

Concentration Gradient and Information Energy for Decentralized UAV Control¹

William J. Pisano²

Department of Aerospace Engineering Sciences
University of Colorado, Boulder, CO 80309-0429
pisano@colorado.edu

Dale A. Lawrence³

Department of Aerospace Engineering Sciences
University of Colorado, Boulder, CO 80309-0429
dale.lawrence@colorado.edu

Kamran Mohseni⁴

Department of Aerospace Engineering Sciences
University of Colorado, Boulder, CO 80309-0429
mohseni@colorado.edu

Abstract

Control techniques for managing a large Sensor Flock are described, consisting of decentralized potential gradient and sensed data gradient guidance laws. Guidance potentials are defined to improve sensed data quality in both a scalar signal-to-noise ratio sense, and in a vector gradient estimation sense. Examples are provided showing that a 147 vehicle flock autonomously produces desired behavior under suitable choices of design parameters. A portion of the flock clusters in regions of high data signal-to-noise ratio, while other vehicles spread out to cover un-sensed areas within a specified coverage volume. Clustered vehicles use measured data gradients to seek the margins and the source of a simulated toxic plume emission.

¹ Research supported by NSF ITR grant no. 042574744

² Graduate Student, Member AIAA

³ Associate Professor, Senior Member AIAA

⁴ Assistant Professor, Member AIAA

Introduction

Dangerous atmospheric contamination plumes can be caused by many sources such as tanker truck or railroad spills, industrial release accidents, or chemical/biological/nuclear terrorist attacks. Depending on prevailing winds and plume buoyancy, contaminants can be transported over large distances or remain in some regions for long periods. Hazards to the public and the environment are highly variable, since severity of exposure depends on the type of contamination, its concentration, duration of exposure, mode of contact, and natural toxin degradation rates. Accordingly, responses to a release event may range from simple broadcast advisories to remain indoors or avoid the area, to large-scale emergency evacuation and subsequent quarantine and decontamination.

In urban areas, large structures create complex wind patterns, increasing the difficulty of predicting plume dispersion. At the same time, high population density increases the stakes in making appropriate response decisions. Timely determination of toxin species, plume source, plume structure, and predictions of plume evolution would have significant benefits in preventing exposure and planning cost-effective remediation. While the development of sophisticated numerical dispersion models is on-going, the accuracy of these models depends on the quality and quantity of sensed toxin concentration and wind velocity data in a large volume of air above the protected area. Data at the lower boundary of this volume can be provided by sensors fixed to existing structures, providing relatively dense, high-quality information about instantaneous toxin concentrations where contact with personnel and property occurs. However, the predictive capability of such a sensor array is limited. Remote sensing of the atmosphere above a city using passive or active scanning optical sensors can provide additional information [16,13], but the resulting spatial resolution, specificity, and concentration accuracy cannot compete with in-situ sensing.

The prospects of high-density in-situ sensing using conventional means, such as radiosonde balloons, are daunting: the volume of interest over a mid-sized city is on the order of 20km in diameter and 2km in height. For a spatial resolution of $(300m)^3$, some 2,300 balloon-carried sensors are required. On the other hand, if the toxic plume occupies only 0.1% of this volume, only 23 sensors would be needed, provided they are somehow located in the vicinity of the plume. This suggests that a fleet of approximately 100 mobile airborne sensors, which can rapidly move to areas with significant toxin concentrations, could suffice. UAVs are a natural choice for the sensor platform, since they could be produced and operated at lower cost than a fleet of 100 conventional piloted aircraft, and toxic exposure to pilots could be avoided.

Large UAVs have high payload capability and can carry significant on-board computing, communication, and sensing resources. They are also high value assets, with a high cost of failure including danger to other aircraft and personnel on the ground. This requires sophisticated flight control and safety systems, and expanded FAA authority to operate over populated areas.

In contrast to work using large UAVs for toxin dispersion characterization, e.g. [6], we suggest that a *Sensor Flock* composed of bat-sized micro aerial vehicles (MAVs) would provide more appropriate technology [19]. (See Figure 1 for a MAV vehicle being developed at the University of Colorado). MAVs would pose little danger to personnel and property on the ground or other air vehicles. They do not need specialized take-off or landing facilities. They could be reusable, expendable, and could be produced in large numbers at low cost. While technology developments are needed, such small vehicles could potentially remain in flight for periods of about

two hours, sufficient to provide highly accurate data for decisions in the critical initial period after a toxin release event. Subsequently, fewer numbers might be used to monitor dispersions over longer periods.



Figure 1. Micro Air Vehicle under development at the University of Colorado. See the companion paper [20] for more detail about its design and control.

Such a Sensor Flock would provide benefits in other atmospheric sensing applications. Examples include: modeling the local weather produced by wildfires to better predict their evolution and improve the deployment of firefighting resources; sensing and modeling of thermodynamic plumes over open ice leads in polar regions to better understand interactions between sea, ice, and atmosphere which contribute to climate change; and in-situ data collection in storms to provide better storm track predictions and understanding of tornado genesis and evolution.

Deploying a 100 vehicle Sensor Flock raises unique challenges in command and control. Each MAV carries very limited on-board power and computing resources. Flight control, toxin sensing, information processing, communication, and decision making must be extremely simple and decentralized. Yet rather sophisticated aggregate behavior is desired, so that the flock can semi-autonomously seek out plumes, guided by supervisory human operators and real-time models of plume evolution.

