

An Intelligent, Heuristic Path Planner for Multiple Agent Unmanned Air Systems

C. Crispin* and A. Sóbester†

University of Southampton, Southampton, England.

In situ observations of atmospheric variables are currently obtained with radiosondes, collecting data along uncontrolled trajectories. As an alternative, we propose an unmanned air system comprising a swarm of unmanned aerial vehicles, released from high altitude weather balloons. Their trajectories are optimised for efficient sampling, with an objective function measuring the space-filling properties of the entire swarm. The dynamics of the aircraft swarm are captured in a number of primitive manoeuvres simulated with JSBSim. A combination of these paths forms a set of moves available to a greedy heuristic algorithm, which determines which of the flyable paths is optimal according to the objective function. At each heuristic optimisation step a decision is made upon the next single move and, once the move is complete, the heuristic repeated, resulting in a stitching together of optimal moves from the flyable set. The proposed path planner comprises of a centralised algorithm, which executes offline, thus, each aircraft executes the heuristic synchronously and the result is a cloud of waypoints to be flown by the swarm. A case study based on orographic flows over South Georgia is used to test the performance of the algorithm in windy environments. The results indicate good performance of the algorithm, even in high, unsteady wind fields.

I. Motivation

Most commonly, balloon and aircraft launched radiosondes are utilised to collect physical, chemical and biological measurements in the atmosphere. These are predominantly used to assist the modelling of atmospheric processes from large global systems to small localised phenomena. Although relatively easy to deploy, balloon and aircraft launched radiosondes have disadvantages. The sonde is usually lost, which incurs an expense and only data transmitted during the flight can be recovered. The most notable disadvantage is the sparsity of the data collected. The sondes collect their information on an uncontrolled trajectory, determined by the flight of the (on the way up) balloon or the parachute (on the way down), which provides for limited control of the sampling locations.¹

Autonomous observation platforms are gaining popularity in the scientific community. Systems previously operated by humans in the air, under water, in space or on land are slowly being replaced by systems with increasing levels of autonomy. This is particularly true of tasks associated with danger to a human operator, such as military operations or systems that involve the handling of hazardous materials, and monotonous tasks such as reconnaissance or surveillance.² Of most interest in this investigation is the automation of environmental monitoring or sampling. An initiative named the Maldives Autonomous Unmanned Aerial Vehicle Campaign (MAC)^{3,4} uses multiple unmanned aerial vehicles for the measurement of aerosol, cloud and solar radiation. The National Institute of Polar Research (NIPR) in Japan has a long running project, which successfully utilises numerous small autonomous unmanned aerial vehicles (UAVs) to perform Antarctic aeromagnetic and aerial photographic surveys.^{5,6}

*Post Graduate Researcher, Institute for Complex Systems Simulation, Computational Engineering Design

†Associate Professor, Computational Engineering Design, AIAA Senior Member.

II. Designing an Unmanned Air System for Atmospheric Research

The approach considered here, closely linked with the Massive Airspace Volume Instrumentation Systems (MAVIS) research project, aims to optimise the trajectories of multiple unmanned aerial vehicles (UAVs), components of a complete unmanned air system (UAS) for atmospheric research. The work presented here develops a centralised, offline path planner to be executed before flight to produce an achievable set of waypoints.

A. Objectives and Evaluating Fitness

The optimality of a trajectory requires multiple objectives to define, including the *space-filling* qualities of the waypoint cloud, the accuracy of the concentration mapping, if applicable, and, in complex environments, will consider the affect of wind and terrain.




Figure 1. Multiple trajectories traversing a domain represented by a dense waypoint cloud

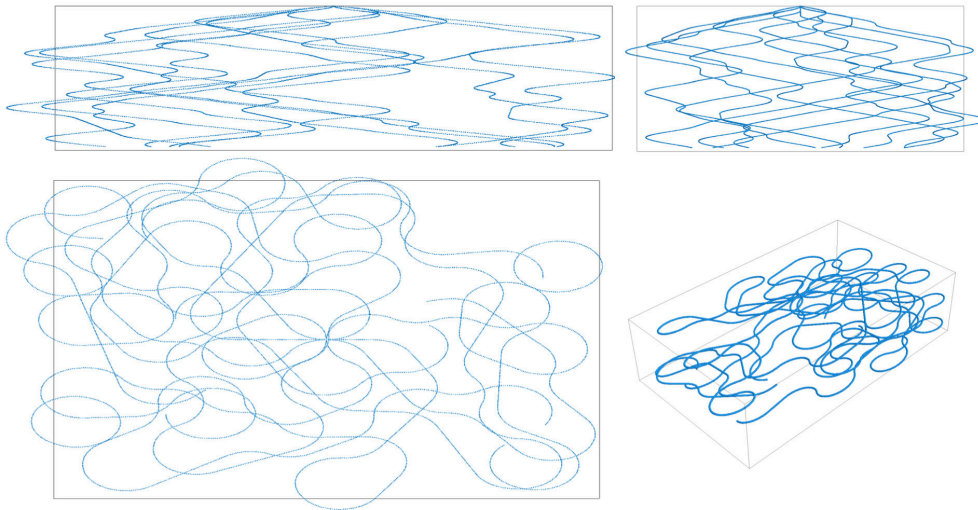


Figure 2. An orthographic projection of multiple trajectories traversing a domain

Figure 1 displays ten simultaneous trajectories exploring a three dimensional domain and the respective waypoint cloud on a zoomed in section, the resolution of which depends on the chosen time step between each successive point. In this case we are only concerned with the space-filling properties of the combined trajectories in three dimensional space.

