moves.

Further modifications were required to the constricted algorithm to make it suitable for the control of aircraft flight. The UAV velocity was set to a conservative maximum UAV speed and the change of direction limited by a maximum turn rate. To prevent swarm members flying away from the search area a feature inspired by Heppner and Grenander's (1990) roosting algorithm was introduced. Each UAV was assigned a ground position to defend (the roost); within the search area they flew a stochastic search pattern looking for intruders, but if they exited the defined search area their target became the roost, drawing them back into the search space, whereupon stochastic search is resumed. At all times, detected targets took priority.

With the introduction of multiple targets, it was observed that when a UAV was tracking one target and another target was discovered by another swarm member, the UAV would be drawn away from the target it was tracking. This was because the 'full' PSO algorithm produces a velocity vector that is a balance between individual and group best positions. Where the strength of these influences is the result of a fitness function and the optimal positions remain static this is not problematic. In the air defence application the target is dynamic and if the UAV is drawn away, there is a possibility that RADAR lock will be lost prematurely. To overcome this, inspiration was taken from Kennedy's (1997) social models. Initially, the swarm operated in *search mode* with no swarming behaviour taking place. When in *search mode*, if a UAV receives information from another swarm member regarding a detected target it enters *swarm mode* (using the "social-only" model). Upon detecting its own target, it enters *engage mode* and removes influences from other UAVs through switching to the "cognition-only" model. In Kennedy's work this model suffered because individuals failed to find optimal regions. In this implementation there is no such problem since the UAV is already in an optimal region (*i.e*. it was detecting a target). Switching modes in this way is not only viable within the human metaphor, it produces behaviour in keeping with animal based teams (Anderson and Franks, 2001).

A further problem presented by multiple targets occurred when a UAV detected more than one target simultaneously. With targets of equal perceived value a human operator might be expected to target the aircraft that they had the highest probability of destroying with the missile type fitted to the aircraft. This rationale was employed in a missile dependent fitness function (in this scenario: 1 + the cosine of the UAV heading with respect to the reciprocal of the target heading) that allowed the UAV to discriminate between simultaneously detected targets. In addition, the problem of maintaining search coverage (swarm diversity) after the initial target detection must be addressed. One could also argue for maintenance of diversity on the grounds of efficiency, since there is no point in a UAV joining in a swarming action against a target it cannot possibly reach. To support diversity, a dynamic neighbourhood was developed based on the UAVs ability to position itself such that it would be able to launch a missile attack. This was termed the Launch Success Zone (LSZ) neighbourhood. A UAV is considered a member of an LSZ neighbourhood if its bearing from target is $\pm 40^{\circ}$ and the target distance is less than the RADAR range multiplied by the relative target speed. The application of the LSZ neighbourhood is illustrated in Figure 1. In the scenario shown, UAV 1 detects the hostile aircraft, calculates the velocity vector and transmits this to the other UAVs. UAV 4 is not able to manoeuvre into a suitable position to engage the hostile aircraft and so does not waste resources attempting to do so. UAVs 1, 2 and 3 are in the LSZ and are able to manoeuvre into a firing position.

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

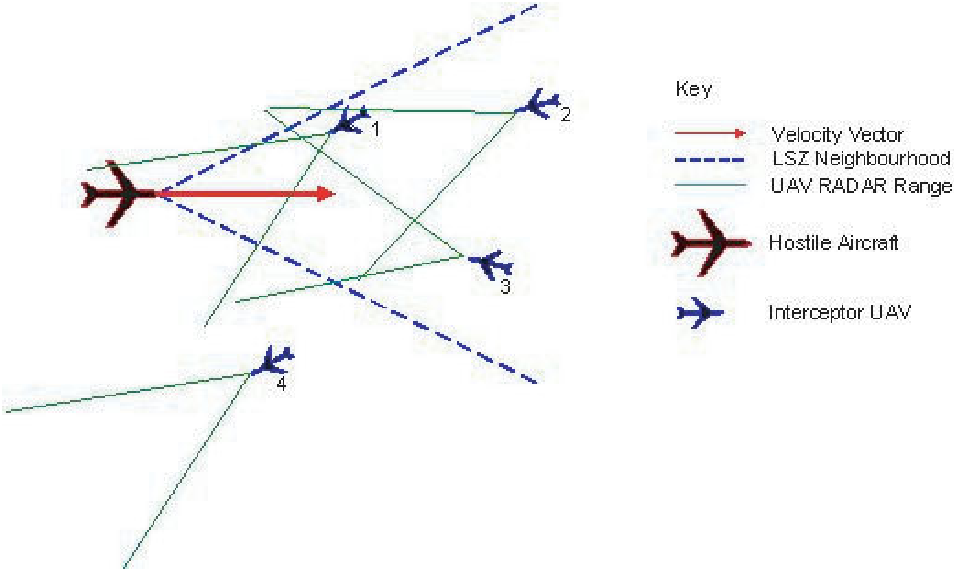

Figure 1. LSZ neighbourhood scenario.

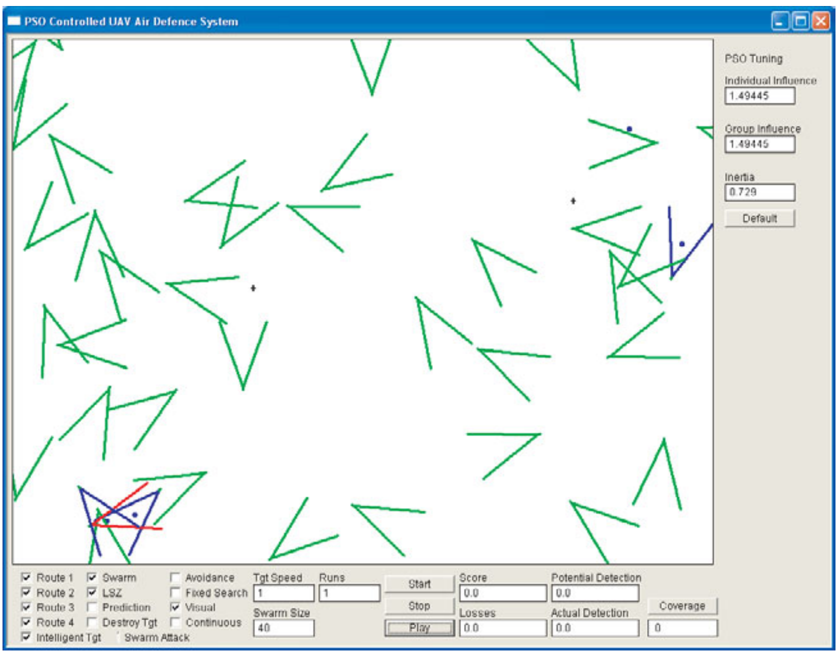

Figure 2. The simulation environment.

5. SIMULATION ENVIRONMENT. To enable observational and parametric studies, a simulation environment (illustrated in Figure 2) was constructed in Ada 95 using the John English Windows Library (JEWL) (English, 2000) and AdaCore's Gnat Programming System (AdaCore, 2005).

