Localizing Vapor-Emitting Sources by Moving Sensors*

Boaz Porat

Department of Electrical Engineering
Technion—Israel Institute of Technology
Haifa 32000, Israel

Arye Nehorai

Department of Electrical Engineering
Yale University
New Haven, CT 06520

Abstract

In a recent paper, the authors have explored the use of novel concentration sensors for detecting and localizing vapor-emitting sources. As was shown there, an array of sensors (at least five) is needed for this purpose in general. In the present paper we propose to replace stationary sensors by moving sensors, thus gaining two important advantages: 1) A single moving sensor can accomplish the task of an array of stationary sensors, by exploiting spatial and temporal diversity. 2) The sensor motion can be planned in real time to optimize localization performance, based on past measurements and minimization of the expected localization error as a function of the future sensor's position. The paper describes the details of this approach and illustrates it by an example.

1 Introduction

Recently, there has been considerable progress in the development of sensors for low concentration vapors. Sensors of this kind are potentially useful in applications such as explosive detection, drug detection, sensing leakage of hazardous chemicals, pollution sensing and environmental studies.

In a previous paper we have examined the problem of detecting and localizing vapor-emitting sources by stationary sensor arrays [1]. Based on the diffusion equation, we developed models for spatial distributions of vapors in various environments, and used them to develop measurement models. We then derived detection and estimation algorithms, analyzed their performance and performed some simulation studies.

The present paper looks at the problem of vapor source localization when the sensors are permitted to move. Such a sensor can achieve spatial diversity by taking measurements at different locations at different times. In principle, a single sensor can thus perform tasks that would need an array of stationary sensors (for example, a minimum of five stationary sensors was shown in [1] to be necessary for localization in three dimensions). As a possible scenario, consider a vehicle, equipped with vapor sensors, moving in an open area where hazardous materials may be present, and whose task is to detect and localize such materials. The vehicle may be driven by a protected person or be remote controlled. Assume that the vehicle's path can be selected in real time, subject to speed and maneuvering limitations, and possible geometric constraints (e.g., obstacles). The goal is to use the freedom in the vehicle's path to optimize localization under a properly defined criterion. In a more general setting, we can consider several vehicles performing a coordinated search. One can liken this scenario to a search by trained dogs, since dogs use their sensing ability and freedom of motion in a similar manner. The great difference is, of course, that dogs function individually and do not share their respective information, while an automated system of the type we propose optimizes performance based on all available information.

The optimality criterion used in this paper for planning the sensors' motion is to reduce as much as possible the expected location estimation error after the next measurement is taken. This is accomplished by computing the gradients of the Cramér-Rao bound (CRB) on the location error with respect to the sensors' coordinates, and moving the sensors opposite to the corresponding gradient directions. Since the location estimation algorithm we use is maximum likelihood, its actual accuracy is expected to be close to the CRB, so by minimizing the CRB we can hope to minimize the actual estimation error. We have named this algorithm the "DOG algorithm", an acronym for *direction of gradient.*

*This work was supported by the Air Force Office of Scientific Research under Grant no. F49620-93-1-0096, the National Science Foundation under Grant no. MIP-9122753, and the Office of Naval Research under Grant no. N00014-91-J-1298.

The paper describes the mathematical details of the proposed localization method, and illustrates it by an example.

2 Mathematical Modeling

In this section we briefly describe the physical and measurement models. Here we consider only the simplest case: that of a point source in an infinite volume, emitting vapors at a constant rate. More general models are developed in [1].