Figure 2 projects the ten trajectories of Figure 1 onto combinations of two dimensions, enabling easier visual inspection of the associated waypoint cloud and its space-filling quality. Although it is obvious from inspection that some regions are less searched than others, it is impossible to visually evaluate a waypoint cloud consistently and accurately. Consider the two, equal length, two dimensional trajectories in Figure 3. With simple statistical measures indicative of space-filling, such as variance, we can suggest that these are two equally space-filling curves. For example the uncoupled variances in x and y positions for 36 discretely sampled coordinate points (at the centre of each grid square) are equal for each curve. However, if more samples are introduced along the curve, the uncoupled variances of the spiral trajectory coordinates become superior. If asked to visually determine which is more space-filling it would be difficult to choose, moreover, it would be impossible to give a quantifiable reason why one is better than the other. Extend this to a third dimension and visually assessing space-filling is compounded, necessitating a method for calculating a sophisticated and reliable scalar describing the *fitness* of a waypoint cloud.




Figure 3. Traversing a two dimensional domain and quantifying space-filling

The objective function captures the space-filling properties of a trajectory. However, applications may also require that the UAVs collect atmospheric measurements, which vary in concentration in a volume of airspace. In these situations the UAVs favour areas of high concentration, as these are of greater interest.

Here we propose a method whereby the volume of airspace to be observed is discretised to form a three dimensional matrix, where each point in the matrix contains a scalar describing how uncertain the system is about the environment at that point⁷ and how interesting that point is, based on concentrations of an measurable atmospheric quantity.

Since no information about the environment is known preflight, this matrix is initialised uniformly. At each point along a trajectory, all points in the matrix within a sphere of designed radius, are reduced using an exponential function inspired by work done by Stone on target detection.⁸ A superposition of these spheres produces a cylindrical tube, Figure 4, which follows the trajectory and influences any discrete points in the domain within the designed radius. Note that in Figure 4 where the domain is clear the uncertainty is unchanged, thus equal to one, indicating a region of complete uncertainty.

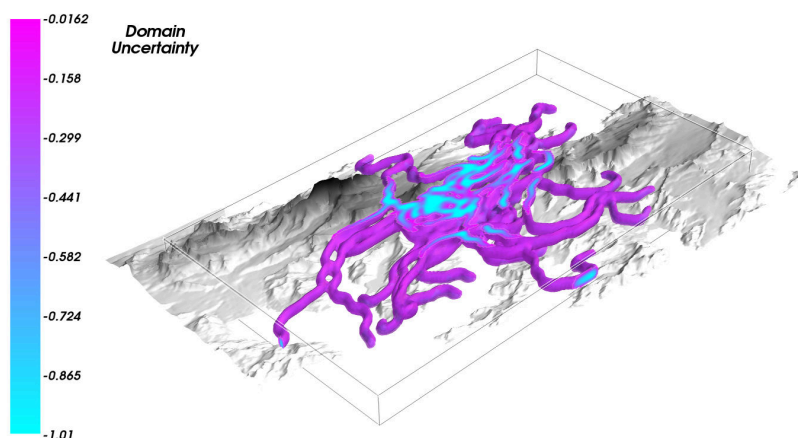


Figure 4. Multiple trajectories carving cylindrical tubes through a uncertainty domain above South Georgia.

The exponential detection function has a diminishing effect the more often the same cells are visited within the three dimensional matrix. The mean reduction across the matrix, thus, equates to the reduction of uncertainty in the three dimensional space. This method, however, is flawed as, although it is often assumed in classical search theory, uncertainty and space-filling cannot always be directly comparable. Consider two trajectories of equal length traversing a discretised two dimensional domain, as in Figure 5.




Figure 5. Traversing a two dimensional grid; reducing uncertainty versus improving space-filling

Each trajectory reduces the uncertainty of the entire domain equally, as each trajectory visits exactly twenty discrete cells, however, one is certainly more space-filling than the other. Measures of uncertainty and space-filling are not always necessarily interchangeable but may be more synonymous when the number or length of trajectories is sufficiently large compared to the number of discrete cells in a domain. Furthermore, when dynamically deciding upon the space-filling properties of a trajectory restricted to the flyable kinematics of an aircraft, the uncertainty of the domain is fairly indicative and, in uncertain environments, the only available information suitable for fast, intelligent decision making.

B. An Optimisation Scheme

The objective function is optimised by a bespoke heuristic algorithm. This route planning problem can be likened to a travelling salesman problem (TSP). There are a number of nodes connected by paths, which have an associated weight determined by a scalar interpolated from the matrix described in Section A. Much like the TSP, this problem is NP-hard, therefore a sub-optimal solution may be the only practical solution.