Control of groups of UAVs has been pursued in various ways, from individually remote-piloted vehicles, such as the Predator, to slaving a sortie of UAVs to a lead piloted vehicle [5], to various formation flight control algorithms, e.g. [11,10,8], to sophisticated optimization-based path planning, task allocation, and coordination approaches, e.g. [18,9,3]. These approaches are ill-suited for the large numbers of vehicles needed in atmospheric sensing, and for the corresponding flocking behavior that allows vehicles to explore freely, guided by the quality of data each vehicle obtains.

In this paper, we discuss the concept of *Information Energy* for command and control [19], and examine the resulting behavior of a 147 vehicle Sensor Flock as it samples simple toxin plume dispersion in simulation. Vehicles communicate with each other to share sensed data and to plan their own motion, and communicate with a ground control center which maintains a model of the plume and predicts dispersion

over time. Each vehicle flies in an energy conserving loiter circle, whose centroid is self-adjusted over time to improve measures of data quality. Information Energy is a potential function maintained on each vehicle, containing terms for sensor data quality, communication quality of service, and the energy cost of motion. Vehicle motion is guided by gradients of the Information Energy potential, which can be computed by on-board flight control processors. Also investigated is the use of concentration field gradient information for vehicle motion. This provides a highly agile and data-reactive sensing network in the air above protected areas which can be monitored and controlled by a single supervising user.

Flock Control Structure

The Sensor Flock has a hierarchical control structure:

1. At the lowest level, simple high-rate control of propulsion and vehicle attitude is used to maintain steady, efficient flight. Instantaneous flight direction and speed are controlled to follow Lyapunov vector fields [7] which provide globally attractive trajectories which are compatible with MAV flight. See Figures 2 and 3.
2. As toxin concentration data is collected, the mid-level control alters loiter circle centers to cause circling vehicles to cluster in regions of high quality data (e.g. sufficient toxin concentration signal to noise ratio), while maintaining multi-hop communication relays to the central data processing facility, and avoiding excessive repositioning that shortens loiter time.
3. High-level control arranges loiter circles within clusters to improve data gradient estimates, and to better distribute clusters within the toxin plume.
4. At the highest level, humans supervise flock behavior based on detailed, centralized plume modeling from transmitted MAV data, providing intelligent decision making where automation cannot.

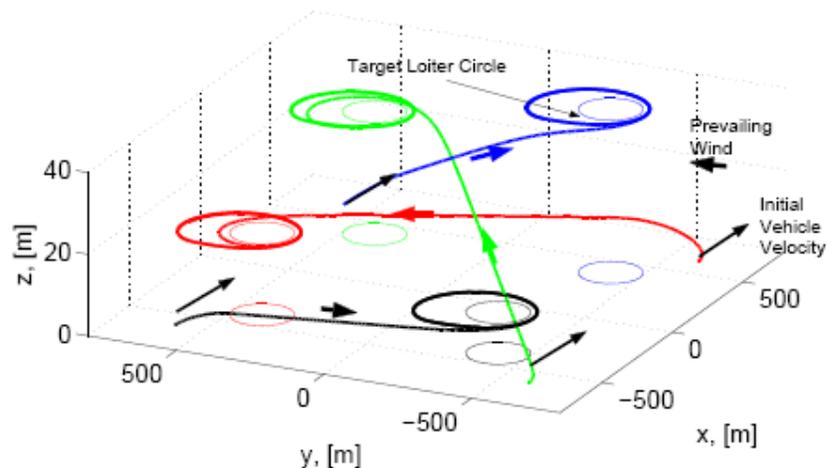


Figure 2. Trajectories of four UAVs resulting from a Lyapunov guidance law providing globally attractive loiter circles at different locations.

The low level loiter-capable autonomy mentioned in tier 1 of the control hierarchy has

been demonstrated on a small (20 oz) MAV at the University of Colorado. Flight data corresponding to an aircraft under complete tier one autonomy is shown in Figure 3.

