An Information Theoretic Radar Tracker Performance Model

Chris Kreucher, Matthew Masarik, John Valenzuela, and Rebecca Malinas
Integrity Applications Incorporated
Ann Arbor, Michigan, 48108
Email: ckreuche@umich.edu

***Abstract*—This paper describes a new information theoretic approach for modeling the performance of a radar Tracker. Our approach is based on computing the Posterior Cramér-Rao Lower Bound (the so-called "Tracker bound") as a function of radar capabilities, locations, and target state. The main contribution of this paper is an extension to more closely model fielded trackers which exhibit a series of track-drop and re-initialization events during the target lifetime. We account for this phenomena by marginalizing the Posterior (Bayesian) information matrix with respect to the tracker start time. We show that this new approach allows us to very closely predict the behavior of a radar tracker by comparing the new performance prediction to Monte Carlo runs of a tracker on a model 3-radar tracking problem with centralized fusion.**

I. Introduction

This paper describes an information theoretic approach for modeling the performance of a radar Tracker. Like earlier work [3], [4], [7]–[10], we use the Posterior Cramér-Rao Lower Bound (the so-called "Tracker bound") to capture the best possible performance of a tracker as a function of radar capabilities, locations, and target state.

In this paper, we modify the optimistic Posterior Cramér-Rao Lower Bound (PCRLB) to account for the fact that a practical tracker is subject to track-drop and re-initialization events. Failure to account for this phenomenon leads to significant discrepancies between the bound and actual tracker performance. Our main contribution is the introduction of an information marginalization approach which conditions on when the track was initialized. In short, instead of computing the information matrix assuming the track was initiated at scan 1 and held to the present scan, we instead compute it under all possible initialization times and marginalize using the probability of each.

Our approach produces an expression which predicts tracker performance as a function of the scenario parameters, but is now far less optimistic then the nominal PCRLB. There are two main benefits of this model. First, it provides an analytical expression for the expected tracker performance that is a function of radar parameters (bandwidth, frequency, power, etc.), as well as the physical location of the radars. It does not require Monte Carlo runs of an actual tracker to predict how a tracker would perform. As such, it can be used in an analytical optimization approach to select radar parameters and locations to meet a desired performance metric. Second, the model makes minimal assumptions about the specifics of the tracker while still achieving improved fidelity over the most optimistic model. The parameters used in our computation are limited to what measurements the tracker exploits, the fusion architecture, and the track start and drop logic.

The rest of this paper is organized as follows. In Section II, we describe how information theory is used to develop a model for the utility of radar measurements. This is done by computing the Fisher Information Matrix (FIM) for the measurements. Next, in Section III, we describe a high-level model for track initiation and drop. Section IV introduces the posterior (Bayesian) information matrix (BIM) and how it models the performance of a tracker by synthesizing FIMs over time. Section V then describes our new marginalization approach to computing the BIM which accounts for track drop and reinitiation. Section VI provides a comparison of the new performance prediction to an EKF tracker. Finally, Section VII concludes.

II. Information Theory for Measurement Utility Modeling

In this section, we define the information available from measurements made by a collection of radar to estimate the state of a moving target. We assume a constellation of S stationary radars with locations $\mathbf{x}_s$, $s = 1, \ldots, S$ and a target with state at time k denoted $\boldsymbol{\alpha}^k$. We will suppress the time dependence for notational clarity until it is needed. The unknown target state is a six-vector capturing the $3D$ position and velocity, i.e., $\boldsymbol{\alpha} = [x\ \dot{x}\ y\ \dot{y}\ z\ \dot{z}]$.

We assume each radar processes data so as to generate measurements of range (r), range-rate ($\dot{r}$), azimuth (θ) and elevation (ϕ) angle at each measurement epoch. The measurement vector is therefore $\mathbf{z} = [r\ \dot{r}\ \theta\ \phi]$. This is a standard model for a pulsed-Doppler radar which radiates multiple pulses, performs a matched filter operation, and uses Fourier processing over a coherent processing interval (CPI) to generate a Range/Doppler Map at a pointing azimuth/elevation [6], [11].

Each radar is characterized by a set of system parameters, including center frequency, bandwidth, transmit powers and antenna gains. These parameters define the resolution of the radar measurements and when coupled with target range and RCS, define the measurement signal to noise ratio (SNR).