The greedy heuristic algorithm can decide upon the optimality of a path based on any number of steps ahead in the volume of airspace, however it only ever performs a single move at each evaluation. This heuristic algorithm introduces intelligence into the system, without introducing extreme computational cost, which is particularly important when the algorithm is to be run repeatedly throughout the optimisation process.

The proposed method will be constrained by the kinematics of a specific aircraft. In this case a glider will be modelled, not only inhibiting the sampling to *flyable* points but also restricting the flight or sampling time, which is dependant on the glide ratio of the aircraft and its release altitude or the total energy (potential) available per flight. Other factors include wind field and altitude affects on density, which alter the handling characteristics of an aircraft. Further to this, upon deployment, this swarm of UAVs will have no previous knowledge of the airspace, though there may be information of varying fidelity regarding terrain and typical wind fields.

Computational time is further reduced by pre-computing the flyable paths. Hwangbo et. al⁹ uses the combination of a greedy, coarse global heuristic combined with a much denser local planner to produce an algorithm capable of producing a reasonable suboptimal trajectory for aircraft intended to navigate a number of obstacles. Here the optimal trajectory is one of shortest length that reaches the target area without hitting an obstacle. In order to improve the efficiency of the dense, local planner a look-up table is used to represent the dynamics of an aircraft. For Hwangbo et. al this look-up table contains a precomputed set of feasible motions available to the aircraft, holding the dynamic characteristics of the aircraft at any of the possible states throughout the flight. Stitching any of these possible motions together from each initial and final state will produce more advanced manoeuvres and allows for stepping through a complete trajectory. Motion primitives are less mature in three-dimensions but their use is still common practice in the path planning of aerial vehicles. Pachikara et. al¹⁰ explores the use of sequential motion primitives to produce a high level motion planner for an aircraft, under environmental constraints.

Unlike Hwangbo et al.⁹ who use dynamically defined Dubins curves to generate motion primitives, we propose the use of a flight dynamics model (FDM) alongside a waypoint generator. The JSBSim FDM can be programmed to use any aircraft, provided a configuration file is supplied. It also permits the coding of custom autopilots and scripts, which allows for waypoint navigation and pitch hold.

First, we generate suitable waypoints via a recursive algorithm and generate a JSBSim script for said waypoints at the desired altitude. JSBSim then flies the waypoints to a designed tolerance. If the tolerance is broken the turn radius is increased and the procedure repeated. This is preformed for any number of altitude levels. The more levels are included the more accurately the flight primitives will model density variation with altitude. Figure 6 shows the precomputed waypoints for an altitude of 10000m superimposed with the JSBSim output.

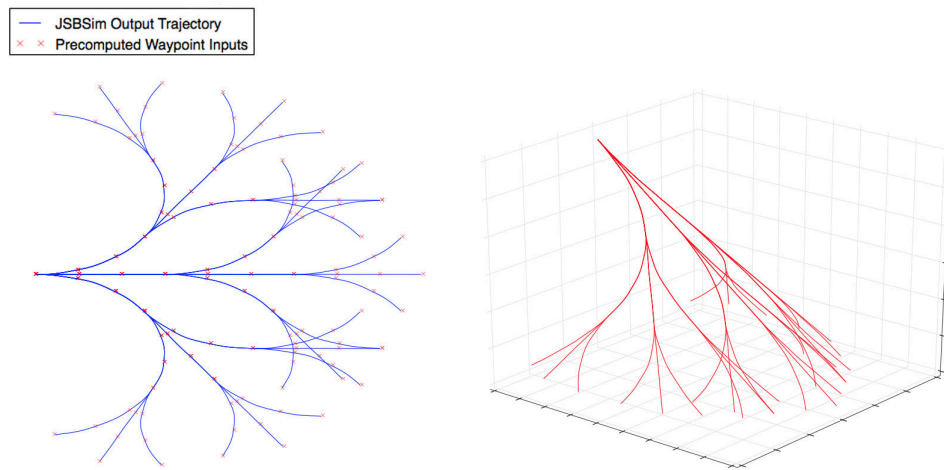


Figure 6. JSBSim output for flight primitives

Each of these primitives is then stored in a look-up table to be used at each step in the heuristic path planner. The result is a stitching together of possible flight paths, show in the example depicted in Figure 7.




Figure 7. Stitching of flight primitives

C. Case Study: South Georgia

A case study, in which the physical atmospheric property to be mapped is wind velocity, is used in order to display the effectiveness of the optimisation scheme. In this case, we are only concerned with the uniformity

of sampling, thus, space-filling is the single objective for optimisation. Wind data is supplied by the British Antarctic Survey (BAS) and the terrain data is sourced from the Shuttle Radar Topography Mission (SRTM).

C.1. The Data

The SRTM supplies 90m by 90m resolution terrain data as a binary file, which is plotted in Figure 8.

