Multi-Objective Evolutionary Guidance for Swarms

Evan J. Hughes
Department of Aerospace, Power and Sensors,
Cranfield University,
Royal Military College of Science,
Shrivenham, Swindon, England. SN6 8LA
ejhughes@iee.org

Abstract- A novel approach for generating optimal flight paths for a swarm of missiles using an on-line, noise-aware, multi-objective evolutionary algorithm is introduced. The trajectory shape is controlled by flying via intermediate points which are adjusted dynamically by the evolutionary algorithm, then towards the point of impact with the target.

I. Introduction

High value and high threat targets are often defended well or difficult to intercept. Firing a salvo of missiles can improve kill probability, but each of the missiles will follow the same trajectory. Alternatively, by adding communications between the missiles, the missiles could act as an intelligent salvo, sharing data to improve countermeasure rejection and to improve target parameter estimates, or better still could co-ordinate their flight paths in order to create a swarm. The swarm concept is defined as having the guided platforms perform a coordinated attack, but not necessarily flying in a formation of any form.

In conventional guidance algorithms, data from the sensor is fed directly to the guidance algorithm, generating lateral acceleration demands which are fed to the platform autopilot, which moves the control surfaces accordingly. Thus platforms flown as a salvo will need to have different guidance algorithms if different flight profiles are required [1].

In evolutionary guidance, the platform is first flown via a sequence of one or more points in space, before flying towards a predicted impact with the target, or a required rendevouz point. The points in space are evolved to generate a flight profile that is an optimal solution to a set of objectives and constraints. With multiple platforms, the flight profiles can be evolved simultaneously, each flight profile being evolved while accounting for the intended flight paths of the other missiles.

With the evolutionary guidance approach, for most of the engagement there is no direct, deterministic path between the sensors and the autopilot. Thus the initial stages of the flight path can be independent of the target position and motion, allowing different trajectories to be generated easily. In this application, the scenario will be changing as the platforms fly, therefore changing the characteristics of the objective functions. With a platform with a highly non-linear flight characteristic, there are likely to be many sub-optimal flight paths. Evolutionary algorithms are global optimisation techniques and are robust to noise in the objective functions and have been shown to be well suited to the single missile path planning problem [2].

This problem is very different to that solved by conventional evolutionary algorithms. The objectives are noisy and so repeated evaluations of the same chromosome will give different answers. The objectives are dynamic and so the optima move slowly with time. The optima may also move rapidly (appear discontinuous) if the target manoeuvres. There are multiple objectives to satisfy, leading to a Pareto set of solutions. When combined with noise, the Pareto set is no longer crisp and solutions now have probabilities of domination associated with them, blurring the Pareto surface and preventing a 'perfect' decision. As a complete missile flight is less than 20 seconds, a decision process is needed to choose a single operating point from the Pareto set.

The noise-aware multi-objective evolutionary algorithm [3] is used to maintain a noisy Pareto optimal set of flight profile sets, based on the primary objectives of the engagement. For this paper two primary objectives are used:

- To maximise the smallest latax capability at impact of all the missiles.
- To maximise the smallest difference between the missiles' impact angles.

The Pareto surface is maintained to allow the guidance algorithm to respond rapidly to changes in the scenario. Minor objectives and constraints are then used to aid the decision making process where one solution from the Pareto surface is used to generate the guidance information for the platforms.

The following sections first describe typical guidance heuristics, then the evolutionary algorithm and decision making process. The missile model used in the examples is detailed and example results presented. Finally the paper concludes and indicates areas of future research.

II. Guidance Heuristics

Proportional Navigation (PN) [4] has been used for many years and is well proven as a guidance algorithm. The main essence of the technique is to form a collision triangle, based on target position and velocity, and use it to estimate an impact point where the missile could first intercept the target. The missile then flies towards the impact point, rather than at the target and uses seeker angle rate to derive the lateral acceleration required to correct the position of the estimated impact point.

The impact point and required flight direction are implicit within the PN formulation and are not calculated explicitly.

The process of estimating an impact point can be generalised to any predicted target manoeuvre, where the minimum time trajectory for the missile to fly is a straight line to the impact point. Proportional navigation has been enhanced to achieve this, such as Augmented Proportional Navigation (APN) [4] where target lateral acceleration is taken into account.

