DISTRIBUTED CHEMICAL PLUME PROCESS DETECTION: MILCOM 2005 #1644

Glenn Nofsinger and George Cybenko
Thayer School of Engineering
Dartmouth College
Hanover, NH 03755
gtn@dartmouth.edu

***Abstract*— We have developed a new approach for detecting and tracking chemical or biological plumes in distributed sensor networks, with the objective of solving the inverse location problem. The canonical plume tracking problem suffers from the challenge of a large state space, and we seek reduced dimensionality using information theoretic and stochastic methods. Working with an airborne plume we model a plume using multiple hypothesis tracking (MHT) techniques as opposed to transport based methods rooted in solutions to differential equations. The simple plume model attributes include: diffusion constant, wind direction, and wind magnitude. The location of the plume prior to current observations is calculated statistically with the use of an estimator-based joint probability.**

The main contribution of this work is the predictor model - a required step of the MHT algorithm. A customized predictor for plumes (as opposed to Kalman filtering) allows the MHT-like algorithm to treat the plume tracking problem as the extreme instance of the multi-target tacking (MTT) problem. The central question: how can a MHT-like method be implemented for plumes in a sensor network of simple sensors capable of rudimentary binary detection, wind speed, and wind direction. The predictor must handle the problem of data association for plume observations. The context for this work is the development of multiple competing models which will be correlated to incoming observations in real time. The models will run in a generic multi-purpose framework called PQS (Process Query System).[1]

Simulations were performed demonstrating the viability of the MHT approach with the use of a customized predictor for plume target tracking.

***Index Terms*— Sensor networks, multiple-target tracking, plume tracking, data association, process query systems**

I. PROBLEM STATEMENT

Rapid detection of plume sources is a problem of intense interest to national security as well as environmental monitoring. Many of the current biological and chemical monitoring systems have detection lag times of hours or days due to the manual data analysis required to identify and attribute harmful agents. Unlike finite targets, a non-rigid plume spans a region and requires the assimilation of information about a non-localized continuum of targets [1], [2], [3]. Quickly assessing the current state of an effected region thus demands the automated processing of large numbers of observations yielding potentially competing hypotheses. Desired information about the harmful agents include source attribution (location), source distribution, and source magnitude. We seek to develop a method capable of handling large numbers of chemical observations by assignment to tracks and hypotheses which are constantly updated, pruned, and ranked on the basis of likelihood.

[1]For more information about PQS work at Dartmouth, see http://www.pqsnet.net/

A. General Problem

Given an observation sequence O^T of T sequential observations from a network of N sensors, what is the probability $P(O^T|S)$ where S is the current or initial state of the surrounding environment. We wish to estimate this likelihood. Given this current state estimation at time t_i what is S_{t_0}or $S_{t-\epsilon}$ where $S_{t-\epsilon}$ is the state at some previous time. S_{t_i} can be considered as as two dimensional matrix S of with width m and height n, and containing $m * n$ concentration values that evolve according to a diffusion process as a function of time. An estimate of S_{t_0} provides a description of the plume origins.

As opposed to the traditional subjects of target tracking such as mechanical vehicles, plume signal propagation is largely determined by wind and diffusion instead of much faster electromagnetic or sound energy [4]. This highly non-linear problem is not obviously solvable by current sensor network target tracking methods. The fluid dynamics of plumes typically mean chaotic meandering flows. It is common to measure a concentration of zero a majority of the time even in the proximity of the source, while large readings may be present at great distances [5]. Many sensor network applications performing inverse array acoustic localization for ground targets use sound intensity and time of travel, but other methods are required in the case of plumes. Another challenge is that chemical observations are often of extremely low resolution, and often only indicate a positive or negative result. Traditional methods for inverse array signal processing fail to transfer into the plume tracking domain.

B. Application of MHT Problem

We suggest monitoring plume sources with MHT, maintaining tracks from collections of individual observations, and tracing observations back to their origin. Unlike the canonical MTT problem which utilizes Kalman filtering, we can measure all the forces having an impact on the plume structure, whereas with traditional MTT the target may have an intelligent unpredictable component such as a pilot control [6]. Knowing the wind history vector ($[W_{t_{i-\epsilon}}, ..., W_{t_{i-2}}, W_{t_{i-1}}, W_{t_i}]$) for each of the N nodes, the substance of interests diffusion constant D, and the relative location of all the nodes in the network allows for the calculation of a plume predictor value for each new observation. Each O_i at sensor n has a probability of correlation for $t_{i-\epsilon} < t_i$ with event $O_{i-\epsilon}$ observed at a different sensor. The plume predictor estimates this probability, and uses the value in observation assignment to tracks, or track initiation. Essentially we ask, what is the probability that O_i and $O_{i-\epsilon}$ originated from the same plume event?

