

Minimum Risk Planning for Teams of Unmanned Air Vehicles

Nathan T. Depenbusch*

Jack W. Langelaan†

The Pennsylvania State University, University Park, PA 16802, USA

A method for task assignment for persistent, low-altitude surveillance by a team of soaring-capable uninhabited air vehicles is discussed. The task assignment approach is based on minimizing the overall risk of mission failure (defined as failure to maintain continuous surveillance), where the risk to the mission of failure to perform a particular task defines the reward function in task assignment. Tasks include target surveillance, energy exploitation, exploration, and transit to target. Methods for assessing the risk of mission failure if each task is not performed are described and the Munkres algorithm is used for task assignment. Monte Carlo simulations show the utility of the proposed approach and the effect of flock size on the likelihood of mission failure is discussed.

I. Introduction

Small uninhabited air vehicles (UAVs, here small means hand-launchable) have some important advantages over their larger brethren. For example, the cost of an individual vehicle is lower and they require less operational support, so riskier missions can be undertaken. However, range and endurance of small UAVs are limited both by the amount of fuel (or batteries) that can be carried and by the reduction in overall aerodynamic efficiency associated with the lower Reynolds Numbers of these aircraft.

However, with the development of planning and control strategies that exploit energy available in the atmosphere, range and endurance can be significantly enhanced. Most of the literature has focused on thermal soaring: Allen¹ showed that soaring flight with small UAVs is practical which was confirmed by demonstration by Allen and Lin.² Strategies for thermalling flight have been examined by Anderson and Kammer.^{3,4} Performance can be further enhanced if aircraft are capable of mapping their environment,⁵ and significant additional improvement is possible if a flock of cooperating soaring-capable UAVs is deployed.⁶ A flock of aircraft is capable of cooperatively mapping the dynamic environment, and this shared map results in significantly improved endurance of the flock.

The problem now is incorporating coordinated soaring flight into the higher level mission goals. This paper considers the problem of persistent surveillance of a stationary ground target by a team of soaring-capable UAVs. At least one aircraft must be within surveillance range of the target at all times; the other aircraft either explore the environment to identify regions where energy can be obtained or actively harvest energy. The result is a problem that combines task assignment, path planning, and control, with the goal of maintaining surveillance of the target for as long as possible.

Typically mission planning problems such as this one would be considered in terms of a multi-objective optimization problem, where some composite cost function is defined in terms of sub-goals and weights that are assigned to each sub-goal. Assigning weights is often done heuristically or with the guidance of simulation results until intuitively reasonable results are obtained. However, “good” values for weights can depend on specific conditions, and in the case of low-altitude flight conditions can be highly dynamic and highly uncertain. In the mission considered here, for example, one sub-goal may be to maintain a certain level of uncertainty in the explored environment (because map quality affects the flock’s ability to gain energy from thermals), a second sub-goal is to ensure that individual vehicles maintain enough altitude (i.e. energy management) and a third sub-goal is to ensure that in addition to the vehicle actively conducted

*Graduate Student, Department of Aerospace Engineering, Student Member AIAA.

†Associate Professor, Department of Aerospace Engineering, Associate Fellow AIAA.

target surveillance a second vehicle is ready to take over that task. Balancing these goals is key to successful mission completion.

This paper seeks to develop a unified cost function based on risk of mission failure, and this cost function will be used in task assignment. The motivation for using risk as the cost function lies in the unification of the cost (or reward) of each sub-goal: once the risk of performing (or not performing) a task is quantified, the development of an overall cost function follows. The tasks considered here include maintaining surveillance, mapping the environment (addressed in detail in earlier work⁶), exploiting energy available, and transiting to the surveillance target.

The focus of this work is on developing methods to assess the effect that neglecting a task has on the overall risk that the team will fail their mission goal. While the tasks considered here are specific to the problem of persistent surveillance by a flock of UAVs, the overall methods and approach are general.

The various methods by which risk is calculated will be addressed in Section II. The details of the specific behaviors and notes on implementation will be given in Section III. Results of various contrived situations as well as full simulation outcomes will be presented in Section IV. Advantages achieved by using this method for aircraft cooperation are discussed and the application of the presented method to a broader range of problem is shown in Section V. Finally, some closing thoughts will be presented in Section VI.

II. Problem Statement

Risk minimization for autonomous aircraft has been a focus in a number of works. In particular, Goerzen and Whalley discuss risk minimization in the context of motion planning by an autonomous rotorcraft.⁷ Where others have approached the risk minimization problem by associating a level of risk with regions in the environment (e.g. an aircraft should avoid risky areas and favor safe zones), here a risk is associated with an action. Specifically, the probability that the mission will fail given that the agent does not pursue a behavior is computed. This equates to a cost for neglecting a task, and will inform a decision by each agent on which tasks are safest to ignore (from the perspective of the team) and therefore which is the best to pursue in a particular circumstance.

This focus on the successful completion of the mission rather than the safety of an individual aircraft leads to altruistic behavior by the member agents. An aircraft will choose to save itself only if that action leads to a higher probability that the team completes the mission.

In a situation where only one aircraft is operating, the agent would be compelled to take the action with the lowest opportunity cost. This is the behavior that would lead to the highest probability of failure if neglected. In the case where multiple aircraft are operating, a task assignment method such as the Munkres Algorithm^{8,9} can be used to determine which activity allocation is most advantageous to the team. With tasks determined, a low-level controller or autopilot is assumed to direct the aircraft appropriately.

This section begins with an overview of the environment model used in all simulations, then proceeds to explain how the environment is interpreted by the aircraft controller. Methods for representing the aircraft energy state are examined and the approach taken in this work is justified. The tools for assessing the risks incurred by an aircraft in the course of mission related operations are then presented.

A. Convective Atmosphere Simulation

Thermal lift is modeled as extending from the surface to the top of the convective boundary layer. The thermal lift model is based on one presented by Gedeon¹⁰ which has been employed in other research¹¹⁻¹³ and has been used in previous work.⁶ The vertical wind induced by a thermal at a distance d from its core is given by

$$w_z(d) = w_0 e^{-\left(\frac{d}{R}\right)^2} \left[1 - \left(\frac{d}{R}\right)^2 \right] \quad (1)$$

where the strong core of the thermal is surrounded by a broader region of sink. Here w_0 is the maximum vertical wind speed and R is the thermal radius.

The thermals modeled are transient with a mean lifespan of 20 minutes and sigmoid growth and decay behavior. The core strength of a thermal follows a similar growth and decay profile between ground level and cloudbase, with the lift reaching full strength tens of meters above ground level and stretching to tens of meters below the top of the convective boundary layer before dropping off. The parameters for generated lift sources are Gaussian random variables derived from Allen¹⁴ and given by Table 1.

Table 1. Thermal parameters for boundary layer thickness $z_i = 1500$ m.

parameter	symbol	mean	σ	max/min	units
Period	$T_{thermal}$	20	5	60 / 5	minutes
Radius	R	75	20	200 / 30	m
Core strength	w_0	-2.56	1.5	-7 / -1	m/s

Thermal locations are generated as events in a Poisson process. This has the effect of making them equally likely to appear at any location in the environment. This, coupled with the large geographic region over which the wind is simulated, results in a uniform random distribution of lift sources. A minimum separation distance of $3R$ is enforced between thermals, which closely simulates reality where a rising column of air will deplete the warm air immediately surrounding the formation point. If the density of the Poisson process is relatively sparse enough, this separation distance only minimally affects the ergodic assumptions tied to the thermal generation process.

The vertical wind is a sum of the effects of the discrete convective features modeled. Thermals are local atmospheric attributes and thus do not greatly influence wind speeds outside of a small distance from their cores. Within the region affected by a thermal however, they are strong sources of atmospheric energy if utilized in an appropriate manner.

