

Exploiting A Constellation of Narrowband RF Sensors to Detect and Track Moving Targets

Chris Kreucher^a, J. Webster Stayman^b, Ben Shapo^a, and Mark Stuff^c

^aIntegrity Applications Incorporated
900 Victors Way, Suite 220
Ann Arbor, MI 48108
ckreucher,bshapo@integrity-apps.com

^bJohns Hopkins University
Traylor Building, Room #726
720 Rutland Ave, Baltimore, MD 21205
web.stayman@jhu.edu

^cMichigan Tech Research Institute
3600 Green Ct., Ste. 100,
Ann Arbor, MI 48105
mastuff@mtu.edu

Abstract¹ — This paper presents a novel approach to detecting and tracking moving targets using a constellation of narrowband radio frequency (RF) sensors. Our methodology is an innovative combination of nonlinear estimation and information theoretic sensor placement. The nonlinear filtering approach fully exploits bistatic Doppler measurements made by the sensors without thresholding or linear/Gaussian assumptions, thereby improving the detection/false alarm tradeoff and lowering tracking error. The information theoretic sensor placement algorithm, which is based on minimizing the Cramer-Rao bound (CRB) on localization variance, selects sensor positions that lead to the best estimation performance, thereby maximally exploiting the finite sensing resources. We illustrate the efficacy of the algorithm by showing how well we can track a target using sensors placed using the CRB optimal method as compared to sensors placed randomly.

I. INTRODUCTION

This paper describes an innovative method for detecting and tracking moving targets by exploiting a constellation of inexpensive narrowband (NB) Radio Frequency (RF) sensors. In our model problem, we consider a road with on the order of a half-dozen sensors placed around the road. A processing center receives and processes the bistatic radar data. It is the goal of the signal processing approach to exploit these multistatic returns for the effective detection and tracking of vehicles operating in the surveillance region.

A constellation of narrowband sensors has a number of benefits over conventional wideband sensors in this application. First, commercial applications have led to an erosion of the available spectrum meaning often only a small portion is available for other use [1][2]. Furthermore, NB sensors are inexpensive due to their simple electronics, require low energy consumption, are easy to maintain, and it is easy to communicate their data to a centralized processing point. But perhaps most importantly, a constellation of NB sensors provide geometric diversity. By exploiting bistatic returns with advanced signal processing techniques like those described here, this trades costly spectral diversity for cost-efficient spatial diversity, while providing performance improvement.

The paper contains two main contributions. First, we give a Bayes optimal nonlinear filtering method of target tracking which admits the highly nonlinear and non-

Gaussian measurements made by the sensors. Standard approaches address the problem in two sequential phases: detection and tracking. In particular, one algorithm is responsible for generating threshold exceedances (detections) at each epoch (often an FFT block, or “snapshot”) and then these detections are passed to a separate (typically Kalman-based) tracking algorithm. In contrast, our approach performs track-before-detect, i.e., there is no thresholding (or other hard decision) on the received measurements and all data is incorporated softly into a probabilistic estimate. Furthermore, our method directly accommodates the non-Gaussian and non-linear sensor to target couplings that characterize this scenario. Given these distinctions, the method yields improved performance in both false alarm/detection tradeoffs and tracking error.

Second, we describe an information theoretic method for choosing where to place the individual sensors in the constellation. Our approach is based on selecting sensor locations to minimize the Cramer-Rao bound (CRB) on estimator variance. This method automatically favors sensor placements that yield beneficial (bistatic) spatial diversity. As such, it optimizes the ability to detect and track moving targets. The efficacy of the sensor placement is demonstrated by comparing tracking performance with sensors randomly placed versus those placed using the CRB optimal sensor placement algorithm.

The paper proceeds as follows. In Section II, we introduce the bistatic Doppler signal model. In Section III, we show how this model is combined with a nonlinear filtering algorithm to provide a tractable Bayes optimal tracking approach. In Section IV, we describe a Cramer-Rao bound based sensor placement algorithm for optimally choosing sensor locations and illustrate its performance. Finally, section V concludes.