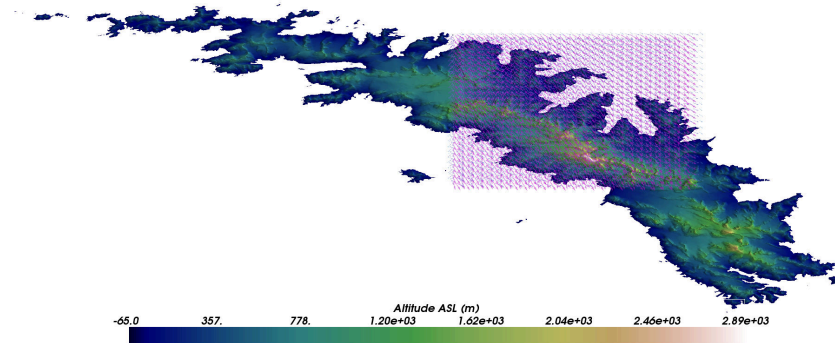


Figure 8. A map of South Georgia depicting altitude above sea level and the region in which high resolution wind data is supplied.

Wind data, supplied by the British Antarctic Survey (BAS), is available in two regions: a high resolution region, 900m by 900m, immediately downstream of the largest peak and a low resolution, 2700m by 2700m, region covering the entire island. The data points are supplied at unevenly spaced altitude levels, for instance there are more levels in the lower troposphere than the upper troposphere, and are devised using a terrain-following hydrostatic pressure vertical coordinate system. Figure 9 shows a single layer of the supplied wind data, exhibiting the effects of the terrain-following coordinate system.

The wind data is the result of the Weather Research and Forecasting (WRF) model, a state-of-the-art weather forecasting tool, which uses a coarse global forecast for boundary conditions. Observations such as those performed by a UAS can be used to validate or improve the WRF model and provide higher resolution data in areas of interest.

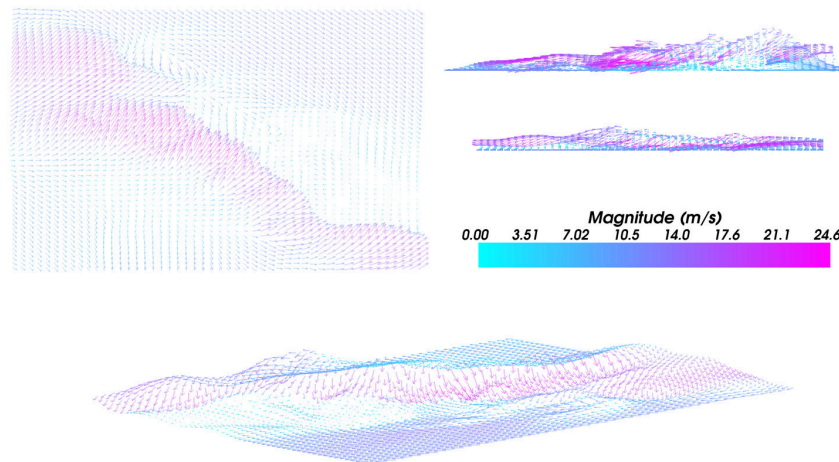


Figure 9. A slice of wind data produced using a terrain-following hydrostatic pressure vertical coordinate system

To allow for simulations in the presence of wind the x , y and z wind component is added to each of the primitives' sampled points, effectively shifting the coordinate frame of reference with the wind. The primitives, thus, stretch and skew with the wind. In high wind, where the wind speed is high relative to the aircraft speed, which is common at altitudes exceeding 5km above South Georgia, the algorithm attempts to keep the trajectories within the volume of airspace to allow for a more optimal trajectory once in areas of reduced wind speed.

III. Algorithm Performance

Figure 10 shows the results for 20 flights in a domain of size 8km by 8km by 6km where wind and terrain are neglected. Each tile displays y slices (first three rows) at 1600m, 3000m, 5800m and 6500m and z slices (bottom three rows) at 1600m, 2900m, 4300m and 5000m containing the original cloud distribution (first row of each block), predicted cloud distribution (second row) and uncertainty domain (third row). The predictor matrix, in both cases, shows good correlation with the original cloud distribution and the uncertainty domain shows the trajectories attempting to cluster near high cloud concentrations, as desired. Further to this, the uncertainty domain also displays a good mixture of space-filling and clustering, giving the a good likelihood that all areas of high pollutant concentration in a region are discovered.

