

Energy-Based Long-Range Path Planning for Soaring-Capable Unmanned Aerial Vehicles

Anjan Chakrabarty* and Jack W. Langelaan†
Pennsylvania State University, University Park, Pennsylvania 16802

DOI: 10.2514/1.52738

To enable long-distance, long-duration flights by small soaring-capable uninhabited aircraft, a graph-based method for planning energy-efficient trajectories over a set of waypoints is presented. It introduces the energy map, which is an upper bound on the minimum energy required to reach a goal from anywhere in the environment while accounting for arbitrary three-dimensional wind fields. The energy map provides the path to the goal as a sequence of waypoints, the optimal speeds to fly for each segment between waypoints, and the heading required to fly along a segment. Trajectories computed using the energy map are compared with trajectories planned using an A^* -based approach. Results are presented for simple wind fields representative of orographic lift. Finally a high-fidelity numerical simulation of a realistic wind field (ridge lift and wave over complex terrain) is used as a test case. The energy-map approach is shown to perform very well without the need for the heuristics associated with A^* .

I. Introduction

ANY small robot (be it ground-, sea-, air-, or space-based) suffers from limited capacity for onboard energy storage. There is often an explicit trade between fuel (for increased mission duration) and sensing payload (for increased data collection). Small aerial robots are subject to additional penalties resulting from their small size. They typically operate at low Reynolds numbers (denoted as Re , it is a dimensionless parameter that defines the ratio of inertial forces to viscous forces). At low Reynolds numbers, viscous forces become very important, and it is very difficult to design a small vehicle with an aerodynamic performance similar to its larger cousins. Small robotic aircraft therefore suffer from both reduced onboard energy capacity and generally lower aerodynamic performance. Together, these two factors greatly reduce the mission capabilities (and hence utility) of small robotic aircraft.

Underwater gliders such as Slocum [1], Spray [2], and Seaglider [3] use variable buoyancy and trim to enable long transects with very little energy use. A similar technique has been proposed for exploration of Venus [4], but this is not practical for Earth-based unmanned aerial vehicles (UAVs); a very large volume of a light gas would be required, resulting in very high drag. However, appropriate flight techniques can permit extraction of energy from the atmosphere by heavier-than-air flight vehicles, greatly extending both range and mission duration of flight vehicles.

Battery technology is continually improving, and this will eventually result in extended range and endurance. However, immediate performance gains are possible by harvesting energy from the atmosphere through soaring flight. Large birds such as eagles, hawks, and condors, as well as human sailplane and hang-glider pilots, routinely use soaring flight to remain aloft for many hours and traverse hundreds of kilometers without flapping wings or the use of engines.

The most common means of atmospheric energy-harvesting exploits vertical air motion (i.e., updrafts). Updrafts have three main causes: uneven heating of the ground, which produces buoyant

instabilities known as thermals; long-period oscillations of the atmosphere, generally called wave, which occur in the lee of large mountain ranges; and orographic lift, where wind is deflected by the slopes of hills and mountains. Typically, updrafts have life spans measured in minutes (for thermals) to hours or days (for ridge and wave lift). Ridge lift and wave are predictable phenomena, and thus one can use trajectory-planning techniques to compute paths that exploit vertical air motion to enable extremely-long-distance or long-duration flights. If the vertical component of wind is strong enough, it may also be possible to recharge batteries by windmilling the propeller or using a ram air turbine, which is a process known as regenerative soaring.

The focus of this paper is on planning long-distance soaring trajectories that harvest energy available from a known wind field (this may be obtained from predictions generated using meteorological forecasting tools such as MM5). Previous research addressed this using a probabilistic-roadmap approach [5], but solutions obtained using probabilistic roadmaps are not necessarily optimal. This paper defines a total-energy map, which computes the total energy required to reach the goal from any starting point in the environment. This total-energy map is computed by first defining nodes (or waypoints) distributed throughout the environment. The speed-to-fly that minimizes the energy expended to fly between neighboring nodes is computed, and wave-front expansion from the goal is used to compute the minimum total energy required to reach the goal from a given node. A trajectory that is optimal within the constraints imposed by the energy map can thus be found. The cost of the constraints imposed by this approach is assessed by comparing trajectories found using the energy map with those found using a heuristic search based on A^* .

II. Previous and Related Work

Rich and varied literature exists in the field of optimal static soaring trajectories with the application of human-piloted soaring flight. Various aspects of optimal static soaring have been addressed, including the optimal speed-to-fly between thermals of known strength (the MacCready problem [6,7], the final glide problem [8], and “dolphin” flight along regions of alternating lift and sink [9–11]. De Jong [12] describes a geometric approach to trajectory optimization and also discusses the optimal deviation from course to minimize time to a goal in a given lift field. Much of this research is directly applicable to the problem of trajectory generation for autonomous soaring flight, but it assumes limited types of known lift distributions (e.g., sinusoidally varying lift [13] or square-wave lift [14]) and generally does not consider the effects of horizontal wind components.

Presented as Paper 2010-8033 at the AIAA Guidance, Navigation, and Control Conference, Toronto, Ontario, Canada, 2–5 August 2010; received 11 October 2010; revision received 7 March 2011; accepted for publication 7 March 2011. Copyright © 2011 by Anjan Chakrabarty and Jack W. Langelaan. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/11 and \$10.00 in correspondence with the CCC.

*Graduate Student, Department of Aerospace Engineering.

†Assistant Professor, Department of Aerospace Engineering. Senior Member AIAA.

Autonomous static soaring is now becoming the focus of more research. Simulation results of thermal flight are reported by Allen [15] and flight-test results are presented in Allen and Lin [16]. Edwards [17] reports very impressive results of autonomous thermal soaring. However, these do not consider the problem of trajectory planning. Dynamic soaring has also become an area of extensive research [18,19], but this is not directly relevant to the problem of trajectory planning for static soaring, because of the different time scales involved.

Wind routing for powered aircraft has been considered for both crewed and uncrewed aircraft. Rubio Torroella [20] describes a planning method based on genetic algorithms; Jardin and Bryson [21] discuss a method based on neighborhood optimal control. Neither of these approaches consider the possibility of harvesting energy from vertical components of the wind field through soaring flight.

III. Graph-Based Approaches to Flight-Path Planning

Cellular-decomposition approaches to robot path planning have been used very successfully. The robot's configuration space is divided into a finite number of regions, and the planning problem is reduced to finding a sequence of neighboring cells between the start and goal regions (e.g., Stentz [22]). These graph-based techniques have been used very successfully in many wheeled-ground-robot path-planning problems and have been used for some UAV planning problems (typically, radar evasion [23]). However, these techniques generally only consider a fixed cost for a transition between nodes in a graph (e.g., time required), and vehicle speed is kept constant. In aircraft applications, total energy can be a critical parameter in trajectory planning (for example, when considering the fuel required to reach the goal). Both environmental and control parameters can affect the energy required for a particular transition: a head wind will increase the required total energy, as will flying at nonoptimal airspeed. Thus, any graph-based planning technique will require a means of accounting for environmental and control conditions in the analysis of costs of transitions between nodes or cells.

Figure 1 shows a schematic of graph-based planning applied to autonomous soaring. Contours show the vertical component of wind, with significant variation caused by the influence of terrain. The environment is first seeded with waypoints (or nodes) and edges. This set of nodes $i = \{i = 0, \dots, m\}$ (with $i = 0$ denoting the goal) and edges $ij = \{i = 0, \dots, m, j = 0, \dots, m\}$ connecting nodes define the allowable paths to the goal. Each edge ij is assigned a constant wind vector w_{ij} with components $w_{x,ij}$, $w_{y,ij}$, and $w_{z,ij}$. Wind-field information is assumed to be available a priori. The problem is now to determine a path to the goal that is feasible, given the aircraft's initial position and initial total energy, and that minimizes some cost of travel. Here, energy expenditure will be used as the cost of travel; for some missions, other parameters (such as time-to-goal) may be more appropriate.

It is assumed that the vehicle is in a trimmed steady-state condition during each transition over an edge ij , and the time required to

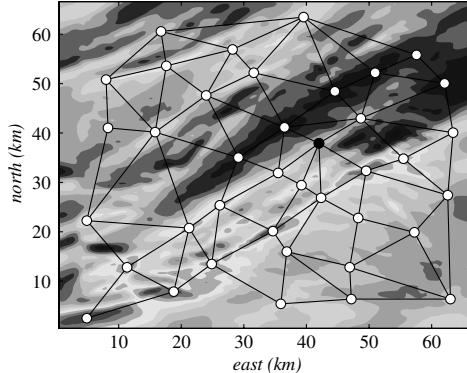


Fig. 1 Schematic of graph-based planning for autonomous soaring over a wind field. Node and edge definition is schematic only.

change from one trim condition to the next as a node is passed is short compared with the length of time required to complete a transition over an edge. Aside from the constraint that the time required to complete a transition at a particular trim condition is long compared with the time required to change trim conditions from one transition to the next (which defines the minimum node spacing), node placement is governed by the assumption that wind is constant over an edge. Higher node density is thus required in regions where spatial gradients in wind field are large (the node placement shown in Fig. 1 is schematic only, not representative of a practical graph). Techniques for determining node placement are beyond the scope of this paper; Cartesian grids will be used in all the examples presented here.

The trajectory-planning problem then becomes a problem of finding the minimum-energy path through the digraph to the goal, i.e., the sequence of nodes \mathbf{n} that minimizes the energy required to reach the goal:

$$e_{\text{goal}} = \sum_{i \in \mathbf{n}, j \in \mathbf{n}} e_{ij} \quad (1)$$

Several techniques have been developed to compute cost-minimizing paths through a graph. However, the unique nature of the problem at hand makes it difficult (if not impossible) to apply them directly. Dijkstra's algorithm is not applicable here, because it is restricted to problems where edge costs are nonnegative, and the Bellman–Ford algorithm encounters problems when negative cycles exist [24]. In this application a negative edge cost implies energy gain, which occurs with flight through an updraft of sufficient strength. Negative energy cycles occur when the aircraft flies repeatedly between two nodes, gaining energy with each trip.