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

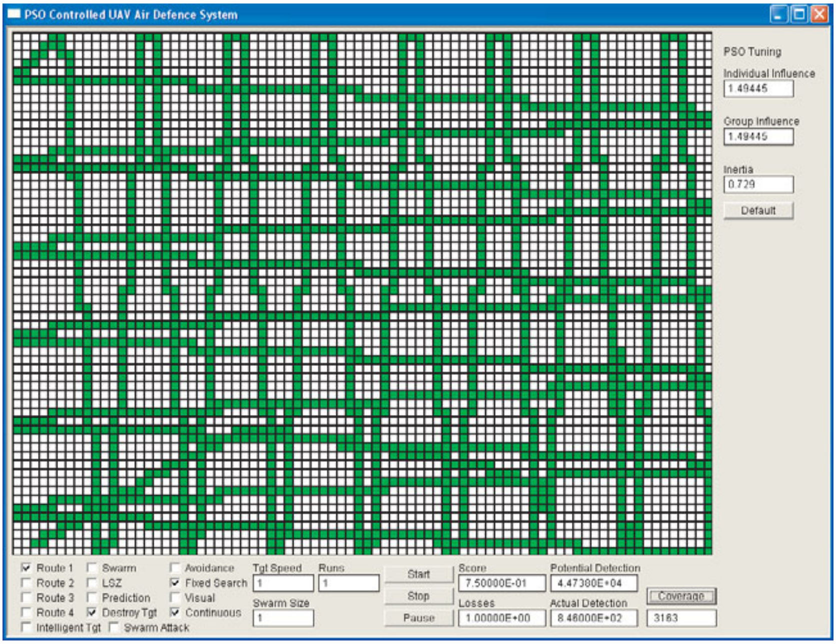

Figure 3. Coverage map arising from fixed search pattern.

The environment offered the following features:

- A two dimensional representation of an area of airspace measuring 400 miles by 300 miles. Each pixel on the display represents a 0·25 square mile. Within the airspace the defended positions are shown by a '+' symbol and the UAVs as small dots with a 'V representing the UAV's RADAR Field Of View (FOV), aligned with the direction of travel. Targets are shown as small discs.
- The ability to launch attack profiles supporting the parametric study. The targets have simple behaviour, but they do have some ability to attempt to avoid being destroyed. This behaviour is made possible through a simulation of a Radar Warning Receiver. In response, the targets turn away from their preset flight path, continuing until they are twice the UAV missile range away, then attempt to fly around the tracking UAVs and outpace them[1].
- Variable swarm configuration, where any size of swarm can be modelled and their performance explored based on either a deterministic search and rescue grid pattern or stochastic search provided by the PSO algorithm.
- A graphical representation of the frequency of occupancy (coverage) by the UAVs of each 25 mile^2 can be produced to aid the visualisation of search and intercept behaviour. Figure 3 illustrates the coverage map following the deterministic search. Frequency of occupancy is indicated by the level of shading (note that this is not RADAR coverage).

[1] This was found to be an effective strategy and, in a sample trial using 4 targets travelling at twice UAV speed against 40 defending UAVs, the defended positions were hit on average more than twice as often (2·85000E-01 against 1·02000E-01).

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

- Three performance metrics are provided:
  - *Score*. This is the average score for that particular experiment. Each time an enemy aircraft is destroyed it is given a score based on the following pseudocode:

    if not bomb_released then
    score:=(distance_to_defended_position − bomb_range)/5.0
    else
    score:=intrinsic_value_of_aircraft_ to_enemy
    end if

    Note: the score is scaled to avoid aircraft being destroyed a long way from the defended position dominating ones that are destroyed closer. The intrinsic value of the aircraft for this project was set to 0·75, which attaches less importance to destroying the aircraft than airspace denial.
  - *Losses*. The average number of times the defended positions have been bombed during the experiment. This is considered the most important indicator of swarm efficacy.
  - *Coverage*. This is a count of the number of map squares that have not been entered by a UAV during the last simulation run.

6. PARAMETRIC STUDY. To ascertain the potential performance of a PSO controlled UAV air defence system a parametric study was performed using the default PSO settings as advised by Eberhart and Shi (2000). This section describes the setup of the study, with results presented in section 7.

6.1. *Experiment Configurations*. Four attack profile experiments were designed to measure the performance of swarm sizes of 20, 40 and 60 against intruder speeds of 1, 2 and 4 times that of the UAV. The profiles were as follows (refer to Figure 4 for illustration of routes):

- *Single Target*. A single attacking aircraft following route 1. This was devised to test against an opportunistic attack such as might be launched by a terrorist organisation (Miasnikov 2005).
- *Two Targets, Two Attackers*. Two aircraft attack two separate targets following routes 1 and 2. This tests the system's ability to deal with diversionary attack.
- *Two Targets, Four Attackers*. Route 1 and 3 constituted a roughly parallel attack against target 1, whilst routes 2 and 4 provide a split attack against target 2. This represents a higher intensity attack against the defended positions, designed to test a more likely conventional attack scenario.
- *Swarm attack*. This profile consists of the four routes described above plus an additional 20 attackers all taking different routes against the targets. This tests the ability of the systems to defend against massive odds.

6.2. *Benchmark Experiment*. The purpose here was to establish a benchmark against which the performance of a swarm of PSO guided UAVs could be compared. Groups of independent UAVs were flown around a fixed search pattern, using avoidance to ensure good search space coverage (see Figure 5). When a target was

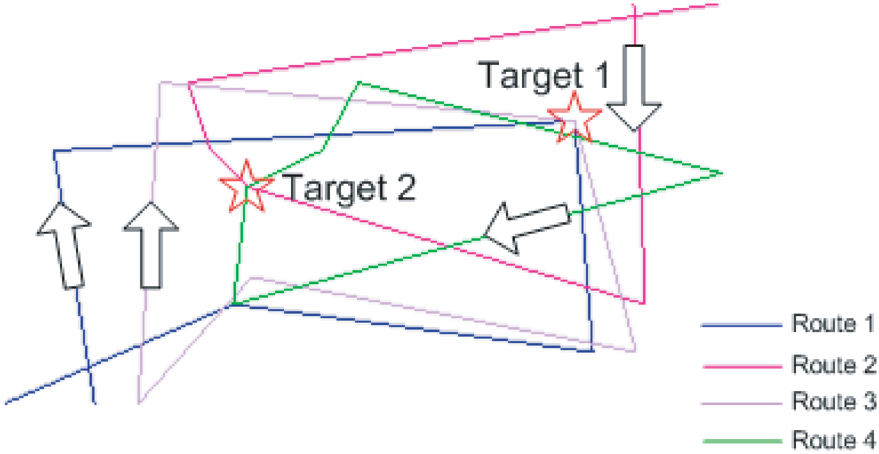

Figure 4. Attack profiles.