Let $c(\vec{r},t)$ be the diffusing substance concentration at a point $\vec{r}=(x,y,z)$ and time t, in Kg/m^3. Let $\vec{f}(\vec{r},t)$ be the flux density vector, in units of Kg/m^{2}s. Fick's law of diffusion asserts that the latter is proportional to the gradient of the former, that is,

$$\vec{f}=-\kappa\nabla c, \tag{2.1}$$

where κ is the *diffusivity*, in units of m^2/s. By the continuity equation, we have for a source-free volume,

$$\frac{\partial c}{\partial t}=-\nabla\cdot\vec{f}. \tag{2.2}$$

Combining (2.1) and (2.2) we get, for a source-free volume and space-invariant diffusivity, the well-known diffusion equation

$$\frac{\partial c}{\partial t}=\kappa\nabla^2 c. \tag{2.3}$$

Consider a point source at $\vec{\rho}=[\chi,\upsilon,\zeta]$ in an infinite medium, releasing a diffusing substance at a constant rate of μ Kg/s, starting at time t_0. The solution of (2.3) is then given by

$$c(\vec{r},t)=\frac{\mu}{4\pi\kappa|\vec{r}-\vec{\rho}|}\mathrm{erfc}\left(\frac{|\vec{r}-\vec{\rho}|}{2\sqrt{\kappa(t-t_0)}}\right), \tag{2.4}$$

where $\mathrm{erfc}(x)=(2/\sqrt{\pi})\int_x^\infty e^{-y^2}dy$ is the complementary error function.

We now discuss the measurement model. Chemical sensors are designed to be as selective as possible, that is, their sensitivity to the substance of interest should be as high as possible, while their sensitivity to foreign (undesired) substances should be as low as possible. Nonetheless, no sensor can be completely insensitive to foreign materials. Therefore, it is reasonable to model the response of a sensor located at point $\vec{r}$ by

$$y(\vec{r},t)=c(\vec{r},t)+b(\vec{r},t)+e(\vec{r},t), \tag{2.5}$$

where

- $c(\vec{r},t)$ is the concentration of the substance of interest, as given by (2.4).
- $b(\vec{r},t)$ is a "bias" (or "clutter") term, representing the sensor's response to foreign substances that may be present.
- $e(\vec{r},t)$ is the sensor's noise.

The source location $\vec{\rho}$, its intensity μ, the diffusivity κ, and the initial time t_0, are assumed unknown.

There are assumed to be m sensors, which are allowed to be in motion. Their locations $\{\vec{r}_i(t), 1\le i\le m\}$ are assumed to be known at all times. The bias term is assumed to be time-invariant and uniformly distributed in space. In other words, we are assuming that substances interfering with the desired measurement are in a steady state and are the result of many sources, so they have reached both temporal and spatial equilibrium. The corresponding bias intensity b is assumed to be unknown. Finally, the noise is assumed to be independent from sensor to sensor, uncorrelated in time, and Gaussian distributed with zero mean and unknown variance σ_e^2. We remark that the Gaussian assumption is not quite physical, since in practice $y(\vec{r},t)$ should be a nonnegative quantity. However, for relatively small σ_e (that is, much smaller than $c(\vec{r},t)+b(\vec{r},t)$), the Gaussian model is a reasonable approximation. In summary, our measurement model for the ith sensor is

$$y(\vec{r}_i(t),t)=\frac{\mu}{4\pi\kappa|\vec{r}_i(t)-\vec{\rho}|}\mathrm{erfc}\left(\frac{|\vec{r}_i(t)-\vec{\rho}|}{2\sqrt{\kappa(t-t_0)}}\right)$$
$$+b+e(\vec{r}_i(t),t),\quad e(\vec{r}_i(t),t)\sim\mathcal{N}(0,\sigma_e^2). \tag{2.6}$$

3 Parameter Estimation and the CRB

Before we describe the moving-sensor algorithm, we briefly recall the maximum likelihood estimation approach and the Cramér-Rao lower bound presented in detail in [1]. We assume that the available measurements are $\{y(\vec{r}_i,t_k), 1\le i\le m, 1\le k\le p\}$, where $p\ge 2$.