For a generalised minimum time guidance heuristic, once the impact point is estimated, the lateral acceleration required to steer the missile towards the impact must be calculated. This is achieved by calculating the angle between the current flight direction of the missile and the direction towards the estimated impact point. The missile must turn through this angle in the shortest time possible, given the current maximum lateral acceleration of the missile. Thus the angular rate required may be established and along with the relationship $a = v\omega_{max}$, where v is the forward velocity and ω_{max} is the maximum angular rate to be applied for the shortest time, the lateral acceleration and the duration of the acceleration event may be generated.

A sub-optimal approach where a lower acceleration is applied for a correspondingly longer time may also be used. This is closer to the operation of traditional PN based techniques. This sub-optimal technique is useful in missiles with high levels of sensor noise and lift-drag coupling, where many large course corrections can have an adverse effect on missile velocity.

$$
\begin{aligned}
\vec{x} &= (T_0 - M_0) \\
M_\tau &= T_0 + \tau\vec{V}_t - \frac{\vec{x}(t-\tau)}{t} \\
\text{at} \quad \tau &= 1, \\
V_m &= |M_1 - M_0| \\
M_1 &= T_0 + \vec{V}_t - \frac{\vec{x}(t-1)}{t} \\
\therefore M_1 &= \vec{V}_t + M_0 + \frac{\vec{x}}{t} \\
\therefore V_m &= |\vec{V}_t + \frac{\vec{x}}{t}| \\
\therefore \quad 0 &= \frac{\vec{x}\cdot\vec{x}}{t^2} + 2\vec{V}_t \cdot \frac{\vec{x}}{t} + (\vec{V}_t \cdot \vec{V}_t - V_m^2) \\
\therefore \quad \frac{1}{t} &= \frac{-\vec{V}_t \cdot \vec{x} \pm \sqrt{(\vec{V}_t \cdot \vec{x})^2 - (\vec{V}_t \cdot \vec{V}_t - V_m^2)\vec{x}\cdot\vec{x}}}{\vec{x}\cdot\vec{x}} \qquad (1)
\end{aligned}
$$

Figure 1 shows the engagement geometry for a constant velocity target. The missile has a current velocity $\vec{V}'_m$ and must pull lateral acceleration in order to obtain a velocity $\vec{V}_m$. The missile must turn through angle θ, and forms the basis of the guidance heuristic described above. The projected point of impact, P_I, is calculated using (1), where t is the predicted impact time, τ is an arbitrary time during the engagement and V_m is the known scalar speed of the missile.

As we always want the shortest impact time, $1/t$ must be as large as possible, therefore (1) may be modified and the impact time may be calculated from (2).

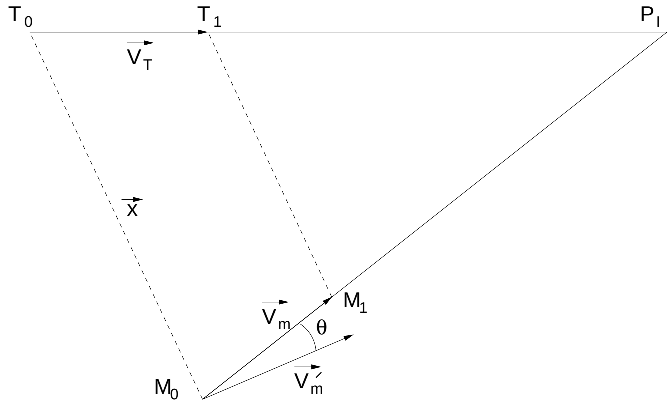

Figure 1: Engagement geometry

$$
\begin{aligned}
q &= (\vec{V}_t \cdot \vec{x})^2 - (\vec{V}_t \cdot \vec{V}_t - V_m^2)\vec{x}\cdot\vec{x} \\
t &= \frac{\vec{x}\cdot\vec{x}}{\sqrt{q} - \vec{V}_t \cdot \vec{x}} \qquad (2)
\end{aligned}
$$

If q or t are negative, the missile is flying too slow and will never catch the target.

Equations to fly at the target in a straight line can also be developed for manoeuvring targets, for example if the target is performing a constant acceleration turn. To use these strategies effectively, good estimates of target lateral acceleration are required.

III. Evolutionary Algorithm

The engagements described in this paper are typically less than 15 seconds in duration. The guidance algorithm may be updated between 10 and 50 times per second. For conventional evolutionary algorithms, the complete run of many generations must be completed within at best 0.1 seconds. An alternative approach is to perform one or two generations for each guidance update, and use the evolutionary algorithm to optimise the resultant noisy and dynamic objective function.