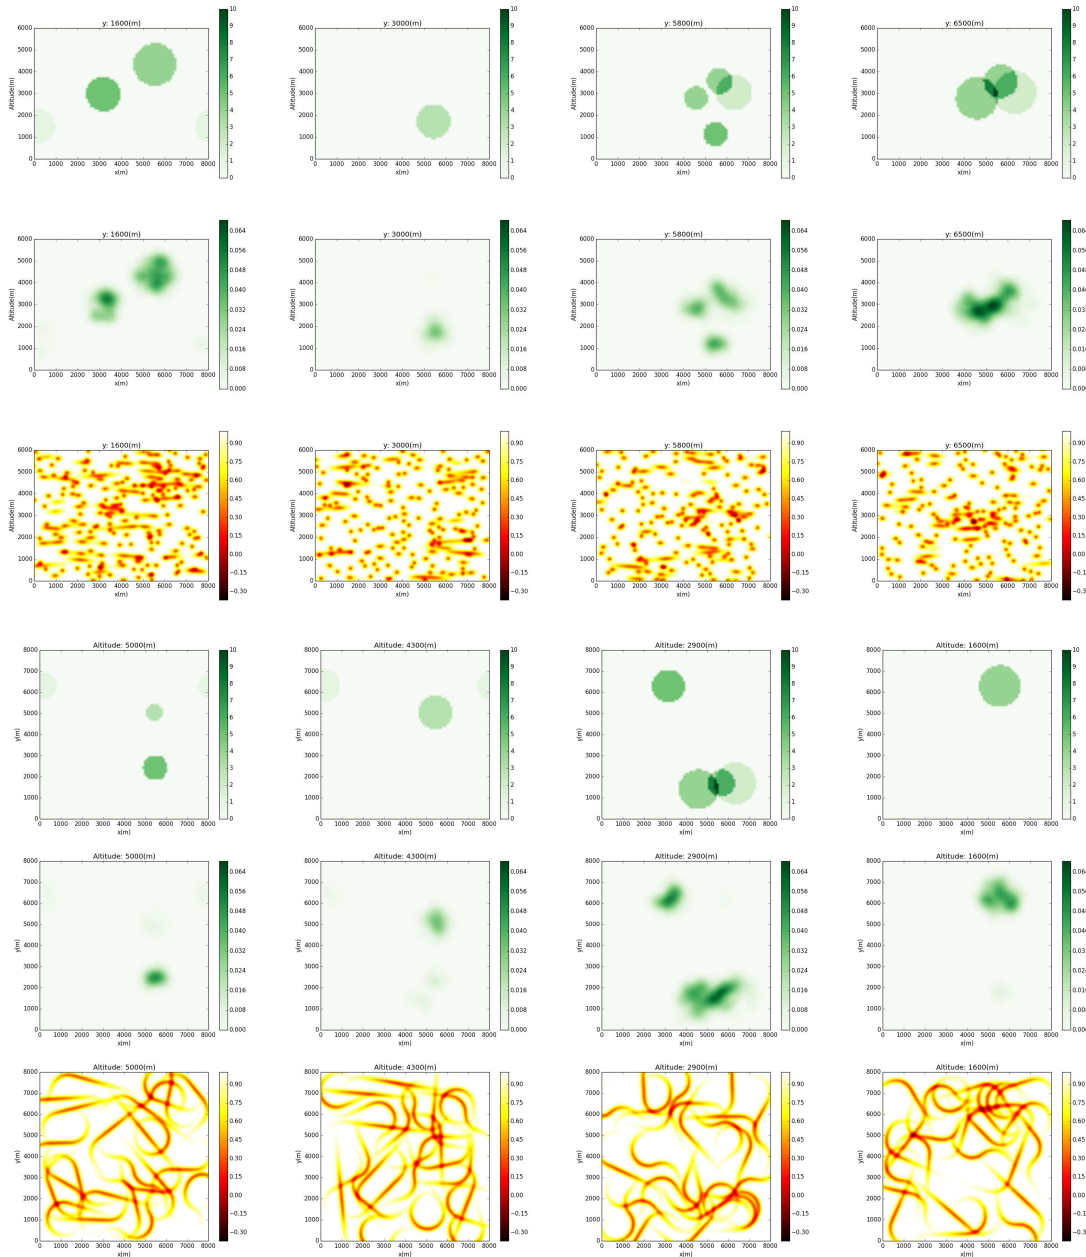


Figure 10. y (top) and z (bottom) slices of the cloud, predictor and sampling domains after 20 flights

To conclude upon the algorithm performance in wind and terrain a number of simulations performed above South Georgia are presented. Here we are only concerned with the *space-filling* properties of the combined trajectories.

The first simulation uses a single step and the second a double step heuristic path planner. Each are otherwise identical in setup with a launch altitude of 5km, a domain size of 10km by 10km, a bounding box of $[-54.4^\circ, -54.3^\circ]$ latitude, $[-36.55^\circ, -36.4^\circ]$ longitude and 40 aircraft.

Figure 11 and Figure 12 display wind profiles for 5km and 1km, respectively, in order to give a reasonable idea of wind direction and magnitude as the flight of the aircraft progresses. To begin with, the aircraft will experience a consistently strong south-easterly wind which transitions to a less consistent north-easterly/easterly wind near landing.