Partition the parameter vector as $[\boldsymbol{\theta}^T,\boldsymbol{x}^T,\sigma_e^2]^T$, where $\boldsymbol{\theta}=[\chi,\upsilon,\zeta,\kappa,t_0]^T$ and $\boldsymbol{x}=[\mu,b]^T$. With this notation we can lump the measurements in the vector form

$$\boldsymbol{y}=A(\boldsymbol{\theta})\boldsymbol{x}+\boldsymbol{e}, \tag{3.1}$$

where $\boldsymbol{y}$ is an (mp)-dimensional vector whose $(m(k-1)+i)$-th component is $y(\vec{r}_i,t_k)$ and similarly for $\boldsymbol{e}$, while $A(\boldsymbol{\theta})$ is a $mp\times 2$ matrix whose $(m(k-1)+i)$-th row is given by

$$[(1/4\pi\kappa|\vec{r}_i-\vec{\rho}|)\mathrm{erfc}(|\vec{r}_i(t_k)-\vec{\rho}|/2\sqrt{\kappa(t_k-t_0)}),\ 1]. \tag{3.2}$$

The maximum likelihood estimate of the parameters is given by

$$\hat{\theta} = \text{argmax}\{y^T P_A(\theta) y\} \quad (3.3)$$

$$\hat{x} = [A^T(\hat{\theta})A(\hat{\theta})]^{-1} A^T(\hat{\theta}) y \quad (3.4)$$

$$\hat{\sigma}_e^2 = (mp)^{-1} y^T P_A^{\perp}(\hat{\theta}) y, \quad (3.5)$$

where $P_A(\theta)$ is the projection matrix on the column space of $A(\theta)$, and $P_A^{\perp}(\theta)$ is the complementary projection matrix, given by

$$P_A(\theta) = A(\theta)[A^T(\theta)A(\theta)]^{-1} A^T(\theta),$$
$$P_A^{\perp}(\theta) = I - P_A(\theta). \quad (3.6)$$

Writing $A(\theta)$ in terms of its two columns, say $[a(\theta), u]$ (where u is a vector of 1's), we can express (3.3) as (omitting the dependence on θ for convenience)

$$\hat{\theta} = \text{argmax} \left\{ \frac{[(y^T a) - (mp)^{-1}(u^T y)(u^T a)]^2}{(a^T a) - (mp)^{-1}(u^T a)^2} \right\}. \quad (3.7)$$

Define

$$\tilde{y} = y - (mp)^{-1}(u^T y)u, \quad \tilde{a} = a - (mp)^{-1}(u^T a)u. \quad (3.8)$$

Using these definitions, we can express $\hat{\theta}$ as

$$\hat{\theta} = \text{argmax} \left\{ \frac{(\tilde{y}^T \tilde{a})^2}{(\tilde{a}^T \tilde{a})} \right\} = \text{argmax} \left\{ \frac{(\tilde{y}^T \tilde{a})^2}{(\tilde{y}^T \tilde{y})(\tilde{a}^T \tilde{a})} \right\}. \quad (3.9)$$

As we see, $\hat{\theta}$ is obtained as the maximizer of the correlation coefficient between $\tilde{y}$ and $\tilde{a}$.

The Cramér-Rao bound (CRB) for unbiased estimates of the parameters is given by [2]

$$\text{CRB}(\theta) = \frac{\sigma_e^2}{\mu^2} \left\{ D^T(\theta) P_A^{\perp}(\theta) D(\theta) \right\}^{-1} \quad (3.10)$$

$$\text{CRB}(x) = \sigma_e^2 \left\{ A^T(\theta)[I - \mu^2 P_D(\theta)]A(\theta) \right\}^{-1} \quad (3.11)$$

$$\text{CRB}(\sigma_e^2) = \frac{2\sigma_e^4}{mp} \quad (3.12)$$

where $D(\theta) = \partial a(\theta)/\partial\theta$, an $(mp \times 5)$-dimensional matrix.