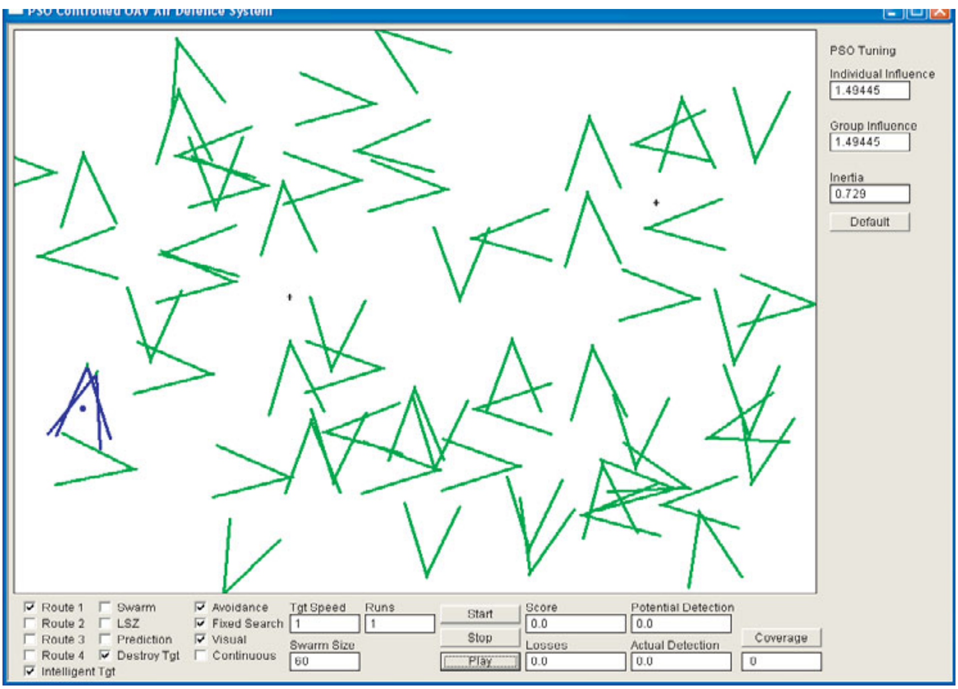

Figure 5. Fixed search pattern with avoidance.

detected, the detector would attempt to get within firing range and destroy it where possible, before resuming the search pattern.

6.3. *Swarm-Only*. Swarms were allowed to engage in stochastic search under the same conditions as the benchmark. When targets were detected the whole swarm responded, as per Figure 6.

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

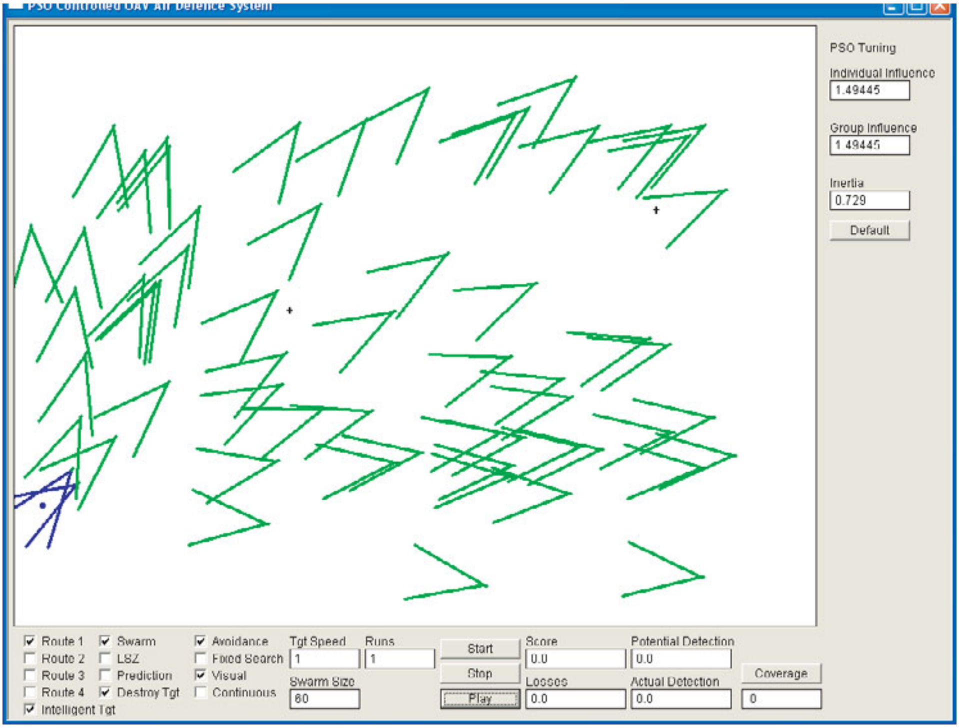

Figure 6. Swarming UAVs attack a single target.

6.4. *LSZ Neighbourhood.* The experiments were repeated under the same conditions as for swarm-only, except the LSZ neighbourhood was employed. Figure 7 illustrates how the LSZ operates in practice. Initially the target was travelling up the screen on route 1. At point A, UAV 1 detected it, causing it to change direction to avoid being destroyed. UAV 1 communicated the target velocity vector to the swarm, and two other UAVs (2 and 3) met the criteria for membership of the LSZ neighbourhood and, believing they could provide assistance, have steered accordingly. Notice that the rest of the swarm are still searching for new targets, since they cannot assist with the current target engagement.

6.5. *Swarm Attack.* A conventional 'swarm attack' does not refer to the use of techniques from the field of swarm intelligence; instead, the term is used to convey the large numbers of simultaneous attacks on a target (see Figure 8). The simulation was modified to allow the inclusion of an additional 20 attackers, all with different attack profiles; half attacking defended position 1 and the other half position 2.

7. COMPARISON OF PERFORMANCE.

The performance of each guidance system (benchmark, swarm-only, and swarm with LSZ) was evaluated on the basis of score and losses. The following section presents graphical representations of performance under varying attack profiles and target velocities (expressed as a multiple of UAV velocity). Swarm sizes of 20, 40 and 60 and target velocities of 1, 2 and 4 were used. Results reported are an average of 1000 trials of each configuration.