The goal is to find source locations, rank likelihoods, which will allow the determination of the number of sources. How well do current observations O_{t_i} in the sensing network correlate to the same original event (S_{t_0})? We propose the monitoring of a two dimensional area $(m * n)$ with a field of N stationary sensors, where N is large. Others have approached the plume tracking problem with groups of autonomous mobile robots [7], [8], [9], however our work examines the problem of static sensor networks.

In previous work [10] we primarily considered the use of the EnKF (Ensemble Kalman Filter) as a mechanism for observation correlation, however this paper considers the implementation of a coordinated MTT system using a much simpler model. EnKF methods are typically used for weather modeling in which high resolution atmospheric data is available from a relatively low number of nodes. Others have also approached the problem using complex models imported from the meteorology community [4]. The result of turbulent flow eddies in particle dispersion events is a highly discontinuous and intermittent distribution of plume particles [7], [5] where gradient following is not practical.

Current sensor network technology allows us to consider solving the problem with less complex models, compensating with a larger number of nodes. Low complexity static networks have the advantage of low energy consumption, lower device cost, and the potential for high density ubiquitous deployment to monitor for long periods of time. The multiple target plume predictor presented here is one such possible simple model that takes advantage of high sensor density.

In our model, the particles making up the plume of interest are all under the influence of the same forces in the surrounding environment. (Wind and diffusion across the the two-dimensional space are uniform). Our model of a chemical plume will consider only diffusion constant D and wind W, which can be observed at each sensor of known location. Given the observation of a known chemical at the sensor allows the sensor to lookup D . Because the wind medium can be measured at each node, the sensor network has complete knowledge of the only two parameters affecting plume dynamics in our model. The implemented wind model is a Markov model, producing a pseudo-random wind direction approximation of real wind.

The goals of our approach:

- Estimating plume source location
- Determination of plume "tracks"
- Few false positives
- Scalable algorithm (to a large sensor system)
- Near real-time tracking

The goal of this work is not improving upon MHT itself, but developing several plume predictors, which can then be inserted into existing implementations of MHT algorithms [6] such as PQS. The function of this plume predictor is the statistical correlation of individual observations, and the assignment of new observations to tracks. This approach supports the eventual development of multiple high level models of a dispersing chemical plume. High level models will enable the development of end users to submit plume query process models. Once these models are developed, several will be simultaneously submitted to a PQS framework. PQS has the ability to correlate incoming data against multiple plume models and rank their likelihoods. PQS queries of interest include: finding the plume source/origin, total mass released, number of distinct sources, and future predicted plume location.

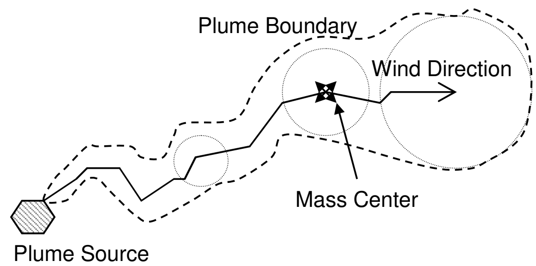

Fig. 1. Forward plume model propagation is a function of diffusion constant D and wind $\vec{W}$, with wind direction determining the dominant direction. The ratio $P_e = \frac{D}{\vec{W}}$, also known as the Peclet number determines the width of the plume region. The Peclet number (Pe) is a measure of the relative importance of advection to diffusion.

II. BUILDING A PLUME MODEL

A. Forward model

Before tackling the inverse plume problem a mathematical model of the forward diffusion process is needed. This forward model provides simulated data of diffusion events until actual field data can be collected for solving the inverse problem. In addition, developing a forward model and running simulations provides fundamental insights into the characteristics involved in the diffusion process.

