Unmanned Aircraft Guidance for Penetration of Pre-Tornadic Storms

Jack Elston * , Eric Frew†

University of Colorado, Boulder, CO, 80309, USA

This paper investigates unmanned aircraft guidance for a severe storm penetrator in order to assess the feasibility of such a mission and to optimize ingress/egress in terms of flight time and exposure to precipitation. Understanding guidance-layer behavior and deriving optimal control strategies will provide tools for developing mission-level concepts of operation and for making vehicle design tradeoffs. A backward propagating wavefront algorithm based on ordered upwind methods is developed for guidance-layer planning of severe storm penetration. Target areas for investigation by the aircraft are identified within simulated storm wind, hail, and precipitation data. The simulated data is used to determine the feasibility of storm penetration to the points of interest, along with calculating airframe exposure to precipitation. Intelligent ingress is also considered for a dynamic wind field using a receding horizon control approach to adapt the flight path in response to environmental changes.

I. Introduction

According to the National Weather Service, 438 fatalities, close to 3000 injuries, and $11.2 billion in damages were incurred due to severe weather in the United States in 2003.[1] These losses could be dramatically reduced with effective advanced prediction and warning systems. Tornadoes are especially violent members of the severe storm family and thus the study of tornado formation and evolution is a public safety necessity. There was an immediate improvement in tornado warning capabilities with the introduction of Doppler radar.[2] However, primitive sensing methods such as human storm chasers and spotters remain the most vital part of the public safety system. The inability to determine the volumetric thermodynamic state of the atmosphere between the ground and the base of the mesocyclone remains a major barrier towards a deeper understanding of tornado genesis. The limitations of remote sensing are evident; one cannot remotely sense the thermodynamic field, these data can only be obtained with *in situ* sensing.

Research into tornadogenesis will not progress significantly until there are measurements of the thermodynamic and microphysical properties aloft in the vitally important rear-flank region of supercell storms. A consensus of research in the last 25 years makes it clear that a small downdraft of a few kilometers width, known as the "rear-flank downdraft" plays a causative role in tornado formation.[1,2] But recent studies have produced a quandary: surface observations from instrumented vehicles beneath this downdraft indicate that it typically arrives at the ground relatively warm and potentially buoyant compared to typical thunderstorm downdrafts, while studies of the flow in and around this downdraft suggest that it is negatively buoyant aloft. It is surmised that this negative buoyancy, if present in sufficient quantities upstream of the location of potential tornado formation, causes the rotation that is eventually reoriented and concentrated into a tornado.[1,2] Unfortunately, no in situ observations are possible in and around this downdraft. Balloons cannot ascend through strong downdrafts and these flows are much too dangerous for penetration using manned aircraft, as evidenced by the inadvertent penetration of a rear-flank downdraft.[2] To understand tornado formation scientists must obtain in situ observations of thermodynamics and the types, amounts, and concentration of hydrometers which, through phase changes, influence the thermodynamics.

Field observations of supercells and tornado formation occur every spring in the Central U.S. Currently these deployments involve two or more Doppler radar stations. The VORTEX-2 experiment currently scheduled for 2009 and 2010 will expand this to five or more mobile Doppler radars with the possible addition of

*Graduate Research Assistant, Department of Aerospace Engineering Sciences. Student Member

†Assistant Professor, Department of Aerospace Engineering Sciences. Member

several UAS.[3] In order to develop vehicle specifications and concepts of operation for a VORTEX-2 deployment, the performance of UAS in severe pre-tornadic storms must be studied. Initial operational constraints point to the use of a small UA to perform such a task. The UA must be easily portable in a road-worthy vehicle, and must be able to be deployed quickly. This mandates that the UA must be able to be transported fully assembled. Furthermore, cost constraints along with considerations for the high potential for vehicle loss also limit its size. Precedence for the use of small unmanned aircraft in the study of atmospheric convection and severe storms has been established by the Aerosonde UA.[4–6] This UA has been used in several field campaigns to study tropical storms and hurricanes.[4–6] However, to date penetration of a pre-tornadic storm, which evolves over quicker and smaller scales than tropical storms, has not been reported by any UAS.

This paper investigates the guidance-layer behavior of a severe storm penetrator in order to assess the feasibility of such a mission and to optimize ingress/egress in terms of flight time and exposure to precipitation in order to minimize overall risk. Understanding guidance layer behavior and deriving optimal control strategies will provide tools for developing mission-level concepts of operation and for making vehicle design tradeoffs. For example, aircraft speed generally increases with aircraft mass, so there is a high-level tradeoff between the size of the initial region from which an aircraft can be deployed (due to its ability to fight against strong winds) and the kinetic energy contained in the vehicle (which relates to system safety).

A. Planning Methods

Intelligent aircraft guidance in the presence of strong winds presents several challenges, particularly when maximum aircraft speeds are less than the peak wind speeds[a]. In these cases the aircraft are unable to travel directly upstream in some locations so large regions of the environment may be inaccessible depending on the initial position of the UA. Furthermore, as these wind fields vary over time, the reachable regions of the environment will change. Thus, planning algorithms must be able to account for the time varying field in the configuration space and adapt quickly in response to changes. Since the environment is dynamic and trajectories are likely to change, near-optimal algorithms that generate plans quickly are preferred to optimal ones that may take significant time.

The guidance layer planning problems considered in this paper are characterized by the presence of strong background winds, making them examples of control affine systems with drift.[7] In general, techniques such as Lie algebra can be extended to determine controllability[7] of systems with drift. On the other hand, optimal motion planning for systems with drift can be extremely challenging. Recent approaches using wavefront expansion via level set[8] and fast marching methods[9] have been applied to linear and control affine systems with drift, especially in the context of autonomous underwater vehicles operating in strong current fields.[10,11]

The wavefront expansion methods are well suited to the planning problems considered here. These methods can generate solutions quickly over discrete grid representations of the environment with the error in the approximation of the wavefront expansion decreasing as the grid size is decreased. In general these methods are not applied to aircraft because wind data is difficult to obtain. However, the Doppler radar systems[3] deployed to study severe storms provide this information at discrete points in the regions of interest and can therefore be incorporated easily by the planning methods. Aircraft kinematic constraints and environmental constraints due to terrain can also be included in these methods.[10,11]

Wavefront expansion algorithms for path planning and optimal control use breadth-first search over a discretized space to approximate the propagation of continuous wavefronts corresponding to solutions of Hamilton-Jacobi equations. A goal region (or point) is indicated and assigned a small value. Each grid point adjacent to the goal is then updated based upon the value of other adjacent points and a penalty function such as Manhattan or Euclidean distance.[12,13] Once the entire space has been evaluated, a path can be generated from any space on the grid to the goal using a gradient descent algorithm. Extensions to the basic propagation steps can account of obstacles and other constraints in the configuration space.[14] These types of planning methods require that the entire space be known *a priori* in order to generate a path, but only need be called once in a static environment in order to calculate the optimal trajectory from every point in the environment.