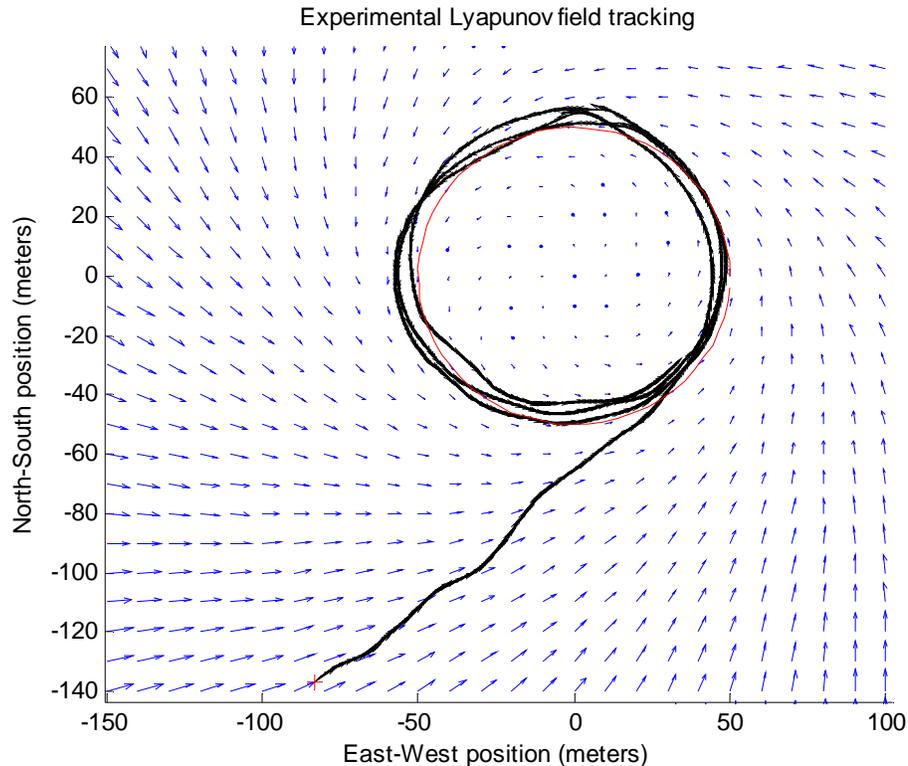


Figure 3. Top-down view of fully autonomous flight demonstrating global Lyapunov stability. Blue arrows correspond to the desired heading according to the Lyapunov vector field. The red + sign indicates the initial position of the aircraft. The red circle indicates the target loiter circle.

The only information required by tier one as demonstrated in this flight is the xyz coordinates of the desired loiter circle center, and the desired loitering radius. On-board computations in an 8-bit microcontroller implement flight control to follow the vector field. Further flight experimentation has shown success in using warping functions to create other globally convergent shapes such as ellipses and racetracks. The reliability of using these vector fields for guidance has been demonstrated by allowing the aircraft to loiter for up to half an hour without any human interaction. Movement of the desired center coordinates has also been successfully demonstrated in flight. Details of this level of control will be presented in a future paper.

Information Energy

Information Energy is used to provide mid-level guidance and control (tier 2) in this paper. Desired behavior includes:

- a) Avoiding collisions between MAVs.
- b) Encouraging some vehicles to cluster in the vicinity of a toxin plume.

- c) Dispersing some vehicles to search for other contaminated areas.
- d) Bounding vehicle motion to a limited coverage volume.
- e) Minimizing energy lost in moving vehicle loiter circle centers.

The Information Energy in each vehicle consists of both auto-potentials and hetero-potentials. The auto-potentials incorporate locally gathered information on each vehicles' own motion, its sensed data, and communication traffic it is relaying. The hetero-potentials include information on neighboring vehicles' relative positions and the quality of their sensed data. Each vehicle computes the local gradient of the composite potential field to determine the desired vector rate of change of it's loiter circle center.

Repelling Potentials

The design of Information Energy functions begins with desired behavior a) above. Consider the following radial hetero-potential function

$$V_{ri} = \frac{1}{2} \ln \frac{r_i^2 + d^2}{r_i^2 + \delta^2} \quad (1)$$

where $r_i = \sqrt{(x-x_i)^2 + (y-y_i)^2 + a_r^2(z-z_i)^2}$ is the (vertically modified) radial distance from the center of a given vehicle's loiter circle (x,y,z) to the i^{th} neighbor vehicle's loiter circle center (x_i,y_i,z_i) . The parameter a_r is used to scale vertical distance differently from the horizontal, in keeping with the relatively large diameter of the volume of interest compared to its vertical extent. The parameters $d > \delta$ control the radial influence of other vehicles: when another vehicle is more distant than d , the corresponding hetero-potential for that vehicle approaches 0. When this other vehicle is closer than d , the potential V_{ri} becomes positive. Taking the gradient of this potential, we obtain:

$$\nabla V_{ri} = \left[\frac{(x-x_i)}{r_i^2 + d^2} - \frac{(x-x_i)}{r_i^2 + \delta^2}, \frac{(y-y_i)}{r_i^2 + d^2} - \frac{(y-y_i)}{r_i^2 + \delta^2}, \frac{a_r(z-z_i)}{r_i^2 + d^2} - \frac{a_r(z-z_i)}{r_i^2 + \delta^2} \right]^T \quad (2)$$

whose magnitude is given by

$$|\nabla V_{ri}| = \frac{(d^2 - \delta^2)r_i}{(r_i^2 + d^2)(r_i^2 + \delta^2)} \quad (3)$$

If the vehicle in question moves in the opposite direction of this gradient, with a rate proportional to its magnitude, the effect is repulsion from the i^{th} vehicle with a strength that is large if the other vehicle is within a distance d but outside a distance δ . At distances large compared to d , the repulsion effect goes to zero. By superimposing the repelling hetero-potentials from all neighboring vehicles, i.e.

$$\nabla V_r = \sum_i \nabla V_{ri} \quad (4)$$

A local aggregate gradient ∇V_r is obtained for each vehicle, causing motion away from all other vehicles—more quickly if they are close and more slowly if they are

well-separated. When used without an additional bounding potential (see below), this causes the flock to disperse, also providing the basic tendency of the flock to produce the desired behavior c) above. Only the nearest neighboring vehicle's information is needed to produce desired results, thus reducing complexity [19].