In a command-guided swarm, a potential solution would be a matrix containing the intermediate aim point vectors for each of the missiles. In this paper, a multi-objective evolutionary strategy was used [5] to evolve good sets of aim points. The initial population of solutions is usually generated entirely at random, within some bounds to ensure that most (but not necessarily all) of the solutions are feasible to fly. If heuristics exist to find good sets of aim points, these could be used to help generate the initial population. In this paper, a single missile was flown at the target without any fixed aim points. This provides a *default* trajectory. The maximum and minimum extents of the default trajectory on each axis are then extended by 20% in all directions, resulting in a bounding box for the locations of the evolved fixed points. The initial points were then generated randomly with a uniform distribution to lie within the bounds.

All the sets of solutions are evaluated by simulating the missiles engaging the target and establishing the overall performance of the trial solution. A set of objectives and constraints are used to guide the optimisation process; for example typical objectives may be to minimise the longest flight time, minimise the difference between the longest and shortest flight (simultaneous time-on-target), maximise the smallest difference between the impact angles etc. If a chosen intermediate aim point causes a missile to miss the target in the simulation, the guidance for the particular missile can default to fly straight towards the closest impact point. In order to allow the results to be displayed clearly in this paper, only two objectives will be considered.

These results will form a *Pareto Optimal Set* [5] where no single solution is better than any of the others when all objectives are taken into account. For example, trajectories with long flight times are undesirable but can achieve a very wide spread of impact angles, while short flight time trajectories will have a small range of impact angles but will have higher impact velocities. For the results shown in section VI., a base population of $M = 50$ trial solutions was used. At each generation, $N = 150$ new trial solutions were created, and all 200 evaluated. The best 50 were then chosen for the next generation. The approach of only retaining a quarter of the solutions helps the algorithm to adapt to rapid changes in the objective functions.

Aim point selection is described in section IV. and allows a single solution from the population to be used for targeting. If a missile passes an aim point then it switches out the evolutionary algorithm and heads for the point of impact with the target. In the simulation, when one missile comes within one kilometre of the target, all the missiles change if necessary to aim at the target and the guidance heuristic, along with any fused target position information, is used to guide the missile.

IV. Multi-Objective Selection and Decision Making Process

With any multi-objective optimisation, the result is not a single solution, but a family of equally attractive solutions. However, for the guidance problem only a single solution must be chosen to provide the aim-points for the missiles. The selection of the aim points is made by combining the primary objectives to form a noisy Pareto surface, and then using secondary objectives and constraints to choose a single set of aim points, which corresponds to a single chromosome, from the resultant noisy Pareto set.

The primary multiple objectives are handled using the Multi-objective Probabilistic Selection Evolutionary Algorithm [3]. Here the multiple objectives are selected, based on the *probability* of non-domination, given that the objectives are noisy. Two objectives were used to form the Pareto front: to maximise the smallest difference between the impact angles, and to maximise the smallest latax capability at impact. These two objectives are often mutually exclusive as manoeuvres must be performed to create a wide set of impact angles. This will slow the platform down, leading to a reduced latax capability.

The selection algorithm was made aware of a small amount of noise on the objectives [3]: $\sigma_n = 1.0$ for the smallest angle objective (in degrees), and $\sigma_n = 0.25$ for the minimum latax (in m/s^2). These values have not been tuned in any way and were used to prevent the Pareto surface from collapsing into a crisp line which would not be representative of the noisy surface where domination can only be described probabilistically.

Next the following rules are used to help choose the single operating point out of the noisy Pareto set:

- Remove any solutions whose worst terminal latax is less than 10g. If all the solutions are worse than 10g, take the solution with the highest terminal latax as being the operating point.

- If short or long engagement, prefer solutions with high terminal latax, otherwise take widest spread of angles.

- If point is not restricted by item (1), trade selection with respect to distance to previous aim point to reduce missile steering demands.

The constraint in item (1) is simple to apply and often leads to the solution being chosen for the first half second after launch when the missile terminal velocity is difficult to predict. When the missile boost phase is over (about 3 seconds), the velocity prediction is often stable enough so most solutions satisfy the constraint.

