

Multiple Source Detection and Localization in Advection-Diffusion Processes Using Wireless Sensor Networks

James Weimer, Bruno Sinopoli, Bruce H. Krogh
 Department of Electrical and Computer Engineering
 Carnegie Mellon University
 Pittsburgh, PA 15213

{jweimer | brunos | krogh}@ece.cmu.edu

Abstract—This paper concerns the use of large-scale wireless sensor networks to detect and locate leaks of specified gases in the presence of time-varying advection (air currents) and diffusion. We show that when leaks are rare but constant for long periods, Kalman filtering combined with binary hypothesis testing provides an effective alternative to full-scale hypothesis testing covering all possible combinations of leaks and leak intensities. To reduce energy consumption and use of communication bandwidth, a two-tiered strategy is proposed in which a reduced number of sensors (Tier 1) provides coarse-grid sensing. When a leak is detected by the Tier 1 strategy, fine-grained grids of sensors (Tier 2) are activated around the vicinities of the detected leak areas to provide more precise detection and localization. Energy consumption is further reduced by applying an information versus energy-based dynamic sensor selection technique. Details of a laboratory implementation are presented and simulation results illustrate the approach and demonstrate its effectiveness.

I. INTRODUCTION

Recent advances in low-power low-cost wireless sensing devices are revolutionizing the way we interact with the physical world. Self organized in networks, such devices deliver an unprecedented amount of information about the environments in which they are deployed, with a high degree of spatial and temporal granularity. Environmental monitoring and control applications, such as plume tracking, fire detection, and control of built environments, have flourished.

Motivated by ongoing research in carbon sequestration [11], this paper concerns the real-time energy-aware detection and localization of sources for a distributed dynamical field in an advection-dominated environment using a wireless sensor network (WSN), where the sources can be at locations other than the sensors. Since the data sampling periods are larger than the time it takes to communicate, a centralized approach to detection and localization is feasible. Since the wireless sensor nodes are incapable of performing computationally intensive calculations, a centralized approach is warranted. However, even a centralized approach to real-time detection and localization is affected by the stringent power and bandwidth constraints imposed by using a WSN for data acquisition. Limited network bandwidth combined with the sufficiently low data sampling

rates needed to guarantee long lasting deployments place a serious constraint on the number of data samples that can be collected and communicated within a sampling period.

These real-time and power constraints have led researchers to develop a variety of strategies for source detection and localization. Fox et al. [6] propose a sensor-level detection scheme as a trigger for global localization. Huang [8] presents a localized detection algorithm with localization through fusion, as an alternative to full-order centralized detection. Michaelides and Panayiotou [13] and Sung et al. [17] take advantage of spatial sensor correlations to perform sensor-level detection and localization. Nofsinger and Cybenko [14] describe a detection-based approach to tracking an airborne plume using sensor networks. Demetriou [3] performs low-power moving-source estimation and detection for two-dimensional processes using sensor scheduling and filtering for a single source.

In all of the above approaches, computational complexity increases with the length of the detection period and explodes combinatorially with the number of disturbances to be detected. Such computational growth prohibits real-time detection and localization. To avoid the growth of the computations with the detection period, Kerr proposes using Kalman filtering to generate a two-ellipsoid overlap test for disturbance detection [9]. Similarly, Brumback and Srinath [1] and Zolghadri [20] extend the seminal work of Kerr to develop other state-estimation-based tests. The principal shortcoming of these state-estimation-based approaches is that their performance cannot be bounded. Additionally, none of the above state-estimation-based approaches addresses the key issue of extending WSN lifetime.

In contrast to several of the papers reviewed above, we consider the problem of detecting and localizing multiple sources. We avoid the combinatorial explosion from multiple hypotheses by separating the issues of detection and localization. This leads to a simple binary hypothesis test for detection and a thresholding problem for localization that is easily implementable in real-time. The initial detection algorithm (Tier 1) uses the Neyman-Pearson detection criterion and a Kalman filter to produce a coarse process estimation using sparse sensor measurements. Upon initial

detection, a fine process estimation (Tier 2) verifies detection and accurately localizes the source(s) using another Kalman filter and Neyman-Pearson detector with full-order sensing in and around the detection area(s). We also introduce a method for performing dynamic sensor selection to further extend the network lifetime. We apply these techniques to the problem of surface monitoring for leak detection on CO₂ sequestration sites [5].

The following section introduces a lumped-parameter state-space model for the advection-diffusion of atmospheric gas concentrations as the basis for WSN monitoring of CO₂ sequestration sites and develops a general strategy for performing source detection. Section 3 discusses the two-tier localization strategy. Section 4 presents source detection and localization simulation results for a field of 961 sensors. Section 5 describes a method to further extend the network lifetime by applying a relaxation-based approach to dynamic sensor selection, and presents simulation results. Section 6 describes our real-time implementation of the detection and localization strategy using the Firefly wireless sensor nodes [12]. The final section summarizes the work presented in this paper and identifies directions for future research.

II. SOURCE DETECTION

We consider the problem of monitoring surface concentrations of specified gases to detect and locate sources of leaks in the presence of time-varying advection (air currents) and diffusion. In the surface monitoring problem, measurements are obtained through sensors connected in a WSN and decisions are based on these noisy measurements and model parameters. We consider the concentration of a single gas as a function of space and time, denoted by $c(p,t)$, where $p = [x, y, z]$ is the location vector and t is time. The concentration dynamics are governed by the following continuous-time partial differential equation (PDE) describing an advection-dispersion process [2]:

$$\frac{\delta c(p,t)}{\delta t} + \phi(p,t) \frac{\partial c(p,t)}{\partial p} = \alpha(p,t) \frac{\partial^2 c(p,t)}{\partial p^2}, \quad (1)$$

where $\phi(p,t) = [\phi_x(p,t), \phi_y(p,t), \phi_z(p,t)]^T$ and $\alpha(p,t) = [\alpha_x(p,t), \alpha_y(p,t), \alpha_z(p,t)]^T$ are the advection and dispersion coefficients, respectively. The surface boundary condition is

$$-\alpha_z(p,t) \frac{\delta c(p,t)}{\delta z} \Big|_{p=(x,y,0)} = \lambda(x,y,t),$$

where $\lambda(x,y,t)$ is the boundary condition at $z = 0$ (the surface).