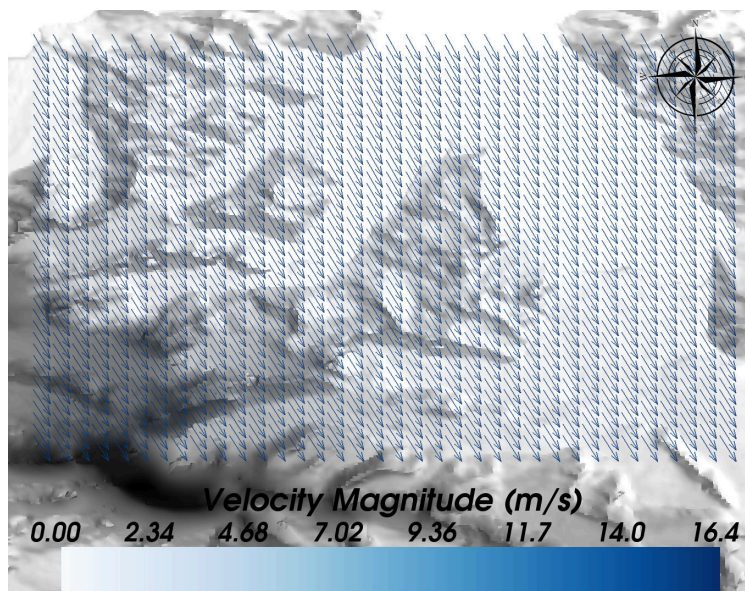


Figure 11. The wind profile at 5km within the bounding box $[-54.4^\circ, -54.3^\circ]$ latitude, $[-36.55^\circ, -36.4^\circ]$ longitude

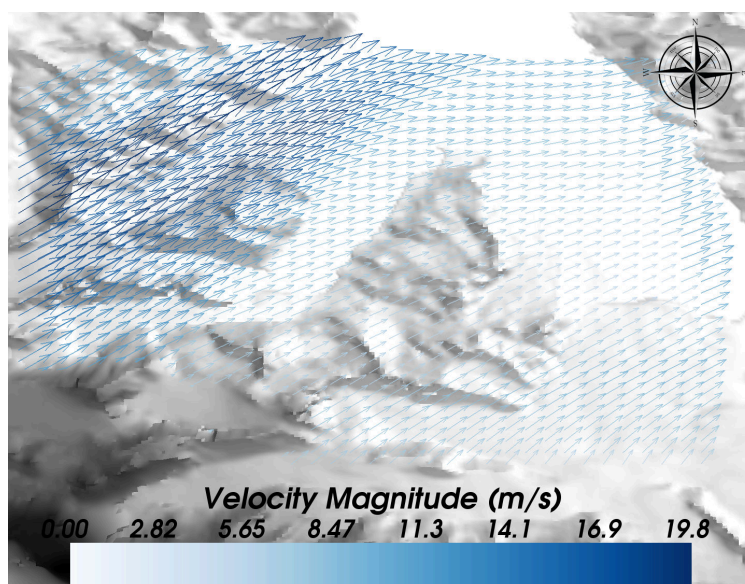


Figure 12. The wind profile at 1km within the bounding box $[-54.4^\circ, -54.3^\circ]$ latitude, $[-36.55^\circ, -36.4^\circ]$ longitude

Figure 13 presents the optimal set of trajectories found using the single step heuristic path planner. The trajectories begin to bunch towards the south-east corner of the domain, owing to the strong south-easterly wind, and do not fully recover before landing, leaving a sizeable portion of the domain unsearched.