The estimated length of the engagement, t_e, is used to select the aim point in item (2). Equation (3) is used to derive a value k_e that lies in the interval [0,1]. When $k_e = 0$ the engagement is either short or long and so the platform tries to optimise the terminal latax. When $k_e = 1$, the range of impact angles is to be maximised.

$$k_e = \exp\left(-\left(\frac{t_e - 12}{6}\right)^8\right) \tag{3}$$

Ideally, it would be best to fly directly at a stable aim point, and then towards the target. This would minimise the amount of manoeuvre the platform must perform and therefore minimise the drag, maximising the terminal kinetic energy and latax capability. Therefore, the distance, d_i, from the i^{th} aim point to the aim point used in the last time interval can be calculated. The best solution should be chosen to minimise d_i.

If the target manoeuvres, then the last aim point may not be optimal anymore. Minimising d_i is no longer necessary and so the aim point can be chosen solely on the main objectives. This allows for a fast response once a manoeuvre has been detected. A manoeuvre detection signal was generated in the simulation using (4), where μ is the manoeuvre detect signal lying in the range $[0, 1]$ with $\mu = 1$ signifying no manoeuvre is expected, and I_k and I_{k-1} are the current and previous impact point predictions.

$$\mu = exp\left(-\frac{|I_k - I_{k-1}|^2}{200^2}\right) \quad (4)$$

Deciding on a single chromosome means resolving the primary objectives of *maximise terminal latax* and *maximise range of impact angles*, subject to the minor objective of *minimise motion of aim point subject to target manoeuvre*. A simple multi-objective combination process can be used, given that the primary objectives and therefore chromosomes of interest are now a noisy Pareto set. The three objectives are calculated for each individual i and three sets of rank orders R_{latax}, R_{angle}, R_{dist} are generated for the three objectives, with the best objective value being rank 1. Three weighted rank values for each individual can then be generated using (5) and the worst taken to represent the score of the individual, F_i. The solution with the lowest F_i is then taken as the best. This non-parametric weighted min-max optimisation is tolerant to the Pareto surface being concave.

$$F_i = \max\left(\frac{1-k_e}{1+\mu}R_{\text{latax}}, \frac{k_e}{1+\mu}R_{\text{angle}}, \frac{\mu}{1+\mu}R_{\text{dist}}\right) \quad (5)$$

The best performing solution is selected and used to supply the aim points for the platforms. This approach is a little crude, as it can be seen in the results that the population forms into clusters around points in space that represent good solutions. By always taking the best solution, the aim point can wander within the cluster, causing small latax demands to be generated. A clustering algorithm could be used to identify the centre of the best performing cluster of individuals, therefore making the aim point more stable.

V. Missile Model

The model is based loosely on a ship-launched boost-glide missile. For simplicity and to increase processing speed, simple single step integration with a coarse 0.1 second interval was used for the main missile models, and a 0.3 second interval for the flight path projection simulations. This approach meant the integration was too coarse to allow the missile control system to be modelled, so the performance characteristics were generated by modelling the limits imposed by the body aerodynamics. The missile is restricted to a two dimensional engagement in the vertical plane and is subject to the initial boost force, changing mass, gravity, forward drag, lift-drag coupling, and changes of air density and speed of sound with altitude.

Equation 6 details the calculation for the lateral acceleration demand of the missile. The equation calculates the acceleration needed to steer towards the impact point, and also calculates the acceleration needed to correct for the effects of gravity. The angle to steer, θ, is as defined in figure 1; δt is the time step of the model; $|\vec{V_m}|$ is the missile speed; $\hat{\vec{V_n}}$ is the unit vector in the direction that the lateral acceleration must be applied; and k is a constant that acts to damp the response of the missile. In the trials shown in this paper, a value of $k = 5$ was used to help prevent many rapid course corrections causing excessive drag on the missile.

$$l_d = \frac{|\vec{V_m}|\theta}{k\delta t} + \hat{\vec{V_n}} \cdot [0,\ 9.8] \quad (6)$$