We assume the sensors are on the surface. Since there are no observations above the surface, an approximation is needed for $\frac{\delta^2 c(p,t)}{\delta z^2}$ at $z = 0$. This is achieved by assuming symmetric diffusion, that is,

$$\frac{\delta^2 c(p,t)}{\delta x^2} = \frac{\delta^2 c(p,t)}{\delta y^2} = \frac{\delta^2 c(p,t)}{\delta z^2}.$$

Applying these assumptions to (1), we use the following estimate for $\frac{\delta^2 c(p,t)}{\delta z^2}$ at the surface:

$$\frac{\delta^2 c}{\delta z^2} = \frac{\frac{\delta c}{\delta t} + \phi_x \frac{\delta c}{\delta x} + \phi_y \frac{\delta c}{\delta y} - \frac{\phi_z}{\alpha_z} \lambda(x,y,t)}{\alpha_x + \alpha_y + \alpha_z}, \quad (2)$$

where, for convenience, we write c to mean $c(p,t)$, and similarly for α and ϕ . Substituting (2), into (1) yields:

$$\begin{aligned} \frac{\delta c}{\delta t} = & \frac{\alpha_x + \alpha_y + \alpha_z}{\alpha_x + \alpha_y} \left(\alpha_x \frac{\delta^2 c}{\delta x^2} + \alpha_y \frac{\delta^2 c}{\delta y^2} \right) \\ & - \phi_x \frac{\delta c}{\delta x} - \phi_y \frac{\delta c}{\delta y} + \frac{\phi_z}{\alpha_z} \lambda(x,y,t), \end{aligned} \quad (3)$$

with boundary and initial concentrations equal to c_0 .

Equation (3) is a 2-D approximation of the 3-D PDE in (1). The approximation in (3) is similar to the 2-D PDE proposed in [3], differing only by scaling. The difference in scaling accounts for the nonmeasurable effects of vertical advection and diffusion inherent in atmospheric monitoring problems where planar sensing is assumed.

For simplicity, we assume the sensors are placed on the surface in a $N \times M$ grid. These points define the locations of the concentrations used as state variables in a lumped-parameter model of the concentration dynamics. We consider sources (leaks) generated by a discrete-time linear system of the form:

$$z_{k+1} = z_k + h_k, \quad (4)$$

where $z_k = \{\bar{\lambda}_1(k), \bar{\lambda}_2(k), \dots, \bar{\lambda}_{(n-1)M+m}(k), \dots, \bar{\lambda}_{NM}(k)\}$ is the set of source inputs corresponding to the values of $\lambda(p, k\Delta_t)$ on the surface regions $p \in \{(n-1)\Delta_x/2, (n+1)\Delta_x/2, [(m-1)\Delta_y/2, (m+1)\Delta_y/2], 0\}$, h_k is the leak process noise, and $n \in \{1, \dots, N\}$ and $m \in \{1, \dots, M\}$ [3]¹. Similarly, we generate the discrete-time state space model for (3) by applying an Euler approximation to (3) and discretizing in time gives

$$\begin{aligned} x_{k+1} &= A_k x_k + \Gamma_k B_k z_k + w_k \\ y_k &= x_k + v_k, \end{aligned} \quad (5)$$

where $x_k = [c_1(k), c_2(k), \dots, c_{(n-1)M+m}(k), \dots, c_{NM}(k)]^T$ is the set of concentration state variables corresponding to $c(n\Delta_x, m\Delta_y, 0, k\Delta_t)$, y_k are the noisy measurements, w_k and v_k are the process and measurement noise respectively, and A_k and B_k are both square matrices of dimension $NM \times NM$ representing the lumped parameter state dynamics governing $x(k)$ according to (3) [7] [15]. Γ_k is a diagonal matrix with binary-values that specify which sources are active at time

¹In this paper we allow for possible leaks at every grid point. The dimension of z_k can be reduced, with a corresponding reduction in the computations, if it is known a priori that leaks cannot occur in some locations.

k :

$$\Gamma_k = \begin{bmatrix} \gamma_k^1 & & 0 \\ & \ddots & \\ 0 & & \gamma_k^{MN} \end{bmatrix}, \quad (6)$$

where $\gamma_k^{mn} \in \{0, 1\}$. When $\gamma_k^{mn} = 0$, a source is considered inactive, while $\gamma_k^{mn} = 1$ equates to an active source. We assume that once a source becomes active at some time $k = k'$ ($\gamma_{k'}^{mn} = 1$), it will remain active indefinitely ($\gamma_k^{mn} = 1$ for all $k \geq k'$). We also assume there is no a priori information on the probabilities of the entries of Γ_k equalling 0 or 1.

