AIAA 3rd "Unmanned Unlimited" Technical Conference, Workshop and Exhibit
20 - 23 September 2004, Chicago, Illinois

AIAA 2004-6530

Information Energy for Sensor-Reactive UAV Flock Control

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

Ryan E. Donahue
Department of Electrical and Computer Engineering
University of Colorado, Boulder, CO 80309-0429
ryan.donahue@colorado.edu

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

Richard Han
Department of Computer Science,
University of Colorado, Boulder, CO 80309-0430,
rhan@cs.colorado.edu

A Sensor Flock is described, consisting of many small UAVs which move according to decentralized potential gradient guidance laws designed to cause improvement in sensed data quality. Information Energy potentials are defined to include terms for sensed data quality, mechanical energy loss due to motion, and quality-of-service for inter-vehicle communication. Examples are provided showing that a 147 vehicle flock autonomously produces desired behavior under suitable choices of parameters in the Information Energy function. A portion of the flock clusters in regions of high data signal-to-noise ratio, while other vehicles maintain communication links to a base station, and still other vehicles spread out to cover unsensed areas within a specified coverage volume. Applications in toxic plume characterization and atmospheric science are discussed.

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,[13,16] 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 23,260 balloon-carried sensors are required. On the other hand, if the toxic plume occupies only 0.01% 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. (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, attritable, 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 several 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.

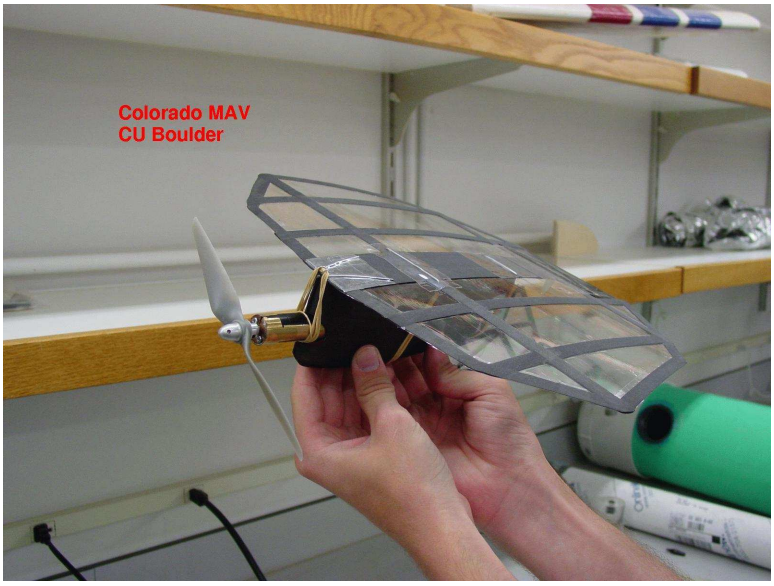

Figure 1. Membrane wing MAV under development at the University of Colorado. Wingspan is 30 cm and mass is 170 g.

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 prediction 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.,[8,10,11] to sophisticated optimization-based path planning, task allocation, and coordination approaches, e.g.[3,9,18] 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 introduce the concept of *Information Energy* for command and control, and examine the resulting behavior of a 147 vehicle Sensor Flock as it samples a 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 are simple to compute in on-board flight control processors. 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 Figure 2.

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.