To maintain the flock in a desired region, a bounding auto-potential is also defined:

$$V_b = \frac{1}{2} \ln(r_c^2 + b^2) \quad (5)$$

where $r_c^2 = (x - x_c)^2 + (y - y_c)^2 + a_c^2(z - z_c)^2$ is the (squared) distance of the vehicle loiter circle center (x, y, z) from the specified center (x_c, y_c, z_c) of the sample volume, and b is the radius of that volume. Note the vertical displacements have been scaled (by a_c) relative to the horizontal, as in the repelling potential earlier. The gradient of this potential is

$$\nabla V_b = \left[\frac{(x - x_c)}{r_c^2 + b^2}, \frac{(y - y_c)}{r_c^2 + b^2}, \frac{a_c(z - z_c)}{r_c^2 + b^2} \right]^T \quad (6)$$

whose magnitude goes to 0 as r_c becomes large, but does so much more slowly than the magnitude of the hetero-potentials V_{ri} . The negative of this gradient is added to the hetero-potential gradient above to produce a vehicle loiter circle rate of change. This produces a bounded flock diameter, satisfying behavior d), whose size can be controlled via the parameter b [19].

To ensure that vehicles remain away from possible contact with structures and personnel on the ground, a hard minimum altitude limit is introduced that acts as a simple barrier on motions produced by the above gradients enforced on each vehicle, so that vehicles (loiter circles) cannot descend below a specified floor altitude.

Plume Model

For the purposes of this paper, data quality will be defined by the scalar toxin concentration value produced by an air-sampling sensor on board each vehicle. We wish vehicles to cluster more closely together in the vicinity of the plume to provide high-quality measurements for estimating plume source and dispersion, and predicting areas on the ground that will be subsequently affected. However, we do not want vehicles to cluster too closely, since that would limit coverage of the plume extent. For the simulations in this paper, a simple steady source dispersive plume (no buoyancy) was modeled to produce toxin concentration data "measurements" from each vehicle, depending on its position in the plume. Concentration values D are determined by the model [14,4]

$$D = \log \left(\frac{m}{\pi(c(x - x_r) + \Delta)} e^{\left(\frac{-0.5[(y - y_r)^2 + (z - z_r)^2]}{(c(x - z_r) + \Delta)^2} \right)} + .001 \right) - \log(.001) \quad (7)$$

where m determines the mass flow rate of the effluent, c is the downstream dispersion coefficient, and Δ is a parameter that bounds the source concentration. For convenience in the simulation, this data concentration is normalized to a maximum value of 1 at the source. Figure 4 shows two partly transparent isosurfaces of data quality values 0.1 (red) and 0.01 (blue). The plume source is at ground level at horizontal coordinates $x_r = -2km$ and $y_r = -2km$, and the prevailing wind is along the positive x direction. The dashed red lines show the bounding box containing the

initial vehicle positions. Dotted blue lines intersect at the plume source, and the blue triangle indicates the ground control center.

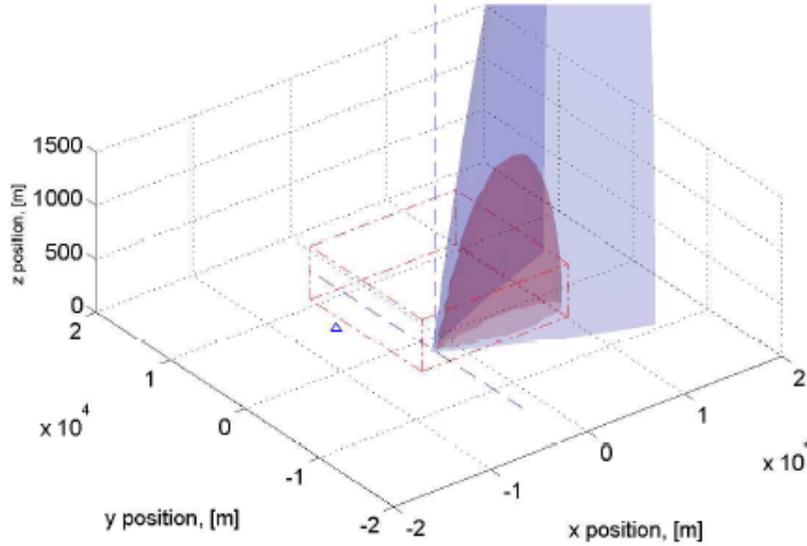


Figure 4. Simulated plume caused by steady toxin emission at ground level under a prevailing x-direction wind. Isosurfaces of data concentration quality values 1.0 (red) and 0.1 (blue) are shown.

Clustering Potentials

Another component is needed in the Information Energy potential to cause a subset of the flock to cluster in the plume where data quality values are high. One approach is to weaken the repelling fields for vehicles with high data quality, so that they are pushed together more closely by the un-weakened potentials of the vehicles outside the plume. The following hetero-potential has been investigated, which replaces the original expression for V_{ri} in (1) with

$$V_{ri} = \frac{1}{2} \ln \frac{r^2 + \left(\frac{d}{\sqrt{1+hD_i}} \right)^2}{r^2 + \left(\frac{\delta}{\sqrt{1+hD_i}} \right)^2} \quad (8)$$

where the parameter h determines the factor by which the original repelling potential is reduced radially. For example, when $h = 100$, the maximum saturated data value of 0.1 causes repelling radii to be reduced by a factor of 10 for vehicles in the high data quality regions of the plume.

Figure 5 shows the resulting motion of vehicle loiter circle centers from the initial grid shown as the red box, using the technique of data-weakened repelling. Since all the hetero-potentials now have gradients with the same sign, the balance between repulsion and clustering cannot be so easily tipped by the number of vehicles that aggregate, and difficulties with stability of the clustering are reduced. This makes the behavior less sensitive to small changes in parameters.