With that as background, let $f(\mathbf{z}; \boldsymbol{\alpha})$ be a scalar PDF on the measurement vector $\mathbf{z}$ given the parameter vector $\boldsymbol{\alpha}$. Our statistical model of the measurements is as follows. First, we

The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

978-1-7281-1679-2/19/$31.00 ©2019 IEEE

assume conditional independence among the measurements, i.e., we assume $f(\mathbf{z};\boldsymbol{\alpha}) = f_r(\mathbf{z};\boldsymbol{\alpha})f_{\dot{r}}(\mathbf{z};\boldsymbol{\alpha})f_\theta(\mathbf{z};\boldsymbol{\alpha})f_\phi(\mathbf{z};\boldsymbol{\alpha})$.

For the measurement of range we assume the radar generates a measurement of target range corrupted by Gaussian noise with variance matched to that implied by the range resolution $c/2BW$ [6], where c is the speed of light and BW is the radar bandwidth, i.e., we assume the range measurement from sensor s is

$$r_s \sim N\Big(||\boldsymbol{\alpha}_p - x_s||, \frac{c}{2BW\sqrt{12}}\Big), \tag{1}$$

where α_p is the position components of α. For convenience we can also write this $r_s = h_r(\alpha; x_s) + w$, where h_r is the nonlinear transformation from target space to measurement space and w is the noise which has standard deviation $\frac{c}{2BW\sqrt{12}}$.

The range-rate, azimuth, and elevation measurements are defined analogously. We use the radar resolutions and beamwidths to define the variance of a Gaussian error on these components as we did for the range measurement.

The Fisher Information Matrix (FIM) is a 6×6 matrix with elements defined as

$$F_{i,j} = -\mathbb{E}\left[\frac{\partial^2}{\partial\alpha_i\partial\alpha_j}\log f(\mathbf{z};\boldsymbol{\alpha})\right]. \tag{2}$$

Subject to some regularity conditions, the Cramér-Rao Bound [12] says that for any unbiased estimator $\hat{\boldsymbol{\alpha}}(\mathbf{z})$, the estimator covariance satisfies

$$cov_{\boldsymbol{\alpha}}\big(\hat{\boldsymbol{\alpha}}(\mathbf{z})\big) \geq F(\boldsymbol{\alpha})^{-1}. \tag{3}$$

And as such, the FIM $F(\boldsymbol{\alpha})$ captures the information available to the tracker from the measurements.

With our Gaussian noise assumption and a nonlinearity h, we can use the definition to write the $(ij)^{th}$ element of the FIM as

$$F_{i,j} = \frac{1}{\sigma^2}\frac{\partial h(\mathbf{z};\boldsymbol{\alpha})}{\partial\alpha_i}\frac{\partial h(\mathbf{z};\boldsymbol{\alpha})}{\partial\alpha_j}, \tag{4}$$

For example, the xy component of the FIM for the range measurement for sensor s is

$$F^r_{xy} = \frac{(x-x_s)(y-y_s)}{\sigma_r^2\sqrt{(x-x_s)^2+(y-y_s)^2+(z-z_s)^2}}. \tag{5}$$

Assuming conditional independence on the measurements as before, the FIM for sensor s is then the combination of the FIMs from the individual measurements, i.e., $F^{s,r} + F^{s,\dot{r}} + F^{s,\theta} + F^{s,\phi}$.

We model a non-unity detection probability (P_d) using the information reduction factor [7], yielding single-sensor FIM

$$F^s = P_d(\boldsymbol{\alpha}, x_s)\big(F^{s,r} + F^{s,\dot{r}} + F^{s,\theta} + F^{s,\phi}\big). \tag{6}$$

Here we have suppressed the other factors effecting detection probability, such as the radar parameters.

Finally, the FIM corresponding to the measurements made by all sensors assuming centralized fusion and time-synchronized measurements is

$$F = \sum_{i=1}^{s} P_d(\boldsymbol{\alpha}, x_s)\Big(F^{s,r} + F^{s,\dot{r}} + F^{s,\theta} + F^{s,\phi}\Big). \tag{7}$$

This matrix captures the information available at a single measurement instant from the measurements of a collection of radars.