II. THE SENSOR MODEL

We assume a constellation of N sensors. To simplify the notation, we assume each sensor is both a transmitter and a receiver, although this assumption is not required by our approach. A sensor transmits a narrowband RF signal, which is reflected off the target and received at each sensor. Then this process repeats $N - 1$ times, with each sensor serving as transmitter. There are then N^2 bistatic pairs, which process the received signal to generate information about bistatic range *rate* from the observed shift in frequency. Note, we assume purely narrowband sensors so

¹ This work was supported by Air Force Research Lab contracts FA8650-09-M-1549 and FA8650-10-C-1718.

no range information is available. We could expand our model to include sensors which provide range information at the cost of more expensive hardware and increased computation.

Given a CPI worth of pulses has been transmitted by sensor A and received by all sensors, Fourier processing yields Doppler compression – under the assumption the target moves at a constant velocity for this short period of time. This slow-time data surface can be used to evaluate the evidence for hypotheses about target bistatic Doppler, which is the instantaneous rate of change of target bistatic range. The measurements made by all N sensors from the energy transmitted by sensor A are corrupted by various types of noise, including clutter, measurement noise, and quantization noise.

We employ the following statistical model to describe the measurements. First, let z_i denote the magnitude measurement in the i^{th} bistatic Doppler resolution cell between sensors A and B . Resolution and the number of cells are determined by the number of pulses and the PRF. The collection of measurements is then the vector of bistatic Doppler correlations in each cell, i.e., the vector

$$z = [z_1 \ z_2 \ \cdots \ z_{L-1} \ z_L] \quad (1)$$

where L is the number of bistatic Doppler cells and the statistics of the complex measurement envelope in each cell depend on whether that cell corresponds to a target-absent cell or a target-present cell.

Let the vector $\theta = [x \ \dot{x} \ y \ \dot{y}]^T$ describe the true 2D position and velocity of the target. We model the statistics in each bistatic Doppler cell as Rayleigh. To be concrete, for each bistatic Doppler cell i

$$f_i(z; \theta) \equiv p(z_i | \theta) = \frac{z_i}{\lambda_i^2(\theta)} \exp\left(-\frac{z_i^2}{\lambda_i^2(\theta)}\right), \quad (2)$$

where λ_i is the pixel-dependent Rayleigh parameter. In the point-target pixilated model, λ_i is

$$\lambda_i(\theta) = \lambda_b + (\lambda_t - \lambda_b)I_i(\theta), \quad (3)$$

where the indicator variable $I_i(\theta)$ is 1 if θ maps into cell i . I is also a function of transmitter and receiver locations, but this is suppressed in the notation for clarity. Therefore, the target containing cell has energy returns which distributed Rayleigh with λ_t and the other (background) cells have energy returns which are distributed Rayleigh with λ_b .

III. NONLINEAR FILTERING

This section describes the track-before-detect approach to joint target detection and localization. The goal is to construct a probability density on target state given the collection of bistatic Doppler measurements taken over time for the constellation of sensors.

Our approach differs from traditional methods in several manners. First, traditional methods operate in two-stages: first performing detection on incoming

measurements and then passing these threshold exceedances to a tracker. This requires measurement thresholding, which is an SNR degrading step. Second, traditional methods use linear (Kalman) filters that sub-optimally model measurements as Gaussian and having linear sensor-to-target coupling. In contrast, our approach does not threshold the measurements but instead uses the raw energy returns, treats the detection and localization problem jointly, and allows both non-Gaussian and non-linear sensor and target statistics.

To be concrete, we denote the 2D position and velocity of a single target at time k as $\theta^k = [x \ \dot{x} \ y \ \dot{y}]^T$. Additionally, H_0^k denotes the hypothesis that no target is present at time k , and H_1^k denotes the hypothesis that a single target is present. Finally, z_i^k is the measurement taken in bistatic Doppler pixel i at time k ; z^k represents measurements in all Doppler cells at time k ; and Z^k denotes the collection of all measurements in all Doppler cells taken up to and including time k . When necessary to specify the particular transmitter and receiver pair we will do so as $z_i^k(t, r)$.