Most wavefront expansion methods for path planning through current fields utilize a forward propagating wavefront planner.[10,11] This is well suited for applications that use small amounts of data and are constantly

[a]the term *strong wind field* is used to denote a field where the peak wind value is greater than the maximum speed of the aircraft

obtaining only local knowledge of the current field with some distant global objective. When considering the problem of storm ingress with a UA, Doppler radar coverage can give a low resolution global wind field. This enables the generation of a globally optimal ingress path using forward wavefront propagation. Unfortunately, real world application of the algorithm will result in differences between the actual and expected trajectories, and will require recalculation of the cost map. Therefore the approach used in this paper utilizes backward wavefront propagation to calculate optimal time-to-go contours from the goal region. The resulting contour map provides information on the feasibility of storm penetration from given locations and yields the optimal trajectory via gradient descent.

The method of wavefront expansion is essential to the optimality of the resulting cost map for use in control. Some methods for cost propagation unnecessarily limit directions of travel due to the discretized nature of the wind data.[10,11] Interpolation can aid in yielding the entire set of feasible paths, but this becomes computationally expensive. A method such as a *sliding wavefront expansion*[11] can be used for forward wavefront propagation, but this method also becomes computationally expensive when used for backward propagation. An approach using *ordered upwind methods*[15] seems to be well suited to this application. It accommodates the previously mentioned limitations by continuously refining cost values on the fringe of the wavefront.

Several fundamental issues remain to be addressed before existing methods developed for underwater vehicles can be applied to aircraft guidance. First, cost functions are needed to represent the ability of the robot to transition from one point in the environment to another in the face of the strong winds. This is nontrivial for complex aircraft dynamics and situations where infeasible steps exist.[10,11] Second, existing methods have only been applied to static fields.[10,11] Time-varying fields complicate the wavefront expansion since the goal region can take on a distorted shape in the configuration space. Third, the incorporation of constraints in backward wavefront propagation is difficult. Important constraints in this work that are not considered under simple wavefront approaches include exposure to precipitation or maximum/minimum elevation gains (e.g., due to terrain or regulatory limits).

This paper presents a backward propagating wavefront algorithm based on ordered upwind methods[15] for guidance-layer planning of severe storm penetration by unmanned aircraft. Feasibility analysis is performed on a static wind field to assess baseline performance expectations. Intelligent ingress is then considered for a dynamic wind field using a receding horizon control approach to adapt the flight path in response to environmental changes.

II. System Models

A. UA Dynamics

The UA model used here assumes the presence of a low-level flight control system capable of altitude-hold, speed-hold, and turn rate-command functions. For basic analysis we assume the flight control system presents to the guidance layer the standard kinematic model:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} V\cos\psi + W_x(x,y,z,t) \\ V\sin\psi + W_y(x,y,z,t) \\ W_z(x,y,z,t) \\ u \end{bmatrix} \qquad |u| \le \omega_{max} \tag{1}$$

where V is the known aircraft speed, u is the turn rate command, ω_{max} is the maximum allowable turn rate, and $[W_x, W_y, W_z]^T$ are the position- and time-dependent wind components derived from storm data given at discrete grid points. For UA locations between grid points, cubic interpolation (using MATLAB's interpn.m) is used to determine the wind speed. Although vertical motion is considered in this work, only planar (i.e. in the x-y plane) motion is controlled. Therefore, to simplify notation let the aircraft planar velocity be $\mathbf{v} = [\dot{x}, \dot{y}]^T$, let the planar wind velocity be $\mathbf{v}_w = [W_x, W_y]^T$ with $V_w = \|\mathbf{v}_w\|$, and let the aircraft velocity relative to the wind be $\mathbf{v}_a = [v_{a,x}, v_{a,y}]^T = [V\cos\psi, V\sin\psi]^T$. Note that $V = \|\mathbf{v}_a\|$ and $\mathbf{v} = \mathbf{v}_a + \mathbf{v}_w$.

B. Storm Data

Unmanned aircraft performance is investigated through flight in several (simulated) characteristic severe storms. The storm data used here were provided by Jerry Straka of the School of Meteorology at the

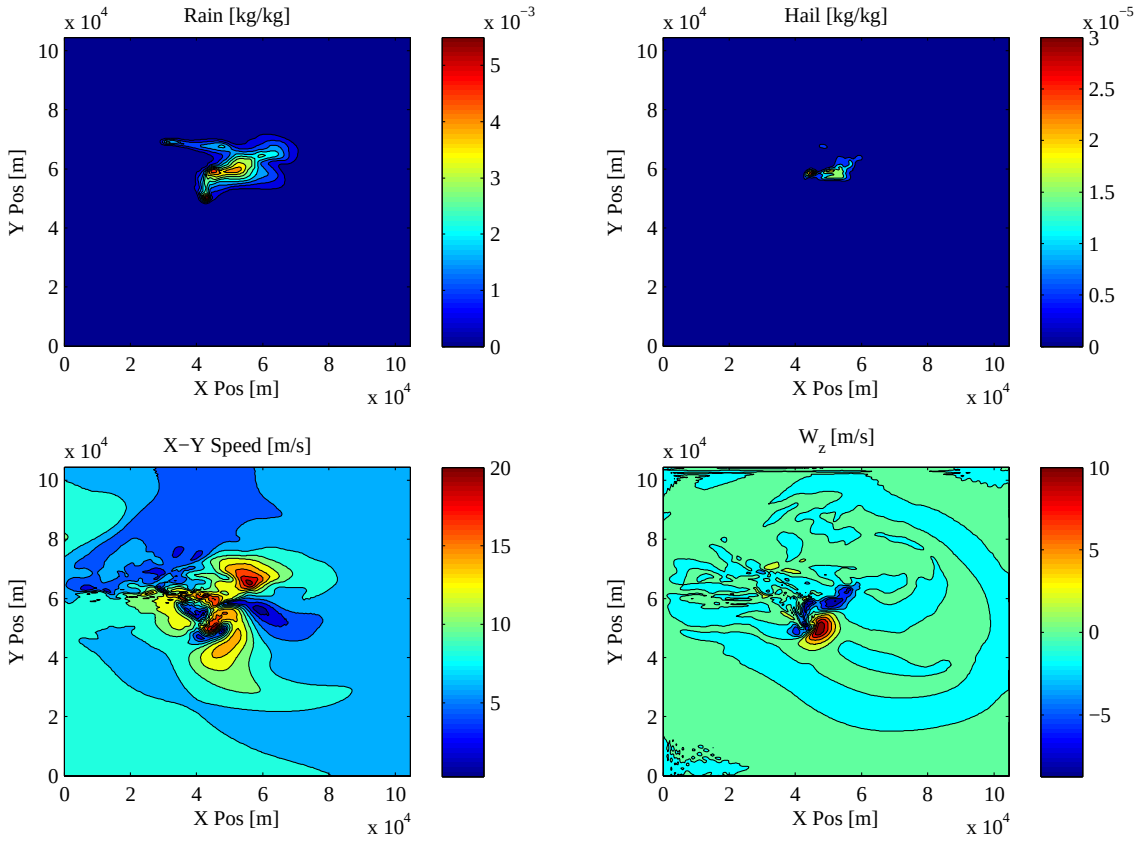

Figure 1. Contours of rain, hail, 2-D speed, and z-component of wind speed for the entire simulation region.