To remove problems associated with cycles, one can add constraints; e.g., A^* only allows a node to be visited once [25]. Another approach is to constrain transitions so that they are always toward the goal. This is a rather severe constraint; however, it enables very fast computation of minimum-energy paths from anywhere in the environment to the goal via wave-front expansion. Wave-front expansion is a breadth-first search that computes the cost-to-go (i.e., the energy required to travel to the goal) for nodes in order of increasing distance to the goal. Constraining transitions so that they are always toward the goal prevents exponential growth in the number of possible paths and ensures that the search is performed in one step, rather than iteratively. This wave-front expansion is discussed in detail in Sec. V.

The set of paths (and associated required initial energy) computed using wave-front expansion is denoted as the *energy map*, and it defines an upper bound on the minimum energy required to reach the goal. Because of the constraint that transitions are always toward the goal, the energy map does not necessarily define the minimum energy required, since a lower-energy path that requires flight away from the goal may exist. However, for any initial energy greater than that defined by the energy map, a feasible path to the goal is guaranteed to exist. After describing the energy map and discussing the energy map for some sample wind fields, this paper applies the energy map to the problem of energy-minimizing flight through a realistic wind field and provides a comparison with a heuristic search to examine the advantages and drawbacks of this approach.

IV. Energy Cost of a Transition

The energy cost of a transition is the minimum energy expended to fly from node i to node j normalized by the distance between the nodes:

$$c_{ij} = -\frac{\Delta e_{ij,\min}}{\Delta s_{ij}} \quad (2)$$

In steady-state flight, this is equal to

$$c_{ij} = -\frac{\dot{e}_{ij,\min}}{v_g} \quad (3)$$

where v_g is the ground speed. In steady-state powered flight, the change in energy Δe_{ij} during flight from node i to node j is a function of the wind vector \mathbf{w}_{ij} , airspeed v_a and throttle setting T . The heading ψ_{ij} required to fly along the desired ground track between the two nodes is a function of the horizontal component of the wind field and the airspeed. Note that in zero wind, the energy change is always negative; for powered aircraft, the change in total energy includes fuel burned. The problem now is to determine the steady-state flight condition that minimizes the energy expended.

A. Vehicle Kinematics

It is assumed that an onboard controller is able to follow heading, airspeed, and throttle commands. Moreover, it is assumed that response to step changes in commands is very fast, compared with the duration of a particular command. Hence, a point-mass model is sufficient to describe vehicle motion for planning purposes (Fig. 2). Vehicle kinematics are given by

$$\dot{x} = v_a \cos \gamma \cos \psi + w_x \quad (4)$$

$$\dot{y} = v_a \cos \gamma \sin \psi + w_y \quad (5)$$

$$\dot{z} = v_a \sin \gamma + w_z \quad (6)$$

where v_a is airspeed; γ is flight-path angle with respect to the surrounding air mass; ψ is heading; and w_x , w_y , and w_z are the three components of the 3-D wind vector.

The flight-path angle γ is a function of airspeed v_a and throttle setting T and can be obtained for steady flight. From Fig. 2, resolving forces parallel and perpendicular to the flight path,

$$mg \cos \gamma = L + T \sin \alpha \quad (7)$$

$$mg \sin \gamma = D - T \cos \alpha \quad (8)$$

where m is mass of the vehicle and α is the angle of attack (note that this implicitly assumes that the thrust axis is aligned with the aircraft's body x axis). Using the standard definition of force coefficients,

$$\cos \gamma = \frac{qS}{mg} (C_L + C_T \sin \alpha) \quad (9)$$

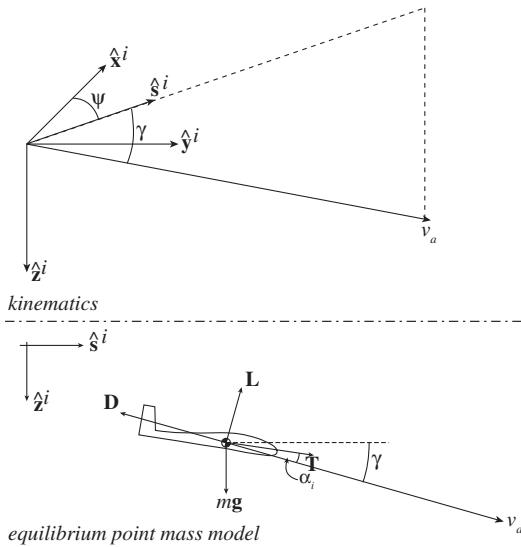


Fig. 2 Point-mass model.

$$\sin \gamma = \frac{qS}{mg} (C_D - C_T \cos \alpha) \quad (10)$$

$$\text{where } q = \frac{1}{2} \rho v_a^2.$$

During cruise, the flight-path angle with respect to the air mass γ is typically small; hence, $\sin \gamma \approx \gamma$ and $\cos \gamma \approx 1$. During trimmed cruise flight, angle of attack is generally small (3 to 6 deg), and thrust is significantly smaller in magnitude than lift. Hence, it is further assumed that $C_T \sin \alpha$ is negligible compared with C_L . From Eq. (9),

$$C_L = \frac{mg}{qS} = \frac{2mg}{\rho v_a^2 S} \quad (11)$$

Here, C_L is the lift coefficient, ρ is density of the air, and S is wing area. A polynomial approximation is used for the aircraft's drag polar:

$$C_D = \sum_{i=0}^n a_i C_L^i \quad (12)$$

Typically, a second-order polynomial is used to represent drag coefficient. However, this is often only valid over a fairly narrow speed range, and a fourth-order polynomial is used here.

Substituting into Eq. (10), the air-mass relative flight-path angle for a particular speed and thrust can thus be computed as

$$\sin \gamma = \frac{qS}{mg} \left(\sum_{i=0}^n a_i C_L^i - C_T \right) \quad (13)$$

Combining Eq. (11) with Eq. (13) and vehicle kinematics, the vehicle's flight path is completely specified by inputs $\mathbf{u} = [v_a \ \psi \ C_T]^T$ and wind speed \mathbf{w} . This model is adequate as long as the length of time of each trajectory segment is large compared with the time constant of the vehicle's step response with respect to the inputs \mathbf{u} .

The inputs v_a and C_T affect the energy expended during the flight; thus, the choice of input is critical to energy-efficient flight. Heading ψ is determined by the ground track between the nodes, airspeed, and the horizontal component of wind velocity.

B. Flight Between Two Nodes

Referring to Fig. 3, the line segment joining two successive nodes is the desired ground track. The velocity of the vehicle is decomposed into in-track v_t and cross-track v_c components to maintain flight along the desired ground track $v_c = 0$. The wind vector is also decomposed into in-track w_t and cross-track w_c components; v_g is the ground speed of the vehicle.

From Fig. 3 the relationship between airspeed, ground speed, heading, and ground track for an arbitrary horizontal component of wind can be obtained:

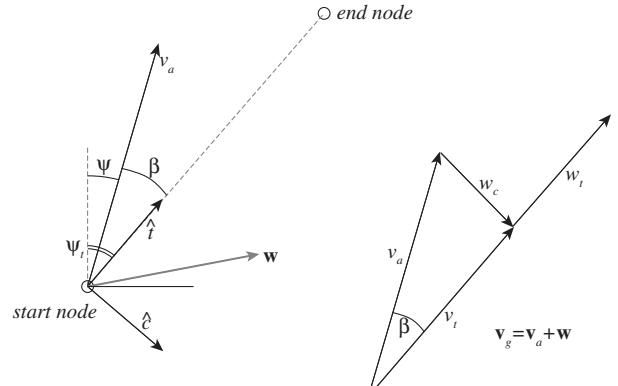


Fig. 3 Track coordinate frames (left) and resolution of airspeed and wind vectors into the track coordinate frame (right). Positive angles are shown.

$$v_t = \sqrt{v_a^2 \cos^2 \gamma - w_c^2} \quad (14)$$

$$v_g = v_t + w_t \quad (15)$$

$$v_a \cos \gamma \sin \beta = w_c \quad (16)$$

where v_g is the magnitude of the ground speed, $v_a \cos \gamma$ is the projection of the airspeed vector onto the horizontal plane, and β is the angle between the airspeed vector and the desired ground track. Recall that γ is generally small; hence, $\cos \gamma \approx 1$ and the ground speed is $v_g \approx \sqrt{v_a^2 - w_c^2} + w_t$. The constraint $v_c = 0$ is expressed in Eq. (16).

The aircraft heading to maintain the desired ground track is $\psi = \psi_t - \beta$. Hence,

$$\psi = \psi_t - \sin^{-1} \frac{w_c}{v_a} \quad (17)$$

Clearly, heading ψ is dependent on airspeed v_a . The problem now is to determine the optimal value of airspeed v_a and thrust coefficient C_T for flight between two nodes. This will be done by analyzing the energy required to fly the path segment between the nodes.

C. Minimizing Energy Loss

The steady-state airspeed that minimizes the energy lost over a segment (or, equivalently, maximizes the energy gained) will be determined. Total energy is

$$E_{\text{tot}} = mgh + \frac{m}{2} v_a^2 + E_s \quad (18)$$

where h is altitude and E_s is onboard stored energy. Specific total energy is

$$e_{\text{tot}} = \frac{E_{\text{tot}}}{mg} = h + \frac{v_a^2}{2g} + \frac{E_s}{mg} = h + \frac{v_a^2}{2g} + e_s \quad (19)$$

This has units of altitude and is sometimes referred to as energy altitude.

Minimizing energy lost over a segment means maximizing $\Delta e_{\text{tot}}/\Delta s$. In steady-state flight this is equivalent to maximizing \dot{e}/v_g : in other words, flying to maximize range. The rate of change of specific energy is

$$\dot{e}_{\text{tot}} = \dot{h} + \frac{v_a \dot{v}_a}{g} + \dot{e}_s \quad (20)$$

In steady flight, acceleration is zero; hence,

$$\dot{e}_{\text{tot}} = \dot{h} + \dot{e}_s = -\dot{z} + \dot{e}_s \quad (21)$$

Recall that z is positive down and \dot{z} is defined in Eq. (6). The quantity \dot{e}_s is the rate of change of onboard stored energy. This is dependent on motor power and the efficiency of energy conversion:

$$\dot{e}_s = -\frac{T v_a}{mg \eta_{\text{ec}} \eta_p} = -\frac{q S C_T v_a}{mg \eta_{\text{ec}} \eta_p} \quad (22)$$

where η_{ec} is the net efficiency of energy conversion from source to shaft (in electrical power systems, this is the product of motor efficiency and speed controller efficiency), and η_p is the propeller efficiency.