By comparison, large scale atmospheric phenomena have a relatively minor effect on the local winds in locations remote from terrain features. If the simulated missions were to be conducted in proximity to the ridges or mountains of Central Pennsylvania, then a model of meso-scale wind would be required for the fidelity of the simulation. Here, large-scale winds as well as cross-winds are ignored, so thermals are assumed to rise as vertical columns.

The wind in the environment may be measured by each agent, and is used to directly affect the agents' dynamics. For simplicity, the wind affecting the aircraft is not subject to any random noise in this work. The wind measurements taken by an aircraft are corrupted with Gaussian random noise however to simulate the behavior of sensors that would be used.

B. Atmospheric Energy Map

The atmospheric energy map is an agent's discrete representation of the wind conditions in the environment. Specifically, the map includes high-density real time wind information allowing the precise the locations of local phenomena such as thermals to be identified. The development of this map was the primary contribution of previous work by the authors⁶ and has since been further explored by Cheng.¹⁵ It was found that the availability of such a map and the knowledge of updraft locations stored within greatly enhanced the capability of a team of aircraft to exploit environmental energy during a mission. The mapping solution previously presented will be employed in this work as a tool for informing the risk computations that follow.

The domain in which the flock of aircraft is tasked to operate is divided into square cells. A cell contains estimates of the wind speeds (w_N, w_E, w_D) for the region that it covers. It is assumed that a cell is small enough that wind is approximately constant throughout. A cell also holds the variances associated with the wind speed estimates, providing a measure of how much those estimates may be trusted.

As the aircraft traverse the mapped region, it is assumed that they can measure local wind conditions (i.e. three components of wind velocity at the aircraft's location). These measurements may be used to augment the knowledge already stored in the map. Communicated measurements from other members of the team further increase the information available. These measurements are incorporated into the map using a Kalman filter. Because the world is assumed to be dynamic (a thermal lifespan is approximately 20 minutes) the the Kalman filter implementation allows the variance of the wind speed estimates in a cell to grow accordingly with time.

The map serves two purposes. It is used by an agent to determine where energy is available if it has been assigned the task of exploiting atmospheric energy. Alternatively, the variance level can drive exploration, encouraging the vehicle will fly towards regions of low confidence to improve its knowledge if it has been assigned the task of exploration. Figure 1 shows a snapshot of a map generated by a flock of 8 aircraft.⁶

In addition to the atmospheric energy map, an elevation map, the height of the atmospheric boundary layer, and knowledge of the performance parameters of all aircraft in the flock constitute the information

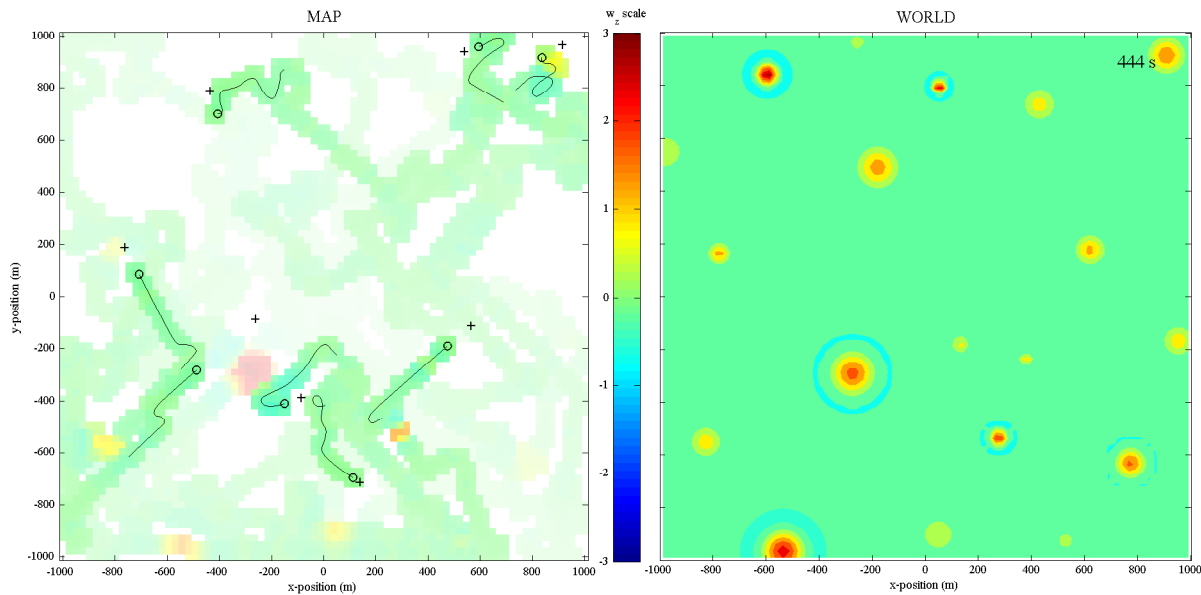


Figure 1. A map generated by eight aircraft.⁶ The figure to the right shows actual wind conditions, the figure to the left is the generated map. Note that many of the thermal features appear on the generated map. Aircraft flight paths for the preceding 30 seconds are shown by the lines, aircrafts' locations are shown by circles, and each assigned waypoint is indicated by a black '+'.⁶

available to the risk-based planner. This is a reasonable amount of information to store on an embedded computer system on even the smallest modern UAVs. The data, if stored as double precision floating-point numbers, requires approximately 0.4 Mb to capture a 2x2 kilometer map with a grid resolution of 25 meters. A larger, 10x10 kilometer map at the same resolution requires nearly 9 Mb of storage space. Both values are well within the capabilities of modern embedded systems.

C. Uncertainty in Aircraft Energy State

With perfect knowledge of the wind across the domain, and with perfect knowledge of how that wind affects the aircraft, propagating an aircraft's state through time would be a deterministic operation. This is clearly not possible. For this reason a stochastic approach is taken, where an aircraft's state is affected by a random process, yielding a probability density function that describes the distribution of possible states after the aircraft is propagated through time.

In its most basic form, the stored atmospheric energy map contains the aircraft's best estimate of the wind conditions at numerous locations. Wind in each cell is approximated by a set of Gaussian random variables along each axis described by a mean w_N, w_E, w_D and the associated standard variance. When this approximated wind acts on a simulated aircraft, it takes the form of the stochastic process introduced above.

The aircraft energy state can be represented as a probability density function. For projections of the energy state into the near future, this pdf is simple, centered about the aircraft's expected altitude with very little spread. The shape and spread of this density function becomes more complex as the state is projected further into the future. With a proper accounting for how wind will affect the aircraft's energy state, we can develop the necessary equations to project this state through the stochastic process of uncertain wind.

For example, we start with the simple case of an aircraft loitering in a single cell of the atmospheric energy map presented previously. Only one set of Gaussian wind components will affect the aircraft in this case. This contrasts the situation where the aircraft transits through multiple cells, and geographically differing stochastic wind processes must be applied. The more complex situation of a cross-country flight will be addressed after the concepts have been developed further.

The aircraft specific energy is typically given by the equation

$$e = \frac{v_a^2}{2g} + h + e_s \quad (2)$$

where v_a is the airspeed, h is the altitude of the aircraft, and e_s is the energy stored on the aircraft. In the following development, e_s will be ignored for the case of a gliding aircraft. This simplification is not necessary, but is done for clarity. The approach is equally valid if stored energy is accounted for and used to prolong flight (e.g. through the use of an electric motor or fuel powered engine).

The rate of change of total specific energy is

$$\frac{de}{dt} = \frac{v_a \dot{v}_a}{g} + \dot{h} + \dot{e}_s \quad (3)$$

In the case of a gliding aircraft maintaining a constant airspeed, the rate of change of total specific energy is

$$\frac{de}{dt} = \dot{h} = -v_a \sin \gamma - w_D \quad (4)$$

where $v_a \sin \gamma$ is the sink rate of the glider with respect to the surrounding air mass and w_D is the vertical component of wind speed (positive down). In steady state flight $\gamma = \arctan\left(\frac{C_D}{C_L}\right)$ is a function of airspeed, hence the sink rate is a function only of airspeed.

