

Coordinated Mapping and Exploration for Autonomous Soaring

Nathan T. Depenbusch*

Jack W. Langelaan†

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

To enable coordinated soaring by a flock of small unmanned air vehicles a method for distributed mapping of the wind field is presented. Presently the work focuses on mapping convective activity (i.e. thermals). The map first discretizes the environment into cells and then uses a Kalman filter to estimate the vertical wind speed and associated covariance in each cell. To improve computational tractability, the wind speed in each cell is assumed to be uncorrelated to all other cells; this results in a set of independent scalar Kalman filters. Measurements of wind speed are available at the location of an aircraft; thermal dynamics are modeled using an exponential decay. The resulting map is combined with a behavior-based controller to enable autonomous soaring. The covariance of wind speed is used to drive exploration and a combination of estimated wind speed and covariance is used to drive exploitation. The utility of the approach is demonstrated using Monte Carlo simulations of a persistent presence task: a flock of UAVs flies in a four square kilometer region and attempts to maximize endurance. Only gliding flight is assumed; with one aircraft the use of the map doubles endurance compared with a no-map case. Increasing flock size to two, four, and eight aircraft results in monotonically increasing performance, with almost all of the eight-aircraft flocks able to remain aloft for the full mission duration.

I. Introduction

Small uninhabited aerial vehicles (UAVs) are commonly used for low altitude reconnaissance or surveillance missions. Typically the endurance of such aircraft (which are frequently hand-launchable and have wingspan ranging from 1 to 4 meters and mass ranging from 1 to 10kg) is about an hour, making missions such as persistent surveillance difficult, if not impossible.

However, the endurance of small UAVs can be significantly improved by exploiting energy available in the environment. Atmospheric energy can be harvested from three phenomena: (a) vertical air motion; (b) spatial wind gradients; (c) temporal gradients (gusts). Energy harvesting from spatial gradients is known as *dynamic soaring*; energy exploitation from vertical air motion has become known as *static soaring* since the time scale of vertical air motion is long compared with vehicle dynamics.

Large birds and human sailplane pilots routinely exploit vertical air motion (lift) to remain aloft for several hours and fly hundreds of kilometers without flapping wings or the use of engines. Exploiting these long duration vertical air motions has been an active area of research for manned glider flight for many years and is now becoming more active for small UAVs as well.¹ Recent research in autonomous soaring has shown that enormous improvement in range and endurance is possible.^{2,3} Further, the stealthiness provided by silent, gliding flight can greatly improve covert operation.

Given the improvement in performance which results from soaring by individual vehicles, it is reasonable to suggest that significant performance improvement is likely when *teams* of soaring-capable UAVs are deployed. In fact, group behavior has been observed in both natural and human soaring flight, with significant improvements in performance.

In 1988 Bednarz reported that Harris hawks (*Parabuteo unicinctus*) hunt in groups of 2 to 6 to capture prey.⁴ Several tactics are described, ranging from shared search (individuals take turns searching for prey

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

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

or spread out to cover more area) to coordinated attacks. In many cases a group of hawks can take prey that are several times larger than an individual bird. Significant improvements in hunting success rates are reported, and he suggests that individual benefits (in terms of net energy intake available per individual) are maximized in a group of approximately 5 birds. Other examples of coordinated hunting by raptors have been reported for Bald Eagles,⁵ Saker Falcons,⁶ Peregrine Falcons,⁷ Lanner Falcons,⁸ and Marsh Harriers;⁹ coordinated hunting is also reported for other birds such as Loggerhead Shrikes¹⁰ and Scrub Jays.¹¹

Both structured and unstructured coordinated soaring has occurred in human sailplane competitions.¹² Structured (i.e. team) soaring in human glider competitions has been found to improve performance significantly.¹³ Tactics generally consist of having team members act as thermal foragers for each other, with formation flight during cruise also a possibility. An unstructured form of coordinated soaring in humans is known as a gaggle, where groups ranging in size from 5 to 50 sailplanes fly together. The key difference between a gaggle and a team is cooperation: significant radio communication occurs in a team, while there is no explicit cooperation in a gaggle. In earlier competitions (1960s) gaggles typically formed in weak conditions and consisted of less-skilled pilots following skilled pilots (a practice called leeching). Flight performance of less-skilled pilots was considerably improved. In the mid 1980s small gaggles of skilled pilots showed good performance in all conditions. The number of pilots involved soon grew, and large, fast gaggles became the norm. Several attempts at limiting team soaring and the gaggle flying through new rules have been made, but so far they have not proven successful: this illustrates the competitive advantage that results from coordinated soaring.

The broad problem considered here is coordinated thermal soaring by a team (or flock) of small UAVs. If the wind field (including vertical components caused by thermals) is known *a priori* then this can be used by the vehicles to increase endurance. Unfortunately, atmospheric convection is still a poorly understood phenomenon, with only coarse, probabilistic predictions available. Tools such as BLIPMAPS generate maps of likely thermal activity at a resolution of approximately two kilometers,¹⁴ but this is too coarse for use by small UAVs.

Hence the focus of this paper is on coordinated mapping of the atmosphere by a flock of small, soaring capable UAVs. The problem of mapping by robots has been well studied (see for example Thrun¹⁵ and references therein), but there are complications when applying techniques directly to atmospheric mapping by small UAVs. First, sensing is only available at the vehicle location (i.e. long-range sensors of vertical air motion are not available). Second, the atmosphere is dynamic. Third, most mapping problems involve solid obstacles, rather than regions to be exploited or avoided.

The approach presented here is inspired by the occupancy grid (used in the mobile robot community for obstacle avoidance), but there are critical differences. An occupancy grid is a numerical implementation of a Bayesian estimator adapted to compute a binary estimation problem (whether a region of the environment, or cell, is occupied or free).¹⁶ These have been used with great success in robot mapping problems.¹⁵ Jakuba and Yoerger¹⁷ utilize an occupancy grid mapping method in the detection and location of hydrothermal vents on a sea floor by an Autonomous Underwater Vehicle. The algorithm starts with no knowledge of the number of vents or vent locations and attempts to find every vent in the search space. Dynamic environments (for example, moving targets) can also be modeled by incorporating “forgetfulness” in the map: an example is described by Lum et al.¹⁸

However, an occupancy grid is a binary estimator. A useful wind map must include both an estimate of the magnitude of the wind and the uncertainty in the estimate. Bower et al. present an occupancy grid inspired approach to thermal mapping for a single aircraft but do not address uncertainty in either measurements or the map.¹⁹ The research presented in this paper: (1) tracks uncertainty in the expected vertical component of wind in addition to the magnitude; (2) is scalable to flocks of aircraft.

It is assumed that each vehicle is capable of autonomous flight and that each vehicle can obtain a measurement of wind speed (e.g. through a measurement of GPS velocity, airspeed and an estimate of orientation). This map is continuously updated as the vehicles explore the environment and includes a model for thermal lifetime. Simulations are used to show the utility of the map for flocks ranging from one to eight aircraft. While fielded soaring-capable UAVs are likely to be equipped with motors (enabling improved foraging for thermals) here only gliding flight is considered. An aircraft is thus considered “down” and unrecoverable when its altitude drops to zero.

The remainder of this paper discusses the atmospheric model (Section II); describes the mapping algorithm (Section III); defines soaring and exploration behavior (Section IV); discusses results of Monte Carlo simulations which show the utility of the approach (Section V); and presents concluding remarks (Section VI).