4 The Moving-Sensor Algorithm

We examine the localization problem in more detail (assuming that detection has already been performed), and suppose that there is a single moving sensor. In practice the sensor will move on a surface (that is, on the ground), but for the sake of generality we will retain three-dimensional notations. Vapor measurements cannot be taken continuously, due to technological limitations, so we assume that there is a minimum interval ΔT between consecutive measurements. The time needed for N measurements is therefore $N\Delta T$. Let $\Delta\vec{\rho} = [\Delta\chi, \Delta v, \Delta\zeta]$ denote the vector of location estimation errors. A reasonable optimality criterion is to estimate the source location within a prescribed accuracy ϵ at a minimum time, that is, to minimize N subject to the requirement $\|\Delta\vec{\rho}\| \leq \epsilon$. The underlying assumption is, of course, that ϵ represents the error below which the source can be identified by other means (e.g., visually). Implementation of such a criterion is, however, very difficult, since the corresponding mathematical problem cannot be solved in closed form. We therefore propose the following approximation.

Suppose we have already taken n measurements, at known sensor locations and instants. Thus we can obtain an estimate $[\hat{\chi}_n, \hat{v}_n, \hat{\zeta}_n]$ of the source coordinates, using the algorithm described in the previous section. The accuracy of $[\hat{\chi}_n, \hat{v}_n, \hat{\zeta}_n]$ (i.e., their standard deviations) can be estimated from the Cramér-Rao lower bound, if we assume that the algorithm is statistically efficient (or a certain small multiple of the bound, if we allow for a certain relative efficiency). For the unknown values of $[\chi, v, \zeta]$ (which are needed in the CRB computation), we substitute the current estimates $[\hat{\chi}_n, \hat{v}_n, \hat{\zeta}_n]$. Let $\vec{r}(n\Delta T) = [x(n\Delta T), y(n\Delta T), z(n\Delta T)]$ be the coordinates of the sensor after the nth measurement, and let $\mathcal{S}_{n+1}$ be the set of all points reachable by the sensor at time $(n+1)\Delta T$, given the speed and maneuvering limitations and the geometrical constraints. Then choose a point

$$\vec{r}((n+1)\Delta T) =$$
$$[x((n+1)\Delta T), y((n+1)\Delta T), z((n+1)\Delta T)] \quad (4.1)$$

in $\mathcal{S}_{n+1}$ so as to *minimize the CRB on the location errors given the measurements up to and including* $n+1$, and move the sensor to that point in order to take the next measurement.

The rationale behind the above tactic is clear: The information in the next measurement depends on the sensor's location while taking that measurement. Of all feasible locations, it makes sense to choose the one that would improve the estimation accuracy as much as possible. If this tactic is repeated at each step, we can hope to minimize the time required for reducing the estimation error to a prescribed level.

A typical shape of the set $\mathcal{S}_{n+1}$ is a sphere centered at $\vec{r}(n\Delta T)$, whose radius is related to the maximum speed, acceleration and deceleration (we assume that

the vehicle has to stop in order to take the next measurement). Physical obstacles are to be excluded from this disk. The next point $\vec{r}((n+1)\Delta T)$ can then be found by searching over a grid on this sphere. At each grid point, the CRB can be computed by the formula (3.10). This, however, may require an unreasonably large number of computations. We therefore propose an approximation: Compute the *gradient* of the CRB after $n+1$ measurements with respect to the prospective new point $[x((n+1)\Delta T), y((n+1)\Delta T), z((n+1)\Delta T)]$. In this computation, all previous points have already been determined, so they are fixed; also, the source location $\vec{\rho}$ is replaced by its current estimate. Then proceed in the opposite direction of the gradient to the boundary of the sphere, or to the furthest feasible point along this direction, and declare this point as $\vec{r}((n+1)\Delta T)$. This procedure is terminated when the CRB becomes smaller than ϵ (or a small multiple thereof). We name this suboptimal approach the DOG algorithm (for *Direction Of Gradient*).