Thus, the rate of change of total energy (in steady-state flight) is

$$\dot{e}_{\text{tot}} = -(v_a \sin \gamma + w_z) - \frac{q S C_T v_a}{mg \eta_{\text{ec}} \eta_p} \quad (23)$$

and minimum-energy loss over a segment occurs when one maximizes the rate of change of energy divided by the ground speed:

$$\frac{\dot{e}_{\text{tot}}}{v_g} = -\frac{v_a \sin \gamma + w_z}{\sqrt{v_a^2 - w_c^2} + w_t} - \frac{q S}{mg \eta_{\text{ec}} \eta_p} \frac{C_T v_a}{\sqrt{v_a^2 - w_c^2} + w_t} \quad (24)$$

Since e_{tot} has dimension distance, \dot{e}_{tot}/v_g is dimensionless.

There may be cases in which onboard energy is especially valuable (for example, to maximize time on station at the goal). In this case, a reward function that weights stored energy more heavily may be appropriate:

$$r_e = \frac{\dot{h}}{v_g} + \mu \frac{\dot{e}_s}{v_g} \quad (25)$$

Here, setting $\mu > 1$ will increase the importance of conserving onboard energy versus conserving altitude. Hence, the general problem of the speed-to-fly over a segment can be solved by finding the airspeed v_a and throttle setting C_T , which solves the optimization problem:

Maximize

$$\frac{\dot{h}}{v_g} + \mu \frac{\dot{e}_s}{v_g} \quad (26)$$

Subject to

$$\dot{h} = -v_a \sin \gamma - w_z \quad (27)$$

$$v_g = \sqrt{v_a^2 - w_c^2} + w_t \quad (28)$$

$$\sin \gamma = \frac{q S}{mg} \left(\sum_{i=0}^n a_i C_L^i - C_T \right) \quad (29)$$

$$\dot{e}_s = -\frac{q S}{mg \eta_{\text{ec}} \eta_p} C_T v_a \quad (30)$$

$$v_{a,\text{min}} \leq v_a \leq v_{a,\text{max}} \quad (31)$$

$$C_{T,\text{min}} \leq C_T \leq C_{T,\text{max}} \quad (32)$$

$$v_g > 0 \quad (33)$$

A generic function minimizer (such as MATLAB's `fmincon`) can be used to find v_a^{opt} and $C_{T,\text{optimal}}$, which minimizes c_{ij} while ensuring that constraints such as airspeed limits (stall and maximum speed) are not exceeded. This formulation for the optimal airspeed computation includes the effect of the 3-D wind vector; once v_a^{opt} has been computed, the required heading to maintain the desired ground track between the start and end nodes is computed from Eq. (17).

Under certain conditions, additional constraints may exist (e.g., gliding flight, constant-altitude flight).

1. Gliding Flight

For gliding flight $C_T = 0$ and $\dot{e}_s = 0$. The only free control input for flight along a path segment is airspeed, and the optimization problem is as follows:

Maximize

$$\frac{-(v_a \sin \gamma + w_z)}{\sqrt{v_a^2 - w_c^2} + w_t} \quad (34)$$

Subject to

$$\sin \gamma = \frac{q S}{mg} \sum_{i=0}^n a_i C_L^i \quad (35)$$

$$C_L = \frac{mg}{q S} \quad (36)$$

$$v_{a,\text{min}} \leq v_a \leq v_{a,\text{max}} \quad (37)$$

$$v_g > 0 \quad (38)$$

In zero wind this will result in flight at the best L/D . Note that this is not the same as the MacCready problem [6]; here, the focus is on computing speed-to-fly to minimize energy loss in steady wind, whereas the MacCready problem is focused on minimizing the time required to glide from a given altitude to a thermal of known strength and then climb back up to the starting altitude.

2. Constant-Altitude Flight

Altitude restrictions may be imposed to ensure separation between manned and unmanned aircraft, or they may occur because of sensor requirements (for example, to ensure that a particular sensor footprint is maintained).

At constant altitude, $\dot{h} = 0$. Hence, $v_a \sin \gamma = -w_z$,

$$\sin \gamma = \frac{qS}{mg} \left(\sum_{i=0}^n a_i C_L^i - C_T \right) = -\frac{w_z}{v_a} \quad (39)$$

and the optimization problem becomes as follows:

Maximize

$$\frac{\dot{e}_s}{v_g} \quad (40)$$

Subject to

$$\frac{qS}{mg} \left(\sum_{i=0}^n a_i C_L^i - C_T \right) = -\frac{w_z}{v_a} \quad (41)$$

$$v_g = \sqrt{v_a^2 - w_z^2} + w_t \quad (42)$$

$$\dot{e}_s = -\frac{qS}{mg} \frac{C_T v_a}{\eta_{ec} \eta_p} \quad (43)$$

$$v_{a,\min} \leq v_a \leq v_{a,\max} \quad (44)$$

$$C_{T,\min} \leq C_T \leq C_{T,\max} \quad (45)$$

$$v_g > 0 \quad (46)$$

Note that constant-altitude flight will often require positive values of thrust. Constant-altitude gliding flight is only possible when the upward component of wind speed exceeds the sink rate of the aircraft. The remainder of this paper is concerned with constant-altitude flight.

D. Regenerative Soaring

For battery-powered aircraft, a windmilling propeller (or a ram air turbine) can be used to recharge batteries at the cost of increased drag. One can thus trade potential energy (altitude) for stored electrical energy. When flying through a strong-enough updraft, it is possible to either 1) gain potential energy by climbing at constant speed, 2) gain kinetic energy by flying at higher speed but at constant altitude, or 3) gain stored electrical energy by flying at constant speed and altitude and windmilling the propeller. Of course, a combination of the three can also occur.

Using a windmilling propeller to gain electrical energy is known as regenerative soaring and was first described by MacCready [26].

In the context of the optimization problem posed above, setting $C_{T,\min} = 0$ means that regeneration cannot occur (as \dot{e}_s , the energy expended from the onboard supply, has a minimum value of zero). Permitting negative values of $C_{T,\min}$ implies that \dot{e}_s can be a positive quantity [Eq. (22)], and thus the energy (potential and/or kinetic) can be transferred to batteries. In practical application, the efficiency of a windmilling propeller may be less than a “propelling” propeller, and efficiency must be accounted for during regeneration. For simplicity, it is assumed that propeller efficiency is constant.

Note that onboard batteries are limited in the total energy that they can store. Regenerative soaring can thus only occur until stored energy reaches its maximum value $e_{s,\max}$.

V. Paths via Wave-Front Expansion: The Energy Map

The energy cost of transitions can now be used to find minimum-cost paths through the graph.

A. Map Definition

After the environment has been seeded with nodes \mathbf{n} , they are ordered by increasing distance from the goal. Transitions that simultaneously satisfy the condition of being to a neighboring node and reducing the distance to the goal are defined as allowable. An example of a regular Cartesian grid with allowable transitions is shown in Fig. 4.

The energy to reach the goal is computed for the group of nodes nearest the goal, and the energy corresponding to each node is defined as their respective costs-to-go. For the next group of nodes the energy required to reach nearest neighbors in the first group is computed, and the cost-to-go for each node is the minimum total energy over all possible paths to the goal. The process continues until energy-to-goal has been computed for each node. This is a breadth-first dynamic programming approach, and with the constraint that transitions must always end in nodes nearer to the goal, the resulting energy map gives an upper bound on the minimum energy required to reach the goal from any point in the environment. It thus provides a means to check the feasibility of a path to the goal for an aircraft with a particular initial position and initial total energy.

In this approach each node is encoded with the total energy required to reach the goal (i.e., the cost-to-go), the next node in the path to the goal, and the control inputs (speed-to-fly and heading) required to reach the next node. The energy map thus encodes a complete path to the goal from anywhere in the environment. Feasibility of paths from a particular starting point and initial energy can be immediately answered.

This raises the question of, “Why solve for the whole environment and not from a single starting point?” The breadth-first search is a means to quickly and deterministically find paths to the goal. Because it is breadth-first it necessarily finds paths from all nodes. In some cases, this may cause this method to be slower than probabilistic approaches that only find paths from a single starting point, but one cannot make conclusions of optimality from randomized approaches. Further, there is a finite likelihood that randomized

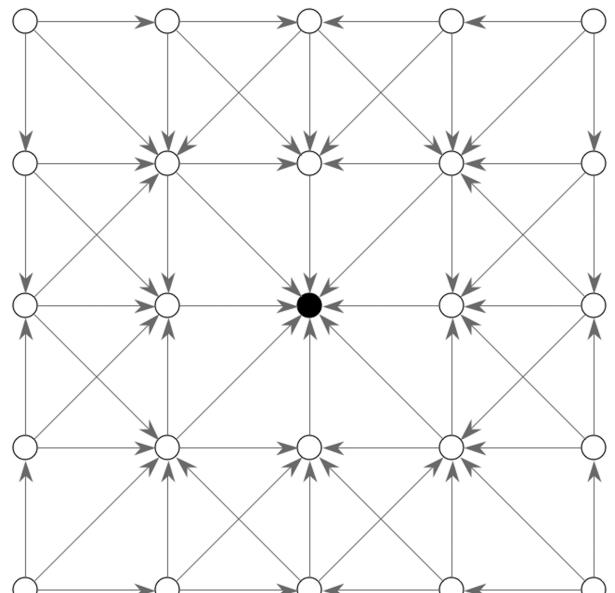


Fig. 4 Sample regular Cartesian grid showing allowable transitions for wave-front expansion. The goal is at the center of the grid.