II. Atmospheric Model

The convective boundary layer is modeled here as noiseless and characterized only by the presence of a number of vertically rising thermals. In this simulation, no lateral winds are modeled and the ground plane is flat and located at $h = 0$. Thermals are equally likely to appear at any point on the world map.

A. Thermals

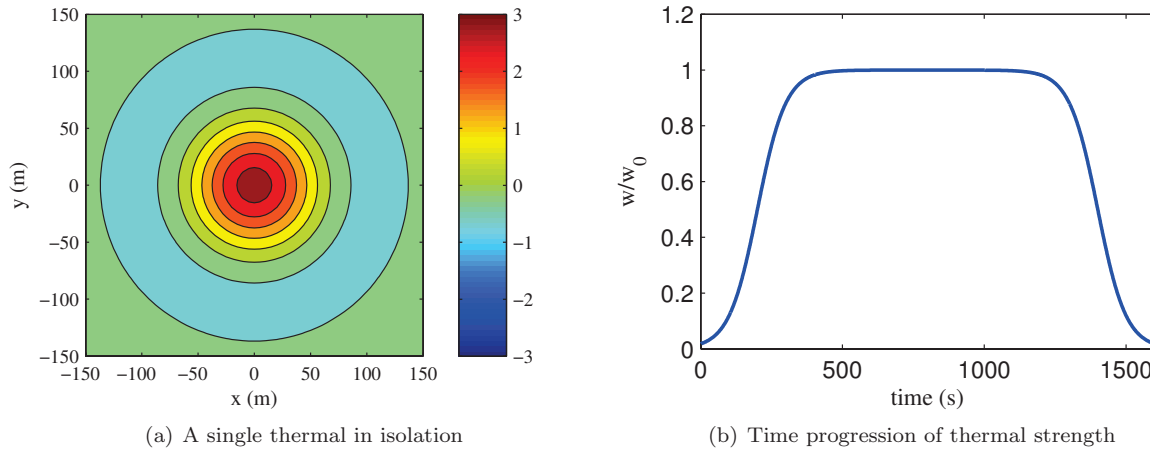


Figure 1. Thermal strength progression with time. The thermal displayed has $R = 50\text{m}$, $w_0 = -3\text{m/s}$ and $(x_0, y_0) = (0, 0)\text{m}$. The thermal period is 1200s and t_0 is 800s. Note that the thermal strength decays on the same time scale as it builds.

The vertical wind speeds induced by a single thermal are based on a model first presented in Gedeon²⁰ and employed in other research:²¹⁻²³

$$w_z(x, y, t_0) = w_0 e^{-\left(\frac{\sqrt{(x-x_0)^2 + (y-y_0)^2}}{R}\right)^2} \left[1 - \left(\frac{\sqrt{(x-x_0)^2 + (y-y_0)^2}}{R}\right)^2 \right] \quad (1)$$

where w_0 is the maximum vertical wind speed, x_0 and y_0 are the coordinates of the thermal center, and R is the thermal radius. The thermal strength, shape, and radius are assumed constant from the ground plane to the top of the convective boundary layer (here set at $z_i = 1500\text{ m}$).

The time-varying quality of the thermal is accounted for by allowing vertical wind speed to increase with a sigmoid function to full strength at $t = t_0$ and to decay similarly afterwards. The period of a particular thermal is given by T .

$$\frac{w_z(x, y, t)}{w_z(x, y, t_0)} = \left(\frac{1}{e^{\eta(t - (t_0 + \frac{1}{2}T))} + 1} + \frac{1}{e^{\eta((t_0 - \frac{1}{2}T) - t)} + 1} - 1 \right) \quad (2)$$

The η term is used to scale the rise time of the thermal. In this work it is taken to be 0.02, yielding a rise time of 180 seconds. This parameter may be easily scaled to simulate thermals with differing rise and decay rates. An example of the time progression for a single thermal is given in Figure 1(b).

The advantage of the Gedeon thermal model is that it yields an area of strongly upwards moving wind at the thermal center, surrounded by a larger, generally downwards moving region of air. This approximates actual thermal behavior that is often observed by glider pilots. The radius of the region affected by the thermal is approximately 3 times that of the upwards moving core of the thermal. The thermal is still a strictly local phenomena where the presence of a thermal has a negligible effect on the wind speed far away from it. This is in contrast to some thermal models which require the "atmospheric sink" cancel the vertical speed of all of the thermals.²⁴ This makes the thermal a global phenomena, where a single thermal affects the wind speed over the entire map.

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

parameter	symbol	mean	σ	max/min	units
Period	$T_{thermal}$	20	7.5	60 / 5	minutes
Radius	R	75	20	200 / 30	m
Core strength	w_0	-2.56	1.5	-7 / 1	m/s

B. World

The world is defined by the sum of the wind speeds contributed by each individual thermal. A maximum number of thermals for a given area (A), boundary layer thickness (z_i), and average radius (\bar{R}) is given by Allen²⁴ as:

$$N = \frac{0.6A}{z_i \bar{R}} \quad (3)$$

In this work thermal generation is a random process, where the probability of generating a thermal is inversely proportional to the number of thermals present with the maximum number of thermals bounded by N . At every time interval Δt_{gen} , the number of active thermals, n , is determined, and new thermals are generated with probability P given by:¹⁹

$$P = \min\left(\frac{\Delta t_{gen} N}{T} (N - n), 1\right) \quad (4)$$

A minimum separation distance between thermals of $3R$ is enforced such that a realistic environment can be approximated. Each thermal is generated with random parameters given in Table 1 and derived from Allen.²⁴ A sample wind field is shown in Figure 2. Note that in this simulation, the majority of the environment has no vertical wind component: this makes soaring somewhat more difficult since the likelihood of encountering a thermal is rather low, but the utility of a map and a coordinated flock will become apparent.

III. Mapping

Planning algorithms typically assume that knowledge of the environment is available. For soaring flight the available knowledge is both incomplete and uncertain: meteorological data can provide information about expected mean wind speeds and about the likelihood of thermal activity, but it cannot give specific information such as high-density real time wind information or precise locations of local phenomena such as thermals. Knowledge of updraft locations (whether thermally or orographically induced) would greatly enhance the capability to exploit environmental energy during a mission.

It is assumed that an aircraft can measure local wind conditions (i.e. three components of wind velocity at the aircraft's location), and this measurement is corrupted by zero-mean Gaussian noise. These measurements are then used to create a map of wind. The wind map consists of estimates of vertical wind speed and the level of confidence in the estimate. The confidence level has two purposes: it can drive exploration, so that the vehicle will fly towards regions of low confidence to improve its knowledge; and it can drive the choice of flight direction when the vehicle needs a thermal to gain altitude.

The environment is divided into $M \times N$ uniform two-dimensional cells that cover the space. The size of a cell roughly corresponds to the turn radius of the aircraft and the wind in each cell is assumed to be constant. The problem now is to compute an estimate $\hat{w}_{z,ij}$ of the vertical wind speed in each cell, and the associated error covariance P_{ij} . The covariance is a measure of the confidence that can be placed on

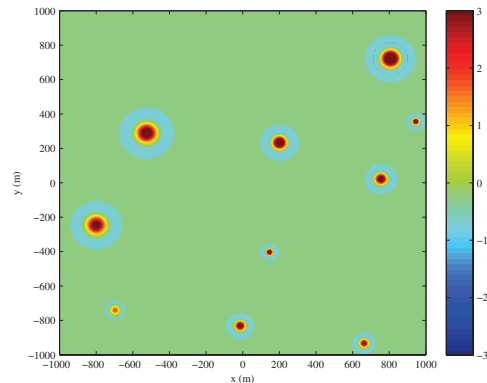


Figure 2. A sample map of the wind speeds over the domain simulated. The total wind speed at any point is the sum of the contributions of each individual thermal.

the wind speed prediction in a particular cell. If the covariance is large, it indicates that there is a large uncertainty in the wind speed prediction. If on the other hand, P_{ij} is small, the wind estimate is assumed to be accurate.