Fundamentally, we wish to estimate the joint probability a target is present (i.e., H_1^k is true) and its state is θ^k , given the measurements. Mathematically, we want the hybrid continuous-discrete density [4]

$$p(\theta^k, H_1^k | Z^k) \quad (4)$$

for all θ^k . This quantity can also be expressed as the product of the target present probability and the target state probability:

$$p(\theta^k, H_1^k | Z^k) = p(H_1^k | Z^k) p(\theta^k | H_1^k, Z^k) \quad (5)$$

Computationally, we represent the target state probability $p(\theta^k | H_1^k, Z^k)$ on a 4D discrete grid (corresponding to the four dimensional state vector θ^k) of $N_x \times N_{\dot{x}} \times N_y \times N_{\dot{y}}$ cells. The surveillance region of interest dictates the spatial extent of the grid.

In the sequential Bayesian approach, we (i) assume a prior estimate of the desired probabilities is present (perhaps completely uninformative), and (ii) generate a recursive formula which relates probabilities at one time step with those at the next.

The target present $p(H_1^k | Z^k)$ and absent $p(H_0^k | Z^k)$ probabilities can be computed recursively using the law of total probability and Bayes' rule, yielding:

$$p(H_1^k | Z^k) = p(H_1^k | Z^{k-1}) \frac{p(z^k | H_0^k)}{p(z^k | Z^{k-1})} \times \int \frac{p(z^k | H_1^k, \theta^k)}{p(z^k | H_0^k)} p(\theta^k | H_1^k, Z^{k-1}) d\theta^k \quad (6)$$

The fundamental data dependent quantity needed to construct the target existence update is the likelihood ratio $p(z^k | H_1^k, \theta^k) / p(z^k | H_0^k)$, which is a function of the sensor measurement statistics and the measured data. The numerator is the likelihood of receiving measurement vector

z^k if state θ^k were indeed true. This is evaluated for each possible state θ^k . The denominator is the likelihood of receiving the measurement vector z^k if in fact there were no target in the region. Notice no thresholding is used. The quantity z^k is the measurement vector received (consisting, for example, of L Doppler values that result from an FFT of L received pulses). This includes all cells, not simply those pixels that correspond to threshold exceedances. Also we use the data from all transmit/receive pairs in this update.

The target state probability is computed similarly:

$$p(\theta^k, H_1^k | Z^k) = p(\theta^k, H_1^k | Z^{k-1}) \frac{p(z^k | H_1^k, \theta^k) p(z^k | H_0^k)}{p(z^k | H_0^k) p(z^k | Z^{k-1})} \quad (7)$$

Therefore, in principle, the update for each possible θ^k proceeds by computing its predicted probability at the next time step (according to a temporal target kinematic model) and then updating using the likelihood ratio $p(z^k | H_1^k, \theta^k) / p(z^k | H_0^k)$. This procedure mirrors the Kalman Filter steps (prediction and update), but with Bayes-optimal processing.

All updates require temporal prediction. For the target absent and present probabilities $p(H_0^k | Z^k)$ and $p(H_1^k | Z^k)$, we use a simple mixing matrix approach, which corresponds to a fixed target arrival/removal probability. Our approach also allows more complicated models with little impact on computational requirements.

The temporal evolution of the probability density on θ^k can be expressed in continuous time using a partial differential equation [7]. For example, if the target is modeled as moving with a constant velocity plus random accelerations (the so-called ‘‘nearly constant velocity model’’). In the following discussion we use this model as an example. Other models may be equally appropriate, including models that use road constraints and higher order motion terms.

Continuing under the assumption that we have a CPI of L pulses from N transmitter/receiver pairs, each of the N^2 pairs (t, r) has a set of I of pixelated bistatic Doppler cells. Then the numerator of the likelihood function can be written (assuming conditional independence given the state) as

$$\begin{aligned} f(z; \theta) &\equiv p(z^k | H_1^k, \theta^k) \\ &= \prod_{t=1}^N \prod_{r=1}^N \prod_{i \in I} p(z_i^k(t, r) | \theta) \\ &= \prod_{t, r} \prod_{i \in I} p(z_i^k(t, r) | \theta) \end{aligned} \quad (8)$$

Figure 1 illustrates the the nonlinear filtering approach. The panels show the target position and velocity marginals after a number of time steps. A white dot shows truth, and the regions of high probability are close to the truth. Notice the non-Gaussian nature of the PDF.