Without the presence of wind, the fundamental process behind plume movement is diffusion described by Fick's Law. In two dimensions, assuming anisotropic diffusion:

$$\frac{\partial C}{\partial t} = D_x \frac{\partial^2 C}{\partial x^2} + D_y \frac{\partial^2 C}{\partial y^2}$$

The diffusion of in x depends only on the distribution in x and the diffusion in y depends only on the distribution in y. The solution of interest [3]:

$$C(x, y, t) = \frac{A_1 A_2}{4\pi t \sqrt{D_x D_y}} e^{(-\frac{x^2}{4D_x t} - \frac{y^2}{4D_y t})}$$

With A_1 and A_2 as constants for the x and y dimensions. The profile of concentration, $C(x, y)$ along any straight line cut through the patch will have a Gaussian distribution. If the concentration profile begins as a point we will see smoothing Gaussian distributions as time progresses. As a direct result of Fick's Law, the flux in any direction is proportional only to gradient in that direction. If the diffusion constants are anisotropic, the cloud will disperse an-isotropically, growing more quickly along the axis with a greater diffusion constant. The propagation length along any axis will be proportional to the diffusion coefficient along that axis:

$$L_x = 4\sigma_x = 4\sqrt{2D_x t}$$

With the introduction of wind the plume process becomes a combination of advection and diffusion. Assuming a uniform constant wind direction, the wind solution is a smoothing Gaussian

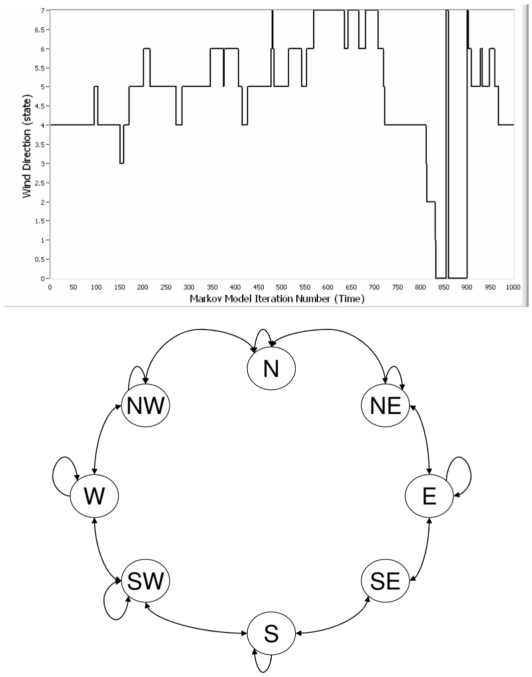

Fig. 2. The generation of pseudo-random wind using a Markov model and state transition probabilities for 8 wind direction states. The output of the wind model approximates real wind variation behavior.

distribution shifted in space linearly as time evolves. The resulting analytical solution to this partial differential equation (Fick's Law) in most situations must be obtained by approximation with numerical methods. Depending on the constraints, initial conditions, and boundary conditions of the diffusion problem a finite-difference may be possible.

The forward simulation allows for creation of arbitrary number of plume sources which may vary in area or intensity. Random plume initial conditions can also be created. The properties of plume number, size, and intensity constitute the initial concentration of the plume state space S_{t_0}. Once S is populated with initial plumes, the wind/diffusion processes iterate on the initial conditions. During the simulation wind can also be set to (0, constant, or variable random). While the forward simulation runs, S can be sampled by set sensors or N randomly placed sensors that report a binary detection when the concentration value exceeds a threshold. (This defines a detection).

In our forward simulation the parabolic partial differential diffusion equation uses the boundary conditions: concentration for all t is set to zero at the boundaries, the initial concentration values in the space are set by the user. The forward numerical simulation implements the forward-difference solution, with a typical grid size of $m = 250$ and $n = 250$. A pseudo-random wind component can be added to this solution, estimated by a Markov wind model with 8 directions (states) [11]. The characteristics of the wind can be altered by changing the transition probabilities between states. Decreasing the chance of staying in the same state ($P_{i \to i}$) increases randomness and frequent changes in wind direction - thus making the plume more difficult to track. In these experiments the Markov model was a uniform transition probability, where the likelihood of transition into the same state, $P_{i \to i} = .95$. The probability of transition into neighbor states was equal in both directions: $P_{i \to i+1} = P_{i \to i-1} = .05$. Two types of contaminant sources are possible in the simulation: limited source in which there is a one time release, constant release source in which additional mass is emitted from the source at every iteration. All the simulations in this paper used a constant release.