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

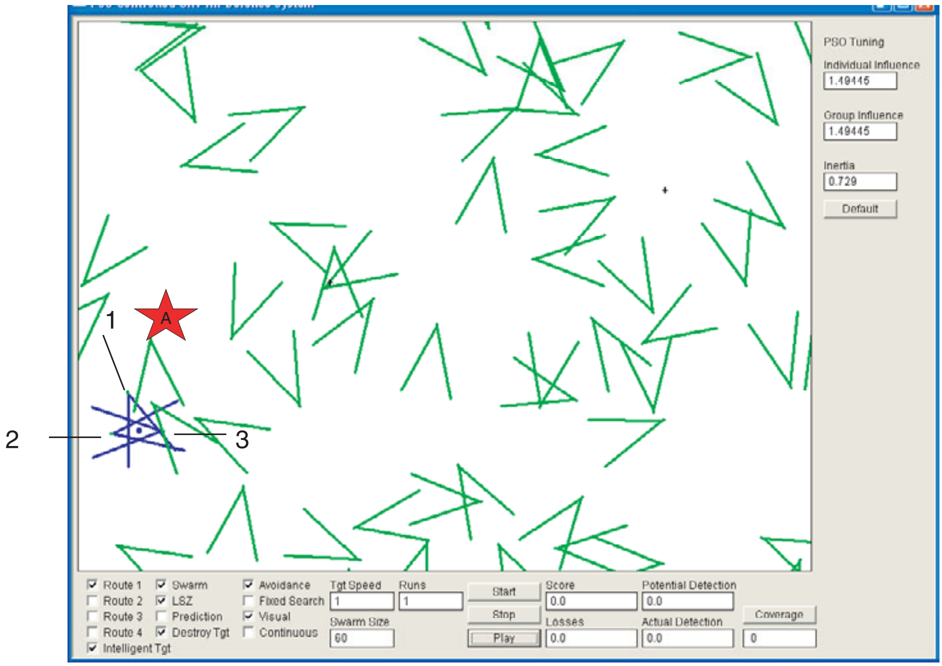

Figure 7. LSZ controlled swarming.

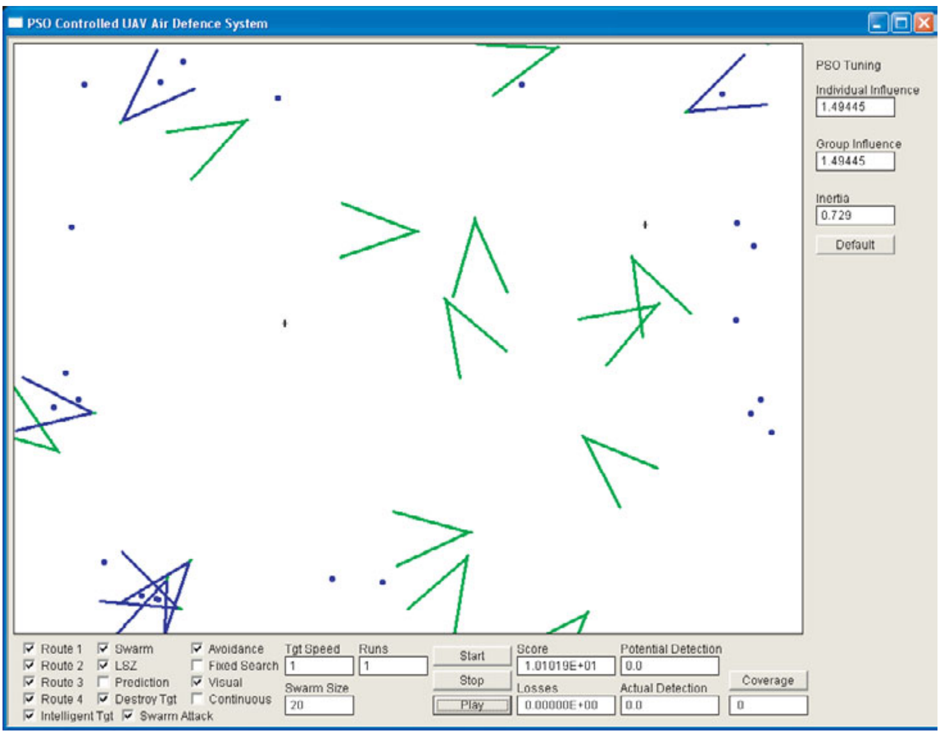

Figure 8. Swarm attack.

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

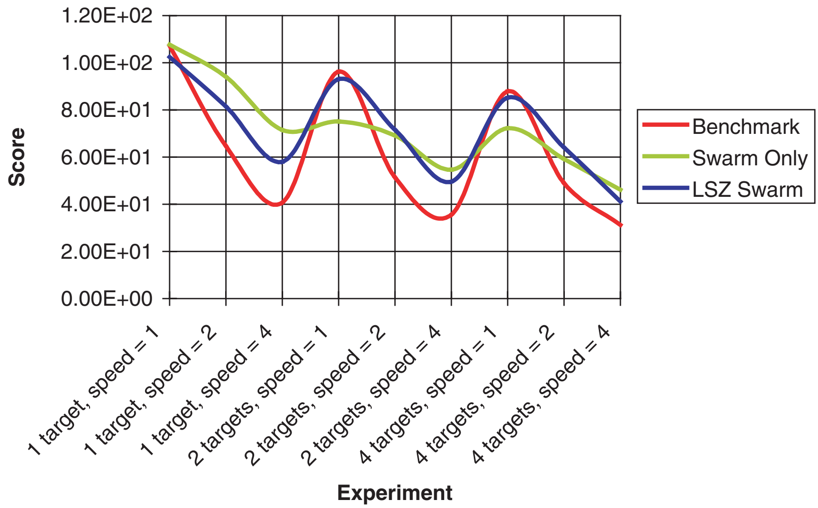

Figure 9. Average score for swarm size = 20.

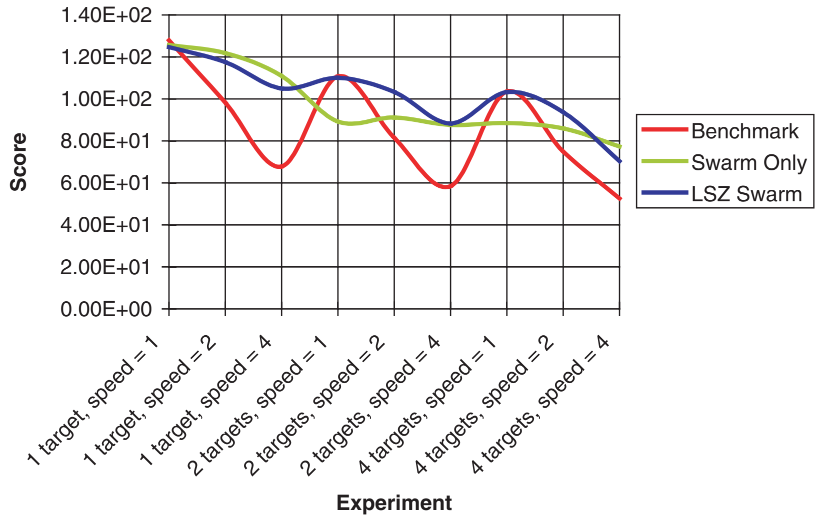

Figure 10. Average score for swarm size = 40.

7.1. *Score*. The relationship between the performance of the guidance systems was consistent in nature across swarm sizes, with the LSZ controlled swarm performing better across the largest range of scenarios (see Figures 9 to 11). With small numbers of UAVs, the benchmark performed well against low speed targets because of the increased likelihood of detection through the efficient deterministic search pattern. Once detection had been achieved, the attacker could not outrun the defender and the benefits of cooperation were not required. At higher target speeds, though, the system performed less well since targets could outrun the detecting aircraft without running into cooperating defenders.