The gradient of the CRB is computed as follows. Let ξ denote any variable with respect to which we want to differentiate [in this case one of the coordinates of $\vec{r}((n+1)\Delta T)$]. Then we get from (3.10)

$$\frac{\partial}{\partial \xi}\mathrm{CRB}(\boldsymbol{\theta}) = \frac{\mu^2}{\sigma_e^2}\mathrm{CRB}(\boldsymbol{\theta})\left[\frac{\partial}{\partial \xi}\left\{D^T(\boldsymbol{\theta})P_A^{\perp}(\boldsymbol{\theta})D(\boldsymbol{\theta})\right\}\right]\mathrm{CRB}(\boldsymbol{\theta}), \quad (4.2)$$

where

$$\begin{aligned}\frac{\partial}{\partial \xi}\left\{D^T(\boldsymbol{\theta})P_A^{\perp}(\boldsymbol{\theta})D(\boldsymbol{\theta})\right\} = \\ \frac{\partial D^T(\boldsymbol{\theta})}{\partial \xi}P_A^{\perp}(\boldsymbol{\theta})D(\boldsymbol{\theta}) + D^T(\boldsymbol{\theta})P_A^{\perp}(\boldsymbol{\theta})\frac{\partial D(\boldsymbol{\theta})}{\partial \xi} \\ -D^T(\boldsymbol{\theta})P_A^{\perp}(\boldsymbol{\theta})\frac{\partial A(\boldsymbol{\theta})}{\partial \xi}[A^T(\boldsymbol{\theta})A(\boldsymbol{\theta})]^{-1}A^T(\boldsymbol{\theta})D(\boldsymbol{\theta}) \\ -D^T(\boldsymbol{\theta})A(\boldsymbol{\theta})[A^T(\boldsymbol{\theta})A(\boldsymbol{\theta})]^{-1}\frac{\partial A^T(\boldsymbol{\theta})}{\partial \xi}P_A^{\perp}(\boldsymbol{\theta})D(\boldsymbol{\theta}).\end{aligned} \quad (4.3)$$

We will omit the expressions for the individual derivatives of $D(\boldsymbol{\theta})$ and $A(\boldsymbol{\theta})$ since they are lengthy, but in principle straightforward to obtain from the explicit expressions (3.2) for the elements of $A(\boldsymbol{\theta})$. In our computer implementation of the algorithm, we obtained those derivatives using the symbolic mathematics capabilities of *Mathematica*.

Note that, in order to determine the direction of motion of the sensor, we only need a small part of the CRB matrix, namely: the diagonal elements corresponding to the three location coordinates $[\chi_n, v_n, \zeta_n]$. There are many ways to combine these elements to obtain a scalar criterion. Since in general there is no reason to prefer one coordinate over the other, we have chosen to use the sum of the three elements. This sum is a measure of the square-magnitude of the location error vector.

An extension to the case of multiple moving sensors (a "Multi-DOG" algorithm) can be suggested as follows. All sensors transmit their locations and measurements to a central unit. Based on the DOG criterion, the central unit computes the next optimal location for each sensor and transmits these locations to the vehicles, which then proceed to their respective destinations for the next measurement. The implementation details are a straightforward extension of the basic DOG algorithm, and we will not elaborate on them here.

5 An Example

The following simulation example illustrates the behavior of the DOG algorithm. A vapor-emitting source is placed at coordinates $(-50, -50)$ meter. Emission is assumed to have started 100 seconds prior to the beginning of the simulation. The source emission-rate μ is normalized to 1 Kg/s. The diffusivity κ is taken as 25 m^2/s. Bias of 10^{-4} $\mathrm{Kg/m}^3$ and random error with standard deviation 10^{-6} $\mathrm{Kg/m}^3$ are assumed to be present in the measurements. The sensor moves on a circle of radius 250 m, centered at $(0,0)$, counterclockwise, starting at time zero. The vehicle's speed is equal to 5 m/s. Measurements are taken every 10 s.

At time $t = 400$ s, detection is assumed to have occurred, and then the DOG algorithm begins, initially using the measurements accumulated during the first 400 seconds. After 19 more measurements (that is, 190 s), the CRB's of the ξ and v coordinates both get below 1 m, and then the simulation stops.

Figure 1 shows the path travelled by the sensor, including the initial circle and the convergence path toward the source. Note how the moving sensor "sniffs around" the source, circling it while gradually decreasing the distance. This behavior is very intuitive: the circling serves the purpose of spatial diversity of the measurements, while decreasing the distance helps increasing the measured concentration, thereby reducing the sensitivity to noise.