Table 1 Parameters for SB-XC glider

Parameter	Value
Mass m	10 kg
Wing area S	1 m ²
$f(C_L)$	$0.1723C_L^4 - 0.3161C_L^3 + 0.2397C_L^2 - 0.0624C_L + 0.0194$
$v_{a,\min}$	12 m/s
$v_{a,\max}$	35 m/s
$C_{T,\min}$	-1
$C_{T,\max}$	0.2
Propeller efficiency η_p	0.8
Net energy conversion efficiency (battery to shaft) η_{ec}	0.855
Maximum onboard energy available $e_{s,\max}$	2500 m

approaches will require an arbitrarily long time to find a feasible path, whereas the time required for the energy-map approach depends only on the number of nodes. Later in this paper it will become evident that definition of admissible cost functions for heuristic approaches to the single-path problem (e.g., A^*) is not obvious, and the energy-map approach obviates the need for heuristics.

It is possible to constrain the search region to reduce computation time. For example, the search area can be constrained to a region of nodes between a starting point and the goal. Alternatively, the search can be terminated as soon as the start node lies on the edge of the wave front. Of course, it is possible to compute the energy map in reverse: i.e., to compute the energy required to reach anywhere in the environment from a particular start node, thus defining the maximum range possible from some initial point. Note that this range map is not simply the reverse of the energy map; the energy required to traverse an edge depends on the direction of travel.

Note that paths found using this method are not necessarily *minimum-energy* paths; since paths are constrained to always approach the goal, trajectories that are more energy-efficient may exist. Such a path would first proceed away from the goal before turning toward it. However, relaxing the constraint that the goal must be approached would make the wave-front-expansion approach computationally intractable for large environments.

B. Energy Maps for Sample Wind Fields

To demonstrate the energy-mapping approach and to provide some intuition of results, energy maps for some simple wind fields are computed. Calculations were performed for an RnR Products

SB-XC ratio-controlled glider; parameters are given in Table 1. Note that a fourth-order polynomial is used to relate C_D to C_L ; this provided a better fit to the computed data over the full speed range. In all cases, flight is constrained to constant altitude; energy thus represents the expenditure of onboard stored energy. In all cases, a Cartesian grid is used and transitions as shown in Fig. 4 are allowable. The effect of the grid on the resulting paths will be apparent.

First, consider a comparison of the zero-wind case with uniform wind blowing along the x axis (Fig. 5). The total energy required to reach the goal is shown as a mesh surface, and flight paths to the goal are shown as black streamlines on the plane $z = 0$.

For the zero-wind case, the expected result is obtained: the energy required to reach the goal increases linearly with distance-to-goal. The streamlines in the graph show the path that the vehicle should follow to satisfy the requirement for the minimum energy. Adding a horizontal wind component tilts the cone, so that starting points downwind of the goal require more initial total energy to reach the goal than starting points upwind of the goal. Following intuition, paths to the goal are modified accordingly.

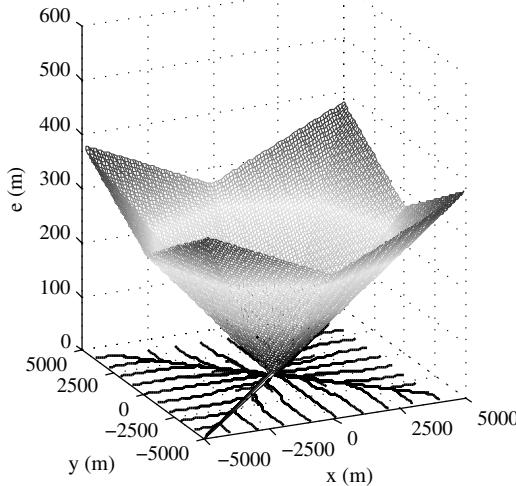
The second case involves a nonuniform wind field consisting of constant wind shear. Here, the horizontal component of wind varies linearly from $w_x = -10$ m/s at $y = -5000$ to $w_x = 10$ m/s at $y = 5000$. The energy map and paths-to-goal are shown in Fig. 6.

Matching intuition, the surface defining the energy map is now twisted. Optimal flight paths show the vehicle maximizing exposure to favorable winds (or minimizing exposure to unfavorable wind); for starting points in the regions $x > 0$, $y > 0$, $x < 0$, and $y < 0$ flight paths remain in the region of high tailwind before turning to approach the goal. For starting points in the regions $x > 0$, $y < 0$, $x > 0$, and $y < 0$ the same is true; flight paths begin with motion toward regions of more favorable wind before turning toward the goal.

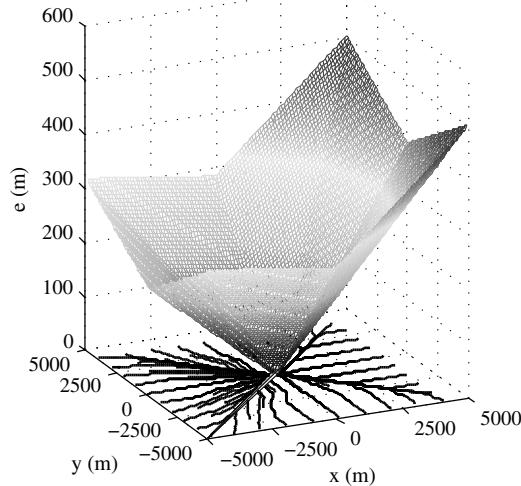
c. Long-Range Flight via Orographic Lift

Orographic lift occurs when wind is deflected upward by terrain: for example, along coastlines with onshore breezes or along hills and ridges. Unlike thermal lift, orographic lift is a predictable phenomenon and tends to be significantly longer-lived. Under the right conditions, extremely-long-duration soaring flight is possible along ridges or mountain ranges.

Consider the problem of flight in an environment with multiple regions of both upward- and downward-moving air, as well as a horizontal wind component (Fig. 7). Two parallel ridges are separated by 12 km. A global coordinate frame is defined with y parallel to the ridges, so that the ridge centerlines are located at $x = 4$ and 16 km. Each ridge is modeled as an infinitely long hemicylinder



a) Zero wind



b) 5 m/s wind in the positive x direction

Fig. 5 Energy maps for a uniform horizontal wind fields. Note the effect of the discretization on results: a polar grid in the left image would result in paths straight to the goal.

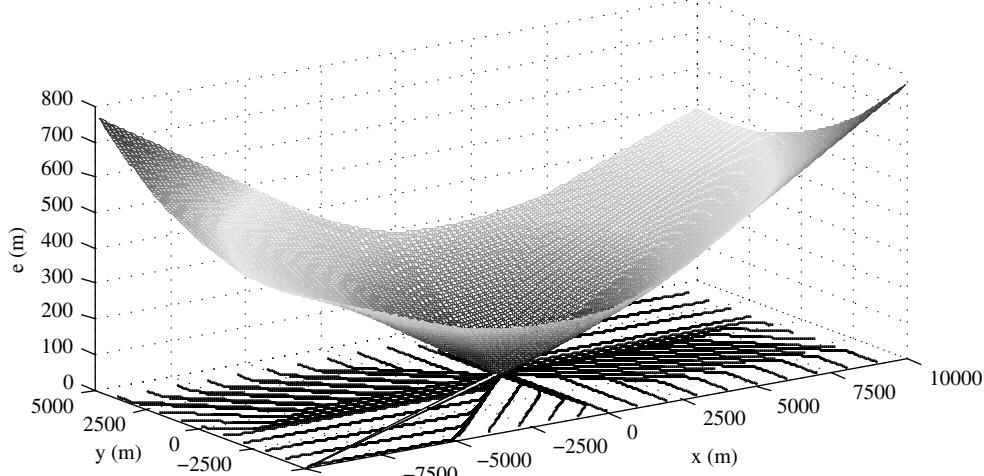


Fig. 6 Constant-wind-shear energy map and minimum-energy flight paths.

with a radius of 300 m, and potential flow is used to compute the wind field. Note that potential flow cannot model flow separation on the downwind side of the ridges, but it is along the upwind side that one finds upward (and thus favorable) air motion. Intuition suggests (and results will show) that the glider will tend to follow the upwind sides of ridges; thus, the flow on the downwind sides of ridges is less critical to trajectory planning (except for the times when the vehicle must traverse these nonfavorable regions).

The origin is located at (0,0) and the energy map is computed for an area defined by $0 \leq x \leq 20$ km and $0 \leq y \leq 100$ km. Flight is constrained to an altitude of 310 m (just enough to clear the ridges). Allowing altitude variation may improve results by reducing overall energy expenditure, but constant-altitude flight also allows the use of graph-based planning without excessive graph size. First, the case of nonregenerative flight is considered.

Recall that the method discussed here does not require uniform grid spacing, and in some cases, uniform grid spacing will either result in inaccurate energy maps (because it is assumed that the wind field is constant over an edge) or excessive computational requirements. For this example, a nonuniform Cartesian grid is used, with finer grid spacing over the ridges (where the wind field changes rapidly) and wider spacing between the ridges (where the wind field is roughly constant). Figure 7b shows a vector plot of the computed wind field at the x coordinates of the grid. Spacing varies from a minimum of 100 m to a maximum of 1000 m. Grid spacing in the y direction is constant at 1000 m.

Figure 8 shows the energy map and paths-to-fly for a wind field resulting from $w_{x,\infty} = -5$ m/s (which results in maximum vertical component of wind of approximately 3 m/s along the ridge). The minimum sink rate of SB-XC glider (i.e., minimum rate of altitude loss in still air) is 0.56 m/s. It is therefore possible to fly with zero net

energy loss along the ridges (indeed, it is possible to employ regeneration to gain energy).

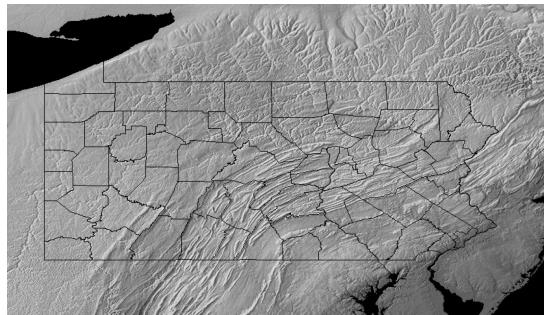
Figure 8 shows the energy map and flight paths when regeneration is not employed. Note first the region near the goal within which the energy map leads the aircraft straight to the goal. Second, there are nearly constant “troughs” of starting energy along the ridges, indicating that energy can be gained very quickly near the ridges. If the glider starts in the region of upward-moving air, only enough energy to begin stable flight is necessary to permit a flight to the goal, even if the goal is more than 100 km distant. In this case, the far ridge does not show a deep trough. Since the flight is restricted to constant altitude and regeneration is not employed, energy (in the form of altitude) cannot be gained and there is no benefit to flying along the far ridge.