Over a time period dt , the incremental change in total energy is

$$de = (-v_a \sin \gamma - w_D)dt \quad (5)$$

The aircraft sink rate $v_a \sin \gamma$ is deterministic, but w_D is a Gaussian random variable (with a mean and variance stored in the atmospheric energy map). Hence de is a Gaussian random variable with mean \bar{de} and variance $\sigma_{w_D}^2 dt^2$.

Combining this with the pdf for the aircraft energy state takes the form

$$e_{t+1} = e_t + de \quad (6)$$

where e_t and de are independent random variables. The density function that results from the sum, e_{t+1} , is the convolution of the component density functions.

$$e_{t+1} = (e_t * de) \quad (7)$$

If we assume that the aircraft energy state distribution is of a simple form such as a Gaussian, this convolution is simple to solve analytically. This may be a valid approximation until the aircraft state is propagated near a boundary. At these points there are hard nonlinear limits in the stochastic wind process caused by the terrain and lack of lift beyond the convective boundary layer. These will make the shape of the aircraft energy state distribution difficult to express parametrically.

D. Particle-based Propagation of Aircraft Energy State

The approach taken in this paper is to propagate the aircraft energy state using a particle filter inspired method. Provided that there are enough particles involved, the filter should approximate the same results as the analytic solution. Additionally, a particle filter is nonparametric, meaning that it is agnostic to the shape of the probability density function (e_t) to be propagated.¹⁶ Similarly, the nonlinearities in the stochastic process de do not introduce difficulties. This approach becomes necessary when an aircraft's projected state is near the ground or at the top of the boundary layer where the shape of the probability density function becomes strongly skewed.

In the absence of the boundary conditions that introduce difficulty computing the parametric shape of the aircraft energy density function, the energy state of a candidate particle (i) is propagated. Mirroring Equation 5,

$$de(i) = (-v_a \sin \gamma - w_z(i)) dt \quad (8)$$

where $w_z(i)$ is a random sample drawn from the distribution W_z . This is summed with the previous energy state of the particle to yield an analog to Equation 6.

$$e_{t+1}(i) = e_t(i) + de(i) \quad (9)$$

There is no true measurement update for this particle filter approach, however viable possible energy states are conditioned on the terrain elevation map carried by the aircraft. The aircraft, given that it is




Figure 2. The energy state of an aircraft with a sink rate of 0.55 meters/second is projected forward in time. The wind through which the state is propagated is a zero mean Gaussian with a variance of one. The aircraft starts at $t = 0$ at an altitude of 750 meters MSL over terrain with an elevation of 350 meters (shown in red). The sink rate of the aircraft in still air is shown by the black dotted line. Five percent of the particles that represent candidate energy states are plotted in blue.

still airborne, must be above the ground and below cloudbase. Thus, inaccessible energy states may be eliminated from consideration and further propagation.

The maximum value for a candidate energy state is set at the top of the convective boundary layer. Using the convective energy from thermal lift alone, it is unreasonable to assume that an aircraft can climb above this level. To deal with the upper boundary, any particles that have been propagated to this level are restricted from passing above cloudbase.

Weights are assigned to candidate particles based on their compliance with the boundary conditions. The weight is a binary value, one if the projected state has not violated the terrain boundary condition (the elevation stored in cell j of the map), and zero if it has.

$$w(i) = \begin{cases} 1 & : e_{t+1} > g(j) \\ 0 & : e_{t+1} \leq g(j) \end{cases} \quad (10)$$

The particles with a weight of zero are discarded from further consideration and the resulting distribution of particles is randomly sampled from those that remain.

The probability that the aircraft survives the loiter period dt is equivalent to the proportion of particles that remain after the terrain boundary condition is applied. With enough particles representing candidate energy states, this is a good approximation of the risk associated with loitering. Similarly, the resulting distribution of particles after resampling will accurately reflect the posterior distribution of the aircraft energy state.

Figure 2 shows an example of the particle filter approach to aircraft state propagation. Note that as the distribution approaches the ground, it becomes skewed. The probability that the aircraft survives a time step given that it has reached the last one is displayed at the top of the plot. The probability that the aircraft is still airborne after 540 seconds is then the product of these independent probabilities, or 0.6466.

E. Reconciling the Discrete Map with Energy State Propagation

The simple case of an aircraft subject to the wind conditions in a single cell of the discrete map was used to develop the approach taken in this work to propagate aircraft energy state through time. A complex situation arises when an aircraft must cross multiple cells to arrive at a destination. The particle filter inspired approach will be used to evolve the energy state once the necessary equations are given.




Figure 3. The energy state of an aircraft with an L/D of 27 is projected through a ridge crossing maneuver. The vertical wind through which the state is propagated is a zero mean Gaussian with a variance of one. There is no horizontal wind. The aircraft starts at $d = 0$ at an altitude of 500 meters MSL over terrain with an elevation of 350 meters (shown in red). The mapped elevation of the terrain rises as the aircraft approaches the summit and drops afterwards. Five percent of the particles that represent candidate energy states are plotted in blue. Note that the irregular distance intervals correspond to the portion of the course that lies in each cell of the wind map.

A flight from one location to another can be represented as a connected set of cell crossings in the stored atmospheric energy map. A fast method for determining which cells are entered while transiting between an agent's start and goal locations is given by Siddon.¹⁷ Once the relevant cells are identified, the length ($l(j)$) of the traversal through each cell is obtained by the same method.

The aircraft energy state is propagated through each crossed cell independently because a different stochastic wind process acts on the aircraft in each cell. The specific energy equation is visited again, however the derivative is taken now with respect to ground distance traveled. The same assumption of constant airspeed is made.

$$\frac{de}{dx} = \frac{de}{dt} \frac{dt}{dx} = \frac{(-v_a \sin \gamma - w_D)}{v_g} \quad (11)$$

v_g is the ground speed of the aircraft. Because the goal is to find the stochastic process representing the change in aircraft energy in crossing a particular cell j , dx is equivalent to the length of the path that lies in cell j , or l_j .

$$de_j = \frac{l_j (-v_a \sin \gamma - w_D)}{v_g} \quad (12)$$

The ground speed is a random variable arrived at by a combination of the two Gaussian random variables representing horizontal wind speed estimates. With knowledge of the path bearing ψ and the horizontal wind speed for a cell, the in-track w_t and cross-track w_c winds can be calculated as a vector projection and rejection of the wind on the track vector ($\hat{t} = [\cos \psi, \sin \psi]$) respectively.

Thus the ground speed can be expressed as in Chakrabarty and Langelaan¹⁸ as

$$v_g = \sqrt{v_a^2 - w_c^2} + w_t \quad (13)$$

The stochastic process affecting the aircraft energy state in the crossing of cell j is the expressed fully as

$$de_j(i) = \frac{l_j (-v_a \sin \gamma - w_D(i))}{\sqrt{v_a^2 - w_c(i)^2} + w_t(i)} \quad (14)$$

where $w_D(i)$, $w_c(i)$, and $w_t(i)$ are randomly sampled from the distributions $w_{D,j}$, w_c , and w_t respectively.

Propagating the energy state of the i^{th} particle across a cell j takes the form

$$e_{j+1}(i) = e_j(i) + de_j(i) \quad (15)$$

where $e_j(i)$ is the energy state of the i^{th} particle before crossing cell j and $e_{j+1}(i)$ is the energy state after crossing the cell.