The system dynamics in (5) are augmented to include the leaks dynamics in (4) yielding:

$$\begin{bmatrix} x_{k+1} \\ z_{k+1} \end{bmatrix} = \begin{bmatrix} A_k & \Gamma_k B_k \\ 0 & I \end{bmatrix} \begin{bmatrix} x_k \\ z_k \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} w_k \\ h_k \end{bmatrix}$$

$$y_k = [I \quad 0] \begin{bmatrix} x_k \\ z_k \end{bmatrix} + v_k. \quad (7)$$

We define the prior distribution on the system state as

$$\begin{bmatrix} \tilde{x}_{0|-1} \\ \tilde{z}_{0|-1} \end{bmatrix} : \mathcal{N} \left(\begin{bmatrix} \hat{x}_{0|-1} \\ \hat{z}_{0|-1} \end{bmatrix}, \begin{bmatrix} P_{0|-1}^x & P_{0|-1}^{xz} \\ (P_{0|-1}^{xz})^T & P_{0|-1}^z \end{bmatrix} \right), \quad (8)$$

where $\mathcal{N}(\mu, \Sigma)$ denotes the multivariate Gaussian distribution with mean vector μ and covariance matrix Σ . Further, we assume the noise variables are distributed as

$$\begin{bmatrix} \tilde{w}_k \\ \tilde{h}_k \\ \tilde{v}_k \end{bmatrix} : \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} W & 0 & 0 \\ 0 & H & 0 \\ 0 & 0 & V \end{bmatrix} \right). \quad (9)$$

Applying the observation model in (7) to the prior in (8), we write the prior distribution on the observations as:

$$\tilde{y}_{0|-1} : \mathcal{N}(\hat{x}_{0|-1}, P_{0|-1}^x + V). \quad (10)$$

Using this formulation, we wish to determine which (if any) sources are active at time $k = K$ based on the observations up to time K , $Y_K = \{y_0, \dots, y_K\}$. It is standard practice to formulate a test for every possible source combination. Assuming a network with 30 sensors and 30 potential source locations, the standard detection problem has $\sim 2.6 \times 10^{32}$ hypotheses to test [6]. Moreover, if each potential source can also have K different intensities the number of tests increases with K factorial. It is obvious that for applications where the number of potential sources and potential intensities is large, the standard hypothesis testing problem becomes intractable.

As an approach that is computationally feasible and can be implemented in real-time, we propose to separate the problems of source detection and source localization by first determining whether *any* source is active (detection), and then deciding *which* sources are actually active (localization). This reduces the detection problem to a *binary* hypothesis testing problem where we write the state dynamics in (5) as:

$$\begin{aligned} x_{k+1} &= A_k x_k + \gamma_k B_k z_k + w_k \\ y_k &= x_k + v_k, \end{aligned} \quad (11)$$

where $\gamma_k \in \{0, 1\}$ represents whether *all* sources are inactive ($\gamma_k = 0$), or *some* source is active ($\gamma_k = 1$). The reason $\gamma_k = 1$ corresponds to *some* of the sources are active is because inactive sources are effectively active sources with a zero-magnitude. This concept plays a key role in the localizing sources once it is determined that some sources are active. We assume that once any source becomes active at time k' ($\gamma_{k'} = 1$), it will remain active indefinitely ($\gamma_k = 1$ for all $k \geq k'$). This new system reduces the detection problem into a binary hypothesis test, where we test the likelihood of no active source versus the likelihood of some active source(s).

In systems with many states and extended observation periods, the optimal Neyman-Pearson detector is not computationally feasible since its dimensionality grows with the observation period. As a heuristic solution, we propose to combine state-estimation techniques and the optimal detector strategy to generate a feasible detector. We begin by designing an optimal detector of dimension $M \times N \times J$, where J defines the window size for which optimal Neyman-Pearson detection can be performed. During an initial transient (when $K \leq J$), the optimal Neyman-Pearson detector is used; once $K > J$, we switch to the fixed-sized detector to make our decision based on the latest J observations, $Y_{K-J+1, K} = \{y_{K-J+1}, \dots, y_K\}$. To generate an estimate of the probability distribution for $Y_{K-J+1, K}$ under each hypothesis, we use a Kalman filter to compute the mean and covariance matrix for y_{K-J+1} from the observations y_0, \dots, y_{K-J} , and then generate the statistics for the complete vector $Y_{K-J+1, K}$ by applying the system dynamics. This is a heuristic that allows us to apply Neyman-Pearson detection over the window of length J using all the available observations.

The Kalman filter prediction and innovation update equations for the augmented system are written as

$$\begin{aligned} \begin{bmatrix} \hat{x}_{k+1|k}^Y \\ \hat{z}_{k+1|k}^Y \end{bmatrix} &= \begin{bmatrix} A_k & \gamma_k B_k \\ 0 & I \end{bmatrix} \begin{bmatrix} \hat{x}_{k|k}^Y \\ \hat{z}_{k|k}^Y \end{bmatrix} \\ P_{k+1|k}^Y &= \begin{bmatrix} A_k & \gamma_k B_k \\ 0 & I \end{bmatrix} P_{k|k}^Y \begin{bmatrix} A_k^T & 0 \\ \gamma_k B_k^T & I \end{bmatrix} + \begin{bmatrix} W & 0 \\ 0 & H \end{bmatrix} \\ K_k^Y &= P_{k|k-1}^Y \begin{bmatrix} I \\ 0 \end{bmatrix} (P_{k|k-1}^{x, \gamma} + V)^{-1} \\ \begin{bmatrix} \hat{x}_{k|k}^Y \\ \hat{z}_{k|k}^Y \end{bmatrix} &= \begin{bmatrix} \hat{x}_{k|k-1}^Y \\ \hat{z}_{k|k-1}^Y \end{bmatrix} + K_k (y_k - \hat{x}_{k|k-1}^Y) \\ P_{k|k}^Y &= (I - K_k [I \quad 0]) P_{k|k-1}^Y, \end{aligned} \quad (12)$$

where

$$P_{k|k-1}^Y = \begin{bmatrix} P_{k|k-1}^{x, \gamma} & P_{k|k-1}^{xz, \gamma} \\ P_{k|k-1}^{zx, \gamma} & P_{k|k-1}^z \end{bmatrix}.$$

By applying our knowledge of the observation equation in (7), we write the distribution on the observed data at time $k = K - J + 1$ in terms of γ as:

$$\tilde{y}_{K-J+1|K-J}^\gamma : \mathcal{N}[\hat{y}_{K-J+1|K-J}^\gamma, P_{K-J+1|K-J}^{\gamma,\gamma}], \quad (13)$$

where

$$\begin{aligned} \hat{y}_{K-J+1|K-J}^\gamma &= \hat{x}_{K-J+1|K-J}^\gamma \\ P_{K-J+1|K-J}^{\gamma,\gamma} &= P_{K-J+1|K-J}^{\gamma,\gamma} + V. \end{aligned} \quad (14)$$

In this approach, the observations are used (via state estimation) to characterize the distribution of y_{K-J+1} . By applying the standard whitening techniques [18], a detection problem can be formulated on the observation set $Y_{K-J+1,K}$.