As expected, the slope of the energy map is significantly steeper for upwind transitions than for downwind transitions. Significantly more energy is required to travel a given distance with a headwind than with a tailwind.

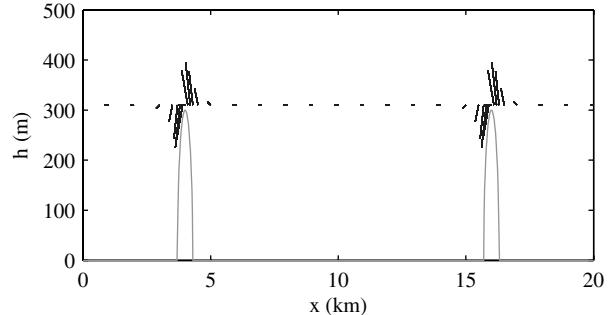
D. Orographic Lift and Regenerative Soaring

Regenerative soaring has the potential to greatly improve the energy efficiency of flight paths. It may also be possible to reduce flight time for a given energy expenditure; with the knowledge that energy can be generated in the future, one can expend energy to fly faster.

Figure 9 shows the energy map and flight paths when regenerative soaring is employed. In this case, the aircraft can gain energy during flight along the ridge, implying that it can begin flight with a nearly empty battery and still reach the goal. Since energy can be gained, it is now energetically favorable (depending on start location) to use the



a) Digital elevation map of Pennsylvania



b) Cross section of parallel ridges and wind field

Fig. 7 Topography of central Pennsylvania (left) and schematic of ridges and potential flow solution of the wind field (right).

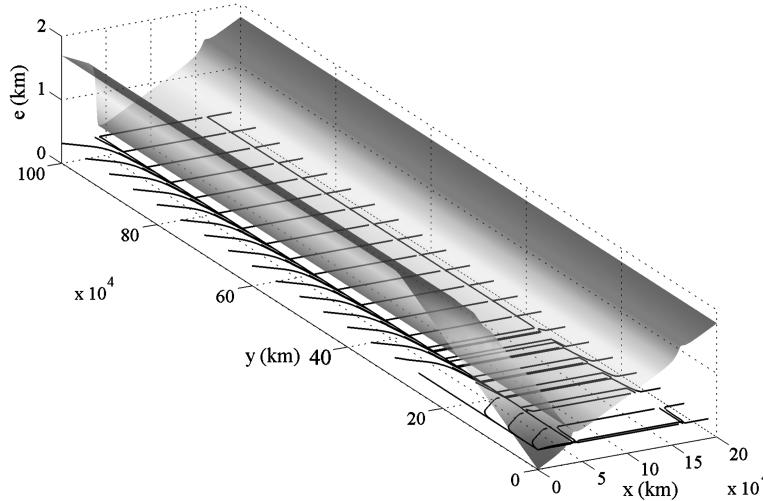


Fig. 8 Energy map (translucent surface) and optimal paths to the goal (streamlines). Regeneration is not employed.

far ridge for portions of the flight. The goal region of attraction is also somewhat smaller, implying that it is now favorable to detour to gain energy from the environment.

E. Computation Time

One method of handling uncertainty in the wind field is to replan trajectories as updated wind information is obtained. This requires real-time trajectory planning. For planning purposes, a trajectory computation is real time if the time required to plan a path is less than the time required to fly one segment of the trajectory. The energy maps above consist of 3700 nodes (37 in the x direction and 100 in the y direction). Generating the energy maps using MATLAB on a 2.67 GHz Dell laptop took 19 s. For grids of 10,000 nodes, energy-map generation took approximately 55 s. For the SB-XC, the best L/D occurs at approximately 16 m/s. Node spacing of 300 to 1000 m will thus allow trajectory computation in less time than required to fly a single segment.

In the implementation described here, the set of nodes and the set of allowable transitions are precomputed, requiring only the wave-front expansion and computation of optimal airspeeds to be done in real time. MATLAB's `fmincon` function was used to compute optimal airspeeds; saving of computation time will occur if this is replaced with either a precomputed lookup table of best speeds to fly and energy change for representative values of horizontal and vertical components of wind or an approximator such as a neural network. Since microSD cards with 16 GB capacity are now available at very low cost, memory is unlikely to be an issue.

VI. Minimum-Energy Paths via Heuristic Search

A search-based approach such as A^* does not impose the transition-to-goal constraint and thus can be expected to give better paths (i.e., with lower required initial energy) than the energy map under some conditions. Transitions are now allowed to all neighboring nodes, and in graph-based terms, the process in which the search process expands is twofold: 1) selecting the next node to visit and 2) planning the best path through the rest of the graph to arrive at the goal. The cost function is a linear combination of two terms: $g(n)$, the cost of the best path found so far, and $h(n)$, a heuristic function that is an estimate of the energy cost from node n to the goal. Here, h is defined as the straight-line distance between the current node and the goal divided by the aircraft's best L/D and is thus an estimate of the energy required to reach the goal. Hence, nodes are expanded according to a cost function:

$$f(j) = \alpha g(j) + (1 - \alpha)h(j) \quad (47)$$

Here, $g(j)$ is the energy cost of the transition to node j :

$$g(j) = c_{ij} \quad (48)$$

where c_{ij} is defined in Eq. (3).

The second component is

$$h(j) = \frac{r_{j0}}{L/D_{\max}} \quad (49)$$

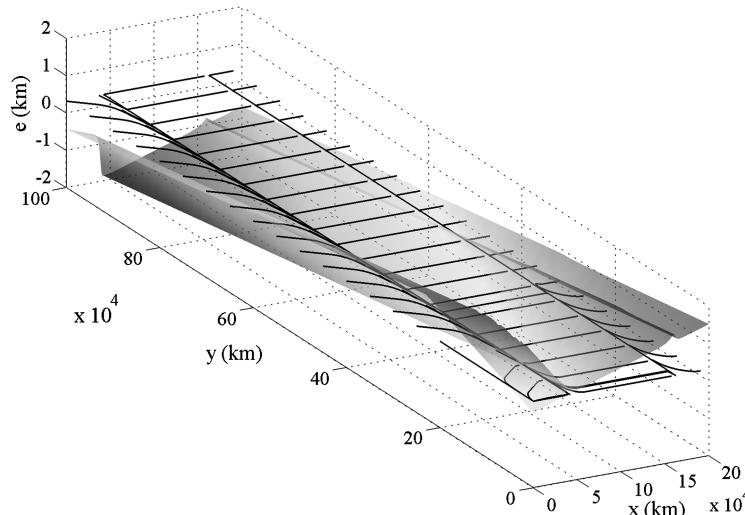


Fig. 9 Energy-map (translucent surface) and optimal paths to the goal (streamlines) when regenerative soaring is employed.

the energy required to reach the goal in straight-line flight in zero wind. This is an estimate of the cost-to-go. Varying the weight α from 0 to 1 allows the computed path to favor path length or energy expenditure, and the key is to find the value of α that gives a good balance between progress-to-goal and the energy required to fly the path.

A. Energy Gain Versus Distance-to-Goal

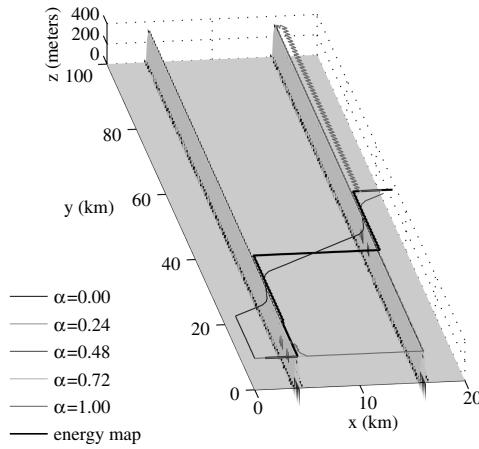
To find a good balance between energy gain and progress-to-goal, the ridge flight case of the previous section is used. Since the heuristic planner allows flight away from the goal, it is most useful to examine starting points that are not on the furthest edge of the problem boundary; various starting points were investigated, but results are only presented for paths starting from $x = 20$ km and $y = 50$ km, for brevity. Similar results were observed for other starting points.

Figure 10a shows flight paths as α is varied from 0 to 1. Clearly, when there is no weight on energy gain ($\alpha = 0$), optimal flight paths should go straight to the goal (constrained by allowable transitions), and this indeed occurs. In this case, regenerative soaring is not employed, and increasing the weight on energy gain causes the vehicle to detour somewhat to enable energy-efficient flight. The amount of the detour (and energy cost) is roughly constant for a large range of α (paths are coincident for $\alpha = 0.24$, 0.48, and 0.72);

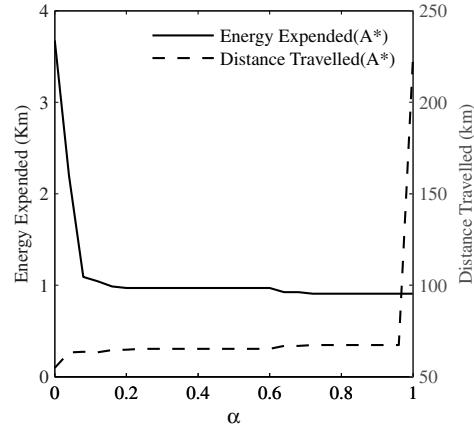
however, once $\alpha = 1$ (and distance-to-goal is no longer part of the path cost), a large detour is flown.

Figure 10b shows the energy expended and the total distance flown as the energy vs progress weight is varied. Results match the intuition that as the importance of energy gain is increased, the net expended energy is reduced. Note that energy *gain* is not possible here (because altitude is constant and regeneration is not employed), but periods of zero-loss energy flight are possible.