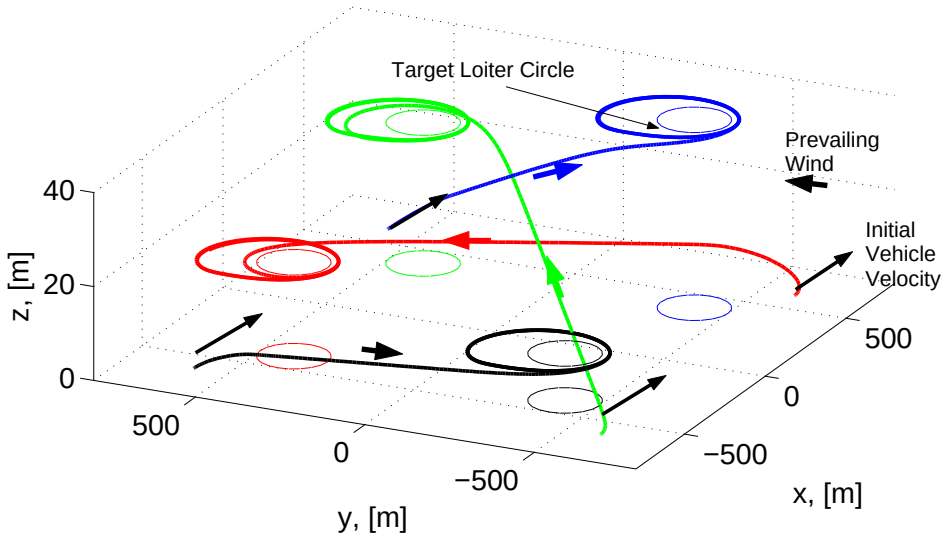

Figure 2. Trajectories of four UAVs resulting from a Lyapunov guidance law providing globally attractive loiter circles at different locations.

Information Energy

Information Energy is used to provide mid-level guidance and control (level 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) maintaining network connectivity between vehicles in the plume and the ground control center

f) 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 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 vector rate of change of its loiter circle center.

Repelling Potentials

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

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

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

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

whose magnitude is given by

$$|\nabla V_{1i}| = \frac{(d^2 - \delta^2) r_i}{(r_i^2 + d^2)(r_i^2 + \delta^2)} \tag{3}$$

If the vehicle in question moves in the opposite direction of this gradient, with a rate proportional to its magnitude, the effect is a 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 all these repelling hetero-potentials, i.e.

$$\nabla V_r = \sum_i \nabla V_{ri} \tag{4}$$

an 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.

Although each of the i^{th} terms in the gradient computation for ∇V_{ri} is simple, there are a large number of terms if the hetero-potential to every other vehicle in the flock is computed. In practice, only vehicles in a neighborhood need be included, since the contribution to the composite gradient from distant vehicles is small. This could be implemented, for example, by only including those other vehicles for which radio signals (communicating their position) have sufficient strength. Distant vehicles are then automatically excluded, saving computation. In addition, since changes in relative position are slow compared to loop rates of this mid-level guidance law, hetero-potentials can be updated with one other-vehicle gradient term per loop, distributing the computation on each vehicle over time.

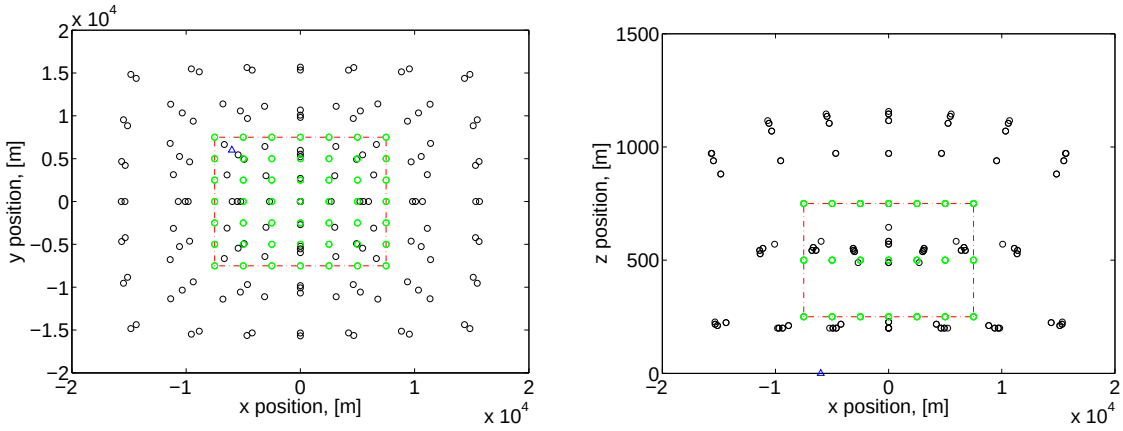

Figure 3. Simulation of flock behaviors a), c) and d), where a regular grid (greed circles) of 147 MAVs autonomously disperses (black circles) to fill an ellipsoidal bounding volume with 100m altitude lower bound.