This binary hypothesis testing problem is written as :

$$\begin{aligned} H_0 : \tilde{Y}_{K-J+1,K} &: \mathcal{N}[\hat{Y}_{K|K-1}^0, P_{K|K-1}^{\gamma,0}] \\ H_1 : \tilde{Y}_{K-J+1,K} &: \mathcal{N}[\hat{Y}_{K|K-1}^1, P_{K|K-1}^{\gamma,1}]. \end{aligned} \quad (15)$$

The heuristic approach in (15) is a binary hypothesis testing problem where the decision rule for choosing H_0 versus H_1 is determined using standard techniques [18].

The optimal binary hypothesis testing problem grows in dimension with time, therefore it has a complexity of $O((MNK)^3)$, while the heuristic binary hypothesis testing problem in (15) has a constant dimension and complexity $O((MN)^3)$. Since the size of the heuristic detection problem does not grow with time, this approach is well-suited for real-time systems that require long observation periods.

III. SOURCE LOCALIZATION

The previous section presented a method for deciding whether no source or some source is active. Once the detector has decided that some source is active at time $k = K$ ($\gamma_K = 1$), we wish to localize which sources are most likely to be active (which $\gamma_K^{mn} = 1$). Additionally, since we are using a WSN to gather observations, we would like to extend the lifetime of the network as much as possible. This section introduces a two-tiered detection strategy and a threshold-based approach to localization that both extends the WSN lifetime and finds the locations of the sources.

A. Two-Tiered Detection Strategy

The detection strategy introduced in the previous section assumes every sensor reports a measurement at each time step. In WSNs containing hundreds of sensors, this is not attractive due to bandwidth and energy constraints. Additionally, due to power constraints on WSNs, it is desirable to perform source detection using as few sensors as possible, especially if the likelihood of an event occurring is small relative to the null hypothesis. These constraints motivate the following two-tier detection and localization strategy where an initial *coarse* detection strategy requires only a subset of the sensors to report measurements at each time step. Then,

when an event occurs, a *fine* detection strategy is employed where all the sensors in the area around the perceived source (as determined by the coarse detector) report measurements to verify and localize the source.

We outline the two-tier detection strategy as :

- 1) Initialize the Kalman filters for coarse detection using rotating subsets of the sensor nodes.
- 2) Perform coarse detection test until H_1 is decided (when $t(y) > \eta$).
- 3) Choose suspect sources from source estimates under H_1 in coarse detection. Determine which suspect sources should be investigated together as a group based on spacial proximity and define test regions to enclose each source grouping.
- 4) Perform fine detection on the regions defined by each test region.
- 5) If *all* fine detectors decide H_0 , go to 1), else, for each fine detection deciding H_1 , record the source location and magnitude, and report an active source.

We assume that all sensors are distributed in a grid format, where each sensor corresponds to a state in the fine detector. From this sensor distribution, we assign sensors in the same spatial proximity to a “sensor group”. The center of each sensor grouping represents a state in the coarse detector, and is chosen such that the coarse detector covers the entire monitoring area. At every time step, the coarse detector takes measurements from one sensor in each sensor group, cycling through the sensors in each group to distribute the energy consumption among the sensors. Errors resulting from the spatial differences in the sensors are considered part of the process noise in the dynamic model for the coarse detector. The energy savings from performing coarse detection comes at the cost of lower detection accuracy.

B. Threshold Localization

In our recursive detection algorithm, the Kalman filter assuming $H_{1,k}$ provides an estimate of the mean and covariance for the sources. We use these estimates to identify potentially active sources through thresholding, where the sources with mean values above the thresholds will be considered active. To determine which sources should be considered active we first determine which source has the largest magnitude. We identify this source as an active source and then define a threshold (η_S) for determining all other active sources according to:

$$\eta_S = \sigma_S \text{erfc}^{-1}(1 - P_S) + m_S, \quad (16)$$

where $\text{erfc}^{-1}(\cdot)$ is the inverse error function, $m_S = \max(\hat{z}_{k-J+1|k-J}^1)$ represents the largest source estimate in $\hat{z}_{k-J+1|k-J}^1$, and σ_S is the corresponding variance of m_S , given in the source covariance matrix, $P_{k-J+1|k-J}^{\gamma,1}$ [18]. $P_S \in [0, 1]$ is an *a priori* probability defining the probability of a leak

being greater than η_S . From this value of η_S , *inactive* and *active* sources are defined as:

$$\hat{z}_{k-J+1|k-J}^1(i) \underset{\text{inactive}}{\overset{\text{active}}{>}} \eta_S, \quad (17)$$

where $\hat{z}_{k-J+1|k-J}^1(i)$ is the i -th element of $\hat{z}_{k-J+1|k-J}^1$.

The following section shows simulation results for the detection strategy of the previous section combined with the localization and energy management strategies described above.