The estimation problem is cast as a set of $M \times N$ 1D Kalman filters: it is thus implicitly assumed that the vertical wind speed in a cell is independent of surrounding cells. Depending on cell size this is not necessarily a good assumption (a thermal may be larger than a cell), but it greatly simplifies the estimation process. This approximation should be conservative: ignoring correlation between cells effectively ignores information available in the environment. In this paper it is further assumed that no *a priori* information is available. In practice knowledge of terrain and sun angle (and potential for orographic lift) will provide *a priori* knowledge; this can be implemented directly in the filter if it is available.

A. Filter Implementation

In Section II it was assumed that thermal strength varies with time according to a sigmoid function. For the purpose of the Kalman filter an exponential decay is assumed, so that

$$w_{z,ij,k+1} = aw_{z,ij,k} + \mathcal{N}(0, Q) \quad (5)$$

Here a is chosen so that w_z decays to 20% of its original value within one thermal lifespan \bar{T} (see Figure 3(a)) and $\mathcal{N}(0, Q)$ represents zero-mean Gaussian process noise. The prediction step of the ij^{th} vertical wind speed estimate is thus

$$\hat{w}_{z,ij,k|k-1} = a\hat{w}_{z,ij,k-1|k-1} \quad (6)$$

$$P_{ij,k|k-1} = a^2P_{ij,k-1|k-1} + Q \quad (7)$$

The process noise Q is chosen so the 1σ uncertainty is 4 m/s after one thermal lifetime (Figure 3(b)). This corresponds roughly to the largest expected thermal strength; this can be changed depending on performance.

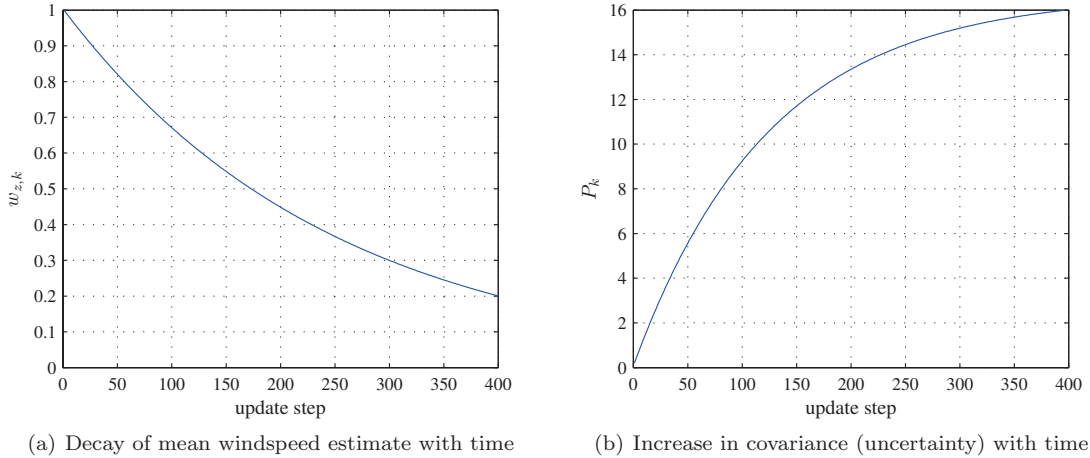


Figure 3. The dynamics of the linear Kalman filter applied to a cell. Step size is $\Delta t_{map} = 3$ seconds.

A measurement z_{ij} of the wind in a cell is made while an aircraft is in the ij^{th} cell:

$$z_{ij,k} = w_{z,ij,k} + \mathcal{N}(0, R) \quad (8)$$

The measurement noise has standard deviation 0.2 m/s.

The measurement update follows the standard linear Kalman filter update:²⁵

$$P_{ij,k|k} = \left(P_{ij,k|k-1}^{-1} + R^{-1} \right)^{-1} \quad (9)$$

$$K_{ij,k} = P_{ij,k|k} R^{-1} \quad (10)$$

$$\hat{w}_{z,ij,k|k} = \hat{w}_{z,ij,k|k-1} + K_{ij,k} (z_{ij,k} - \hat{w}_{z,ij,k|k-1}) \quad (11)$$

This measurement update is also a scalar update, which can be performed extremely quickly.

In addition to the cell which contains an aircraft, a measurement update is also performed for that cell's immediate neighbors. Since a thermal is larger than one cell size, it is likely that vertical wind speed in nearby cells will be similar; to reflect the increased uncertainty of the wind speed in neighboring cells the measurement noise covariance increases linearly from $\sigma = 0.2$ m/s at the aircraft location to $\sigma = 4$ m/s one mean thermal radius away.

As can be seen in Figure 3(b) uncertainty in thermal strength increases sharply without available measurements. The expected wind speed in a cell decreases as well (Figure 3(a)). The uncertainty will be used by the controller to ensure that a cell is visited so that a measurement can be obtained.

B. Aircraft Model



Figure 4. Coordinate frames. Left: top view. Positive directions of vectors shown. ψ is positive from the positive N axis. Right: side view. Positive directions of vectors shown. γ shown is negative.

Gliding flight is modeled here. Vehicle parameters are given in the Appendix. It is assumed that an on-board autopilot is able to follow turn rate and airspeed commands as well as maintain a trimmed flight configuration, thus a kinematic model is sufficient to describe aircraft motion. Standard autopilot instrumentation (GPS, airspeed, acceleration, angular rate, magnetometer) are assumed to be available for locating the aircraft and taking measurements.

Aircraft positions are recorded in a standard NED coordinate system. Position is assumed to be perfectly known at all times. Heading angles are measured from the North axis as shown in Figure 4(a). Airspeed is specified along the glidepath of the aircraft, at angle γ from the horizontal (Figure 4(b)). The D axis is positive downwards, thus wind speed is similarly positive downwards. The resulting state equations are:

$$v_N = v_a \cos \psi \cos \gamma + w_N \quad (12)$$

$$v_E = v_a \sin \psi \cos \gamma + w_E \quad (13)$$

$$v_D = v_a \sin \gamma + w_D \quad (14)$$

where $[w_N \ w_E \ w_D]^T$ is the wind speed vector at the aircraft's location.

Control of the aircraft is enacted through the input vector u , where:

$$u = [v_a \ \dot{\psi}]^T \quad (15)$$

The methodology for determining control inputs is addressed in Section IV.

The aircraft's state is now defined as $[P_N \ P_E \ P_D \ v_a \ \psi \ \gamma]^T$ where glidepath angle (γ) is given as a function of C_L and C_D :

$$\gamma = \arctan \frac{C_D}{C_L} \quad (16)$$

The lift coefficient, C_L is determined from airspeed assuming steady-state flight, and C_D is given as a function of C_L . A final value is appended to the state vector to indicate if the aircraft is currently active.

For obvious reasons, if the aircraft's altitude drops below the ground plane at any time during simulation, this value is permanently set to inactive.

The kinematic model of the aircraft is propagated through time by fourth-order Runge-Kutta integration with a time step of 0.02 s. Low-level control inputs (the determination of the vector u) are computed at each time step, though measurements, behavior determination, and mapping are done at different rates.