III. The In-Track Model

In this section, we develop a two-state model for when the target is in and out of track. Let $P_d^{s,t}$ be the detection probability from sensor s at time t. This number is meant to include the standard probability that the target-originated reflection has energy which exceeds threshold, the probability a detection gates with the tracker (P_G), and the probability it is correctly associated (P_{CA}) [13]. The traditional "Swerling" detection probability is defined by the link budget, target RCS, and target range. The gating probability captures effects like large random measurement errors and target maneuvers. In sum, the composite detection probability is a function of target state, target RCS profile, and radar parameters and as such it will vary from scan-to-scan.

We continue to assume centralized fusion, i.e., all data is fused at a central station. We assume that the central station forms tracks using this information and therefore a "detection" means that at least one node has a detection. We define the "network" detection probability as the probability that at least one node has a threshold exceedance, i.e.,

$$P_d^t = \Big(1 - \prod_i (1 - P_d^{i,t})\Big). \tag{8}$$

Other fusion architectures are possible. For example, we could instead assume that each sensor forms tracks independently and the tracks are fused rather than measurements. In this case, we would use $P_d^s = P_d^{i,s}$ and perform the following analysis at each sensor. We do not discuss this possibility further in this paper.

Our high-level model for the tracker is that it proceeds through a sequence of in-track and out-of-track states. Ideal tracking would have the tracker initiate at time 1 and hold for the entire target life, but in practice there are periods of time where the target is in track and periods where it is out of track.

Fielded trackers like the Multiple Hypothesis Tracker (MHT) use "M-of-N" track confirmation (initiation) and deletion logic. Blackman [1] gives an example wherein a target is initiated with 3 detections in 4 frames (3-of-4) and removed with 4 consecutive misses. Therefore, our model for how a tracker moves in and out of track will be given as follows: A track is initiated when (at least) M detections occur in the last N opportunities, and an existing track is deleted if there are K consecutive misses.

978-1-7281-1679-2/19/$31.00 ©2019 IEEE

We use the following notation to capture this phenomenon:

- P^t_{MN} is the probability that at least M detections have occurred in the last N opportunities at scan t (i.e., scans $t-N+1, t-N+2, \dots t$).
- P^t_K is the probability of K consecutive non-detections occurring at scan t.

Since the base detection probability P^s_d is a function of scan t, computing these quantities cannot be done by appealing to the standard binomial distribution. We use the method of [2] to compute these quantities efficiently.

IV. The Bayesian Information Matrix

The Bayesian Information Matrix (BIM) describes the information available to a tracker by synthesizing FIMs and the motion model over time. If a target was under track from scan 1 until time k, the Bayesian Information Matrix J can be computed using the recursion of [3]. In the special case where the tracker models the target motion as linear and Gaussian (e.g., with the so-called nearly constant velocity model), the recursion takes a particularly simple form,

$$J^{t+1} = (Q + A(J^t)^{-1}A^T)^{-1} + F^{t+1}, \tag{9}$$

where we define J^0 to be some minimal information matrix, and we execute the recursion from scan 1 until scan k. The PCRLB at time t is then given by $(J^t)^{-1}$.

J^0 is chosen to represent minimal information, and here we set it to have an inverse which is very large value, to indicate the uncertainty is the size of the surveillance region. Later simulations will show that this setting does not materially effect the conclusions because when the target is under track the PCRLB quickly stabilizes to a state which is very weakly dependent on this initial condition.

Let the notation $J^k_{1:k}$ denote the BIM at time k if the target was initiated at scan 1, i.e., the BIM computed using the recursion from scan 1 until scan k. This BIM is valid if the target was held in-track for the entire timespan $1, \dots, k$.

V. The Marginalized BIM and PCRLB

This section generalizes the computation to the case where a target is not in track for its entire lifetime. If the tracker initiated at scan $q > 1$, or if the tracker has experienced a sequence of initializations and drops ending with the latest initialization at time q, the BIM should be computed using only measurements from scan q until the current scan. We will denote this BIM $J^k_{q:k}$.

The main idea of this paper is that we should compute the BIM at the current time by marginalizing over the track initialization time, i.e.,

$$J^k = \sum_{q=1}^{k} p^k(q) J^k_{q:k} + (1 - p^k_{in}) J^0, \tag{10}$$

where $p^k(q)$ is the probability the tracker initiated at scan q and held until k, $J^k_{q:k}$ is the corresponding BIM, and $1 - p^k_{in}$ is the probability the target is not in track at time k.

Each quantity $J^k_{q:k}$ is computed by straightforward application of the Tichavsky [3] recursion starting from scan q and running until scan k. The main task remaining is to specify the probability $p^k(q)$ to perform the marginalization.