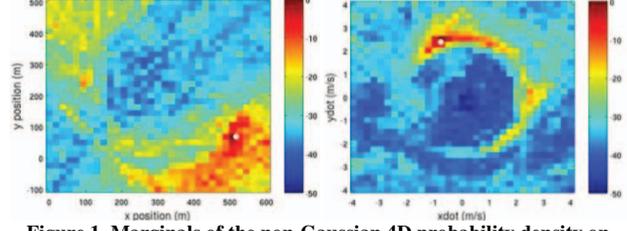


Figure 1. Marginals of the non-Gaussian 4D probability density on target state. It inherits the non-Gaussianity, from the complicated nonlinearities and Rayleigh statistics of the bistatic Doppler measurements.

IV. CRAMER-RAO BOUND SENSOR PLACEMENT

This section discusses the Cramer Rao bound (CRB) approach for optimal sensor placement. Qualitatively, our sensor placement algorithm selects sensor locations that yield the best performance in estimating target location in terms of the CRB on target state estimator variance. In other words, we choose the sensor positions that yield the best *bound* on achievable estimator performance.

While other authors [2] use CRB metrics for gauging the performance of bistatic systems, our approach includes several innovations. First, the approach places sensors to optimize detection and localization performance, rather than just to gauge performance. Second, it accounts for the entire region under surveillance using a novel spatial weighting on the CRB. Finally, it approximates a computationally intractable global optimization in a high dimension space by an iterative-greedy approximation. Simulations show this approach yields intuitive results which perform well in simulations.

We want the CRB corresponding to eq. (2), but are hampered by differentiability issues because the Doppler cell envelope statistics are spatially discontinuous. For this reason, we make the approximation

$$\lambda_i(\theta; t, r) \approx \lambda_b + (\lambda_t - \lambda_b) \exp\left(-\frac{(h(i) - \mu(\theta; t, r))^2}{2\sigma^2}\right) \quad (9)$$

where $h(i)$ represents the Doppler frequency corresponding to pixel i , and $\mu(\theta; t, r)$ relates the target parameter vector to a Doppler frequency:

$$\mu(\theta; t, r) = \frac{c - g(\theta; t_x, t_y)}{c + g(\theta; r_x, r_y)} f_c - f_c \quad (10)$$

where

$$g(\theta; a, b) = \frac{\dot{x}(x - a) + \dot{y}(y - b)}{\sqrt{(x - a)^2 + (y - b)^2}}$$

Assuming the frequency bins are independent yields the joint distribution among the pixels as

$$\begin{aligned}
f(z; \theta) &= \prod_{t,r} \prod_i f_i(z; \theta, t, r) \\
&= \prod_{t,r} \prod_i \frac{z_i}{\lambda_i^2(\theta; t, r)} \exp\left(-\frac{z_i^2}{\lambda_i^2(\theta; t, r)}\right),
\end{aligned} \tag{11}$$

we wish to find the Cramer-Rao bound associated with estimation of the parameter vector θ from the measurements z . For an unbiased estimator this bound is based on the inverse of the Fisher Information Matrix (FIM), whose $(m, n)^{th}$ element is defined as

$$I_{m,n} = \mathbb{E}_z \left[\frac{\partial}{\partial \theta_m} \log f(z; \theta) \frac{\partial}{\partial \theta_n} \log f(z; \theta) \right] \tag{12}$$

The required partials are

$$\begin{aligned}
\frac{\partial}{\partial \theta_n} \log f(z; \theta) &= 2 \sum_{t,r} \sum_i \frac{z_i^2}{\lambda_i^3(\theta; t, r)} \frac{\partial \lambda_i(\theta; t, r)}{\partial \theta_n} \\
&\quad - \frac{1}{\lambda_i(\theta; t, r)} \frac{\partial \lambda_i(\theta; t, r)}{\partial \theta_n}
\end{aligned} \tag{13}$$

Which leads to the FIM elements:

$$\begin{aligned}
I_{m,n} &= 4 \left(\sum_{t,r} \sum_i \frac{\frac{\partial \lambda_i(\theta; t, r)}{\partial \theta_m}}{\lambda_i(\theta; t, r)} \right) \left(\sum_{p,q} \sum_j \frac{\frac{\partial \lambda_j(\theta; p, q)}{\partial \theta_n}}{\lambda_j(\theta; p, q)} \right) \\
&\quad + 16 \left(\sum_{t,r} \sum_i \frac{\frac{\partial \lambda_i(\theta; t, r)}{\partial \theta_m} \frac{\partial \lambda_i(\theta; t, r)}{\partial \theta_n}}{\lambda_i^2(\theta; t, r)} \right)
\end{aligned} \tag{14}$$

The desired CRB for a parameter θ is then

$$\text{cov}_\theta(T(\mu)) \geq I(\theta)^{-1}, \tag{15}$$

where $T(\mu)$ is any unbiased estimator of the θ . In the current setting, we are interested in estimating target state over a known road.