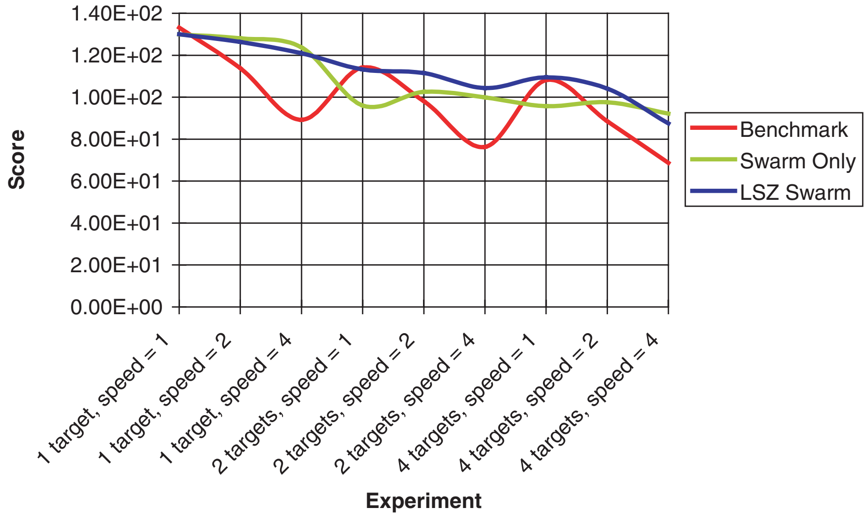

Figure 11. Average score for swarm size = 60.

Once the swarms were large enough (*e.g.* 40 and above) it was good to note the consistency of the LSZ swarm, against varying attack profiles and target velocities. This can be attributed to the search area remaining defended when multiple attacks occurred, and the LSZ was less sensitive to relative velocity because it adapted to faster attacking aircraft and there were sufficient UAVs to make use of the dynamic nature of the neighbourhood.

The effect on the performance of the swarm-only system against multiple targets at slow speed is again highlighted when compared to the performance of the other systems. Particularly in Figures 10 and 11, it can be seen that swarm-only performance drops against slower speed targets whilst the other systems improve because they find the targets easier to detect and destroy without being distracted by early detections.

7.2. *Losses.* The losses metric can be regarded as a measure of the efficacy of the system in defending the ground positions from attack. The results (Figures 12 to 14) followed the same pattern as for the score, with the two swarming algorithms improving markedly when swarm sizes were sufficient (*e.g.* for this size search area at least 40 UAVs). Again the swarm-only system was outperformed by the benchmark for slow speed multiple target attacks but the difference was only marginal and the swarm-only advantage where target speed is higher could make it more attractive.

In Figure 12, where the swarm size is insufficient to exploit the benefits of the cooperative nature of the LSZ it can be seen that it behaves very much like the benchmark, albeit with a slight improvement (where occasional cooperation would occur when there happened to be UAVs close enough). As swarm size increased, in Figures 13 and 14, the behaviour became much more consistent, with the latter showing that the system was only really under pressure with four targets travelling four times their speed. Thus, the LSZ always outperforms the benchmark and, with sufficient UAVs, the swarm-only system (except in the final scenario of 4 targets

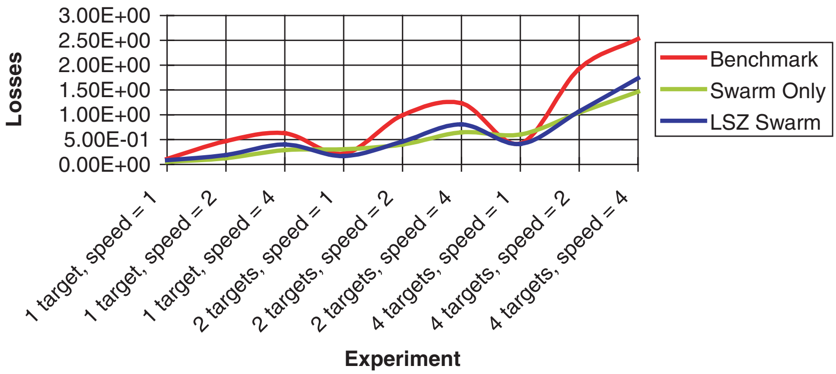

Figure 12. Average losses with swarm size = 20.

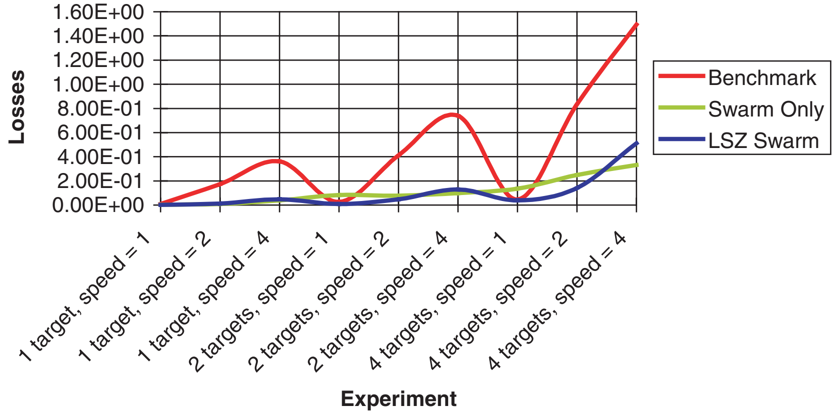

Figure 13. Average losses with swarm size = 40.

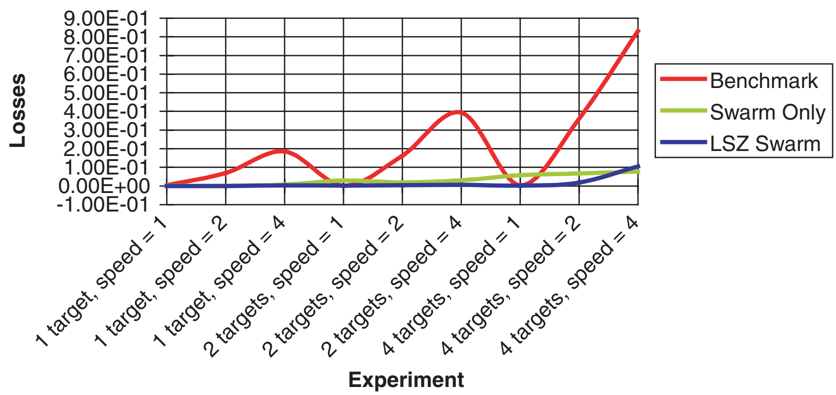

Figure 14. Average losses with Swarm Size = 60.