We first compute the probability of being in track at time t. We notice there is a recursion which can be used to approximately compute if a target is in track: either (i) the target was in track at time $t-1$ and didn't drop, OR (b) the target was out of track at time $t-1$ and we received an $M-of-N$ event at time t. This can be expressed as

$$p^t_{in} = p^{t-1}_{in}(1 - p^t_K) + (1 - p^{t-1}_{in}) p^t_{MN}, \tag{11}$$

where $p^{M-1}_{in} = 0$ and $p^M_{in} = p^M_{MN}$.

Similarly, we can compute the probability that the track has just started at time t as $p^t_{init} = p^t_{in}(1 - p^{t-1}_{in})$. Finally, the probability that a track is initiated at time q and held until scan k is the probability it was initiated and then not dropped,

$$p^k(q) = p^q_{init} \prod_{i=q}^{k} (1 - p^i_K). \tag{12}$$

Eqs. (10) - (12) completely specify the marginalization.

VI. Example

We illustrate the new marginalized performance model by comparison to an Extended Kalman Filter (EKF) tracker which uses $M-of-N$ initiation and K-miss drop logic. Here we selected $M = 3, N = 5$ and $K = 3$. The scenario uses 3 S-band radars with $30MHz$ bandwidth. A target flies around the region at approximately $250m/s$ as illustrated in Figure 1. We assume that each radar has a 1-degree azimuth and elevation beamwidth, a $64ms$ CPI, a $10s$ scan rate, $20dB$ transmit gain, and that the target has a constant $3dBsm$ RCS. We simulate a 60 minute (361 scan) vignette. These values completely define the $r, \dot{r}, \theta, \phi$ measurement uncertainties and received SNR.

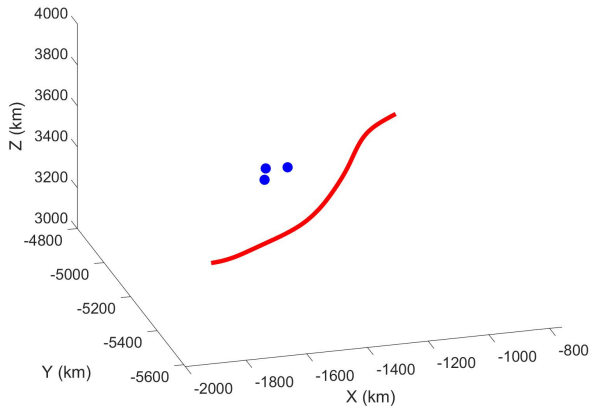

Fig. 1: The simulation scenario we use to illustrate the new performance prediction. A single target flies at constant velocity through a field of 3 surveillance radar.

978-1-7281-1679-2/19/$31.00 ©2019 IEEE

Figure 2 illustrates the RMSE computed from the PCRLB for this scenario for cases where the tracker is assumed to be initialized and held starting at scans $1, 100, 200$, and 300. Each curve shows the PCRLB from the presumed initialization time until the current time.

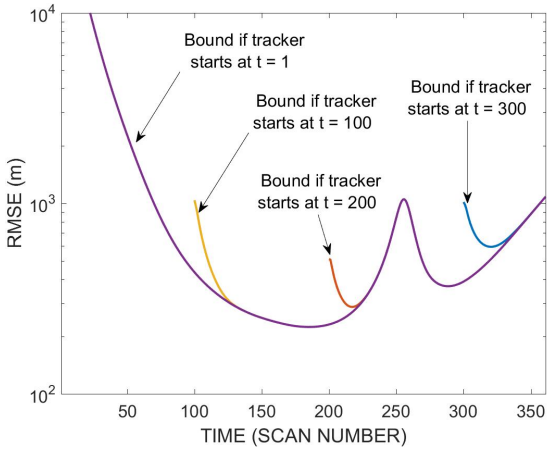

Fig. 2: The tracking performance prediction for several tracker initialization scans, as measured by the RMSE from the PCRLB.

These information curves are combined with the probability that the tracker has initiated at a scan and holding to the current scan to compute the marginalized information. Figure 3 illustrates one such probability curve, showing the probability of tracker initiation and hold until scan 285 for each possible initiation time. The two main peaks mean that there are basically two situations : the tracker initiated around scan 100 and held, or the tracker has only recently initiated. In addition, the probability the target is out of track at time 285 is 10%.