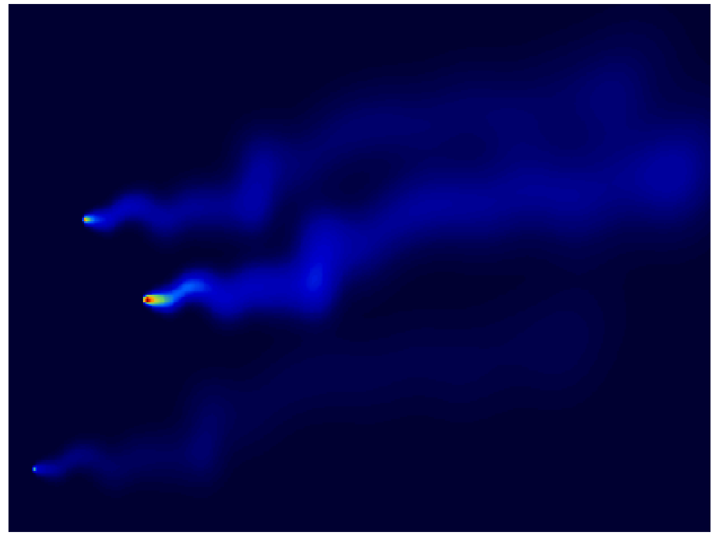

Fig. 3. Forward diffusion process implemented with numerical forward difference method and pseudo-random wind field. Three plume sources of different initial concentrations started this simulation.

B. Multiple Target Inverse Model

A plume may be considered as region with a center of mass, with ideal observations expected to be based on the diffusion equation and distance from the plume center. In practice however, plumes split into separate discontinuous filaments - moving in a chaotic flow regime. The plume may generate low readings near the source, or intermittent high concentration readings to sensors at great distance. Techniques for MHT are very appropriate for such intermittent data availability, and also can handle new track initiation and termination [6]. The use of Gaussian descriptions for atmospheric dispersion models can give rise to very misleading estimates of concentration fields. A plume is frequently not well dispersed, rather consists of a long sinuous volume of material. As a result, a detector with a fast response time will report a series of relatively short bursts of high concentration adjacent to long intervals during which the concentration is close to zero. Detections are essentially binary in nature. This property of plume dispersion called intermittentcy results in concentration readings of zero at a given sampling point.

MHT for plumes addresses the problem of assigning new observations to existing track hypotheses, or new hypothesis creation. (data association). Kalman or Bayesian filter predictors assume a strict temporal arrival order of observations and may have to throw out late-arrival data. The tracker may have have to repeat calculations to integrate the late-arrival in prior calculations, thus degrading performance. MHT maintains a ranking of data associations, based on likelihood. MHT suffers from

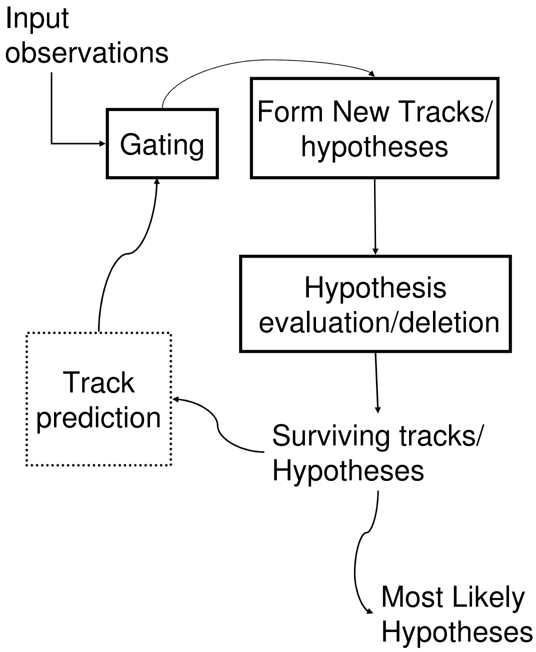

Fig. 4. Track Prediction as a part of the MHT algorithm: assignment of new observations to existing tracks or the decision to create to a new track requires track prediction based on the last known observations.