Let $C = I(\theta)^{-1}$, and C_{xx} denote the x elements from the matrix C . We have chosen to use a minimax approach and choose to select the sensor locations $\hat{x}_1, \hat{y}_1, \dots, \hat{x}_N, \hat{y}_N$

$$\hat{x}_1, \hat{y}_1, \dots, \hat{x}_N, \hat{y}_N = \underset{x_1, y_1, \dots, x_N, y_N}{\text{argmin}} \max_{\theta \in \text{road}} (C_{xx} + C_{yy}), \tag{16}$$

In practice, finding the global optimal solution is difficult due to the large dimensionality of the problem (i.e., determining the (x, y) coordinates of N sensors jointly is an optimization in \mathcal{R}^{2N}). We therefore instead employ an iterative greedy approach, where we always optimize one sensor at a time, by placing it at the position that minimizes

the CRB bound on estimation performance. The method proceeds as follows. First, we randomly select positions for all N sensors. Then we hold $N-1$ sensors fixed, and choose the best location for the remaining sensor by minimizing the CRB for all possible positions. This is a tractable 2D optimization, which we do by discrete enumeration of a selected set of possible positions for the sensor. We then move to the second sensor, holding the other $N-1$ fixed (including the most recently positioned sensor and its new position). We repeat this process for all sensors, and then start back at the first sensor, repeating until no sensors change position.

This has the advantage of recasting the problem as N optimizations in \mathcal{R}^2 (which we further discretize to a finite number of possibilities). The disadvantage is that the joint coupling of the sensor locations is not fully modeled. We are not guaranteed to find the globally optimal solution to (16), and in fact in simulation we find that the final sensor position does depend on the initial choice of sensor location. However, interestingly, the minimum found is typically very close to the actual global minimum in value.

Figure 2 shows an example of the CRB surface for placing sensor 6. The CONOPS entails providing surveillance over a road highlighted in green. Sensors 1 through 5 have already been placed, and we are attempting to place the sixth sensor. Each potential positioning has been evaluated under the CRB metric (low is good). As the figure shows, putting sensors in places already well covered by existing sensors is a poor choice in terms of minimizing the CRB.

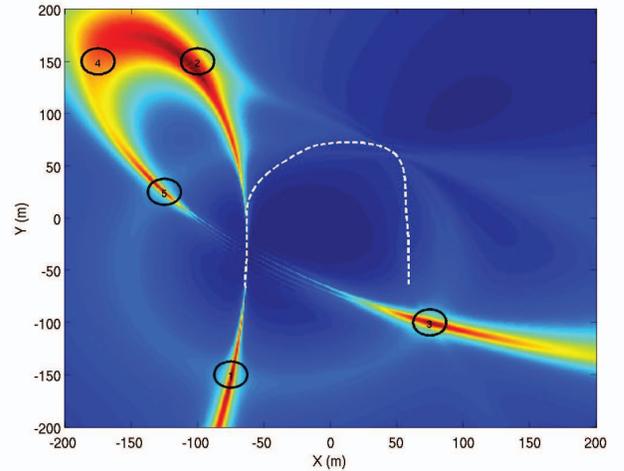


Figure 2. A CRB sensor placement surface. A road is indicated in white. The placement of sensors 1 through 5 are indicated by the black circled numbers. The surface shows the CRB for placing the sixth sensor at each possible location in the surveillance region. Red is high, indicating a poor choice of sensor location and blue is low indicating a good choice of sensor location.

Figure 3 compares tracking performance using sensors placed using the CRB optimal algorithm and sensors placed randomly, for 2, 4, 6, 8 and 10 sensors. As the figure illustrates, CRB optimal placement of the sensors leads to significant improvement in tracking performance.