C. Aircraft Energy

Aircraft specific energy is given here as the sum of potential and kinetic energy normalized by aircraft weight. The second derivative of aircraft specific energy will become important in the thermal centering controller presented in Section IV

$$e_{tot} = h + \frac{1}{2g}v_a^2 \quad (17)$$

$$\dot{e}_{tot} = \dot{h} + \frac{1}{g}\dot{v}_av_a \quad (18)$$

$$\ddot{e}_{tot} = \ddot{h} + \frac{1}{g}(\ddot{v}_av_a + \dot{v}_a^2) \quad (19)$$

No consideration is made for stored onboard energy because of the gliding assumption, however, it is assumed that all aircraft have sufficient onboard energy to actuate control surfaces to follow commands for the duration of each simulated mission

D. Multiple Agent Systems

In a multiple agent system, each aircraft is modeled independently. Measurements are shared between aircraft at a time interval specified by Δt_{map} . Though not specified by the controller, this results in aircraft that typically avoid each other when exploring, as each aircraft seeks to explore the area containing the least information. Also, when seeking a thermal to gain altitude, a single aircraft in a multiple agent system benefits greatly from map information gained by the other aircraft.

One can see this expanded to a multiple agent surveillance strategy, where the aircraft in the system are simultaneously mapping the environment and observing a target. A single aircraft may then be kept in sight of the target at all times, switching out with another aircraft as it loses altitude. The original aircraft can then find the nearest thermal, gain altitude, and return to its surveillance duties when the currently surveying aircraft must switch out.

E. Map generation and exploration

The method of aircraft control and mapping is given here. The steps that are followed to generate the local map are specified below. Note that each step is not necessarily evaluated at the same rate.

1. Update behavior of each aircraft in the world map (high level control):
 - A decision maker selects a behavior for the aircraft based on its energy state, current environment measurements, as well as its location and the information contained in the shared map.
 - If necessary for the behavior assigned, a waypoint on the map is identified.
2. Update aircraft state and determine control inputs (low level control):
 - Control inputs are determined based on the assigned behavior:
 - If a waypoint was assigned, the turn rate required to track a waypoint is determined
 - If the aircraft is assigned to thermal, the turn rate required to remain about the center of the thermal is determined.
 - The correct airspeed to fly is determined
 - The aircraft's state is updated through integration of the kinematic model.
 - The cell that the aircraft is currently in is determined, and distances are recorded to the centers of adjacent cells.

- Measurements of the environment (wind speeds) are taken at the aircraft’s current location.
3. Update the shared map
- \hat{w}_z and P are updated through system dynamics
 - Each aircraft’s measurements \mathbf{z} are incorporated by updating \hat{w}_z and P in the affected cells.

The high level control in the first step is complex and will be detailed in Section IV. It should be clear that an important part of the mapping process is deciding, based on the known wind map and vehicle state, what action should be undertaken at the present time step. Exploration is the goal of each vehicle, however, if the aircraft has passed below a threshold altitude, the aircraft must gain altitude in order to continue its mission.

The low level control step requires that a controller be able to follow the commands specified in the high level control step. It is here that the aircraft state is updated in simulation, however, in hardware implementation it would be this controller that gives commands to the autopilot. The specifics of the low level controllers are given in Section IV. The second step is also where measurements are taken.

The third step is mapping of the measurements whereby the the measurements of each aircraft are combined and knowledge of the environmental dynamics is utilized to produce a useful map of the region of interest.

IV. Aircraft Control and Behavior Determination




Figure 5. A diagram showing behaviors and the transfers between them. Details of each behavior are given in Table 2

With a mapping method in place, a control methodology is now required to determine the aircraft behaviors to best keep the map updated. Several key behaviors and the switches between the behaviors are programmed and shown in Figure 5. The direction of the arrows indicates in which direction a transition is possible. The four behaviors used here are detailed below and summarized in Table 2.

Decisions are made at an interval Δt_{plan} . The order of precedence of the behaviors is shown in Table 2, where the bottommost behavior takes top priority and the top behavior is only followed if none of the other behaviors are triggered.

A. Exploration Behavior

The goal of an aircraft assigned an exploration behavior is to minimize the uncertainty (covariance values) of the wind estimates on the map. No claim is made of the optimality of the presented behaviors, they do however function well in simulation.

1. Local Exploration

When exploring locally, an aircraft attempts to minimize the uncertainty in the wind measurements within a small region around that aircraft. This is the default behavior of each aircraft, and leads to small, well

Table 2. Summary of the behaviors and the switches between them as employed in this simulation.

Behavior	Trigger	Prerequisite	Description
Local Exploration	(none)	Default behavior. Local covariance must be above threshold.	Aircraft explores local cells with high covariance
Global Exploration	$mean(\Sigma_{local}) \geq \Sigma_{thresh}$	Local covariance must be below threshold.	Aircraft flies to the center of the quadrant with the highest covariance.
Cruise to	$e < e_{thresh}$	Aircraft's energy is below a threshold.	Prerequisite to thermalling behavior. Aircraft flies towards region of substantial lift.
Thermal	$w_z \geq v_{MacCready}(h)$	Aircraft must be located in a region with substantial atmospheric lift.	Aircraft circles in an attempt to center the thermal and gain altitude

mapped regions in the local map. A local region is defined by Δi_{local} and Δj_{local} , a small number of cell indexes above and below the aircraft's current cell location $(i_{current}, j_{current})$ in the N and E directions. The target cell for exploration then becomes:

$$Target_{ij} = \max(\Sigma_{i^*, j^*}) \text{ for } \begin{aligned} i^* &\in [i_{current} - \Delta i_{local}, i_{current} + \Delta i_{local}], \\ j^* &\in [j_{current} - \Delta j_{local}, j_{current} + \Delta j_{local}] \end{aligned} \quad (20)$$

The local exploration behavior is the default for the simulated aircraft. An aircraft checks the covariance of the cells within three indexes of its current location ($\Delta i_{local} = \Delta j_{local} = 3$) as shown in Figure 6(a). The cell with the highest covariance becomes the aircraft's target, and a waypoint is placed in its center. The size of Δi_{local} and Δj_{local} , the number of indexes over which to look, must be large enough that an aircraft has impetus to cross the path of an aircraft that has recently mapped an area.

2. Global Exploration

If the mean covariance of the cells within an aircraft's local exploration area is below a set threshold, the aircraft is assigned to explore a much wider region. The global exploration behavior is thusly triggered when the wind is known with a high degree of confidence in the vicinity of the aircraft. The map within $\Delta i_{global} = \Delta j_{global} = 30$ indexes of the aircraft is divided into quadrants (Figure 6(b)), and the mean covariance in each quadrant is determined, $\Sigma_{mean, Q}$. The aircraft is then directed to fly towards the center of the quadrant with the highest covariance and consequently in which the wind measurements are the least certain. This behavior is primarily used to traverse the map to explore large unmapped regions.

$$Target_Q = \max(\Sigma_{mean, Qk}) \text{ for } k \in [1, 4] \quad (21)$$

where

$$\begin{aligned} Q1 &= i \in [i_{current} - \Delta i_{global}, i_{current}], j \in [j_{current}, j_{current} + \Delta j_{global}] \\ Q2 &= i \in [i_{current} - \Delta i_{global}, i_{current}], j \in [j_{current} - \Delta j_{global}, j_{current}] \\ Q3 &= i \in [i_{current}, i_{current} + \Delta i_{global}], j \in [j_{current} - \Delta j_{global}, j_{current}] \\ Q4 &= i \in [i_{current}, i_{current} + \Delta i_{global}], j \in [j_{current}, j_{current} + \Delta j_{global}] \end{aligned} \quad (22)$$