Figure 2 shows the CRBs on the two coordinates (rather, the square-roots thereof, in meters) as a function of time, while Figure 3 shows the corresponding location errors. The CRBs decay monotonically, as dictated by the algorithm. The actual errors decay

as well, but not monotonically, since this is a single random simulation.

6 Conclusions

We have examined the possible use of moving sensors for localizing vapor-emitting sources. A single moving sensor can fulfill the task of an array of stationary sensors and, furthermore, the sensor's path can be planned in real time so as to optimize localization according to a specified criterion. The optimality criterion proposed in this paper, motion opposite to the direction of the gradient of the CRB (with respect to the sensor's coordinates), approximately minimizes the time to reduce the location error to a prescribed level. The CRB-gradient approach is applicable in principle to other areas. It can be used whenever the sensors are movable and their motion times are on the same time scale as the measurement process (or faster).

References

[1] A. Nehorai, B. Porat and E. Paldi, "Detection and Localization of Vapor-Emitting Sources", *IEEE Trans. Signal Processing*, to appear, January 1995.

[2] P. Stoica and A. Nehorai, "MUSIC, Maximum Likelihood, and Cramér-Rao Bound," *IEEE Trans. Acoustic, Speech, Signal Processing*, 37, No. 5, pp. 720–741, 1989.

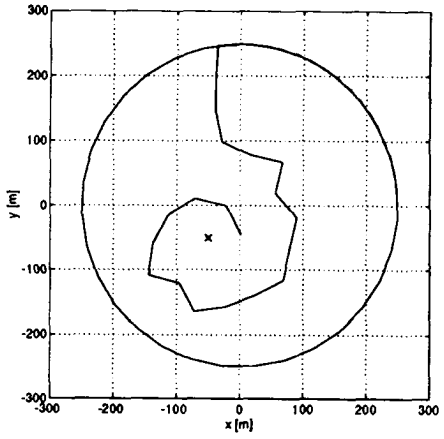

Figure 1. The sensor's path; source is at (-50, -50).

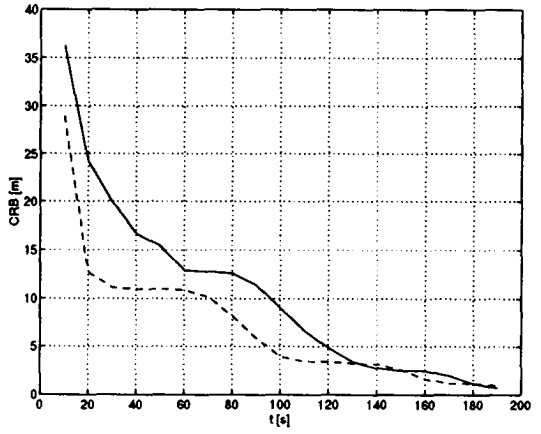

Figure 2. The Cramér-Rao bounds on the source coordinates as a function of time; solid line: x, dashed line: y.

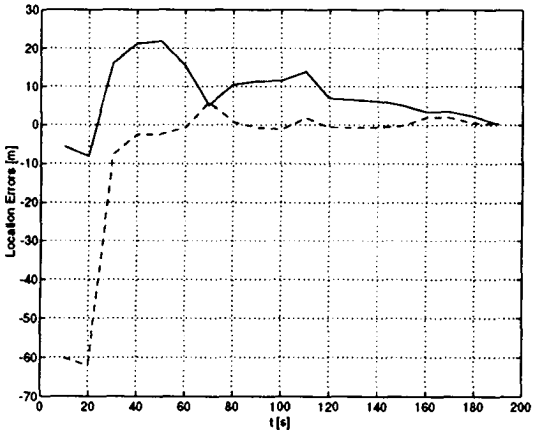

Figure 3. The location errors as a function of time; solid line: x, dashed line: y.