As before, weights are assigned to candidate energy states in a binary fashion. If a state is projected below the mapped elevation in cell j , it is removed from consideration. The probability that the aircraft can pass through cell j is approximated by the proportion of candidate states that remain. The candidate energy states used at the beginning of the next cell crossing are drawn randomly from the remaining candidates. This action makes cell crossings independent events and simplifies the probability computation. Once the aircraft energy state has been propagated through all of the cells, the probability that the route may be traversed is the cumulative probability of successfully crossing each cell independently.

An aircraft's energy state is projected through a ridge crossing in Figure 3. Again, as the energy distribution is propagated in proximity to the mapped ground elevation, it becomes strongly skewed. Displayed on the plot are the probabilities that the aircraft can cross an individual cell, independent of other cell crossings. The probability that the aircraft can successfully complete the nearly two kilometer cross-country traverse over the ridge is the product of these independent transition probabilities and is equal to 0.2408.

F. Probabilistic Glide Range “Amoeba”




Figure 4. The glide footprint of an aircraft with an L/D of 27 at 600 meters MSL located at the center of a terrain map is plotted in solid blue. Terrain elevations are given in meters. This region corresponds to a probability of 0.99 that the aircraft can reach anywhere within the boundary. The solid red line is the region that is observable using a 90° field of view downward pointing camera aboard an aircraft at that altitude in that location. The dotted blue line represents the $P = 0.99$ glide amoeba for the aircraft after its state has been projected forward in time 60 seconds. Note that the irregular boundary on the projected amoeba is due to uncertainties in mapped wind speeds. The projected observation region is outlined by dotted red.

Rather than find the probability that an aircraft can make it to a specific destination, a related approach may be used to compute the expected range of the aircraft flying in a straight line given a starting energy state distribution. With this approach, the path is not terminated upon reaching a destination, but rather when the probability that the aircraft is still airborne passes below a threshold.

Siddon's method¹⁷ for finding the cells crossed by the ray and the distance that the ray travels through each cell is again used. The particles representing candidate aircraft energy states are propagated cell by cell as before. The probability of crossing each cell given that the energy state is known as it enters the cell are

independent. After crossing each cell, the cumulative probability of crossing each cell until that point, which is equivalent to the probability that the aircraft survives the journey from the origin to the end of the current cell, is computed. Once this cumulative probability drops below a predefined threshold, the propagation is stopped. Because the stochastic process of crossing a cell is linear, the fraction of the last cell that could be transited is easy to compute. The range of the aircraft along the ray is recorded.

This computation is performed along several bearings from a starting location and energy state, a range-bearing map can be built. This map shows the “glide footprint” of the aircraft for a given probability of success. This is similar to the glide range calculations done by commercial glide computers used in soaring flight by sailplane pilots. In this application it is sometime referred to as a glide amoeba due to the irregular shape of the glide range footprint projected onto the ground.

The usefulness of this calculation lies in the ability it gives an aircraft to quickly reference which features are within range at any time. Features that pass this filter may be examined further by computing the exact probability that the aircraft can transit to the feature’s location.

G. Probabilistic Observation Region “Amoeba”

A prediction of the observation range of an aircraft equipped with a camera can be made with knowledge of the surrounding terrain, the distribution of possible altitudes, and the camera field of view. Unlike the glide range calculation, a ray representing the edge of the observable region is not affected by any stochastic processes as it travels. The uncertainty in the observation region comes instead from the uncertainty in the observation altitude.

The observation range for a single particle (a particular instance of the energy state probability distribution) is determined along a number of bearings from the initial location. The range is recorded when a ray that lies on the boundary of the camera’s field of view intersects the mapped terrain height as it crosses a cell. What results is a possible observation region as computed from a single particle. When the observation ranges along a bearing for each of the particles that makes up the possible aircraft energy state are computed, the statistics of a probable observation range can be returned.

Like the glide region, the probabilistic observation range provides a convenient way to check if the observation target will be in sight of the aircraft after it transits to a surveillance point.

H. Probability of Finding Lift

If lift sources are taken to be randomly located with a uniform density over the map (they are instances of an ergodic point process), then the probability of encountering lift may be easily computed. The thermals modeled herein are generated by a Poisson point process with an intensity measure of λ_{acc} , thus for this simulation, the uniform density assumption is valid.

Under this assumption, the probability that an area A_i of the map does not contain a source of lift is

$$P^c(B_i) = e^{-\lambda A_i} \quad (16)$$

and the probability that a lift source lies within the area A_i is

$$P(B_i) = 1 - e^{-\lambda A_i} \quad (17)$$

When computing the probability that lift has been identified within the glide range of an aircraft, λ is taken to be the density of identified lift sources on the map, and A_i is the area covered by the glide range footprint of the aircraft. To compute the probability of stumbling across a source of lift, A_i is taken to be the area swept by the aircraft in the course of transiting the map. It is assumed that an aircraft can identify a thermal if it passes within a radius R of its center. Therefore, the area swept is $2Rl$ where l is the longest distance the aircraft can travel before landing out. The density in this case is the expected thermal density at the time of the mission.

I. Soaring and Persistent Surveillance

In the context of persistent surveillance, soaring flight is used to extend endurance of the UAVs. It is therefore simply one of many tasks that can (or must) be performed during the mission. Task definition and task assignment are thus a critical problems that must be addressed

One approach is to implement a set of behaviors that encode the necessary tasks: surveillance, energy exploitation, exploration, and transit to the surveillance region are four possible behaviors. The problem now is to assign tasks to the members of the flock so that overall flock performance is maximized.

In a task assignment matrix

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \dots & a_{mn} \end{bmatrix} \quad (18)$$

a particular entry a_{ij} defines the reward accrued if an agent i performs a task j . Here a set of four behaviors are defined:

1. target observation: the aircraft maneuvers to ensure that the target is in view;
2. transit to target observation: the aircraft flies to a point where target observation can begin;
3. exploration: the aircraft explores the region, mapping regions of upwards air motion;
4. exploitation: the aircraft exploits vertical air motion to increase its total energy.

Note that multiple aircraft can be assigned to a particular task (although for some tasks, such as target observation and transit to target observation, this will be sub-optimal). Further, if there are less than four aircraft in the flock then some tasks will be unfilled at each time step. Task assignment is a rich field of research: here the Munkres algorithm⁸ is used. Necessary tasks and a method for computing the reward of performing a task can now be defined.

III. Behavior Descriptions and Associated Risk

The tools and methodology developed in the previous section will now be used to aid in the computation of risk associated with conducting behaviors. It is important that the definition of the behaviors be clear and complete such that risks may be associated with them. Here several behaviors are defined that are used in this investigation, but the methods applied to the risk definition may be used with a much broader group of behaviors.

As was shown in earlier work by the authors⁶ the endurance and efficiency limitations can be mitigated, and prolonged missions conducted, when a team of multiple aircraft is used. Information shared between the aircraft allowed each to build an energy map of the environment, which was utilized when an aircraft found itself in need of energy. Because the mission in the prior work was only to build and maintain the most accurate atmospheric energy map possible, the behaviors were limited to an exploration and an exploitation mode. An aircraft deciding which behavior to pursue was a simple matter of assessing its energy state at the beginning of the planning period. The aircraft would then only engage in exploitation when energy was needed, the default mode being exploration.

In defining behaviors and associated risks, all aircraft begin at a high altitude and are afterwards assumed to be in gliding flight. Some aircraft will be assigned to explore the environment and search for sources of atmospheric energy while at least one aircraft will be assigned to observe the target. Each aircraft is continuously taking measurements of wind speed at its location and broadcasting those measurements to the other members of the flock. Because no communication dropouts are presently assumed, the maps that each aircraft build will be identical as the same information will be stored in each.

A interval of time is specified over which an aircraft is expected to hold the same behavior. At the beginning of an interval, tasks will be assigned and held for the duration of the interval. Task switching may be done at the start of the next interval. Taking time in discrete intervals makes the risk computation a more straightforward process and ensures that a task is performed for a reasonable length of time before re-assignment occurs.