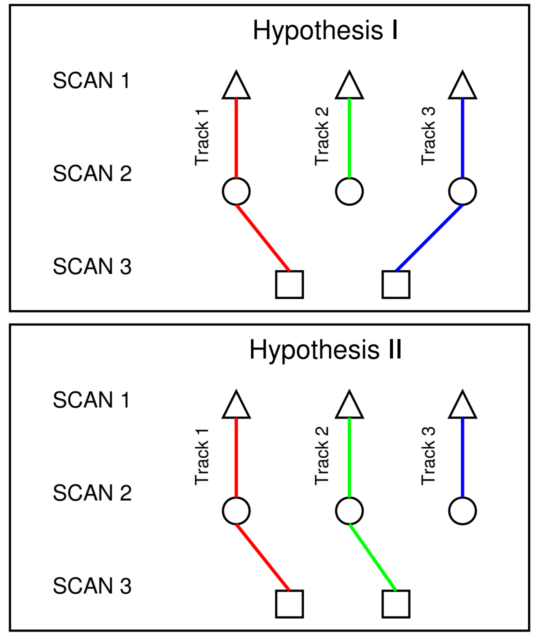

Fig. 5. The role of the predictor in MHT is observation assignment. Given the same set of observation scans multiple track permutations and hypothesis sets are possible.

combinatorial explosion in data association, however additional knowledge about target kinematics or the environment can help prune hypotheses [1]. In the case of plume tracking with variable wind, we have the advantage of having a complete knowledge of the wind vector history, which is essentially the propagation history for all particles in the region.

C. PLUME PREDICTOR

The standard Kalman predictor used in target tracking is not designed for large continuous mass - we require a custom predictor designed on simple binary sensor information. For each observation, the particle event will be projected backward in time onto a probability distribution based on wind history, substance type (D), and location. (The plume predictor).

Observing a single release plume event - one can calculate a total threshold radius - which is the total surface area in which the event can ever be detected. (Assuming a minimum threshold of detectability $Thresh_{min}$). Beyond a certain time maximum $T_{max}$the peak of the Gaussian distribution will sit beneath $Thresh_{min}$. This time T_{max} can be estimated, and the total area of detectability (TAD) calculated. After performing this analytical derivation for TAD and making a graphical plot of the TAD, it can be adjusted for various wind and D values. TAD produces a cone shape, with width determined by the Peclet number P_e , based on the ratio $P_e = \frac{D}{W}$. Within this cone area, the likelihood of detection at a given location and time can also be estimated. The geometric constraint results in a cone-shaped area that represents the entire zone of detectability over all time, given θ_i. The total area of detectability is a function of:

- wind, w
- diffusion constant, D
- total mass released, M

Distorted ice-cream cone geometries result when TAD is projected back-wards into the oncoming wind direction if wind varies over time. By estimating joint probabilities between multiple sensors and different TAD zones, overlap regions within boundaries provide additional information. By looking at the joint probabilities of two or more TAD zones (binary "YES" detections at two or more sensor locations) we can estimate joint probabilities of a plume origin to higher certainty. For example if multiple sensors have detections from overlapping TAD regions, we expect a lower probability of false detection for all current observations. For the sensors A and B, and the plume origin x:

$$P(B|A,x) = \frac{P(A,B|x)}{P(A|x)}$$

For three sensors (with corresponding cones) A, B, and C:

$$P(A,B,C) \quad = \quad P(A,B|C)P(C)$$

$$= P(A|B,C)P(B,C)$$

$$= P(A|B,C)P(B|C)P(C)$$

When observations are available from the sensor network, the current wind vector history for the relevant node is weighted with the observation. Thus, a large number of observations at a particular nodes gives more weight to its particular plume predictor. Once a series of these weighted plume predictors are compared against neighbor sensor nodes that lie within its overlaid plume predictor regions, relationships can be formed between nodes. A map of these integrated node-node relationships produces a connectivity map among nodes. This map represents