University of Oklahoma.[16–18] The set is planar two-dimensional data for a single instant in time for an altitude of 500 meters. Storm data is given on a 210 by 210 grid with a resolution of 500 meters between grid points in a reference frame that moves with the storm (with velocity [18, 4] m/s). Data includes rain, hail, and three-dimensional wind velocity components (relative to the storm frame of reference). The storm simulation is aligned such that the positive y-axis corresponds to North. Figure 1 shows contours plots of the storm data including rain, hail, planar (x-y) wind speed, and the z-component of wind vector. Figure 2 shows a close up of the main portion of the rain contour overlaid with planar wind velocity vectors.

The rear-flank downdraft, which is the goal region for pre-tornadic storm penetration (see Fig. 2), is characterized by several traits, with some observable in the simulation data. First, as the name implies, the wind velocity in the region has a strong downward component (Fig. 1d). Second, the region sits behind a gust front which is observable in the wind data (Fig. 2). Third, the region contains a large amount of rain (also observable in Fig. 2). Finally, the rear-flank downdraft sits behind a Cumulonimbus cloud base that is not visible in the data, but is visible in the field (and marked on Fig. 2). These traits of the rear-flank downdraft allow both the operator and the intelligent ingress system to identify and track the moving goal region during the penetration mission. For the work presented here it is assumed that the goal region does not move in the storm frame of reference.

III. Path Planning using Wavefront Expansion Methods

Guidance layer path planning algorithms for storm penetration are developed based on ordered upwind methods for wavefront propagation. Ordered upwind methods (OUM) are a class of non-iterative algorithms that solve continuous wavefront propagation problems on a discretized grid representation of the configuration space. These Dijkstra-like schemes maintain an ordering of points and systematically compute the solution by relying on known, previously computed information which lies "upstream".[15] In the context of path

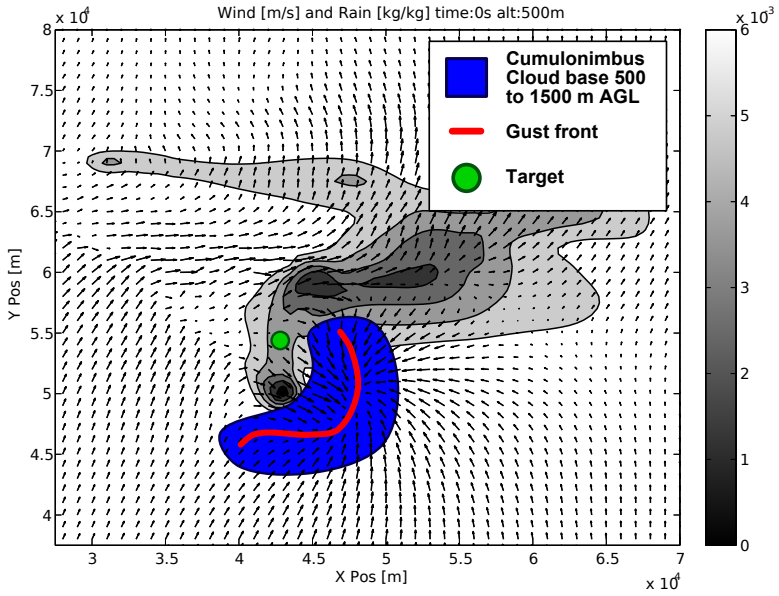

Figure 2. Contours of rain with planar wind vectors around region of interest.

planning and optimal control, OUM is used to approximate the solution of the Hamilton-Jacobi equations that describe optimal aircraft trajectories. The ordered upwind methods use partial information about the characteristic directions of the wavefront propagation to decouple the nonlinear systems described by the Hamilton-Jacobi equations, producing one-pass algorithms of greatly reduced computational labor compared to other numerical techniques.[15]

Application of OUM to the problem of severe storm penetration is complicated by two primary factors. First, the presence of a strong wind field can restrict the motion of the aircraft. There will be regions in the environment where the aircraft cannot head against the wind. Second, the wind field leads to anisotropic drift terms that complicate the wavefront expansion. Fast Marching methods for wavefront expansion are designed for isotropic wavefronts that expand with equal speed in all directions. When these methods can be applied they are very efficient and computationally fast. Unfortunately they cannot be applied directly here. These two issues are addressed by the proper development of a time penalty cost function) and wavefront expansion algorithm, respectively.

A. Cost Function

Utilizing a cost function that ensures that all resulting paths are feasible is pertinent to the success of the presented mission. Previous cost functions, while acceptable given there are no strong winds ($V_a >> V_w$), yield infeasible paths in strong winds.[10,11] Given the design constraints imposed on the unmanned aircraft, it is assumed that the environment encountered by a UA when approaching a pre-tornadic storm will contain winds that will likely be larger than the maximum air speed of the vehicle.

To ensure that all generated paths are feasible, travel time for the aircraft to reach a target point is used as the cost metric.[11] To generate the cost function it is assumed that the aircraft can move in any direction with speed V and that a path following controller can be implemented such that the aircraft described by (1) can stay on the resulting path. Given a goal point a distance $\mathbf{d}$ away, aircraft velocity $\mathbf{v}_a$, and wind field velocity $\mathbf{v}_w$ the travel time τ is related to these variables by $\mathbf{d} = (\mathbf{v}_w + \mathbf{v}_a) \cdot \tau$. Projecting this on the x and y axis yields

$$\begin{aligned} d_x &= (v_{a,x} + W_x) \cdot \tau \\ d_y &= (v_{a,y} + W_y) \cdot \tau \end{aligned} \tag{2}$$

Aircraft speed components $v_{a,x}$ and $v_{a,y}$ can be eliminated by substituting $V_w^2 = v_{a,x}^2 + v_{a,y}^2$ into (2) to give

$$(d_x - W_x \cdot \tau)^2 + (d_y - W_y \cdot \tau)^2 = v_w^2 \cdot \tau^2. \tag{3}$$

Solving this quadratic equation for travel time yields:

$$\tau = \frac{-(W_x \cdot d_x - W_y \cdot d_y) + \sqrt{\Delta}}{v_w^2 - c^2} = \frac{\sqrt{\Delta} - \langle \mathbf{d} \cdot \mathbf{c} \rangle}{v_w^2 - c2} \tag{4}$$

where $\Delta = v_w^2 \cdot (d_x^2 + d_y^2) - (W_x \cdot d_y - W_y \cdot d_x)^2$. Considering the case where $V_a = V_w$, the equation can be simplified to:

$$\tau = \frac{d^2}{2\langle \mathbf{d} \cdot \mathbf{c} \rangle} \tag{5}$$

The time τ is only defined when $\Delta \geq 0$. The union of the set where τ is undefined, and when $\tau < 0$ identifies infeasible regions of travel.