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

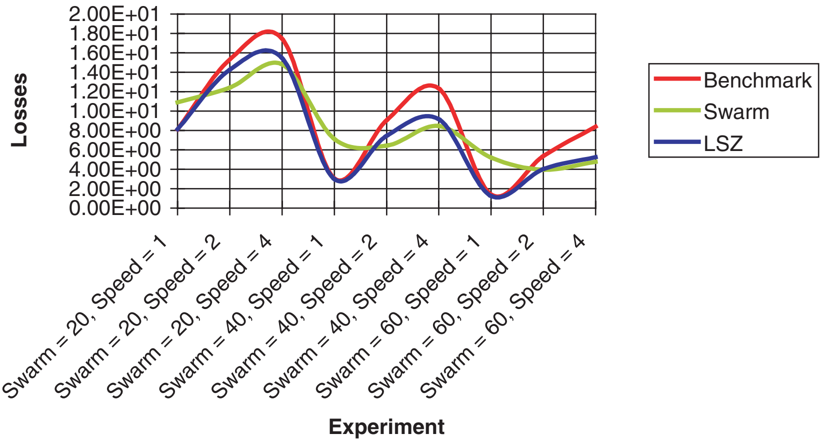

Figure 15. Performance of systems against a 24 aircraft swarm attack.

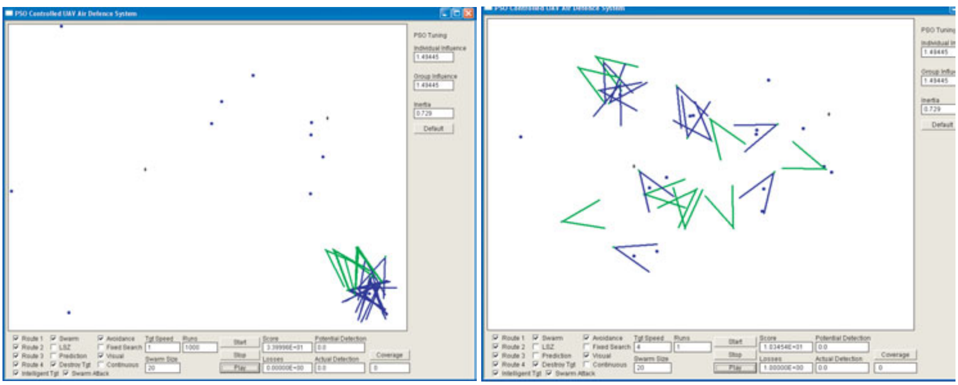

Figure 16. The swarm-only system produced a poor diversity of response at speed = 1, leaving defended positions unprotected (left). At speed = 4, the system is gradually drawn toward the middle of the search space (right).

at 4 times UAV velocity). One might expect that tuning of the LSZ (e.g. in terms of defining angle and distance) to improve matters further – i.e. extending the neighbourhood when faced with faster moving targets. This consistency against different attack profiles is very attractive when considering a defensive system.

7.3. *Swarm Attack*. Effectiveness against swarm attack can be shown through the losses metric. Again the LSZ swarm consistently outperformed the Benchmark because, whilst both systems maintained airspace coverage, the cooperative nature of the swarm allowed the LSZ UAVs to destroy targets that would have otherwise escaped by outpacing the defenders. This is highlighted by the increase in performance of the LSZ swarm at higher speed differentials (see Figure 15).

However, although the swarm-only system generally performed poorly due to homogeneous response (see Figure 16), it did perform best at higher speeds. This was due to the targets spending less time within the UAV RADAR range, thereby reducing the amount of swarming time against a single target, causing the UAVs

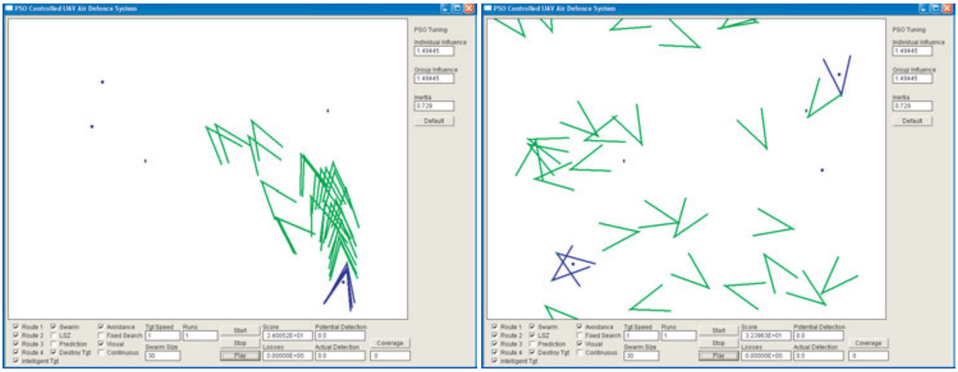

Figure 17. Lack of diversity in the swarm-only system (left). Diversity of response is maintained with LSZ enabled (right).

to concentrate toward the centre of the defended area. This could be regarded as a 'sensible' behaviour since at times of intense pressure it is a good defensive tactic to withdraw outlying defenders and concentrate them round the point to be defended.

7.4. *Diversity Of Response*. The lack of response diversity in the swarm-only system against multiple targets has been highlighted. Four aircraft, travelling at UAV speed, were set to attack the two defended positions, which were guarded by 30 swarming UAVs. Figure 17 illustrates that the aircraft detected first have the ability to draw the swarm away from the defended position and the other attacking aircraft. This homogenous response (or lack of diversity) renders the swarming system open to the exploitation of its emergent behaviour. With the LSZ enabled, however, diversity is maintained and early detections do not distract the whole swarm.

8. DISCUSSION.

The following discussion explores a range of issues based on swarming using the LSZ neighbourhood as this was the most effective system overall.

8.1. *Categorisation As Swarm Intelligence*. Whilst visually the UAVs appear to exhibit properties of a swarm, it is useful to apply criteria that assess whether swarm intelligence exists and thereby establish whether the advantages offered by swarming are available. Millonas (1994) identified five basic principles that could be used to identify swarm intelligence, and these principles are shown to exist within the LSZ system as follows:

- *Proximity*. All swarm members should be able to perform elementary space and time computations. These are fundamental behaviours for the autonomous swarming UAVs. Just as natural organisms use such calculations to perform activities such as searching for food, the UAV sensors detect the velocity vector of targets and friendly aircraft, and, combined with knowledge of their own velocity vector, attempt to compute an immediate course of action that will have the highest likelihood of achieving the group goal.
- *Quality*. The swarm members should be able to respond to quality measures with respect to the attainment of their individual goals. This was achieved through the application of a fitness function, which considers the relative positions of

the UAV and the target aircraft, their bearing, and the missile characteristics. With the ability to ascertain the quality of a given solution to its goals, the UAV can select the most appropriate target from several possibilities. A further response to a quality measure was the UAVs self-determination with respect to LSZ membership – the quality measure being the relative target and UAV positions and the velocity vector of the target, the goal being an efficient response to attack and maintenance of airspace coverage.