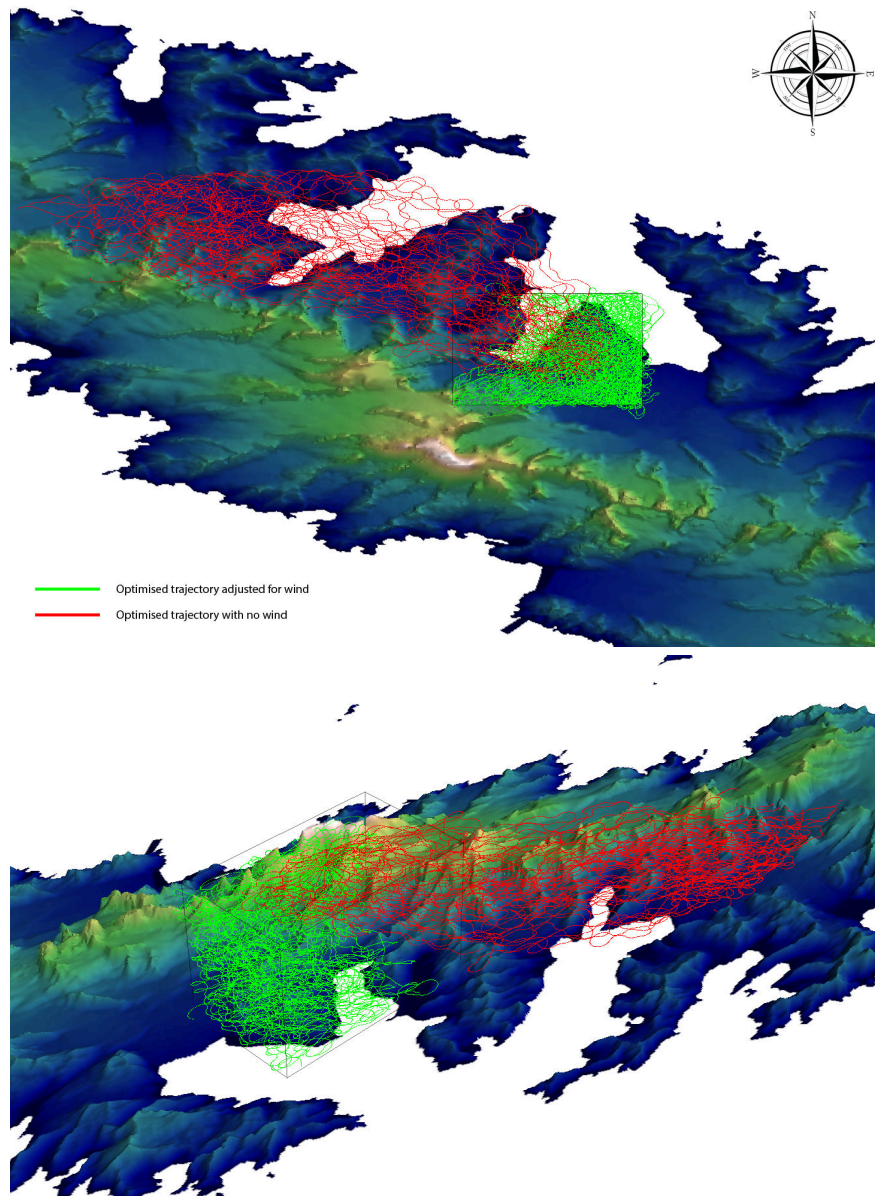


Figure 13. Optimised trajectories for 40 aircraft with a launch altitude of 5km, a domain size of 10km by 10km, a bounding box of $[-54.4^\circ, -54.3^\circ]$ latitude, $[-36.55^\circ, -36.4^\circ]$ longitude and a single step heuristic path planner

Figure 14 presents the optimal set of trajectories found using a double step heuristic path planner. This set of trajectories also appears to bunch due to the prevailing wind, however, unlike that of the single step, the double step algorithm manages to recover somewhat, thus, exploring much more of the domain. This is further seen in the uncertainty domains plotted in Figure 15 and Figure 16, where the resulting uncertainty of all sides of the domain is visualised. Although very similar, it is obvious from the side projections that the double step algorithm performs much better and reaches regions in the domain that the single step algorithm never searches. In spite of the superior performance of the double step, the computational effort required scales with $O(N^s)$, where N is the number of motion primitives and s the number of steps, making it a much more computationally expensive search. It is apparent that the algorithm needs to be developed to reduce the computational time required for multiple steps, enabling its full potential to be realised. This is particularly important when wind is considered as it seems much more likely a trajectory will get stuck in a local optimum when not allowed a reasonable number of steps to discover globally optimal paths.