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

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

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 protected 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

$$\nabla V_b = \left[\begin{array}{ccc} \frac{(x - x_c)}{(r_c^2 + b^2)} & \frac{(y - y_c)}{(r_c^2 + b^2)} & \frac{a_c(z - z_c)}{(r_c^2 + b^2)} \end{array} \right]^T \tag{6}$$

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. Thus, vehicles that are repelled to distances approaching the radius b experience increasingly strong attractive "forces" toward the flock center, which eventually overcome the repelling effects from the V_{ri} potentials. This produces a bounded flock diameter, satisfying behavior d), whose size can be controlled via the parameter b. Vehicle loiter circle centers then become approximately equally distributed (due to the repelling potentials) throughout an ellipsoidal region centered at (x_c, y_c, z_c).

To ensure that vehicles remain away from possible contact with structures and personnel on the ground, a simple barrier on motions produced by the above gradients is enforced on each vehicle, so that vehicles (loiter circles) cannot descend below a specified floor altitude. Horizontal motion at this boundary remains unconstrained. Figure 3 shows results of a simulation using the repelling and bounding potential gradients, together with the floor barrier, where vehicles (loiter circle centers) are initialized at a rectangular 7x7 array of grid points spaced 2.5 km horizontally, and in three layers spaced by 250 m vertically. The vehicles disperse to equi-distant stations within a flat-bottomed ellipsoid approximately 20 km in diameter and 1 km thick, bounded below by a 100 m floor. This provides autonomous low-resolution coverage of a large region, but does not incorporate desirable behaviors b) and e) supporting high-resolution plume sampling and high-rate networking of this information to the ground control station.

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. As long as sensed data is above a specified signal-to-noise threshold, data quality will be defined to be "high". I the maximum concentration is normalized to 1, we will presume that high data quality is obtained with normalized concentrations larger than 0.1. On the margins and outside the plume, sensed signal quality will decrease with concentration to 0. Clustering tendencies of vehicles should be equivalent anywhere in a plume where the data quality value is larger than 0.1, and this tendency should be reduced according to the quality value outside the plume. 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[4, 14]

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

where m determines the mass flow rate of the effluent, and 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, and then saturated so that locations in the plume with concentration larger than 0.1 have D set to 0.1. 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.