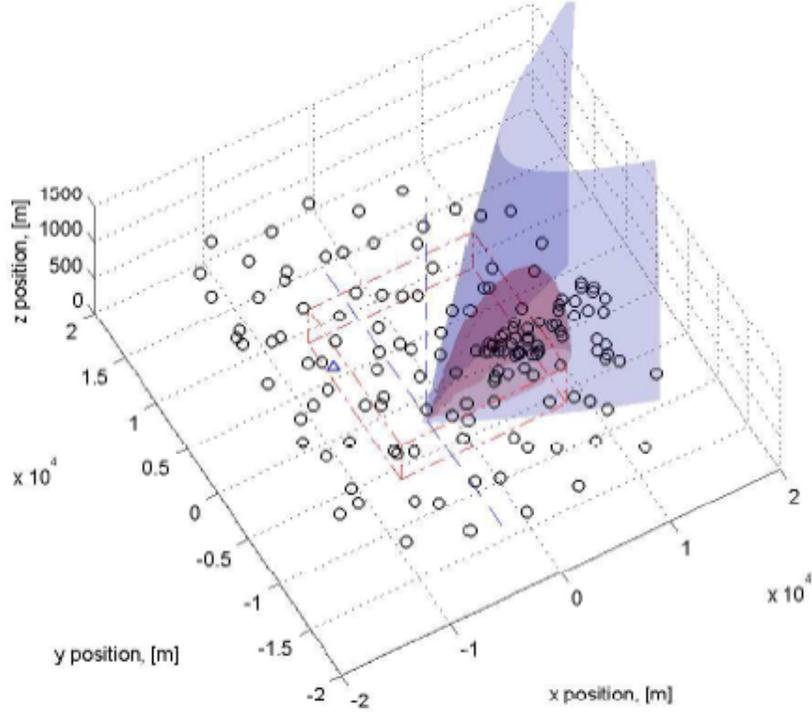


Figure 5. Clustering in the vicinity of the plume achieved by weakening of repelling potentials for vehicles with high data quality.

Motion Energy Conservation

Desired behavior e) is addressed by adjusting the effective viscosity of the medium in which the vehicles are moving. Larger viscosities are achieved by reducing the step size μ in the position updates, e.g. for the x-direction update of each vehicle loiter circle center

$$x_{k+1} = x_k + \frac{\mu \nabla V_x}{(1 + \mu |\nabla V_x|)} \quad (9)$$

This produces a rate of change in x which depends on the composite potential gradient “force” ∇V . The normalization $(1 + \mu |\nabla V_x|)$ limits these rates of change. This effective viscosity acts to suppress unnecessary activity in the vehicles, which helps to minimize excessive energy expenditure in moving loiter circle centers and thereby maximize time aloft.

Concentration Gradient Potentials

The above potential functions produce a rough clustering behavior in a plume, but additional control over vehicles within the plume may be needed to provide more accurate modeling of the plume source and extent. A more refined distribution can be obtained by using locally determined values for the plume concentration gradient. A 3-dimensional gradient for a given volume can be found numerically by finding the difference in concentration at several points within the volume. For a vehicle flying in a straight line, gradient determination is very inaccurate normal to the line, because no

spatial variation can be sensed in the resulting 2-dimensional null space.

For vehicles within the plume, the center of a loiter circle is the information that describes the position of the vehicle. If it assumed that the rate of change of the plume concentration is slow (on the order of several minutes) relative to the loiter period of the aircraft (on the order of 30 seconds for a 100 m diameter circle and a 10m/sec flight speed), a local gradient can be determined on an individual plane contained on the loiter circle. This still leaves a data null space normal to the circle (in the z-direction for a horizontal loiter).

A possible solution to this is a tilted, precessing loiter circle that climbs on one side and dives on the other. Two or more orbits, precessed relative to each other, eliminate the data null space, providing improved estimates of the 3-dimensional concentration gradient. In addition to the assumptions above, it is also necessary that the projected center of the loiter circle be moving slowly enough so the vehicle is circling as opposed to flying straight to “distant” transfer orbits (see Figure 3).

In order to simulate the above described determination of concentration gradients, several motion guidelines were put into place. In order to assure that the vehicle is circling in adequate data, it is assumed that vehicles above a certain concentration threshold, i.e. in the plume, have gradient information available. Vehicles outside the plume are assumed to not have adequate information, or have their loiter centers moving too fast to find gradient information; thus they have none available in the simulation. The aircraft determined to have plume gradient information available were given the analytic gradient

$$\nabla \mathbf{D} = \frac{\partial \mathbf{D}}{\partial x} \hat{x} + \frac{\partial \mathbf{D}}{\partial y} \hat{y} + \frac{\partial \mathbf{D}}{\partial z} \hat{z} \quad (10)$$

of the plume function \mathbf{D} , Eqn. (7), effectively assuming perfect gradient estimation for the purposes of this simulation.