Precipitation can also be accounted for using the cost function. It is assumed that exposure to precipitation will have a upper bound for safe aircraft flight, rather than a cumulative effect. Given this assumption, the rain can be treated as an obstacle, and a rain term is added to the cost function:

$$\tau = \frac{\sqrt{\Delta} - \langle \mathbf{d} \cdot \mathbf{c} \rangle}{v_w^2 - c2} + \alpha_r \tag{6}$$

$$\begin{cases} \alpha_r(x, y, z, t) = 0 & : rain(x, y, z, t) \leq rain_{max} \\ \alpha_r(x, y, z, t) = \infty & : rain(x, y, z, t) > rain_{max} \end{cases} \tag{7}$$

Should the case present itself where it is also advisable to avoid precipitation when possible, a second rain term, β_r can be added to the cost function:

$$\tau = \frac{\sqrt{\Delta} - \langle \mathbf{d} \cdot \mathbf{c} \rangle}{v_w^2 - c2} + \alpha_r + \beta_r \tag{8}$$

$$\beta_r(x, y, z, t) = k_{rain} \cdot rain(x, y, z, t) \tag{9}$$

B. Wavefront Expansion

Wavefront expansion is performed using ordered upwind methods to propagate cost-go-to values over the discretized configuration space of the unmanned aircraft. Unlike standard Fast Marching expansion methods, OUM accounts for the anisotropic nature of the wind field during wavefront propagation. Ordered upwind methods are implemented via finite difference update rules that are proven to converge to the viscosity solution of the Hamilton-Jacobi PDE corresponding to the optimal control problem, defined here based on the cost function (8).[15]

The optimal cost-to-go map is generated by dividing the considered space into three sets of points, *Far*, *Considered*, and *Accepted*.[b] Figure 3 shows a division of the configuration space into the various sets during a calculation of the cost-to-go map. The *AcceptedFront* is defined to be the set of *Accepted* points adjacent to at least one member of the set of *Considered* points. Let AF be the set of line segments $\mathbf{x}_j\mathbf{x}_k$ where $\mathbf{x}_j$ and $\mathbf{x}_k$ are adjacent grid points on the *AcceptedFront* such that there is a *Considered* grid point $\mathbf{x}_i$ adjacent to both points. For each *Considered* grid point $\mathbf{x}_i$ the set of points on the *AcceptedFront* "near" $\mathbf{x}_i$ is defined

$$NF(\mathbf{x}_i) = \left\{ (\mathbf{x}_j, \mathbf{x}_i) \in AF | \exists \mathbf{x} \text{ on } (\mathbf{x}_j, \mathbf{x}_k) \text{ s.t. } \parallel \mathbf{x} - \mathbf{x}_i \parallel \leq h \frac{F_2}{F_1} \right\} \tag{10}$$

where h is the grid size, and F_1 and F_2 are the upper and lower bounds on the speed of the AF respectively and make up the anisotropy ratio $\Gamma = F_2/F_1$. Let $U(\mathbf{x}_i)$ represent the optimal cost-to-go from point $\mathbf{x}_i$ and let $V_{\mathbf{x}_j,\mathbf{x}_k}(\mathbf{x}_i)$ be a consistent upwinding approximation when the characteristic $\mathbf{x}_i$ lies in the simplex $\mathbf{x}_i\mathbf{x}_j\mathbf{x}_j$, which we take here to be

$$V_{\mathbf{x}_j,\mathbf{x}_k}(\mathbf{x}_i) = \min_{\zeta \in \{0,1\}} (\zeta(\tau_{ij} + U(\mathbf{x}_j)) + (1 - \zeta)(\tau_{ik} + U(\mathbf{x}_k))) \tag{11}$$

where τ_{ij} is the value of the penalty function when moving from $\mathbf{x}_j$ to $\mathbf{x}_i$, e.g. (8). Finally, let Ω be the domain representing the environment, $\delta\Omega$ be the goal region in the environment, and $q(\mathbf{x}_i)$ be the penalty for exiting the domain at $\mathbf{x}_i$ (which is taken here to be $q(\mathbf{x}_i) = 0 \ \forall \mathbf{x}_i$) .

The following is the method for generating a cost map using the ordered upwind method:

[b]The nomenclature and description presented here closely follow that presented in Ref. 15

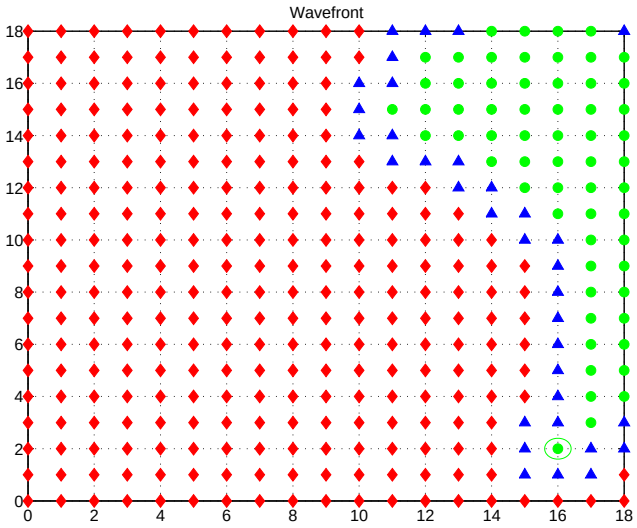

Figure 3. Set of *Far* (red) *Considered* (blue) and *Accepted* points in OUM planner partially through generation of the cost-to-go map. The storm data used is the test data set depicted in Figure 4.

1. Start with all grid points in Far ($U_i = \infty$).
2. Move the boundary grid points ($\mathbf{x}_i \in \partial\Omega$) to *Accepted* ($U_i = q(\mathbf{x}_i)$).
3. Move all the grid points $\mathbf{x}_i$ adjacent to the boundary into *Considered* and evaluate the tentative value of $V(\mathbf{x}_i) = \min_{(\mathbf{x}_j,\mathbf{x}_k)\in NF(\mathbf{x}_i)} V_{\mathbf{x}_j,\mathbf{x}_k}(\mathbf{x}_i)$.
4. Find the grid point $\mathbf{x}_r$ with the smallest value of $V(\mathbf{x}_i)$ among all the *Considered*
5. Move $\mathbf{x}_r$ to *Accepted* and update the *Accepted Front* ($U(\mathbf{x}_r) = V(\mathbf{x}_r)$).
6. Move the *Far* grid points adjacent to $\mathbf{x}_r$ into *Considered*.
7. Recompute the value for all the *Considered* $\mathbf{x}_i$ within the distance $h\frac{F_2}{F_1}$ from $\mathbf{x}_r$. If the new computed value is less than the previous tentative value for $\mathbf{x}_i$ then update U_i.
8. If *Considered* is not empty then goto 4.