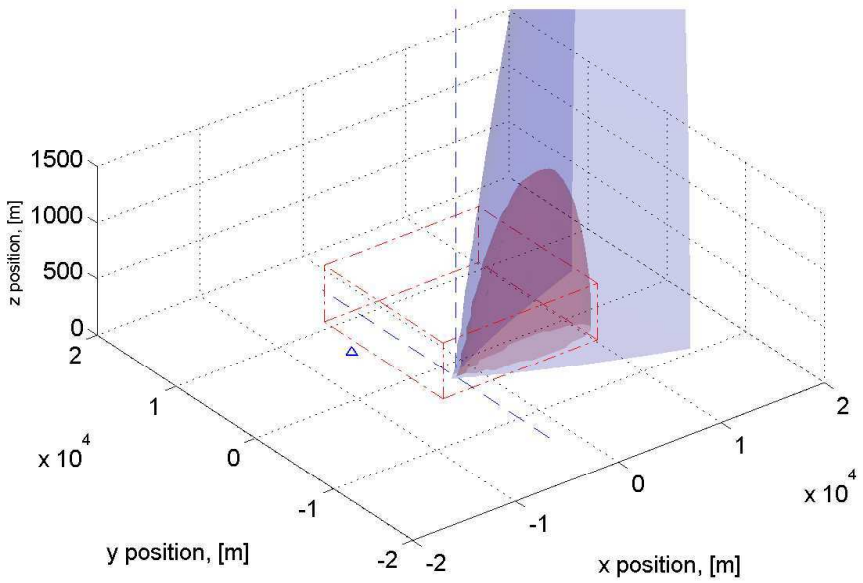

Figure 4. Simulated plume caused by steady toxin emission at ground level under a prevailing x-direction wind. Isosurfaces of data concentration quality values 1.0 (red) and 0.1 (blue) are shown.

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 counter the repelling potentials with attractive potentials for vehicles which are already in the plume, causing other vehicles to be attracted to their vicinity. For example, the clustering hetero-potential

$$V_{ci} = \ln \frac{r_i^2 + d_1^2}{r_i^2 + D_i^2 + d_1^2} \tag{8}$$

has been investigated, where d_1 is used to set the data quality threshold, and D_i is the data concentration value communicated from other vehicles. This produces a potential well around other vehicles which have high-quality sensor data. Together with the large repelling gradient close to other vehicles, this produces a kind of 3-dimensional moat around each vehicle in the plume, attracting other vehicles from distances on the order of d_1, but retaining separation distances on the order of δ from the repelling potential (1). Unfortunately, as vehicles begin to cluster, the aggregation of all these potential wells produces a deep "super-well", which attracts more vehicles, etc. While it is possible to choose parameters to keep all vehicles in the flock from clustering to the plume, retaining some for low-resolution coverage of the rest of the protected volume, parameter choice is very sensitive. The difficulty is that the repelling and attracting potentials must balance, and this balance is affected by the

number of vehicles that aggregate. As a result, the required parameter values can be difficult to determine. Figure 5. shows a plot of a 2D potential of this type, including repelling and attracting terms. Note the deep well where attracted vehicles have gathered. As additional vehicles fall into this well, the lip of the basin of attraction expands, attracting more vehicles, out to a radius where the attraction is balanced by repelling potentials.

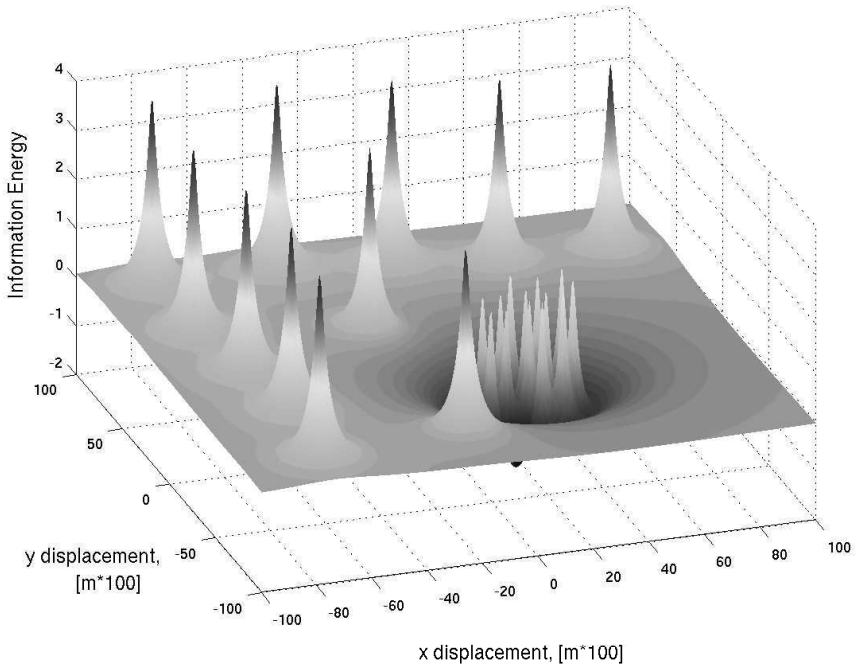

Figure 5. Illustration of a an aggregate potential seen by one vehicle on a 2D domain containing both repelling and attractive terms, causing clustering which is very sensitive to potential function parameters.