A. The Target Observation Behavior

The observation behavior is directly responsible for the success of the mission. Mission failure has been defined as losing sight of the target for any length of time: thus at least one aircraft must be engaged in this

behavior at all times. All other behaviors can be looked at as supporting this primary mission goal.

If the target is located within the observation region of one of the aircraft, it is under surveillance. In this case there is no risk to the mission. An aircraft that is observing the target will loiter in the airspace above the target for the duration of its assigned time on station. This behavior is conducted at the airspeed for minimum sink to prolong the time that may be spent over the target.

The risk to the successful completion of the mission incurred by an aircraft not watching the target is a function of both that aircraft's location relative to the target and the locations of all other aircraft in the team relative to the target. Computing the risk due to not watching the target is an irrelevant calculation if the vehicle is first not in a position to do so. Therefore, the observation region is first checked to ensure that the target is located within. If not, this computation does not proceed further.

If the aircraft is able to observe the target from its present location, the risk to the mission should it ignore this task is determined as follows. We examine the risk from the perspective of a single aircraft n in the team of N aircraft:

$$P(D) = P(D|E_{t+1})P(E_{t+1}) + P(D|E_{t+1}^c)P(E_{t+1}^c) \quad (19)$$

where D is the event that the mission is failed given that aircraft n does not observe the target for the current period. Event E_{t+1} is the event that the target is under observation by any of the other aircraft after a single period. The complement of event E_{t+1} , or E_{t+1}^c , is the event that the target is not under observation for the current period. As defined in the mission statement, $P(D|E_{t+1}^c)$ is equal to one, implying that if the target is not under observation, the mission will be failed.

The subscript $t + 1$ is used because all aircraft are assumed to hold their behavior for the length of one period. Thus, an observing aircraft spends the entire period attempting to maintain watch over the target. The probability that the target is still under surveillance at the end of the period represents the lowest probability of observation for the whole period. This is a conservative estimation of the risk.

The probability that the target is not under observation by any of the other aircraft for the duration of the current period takes the form

$$P(E_{t+1}^c) = \prod_{i=1, i \neq n}^N (1 - P(E_{i,t+1})) \quad (20)$$

with $E_{i,t+1}$ being the event that the target lies within the observation region of aircraft i for that time. The value $P(E_{i,t+1})$, the probability that the target is within the observation range of aircraft i , is:

$$P(E_{i,t+1}) = P(E_{i,t+1}|G_{i,t+1})P(G_{i,t+1}) + P(E_{i,t+1}|G_{i,t+1}^c)P(G_{i,t+1}^c) \quad (21)$$

The target will only be under observation if the aircraft is still airborne at $t + 1$. Therefore, the probability that aircraft i is observing the target is conditioned on the probability that the aircraft is still airborne at that time, $P(G_{i,t+1})$. This also means that the probability $P(E_{i,t+1}|G_{i,t+1}^c)$ is equal to zero, as if the aircraft does not survive, it can no longer observe the target.

To properly account for the risk imposed on the mission due to the energy lost by an agent while observing the target, the probability $P(D|E_{t+1})$ must be evaluated. Defining event C_{t+1} as the event that the observing aircraft are unable to find a source of lift after loitering over the target for a period, the risk is equivalent to:

$$P(D|E_{t+1}) = P(D|C_{t+1})P(C_{t+1}) = \sum_{i \in obs} P(D|C_{i,t+1})P(C_{i,t+1}) \quad (22)$$

Aircraft i is in the observing set if $P(E_{i,t+1}|G_{i,t+1})$ is greater than zero. This value serves to differentiate candidate observers who may be at different altitudes over the target. The calculation of $P(C_{i,t+1})$ and $P(D|C_{i,t+1})$ will be discussed further.

B. Transit to Observation Target Behavior

The purpose of the transit behavior is to drive an aircraft, which may be in a remote region of the map, to an area from which it can observe the target. This transit should take no more than one time step to complete. An aircraft assigned to transit will depart immediately for the target observation region. So that

the aircraft is in the most advantageous possible energy state on arrival, the transit maneuver is conducted at the airspeed for best glide. Once over the target, the recently arrived aircraft will begin to circle in such a way that the target remains in sight.

The risk to the mission here is a function of the current energy state of the observing agent, and the projected energy state of the agent when it reaches the observation area. Essentially this is the probability that the current observing agent is compelled to seek a source of lift before a replacement arrives to continue surveillance of the target if the agent does not depart to fulfill that role now.

It is assumed here that the aircraft can reach the observation target within one period of committing to do so. This makes the length of the decision period contingent on the performance of the aircraft as well as the size of the map. If the aircraft does not transit to the observation region at the current time, then the current observer must remain on station for at least two more periods ($t + 2$) before relief arrives.

Again, we examine the risk as calculated by a single aircraft n among a team of N aircraft. If aircraft n is already in a position from which it can observe the target, the transit behavior is not an available task and the computation of the risk associated with ignoring this task is irrelevant. If the aircraft is not in a position above the target, the risk to the successful completion of the mission imposed by ignoring the task is determined as follows.

If event D is the event that the mission is failed given that aircraft n does not transit to a point from which it can observe the target, then it may be expressed as:

$$P(D) = P(D|E_{t+2}) P(E_{t+2}) + P(D|E_{t+2}^{\mathbb{C}}) P(E_{t+2}^{\mathbb{C}}) \quad (23)$$

where E_{t+2} is the event that the observation target is still under observation by any aircraft after two periods. The probability that the mission is failed given the target is not in sight of any of the observing aircraft after two periods $P(D|E_{t+2}^{\mathbb{C}})$ is defined by the mission objective to be one. As before, the probability that the target is not under observation may be expressed as:

$$P(E_{t+2}^{\mathbb{C}}) = \prod_{i=1}^N (1 - P(E_{i,t+2})) \quad (24)$$

this time, aircraft states are projected two periods into the future. Similarly, the probability that each aircraft i of the team is still able to watch the target after two periods is:

$$P(E_{i,t+2}) = P(E_{i,t+2}|G_{i,t+2}) P(G_{i,t+2}) + P(E_{i,t+2}|G_{i,t+2}^{\mathbb{C}}) P(G_{i,t+2}^{\mathbb{C}}) \quad (25)$$

which is conditioned on $G_{i,t+2}$, the event that aircraft i is still airborne at the end of two periods. The probability that the target is still within sight of an aircraft that is no longer airborne $P(E_{i,t+2}|G_{i,t+2}^{\mathbb{C}})$ is equal to zero. The probability of aircraft i remaining airborne, $P(G_{i,t+2})$, for two periods is determined by propagating its state forward through a loitering process as presented in Section D.

The probability that the target remains within sight of the observing aircraft after two periods is determined by computing the probabilistic observation region of the aircraft after it has loitered for two periods (the density function resulting from the computation of $P(G_{i,t+2})$), and checking the certainty that the target will still be observable.

To account for the risk that any observing aircraft are taking by losing energy while staying on station, and therefore contributing a level of risk to the completion of the mission, the value $P(D|E_{t+2})$ is defined. This is the risk that, given an aircraft stays on station to observe the target for two periods, it can no longer find a source of lift from its new (likely lower) energy state. Similar to the observation case, here $P(D|E_{t+2})$ is expressed as:

$$P(D|E_{t+2}) = P(D|C_{t+2}) P(C_{t+2}) = \sum_{i \in obs} P(D|C_{i,t+2}) P(C_{i,t+2}) \quad (26)$$

This equates to the probability that the mission will still fail given the aircraft involved in observing the target are still airborne after two periods. If an observer were required to stay on station observing the target for too long before being relieved, it may not be able to reach a reliable source of lift from its resulting energy state. Thus the importance of an observing aircraft to the mission $P(D|C_{i,t+2})$ is taken as the probability that the mission will fail in the event of $C_{i,t+2}$, that the observer is unable to find lift from its energy state after two periods.