Figure 4(a) shows a test case which was shown to fail in an isotropic wavefront propagation given certain vehicle speeds due to the fact that the *corridor of feasibility* does not intersect any grid points.[11] Figure 4(b) shows the optimal cost-to-go map generated from a wavefront using the ordered upwind method. It is apparent from the figure that a feasible solution can be found by descending the gradient of this map.

IV. Application

The ordered upwind methods described in the previous section can be applied directly to the set of data generated from the storm simulation. The results presented here currently ignore altitude and integrated precipitation exposure constraints, but still provide a valuable tool for determining the feasibility of UA ingress and can be effectively used for intelligent path planning during ingress.

A. Feasibility, Optimal Ingress Time, and Aircraft Design

Feasibility and optimal ingress time can be evaluated by using (4) as the penalty function in the ordered upwind expansion process. The result of the OUM using this cost gives an approximation of the optimal time-to-go from every point on the grid. The maximum nearby grid value in the cost map to any (non-grid) location represents an upper bound on the time for the aircraft to reach the goal when deployed from that

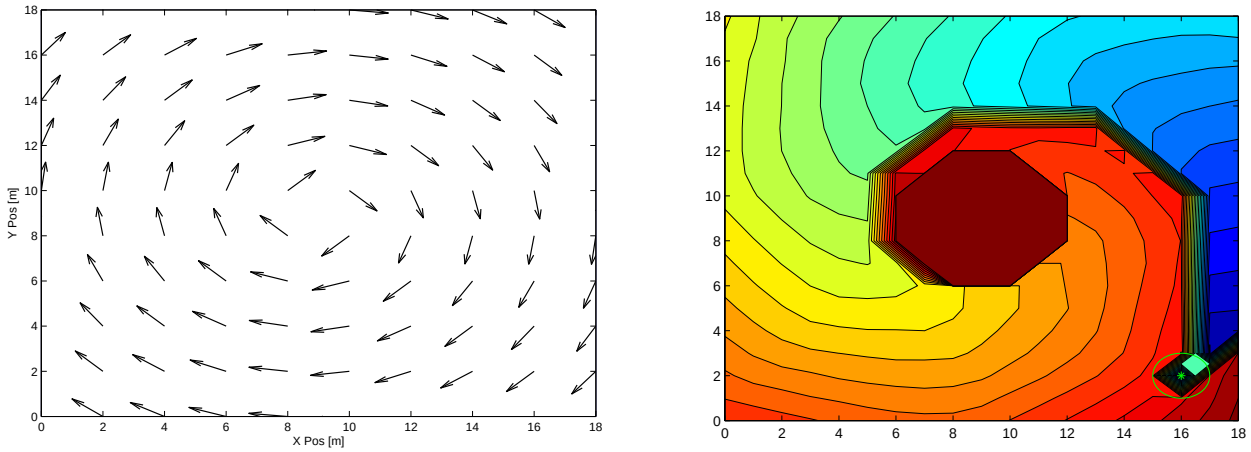

(a) (b)

Figure 4. a.Test wind field, constant speed of $30m/s$; b. Cost map obtained by using an ordered upwind method for backward wavefront expansion with $V_a = 20m/s$

location. Regions from which the aircraft cannot possibly reach the goal location will be denoted by an undefined (or very large) value in the map.

Figure 5 contains examples of the time-to-go maps generated using simulated storm data (Fig. 5a) with the goal location at $\mathbf{x}_{goal} = [4.25, 5.4]^T \cdot 10^4$ m/s and aircraft with constant air speeds of $V_a = \{7.5, 10.0, 20.0\}$ m/s (Fig. 5b-d). For a slow moving aircraft (Fig. 5b), the wind field in the storm separates the environment into two distinct regions: the southern region where the aircraft is able to reach the goal location quickly (within 10 min or less) and the northern region where optimal ingress time is significantly larger (30 minutes or greater). Note, the simulation extends beyond the borders of Fig. 5b, which just shows results near the goal region, so trajectories from the northern region can in fact flow down into the southern region. The next plot (Fig. 5c) shows that increasing the airspeed slightly to $V_a = 10.0$ m/s changes the time-to-go map significantly. Ingress paths from the southern region are opened up along corridors on the eastern and western edges. The map further shows that the western corridor is generally better since the aircraft flies with the wind field there. The center of the environment is still marked by a boundary that cannot be crossed. Finally, for air speed of $V_a = 20.0$ m/s, Fig. 5d shows that the aircraft can fight the wind from any location to reach the goal without having to go around the center of the environment. The optimal time-to-go is still varied across the environment, but all optimal paths are reasonably direct.

The optimal time-to-go maps can be used in several different ways. As a design tool, the maps can be used to evaluate the tradeoff between aircraft speed (which usually relates to mass) and mission-level performance objectives like ingress time and deployment feasibility. Figure 5 shows wide variation in the general shape of the time-to-go map, implying very different results for different air speeds. A salient example is the map with $V_a = 7.5$ m/s which shows that the entire northern half of the environment has optimal trajectories at least 3 times slower than the southern region. Thus, if mission-level constraints are going to force the aircraft to deploy north of the storm, the airspeed should be increased.

A second use of the time-to-go maps is the selection of deployment locations. The figures show that given vehicles with slower speeds, deployment location becomes more important. In some cases small movements in deployment location (less than 5 km) can significantly affect the time for the UA to reach the goal location. Consider the map for $V_a = 10.0$ m/s (Fig. 5c) which shows two corridors along the edges. Deployment from $[5.5, 6.2]^T \cdot 10^4$ m leads to a path that has to travel around the center and into the western corridor, taking approximately 15 minutes to reach the goal. Moving the deployment to the location $[5.5, 6.3]^T \cdot 10^4$ m puts the aircraft on the other side of the central obstruction and lets the aircraft head straight toward the goal region, reaching it in close to 8 minutes. From an operational perspective, this shows that the deployment team should not stop alongside the Doppler radars that must be stationary to collect wind data. Instead, the deployment vehicle should continue moving until an appropriate location is reached. The time-to-go map implies that mission-level optimization should be performed that combines time required to reach a deployment location as well as airborne time-to-go.

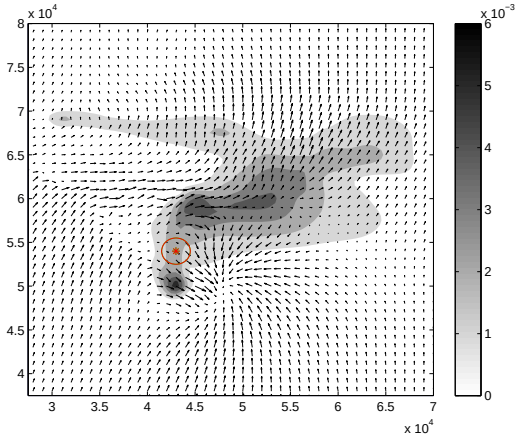

(a) Wind field and precipitation values from simulated storm data.