- *Diversity Of Response*. The swarm should not excessively commit its resources in a manner that would leave it vulnerable. In previous work (Kennedy 1999, Kennedy and Mendes 2002), the sociometric structure of the swarm was seen to effect the rate of convergence on a particular solution. Where the solution space had a simple landscape this was not problematic. However, where the landscape was complex, convergence on early solutions had a large probability of being suboptimal. In the air defence scenario, if the whole swarm converges on the first enemy aircraft detected this too could be premature, since other aircraft could enter the controlled airspace vacated by the swarming UAVs. By maintaining a diversity of response, as achieved with LSZ neighbourhood, this can be avoided. The LSZ neighbourhood originated from the idea that it would only be useful for a UAV to attempt to converge on a target if it would be able to manoeuvre into a position where a missile could be launched successfully. Once implemented, it was found to have the more useful property of maintaining diversity of response, and without it the swarm could over-commit to a single, possibly diversionary, target, leaving the defended position vulnerable to other attacks.
- *Stability*. This principle is important to UAVs, because an unstable swarm, that is one that changes its behaviour too often, is inefficient. The inefficiency is that, when a goal changes for a swarm member, it will usually have to move to another part of the search space or radically change direction. Utilising the LSZ neighbourhood increases swarm stability since it gives the UAV the ability to decide for itself whether it is efficient use of its resources to change mode.
- *Adaptability*. All the swarm members should be able to switch to a different mode of operation, but only when there is a sufficient return on investment. The tuning of the LSZ is critical to this principle because if it set too small then it is less likely that UAVs finding themselves included will have sufficient time to manoeuvre into a firing position. Conversely, if the LSZ neighbourhood is too large then UAVs may be switching mode unnecessarily, either because other UAVs are likely to have destroyed the target long before the more distant ones get into firing range, or because the target will have changed direction, causing unnecessary mode changes.

8.2. *Network Topology*. For optimisation problems, selection of a suitable topology is usually dependant on the fitness landscape (Kennedy and Mendes, 2002), which for the airspace defence problem is dynamic and unpredictable. If it were possible to predict a unimodal landscape, total connectivity with all UAVs swarming toward the target would be most appropriate since it provides a homogenous response, but in a multimodal landscape this strategy could leave large areas of the airspace open to exploitation by other attackers. This problem was alleviated by

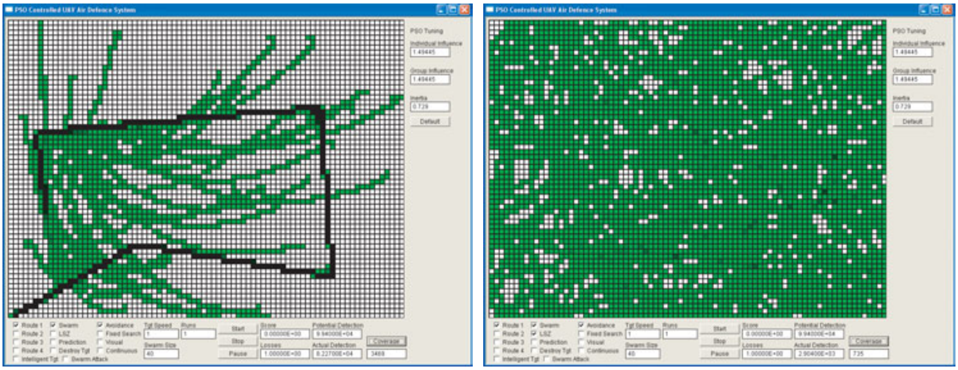

Figure 18. Swarm Coverage without LSZ neighbourhood (left) and with LSZ neighbourhood (right).

the introduction of the LSZ neighbourhood topology. Whilst all of the swarm members were connected due to the broadcast communication strategy employed in the simulation, each individual decided whether it was profitable to react by calculating its fitness with respect to the received target information. The importance of this to the quality of defence system is illustrated in section 7·4 and by comparing the levels of coverage shown in Figure 18. In this experiment a medium size swarm (40 members) was challenged by a single, same speed, target that was allowed to fly for one complete route. Without the LSZ neighbourhood 3469 squares remained untravelled by UAVs in contrast with 735 with LSZ selected.

8.3. *Rules*. Examination of the swarm members' behaviour reveals that each follows a simple set of rules. This may appear contradictory to the philosophy of autonomous behaviour, but without these rules there would be no system; each UAV would fly aimlessly around the search space, and engagement with targets would be purely on chance. The defining aspect of the rules is that they are low-level; an external leadership hierarchy does not govern each individual's behaviour, but the combination of individual behaviours determines the emergent group behaviour. The rules that each individual must follow are:

- Record and communicate locations of detected targets.
- Fly toward the source of an LSZ neighbourhood that the UAV finds itself part of.
- Where an individual finds itself in multiple LSZ neighbourhoods it must decide which would be the most profitable source to pursue.
- When an individual is not part of a LSZ neighbourhood it should engage in a stochastic search, staying within a given area to maintain swarm density.

8.4. *Feedback*. Without feedback, systems cannot adapt to dynamic environments. The feedback in particle swarm systems cause the UAVs in this implementation to acquire a heading toward the target and maintain the heading toward it whilst the target is being detected. Loss of detection breaks the feedback loop causing the swarm members to resume the search mode.

8.5. *Determinism*. An implementation of the UAV defence system would be regarded as safety critical; its failure is unacceptable on the grounds of risk to human

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

life. Currently such systems must be deterministic in the sense that it must be possible to predict its behaviour given a certain set of circumstances[2]. To an extent each UAV's behaviour is predictable; they all follow a set of low-level rules. Their flight trajectories appear to be stochastic in that they are not predictable to an observer, without knowledge of the algorithms and the random number sequence. One would not, perhaps, expect such a system to be fully autonomous, and some form of oversight might be expected. This work is really about the guidance system, rather than the process of destroying a target. As artificial intelligence paradigms take more control of defence systems, maybe the next generation of electronic warfare systems will be behaviour analysts and predictors?

9. CONCLUSIONS AND FURTHER WORK.

This research has examined the possibility of employing Particle Swarm Optimisation (PSO), in the development of a guidance strategy for UAVs in an air defence role. The precise nature of the task is secondary to the assessment of the utility of PSO in facilitating cooperative search, tracking and interception, which may find more general application.