IV. SIMULATION RESULTS

We implemented the detection strategy presented in the previous sections in MATLAB. For each simulation, concentration fields are generated using a fine-grained finite-difference approximation of (1) for a monitoring field of $300m$ by $300m$. For all simulations, we define the detection problem parameters as :

- $N = M = 31$, (961 sensors)
- Number of sensor groups = 100
- $P_{FA} = .10$ (coarse detector)
- $P_{FA} = .01$ (fine detector)
- $J = 5$ (observation window)

In the first simulation, we define the normalized PDE parameters as:

- $\alpha_x(p, t) = \alpha_y(p, t) = \alpha_z(p, t) = \frac{1}{50} \frac{(10m)^2}{5min}$
- $\phi_x(p, t) = \frac{1}{5} \cos(\frac{\pi t}{10}) \frac{10m}{5min}$
- $\phi_y(p, t) = \frac{1}{5} \frac{10m}{5min}$
- $\Delta_x = \Delta_y = \frac{1}{30} (10m)$
- $\Delta_t = 1(5min)$

with three sources:

- $\lambda(\frac{1}{5}, \frac{1}{2}, t) = 1 \frac{kg}{(10m)^2 5min}$ for $t > 20$
- $\lambda(\frac{2}{5}, \frac{3}{10}, t) = \frac{3}{4} \frac{kg}{(10m)^2 5min}$ for $t > 20$
- $\lambda(\frac{3}{5}, \frac{7}{10}, t) = \frac{1}{2} \frac{kg}{(10m)^2 5min}$ for $t > 20$
- else $\lambda(x, y, z, t) = 0$

This simulation contains three sources separated spatially, but with two of the sources significantly affecting the region near another source through advection and diffusion as shown by the concentration profile at $k = 100$ in Fig. 1. This figure shows three regions of the monitoring field with elevated concentration levels.

The coarse detection provides six source estimates as shown in Fig. 2 by squares, where the dots represent the sensors and x's represent the centers of the course sensor groups. At time $k = 54$, the coarse detector source estimates identify four fine detection regions as illustrated in Fig. 2 by the dashed rectangles. A fine detection is performed on each fine detection region separately and three sources are predicted (spatially) as the circles in Fig. 2. The circles in Fig. 2 match the actual spatial position of the sources listed above. The predicted magnitudes of the sources at time




Figure 1. Concentration values in field monitoring area at $k = 100$ for simulation 1.

$k = 74$ (from the fine detector) are shown in Fig. 3. From Fig. 3, we can tell that the fine detector accurately predicts the magnitudes of the leaks as specified in Fig. 2. Additionally, Fig. 3 also illustrates that in some fine detection regions, no sources were present.




Figure 2. Sensor grid vs. detection outputs for simulation 1.

In the above simulation, all the sources were predicted to the same spatial accuracy as the full-order detection problem (containing 961 sensors) using only 100 sensors at each time step during coarse detection and 324 different sensors during fine detection. Since the probability of source being *active* is assumed to be small, using only 100 sensors at each time step represents a significant savings in both energy and bandwidth. When a leak occurs, the strategy identifies the leak locations accurately in an energy efficient manner by requesting data from only the course sensor groups until some possible leaks are detected.




Figure 3. Leak estimates vs. coarse detection regions for simulation 1.

In the second simulation, we define the normalized PDE parameters the same as the first simulation except:

- $\phi_x(p, t) = 0$;
- $\phi_y(p, t) = \frac{1}{5} \cos\left(\frac{\pi t}{10}\right) \frac{10m}{5min}$

We define the following three sources:

- $\lambda\left(\frac{3}{10}, \frac{3}{10}, t\right) = 1 \frac{kg}{(10m)^2 5min}$ for $t > 20$
- $\lambda\left(\frac{2}{5}, \frac{3}{10}, t\right) = \frac{1}{2} \frac{kg}{(10m)^2 5min}$ for $t > 20$
- $\lambda\left(\frac{7}{10}, \frac{3}{5}, t\right) = 1 \frac{kg}{(10m)^2 5min}$ for $t > 20$
- else $\lambda(x, y, z, t) = 0$

This is the same number of sources as in the previous simulation, but two of the sources are so close to each other that the effects of each source merge in the monitoring field. Additionally, one source is 50% smaller than the others. These sources (and the advection-diffusion dynamics) generate a concentration profile at $k = 100$ as shown in Fig. 4. This figure shows only two regions of the monitoring field with elevated concentration levels, despite the existence of three sources.

The coarse detection provides five source estimates as shown in Fig. 5 by squares. The coarse detector source identifies two fine detection regions at $k = 59$ as illustrated in Fig. 5 by the dashed rectangles. A fine detection is performed on each fine detection region separately and three sources are identified at time $k = 79$, as indicated as the circles in Fig. 5. As in the previous simulation, the circles in Fig. 5 again match the actual spatial position of the sources listed above and the predicted magnitudes of the sources (shown in Fig. 6) match the actual source magnitudes.

This second simulation shows that even sources near one another can be detected accurately, despite significant differences in magnitudes. As in the previous simulation, the coarse detector uses 100 sensors to report at each time step. However, only 185 sensors were needed to perform fine detection, resulting in an accurate source prediction.




Figure 4. Concentration values in field monitoring area at $k = 100$ for simulation 2.




Figure 5. Sensor grid vs. detection outputs for simulation 2.




Figure 6. Leak estimates vs. coarse detection regions for simulation 2.

The results in this section show that the proposed strategy can accurately detect and localize sources. To increase the WSN lifetime, the following section presents a method of performing reduced order sensing using dynamic sensor selection which can significantly increase the WSN lifetime at a cost in time to detection.

V. DYNAMIC SENSOR SELECTION

As an independent method of prolonging the network lifetime, dynamic sensor selection (DSS) can be applied to use fewer sensors, possibly with a reduction in state estimation accuracy. DSS can be used in addition to the detection and localization scheme previously introduced by performing DSS over the tier 1 sensor groups. We proposed a method of sensor selection in [19]. In this previous work, we introduced a heuristic relaxation approach for dynamic sensor selection assuming a constant number of sensors would be selected at each time step. This section extends the previous work to consider both channel reliability and energy consumption in performing DSS.