C. The Exploration Behavior

The probability that the current level of knowledge will ultimately cause mission failure drives the exploration behavior. An aircraft that is exploring the environment does so to improve the shared map of atmospheric energy and to decrease the uncertainty in the mapped wind estimates. This equates to making the map more closely resemble reality. Adequate exploration by the team ensures that the knowledge of lift locations stored in the map is accurate as well as current, and is useful to an aircraft in need of a lift source. The exploration behavior is a way for the flock to plan ahead for times when energy will be required in the future.

Many of the wind features that can be used for lift (eg. thermals) are transient in nature, and thus the map needs to be refreshed frequently. This is done by returning as often as possible to regions that have been previously measured to track changes in wind over time. Because of its role in prolonging mission duration, exploration is the default behavior for any of the team members, and is pursued when no other tasks are pertinent. The mechanics of the exploration behavior were presented by Depenbusch⁶ and further refined by Cheng¹⁵ who presented methods for more efficiently searching the map for lift when the geography is known.

Essentially, an aircraft should visit cells in the map where the uncertainty in the wind estimates is highest. Because an aircraft is continuously taking measurements, the act of visiting a region on the map (or simply passing through) adds measurements to the estimate of wind speeds in that region, decreasing uncertainty. The map knowledge decays with time, so a region that has recently been visited by any of the team members will be less attractive as a destination than one that has not been visited for some time. For the purposes here, it suffices to say that exploration decreases uncertainty and increases the likelihood that an aircraft in need of atmospheric energy will be able to find it.

If an aircraft encounters a strong source of lift while exploring, it may temporarily abandon the behavior to climb. This helps to positively identify lift sources. During the climb, many measurements of the wind speed are taken and the map is updated to reflect the certainty that lift can be found at this location. The threshold for this diversion is a function of the aircraft's altitude and the strength of the lift encountered.

To compute a risk associated with not exploring the map, the energy state of each aircraft in the team is considered. The map is used as a way to store information on sources of available lift, if the map does not contain any information on available lift, a large risk is posed to the team.

For each aircraft we compute the probability that the aircraft will not find a source of lift before it lands out. An aircraft may find lift by examining a lift source within its reachable glide range, or an aircraft may "stumble upon" a source of lift while traversing the map. The first scenario is more likely if the aircraft is high above the ground (it can glide far) and if more lift sources have been identified. The second scenario, which does not rely on mapped information, is also more likely to occur if the aircraft is high due to the longer range it can fly from that altitude.

The exploration risk is computed identically by each aircraft, the only differences may arise from different maps. The risk is an approximation of the probability that the team of N aircraft will fail the mission due to the current knowledge of the wind conditions stored in the map. The probability that the mission is failed due to a lack of accurate knowledge is then:

$$P(D) = \sum_{n=1}^N P(D|C_i) P(C_i) \quad (27)$$

where D is the event of mission failure if an aircraft does not pursue the exploration behavior with the goal of adding to the knowledge stored in its map. The importance of aircraft i to the completion of the mission is given by $P(D|C_i)$, the probability that the team will fail the mission given aircraft i lands out.

Event C_i is the event that aircraft i cannot find lift before it lands out, then it may be defined as:

$$P(C_i) = 1 - \left[P(A_i|B_i) P(B_i) + P\left(A_i|B_i^c\right) P\left(B_i^c\right) \right] \quad (28)$$

where A_i is the event that aircraft i finds lift and B_i is the event that lift is present within the glide range of aircraft i .

The value of $P(A_i|B_i)$ is conservatively taken to be the threshold probability used to define the glide range in Section II.F, and should always be a value near one. The probability that lift is found within the glide range of the aircraft $P(B_i)$ is a function of the known density of lift on the map, and the area that the aircraft can reach (a function of the aircraft's altitude). This is the value that drives exploration, if the density of known lift sources is very low, or zero, then then probability of a known lift source lying inside of the glide range of the aircraft will be very small.

The probability that the aircraft will stumble upon lift given that it cannot locate a suitable source within glide range $P(A_i|B_i^c)$ will always be a relatively small value. This value is related to the area that is swept out by the glider as it crosses the environment. The expected density of lift sources also factors into this calculation.

D. Exploit Energy Behavior

If an agent is in a low energy state, it needs to exploit lift available to gain energy and continue the mission. The energy exploitation behavior provides a means for accomplishing this. When commanded to exploit atmospheric energy in the environment, an agent will first examine the energy map for sources of reliable lift. The map, which has been updated continuously by the aircraft engaged in the exploration behavior, and will now be used. When a suitable area of lift is identified, the agent will transit to that location and proceed to climb. A thermal centering controller is needed to take full advantage of the lift available, and is covered in prior work by Andersson and Kaminer.³

In the case that no suitable lift sources can be identified, the aircraft will take the bearing for which the longest probable flight is possible. This computation was done previously when computing the glide footprint of the aircraft.

The risk to the mission of losing an individual agent is a function of the number of agents that remain active. The problem then becomes to determine the probability that an agent will be within range of a known source of lift if it does not choose to exploit now.

Unlike the other risks computed, the risk incurred if an aircraft does not choose to exploit the energy mapped is solely dependent on the aircraft's own energy state. All other risks take into account the energy states of others. This serves as a primary distinguishing value that differentiates the actions of the aircraft comprising the flock.

Recall that all agent behaviors are assumed to be held for the duration of a planning period. If the aircraft chooses not to take advantage of atmospheric lift at this time, it will be compelled to wait for a full period before seeking out a source of lift. To define the probability that the aircraft lands out before finding a suitable source, we examine the actions available to the aircraft. The probability being computed is the risk incurred for neglecting the current action, seeking lift. Therefore, the probability of the aircraft landing out is the probability that the aircraft lands out, given that it does not seek out lift at the present time. Therefore, all computations are done from the aircraft's projected energy state at $t + 1$.

The time until the aircraft is able to change its behavior is important here. If the aircraft does not pursue the exploitation task now, it is assumed that the next opportunity it would have to do so is at the next behavior change. For this reason, the aircraft's energy state is propagated forward in time for this duration. All assessments will be done from this propagated state.

Again, event D is the event that the mission is failed given that the behavior is ignored, and event $C_{n,t+1}$ is the event that the current aircraft (n) lands out before it encounters a suitable source of lift after waiting for one period before seeking that lift. Thus the probability of mission failure for neglecting the energy exploitation behavior is:

$$P(D) = P(D|C_{n,t+1}) P(C_{n,t+1}) \quad (29)$$

Note that the probability that the mission is failed given that the aircraft does find a source of lift before landing out ($P(D|C_{n,t+1}^c)$) is a difficult value to conceptualize or define, and should be very small, so it is ignored.

The probability that the aircraft does not find lift from its energy state projected one period into the future is

$$P(C_{n,t+1}) = 1 - \left[P(A_{n,t+1}|B_{n,t+1}) P(B_{n,t+1}) + P(A_{n,t+1}|B_{n,t+1}^c) P(B_{n,t+1}^c) \right] \quad (30)$$

Event $A_{n,t+1}$ is still the event that the aircraft encounters a suitable lift source. Event $B_{n,t+1}$ is the event that there is an identified source of suitable lift located within the glide range of the aircraft after it has pursued a different task for the duration of one planning period.

IV. Simulation Results

To provide a realistic simulation environment a terrain map of the State College, Pennsylvania area is used in the persistent surveillance scenario. Terrain elevation data is obtained from the USGS^a and modeled at a resolution of 30 meters. An aircraft is deemed lost once its altitude is less than terrain height at its location.

The task consists of maintaining surveillance of a target at the center of a $2 \text{ km} \times 2 \text{ km}$ region. Flock sizes of three, five, and seven aircraft are simulated to examine differences in behavior as flock size increases. Four environmental thermal densities are simulated: one thermal every two square kilometers, one thermal every square kilometer, two thermals per square kilometer, and four thermals per square kilometer. The differing environmental conditions are utilized to assess the effect that the abundance of lift has on the success of the observation mission. Because of the random distribution of lift sources, the thermal densities do not guarantee that thermals will be present, but rather give expected values.