Figure 11 shows energy expended, velocity, and thrust coefficient along the flight path as α is varied. Regions where thrust coefficient is zero show where the aircraft is flying in upward-moving air. For the case $\alpha = 0$ the energy expenditure increases steadily along the path and the distance flown is shortest. However, note that the energy expended to reach the goal (at 55 km) is 3600 m, significantly greater than the maximum battery capacity of 2500 m. Thus, direct flight-to-goal is not feasible for this case. Increasing the cost of energy expenditure quickly decreases the energy required to reach the goal at the cost of increased distance. Since regeneration is not employed, the aircraft flies at high speed in lift along the upwind sides of ridges while expending zero net energy. It is evident that placing too large a weight on energy cost (when $\alpha = 1$) is not beneficial; the distance traveled is extremely large, but the energy expended is almost the same as the $\alpha = 0.48$ case. Note that if the goal is maximizing time aloft, then this would be a useful approach, but the objective here is flight to a goal.



a) Comparison of A* paths



b) Initial Starting energy and distance travelled for varying α

Fig. 10 Different paths to the goal starting from $x = 20$ km and $y = 50$ km (left) and comparison of the initial starting energy for $x = 20$ km and $y = 50$ km (right). The energy-map-computed flight path has energy expended equal to 1.78 km and distance traveled equal to 73.1 km.

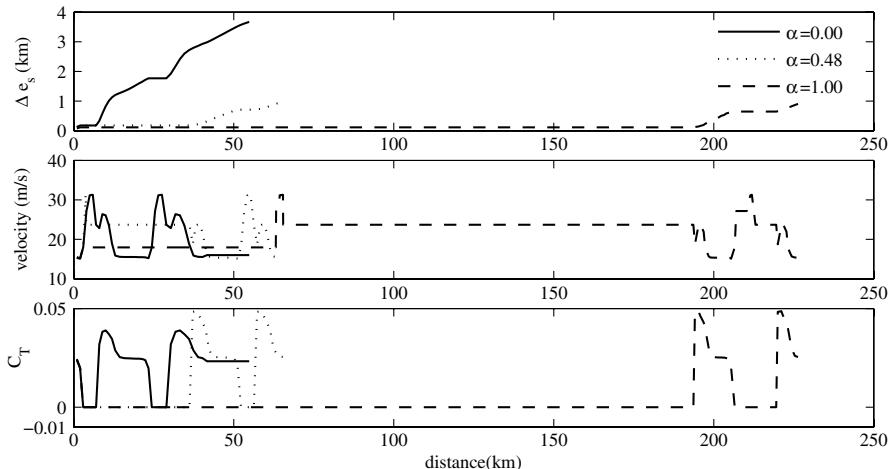


Fig. 11 Energy, velocity, and thrust coefficient along the paths for different values of α when regeneration is not employed.

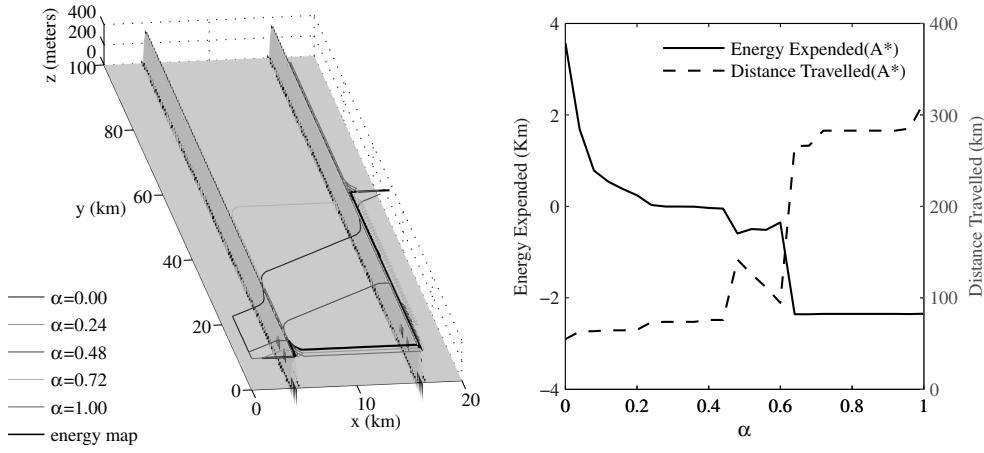


Fig. 12 Different paths to the goal starting from $x = 20$ km and $y = 50$ km (left) and comparison of the initial starting energy for $x = 20$ km and $y = 50$ km (right). The energy-map-computed flight path has energy expended equal to 0.48 km and distance traveled equal to 85.3 km.

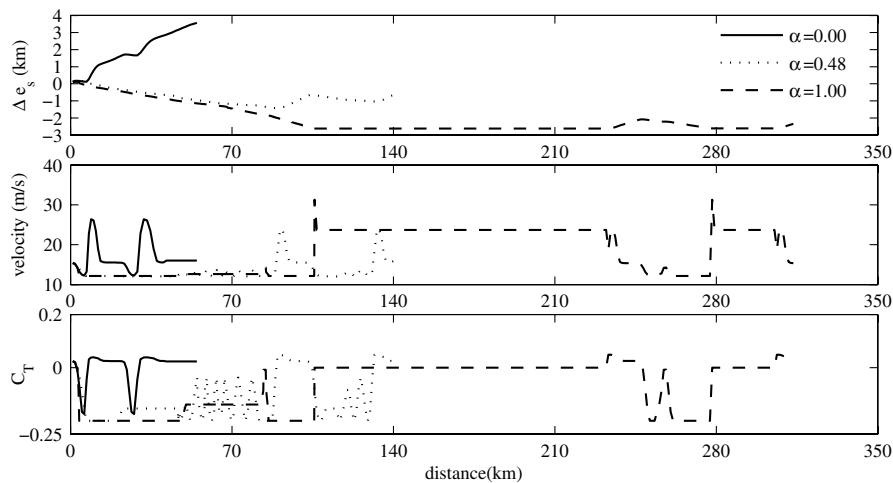


Fig. 13 Energy, velocity, and thrust coefficient along the paths for different values of α when regeneration is employed.

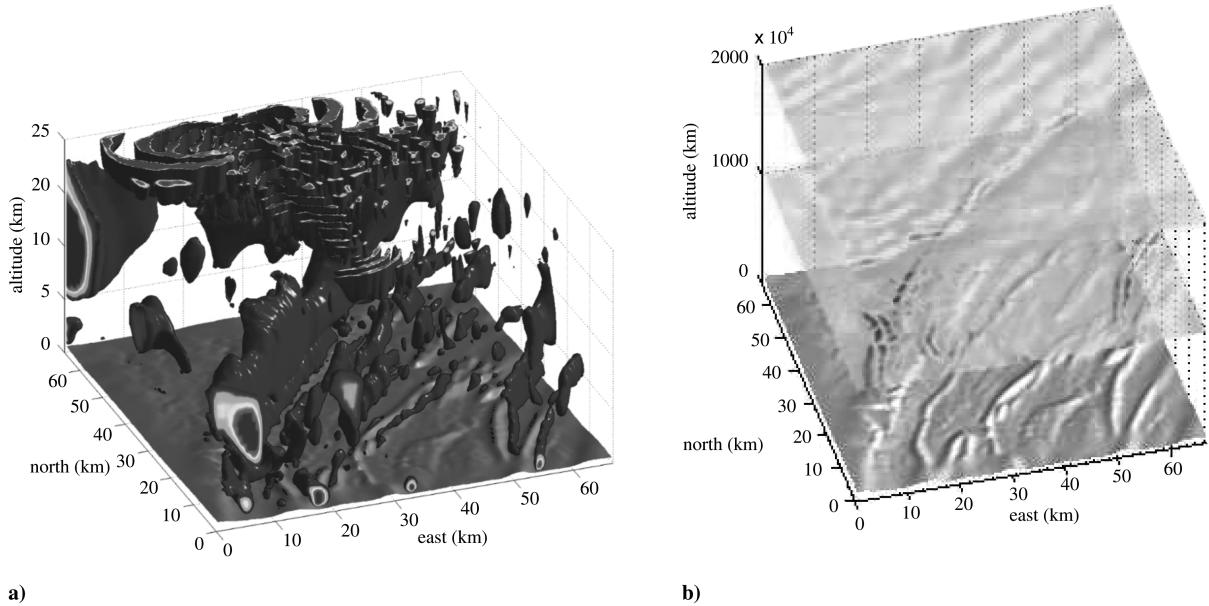


Fig. 14 Plots of a) isosurface of wind field where the vertical speed of air is greater than 0.56 m/s (minimum sink rate of the glider) and b) contours showing magnitude of vertical component of wind.

B. Regenerative Soaring and the Heuristic Search

Enabling regenerative soaring changes results significantly. Now that energy can be gained from the environment, flight paths that exploit energy become significantly more favorable, and the degree of detour (and the value of the detour) is significantly affected by the weighting α .

Figure 12a shows flight paths as α is varied from 0 to 1. Increasing the cost of energy expended causes paths to take progressively longer detours. Eventually (when $\alpha \approx 0.24$), the net energy expenditure is zero, whereas the distance traveled has not increased significantly (Fig. 12b). Further increasing α provides greater reduction in net energy expended, until regeneration allows the battery to become fully charged during the flight. This implies that the aircraft can begin the flight with almost zero charge (only enough to reach the first ridge is required), and it will reach the goal almost fully charged (a small amount must be expended to reach the goal from the near ridge at $x = 4$ km).

There is a small range of α from 0.5 to 0.6 at which the cost function used by the heuristic search appears ill-conditioned. Here, an increase in α *increases* the required energy and *decreases* the distance traveled; the opposite of the behavior observed outside of this range. Note that this only occurred in cases in which regeneration is employed. This behavior was observed for all starting points; it is likely due to the cost function becoming inadmissible over this range of α . In the context of A^* , an inadmissible function is one that is not guaranteed to underestimate the cost-to-go. As a consequence, the A^* algorithm is not guaranteed to return an optimal path [27]. One must therefore take care in choosing a cost function; using time-to-goal rather than energy-to-goal may also result in an inadmissible function, since flight speeds can be considerably higher in regions of upward-moving air (or with a tailwind) than in regions with no wind or unfavorable wind.