A major concern is the positioning of vehicles within the plume. Intuitively it makes sense to have more vehicles inside the plume rather than outside, but how should the vehicles be distributed throughout the plume to best facilitate mapping and modeling of the data? The first concern is finding the source of the agent that is entering the atmosphere. With concentration data only this can be difficult because it is hard to know when a maxima in the distribution field is reached. It is also hard to determine the direction a vehicle should move to increase its sensed concentration. The addition of gradient field information at the vehicle level allows for direct gradient ascent, e.g. the vehicle flies in the direction specified by the gradient vector. Analogous to Eqn. (9), x-components evolve according to

$$x_{k+1} = x_k + \mu \nabla \mathbf{D}_x \quad (11)$$

Given this control architecture, the aircraft (loiter circle center) will always move in the direction of increasing concentration. In simple plume dispersion cases such as the one being studied here, this will always converge to the point of greatest concentration within the volume (assuming the step size μ is small enough), which is the best estimate of the chemical source given the information available. When combined with the repulsion field equations outlined above, the motion for the x-component is given by:

$$x_{k+1} = x_k + \frac{\mu(\nabla\mathbf{V}_{rx} + \nabla\mathbf{V}_{bx} + \nabla\mathbf{D}_x)}{1 + \mu|\nabla\mathbf{V}_{rx} + \nabla\mathbf{V}_{bx} + \nabla\mathbf{D}_x|} \quad (12)$$

Where $\nabla\mathbf{V}_r$ is from equation 8 above and $\nabla\mathbf{V}_b$ is from equation 6 above. The addition of this “source finding” algorithm results in a high concentration of vehicles at the source with a progressively lower concentration throughout the rest of the plume. This results in a large number of data points near the source of the plume, and a small number farther away.

One difficulty in having so many vehicles so tightly concentrated at the source is that the quality of the map or model generated by the data will be poor farther away from the source. If the goal of the data collection is to identify the source, this is adequate. If downstream propagation predictions such as the direction, speed, and concentration with which the plume will move are desired, the source-centric model will be far from ideal. An alternative approach is to use a team-cooperative strategy that designates a small number of vehicles as “source finders”, while the rest of the vehicles act as “margin finders” for the plume. The distribution of the second set of vehicles in the current simulation is based on finding a level curve of the plume concentration. The vehicles use the gradient information (if any) that is available to them to locate a 3-dimensional level curve, or isosurface, of the plume. The vehicles also obey the bounding ellipsoid and vehicle repulsion rules outlined above. The end result is a relatively even dispersion of many vehicles on the isosurface of interest. This set of data points gives a 3-dimensional map of that particular isosurface which, for the simple plume models currently under investigation, can be used to further extrapolate the remaining volumetric distribution of the plume.

The “margin-finding” behavior is produced as follows. A plume concentration of interest is specified a priori to which the vehicles will be attracted. The logic on board the vehicles specifies that if the vehicle is in a concentration lower than the specified value, it will ascend the gradients. Conversely, if the concentration is higher than the specified value, it will descend the gradients. Once a vehicle finds itself within a few percent of the desired isosurface concentration, it will modify its flight algorithm once again to 1) stay on the isosurface and 2) distribute itself evenly with respect to the other vehicles present on the surface. To accomplish condition one, the net movement from potential field summation in the direction normal to the isosurface is set to zero. This is done using the loiter circle force

$$F_t = (\mathbf{I} - \nabla\hat{D} * \nabla\hat{D}^T) * \frac{(\nabla\mathbf{V}_r + \nabla\mathbf{V}_b)}{1 + \mu|\nabla\mathbf{V}_r + \nabla\mathbf{V}_b|} \quad (13)$$

where F is the potential field summation for the given vehicle, $\nabla\hat{D}$ is the normalized gradient vector at the current vehicle position, \mathbf{I} is the 3x3 identity matrix, and F_t is the potential field summation vector in the tangential direction. Condition two is satisfied by ignoring all potential field effects on the vehicle except for the repulsion of other vehicles once in the concentration margin.

The result of all the above vehicle control terms is five distinct vehicle behaviors:

1. Locate plume source (Green circles in simulation)
2. Find isosurface from high concentration (Red triangles in simulation)
3. Remain on isosurface (Blue circles in simulation)
4. Find isosurface from low concentration (Red triangles in simulation)
5. Maintain coverage over non-plume area (Black stars in simulation)

The task of each vehicle is determined by its own sensed scalar concentration, and is not determined prior to deployment. In fact, all aircraft are deployed in the same mode and determine their appropriate task in flight, based on measured data concentration values. The corresponding motion update equations for each vehicle task are as follows:

$$\begin{aligned}
 1. \quad x_{k+1} &= x_k + \frac{\mu(\nabla \mathbf{V}_{rx} + \nabla \mathbf{V}_{bx} + \nabla \mathbf{D}_x)}{1 + \mu|\nabla \mathbf{V}_{rx} + \nabla \mathbf{V}_{bx} + \nabla \mathbf{D}_x|} && (\nabla \mathbf{V}_r \text{ given by equation 8}) \\
 2. \quad x_{k+1} &= x_k + \frac{\mu(\nabla \mathbf{V}_{rx} + \nabla \mathbf{V}_{bx} - \nabla \mathbf{D}_x)}{1 + \mu|\nabla \mathbf{V}_{rx} + \nabla \mathbf{V}_{bx} - \nabla \mathbf{D}_x|} && (\nabla \mathbf{V}_r \text{ given by equation 8}) \\
 3. \quad x_{k+1} &= x_k + \mu(\mathbf{I} - \nabla \hat{\mathbf{D}} * \nabla \hat{\mathbf{D}}^T) * \frac{(\nabla \mathbf{V}_{rx})}{1 + \mu|\nabla \mathbf{V}_{rx}|} && (\nabla \mathbf{V}_r \text{ given by equation 4}) \\
 4. \quad x_{k+1} &= x_k + \frac{\mu(\nabla \mathbf{V}_{rx} + \nabla \mathbf{V}_{bx} + \nabla \mathbf{D}_x)}{1 + \mu|\nabla \mathbf{V}_{rx} + \nabla \mathbf{V}_{bx} + \nabla \mathbf{D}_x|} && (\nabla \mathbf{V}_r \text{ given by equation 8}) \\
 5. \quad x_{k+1} &= x_k + \frac{\mu(\nabla \mathbf{V}_{rx} + \nabla \mathbf{V}_{bx})}{1 + \mu|\nabla \mathbf{V}_{rx} + \nabla \mathbf{V}_{bx}|} && (\nabla \mathbf{V}_r \text{ given by equation 8})
 \end{aligned}$$