B. Exploitation Behavior

In order to utilize energy in the atmosphere to extend endurance of a gliding aircraft two things must happen: a thermal of adequate strength must be located, after which an aircraft must maneuver such that it can take advantage of the thermal lift. The first step is accomplished in this research by utilizing the generated

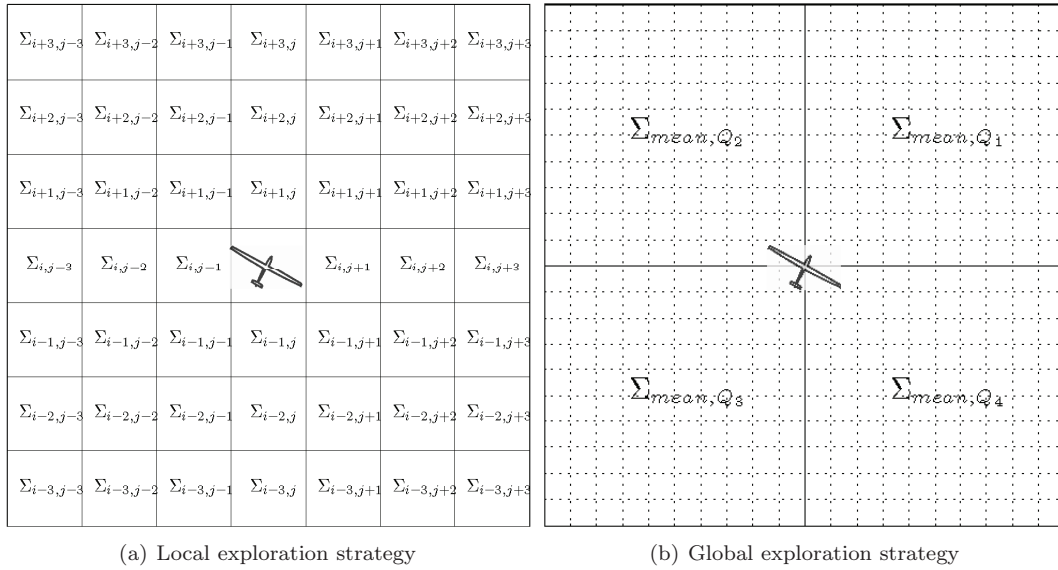


Figure 6. Two strategies for exploring the environment are presented here. In the local exploration behavior, a waypoint is set at the center of the local cell with the highest uncertainty. In the global exploration behavior, a waypoint is set at the center of the quadrant with the highest uncertainty.

wind map and will be discussed in subsection 1 and subsection 2. The second step is accomplished through the application of an appropriate thermal centering controller and is detailed in subsection 3. Entering and exiting a thermal is not addressed, as it is sufficient to follow the rules in Table 2 to accomplish these tasks.

Exploitation behavior is triggered in two ways. If an aircraft happens across a region where the vertical wind speed is greater than its current MacCready value, the aircraft will stop to take advantage of the thermal. Alternatively, if an aircraft passes below a threshold altitude, it becomes necessary to exploit the environment to regain energy.

1. A Thermal is Found Inadvertently

The MacCready value, or “MacCready setting” is important to competition glider pilots. Originally proposed by Paul MacCready²⁶ and further described by Reichmann,²⁷ much of the examination of this value has been done with the goal of improving performance in soaring competitions. Classically, the MacCready value is set to the expected rate of climb in the next thermal on the course. However, in our current problem of a small domain and available thermal strength estimates, the MacCready value corresponds to the minimum strength thermal that a glider should take if found. This value varies with altitude and the shape of the function is determined for an expected distribution of thermal strengths.

The MacCready number encourages an aircraft not to settle for a weak thermal when it is likely to find a stronger thermal if it continues to search. However, if a strong thermal is encountered, the aircraft is encouraged to take advantage of it because it is unlikely a stronger one will be found. The inherent uncertainty is addressed in a comprehensive paper by Cochrane,²⁸ and his method is adapted to the current situation. Because the goal is endurance, it is assumed that the aircraft has indefinitely




Figure 7. The MacCready value given as a function of altitude for the distribution of thermal strengths modeled here. The blue line shows the MacCready value as determined by Cochrane’s method, the green line shows the function fit used to approximate the MacCready value.

far yet to travel “on course,” and all calculations regarding a final glide are ignored.

The calculated MacCready value function for the modeled aircraft polars and thermal strength distribution is shown in Figure 7. At every time step Δt_{plan} , each aircraft is assigned a MacCready value based on its altitude. If the aircraft registers local lift greater than its assigned MacCready value, the aircraft is directed to take advantage of that thermal and begins the centering process.

If an aircraft is engaged in any behavior except thermalling and encounters a region in which w_D is greater than the MacCready value of that aircraft given its current altitude, the aircraft will take the encountered thermal. Similarly, if an aircraft is thermalling and the current thermal strength drops below the aircraft’s present MacCready value, the aircraft is directed to exit the thermal and seek a stronger one elsewhere. This limits the time an aircraft spends in weak thermals.

2. A Thermal Must be Sought

If the “cruise to” behavior given in Table 2 is triggered, the aircraft is directed to find a nearby thermal of adequate strength to gain altitude. In this process, the map is used rather than built. The region around the aircraft within gliding range is checked for the likely presence of thermals. Gliding range is approximately determined by:

$$Range = \frac{P_D - h_{min}}{\tan \gamma} \quad (23)$$

If a cell that likely contains a thermal is identified within this range, the distance from the aircraft to that cell’s center (d) is computed. The measure of attractiveness of that cell, $F_{i,j}$, is determined as:

$$F_{i,j} = \frac{\mu_{i,j}}{d} \quad (24)$$

and the cell with the greatest value of $F_{i,j}$ becomes the aircraft’s destination. If a sufficiently strong thermal (determined by the aircraft’s MacCready value [Figure 7]) is encountered en-transit, the aircraft will take the new thermal instead as specified in Section 1.

3. Thermal Centering Control

Though not central to the success of the current research, a method for thermalling is required for completeness of the system model. The thermal centering controller employed here was proposed by Andersson and Kammer.^{29,30} The controller returns a commanded turn rate $\dot{\psi}_c$ for a specified steady-state turn radius r .

$$\dot{\psi}_c = \frac{1}{r} v_a - k \ddot{e}_{cv_a} \quad (25)$$

The coefficient k is a gain used to scale the importance of energy change to the commanded turn rate.

Because a change in airspeed while attempting to locate a thermal center will give a false reading of the energy gained from that thermal, \ddot{e}_{cv_a} (the value of Equation 19 when airspeed is held constant) is used.

$$\ddot{e}_{cv_a} = \ddot{h} \quad (26)$$

Because of the constant airspeed requirement:

$$\dot{h} = -v_D = -w_D \quad (27)$$

thus the second derivative of aircraft specific energy is equivalent to the change of the wind speed in the D direction.

$$\ddot{e}_{cv_a} = -\dot{w}_D \quad (28)$$




Figure 8. The aircraft’s final position after 120 seconds of simulation is circled. The flight path for previous time steps is given by the blue line. A thermal is located at [0,0].

Values for r and k are taken from Andersson and Kammer as 30 meters and 25 respectively.