Figure 13 shows energy expended, velocity, and thrust coefficient along the flight path for three sample values of α . Since regeneration

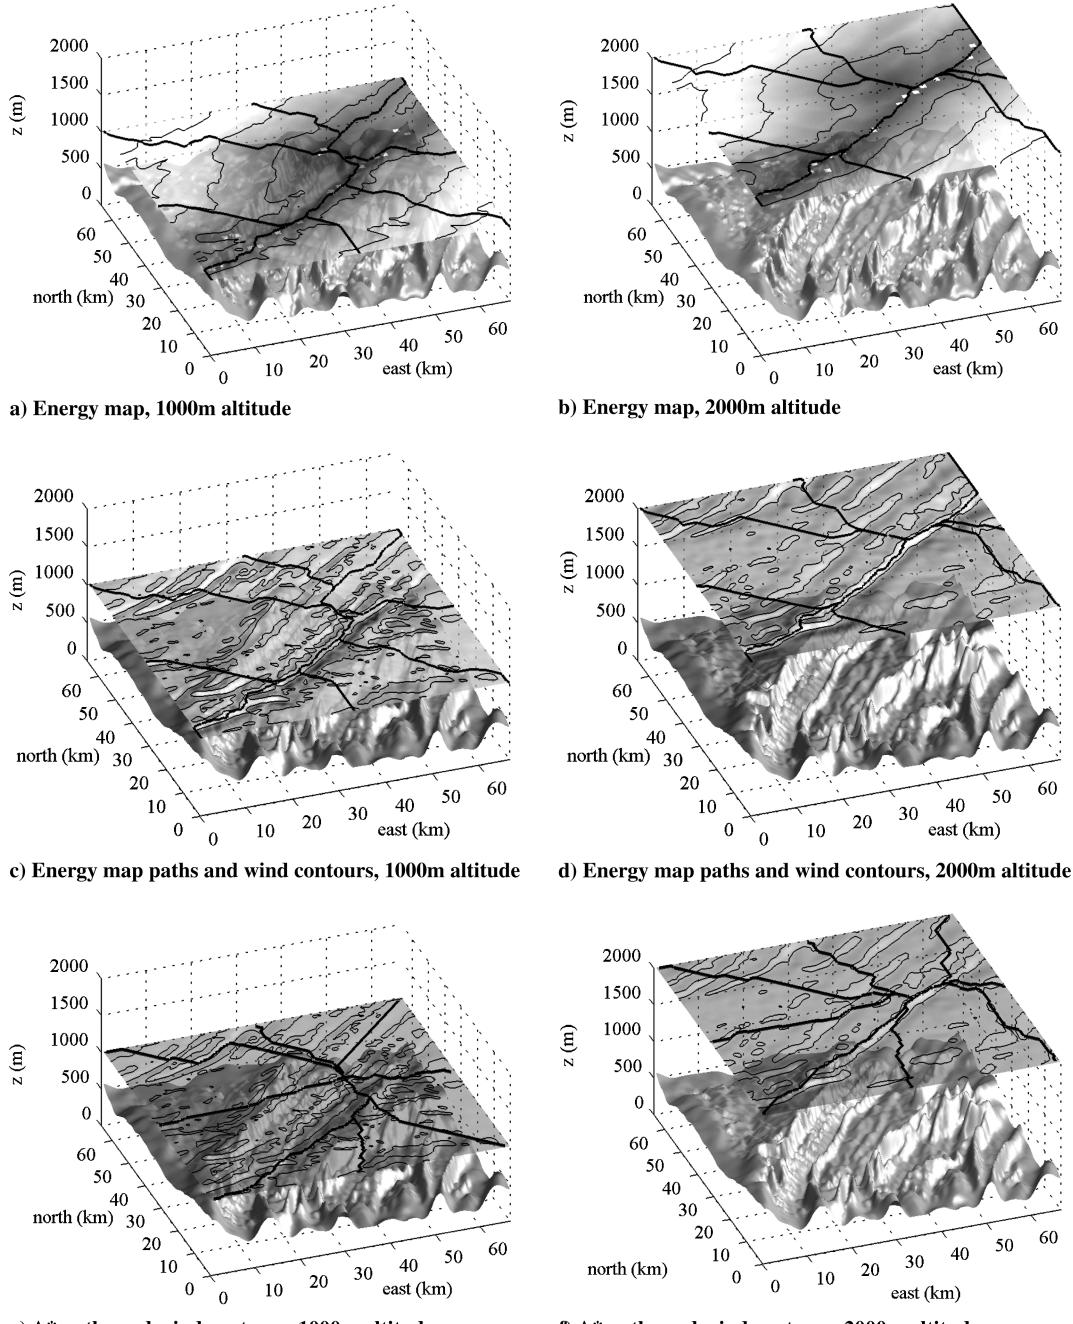


Fig. 15 Flight paths for energy map and A^* ; paths at 1000 m altitude (left) and paths at 2000 m altitude (right).

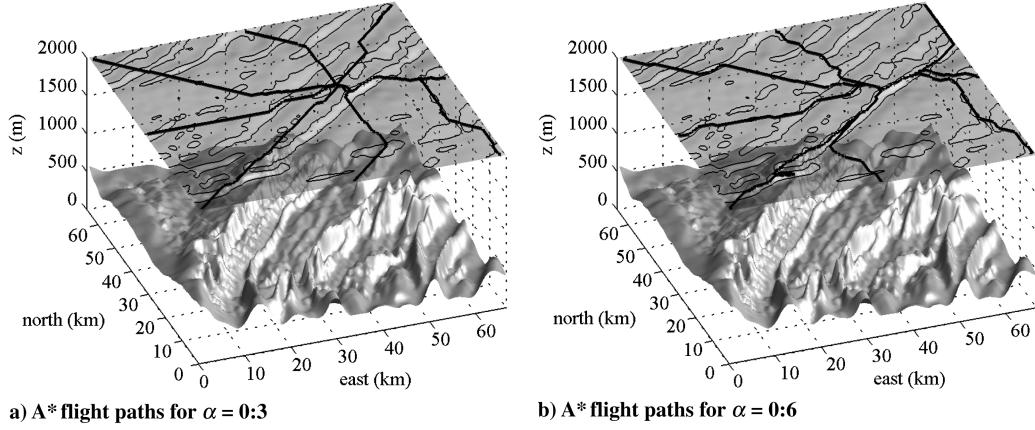


Fig. 16 Effect of varying α on A^* -computed flight paths. Note the detour into regions of favorable wind for $\alpha = 0.6$.

is now employed (computed as part of the speed-to-fly), even the case in which zero weight is placed on energy expenditure shows regeneration occurring. Areas where the slope of Δe vs distance (the top plot) is negative shows where regeneration occurs. It is clear that for the $\alpha = 1$ case the battery reaches full charge, but this results in a flight of over 300 km in length, over twice as far as when energy and distance are balanced more evenly.

It is thus clear that balancing energy expenditure and distance-to-goal is critical in obtaining useful paths. When the objective is endurance, placing significant weight on energy expenditure is likely to give best results, but the current focus on travel to the goal setting $\alpha = 0.48$ gives good results for both nonregenerative and regenerative soaring. This value will be used for all further examples.

VII. Comparing the Energy Map and the Heuristic Search in Realistic Wind Fields

Real wind fields can be extremely complex, exhibiting significant temporal and spatial variation. Trajectory generation to exploit atmospheric energy becomes correspondingly complex. To show the utility of the approaches described here, a high-fidelity simulation of a wind field is used as a test case. Wind-field data were generated using WRF-ARW (Weather Research and Forecasting-Advanced Research WRF) version 2.2 and simulate the development of ridge lift and mountain wave over central Pennsylvania [28]. The wind field's structure is shown in Fig. 14. The isosurfaces in Fig. 14a bound regions where energy can be gained from the atmosphere; the contour slices of Fig. 14b show the vertical component of wind speed at two altitudes: 1000 and 2000 m above sea level.

Paths generated using the energy map and the heuristic search are compared, and the sensitivity of heuristic-search paths to changes

in weight α will be shown. The goal is located at $(x, y) = (40 \text{ km}, 45 \text{ km})$.

Wind-field data are available at a node spacing of 444 m, and the grid used for the wind field is used as the graph for planning. At a flight speed of 35 m/s it takes approximately 12 s to fly along an edge between two nodes; this is approaching the minimum flight time required for the assumption of steady flight to be valid. At a speed of 16 m/s the time is approximately 28 s, which is significantly larger than the expected time constant of 1–2 s associated with step changes in airspeed.

Energy maps were generated for two altitudes: 1000 and 2000 m. Using the heuristic search with $\alpha = 0.48$, paths were generated from eight starting points distributed around the boundary of the environment. Results are shown in Fig. 15. The plots show terrain with vertical relief exaggerated due to scaling, energy-map contours, flight paths (Figs. 15a and 15b), energy-map-derived flight paths superimposed on wind contours (subplots 15c and 15d), and heuristic-search-computed flight paths superimposed on wind contours (subplots 15e and 15f).

The energy-map plots in Fig. 15 for 1000 and 2000 m are similar, indicating that for this case, altitude does not have a large effect on the energy required to reach the goal (note that this is not generally true). Start position does have a large effect on energy required (in Figs. 15a and 15b, lighter shades indicate greater energy required).

Paths clearly move through regions of upward air motion, with a preferred corridor that runs diagonally from the southwest at $(x, y) = (0, 0)$ to the northeast at $(x, y) = (65, 65)$ (Figs. 15c–15f). Flight paths computed using the energy map are superficially similar to the A^* flight paths; however, close examination shows that the energy-map paths follow regions of favorable wind more effectively than the A^* paths.