Following a similar formulation as presented in [19], we rewrite the observation equation in (7) to incorporate sensor selection as:

$$y_k = \begin{bmatrix} U_k Z_k & 0 \end{bmatrix} \begin{bmatrix} x_k \\ z_k \end{bmatrix} + v_k, \quad (18)$$

where $Z_k \in \mathcal{Z}$ is the DSS matrix at time step k . \mathcal{Z} is the set of all possible DSS matrices defined as:

$$\mathcal{Z} = \left\{ Z \in \{0, 1\}^{s \times l} \mid \mathbf{1}_s^T Z \in \{0, 1\}^l, Z \mathbf{1}_l = \mathbf{1}_s \right\},$$

and $s \in \{0, 1, \dots, l\}$ is the number of sensors selected to report measurements. U_k is the channel selection matrix. U_k^T is defined as the basis of the column space of $Z_k \Gamma_k Z_k^T$, where Ω_k is the channel reliability matrix defined as the binary diagonal matrix:

$$\Omega_k = \begin{bmatrix} \omega_1(k) & & 0 \\ & \ddots & \\ 0 & & \omega_l(k) \end{bmatrix},$$

where each $\omega_i(k) \in \{\omega_1(k), \dots, \omega_l(k)\}$ is a bernoulli random variable with $P[\omega_i(k) = 1] = \bar{\omega}_i$. Since the selected measurements are gathered using a lossy WSN, some selected measurements will not be received. In the model, U_k accounts for packet loss within the network.

Applying the updated observation equation in (18) to the Kalman filter equations in (12), we write the next step *posteriori* inverse error covariance matrix as :

$$P_{k+1|k+1}^{-1} = \left(\begin{bmatrix} A_k & B_k \\ 0 & I \end{bmatrix} P_{k|k} \begin{bmatrix} A_k^T & 0 \\ B_k^T & I \end{bmatrix} + \begin{bmatrix} W & 0 \\ 0 & H \end{bmatrix} \right)^{-1} + (V_{k+1}^{-1} Q_{k+1} \Gamma_{k+1})^{-1}, \quad (19)$$

where $Q_{k+1} = Z_{k+1}^T Z_{k+1} \in \mathcal{Q}$ is the square sensor selection matrix. \mathcal{Q} is the set of all possible square sensor selection matrices, defined as:

$$\mathcal{Q} = \left\{ Q \in \{0, 1\}^{l \times l} \mid \mathbf{1}^T Q \mathbf{1} = \text{trace}(Q) \right\}. \quad (20)$$

Consequently the number of sensors selected, s , is equal to the $\text{trace}(Q_{k+1})$ and Z_{k+1}^T is the basis of the column space of Q_{k+1} .

In DSS, we assume $H_{1,k}$ ($\gamma_k = 1$) so that we will choose sensors that will improve the Kalman filter estimation for both x_k and z_k . Assuming $\gamma_k = 1$ only applies to dynamic sensor selection and does not reflect the true value of γ_k in the detection and localization strategy.

A. DSS using Information versus Energy

We wish to maximize the sensor network lifetime while still providing enough observations to perform detection and localization. Thus, the DSS strategy must trade-off between the accuracy of the state estimate and the energy required to obtain the selected measurements. One common means to improving the state estimate accuracy is to select measurements such that the mean squared error is minimized. Equivalently, this corresponds to minimizing the trace of the error covariance matrix. The energy required to obtain the state estimate is a function of the sensor selection matrix, routing protocol, medium access channel (MAC), transmission power, channel reliability, and potentially more factors, depending on the application. To determine which sensors will produce the "best" system state estimate, we introduce the objective function:

$$J_\infty = (1 - \mathcal{C}) \text{trace}(P_{\infty|\infty}) + \mathcal{C} E_\infty, \quad (21)$$

where $\text{trace}(P_{\infty|\infty})$ is the steady state mean squared error, $\mathcal{C} \in [0, 1]$ is the energy cost factor, and E_∞ is the energy cost function at steady state which is assumed to be convex with respect to the sensor selection matrix. By increasing \mathcal{C} , the objective function will increase the weight on the cost of the energy required to obtain the state estimate.

For applications with time-varying dynamics, the $\text{trace}(P_{\infty|\infty})$ cannot be calculated since the future dynamics are unknown. Therefore, as a greedy approximation, we introduce the one-step objective function as:

$$J_{k+1} = (1 - \mathcal{C}) \text{trace}(P_{k+1|k+1}) + \mathcal{C} E_{k+1}, \quad (22)$$

where E_{k+1} is the next state energy cost. E_{k+1} can be determined several ways. Here we assume that the next state energy cost is proportional to the number of hops required to transmit the observations to the base station where centralized computing occurs. The one-step optimal DSS strategy at time step k is then the collection of sensors minimizing J_{k+1} :

$$Q_{k+1} = \arg \min_{Q_{k+1} \in \mathcal{Q}} J_{k+1}. \quad (23)$$

The minimization problem in (23) is a 0-1 integer programming problem with exponential complexity. For large-scale systems, this optimization problem is not feasible. Following the relaxation to a convex optimization problem in [19], the optimization problem is formulated for the objective function in (22).

B. Simulation Results

To perform the optimization step in the proposed sensor selection strategy, we use the MPT toolbox [10]. The following simulation results demonstrate the performance of the proposed DSS strategy in terms of both extending network lifetime and the effect DSS has on the detection and localization strategy. For these simulations, we used the same parameters from the second simulation study in the previous section.

We apply the DSS strategy to the detection and localization strategy and simulated with the parameters:

- Base station located at position $[0, 0]$
- $\bar{\omega}_i = (0.9)^{\pi(i)}$
- $\mathcal{C} = 0.2$
- $J = 15$

where $\pi(i)$ is the number of hops required for the i^{th} sensor to transmit to the base station. Additionally, we assume the base station has sufficient power to communicate to all sensors in one hop.




Figure 7. Number of sensors selected and percentage of energy used at each time step.