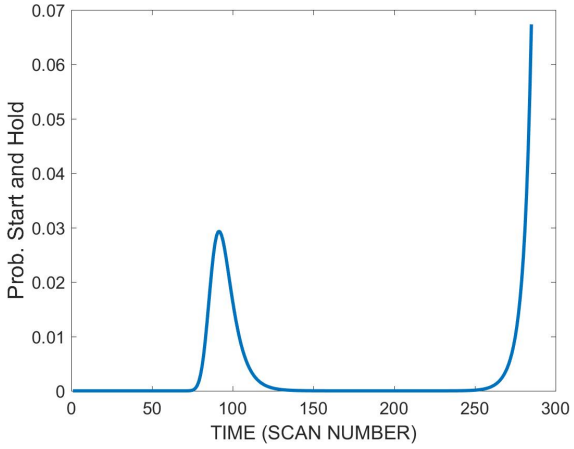

Fig. 3: The probability of starting at each scan and holding until scan 285. . The probability the target is out of track at time 285 is 10%.

Figure 4 amplifies this point by showing the PCRLB for all possible initialization scans from scan 1 to scan 361. Notice that there is always a non-zero probability the target is out of track. In this case the tracker error corresponds to J^0 (i.e., a very large value meant to indicate the uncertainty is the size of the surveillance region). This possibility enters prominently into the marginalization as seen by the solid black line in the figure, which is the marginalization of the BIMs and thus includes the probability of being out-of-track.

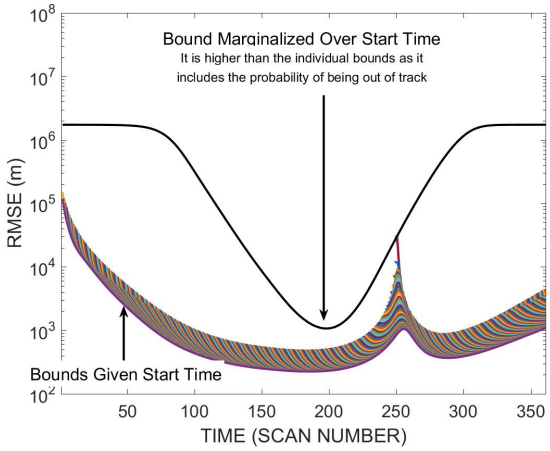

Fig. 4: The tracking performance prediction for all tracker initialization scan times and the marginalization. Note that the marginalization is computed using the individual BIMs *and* the probability that the target is out of track at a scan. Therefore, the marginalized prediction must be higher than any of the individual bounds. The prediction dips at scan 200 mainly due to the probability of being out of track very low.

We now compare the performance prediction to a tracker. We implemented an EKF with the same MNK logic used to predict the performance. 200 instantiations of the tracker are shown in figure 5. These include track-drop events and runs where the tracker is out of track for large periods of time.

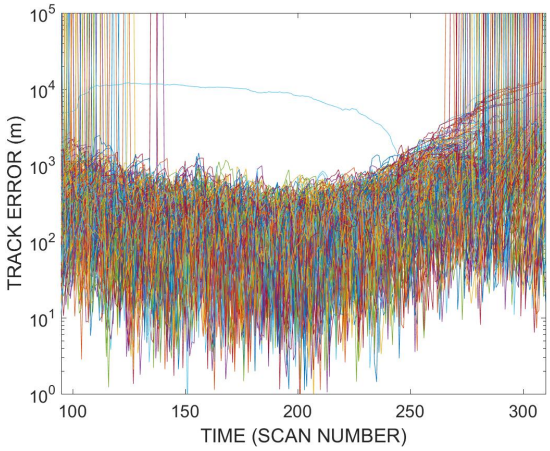

Fig. 5: 200 instantiations of an EKF in the model problem.