An alternate approach is to weaken the repelling fields for vehicles with high data quality, so that they are pushed together more closely by the unweakened 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

$$V_{ri} = \frac{1}{2} \ln \frac{r^2 + (\frac{d}{1+hD_i})^2}{r^2 + (\frac{\delta}{1+hD_i})^2} \tag{9}$$

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. Note that the separate clustering hetero-potentials V_{ci} are not needed in this approach.

Figure 6 shows the resulting motion of vehicle loiter circle centers from the initial grid of Figure 3, 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 avoided. This makes the behavior less sensitive to small changes in parameters.

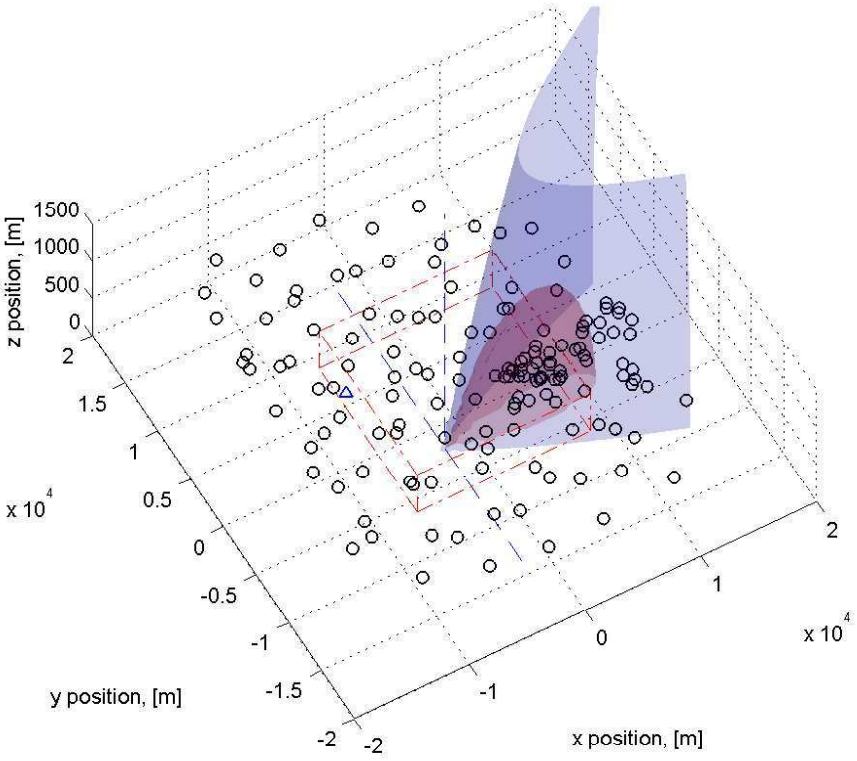

Figure 6. Clustering in the vicinity of the plume achieved by weakening of repelling potentials for vehicles with high data quality.