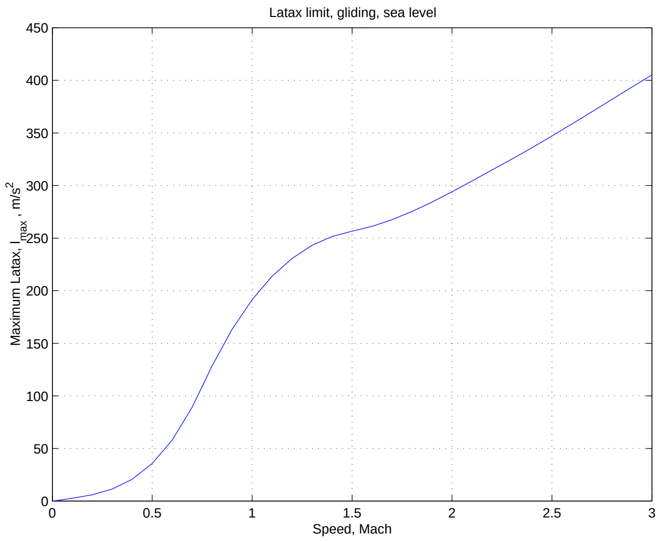

Figure 2: Curve of maximum latax with respect to forward speed

The lateral acceleration demand, l_d, may not be achievable though and so must be limited by calculating the maximum possible demand for the given conditions as shown in (7).

$$l'_d = \begin{cases} \min(|l_d|, l_{max}) & l_d \geq 0, \\ -\min(|l_d|, l_{max}) & l_d < 0, \end{cases} \quad (7)$$

Figure 2 shows the non-linear relationship used to determine the maximum latax demand, l_{max}, for a given speed. If the required latax exceeds the maximum value, the demand is cropped. The graph shows the performance for the missile in the glide phase where the mass is at the minimum value and the missile at sea level.

Equations 8, 9 &10 give the approximations for the speed-of-sound, V_s, air density, ρ, and mass m. Speed of sound and air density vary with respect to altitude, h.

$$V_s = 340.3 - 0.0041\,h \quad (8)$$

$$\rho = 1.375 \exp\left(\frac{-h}{11000}\right) - 0.150 \quad (9)$$

$$m = 34(\tanh(1.5 - t) + 1) + 75 \quad (10)$$

The forward acceleration, a, is calculated using (11), boost force using (12), change in forward velocity using (13), and change in position with (14).

$$a = \frac{f - \frac{1}{2}\rho|\vec{V_m}|^2 AC_d}{m} - \left|\frac{l'_d}{C_{ld}}\right| \quad (11)$$

$$f = 17000(\tanh(5-2t)+1) \quad (12)$$

$$\delta\vec{V_m} = \delta t\left(a\hat{\vec{V_m}} + [0, -0.98] + \hat{\vec{V_n}} l'_d\right) \quad (13)$$

$$\delta P = \delta t\vec{V_m} \quad (14)$$

A missile cross-sectional area of $A = 0.0254$, mean drag coefficient of $C_d = 0.45$ and mean lift-drag ratio of $C_{ld} = 3.5$ were used in the simulations. The actual values for C_d and C_{ld} were different for each missile by up to $\pm 10\%$.

VI. Results

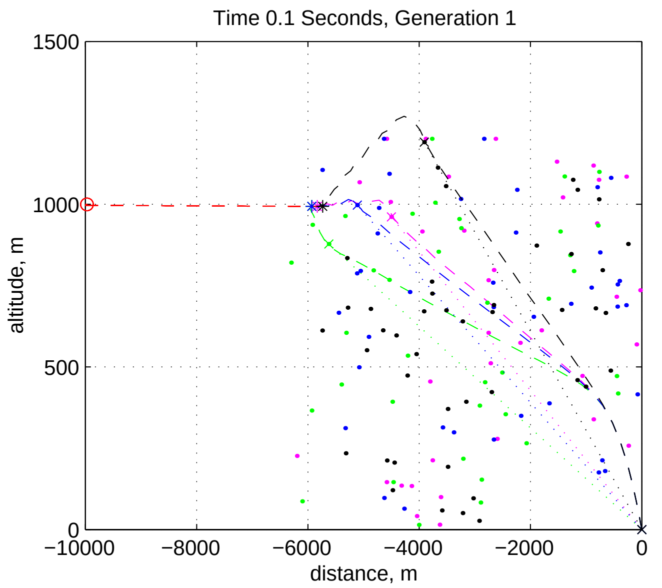

Figure 3: Initial predicted trajectories showing location of fixed points in evolutionary algorithm population at generation 1

A sample swarm with four missiles, launched from coordinate [0, 0], was simulated engaging a target flying with constant velocity at Mach 1. The missile sensor system was simulated by corrupting the exact target position and velocity with noise, causing the estimated impact point to wander. An evolutionary algorithm with a working population of 50 trial aim-points was used, with one generation being simulated as 0.1 seconds.