Though developed for a substantially differently shaped thermal model,²⁹ the controller is easily deployed in the present simulations. The result is an effective thermal centering strategy shown in Figure 8. The aircraft tightens its circle around the thermal center as it attempts to track the prescribed 30 meter steady-state turn radius. Advantageously, no complex calculations such as determining an induced roll rate are required for the implementation of this controller.

C. Low-level Control

The low-level controller is in place to follow the commands specified by the high-level controller. This means tracking specified waypoints or the turn rate required to center a thermal (Equation 25). The airspeed is also controlled such that the aircraft flies at the most efficient speed for a given behavior.

1. Turn Rate Control

With a specified waypoint as a destination, the turn rate is determined to be the difference between the aircraft's heading (ψ) and the heading required to reach the waypoint (ψ_{wp}).

$$\dot{\psi}_c = \psi_{wp} - \psi \quad (29)$$

This difference is limited such that the load factor caused by the turning maneuver is no greater than 2 (a 60° bank angle) yielding a $\dot{\psi}_{max}$ for a given airspeed. This limit is in place to preserve the assumption that sink rate does not vary with bank angle. At every integration time step, this difference is computed and the turn rate is output to correct the aircraft's heading in order to remain on course.

If an aircraft is directed not to follow a waypoint but rather to center a thermal, the turn rate returned is that given by Equation 25. This turn rate is also limited in the same way described above.

2. Airspeed Control

When in a thermal, the airspeed to be flown is that which yields the minimum sink rate for the aircraft. This results in the greatest climb rate,³¹ and because the aircraft is not traversing the world, no penalty is paid for the relatively slow speed. This airspeed is determined as the maximum of the sink-rate polar of the modeled aircraft shown in Figure 12(a). Typically the airspeed for minimum sink will depend on bank angle, but because shallow bank angles are being assumed in the kinematics model, a single value is determined here.

The MacCready value may also be used to determine the optimum airspeed to fly to achieve the fastest cross-country cruise in glider competitions. Because the goal here is exploration, it is not clear that the MacCready value gives the optimal airspeed to fly while exploring the map. Indeed, early simulation results show quantitatively degraded performance when the MacCready speed to fly is used rather than maintaining the airspeed for best L/D. Thus, when cruising between thermals or exploring the world, the airspeed flown is that for best L/D given the measured local vertical wind speed. This airspeed to fly, v_{stf} , is given by Reichmann²⁷ as:

$$v_{stf} = \sqrt{\frac{c + w_z}{a}} \quad (30)$$

where a and c are coefficients of a quadratic fit of the aircraft's sink-rate polar and are given in the Appendix.

V. Results and Discussion

Monte Carlo simulations were conducted for five different cases. The first case (the baseline) consists of a single glider which does not generate a map. Subsequent cases consisted of flocks of 1, 2, 4, and 8 aircraft using a map generated during the flight. For each case 50 runs were conducted, with the same wind field used for each case and different wind fields used for each run.

For all cases the objective was to maximize time aloft. A mission is deemed to end when the last aircraft touches down or the simulation reaches maximum time; here maximum time is ten hours.

Each run is initialized with aircraft located at the center of the environment at an altitude of 1500 meters (defined as, z_i , the height of the convective boundary layer).

A. The Baseline Scenario

The baseline case involved a single aircraft attempting to stay aloft as long as possible without generating a map. In this scenario, the only way for the aircraft to gain altitude is to randomly happen across a thermal. No memory of previously encountered thermals, nor knowledge of where thermals are likely to exist is available. This roughly represents the actions of model sailplane pilots searching for thermal lift on a “blue” day, i.e. a day where there are no visual indications of thermal activity.

In this baseline scenario the aircraft chooses a random waypoint in the map. If wind is encountered of greater magnitude than the aircraft’s current MacCready value, the aircraft takes that thermal. If not, a new waypoint is selected. This amounts to a random search of the flight domain.

The single aircraft stayed aloft for an average of 49 minutes over 50 simulated missions. An aircraft in a steady glide at best the airspeed for best L/D from an altitude of 1500 m will stay aloft for 45 minutes. This implies that though there is an advantage to be gained from taking thermals if they are available, the lack of a map means that a random search is an ineffective strategy for finding thermals. The majority of aircraft do not encounter any thermals during the course of their mission.

These results suggest that the wind field as used here is actually conservative: flight test results reported by Allen,² Edwards,³ and Andersson³⁰ have all shown that significantly increased endurance is possible, and that thermals are actually encountered more frequently than simulated here. The results presented here are thus a conservative representation of performance, and it will be shown that even in this set of scenarios significant improvement is possible using multiple aircraft.

In every simulated wind field, the baseline case performed worse than all of the other aircraft systems.

B. One Aircraft

A single aircraft is simulated exploring and mapping the environment. With one aircraft, the generated map acts as a form of memory for the aircraft rather than a means of communicating information. Because the domain is large relative to the area observable by the aircraft, regions must be revisited frequently if the map is to be of any merit. Because of this, the aircraft is almost always in the local exploration mode leaving large areas of the map unexplored.

It can be assumed that a system of multiple aircraft in which each aircraft is able to map the environment but lacks a method for communicating its observations will perform similarly to the single aircraft case. Thus the single aircraft case may be used as a second baseline for comparison.

The single aircraft system had an average endurance of 1.66 hours over a simulated 50 missions, more than double the no-map case. The shortest modeled mission flown by one aircraft while generating a map was 49 minutes in length, longer than the mean of the missions in which no map was generated. In fact, the availability of a map resulted in improved performance for all runs in the Monte Carlo simulation.

C. Multiple Aircraft

In a system with two or more aircraft the map becomes meaningful as a tool for communication and unplanned (or emergent) cooperation becomes evident. The aircraft tend to avoid each other while exploring the map: this avoidance is not encoded directly into aircraft behavior, but results from the “desire” to

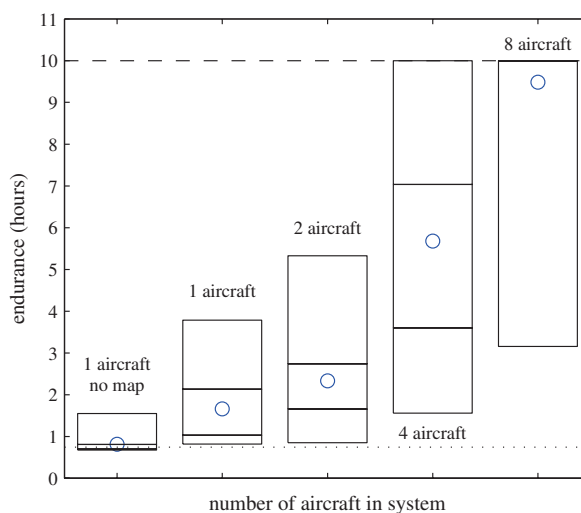


Figure 9. A summary of the Monte Carlo simulations. The dotted line at the bottom of the figure represents the flight time of an aircraft gliding at best L/D from an altitude of 1500 m. The dashed line at the top of the figure shows the maximum length a mission was simulated for. The bars indicate the minimum, 25th percentile, 75th percentile, and maximum survival times for each mission scenario. For 8 aircraft the 25th and 75th percentile flock survivals are 10 hours.

explore regions that have high uncertainty. As one would expect, multiple aircraft result in more up-to-date maps of the wind field, and this improved accuracy results in significantly improved overall endurance.

Results of the Monte Carlo simulations are summarized in Figure 9 and Table 3. As the number of aircraft in the flock increases the overall mission endurance (defined as the length of time until the last aircraft touches down) increases dramatically.