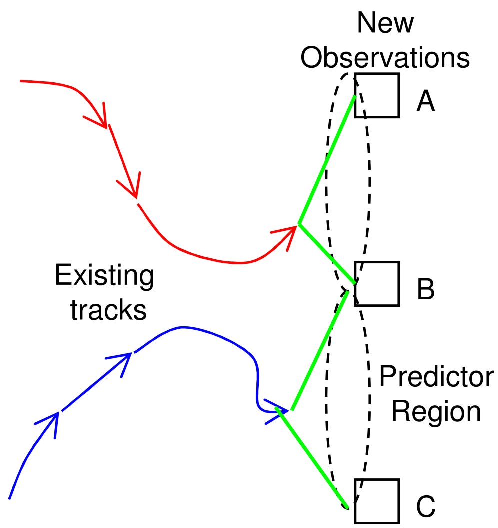

Fig. 6. Data association and plume predictor for MHT

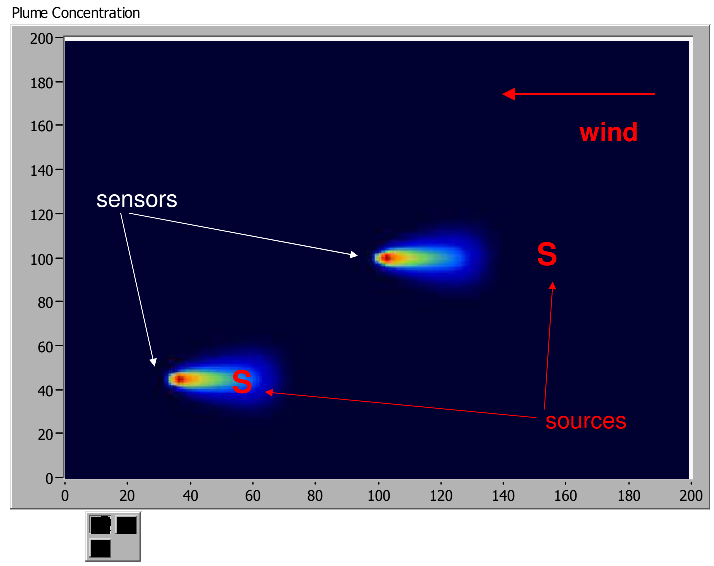

Fig. 7. Single plume predictors for 2 sensors, based on a single observation. Constant wind from the right. For this single observation the most likely attribution region can be seen to the right. No sensors are located in this upstream direction, therefore no connectivity relationships in the network would be affected.

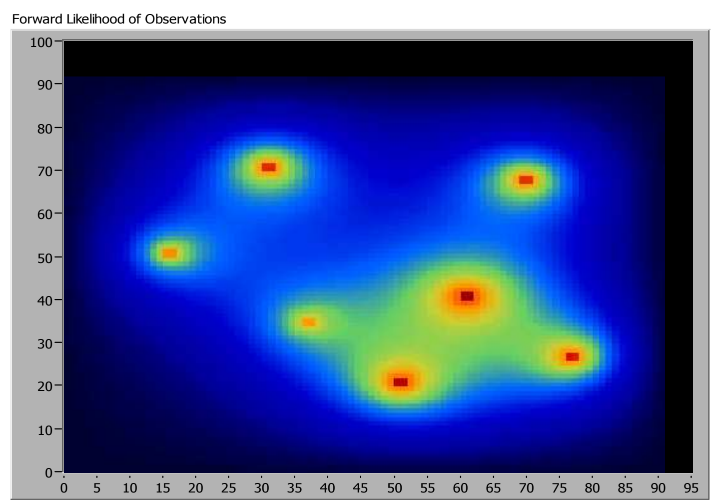

Fig. 8. Accumulated and overlaid plume predictors $\sum P_i$ for 7 sensors . The unidirectional probability distribution results from no wind. For a single observation, no information on direction is available. The predictors allow the calculation of connectivity values between neighboring sensor nodes.

the likelihood of 2 sensors being passed across by the same plume event. When large numbers of sensors are available in the region of interest, and as N increases, the connectivity map converges to an approximation of the relevant tracks that approximate a plume dispersion event.

III. Results

A. Plume Predictor for Single Observations

Using the plume predictor P_i at time t_i, given a single O_i at node n, O_i is correlated with W_i (wind history vector) to produce a probability distribution in two dimensions at t_i of Γ_i . The derivation of this space is based on the TAD, and represents the likelihood that an individual observation O_i originated at a given upstream location. In the color intensity plots, the color intensity represents this likelihood of attribution for O_i . Probabilities are higher along the axis of W, fading tangentially in a Gaussian distribution. Locations outside the TAD for a given W_i and D have an intensity value of 0 (black).

For the plume predictor with multiple accumulated observations, a gradual overlaid predictor map is created. The correlation output of $P_i * W_i$ for each O_i can be summed over all O, where each Γ_i is summed linearly with previous or future Γ. As more O are detected at sensor nodes, a greater degree of situational awareness develops, allowing for the weighting of the sensor connectivity map edges. Edges intensities between sensor nodes represent the probability of observation correlation.