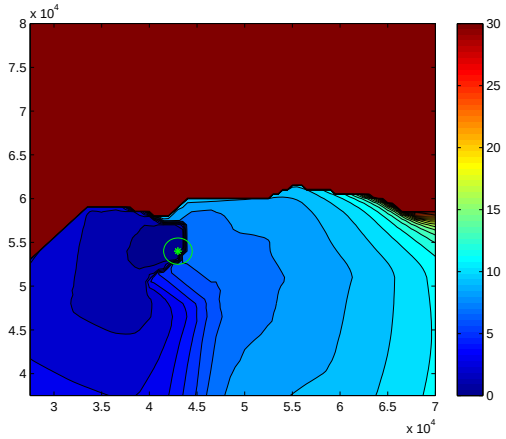

(b) Cost map obtained by using an ordered upwind method for backward wavefront expansion, target indicated by green circle. Color correlates with time in minutes required to reach goal location given $V_a = 7.5m/s$

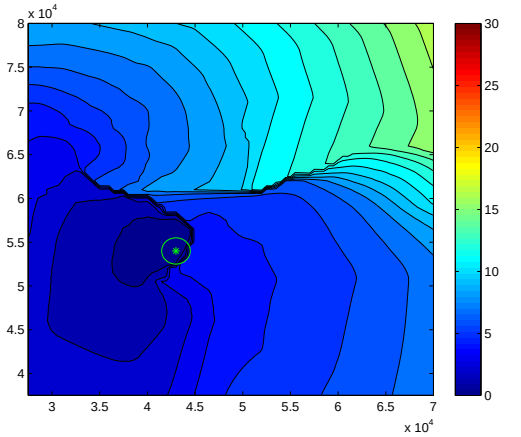

(c) Cost map obtained by using an ordered upwind method for backward wavefront expansion, target indicated by green circle. Color correlates with time in minutes required to reach goal location given $V_a = 10m/s$

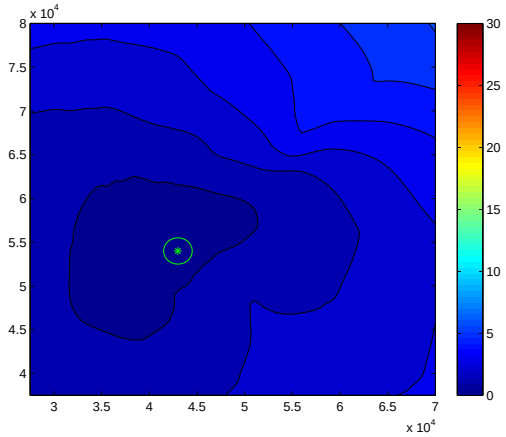

(d) Cost map obtained by using an ordered upwind method for backward wavefront expansion, target indicated by green circle. Color correlates with time in minutes required to reach goal location given $V_a = 20m/s$

Figure 5.

The time-to-go maps can also be used to derive the optimal ingress trajectory. A key trait of OUM is that the propagated wavefront is an approximation of the viscosity solution of the Hamilton-Jacobi equations associated with the optimal trajectory design problem.[15] The optimal trajectory is determined by following the gradient of the time-to-go map. Once the map is determined for a given goal and static wind field, it can be used at any time to determine the optimal trajectory from any location. From an operational perspective, this enables instantaneous response once the deployment location is determined. Having the complete map also allows for instantaneous adaptation when disturbances push the aircraft off the optimal trajectory. In fact, the next section shows how the time-to-go map can be viewed as a feedback control law that can be used for adaptive intelligent ingress.

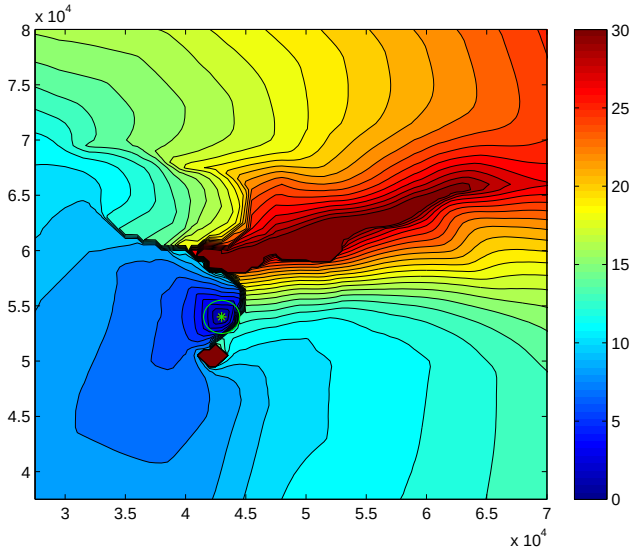

Figure 6. Cost map obtained by using an ordered upwind method for backward wavefront expansion, target indicated by green circle. Cost function includes both a rain exposure limit $\alpha_r = 4 * 10^{-3} kg/kg$ and a integrated rain exposure avoidance term $\beta_r = 500$. Color correlates with time in minutes to reach goal location given $V_a = 10 m/s$

A second example shows how the OUM can be used to account for precipitation. Figure 6 shows the resulting cost-to-go map derived using (8) with a rain exposure limit $\alpha_r = 4 * 10^{-3} kg/kg$ and a integrated rain exposure avoidance term $\beta_r = 500$. In this example the rain exposure limit creates two "obstacles" that are denoted as infeasible regions in Fig. 6. The resulting optimal paths move away from these regions before curving towards the goal. The parameter β_r was heuristically set to weight the relative importance of minimum time and minimum rain exposure for this scenario. As such, the value of the cost-to-go loses its direct physical interpretation. Nevertheless, the resulting map still yields trajectories that are optimal according to the given cost function.

B. Intelligent Ingress

Intelligent ingress can be realized using the cost maps generated previously for feasibility studies. By following the gradient of the cost map, the UA will reach the goal location in the minimal amount of time, assuming that the wind field does not change. Given the map, the optimal path to the goal point can be defined by the set of points, $\mathbf{x}_i, i = 1, \ldots, n$ where x_0 is the deployment location of the UA and:

$$\mathbf{x}_{i+1} = \mathbf{x}_i - \epsilon \nabla U(\mathbf{x}_i) \tag{12}$$

where ϵ is the step size. Although knowing the complete optimal trajectory can be beneficial for higher level reasoning, it is not necessary for aircraft guidance. Instead, the aircraft at $\mathbf{x}_i$ simply uses (12) directly to determine the local direction of its next step.