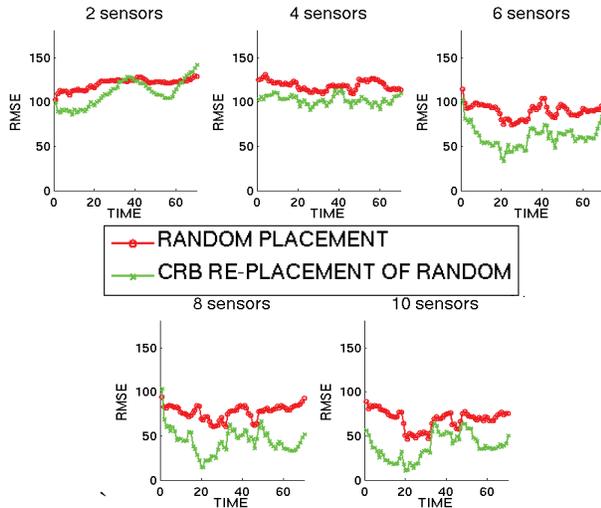


Figure 3. A comparison of tracker performance when using sensors placed randomly and sensors placed using the CRB optimal method.

Figure 4 illustrates the CRB surface when the road constraint is not used to bias the sensor locations. It uses the same five sensor locations as the earlier figure, but does not use the road to bias the placement of the sixth sensor.

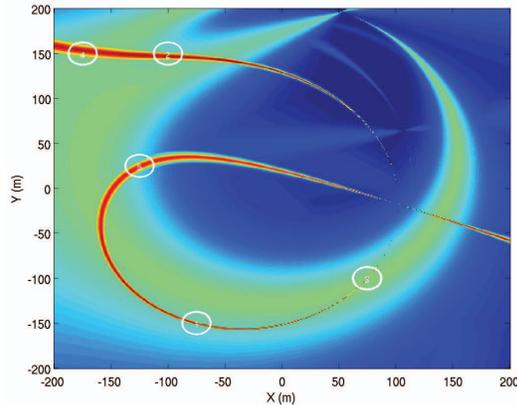


Figure 4. CRB sensor placement surface without road constraints for the sixth sensor (about to be placed). Red is high, indicating poor sensor locations and blue is low, indicating good sensor locations.

V. CONCLUSION

This paper presented a novel approach to detecting and tracking moving targets using a constellation of narrowband radio frequency (RF) sensors. The methodology is an innovative combination of nonlinear estimation and information theoretic sensor placement. The nonlinear filtering approach fully exploits bistatic Doppler measurements made by the sensors without thresholding or linear/Gaussian assumptions, thereby improving the detection/false alarm tradeoff and lowering tracking error. The information theoretic sensor placement algorithm, which is based on minimizing the Cramer-Rao bound on localization variance, selects sensor positions that lead to the best estimation performance, thereby maximally exploiting the finite sensing resources.

VI. REFERENCES

- [1] B. Himed, H. Bascom, J. Clancy and M. Wicks, "Tomography of moving targets", Proceedings of SPIE, the International Society for Optical Engineering, 2001, vol 4540, pp. 608-619.
- [2] L. Lo Monte, R. Ansari, D. Erricolo, and M. C. Wicks, "The use of geometric diversity for spectral dominance in underground imaging", 2009 IEEE International Waveform Diversity and Design Conference, pp. 55-61.
- [3] Haimovich, A. et. al. "CRB on Target Localization Estimation in MIMO Radar Systems" 42nd Information Sciences and Systems Conf., 2008.
- [4] C. Kreucher, B. Shapo, and R. Bethel, "Multitarget Detection and Tracking using Multi-Sensor Passive Acoustic Data", Proceedings of The Thirtieth Annual IEEE Aerospace Conference, March 7 – 14, 2009.
- [5] C. Kreucher, K. Kastella, and A. Hero, "Multitarget Tracking using the Joint Multitarget Probability Density", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1396-1414, October 2005.
- [6] M. Morelande, C. Kreucher, K. Kastella, "A Bayesian Approach to Multiple Target Detection and Tracking", *IEEE Transactions on Signal Processing*, vol. 55, no. 5, pp. 1589-1604, May 2007.
- [7] K. Kastella and C. Kreucher, "Multiple Model Nonlinear Filtering for Low Signal Ground Target Applications", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 2, April 2005, pp. 549-564.