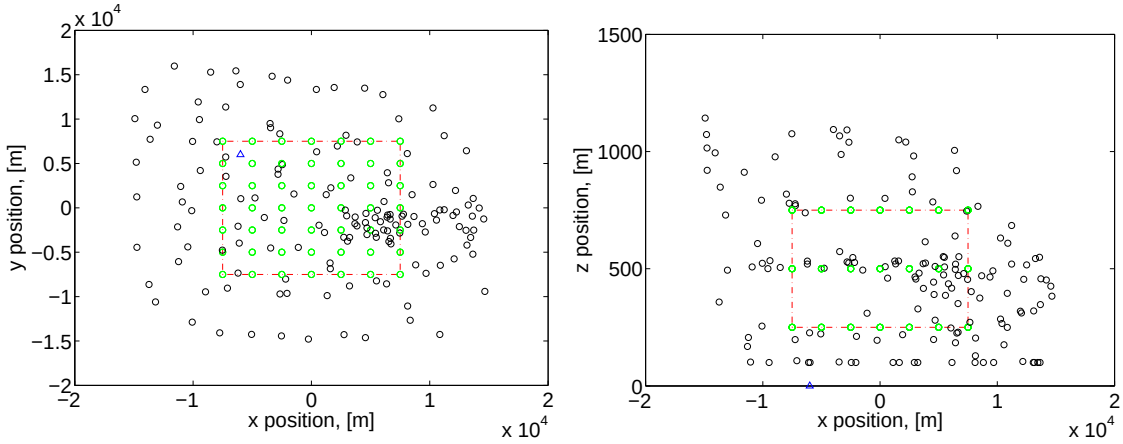

Figure 7. MAV motions using information energy with data clustering effects, shown projected onto the xy plane (left) and the xz plane (right)

Networking Potentials

In order to emulate the data traffic from an ad-hoc communication network among MAVs, a weighted average of unit vectors pointing from the base station to each vehicle was computed. These vectors were weighted based upon the the data quality at the location of that vehicle. As the vehicles cluster in the plume, we obtain an average vector which points from the base station in the direction of the approximate centroid of the good-data cluster. The ad-hoc networking is assumed to primarily route sensor data traffic through vehicles which are positioned in the vicinity of this line connecting the base station to the cluster. Desired behavior e) requires that a sufficient number of vehicles remain in this communication corridor in order to provide adequate channel capacity by relaying sensor data along multi-hop links between vehicles in the flock. For the present purposes, we presume that the data relay traffic C that each vehicle is carrying is related to its distance Dc away from the axis of the communication corridor via

$$C = g\frac{dr^2}{dr^2 + Dc^2} \tag{10}$$

which has a maximum of g along the axis of the corridor, and falls off to zero for vehicles away from the corridor. This communication traffic measure C is used in a similar manner as the data quality D in (9) to weaken the repelling potentials, causing vehicles to cluster along the communication corridor. Figure 8 shows the results of the combined data and communication clustering effects. This has the desired dispersion of some vehicles and clustering of other vehicles in the plume and along the communication corridor.

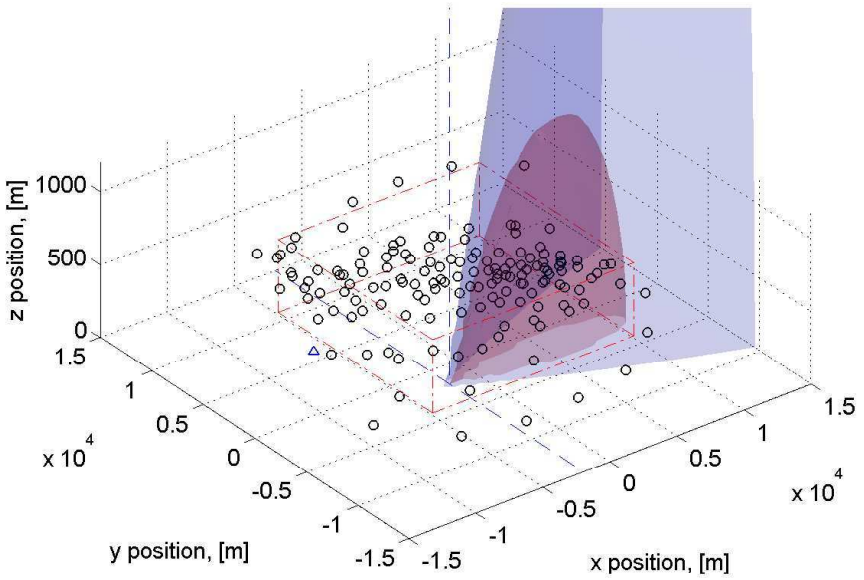

Figure 8. MAV motions using information energy with both data and communication clustering effects.