Figure 7 shows the number of sensors selected using DSS and the energy used to gather the observations versus the time step. The DSS strategy weights the cost of energy required to gather data against the improvement in estimation accuracy. For $\mathcal{C} = 0.2$, Fig. 7 shows that the

average number of sensors selected is 13 sensors (out of a possible 100 sensors). The energy savings when DSS is employed as opposed to full-order sensing is significant. In this simulation, the DSS approach to detection and localization used approximately 5% of the energy as the full-order approach (as quantified by the energy function). In this example, the network lifetime can be extended almost 20 times by using DSS instead of full - order sensing.

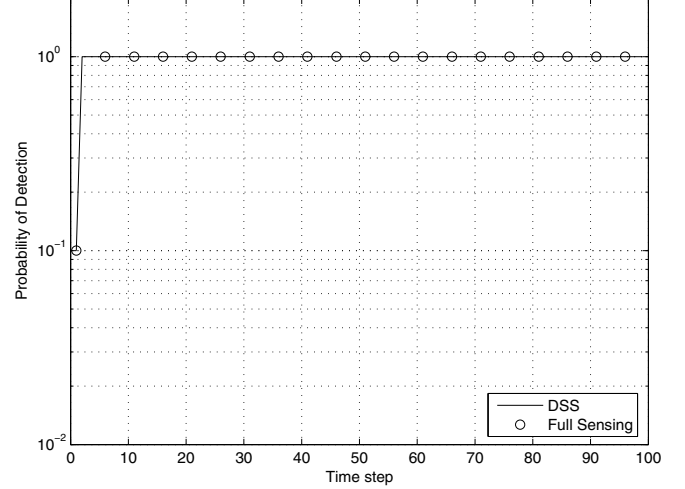


Figure 8. Probability of detection (P_D) vs. time step.

Figure 8 shows the probability of detection (P_D) versus the time step. These results show that the probability of deciding an active source is present when a source is active is not significantly affected by using DSS. This performance was achieved by increasing the window size when the DSS approach is used. In the simulations, when full-order sensing performed, the window size was 5 observation periods ($J=5$); however, when DSS is applied the window size can increase since each observation period contains fewer observations. Since the optimal detector scales with both time and the number of observations, if the number of observations per period is decreased, the number of periods can increase. Additionally, due to the stable plant dynamics, the observations under the DSS strategy are not as correlated as the observations under full-order sensing, since the observations under the DSS strategy are gathered over a longer period of time. This reduced correlation also serves to improve the detector's performance.

The test statistics for the full-order sensing and DSS strategies are shown in Fig. 9. In this figure, when the statistic is positive, we declare a source is active and similarly, we declare no active sources when the statistic is negative. The results show that it is possible for the DSS strategy to outperform the full-order sensing strategy and claim a detection earlier. This non-intuitive effect is caused by the DSS strategy using a larger observation set window that provides a similar number of observations as the full-




Figure 9. Normalized test statistic vs. time step.

order sensing. Additionally, since the observations are less correlated, fewer observations of leaks are needed to claim a leak. In this simulation, the DSS strategy selected sensor groups whose observations indicated a leak was present. This selection of sensors coupled with the lower correlated observations resulted in the DSS strategy detecting a leak before the full-order strategy. In general, Fig. 9 indicates that the full order sensing statistic increases at a faster rate than the DSS statistic, however there are outlier cases where the DSS strategy can outperform the full-order strategy.

VI. IMPLEMENTATION

We have developed a testbed of 30 wireless sensors, to implement and evaluate the proposed detection and localization strategy. To fully analyze the performance of the detection strategy, a back-channel monitoring system exists to gather all the sensor measurements. This system allows a comparison between the WSN detection capabilities versus a standard wired network. A 3.6 GHz machine running Windows XP is used as a central processor and the detection and localization strategy as well as the DSS strategy are written for MATLAB. The windows machine CPU clock is used to determine when a request is sent to the WSN gateway for data acquisition. To send a request for data, MATLAB communicates via a socket to a Linux machine attached to the gateway of the WSN. The WSN consists of Firefly nodes [12] and a stable networking protocol called SAMPL [16] has been implemented using the Nano-rk real-time operating system [4]. The process of gathering observations and performing source detection and localization must be timed such that the sensors take measurements at the same instant. This requires time synchronization between nodes and the information processing algorithms to occur at specific times, as outlined by the timing diagram in Fig. 10. Time synchronization for sampling is achieved for by

delaying sensor measurement based on the routing tree depth provided by SAMPL, where the larger the depth, the shorter the delay. Exact time synchronization is not essential since the sampling rate is very large when compared to the error in sampling times. The remaining part of this section uses Fig. 10 to illustrate key points about data gathering and central processing.

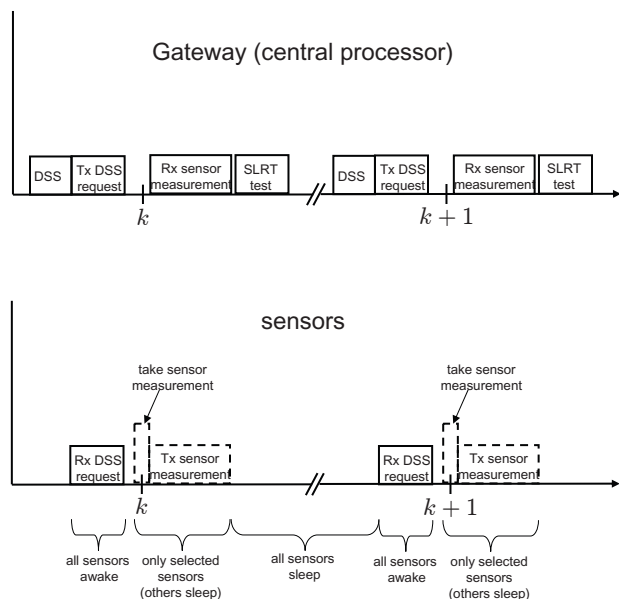


Figure 10. Timing diagram.