In practice the aircraft will not follow the optimal direction but will instead move to

$$\mathbf{x}_{i+1} = \mathbf{x}_i - \epsilon \nabla U(\mathbf{x}_i) + \mathbf{e}_i \tag{13}$$

where $\mathbf{e}_i$ is a disturbance vector due to the fact that the wind field measurements will have errors, the goal location may be moving, and the value function $U(\mathbf{x}_i)$ calculated by the OUM is an approximation. In this case the aircraft simply uses the time-to-go map at the new location and moves in the direction of $\nabla U(\mathbf{x}_{i+1})$.

In an operational deployment Doppler radar systems will provide wind field estimations and storm structure at periodic times. Let T_D represent the time between Doppler measurements and let U^t be the cost-to-go map determined from wind field data collected at time t. In this case it is expected that the wind field changes over time so the value function U^k for $k \leq t \leq k + T_D$ will accrue errors as the actual wind field evolves. The feedback control structure of the backward propagation solution is used as part of a receding horizon controller for intelligent ingress. In particular, over the time interval $k \leq t \leq k + T_D$ the aircraft moves in the direction of $-\nabla U^k$. At time $t = k + T_D$ a new measurements of the wind field are collected and the new cost-to-go map U^{k+T_D} is calculated.

Depending on model information available to the guidance layer, an additional term $\hat{\mathbf{e}}_i$ can be added to the update rule

$$\mathbf{x}_{i+1} = \mathbf{x}_i - \epsilon \nabla U(\mathbf{x}_i) + \mathbf{e}_i - \hat{\mathbf{e}}_i \tag{14}$$

where the estimate $\hat{\mathbf{e}}_i$ is derived from a simple update law like

$$\hat{\mathbf{e}}_{i+1} = \hat{\mathbf{e}}_i + \gamma(\mathbf{x}_{opt,i+1} - \mathbf{x}_{i+1}) \tag{15}$$

where γ is a weight used to determine the frequency response of the update and $\mathbf{x}_{opt,i+1}$ is the location the aircraft would have moved to had the wind field been exact.

So far derivation of the intelligent ingress planner has assumed that the aircraft can travel instantaneously in the direction of $-\nabla U^k$. Unfortunately the kinematic aircraft is constrained to move in the heading direction only with a bounded turn rate. In this case an additional proportional control law is wrapped around the ingress planner to calculate the aircraft turn rate $u = k_u(\psi_{opt} - \psi)$ where $\psi_{opt} = \angle(-\nabla U(\mathbf{x}_i))$ is the direction of the gradient of the cost-to-go map. Again, because the backward propagation using OUM provides $U(\mathbf{x}_i)\ \forall \mathbf{x}_i$, the resulting ingress planner performs like a feedback control law.

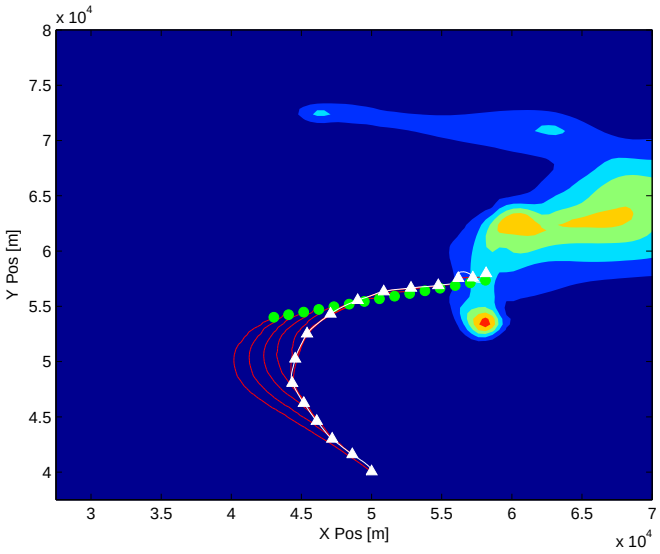

Figure 7. Rain contours of the simulated storm data at the end of the simulation. The track of the vehicle is shown in white, with triangle at the points where the "radar" data was updated. Green circles show the corresponding goal points at each update. The red lines show the calculated optimal path at the start each iteration.

A simulation was conducted using the cost-to-go map derived for the storm data. Given that the storm data is depicted in a moving reference frame, its location over time with respect to the inertial reference frame moves with a constant speed, (v_u, v_v). A linear interpolation between data points in both space and time is used to generate the "actual" wind field that affects the dynamics of the vehicle. The onboard planner only receives wind field updates at the rate of once per minute, corresponding to a simulated update

from Doppler radar information. It is from these updates that the planner calculates the cost map used for guidance for the next minute.

Figure 7 shows the track of the vehicle for the entire simulation, and the final location of the storm rain contours. The points where the onboard "radar" data was updated are depicted as white triangles. At each of these points, the path generated from the cost-to-go map is shown in red, and the corresponding goal location is shown as a green circle. Figure 8 shows the error between the expected location of the vehicle and the actual position during the simulation. The error value in the graph generally increases over time until a "radar" update is received by the controller. This validates the usefulness of the addition of an error term (which was not included in this simulation) to the controller.

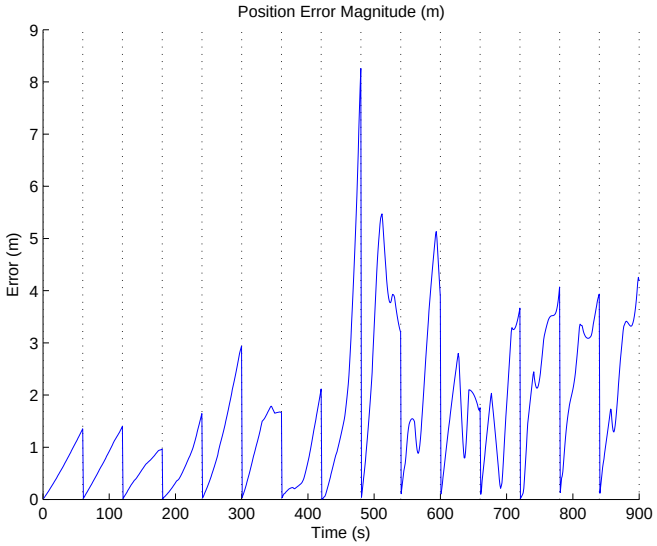

Figure 8. Error between expected and actual position during the simulation.

Figure 9 shows the altitude for the vehicle for the course of the simulation. This aspect was not controlled, but is included to demonstrate the necessity for considering an altitude limit as part of the cost function. It can been seen in the figure that the aircraft reaches the goal location $300m$ lower than its starting altitude (as a result of entering the rear-flank downdraft). This term will be difficult to be included into the cost function, as it must maintain an integral of the vertical speed over the entirety of the wavefront propagation. Absolute vertical speeds (downdrafts and updrafts) can be considered as obstacles and therefore included much more easily into the cost function.