A. Simulation Setup

A homogeneous flock of aircraft representative of the RnR Products SB-XC (a 4.1m wingspan radio-controlled glider) is used: maximum L/D is 27 at an airspeed of 15.8 meters per second and minimum sink rate is 0.55 m/s at an airspeed of 14.5 m/s. No motor model was employed as only gliding flight is simulated.

Risks associated with neglecting behaviors are computed as presented in Section III. The task assignment is performed using an implementation of Munkres algorithm. The observation task is assigned first, afterwards all of the other tasks are allocated. There are $N - 1$ exploration tasks available to the aircraft (where N is the number of aircraft in the flock). Similarly, there are $N - 1$ energy exploitation tasks up for assignment. There is a single transit task available for assignment. This ensures that only one aircraft is tasked to take over for the current observer while the rest of the flock either explores the environment or exploits the energy available.

The interval between task assignments is defined to allow an aircraft at the most distant location in the map to return to the target's location in one step. For the modeled aircraft and the size of the region mapped here, this interval is 95 seconds.

B. Monte Carlo Simulation


Twelve simulated missions, each lasting one hour for each set of flock size and lift abundance were conducted. A one hour mission is long enough that an aircraft must encounter several sources of lift over the course of the mission to survive the whole time. This duration is sufficient to demonstrate that the cost function results in effective behavior assignment.

The aircraft are started at approximately 1000 meters MSL, randomly dispersed over the region to be mapped. No *a priori* information is used to initialize the atmospheric energy map. Care is taken to ensure that at least one aircraft begins the mission in a position from which it can observe the target. This initial condition for the flock roughly approximates the situation where a flock of aircraft have been launched, their motors have been used to climb to altitude, they have transited to an observation region, and now they are commanded to begin their persistent surveillance mission.

Figure 5 shows a summary of the mission simulations. Figure 5(a) shows the thermal density plotted against the mean percentage of the mission duration (one hour) for which the target is under observation for each flock size. The horizontal dotted line indicates the baseline scenario where a single gliding aircraft watches the target with no atmospheric lift present. In this case the aircraft will land out after approximately 32 percent of the mission time. Any improvement above this baseline can be attributed to the effectiveness of the cooperation strategy in finding lift, sharing wind measurements, and dividing tasks in such a way that the team is more capable.

The mean percentages of aircraft that survive until the end of the same missions are plotted in Figure 5(b) against thermal density. As mentioned, a single aircraft would be unable to last for the duration of the mission. For an aircraft to survive, it must have located at least one, and likely more, lift sources from which to gain altitude.

^a<http://ned.usgs.gov>



(a) Mean percentage of the mission duration that the target is under observation. The horizontal dashed line represents the capability of a single gliding aircraft starting from a similar initial altitude. (b) Mean percentage of the original flock members that survive until the end of the simulated mission.

Figure 5. Summary of the Monte Carlo simulations. The performance of the seven aircraft flock is plotted in red, the five aircraft flock plotted in green, and the three aircraft flock plotted in blue. All values are a mean over 12 simulations.

C. Behavior Assignment Examined

Figure 6 shows a plot of the path of the aircraft in a five agent flock over the duration of a mission. A single aircraft is highlighted, and its path is colored by the behavior assigned to it. The time history of the aircraft altitudes in the simulation depicted in Figure 6 is shown in Figure 7. Note that two aircraft from the original flock of five are lost over the course of the simulation. The first lands out approximately 1200 seconds into the simulation after failing to find a lift source.

The highlighted aircraft is first assigned the task of exploring the environment. This portion of its path in Figure 6 and Figure 7 is colored blue. The aircraft is fortunate to find a source of thermal lift shortly after beginning its exploration task. Because the strength of this thermal is great enough, the aircraft takes advantage of the lift available. While it has still been assigned the task of exploration, the aircraft temporarily abandons this mission to climb.

The spiral column in Figure 6 indicates that the aircraft is taking advantage of thermal lift. This corresponds to a steep positively sloped section on the barogram plot (Figure 7). The varied radius of the spiral path is due to the aircraft seeking to maximize its climb rate in a dynamic thermal. It may take several turns for the aircraft controller to “center” the thermal and find the optimal radius. All thermals become weak at the top of the boundary layer (here 1500 m MSL), explaining the drastically decreased radius of the turns at the top, and the decreased lift experienced. An aircraft will leave a lift source if the lift experienced is below a threshold value or if it has been assigned a task that disallows thermalling. The threshold value increases with altitude, so aircraft may out climb a thermal’s usefulness. The highlighted aircraft does this as it passes above 1500 meters, the top of the modeled convective boundary layer.

At $t = 380s$ the highlighted aircraft has the highest altitude of any of the team members. From its location it is able to observe the target. In this way it is qualified to and is consequently assigned to pursue the observation task. When the aircraft is observing the target (in Figure 6 shown in green), it circles in the target’s location. As the aircraft sink lower, the circle that it makes is necessarily constricted due to the observation cone. Note that the aircraft is often assigned the observation task when it is the team member with the highest altitude, though other factors such as location have some influence.

As the highlighted aircraft observes the target, it sinks. After some time the aircraft is again assigned




Figure 6. Plotted are the paths taken by members of a five aircraft flock over the course of a one hour simulated observation mission. The thermal density in the environment is four thermals per square kilometer. The height of the convective boundary layer is 1500 meters MSL. The path of a single aircraft is highlighted and its behaviors are indicated by the color of the line: blue = exploration, green = observation, black = transit, red = exploit.




Figure 7. Time history of the altitudes of members of a five aircraft flock over the course of a one hour simulated observation mission. The thermal density in the environment is four thermals per square kilometer. The height of the convective boundary layer is 1500 meters MSL. The barogram of a single aircraft is highlighted and its behaviors are indicated by the color of the line: blue = exploration, green = observation, black = transit, red = exploit.

to explore the environment as another member of the team takes over observation duties. In the course of exploring the environment, the highlighted aircraft passes outside of the region from which it can observe the target. When it becomes too risky for the observing aircraft to continue this task, the highlighted aircraft is assigned the transit behavior (colored black in Figure 6 and Figure 7). When pursuing the transit behavior, the aircraft flies towards the target and begins to circle.

The highlighted aircraft then observes the target for a short stint before again being assigned to the exploration task. As the aircraft's altitude drops below approximately 500 meters MSL (approximately 150 meters AGL over this terrain), it is assigned to exploit the available atmospheric energy. The path of the aircraft while pursuing the exploitation behavior is colored red. This behavior is assigned when the aircraft is low and when another member is able to watch the target. The highlighted aircraft is able to take advantage of the knowledge of a thermal location stored on the map to climb. The aircraft continues to climb even after it has been again assigned to explore because the measured strength of the lift is strong enough to justify temporarily neglecting exploration.

Figure 6 shows the myriad of behaviors that must be pursued by each aircraft, sometimes on multiple occasions, to successfully complete a mission. In the plotted mission, the target was within sight of at least one aircraft for the entire duration of the simulation. The complexity of the delegation of tasks to individual team members shows the utility of the risk-based cost function.

V. Discussion

Figure 5(a) shows that the mean time that the target is under surveillance is improved over the baseline case of a single aircraft watching the target for as long as possible before landing out. Cooperation allows the teams to observe the target for a greater percentage of the mission duration for all flock sizes and thermal densities. When results are examined, the target is continuously under observation as long as at least one of the original team members is still airborne. This implies that the failure to observe the target for the duration of the entire mission is due to the loss of every aircraft in the team.