Before the sampling instant at time k , the central processor transmits which sensors should be collected. During this period, all sensors turn on their receivers and listen for the central processor broadcast. Once the broadcast is received, the sensor determines whether it will be used for either taking a measurement or routing of the measurements back to the central processor. If the sensor is needed, the measurement is taken and information sent back to the central processor. If the sensor is not needed, the sensor turns off the radio and goes into a sleep mode until the next time it needs to listen for a DSS request. After the observations are gathered at the central processor, the detection and localization strategy is performed for the system at time k . Then just before the DSS request is transmitted to the nodes, the central processor performs DSS. The reason the central processor waits until the last moment to perform DSS is because the DSS process requires knowledge of the system dynamics evolving the states from time k to time $k+1$. These dynamics describe real world phenomena which are known to be time varying. By waiting until the last possible moment to perform DSS, we are able to monitor the dynamics for the majority of the time period over which they govern. This leads directly to better models of the dynamics relating the system state at time k to the system state at time $k+1$.

VII. DISCUSSION AND FUTURE WORK

This paper presents a new approach to the real-time detection of multiple sources in dynamic fields governed by partial differential equations. Kalman filters provide the mean and covariance data for a binary hypothesis test. When active sources are detected, the Kalman filter estimates of the source values identify the locations of the active sources. A two-tiered approach reduces the number of sensors required to perform fine-grained detection and localization which serves to extend the network lifetime. To further extend the network lifetime a previously introduced DSS strategy is improved to consider both channel reliability and energy consumption. A real-time implementation of the detection and localization strategy is presented and timing constraints outlined. Simulation results demonstrate the effectiveness of the proposed method for wide-area detection and localization using wireless sensor networks.

Future research directions on dynamic sensor selection involve improved modeling of energy consumption and channel reliability based on routing and MAC protocols, and dynamically determining the energy cost to achieve a desired performance. Directions for future research in the realm of detection and localization include the use of non-uniform sensor deployment, and systematic procedures for choosing parameters in the detection scheme to achieve desired levels of performance.

VIII. ACKNOWLEDGMENTS

This work was performed in support of ongoing research in sensor systems and diagnostics at the National Energy Technology Laboratory under RDS contract DE-AC26-04NT41817

REFERENCES

- [1] B. Brumback and M. Srinath. A chi-square test for fault-detection in kalman filters. *IEEE Transactions on Automatic Control*, June 1987.
- [2] Ashim K. Datta. Biological and bioenvironmental heat and mass transfer. marcel dekker, inc., 2002.
- [3] Michael A. Demetriou. Process estimation and moving source detection in 2-d diffusion processes by scheduling of sensor networks. *American Control Conference, 2007. ACC '07*, pages 3432–3437, 9-13 July 2007.
- [4] A. Eswaran, A. Rowe, and R. Rajkumar. Nano-rk: an energy-aware resource-centric rtos for sensor networks. In *Real-Time Systems Symposium, 2005. RTSS 2005. 26th IEEE International*, pages 10 pp.–265, 2005.
- [5] H. S. Nance et al. Surface environmental monitoring at the Frio CO₂ sequestration test site. In *Fourth Annual Conference On Carbon Capture and Sequestration DOE/NETL*, Alexandria, Virginia, May 2005.
- [6] E. B. Fox, J. W. Fisher, and A. S. Willsky. Detection and localization of material releases with sparse sensor configurations. *Signal Processing, IEEE Transactions on*, 55(5):1886–1898, May 2007.
- [7] V. Gupta, T. Chung, H. Babak, and R. Murray. On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage. *Automatica*, 42:251–260, Feb 2006.
- [8] Yinlun Huang. On distributed fault-tolerant detection in wireless sensor networks. *IEEE Trans. Comput.*, 55(1):58–70, 2006. Student Member-Xuanwen Luo and Member-Ming Dong.
- [9] T. Kerr. Statistical analysis of a two-ellipsoid overlap test for real-time failure detection. *IEEE Transactions on Automatic Control*, August 1980.
- [10] M. Kvasnica, P. Grieder, and M. Baotić. Multi-Parametric Toolbox (MPT), 2004.
- [11] K. S. Lackner. Climate change: A guide to CO₂ sequestration. *Science*, 300(5626):1677–1678, June 2003.
- [12] Rahul Mangharam, Anthony Rowe, and Raj Rajkumar. Firefly: a cross-layer platform for real-time embedded wireless networks. *Real-Time Systems*, 37(3):183–231, December 2007.
- [13] M.P. Michaelides and C.G. Panayiotou. Event detection using sensor networks. *Decision and Control, 2006 45th IEEE Conference on*, pages 6784–6789, 13-15 Dec. 2006.
- [14] G. T. Nofsinger and G. V. Cybenko. Airborne plume tracking with sensor networks. In *Proceedings of the SPIE, Volume 6231, pp. 623112 (2006)*, volume 6231 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, June 2006.
- [15] C. Phillips and H. Nagle. *Digital Control System Analysis and Design*. Prentice Hall, 1995.
- [16] A. Rowe, K. Lakshmanan, and R. Rajkumar. Sampl: A simple aggregation and message passing layer. In *Real-Time Systems Symposium*, Nov 2008.
- [17] Youngchul Sung, Lang Tong, and H. Vincent Poor. Sensor configuration and activation for field detection in large sensor arrays. In *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*, page 22, Piscataway, NJ, USA, 2005. IEEE Press.
- [18] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory*. John Wiley & Sons, Inc., New York, 1968.
- [19] J.E. Weimer, B. Sinopoli, and B.H. Krogh. A relaxation approach to dynamic sensor selection in large-scale wireless networks. *Distributed Computing Systems Workshops, 2008. ICDCS '08. 28th International Conference on*, pages 501–506, June 2008.
- [20] A. Zolghadri. An algorithm for real-time failure detection in kalman filters. *IEEE Transactions on Automatic Control*, October 1996.