A flock of four aircraft is able to survive to the end of the simulation 12.5% of the time, an eight aircraft flock survived 10 hours in 88.0% of simulations. The mean survival in the eight aircraft flock is reduced by a few outliers– 90% of the runs resulted in greater than 8.2 hours endurance.

Table 3. Summary of Monte Carlo simulations. \bar{t}_1 is the average time (over the 50 runs) before the first aircraft touches down; \bar{t} is the average total mission time; $t_{10\%}$ and $t_{90\%}$ are the 10th and 90th percentile mission times, respectively.

# aircraft	\bar{t}_1	\bar{t}	$t_{10\%}$	$t_{90\%}$
baseline	49 min	49 min	41 min	1.03 h
1	1.66 h	1.66 h	55 min	2.62 h
2	2.04 h	2.33 h	1.13 h	4.01 h
4	4.04 h	5.68 h	2.55 h	10.0 h
8	8.90 h	9.55 h	8.23 h	10.0 h

A snapshot of results for a representative run for a flock of eight aircraft is shown in Figures 10 and 11. Figure 10 show the wind map and the actual wind field for an instant in time, with estimated vertical wind speed represented by hue (color) and the certainty of the estimate (i.e. inverse of covariance) represented by saturation: high certainty corresponds to high saturation. Pale regions of the map are thus areas that have not been visited in some time.

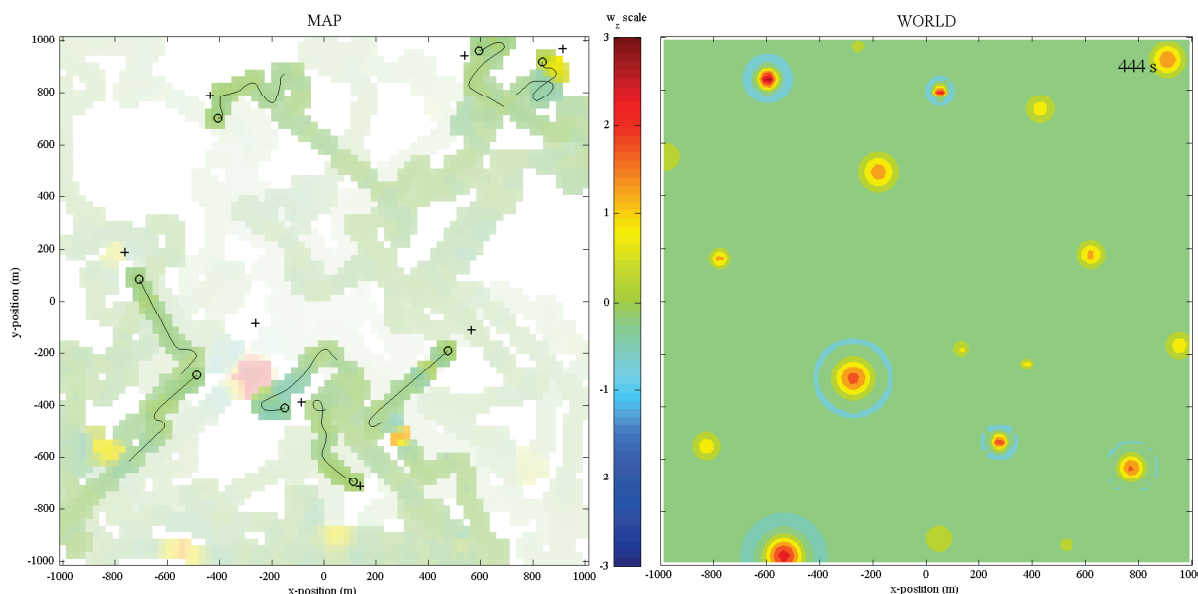


Figure 10. A map generated by eight aircraft. The figure to the right shows actual wind conditions, the figure to the left is the generated map. Note that many of the thermal features appear on the generated map. Aircraft flight paths for the preceding 30 seconds are shown by the lines, aircrafts' locations are shown by circles, and each assigned waypoint is indicated by a black '+'.

A 45 minute segment of the flight history of one of the aircraft in an eight-vehicle flock is shown in Figure 11. Global exploration can be seen in the long straight segments at high altitude; local exploration by the meandering segments; cruise to a thermal by long straight segments at low altitude and thermal flight by spirals that increase in altitude. The slowly tightening spiral at $(P_E, P_N) \approx (0, 500)$ indicates that

the thermal was decaying in strength: the slightly negative \ddot{e} which results from a decaying thermal caused steadily increasing bank angle.

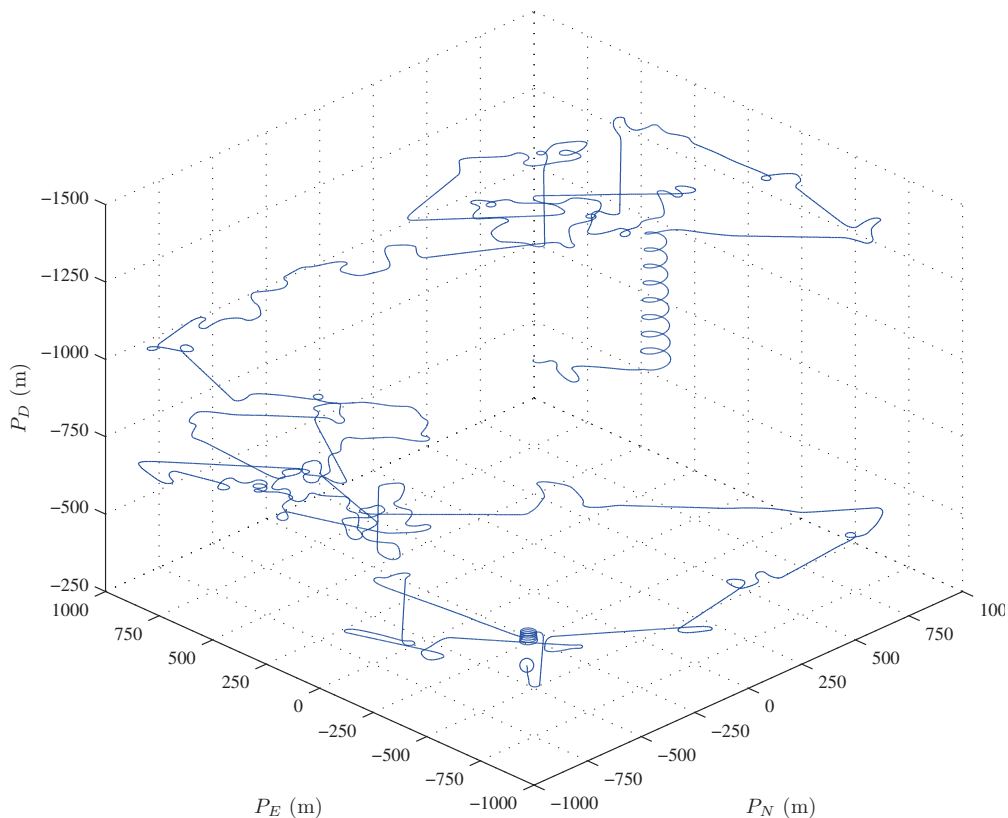


Figure 11. Shown is the flight path for 45 minutes of simulated flight of a single aircraft as part of an 8 aircraft system.

VI. Conclusion

This paper has described a method for coordinated autonomous soaring based on a distributed mapping algorithm. The map consists of a discretized model of the environment and uses a set of one dimensional Kalman filters to compute the expected vertical wind speed and associated covariance at the centroid of each cell. Currently only convection is modeled, with thermal dynamics modeled using an exponential decay. Measurements of wind speed are assumed to be available only at the location of each aircraft and are assumed to be noisy.