The task 3 algorithm uses $\nabla \mathbf{V}_r$ given by equation 4 because it includes no concentration clustering information, thus its only effect is a repulsion field between the vehicles. Figure 6 shows the resulting simulated vehicle distribution with these modifications in place.

The vehicles in Figure 6 are color coded according to their current assignments. The green vehicles are locating the source of the plume, blue are within a few percent of the isosurface of interest (shown as the red shell). Red vehicles are inside the plume and searching for the isosurface, and the black vehicles are spread outside of the plume in case a new source of concentration emerges.

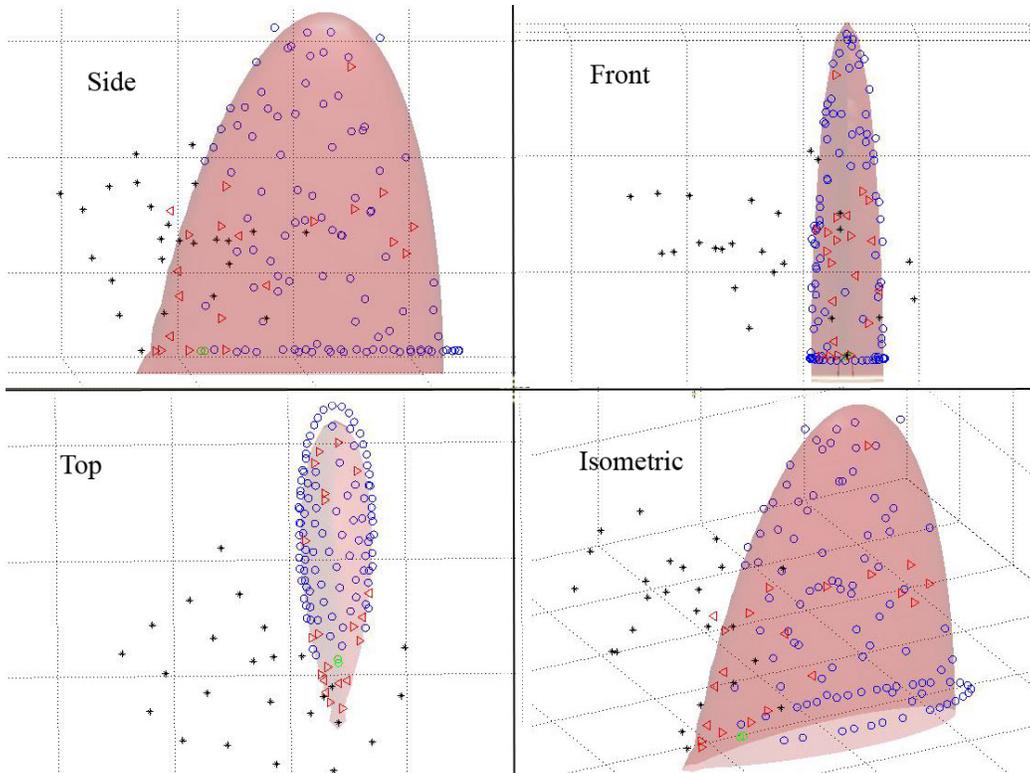


Figure 6. Vehicle distribution with gradient field modifications. Green circles indicate vehicles finding the chemical source. Red triangles indicate vehicles attempting to find the isosurface margin of interest. Blue circles indicate vehicles within the isosurface margin of interest. Black stars indicate vehicles searching for unknown areas of chemical concentration.

Conclusion

This paper discussed the concept of information energy and introduced the concept of concentration gradient determination for the purpose of locating and mapping a toxic plume dispersed over a dense urban area with a sensor flock. Using simple potential fields and the concentration gradient information, it was demonstrated that clustering can be autonomously achieved around the toxic plume in a manner advantageous to finding the contamination source and estimating plume margins.

At each level, gradient descent of energy functions provides a local optimization approach, easily computed by each vehicle. Lyapunov vector fields provide a rigorous way to supply compatible dynamics with global stability properties when the desired equilibrium set is known (e.g. loiter circles). Dynamics at the mid level require the construction of appropriate energy functions, which is not as simple in this application as in formation control [11,10,8] or robot path planning [17,15], where stability to known equilibria are desired. Also, this approach seeks to include not only data-dependent motions as in [1,6], but also network traffic and internal energy factors. We also seek to provide analytic Lyapunov stability support for control law behavior, unlike general intelligent agents [2,12] which are more flexible, but more difficult to characterize. The information energy and data gradient approach provides extremely simple, decentralized solutions for computation and communication feasibility in a flock of hundreds of low-capability vehicles.