V. Conclusion

Ordered upwind methods were used to develop guidance-layer planning algorithms for severe storm penetration by small unmanned aircraft. The ordered upwind methods are a class of non-iterative algorithms that solve continuous wavefront propagation problems on a discretized grid with significantly less computational effort than other numerical techniques. In the context of path planning and optimal control, ordered upwind methods approximate the solution of the Hamilton-Jacobi equations that describe optimal aircraft trajectories. They are well suited to the planning problems considered here since they are designed for anisotropic flow fields and because the Doppler radar systems deployed to study severe storms provide wind field information at discrete points in the environment.

A backward propagating wavefront algorithm based on ordered upwind methods was created to aid development of mission-level concepts of operation for storm penetration. In particular, the algorithm generates cost-to-go maps of the environment that i.) can be used to determine the feasibility of storm penetration from different locations; ii.) can be used as a design tool to understand aircraft speed requirements; iii.) show the sensitivity of ingress time to deployment location; iv.) yield the optimal ingress trajectories; and v.) form a component of a receding horizon intelligent ingress planner. Unlike previous work that uses forward

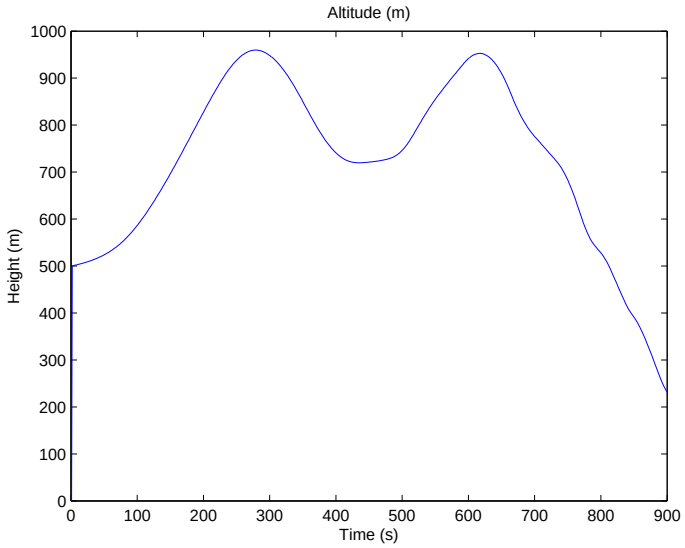

Figure 9. Altitude profile for the simulated flight.

propagating wavefront planners in current fields, the global cost-to-go maps generated here using backward propagation are used as feedback control laws to simplify the overall computational of the guidance layer inputs.

The planning algorithms were applied to simulated severe storm data in order to highlight their utility. The simulated data emulates measurements that can be taken in the field by Doppler radar systems. Several different examples showed how storm winds create distinct regions in the environment such that vehicle design and deployment location can have significant impact on overall mission performance. The inclusion of precipitation as a constraint and as an additional term in the optimization objective function were shown. Finally, intelligent ingress in the face of unknown disturbances was also demonstrated.

Acknowledgments

The Authors would like to thank Jerry Straka and Erik Rasmussen for providing the simulated storm data and helping us interpret it. The Authors also thank Adam Houston and Brian Argrow for supporting the overall objectives of this effort.

References

[1]National Weather Service Office of Climate Weather and Water Services, "NWS weather fatality, injury, and damage statistics," http://www.nws.noaa.gov/om/hazstats.shtml, 2008.

[2]Argrow, B., Lawrence, D., and Rasmussen, E., "UAV Systems for Sensor Dispersal, Telemetry, and Visualization in Hazardous Environments," *43rd Aerospace Sciences Meeting and Exhibit*, AIAA, January 10-13 2005.

[3]"The VORTEX-2 Website," http://www.vortex2.org/, 2008.

[4]Hipskind, S., Tyrell, G., Holland, G., and Curry, J., "Use of the Aerosonde Unihabited Aerial Vehicle (UAV) in the Fourth Convection and Moisture Experiment (CAMEX 4)," *AIAA's 1st Technical Conference and Workshop on Unmanned Aerospace Vehicles*, Portsmouth, VA, May 2002.

[5]Lin, P.-H. and Lee, C.-S., "Fly into typhoon Haiyan with UAV Aerosonde," *12th Symposium on Meteorological Observations and Instrumentation*, 2003.

[6]Lin, P.-H. and Lee, C.-S., "The eyewall-penetration reconnaissance observation of typhoon Longwang (2005) with unmanned aerial vehicle, Aerosonde," *Journal of Atmospheric and Oceanic Technology*, Vol. 25, No. 1, 2008, pp. 15 – 25.

[7]LaValle, S. M., *Planning Algorithms*, Cambridge University Press, New York, NY, 2006.

[8]Mitchell, I. M., "The Flexible, Extensible and Efficient Toolbox of Level Set Methods," *Journal of Scientific Computing*, 2007.

[9]Philippsen, R. and Siegwart, R., "An interpolated dynamic navigation function," *Proceedings - IEEE International Conference on Robotics and Automation*, Barcelona, Spain, 2005, pp. 3782 – 3789.

[10]Petres, C., Pailhas, Y., Patron, P., Petillot, Y., Evans, J., and Lane, D., "Path planning for autonomous underwater vehicles," *IEEE Transactions on Robotics*, Vol. 23, No. 2, 2007, pp. 331 – 41.

[11]Soulignac, M., Taillibert, P., and Rueher, M., "Adapting the wavefront expansion in presence of strong currents," *2008 IEEE International Conference on Robotics and Automation*, Pasadena, CA, 2008, pp. 1352 – 8.

[12]Jarvis, R. A., "Collision-free trajectory planning using the distance transforms," *Mechanical Engineering Transactions of the Institution of Engineers*, Vol. 3, 1985, pp. 187–191.

[13]Dorst, L. and Trovato, K. I., "Optimal path planning by cost wave propagation in metric configuration space," *Proc. SPIE, Mobile Robots III*, edited by W. J. Wolfe, Vol. 1007, Jan. 1988, p. 186.

[14]Barraquand, J., Langlois, B., and Latombe, J. C., "Numerical Potential Field Techniques for Robot Path Planning," *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 22, No. 2, April 1992, pp. 224.241.

[15]Sethian, J. A. and Vladimirsky, A., "Ordered upwind methods for static Hamilton-Jacobi equations: Theory and algorithms," *SIAM Journal on Numerical Analysis*, Vol. 41, No. 1, 2003, pp. 325 – 363.

[16]Straka, J. M. and Anderson, J. R., "The numerical simulations of microburst producing thunderstorms: Some results from storms observed during the COHMEX experiment," *Atmospheric Science*, Vol. 50, 1993, pp. 1329–1348.

[17]Gilmore, M. S. and Straka, J. M., "Precipitation uncertainty due to variations in precipitation particle parameters within a simple microphysics scheme," *Mon. Wea. Rev.*, Vol. 132, 2004, pp. 2610–2627.

[18]Gilmore, M. S., Straka, J. M., and Rasmussen, E. N., "Precipitation and evolution in simulated deep convective clouds between liquid-only and simple ice and liquid phase microphysics," *Mon. Wea. Rev.*, Vol. 132, 2004, pp. 1897–1916.