The utility of the map is demonstrated using Monte Carlo simulations of a persistent presence task using gliding UAVs. A set of behaviors and a switching logic is used so that an aircraft can exploit the map to improve individual endurance. The availability of a map doubles the endurance compared with a no-map case. Increasing the number of aircraft in the flock leads to improved overall accuracy of the map and leads to overall improvement in mission endurance. Monte Carlo simulations show the improvement in performance provided by the map and by increasing the flock size.

References

- ¹Ostbo, M., Osen, P., Rokseth, G., Homleid, O., and Sevaldrud, T., "Exploiting Meteorology to Enhance the Efficiency and Safety of UAV Operations," Tech. Rep. 2004/00981, Norwegian Defence Research Establishment, Kjeller, Norway, March 2004.
- 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.
- ³Edwards, D. J., "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," *AIAA Guidance, Navigation and Control Conference*, American Institute of Aeronautics and Astronautics, Reston, Virginia, August 2008.
- ⁴Bednarz, J. C., "Cooperative Hunting in Harris' Hawks (*Parabuteo unicinctus*)," *Science*, Vol. 239, No. 4847, March 25 1988, pp. 1525–1527.
- ⁵Folk, M. J., "Cooperative Hunting of Avian Prey by a Pair of Bald Eagles," *Florida Field Naturalist*, Vol. 20, No. 4, 1992, pp. 110–112.
- ⁶Eakle, W. L., Millier, C., Mineau, P., and Vilagosi, J., "An Example of Cooperative Hunting by Saker Falcons in Hungary," *Journal of Raptor Research*, Vol. 38, No. 3, September 2004, pp. 292–293.
- ⁷Dekker, D. and Taylor, R., "A Change in Foraging Success and Cooperative Hunting by a Breeding Pair of Peregrine Falcons and their Fledglings," *Journal of Raptor Research*, Vol. 39, No. 4, December 2005, pp. 394–403.
- ⁸Leonardi, G., "Cooperative Hunting of Jackdaws by the Lanner Falcon (*Falco biarmicus*)," *Journal of Raptor Research*, Vol. 33, No. 2, June 1999, pp. 123–127.
- ⁹Kitowksi, I., "Social Learning of Hunting Skills in Juvenile Marsh Harriers *Circus aeruginosus*," *Journal of Ethology*, Vol. 27, October 2009, pp. 327–332.
- ¹⁰Frye, G. G. and Gerhardt, R. P., "Apparent Cooperative Hunting in Loggerhead Shrikes," *The Wilson Bulletin*, Vol. 113, No. 4, December 2001, pp. 462–464.
- ¹¹Bowman, R., "Apparent Cooperative Hunting in Florida Scrub-Jays," *The Wilson Bulletin*, Vol. 115, No. 2, June 2003, pp. 197–199.
- ¹²Brigliadori, L. and Briigliadori, R., *Competing in Gliders*, Bellavite, Missaglia, Italy, 2005.
- ¹³Petteson, A., "Team Flying and Gaggles in Soaring Championships," .
- ¹⁴Glendening, J. W., "Boundary Layer Information Prediction MAP," .
- ¹⁵Thrun, S., Burgard, W., and Fox, D., *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*, The MIT Press, 2006.
- ¹⁶Elfes, A., "Using Occupancy Grids for Mobile Robot Perception and Navigation," *Computer*, Vol. 22, No. 6, June 1989, pp. 46–57.
- ¹⁷Michael V. Jakuba, D. R. Y., "Autonomous Search for Hydrothermal Vent Fields with Occupancy Grid Maps," *Australasian Conference on Robotics and Automation*, 2008.
- ¹⁸Christopher W. Lum, Rolf T. Rysdyk, A. P., "Occupancy Based Map Searching Using Heterogeneous Teams of Autonomous Vehicles," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2006.
- ¹⁹Geoffrey C. Bower, Tristan C. Flanzer, A. D. N., "Dynamic Environment Mapping for Autonomous Thermal Soaring," *AIAA Guidance, Navigation, and Control Conference, 2-5 August 2010, Toronto, Ontario Canada*, August 2010.
- ²⁰Gedeon, 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.
- ²¹Ying Celia Qi, Y. J. Z., "Energy-Efficient Trajectories of Unmanned Aerial Vehicles Flying through Thermals," *Journal of Aerospace Engineering*, Vol. 18, No. 2, April 2005, pp. 84–92.
- ²²Langelaan, 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.
- ²³Andrew T. Klesh, Pierre T. Kabamba, A. R. G., *Optimization and Cooperative Control Strategies*, Vol. 381, chap. Optimal Cooperative Thermalling of Unmanned Aerial Vehicles, Springer Berlin / Heidelberg, 2009, pp. 355–369.
- ²⁴Allen, M. J., "Updraft Model for Development of Autonomous Soaring Uninhabited Air Vehicles," *44th AIAA Aerospace Sciences Meeting*, AIAA Paper 2006-1510, American Institute of Aeronautics and Astronautics, January 2006.
- ²⁵Simon, D., *Optimal State Estimation*, Wiley Interscience, Hoboken, New Jersey, 2008.
- ²⁶MacCready Jr., P. B., "Optimum Airspeed Selector," *Soaring*, January-February 1958, pp. 10–11.
- ²⁷Reichmann, H., *Cross-Country Soaring*, Soaring Society of America, 1978.
- ²⁸Cochrane, J. H., "MacCready theory with uncertain lift and limited altitude," *Technical Soaring*, Vol. 23, 1999, pp. 88–96.
- ²⁹Klas Anderson, I. K., "On Stability of a Thermal Centering Controller," *AIAA Guidance, Navigation, and Control Conference*, August 2009.
- ³⁰Anderson, K., Kammer, 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.
- ³¹Administration, U. S. F. A., *Glider Flying Handbook*, Skyhorse Publishing, Inc., 2007.

Appendix: Vehicle Properties


All aircraft simulated have parameters based on the SBXC radio controlled motor gliders. The controller parameters given in Table 5 were hand tuned.

Table 4. Parameters for SBXC glider.


variable	value	description
m	10 kg	mass
S	1 m ²	wing area
$f_{LD}(\varphi)$	$0.1723\varphi^4 - 0.3161\varphi^3 + 0.2397\varphi^2 - 0.0624\varphi + 0.0194$	$\varphi = C_{L0} + C_{L\alpha}\alpha$
a, b, c	-0.0059, 0.1507, -1.4833	sink rate polar fit ($\frac{m/s}{v_\alpha}$)

Table 5. Controller parameters

parameter	value	description
Δt_{map}	3 s	interval for map update
Δt_{plan}	1 s	interval for determining behavior/ assigning new waypoints
$\Delta i, j_{local}$	3	number of indexes to examine for local exploration
$\Delta i, j_{global}$	30	number of indexes to examine for global exploration
h_{min}	100 m	minimum safe altitude
Σ_{thresh}	$5.2 \frac{m^2}{s^2}$	threshold covariance for a local region to be considered well-known
e_{thresh}	400 m	threshold specific energy before a thermal is sought
Δt_{gen}	60 s	interval for generating new thermals
z_i	1500 m	convective boundary layer depth



(a) Sink rate versus airspeed for the SBXC motor glider. The green line shows the actual sink rate polar while the red line shows the quadratic approximation.



(b) L/D versus airspeed for the SBXC motor glider.

Figure 12. Vehicle properties.