A simulation environment based on an abstract world model was created to facilitate the exploration of PSO concepts in a representative, but much simplified, air defence scenario. A parametric study was undertaken to compare the performance of three defensive strategies: a benchmark configuration where UAVs adopted a deterministic search pattern, a swarm-only system, and a swarm with LSZ neighbourhood. At slow target speeds the benchmark system worked well due to its efficient systematic search and the ability of individuals to destroy the targets without cooperation. Once target speeds increased, performance declined because targets were able to escape from the detecting aircraft. The swarm-only system worked well against single targets because, once detection occurred, the whole swarm worked cooperatively to defend the area. This homogeneous response hindered the system when multiple targets were involved, especially at slower speeds, because once the first target was detected the whole swarm would be drawn toward it leaving large areas of the airspace vacant, to be exploited by other targets. The dynamic neighbourhood of the LSZ system overcame this problem to a large extent. When a target was detected, only those UAVs that considered, through the application of a set of low-level rules, that they were in a position to be able to destroy the target actually reacted to it. This meant that those who did not consider themselves part of the neighbourhood could continue to search for other targets. The stochastic search of both swarm systems meant that smaller swarms did not perform well. This was due to the inefficiency of searching randomly over a large area using few resources. The performance of all three systems in defending against a large scale attack was also explored. Interestingly, the swarm-only system performed well against high-speed attack, because the swarm tended to be drawn toward the focus of the attack and the concentration of UAVs resulted in better defence of that area.

For the simplified application under investigation the particle swarm approach to cooperative guidance improved the ability of a group of UAVs to defend locations within a specified area, when compared with a more traditional fixed search pattern.

[2] Although BAE Systems have recently announced its commercial development of autonomous UAVs (Robinson, 2006).

https://doi.org/10.1017/S0373463307004444 Published online by Cambridge University Press

This remained true for isolated attacks as well as for overwhelming swarm attacks. Critically, the proposed dynamic neighbourhood scheme has not only allowed diversity of response to be maintained, but also enabled the defending UAVs to intercept much faster targets. Thus, the development of the LSZ is considered a valuable contribution, enabling a higher level of swarm intelligence to be realised.

This work has demonstrated the potential for PSO to contribute to cooperative guidance in complex high-speed multimodal domains. However, this is at the earliest stage of development and further work is in progress to explore the use of biologically plausible solutions to the problems of search efficiency and maintenance of target pursuit post detection loss.

ACKNOWLEDGEMENTS

The authors would like to thank Officer Commanding Tornado In-Service Software Maintenance Team, Squadron Leader Darren Goldstraw for his support of this work. The views expressed herein are solely those of the authors.

REFERENCES

AdaCore, (2005). Ada Academic Initiative. Available from: http://www.adacore.com/academic_overview.php [accessed 26 January 2005].

Alighanbari, M., Kuwata, Y. and How, J. P. (2003). Coordination and Control of Multiple UAVs with Timing Constraints and Loitering. IEEE American Control Conference 2003.

Anderson, C. and Franks, N. R. (2001). Teams in animal societies. Behavioral Ecology Vol. 12, No. 5, pp. 534–540.

Chandler P, Pachter, M., Swaroop, D., Fowler, J, *et al*. (2002). Complexity in UAV cooperative control. IEEE American Control Conference 2002.

Clerc, M. and Kennedy, J. (2002). The particle Swarm: explosion, stability and convergence in a multi-dimensional complex space. IEEE Transactions on Evolutionary Computation No.6, 2002

Doctor, S. and Venayagamoorthy, G. K. (2004). Unmanned Vehicle Navigation Using Swarm Intelligence. International Conference on Intelligent Sensing and Information Processing, Chennai, India, pp. 249–253, 2004.

Eberhart, R. C. and Shi, Y. (2000). Comparing inertia weights and constriction factors in particle swarm optimization. Proc. 2000 Congr. Evolutionary Computing, pp 84–88, 2000.

English, J. (2000). About JEWL. Available from: http://www.it.bton.ac.uk/staff/je/jewl/ [Accessed 22 June 2005].

Frew, E. W. and Lawrence, D. A. (2005). Cooperative Stand-off Tracking of Moving Targets by a Team of Autonomous Aircraft. AIAA Guidance, Navigation, and Control Conference, San Francisco, CA, August 2005.

Heppner, F. and Grenander, U. (1990). A stochastic nonlinear model for coordinated bird flocks. In S. Krasner, Ed., The Ubiquity of Chaos. Washington DC, USA: AAAS Publications, 1990.

Jin, Y., Minai, A. A. and Polycarpou, M. M. (2003). Cooperative real-time search and task allocation in UAV Teams. In Proceedings of the IEEE Conference on Decision and Control, pp. 7–12.

Kennedy, J. and Eberhart, R. C. (1995). Particle swarm optimization. Proceedings of IEEE International Conference on Neural Networks, Piscataway, NJ. pp. 1942–1948, 1995.

Kennedy, J. (1997). The particle swarm: Social adaptation of knowledge. Proceedings of the 1997 International Conference on Evolutionary Computation, pp. 303–308. IEEE Service Center, Piscataway, NJ, 1997.

Kennedy, J. and Eberhart, R. C. (1997). A discrete binary version of the particle swarm algorithm. Proceedings of the World Multiconference on Systemics, Cybernetics and Informatics 1997, Piscataway, NJ. pp. 4104–4109, 1997.

Kennedy, J. (1999). Small Worlds and Mega-Minds: effects of neighborhood topology on particle swarm performance. *Proc. Congress on Evolutionary Computation 1999*, pp. 1931–1938. Piscataway, NJ: IEEE Service Center, 1999.

Kennedy, J. and Mendes R. (2002). Population structure and particle swarm performance. Proceedings of the Congress on Evolutionary Computation, 2002. CEC 02., Volume 2, pp. 1671–1676, 2002.

Krishna, K. M., Hexmoor, H., Pasupuleti, S. and Llinas, J. (2005). Parametric control of multiple unmanned air vehicles over an unknown hostile territory. In 2005 Proceedings of Knowledge Intensive Multiagent Systems (KIMAS-05), pages 117–121. Boston, MA

McCarley, J. S. and Wickens, C. D. (2004). Human factors concerns in UAV flight. University of Illinois at Urbana-Champaign Institute of Aviation, Aviation Human Factors Division, 2004.

Miasnikov, E. (2005). Threat of Terrorism Using Unmanned Aerial Vehicles: Technical Aspects. Center for Arms Control, Energy and Environmental Studies Moscow Institute of Physics and Technology 2005.

Millonas, M. M. (1994). Swarms, Phase Transitions, and Collective Intelligence. Artificial Life III. Ed., Langton, C.G., SFI Studies in the Sciences of Complexity, Proc. Vol. XVII, Addison-Wesley, 1994.

Reeves, W. T. (1983). Particle systems – a technique for modeling a class of fuzzy objects. ACM Transactions on Graphics, Vol. 2, No. 2, pp 91–108, April 1983.

Reynolds, C. W. (1987). Flocks, Herds, and Schools: A Distributed Behavioral Model, in Computer Graphics, 21(4) (SIGGRAPH '87 Conference Proceedings) pp 25–34. 1987.

Robinson, T. (2006). Out of the black. Aerospace International, April 2006. Royal Aeronautical Society.

Shi, Y. and Eberhart, R. C. (1998). A modified particle swarm optimizer. Proceedings of the IEEE International Conference on Evolutionary Computation, 69–73. Piscataway, NJ: IEEE Press, 1998.