Figures 3 & 4 show the initial population and initial Pareto set at instant of launch. The missile velocity is still zero at this instant, so a number of generations can be executed before a significant speed and change in position has been attained, allowing the algorithm to begin to converge. By 2 seconds (20 generations) clusters of aim-points are forming and the Pareto set is stabilising, yet the missile has only travelled 500 metres.

The missile positions are marked with crosses, the proposed fixed points are dots, the target position is a circle, the predicted trajectories are dashed lines, and the predicted impact points are stars. On the Pareto set, the chosen operating point is ringed.

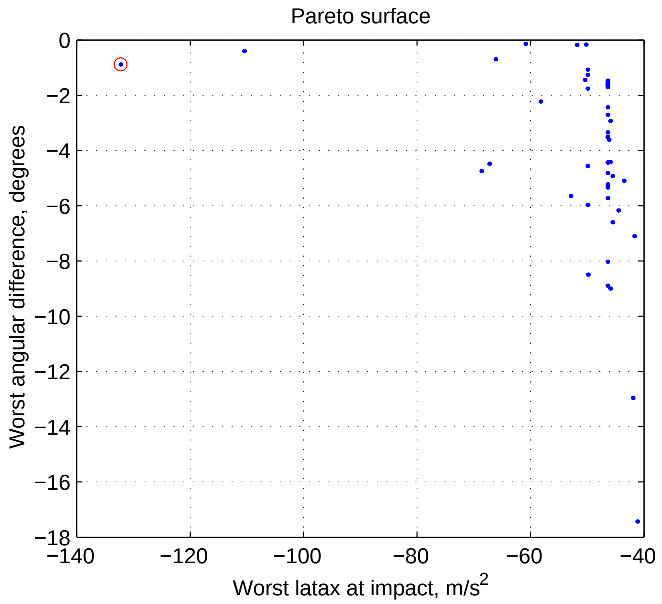

Figure 4: Initial Pareto set at generation 1

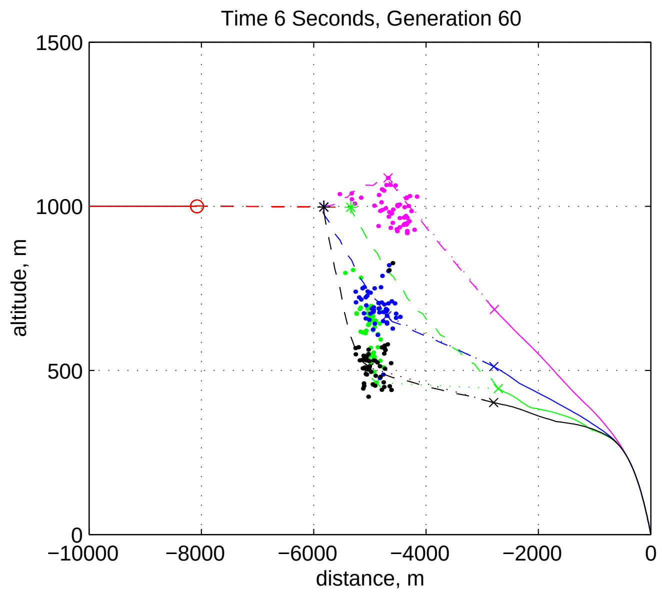

Figure 5: Partial and predicted trajectories showing location of fixed points in evolutionary algorithm population at generation 60

Figures 5 & 6 show the state at 6 seconds (60 generations). Here the aim points are forming tight clusters and the trajectories are forming well. The chosen operating point appears to be good on angle, but poor on impact velocity. The particular point has been chosen as the range is not short or long and it is close to the last operating point. It must be remembered that as the objectives are noisy, the Pareto surface is no longer crisp with all points being 100% non-dominated. The noisy Pareto optimisation ensures that the point has a good probability of being non-dominated.