References

- [1] R. Bachmayer and N. E. Leonard, "Vehicle Networks for gradient Descent in a Sampled Environment", *Proc. IEEE Conf. on Decision and Control*, Las Vegas, NV, Dec., 2002, pp. 112–117.
- [2] P. R. Chandler and M. Pachter, "Hierarchical Control for Autonomous Teams", *Proc. AIAA Guidance, Navigation, and Control Conf.* Montreal, Canada, Aug., 2001. AIAA paper no. 2001-4149.
- [3] P. R. Chandler, M. Pachter, D. Swaroop, F. M. Fowler, J. K. Howlett, S. Rasmussen, C. Schumacher, and K. Nygard, "Complexity in UAV Cooperative Control", *Proc. American Control Conference*, Anchorage, AK, May, 2002.
- [4] F.A. Gifford, "Turbulent diffusion-typing schemes: A review", *Nuclear Safety*, Vol. 17, No. 1, pp. 68–86, 1976.
- [5] C. E. Hanson, M. J. Allen, J. Ryan, and S. R. Jacobson, "Flight Test Results for an Autonomous Formation Flight Control System", *Proc. AIAA's 1st Tech. Conference & Workshop on UAV, Systems, Technologies, & Operations*, Portsmouth, VA, May, 2002. AIAA paper No. 2002-3431.
- [6] M. A. Kovacina, D. Palmer, G. Yang, and R. Vaidyanathan, "Multi-Agent Control Algorithms for Chemical Cloud Detection and Mapping Using Unmanned Air Vehicles", *Proc. IEEE/RSJ Conf. Intelligent Robots and Systems*, Lausanne, Switzerland, Oct., 2002, pp. 2782–2788.
- [7] D. A. Lawrence, "Lyapunov Vector Fields for UAV Flock Coordination", *Proc. 2nd AIAA "Unmanned Unlimited" Systems, Technologies, and Operations Conf.* San Diego, CA, Sept., 2003, AIAA paper no. 2003-6575.
- [8] N. E. Leonard and E. Fiorelli, "Virtual Leaders, Artificial Potentials and Coordinated Control of Groups", *Proc. IEEE Conf. on Decision and Control*, Orlando, FL, Dec., 2001, pp. 2968–2973.
- [9] W. M. McEneaney and B. Fitzpatrick, "Control for UAV Operations Under Imperfect Information", *Proc. AIAA's 1st Tech. Conference & Workshop on UAV, Systems, Technologies, & Operations*, Portsmouth, VA, May, 2002. AIAA paper No. 2002-3418.
- [10] E. A. Mehiel and M. J. Balas, "A Rule-Based Algorithm that Produces Exponentially Stable Formations of Autonomous Agents", *Proc. AIAA Guidance, Navigation, and Control Conf.* Monterey, CA, Aug., 2002. AIAA paper no. 2002-4591.
- [11] R. Olfati-Saber and R. M. Murray "Distributed Structural Stabilization and Tracking for Formations of Multiple Dynamic Agents", *Proc. IEEE Conference on Decision and Control*, Las Vegas, NV, Dec. 2002, pp. 209–215.
- [12] Y. Yang, M. M. Polycarpou, and A. A. Minai, "Opportunistically Cooperative Neural Learning in Mobile Agents", *Proc. Int. J. Conf. Neural Networks*, May, 2002, pp. 2638–2643.
- [13] C.J. Senff, R.M. Hardesty, R.J. Alvarez II, S.D. Mayor, "Airborne lidar characterization of power plant plumes during the 1995 Southern Oxidants Study", CIRES/University of Colorado/NOAA, Environmental Technology Laboratory Report R/E/ET2, Boulder, CO, 1998.
- [14] D. H. Slade, "Meteorology and atomic energy", U.S. Atomic Energy Commission report TID-24190, Clearinghouse for Federal Scientific and Technical Information, National Bureau of Standards, 1968.
- [15] P. Song and V. Kumar, "A Potential Field Based Approach to Multi-Robot Manipulation", *Proc. IEEE Conf. on Robotics and Automation*, Washington, DC, May, 2002, pp. 1217–1222.
- [16] R.L. Wayson, G.G. Fleming, B. Kim, Wynn L. Eberhard, and W.A. Brewer, "The use of LIDAR to characterize aircraft initial plume characteristics", FAA report FAA-AEE-02-04, 2002.
- [17] H. Yamaguchi and T. Arai, "Distributed and Autonomous Control Method for Generating Shape of Multiple Mobile Robot Group", *Proc. Int. Conf. Intelligent Robots and Systems*, Sept., 1994, pp. 800-807.
- [18] I. H. Yang and Y. J. Zhao "Real-Time Trajectory Planning for Autonomous Aerospace Vehicles Amidst Static Obstacles", *Proc. AIAA's 1st Tech. Conference & Workshop on UAV, Systems, Technologies, & Operations*, Portsmouth, VA, May, 2002. AIAA paper No. 2002-3421
- [19] DA Lawrence, K Mohseni, R Han "Information energy for sensor-reactive UAV flock control", *3rd AIAA Unmanned Unlimited Technical Conference, Workshop and Exhibit*, 20-23 September 2004, AIAA paper 2004-6530
- [20] J. Hall, K. Mohseni, and D. Lawrence, "Lateral Control of a Tailless Micro Aerial Vehicle", *Proc. AIAA Guidance, Nav., and Cont. Conf.* Keystone, CO, Aug., 2006. AIAA-2006-6689