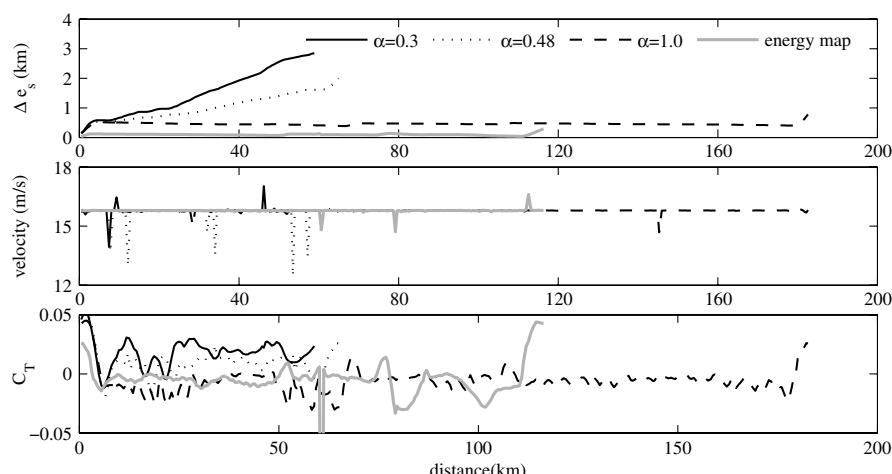


Fig. 17 Flight-path history starting from $(x, y) = (0, 0)$. In this case, the energy map defines the energetically most efficient path.

Varying the weight parameter α in the A^* cost function shows the sensitivity of the solution to changes in the importance of energy expenditure. Paths are shown in Fig. 16; careful examination shows important differences in which a detour allows some energy gain. Increasing α to 1 results in a path that looks identical to the $\alpha = 0.6$ path at the scale of the plot, but that meanders significantly at smaller scale to maximize energy gain.

Comparing the path history of the energy-map-computed paths and A^* -computed paths shows both the benefits of each approach and the effect of varying the cost of energy expenditure. This is shown in Fig. 17 for paths at a 2000 m altitude starting from (0,0). Small values of α lead to energetically inefficient paths that lead quickly to the goal; increasing α results in paths that are energetically more efficient, but fly over a longer distance. In this case, the energy-map path is most efficient energetically, requiring very little initial energy to reach the goal.

Note that in this case, even with $\alpha = 1$, the A^* approach does not compute the energetically optimal path to the goal; the energy map gives a lower-energy path (with less distance traveled). Since A^* is only guaranteed to give optimal paths if the cost function is admissible (meaning that the cost-to-go is always underestimated), this indicates that the cost-to-go is sometimes overestimated for this wind field. However, since A^* is guaranteed to find a path if a feasible one exists, it is still useful for comparison purposes.

The main difference between A^* and the energy map for this path occurs at the beginning of the trajectory. The energy-map path goes quickly to a region of rising air, where the aircraft can fly with nearly zero energy expenditure. In contrast, the A^* paths lead more toward the goal before finding the ridge.

VIII. Conclusions

This paper introduces the energy map for planning energy-efficient trajectories through a complex wind field. This is a graph-based approach to path planning, defining a path through a sequence of nodes (waypoints) in the environment. The speed-to-fly over an edge in the graph is computed by minimizing the energy expenditure for that segment, including the effects of the three-dimensional wind field. Next, wave-front expansion from the goal is used to compute minimum-energy paths. Each waypoint (node in the graph) encodes the cumulative minimum total energy required to reach the goal, the next node to be traversed to the goal, the speed-to-fly, and the required heading. The total energy required to reach the goal defines the energy map, and it is an upper bound on the minimum energy required to reach the goal from anywhere in the environment.

Regeneration (conversion of potential or kinetic energy to stored electrical energy using a windmilling propeller or ram air turbine) is included in the computation of the speed-to-fly; thus, the net required energy for flying a path to the goal can be computed. If there is sufficient upward-moving air (in the form of orographic or wave lift), it is possible to reach the goal with negative net energy cost (i.e., energy gain).

The major constraint in generating the energy map is that paths must always proceed toward the goal. As a result, a lower-energy trajectory may exist that proceeds away from the goal for some portion of the path. The computation of minimum energy to fly between waypoints is used with a heuristic planner (A^*) to determine the effect of the progress-to-goal constraint. Although paths that are energetically more efficient can be found in some cases, the cost is often a greatly increased distance flown to the goal. Moreover, defining an admissible cost function for the A^* algorithm is difficult and may be dependent on the particular wind field. The energy map thus defines a useful means of trajectory planning. Paths obtained using the energy map could also be used as an initial guess for subsequent trajectory optimization.

Acknowledgments

This research was funded by the National Science Foundation under grant IIS-0746655. The mountain wave wind-field data were provided by George S. Young, Brian J. Gaudet, Nelson L. Seaman,

and David R. Stauffer of the Pennsylvania State University, Department of Meteorology.

References

- [1] Webb, D. C., Simonetti, P. J., and Jones, C. P., "Slocum: An Underwater Glider Propelled by Environmental Energy," *IEEE Journal of Oceanic Engineering*, Vol. 26, No. 4, Oct. 2001, pp. 447–452. doi:10.1109/48.972077
- [2] Sherman, J., Davis, R. E., Owens, W. B., and Valdes, J., "The Autonomous Underwater Glider 'Spray,'" *IEEE Journal of Oceanic Engineering*, Vol. 26, No. 4, Oct. 2001, pp. 437–446. doi:10.1109/48.972076
- [3] Eriksen, C. C., Osse, T. J., Light, R. D., Wen, T., Lehman, T. W., Sabin, P. L., Ballard, J. W., and Chiodi, A. M., "Seaglider: A Long Range Autonomous Underwater Vehicle for Oceanographic Research," *IEEE Journal of Oceanic Engineering*, Vol. 26, No. 4, Oct. 2001, pp. 424–436. doi:10.1109/48.972073
- [4] Woolsey, C., Hagerman, G., and Morrow, M., "A Self Sustaining Boundary Layer Adapted System for Terrain Exploration and Environmental Sampling," Phase I Final Report, NASA Institute for Advanced Concepts, Atlanta, 2005.
- [5] Langelaan, J. W., "Tree-Based Trajectory Planning to Exploit Atmospheric Energy," *Proceedings of the American Control Conference*, IEEE, Piscataway, NJ, June 2008, pp. 2328–2333.
- [6] MacCready, P. B., Jr., "Optimum Airspeed Selector," *Soaring*, Jan.–Feb. 1958, pp. 10–11.
- [7] Cochrane, J. H., "MacCready Theory with Uncertain Lift and Limited Altitude," *Technical Soaring*, Vol. 23, No. 3, July 1999, pp. 88–96.
- [8] Reichmann, H., *Cross-Country Soaring*, Thomson, Santa Monica, CA, 1978.
- [9] Arho, R., "Optimal Dolphin Soaring as a Variational Problem," *Organisation Scientifique et Technique Internationale du Vol à Voile*, Publ. 13, 1974.
- [10] Metzger, D. E., and Hedrick, J. K., "Optimal Flight Paths for Soaring Flight," *Journal of Aircraft*, Vol. 12, No. 11, 1975, pp. 867–871. doi:10.2514/3.59886
- [11] Sandauer, J., "Some Problems of the Dolphin-Mode Flight Technique," *Organisation Scientifique et Technique Internationale du Vol à Voile*, Publ. 15, 1978.
- [12] de Jong, J. L., "The Convex Combination Approach: A Geometric Approach to the Optimization of Sailplane Trajectories," *Organisation Scientifique et Technique Internationale du Vol à Voile*, Publ. 16, 1981, pp. 182–201.
- [13] Pierson, B. L., and Chen, I., "Minimum Altitude Loss Soaring in a Specified Vertical Wind Distribution," *Science and Technology of Low Speed and Motorless Flight*, CP-2085, edited by P. W. Hanson, NASA, Hampton, VA, March 1979, pp. 305–318.
- [14] Sander, G., and Litt, F. X., "On Global Optimal Sailplane Flight Strategy," *Science and Technology of Low Speed and Motorless Flight*, CP-2085, edited by P. W. Hanson, NASA, Hampton, VA, March 1979, pp. 355–376.
- [15] Allen, M. J., "Autonomous Soaring for Improved Endurance of a Small Uninhabited Air Vehicle," 43rd AIAA Aerospace Sciences Meeting and Exhibit, AIAA Paper 2005-1025, Reno, NV, Jan. 2005.
- [16] Allen, 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, Reno, NV, Jan. 2007.
- [17] Edwards, D. J., "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," AIAA Guidance, Navigation, and Control Conference, AIAA, Paper 2008-7244, Reston, VA, Aug. 2008.
- [18] Sachs, G., and Mayrhofer, M., "Shear Wind Strength Required for Dynamic Soaring at Ridges," *Technical Soaring*, Vol. 25, No. 4, Oct. 2001, pp. 209–215.
- [19] Zhao, Y. J., "Optimal Patterns of Glider Dynamic Soaring," *Optimal Control Applications and Methods*, Vol. 25, No. 2, 2004, pp. 67–89. doi:10.1002/oca.739
- [20] Rubio Torroella, J. C., "Long Range Evolution-Based Path Planning for UAVs Through Realistic Weather Environments," M.S. Thesis, University of Washington, Seattle, WA, 2004.
- [21] Jardin, M. R., and Bryson, A. E., "Neighboring Optimal Aircraft Guidance in Winds," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 4, 2001, pp. 710–715. doi:10.2514/2.4798
- [22] Stentz, A., "The Focussed D^* Algorithm for Real-Time Replanning," *Proceedings of the International Joint Conference on Artificial Intelligence*, Aug. 1995.

- [23] Anisi, D. A., Robinson, J. W. C., and Ogren, P., "On-Line Trajectory Planning for Aerial Vehicles: A Safe Approach with Guaranteed Task Completion," AIAA Guidance, Navigation, and Control Conference, AIAA Paper 2006-6107, Keystone, CO, Aug. 2006.
- [24] Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C., *Introduction to Algorithms*, 2nd ed., MIT Press, Cambridge, MA, 2001.
- [25] Latombe, J.-C., *Robot Motion Planning*, Kluwer Academic, Norwell, MA, 1991.
- [26] MacCready, P. B., Jr., "Regenerative Battery Augmented Soaring," *Technical Soaring*, Vol. 23, No. 1, Jan. 1999, p. 28.
- [27] LaValle, S. M., *Planning Algorithms*, Cambridge Univ. Press, New York, 2006.
- [28] Young, G., Seaman, N. L., Gaudet, B., and Stauffer, D. R., "Interaction of a Mountain Lee Wave with a Basin Cold Pool," *13th Conference on Mesoscale Processes*, American Meteorological Society, Chicago, 2009.