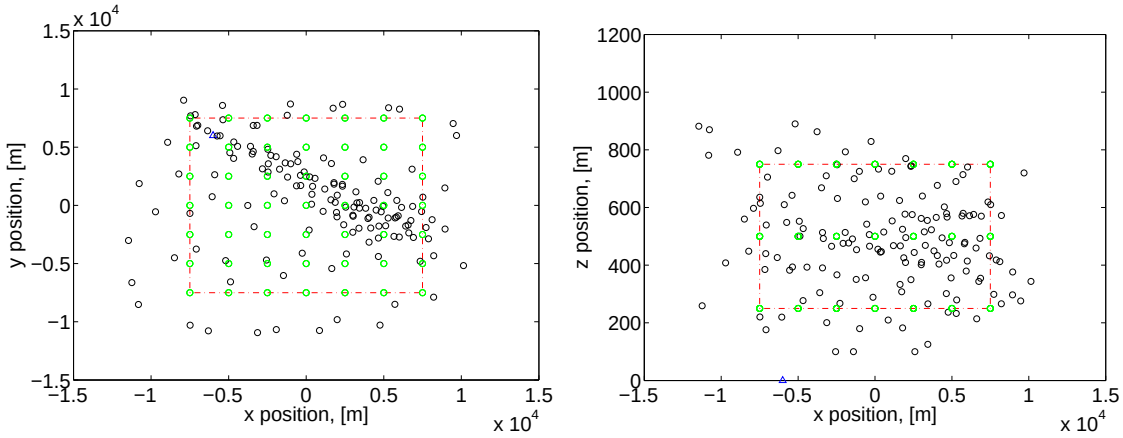

Figure 9. MAV motions using information energy with both data and communication clustering effects, shown projected onto the xy plane (left) and the xz plane (right)

Motion Energy Conservation

Desired behavior f) 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

$$x_{k+1} = x_k + \mu \nabla V_x / (1 + \mu |\nabla V_x|) \tag{11}$$

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.

Conclusions and Future Work

This paper presented the concept of information energy for the purpose of locating and mapping a toxic plume dispersed over a dense urban area with a sensor flock. Using simple potential fields, it was demonstrated that clustering can be autonomously achieved around the toxic plume, while also maintaining a multi-hop path for communications to the base station.

While the information energy approach shows considerable promise, some important questions remain:

- How should desired behavior be precisely defined? At the lower levels, behavior definition is easier. e.g. global loiter circle stability. At the supervisory level, behavior will be situation dependent and user controlled, hence adequate command and control must be provided (see below). For mid- and high-level autonomy, desired behavior is not well-defined. Partly, this is because it will be data-dependent, and partly because precise motions are not required. However, a precise *notion* of desired behavior in specific dispersion scenarios will be needed to develop meaningful theory and to more fully evaluate or compare candidate autonomous algorithms. This can be addressed

by constructing representative test case scenarios and corresponding behavior metrics, ranging from analytic dispersions for which rigorous results may be obtained, to representative dispersions for realistic performance assessment.

- How can desired behavior be produced with assured performance? Complexity is reduced by the proposed control hierarchy, which aggregates similar behavior in each level, and limits interaction between levels by taking advantage of several features of the application: time scale separation (fast attitude control vs. slow loiter circle re-positioning); domain splitting (clustering only produces radial constraints, perhaps allowing independent arrangement for better gradient estimates); length scale separation (local support of potential functions limits the number of neighbors contributing to potential gradients); and parametric interaction (inter-level interaction is via variation of potential function parameters, instead of explicit motion commands).

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[8,10,11] or robot path planning,[15,17] were 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. Finally, we must provide 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.* Monterrey, 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