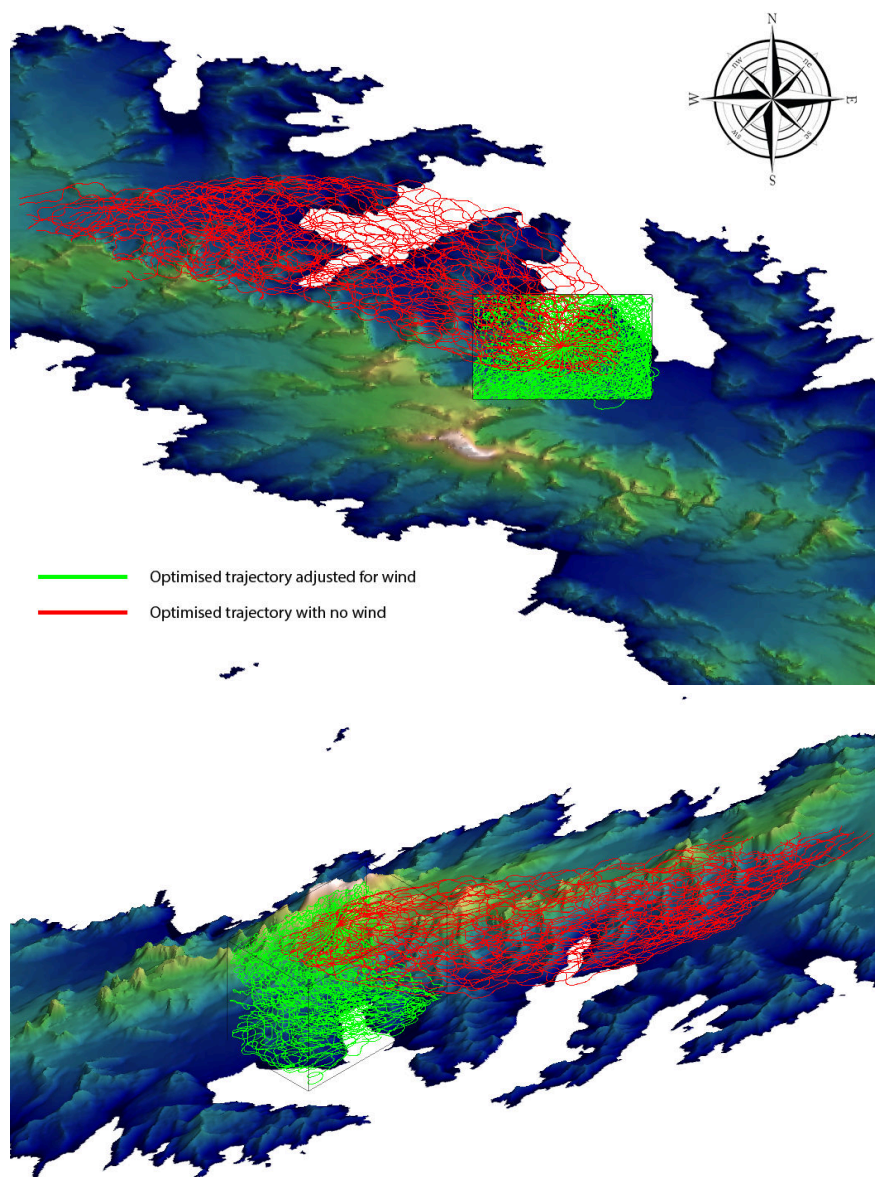


Figure 14. Optimised trajectories for 40 aircraft with a launch altitude of 5km, a domain size of 10km by 10km, a bounding box of $[-54.4^\circ, -54.3^\circ]$ latitude, $[-36.55^\circ, -36.4^\circ]$ longitude and a double step heuristic path planner




Figure 15. The domain uncertainty for 40 aircraft with a launch altitude of 5km, a domain size of 10km by 10km, a bounding box of $[-54.4^\circ, -54.3^\circ]$ latitude, $[-36.55^\circ, -36.4^\circ]$ longitude and a single step heuristic path planner




Figure 16. The domain uncertainty for 40 aircraft with a launch altitude of 5km, a domain size of 10km by 10km, a bounding box of $[-54.4^\circ, -54.3^\circ]$ latitude, $[-36.55^\circ, -36.4^\circ]$ longitude and a double step heuristic path planner

IV. Conclusions

Overall, the algorithm has performed well. It has been successful in producing trajectories that are space-filling, in still and windy environments. Moreover, the algorithm has also demonstrated its ability to detect regions of high concentration pollutant in the atmosphere.

However, the heuristic requires further development to reduce its speed of execution, allowing for a larger number of steps to be searched. In the case presented, the double step algorithm greatly outperforms the single step algorithm, suggesting that exploration is particularly important when considering windy environments. This ensures that aircraft do not get stuck in a local optimum. A trade off between exploration and optimality should be investigated. Perhaps the resolution of the search could be sacrificed to improve the scaling of the algorithm, whilst sacrificing optimality that may be lost due to the reduced search space.

In addition, the decentralisation of the algorithm should be investigated. Although this path planner works well offline, a more sophisticated dynamic algorithm should be employed for online computation. This algorithm should incorporate a model for communication between aircraft and consider the effect of incomplete, distributed information on the optimality of an aircraft's trajectory.

Acknowledgments

We would like to thank the British Antarctic Survey for providing wind data above the island of South Georgia.

References

- ¹Sóbestor, A., Johnston, S., Scanlan, J., O'Brien, N., Hart, E., Crispin, C., and Cox, S., "High Altitude Unmanned Air System for Atmospheric Science Missions." Tech. rep., Faculty of Engineering and the Environment, University of Southampton, 2011.
- ²Blyenburgh, P. v., "UAVs: an overview," *Air and Space Europe*, Vol. 1, 1999, pp. 43 – 47.
- ³Roberts, G. C., Ramana, M. V., Corrigan, C., Kim, D., and Ramanathan, V., "Simultaneous observations of aerosol–cloud–albedo interactions with three stacked unmanned aerial vehicles," *Proceedings of the National Academy of Sciences*, Vol. 105, No. 21, 2008, pp. 7370–7375.
- ⁴Corrigan, C. E., Roberts, G. C., Ramana, M. V., Kim, D., and Ramanathan, V., "Capturing vertical profiles of aerosols and black carbon over the Indian Ocean using autonomous unmanned aerial vehicles," *Atmospheric Chemistry and Physics*, Vol. 8, No. 3, 2008, pp. 737–747.
- ⁵Funaki, M. and Hirasawa, N., "Outline of a small unmanned aerial vehicle (Ant-Plane) designed for Antarctic research," *Polar Science*, Vol. 2, No. 2, 2008, pp. 129 – 142.
- ⁶Funaki, M., Higashino, S.-I., Sakanaka, S., Iwata, N., Nakamura, N., Hirasawa, N., Obara, N., and Kuwabara, M., "Small unmanned aerial vehicles for aeromagnetic surveys and their flights in the South Shetland Islands, Antarctica," *Polar Science*, Vol. 8, No. 4, 2014, pp. 342 – 356.
- ⁷Sujit, P. B. and Ghose, D., "Search using multiple UAVs with flight time constraints," *Aerospace and Electronic Systems, IEEE Transactions on*, Vol. 40, No. 2, 2004, pp. 491–509.
- ⁸Stone, L., *Theory of optimal search*, Computational Methods for Modeling of Nonlinear Systems, Academic Press, 1975.
- ⁹Hwangbo, M., Kuffner, J., and Kanade, T., "Efficient Two-phase 3D Motion Planning for Small Fixed-wing UAVs," *Robotics and Automation, 2007 IEEE International Conference on*, april 2007, pp. 1035 –1041.
- ¹⁰Pachikara, A., Kehoe, J., and Lind, R., "Path-Parameterization Approach Using Trajectory Primitives for Three-Dimensional Motion Planning," *Journal of Aerospace Engineering*, Vol. 26, No. 3, 2013, pp. 571–585.