Figure 6 shows the new marginalized prediction along with the RMSE of the EKF taken over many trials. Notice that the prediction averages over both in and out-of-track target

978-1-7281-1679-2/19/$31.00 ©2019 IEEE

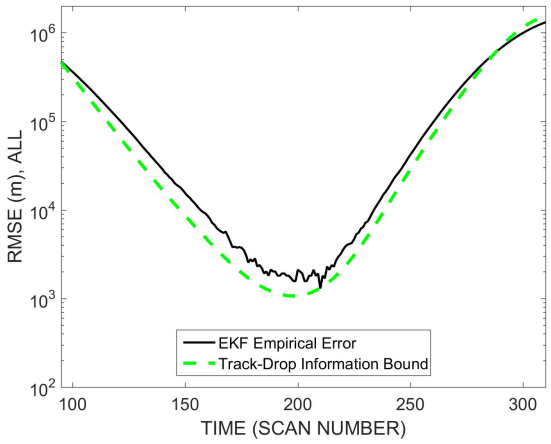

Fig. 6: A Comparison of the tracker RMSE and the track-drop-BIM. The dashed curve is the new track-drop-sensitive marginalized prediction, and the solid curve is the RMSE for many runs of an EKF in a matched scenario.

states and the EKF RMSE calculation includes times where the tracker is in and out of track. When the EKF is out-of-track we mark its squared error as the size of the surveillance region. Analogously, when the model predicts there is some probability of the target being out of track, the marginalization includes a J^0 term which corresponds to this same size.

VII. Conclusion

This paper has described an approach to modeling the performance of a radar Tracker using information theory. Like earlier work, our approach is based on computing the Posterior Cramér-Rao Lower Bound (the so-called "Tracker bound) as a function of radar capabilities, locations, and target state. A new aspect of this work is we modify the bound to more closely model how fielded trackers work, which includes a sequence of track-drop and reinitialization events.

Our approach is to perform a marginalization of the Posterior (Bayesian) information matrix. We show the utility of this new performance prediction by comparison to a large set of Monte Carlo runs of a tracker.

Acknowledgment

This research was funded by DARPA under contract FA8750-17-C-0033. We are grateful to Lt. Col. Jimmy Jones for his support.

References

[1] S. Blackman, "Multiple Hypothesis Tracking for Multiple Target Tracking", *IEEE Aerospace and Electronics Systems Magazine*, vol. 19, no. 1, January 2004, pp. 5-18.

[2] Y. Hong, "On computing the distribution function for the Poisson binomial distribution". *Computational Statistics and Data Analysis*, vol. 59, pp. 41-51, 2013.

[3] P. Tichavsky, C. Muravchik, and A. Nehorai, "Posterior Cramér-Rao Bounds for Discrete-Time Nonlinear Filtering", *IEEE Transactions on Signal Processing*, vol. 46, 1998.

[4] C. Kreucher, "Optimal Sensor Placement for a Constellation of Multi-static Narrowband Pixelated Sensors", *IEEE Transactions on Systems, Man, and Cybernetics – Part C: Applications and Reviews*, 42(6):1374-1383, Nov. 2012.

[5] C. Kreucher and K. Bell, "A Geodesic Flow Particle Filter for Non-Thresholded Radar Tracking". *IEEE Transactions on Aerospace and Electronic Systems*, 2018.

[6] M. Richards, J. Scheer and W. Holm, *Principles of Modern Radar: Basic Principles*, 2010.

[7] M. L. Hernandez, A. D. Marrs, N. J. Gordon, S. R. Maskell, and C. M. Reed, "PCRLB for Tracking in Cluttered Environments: Measurement Sequence Conditioning Approach", *IEEE Transactions on Aerospace and Electronic System*, vol. 42, no. 2, April 2006.

[8] M. L. Hernandez, B. Ristic, A. Farina, L. Timmoneri, "A Comparison of Two CramrRao Bounds for Nonlinear Filtering with Pd less than 1", *IEEE Transactions on Signal Processing*, vol. 52, no 9. September 2009.

[9] P. Stinco, M. Greco, F. Gini, and A. Farina, "Posterior Cramér-Rao lower bounds for passive bistatic radar tracking with uncertain target measurements", *Signal Processing*, vol. 93, 2013.

[10] H. Meng, M. Hernandez, Y. Liu, and X. Wang, "Computationally efficient PCRLB for tracking in cluttered environments: measurement existence conditioning approach", *IET Signal Processing*, 2008.

[11] J. Toomay, and P. Hannen, Radar Principles for the Non-Specialist, 3rd. edition, 2004.

[12] H. Van Trees and K. Bell (eds), Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking, Wiley, 2007.

[13] Y. Bar-Shalom and T. Fortmann, Tracking and Data Association, Academic Press Professional, San Diego CA, 1987.

978-1-7281-1679-2/19/$31.00 ©2019 IEEE