The effect of an increased flock size is that it allows for the observation mission to be conducted in a sparser environment. Figure 5(a) shows that the performance of the three aircraft flock increased with increasing thermal density. The same trend is evident with a five aircraft flock until the density of lift sources in the environment rises above two per square kilometer. After this point, a denser thermal environment does not meaningfully impact the mission success rate. Similarly, with a seven aircraft team thermal densities over one thermal per square kilometer do not result in improved performance.

This shows that in a small team of aircraft, where exploration resources are limited, the team benefits from abundant lift sources. Less exploration is required to find the lift needed to prolong a mission. A large team is able to invest enough aircraft in exploration behavior to find the necessary lift, even when lift is sparse.

The survival rates of the aircraft displayed in Figure 5(b) are rather poor. These low values can be attributed to two sources. The first is that, in comparison to similar simulations in prior work,⁶ the density of thermals in the environment is low for the sparsest three scenarios. Even the most thermal-rich environment modeled roughly corresponds to the least dense situations encountered previously. The lack of lift sources will obviously have an effect on the endurance of a team of aircraft. Regardless, a valuable examination of the performance of the flock in the face of a poor soaring day is made. Note, however, that the duration of soaring flights reported by other researchers suggests that actual thermal density is significantly higher than that used in simulations here.^{2,4}

A systematic undervaluing of the contribution to the team's success of a single aircraft is also to blame for the high numbers of aircraft lost through the simulations. The value $P(D|C_n)$, or the probability that the mission will be failed if aircraft n lands out, is taken here to be $1/N$, inversely proportional to the number of aircraft remaining. This value is almost certainly too low in most situations. As a result, the risk to the mission incurred if an agent does not choose to explore is underestimated, discouraging this behavior until an aircraft is very low. A more thorough examination of the contribution of a single aircraft to the success of the team will doubtless improve the survival rates of the flock members.

The tools and the procedure developed in Section II are generally applicable to propagating aircraft energy state forward in time through an uncertain environment and computing the risk associated with a number of available actions. Risk is a quality associated with any mission. Using UAVs, particularly small ones as were examined in this work, over piloted aircraft allows for a higher acceptable level of risk. Unnecessary

risks should always be avoided. This research makes steps towards quantifying the approximated level of risk associated with mission tasks in an explicit manner. As long as the risk associated with performing a task (or not performing a task, as in this paper) can be quantified, heuristic task assignment can be avoided.

It is useful to have a method for propagating the aircraft energy state with consideration given to the stochastic environment in which it operates. This allows the mission planner to account for uncertainties in the knowledge of the environment and in aircraft performance parameters.

VI. Summary and Conclusions

The primary contribution of this work is to introduce a cost function to aid the assignment of behaviors to members of a flock of aircraft. In a small team, the task assignment problem may be straightforward. However, a method for differentiating between potential behaviors in a situation with a large flock size, a complex or uncertain environment, and strict mission goals is needed.

We approach this by assessing risk in a systematic manner. The team's mission is taken into account as well as the individual states of the member aircraft. Aircraft energy states are projected through the perceived environment. What results is a way for each team member to approximate the risk to the mission for neglecting each of a set of candidate tasks. With this knowledge, the task assignment problem can be solved in an informed manner.

The risk-based cost function was shown to be a valuable aid in deciding the appropriate task to be pursued by each aircraft engaged in an observation mission. Performance improvements were seen for larger flock sizes where the task assignment problem becomes more complex but the capabilities of the flock are correspondingly greater. This method is applicable to any situation where limited knowledge is available about the environment and decisions must be made regarding what action or behavior will limit the risk of mission failure.

VII. Acknowledgments

This research was conducted with Government support under and awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a and by the Office of Naval Research under Grant N000141110656.

References

- ¹Allen, M. J., "Autonomous Soaring for Improved Endurance of a Small Uninhabited Air Vehicle," *43rd AIAA Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics, Reno, Nevada, January 2005.
- ²Allen, M. J. and Lin, V., "Guidance and Control of an Autonomous Soaring Vehicle with Flight Test Results," *AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2007-867, American Institute of Aeronautics and Astronautics, Reno, Nevada, January 2007.
- ³Anderson, K. and Kaminer, I., "On Stability of a Thermal Centering Controller," *AIAA Guidance, Navigation, and Control Conference*, August 2009.
- ⁴Anderson, K., Kaminer, I., and Jones, K. D., "Autonomous Soaring; Flight Test Results of a Thermal Centering Controller," *AIAA Guidance, Navigation and Control Conference*, AIAA Paper 2010-8034, American Institute of Aeronautics and Astronautics, Toronto, Canada, August 2010.
- ⁵Bower, G. C., Flanzer, T. C., and Naiman, A. D., "Dynamic Environment Mapping for Autonomous Thermal Soaring," *AIAA Guidance, Navigation, and Control Conference, 2-5 August 2010, Toronto, Ontario Canada*, August 2010.
- ⁶Deppenbusch, N. T. and Langelaan, J. W., "Coordinated Mapping and Exploration for Autonomous Soaring," *AIAA Infotech@Aerospace Conference*, St. Louis, Missouri, March 29-31 2011.
- ⁷Goerzen, C. and Whalley, M., "Minimal risk motion planning: a new planner for autonomous UAVs in uncertain environments," *AHS International Specialists' Meeting on Unmanned Rotorcraft*, Tempe, AZ, January 25-27 2011.
- ⁸Munkres, J., "Algorithms for the Assignment and Transportation Problems," *Journal of the Society for Industrial and Applied Mathematics*, Vol. 5, No. 32, 1957.
- ⁹Bourgeois, F. and Lassalle, J. C., "An Extension of the Munkres Algorithm for the Assignment Problem to Rectangular Matrices," *Communications of the ACM*, Vol. 14, No. 12, 1971.
- ¹⁰Gedeon, J., "Dynamic Analysis of Dolphin Style Thermal Cross Country Flight," *Proceedings of the XIV OSTIV Congress, Organisation Scientifique et Technique Internationale du Vol à Voile*, 1974.
- ¹¹Qi, Y. C. and Zhao, Y. J., "Energy-Efficient Trajectories of Unmanned Aerial Vehicles Flying through Thermals," *Journal of Aerospace Engineering*, Vol. 18, No. 2, April 2005, pp. 84-92.
- ¹²Langelaan, J. W. and Bramesfeld, G., "Gust Energy Extraction for Mini- and Micro- Uninhabited Aerial Vehicles," *46th*

AIAA Aerospace Sciences Meeting and Exhibit, AIAA Paper 2008-0223, American Institute of Aeronautics and Astronautics, Reston, Virginia, January 2008.

¹³Klesh, A. T., Kabamba, P. T., and Girard, A. R., *Optimization and Cooperative Control Strategies*, Vol. 381, chap. Optimal Cooperative Thermalling of Unmanned Aerial Vehicles, Springer Berlin / Heidelberg, 2009, pp. 355–369.

¹⁴Allen, M. J., “Updraft Model for Development of Autonomous Soaring Uninhabited Air Vehicles,” *44th AIAA Aerospaces Meeting*, AIAA Paper 2006-1510, American Institute of Aeronautics and Astronautics, January 2006.

¹⁵Cheng, K. K., *Guided Exploration for Coordinated Autonomous Soaring Flight*, Master’s thesis, The Pennsylvania State University, University Park, PA 16802, August 2013.

¹⁶Thrun, S., Burgard, W., and Fox, D., *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*, The MIT Press, 2006.

¹⁷Siddon, R., “Fast calculation of the exact radiological path for a three-dimensional CT array.” *Medical Physics*, Vol. 12, No. 2, 1985, pp. 252.

¹⁸Chakrabarty, A. and Langelaan, J. W., “Flight Path Planning for UAV Atmospheric Energy Harvesting using Heuristic Search,” *AIAA Guidance, Navigation and Controls Conference*, Toronto, Canada, August 2010.