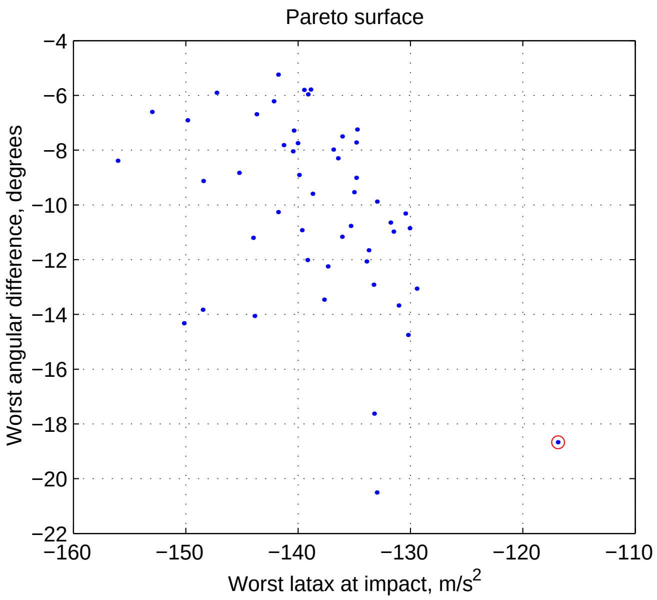

Figure 6: Pareto Front at generation 60

The first impact occurs at 13.0 seconds, and is shown in figure 7. It can be seen that the missiles each took a different path to the target, with quite a good spread of impact angles. In the early phase of the engagement, all the paths are the same as the missile manoeuvre is restricted by the maximum lateral acceleration limit of the platforms. As the process is stochastic, each run of the algorithm will produce different flight paths.

VII. Conclusions

The multi-objective evolutionary guidance approach provides a comprehensive framework allowing multiple missiles to coordinate attacks on single or multiple targets. The framework also allows for data from multiple sensors to be fused easily, as the guidance requires estimates of missile and target positions etc. in absolute coordinates to be used. The results were generated using a highly non-linear missile, combined with noisy measurements and uncertain system models. This shows clearly that the method is tolerant of complexity and many sources of error. The method also demonstrates how a simple decision making process can be used to select a single point from a Pareto set in a dynamic system.

The technique does however require significant processing resources. On todays fast machines, evolutionary guidance as demonstrated will be just realisable in real time, albeit with relatively small population sizes. The larger the population size used, and the more accurate the missile simulations, the better the guidance will perform.

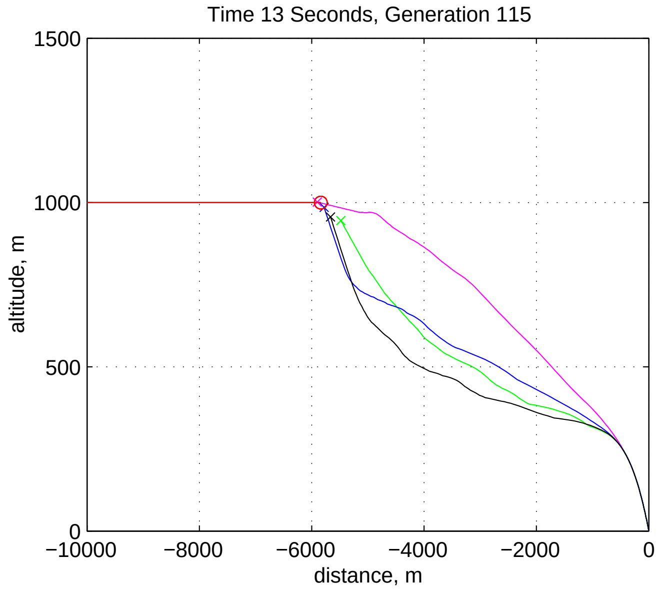

Figure 7: Final trajectories, first missile impact at 13.0 seconds

Acknowledgements

The author would like to acknowledge the use of the Department of Aerospace, Power, and Sensors DEC Alpha Beowulf cluster for this research.

References

[1] P. A. Creaser and B. A. Stacey, “Generation of dual missile strategies using genetic algorithms,” in *CEC99 conference*, vol. 1, (Washington, DC), pp. 117–124, IEEE, 6–9 July 1999.

[2] L. Quing, G. Wei, L. Yuping, and S. Chunlin, “Aircraft route optimization using genetic algorithms,” in *GALESIA ’97 Conference*, pp. 394–397, IEE, 2–4 September 1997. Conf. Pub. No. 446.

[3] E. J. Hughes, “Constraint handling with uncertain and noisy multi-objective evolution,” in *Congress on Evolutionary Computation*, vol. 2, pp. 963–970, IEEE, May 27–30 2001.

[4] P. Zarchan, *Tactical and Strategic Missile Guidance*. American Institute of Aeronautics and Astronautics Inc., 3rd ed., 1997. ISBN 1-56347-254-6.

[5] K. Deb, *Multi-objective optimization using evolutionary algorithms*. John Wiley & Sons, 2001.