B. Track Formation

The sensor network connectivity map is generated allowing the estimation of neighboring sensor relationships, and the estimation of plume tracks from the sampled observations series. Each sensor node is randomly placed on the space $S_{m,n}$ and generates weighted connectivity values depending on wind history, relative sensor location, and the total number of observations received that support a given node-node connection. These weighted edges represent potential tracks, and the sum of a probability along a series of tracks can be used to estimate the total likelihood for a hypothesis combination of tracks. The figure illustrates the output of the two most likely tracks in a given simulation run, as well as the associated edge likelihoods, ranked into 4 categories of probability. Comparison of these tracks demonstrates a high degree of correlation with the true plume path traversed.

IV. Conclusion

The system performance illustrates the viability and potential of using PQS process models within sensor networks for solving the inverse plume location problem. Moreover, the speedy development of this application illustrates the potential power of the PQS framework. Our group has been working with sensor

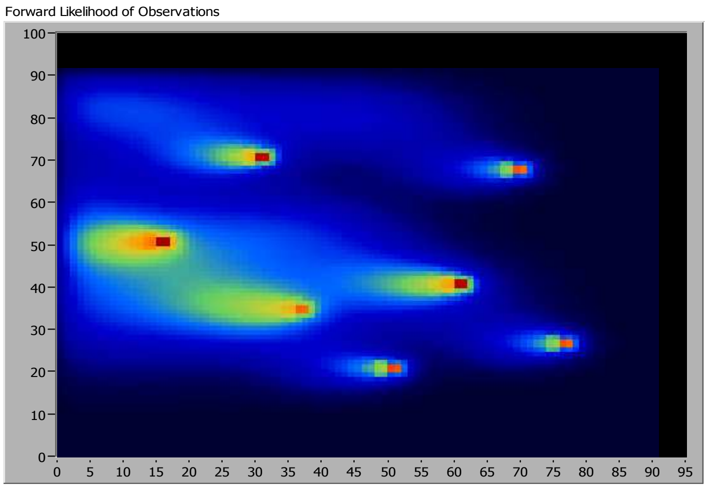

Fig. 9. Accumulated and overlaid plume predictors for 7 sensors, based on the changing wind vector history. $\sum \Gamma_i$. In this scenario the plume sources and wind originate from the middle left. Because several sensors are positioned within the plume predictor regions of neighbor nodes, connectivity relationships will be created or augmented. The predictors allow the calculation of connectivity values between neighboring sensor nodes.

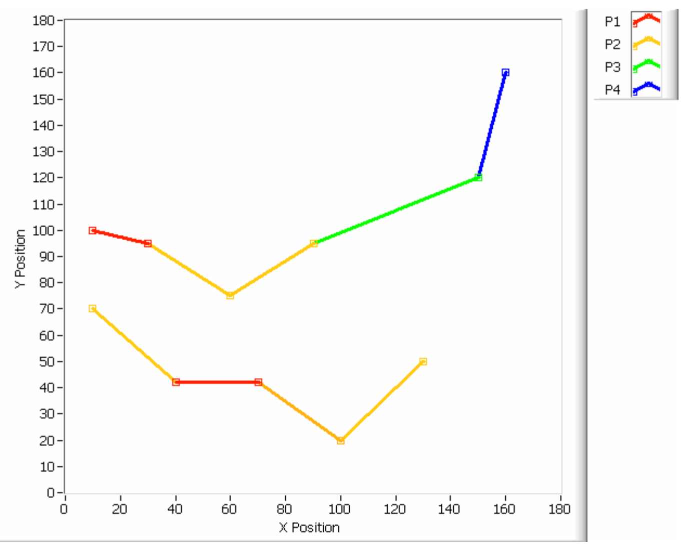

Fig. 10. The construction of plume track estimates showing four levels of association probability between sensor nodes, where $P_1 > P_2 > P_3 > P_4$. These are the 2 most most likely tracks and correctly identity two plume sources near the middle left edge of the plane. This experiment took place with variable wind from the left, m=n=250, N=50 randomly placed sensors.

networks for the past four years and has shifted its efforts away from the physical layer into the realm of sensor network queries, applications, and building a generic query processing infrastructure called the Process Query System (PQS). PQS is designed as a core engine that can handle sensor network data from a wide range of applications. So far we have implemented sensor network applications in the areas of physical vehicle multiple target tracking, live fish movement tracking, computer security monitoring, and chemical plume tracking in a sensor network. All these applications use the same core PQS engine.

The chemical plume tracking application takes real time observations of a chemical in the air environment, and allows a group of

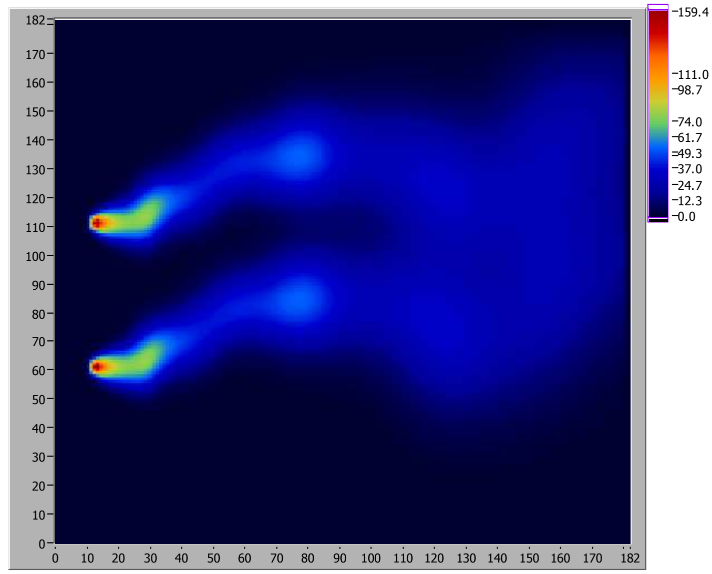

Fig. 11. Plume truth with two sources located at the mid-region of the left side.

sensors make statements about the chemical plume origin. These generated hypotheses about the plume origin, size, history, and shape are based on low resolution measurements of concentration in conjunction with wind and wind history. By correctly assigning observations to tracks with a plume predictor, the implementation of PQS and MHT is possible for plumes. Maintaining and pruning multiple tracks with MHT allows new hypotheses to be generated from old observations if data arrives non-sequentially. The novelty of this experiment is not the specific application, however the fact that a generic tracking framework can be applied to such a diverse set of domains. A domain expert can develop with PQS a description of a high level "process" in the environment, and very quickly have the system running, searching for instances of this high level process. This allows end users more time to focus on model development, as opposed to low level sensor network hardware and data acquisition details.

References

[1] Feng Zhao, Jie Liu, Leonidas Guibas, and James Reich, “Collaborative signal and information processing: An information directed approach,” *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1199–1209, 2003.

[2] S. Kazadi, R. Goodman, D. Tsikata, D. Green, and H. Lin, “An autonomous water vapor plume tracking robot using passive resistive polymer sensors,” *Autonomous Robots*, vol. 9, no. 2, pp. 175–188, 2000.

[3] Bernoit R. Cushman-Roisin and E. Deleersnijder, *Introduction to Geophysical Fluid Dynamics - Physical and Numerical Aspects*, under contract with Academic Press, Dartmouth College Hanover, NH 03755, 2005.

[4] Patrick D. French, John Lovell, and Nelson L. Seaman, “Distributed sensor network for local area atmospheric modeling,” *Unattended ground sensor technologies and applications V, proceedings of SPIE*, vol. 5090, pp. 214–224, 2003.

[5] C.D. Jones, “On the structure of instantaneous plumes in the atmosphere,” *Journal of Harardous Materials*, 1983.

[6] Donald B. Reid, “An algorithm for tracking mutiple targets,” *IEEE Transactions on Automatic Control*, vol. AC-24, no. 6, pp. 843–854, December 1979.

[7] Jay A. Farrell, Shuo Peng, and Wei Li, “Plume mapping via hidden markov methods,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 33, no. 6, 2003.

[8] A. Hayes, A. Martinoli, and R. Goodman, “Swarm robotic odor localization,” 2001.

[9] Adam T. Hayes, Alcherio Martinoli, and Rodney M. Goodman, "Distributed odor source localization," *IEEE Sensors Journal*, vol. 2, no. 3, 2002.

[10] Glenn Nofsinger and Keston Smith, "Plume source detection using a process query system," in *Defense and Security Symposium 2004 Proceedings*, Bellingham, Washington, April 2004, SPIE.

[11] Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, 1989.