The Modified Riccati Equation for Amplitude-Aided Target Tracking in Heavy-Tailed Clutter

EDMUND BREKKE, Member, IEEE
National University of Singapore

ODDVAR HALLINGSTAD, Member, IEEE
Norwegian University of Science and Technology

JOHN GLATTETRE
Kongsberg Maritime AS

The performance of tracking methods can most often only be assessed by means of Monte-Carlo simulations. An exception to this rule is the popular probabilistic data association filter (PDAF), whose root mean square error (RMSE) can be predicted by means of the modified Riccati equation (MRE).

To the best of our knowledge, the first treatment along these lines for the PDAF with amplitude information (PDAFAI) is presented here. We evaluate the MRE with amplitude information (AI) for the case of a Swerling I target in heavy-tailed, or more precisely K-distributed, background noise.

The MRE can be used to determine an optimal nominal false alarm rate. To the best of our knowledge, the first systematic approach to the determination of false alarm rates in heavy-tailed clutter is presented here. In particular, it is indicated that the PDAFAI can safely operate in the presence of very abundant clutter, while the PDAF only can cope with limited amounts of clutter.

Manuscript received March 1, 2010; revised July 23, 2010; released for publication October18, 2010.

IEEE Log No. T-AES/47/4/942907.

Refereeing of this contribution was handled by W. Koch.

Authors' current addresses: E. Brekke, Acoustic Research Laboratory, National University of Singapore, Singapore, E-mail: (edmundfo@gmail.com); O. Hallingstad, University Graduate Center, Norwegian University of Science and Technology, Postboks 70, Kjeller, 2027, Norway; J. Glattetre, Kongsberg Maritime AS, NO-3191 Horten, Norway.

I. INTRODUCTION

The minimum mean square error (MMSE) estimator for a linear filtering problem is the Kalman filter, and the covariance of its error is governed by the Riccati equation. For target tracking the Kalman filter alone is no longer an adequate solution even if the kinematics and the measurement model are linear. The presence of clutter as well as the sporadic absence of target measurements necessitates data association, for example through the probabilistic data association filter (PDAF) [1, p. 129, 2] or the PDAF with amplitude information (PDAFAI) [1, p. 249, 3].

Our intuition tells us that since neither clutter nor missed detections are beneficial, larger estimation errors must be expected when solving a tracking problem than when solving the corresponding filtering problem. The modified Riccati equation (MRE) [4] was developed to make this intuitive notion precise. It is a generalization of the Riccati equation [5, p. 211] to single-target data association problems, or more precisely to the PDAF. The key idea of the MRE is a so-called information reduction factor which quantifies how much information is lost due to nonunity detection probability P_{D} and false alarms with rate P_{FA}.

The MRE is an important tool because it provides a connection between the detection process and the tracking process. For any detection scenario the probability of detection P_{D} is a function of the false alarm rate P_{FA}. This function is known as the receiving operating characteristic (ROC) and traces a curve in the plane spanned by P_{FA} and P_{D}. There are at least two ways in which the MRE can be used to determine an optimal false alarm rate for a given tracking scenario. Either one can minimize the output of the MRE along the ROC curve, or one can use its output in the hybrid averaging technique of [6] in order to minimize the expected track loss probability.

The lowest possible error covariance that can be attained for a given estimation problem is known as the Cramer-Rao lower bound (CRLB). While the CRLB originally only applied to parameter estimation problems, it was extended by [7] to state estimation problems as well. The version of [7] is often known as the posterior Cramer-Rao lower bound (PCRLB). The PCRLB is a popular tool in the analysis of nonlinear filtering problems, and has been applied to several filtering problems including track-before-detect (TBD) in [8]. It has been evaluated for conventional single-target tracking in [9] and multi-target tracking in [10].

The PCRLB does not tell us how well the PDAF or any other practical tracking method can be expected to perform, but only how an optimal method would perform. This is kind of redundant in the single-target setting, since one has good reason to expect that this optimal performance anyway

is reached by a TBD method, whose PCRLB was analyzed in [8]. The MRE does on the other hand predict the expected performance of the PDAF, and is therefore from a practical perspective just as important as the PCRLB. Since the PDAF is suboptimal, the MRE should be viewed as an upper bound of the PCRLB. Recent research on simplified versions of the MRE has attempted to make this relationship more precise [11].

Despite the extensive treatments of the MRE and the PCRLB in the tracking literature, we have seen very little discussion regarding the impact of amplitude information (AI) in this context. To the best of our knowledge, AI has only been included in such treatments for TBD [8, pp. 251–257] and for the maximum-likelihood probabilistic data association (ML-PDA) [12, pp. 157–179].

A prerequisite for the usage of AI is adequate modeling of the amplitudes of both clutter and target measurements. The simplest alternative is to model both as Gaussian or Rayleigh distibuted with different parameters. However, it is often observed that the background noise is more heavy-tailed or target-like than one would expect under these assumptions [13]. An alternative background model which is more adequate under such circumstances is the K-distribution [14]. The reader is referred to [15, sect. 2.3.1] for a comprehensive discussion regarding the benefits of this model. It should be noted that the Rayleigh distribution is a limiting case of the K-distribution.

While heavy-tailed clutter has been a major focus area for the radar and sonar signal processing communities, it has barely been treated from a target tracking perspective. To the best of our knowledge, the first such treatment is found in [16], in which the PDAFAI was tailored to deal with K-distributed clutter. The performance of this tailored PDAFAI together with several other PDAF-based trackers was investigated in [16] using rather realistic Monte-Carlo simulations.

This paper continues the investigation of target tracking in heavy-tailed clutter by predicting the performance of the PDAFAI in the same way as the performance of the PDAF was predicted in [4]. The MRE with AI is in this paper evaluated for a Swerling I target in K-distributed clutter.

The paper is organized as follows. In Section II the problem to be solved by the PDAFAI is formally described in its proper Bayesian setting. The PDAF and PDAFAI methods are briefly summarized in Section III. The MRE with AI is then presented in Section IV. In Section V we discuss results from the evaluation of the MRE, and how these results are in agreement with experimental performance evaluation using Monte-Carlo simulations. A conclusion is given in Section VI. Details regarding the numerical evaluation of the MRE are left for the Appendix. The derivation of the MRE with AI follows along the lines of the derivation of the original MRE as given in [17]. For reasons of brevity the full derivation is not given here. The reader is instead referred to [15].

II. CONCEPTUAL FRAMEWORK

The aim of Bayesian single-target tracking is to evaluate the posterior probability density function (pdf) $p(\mathbf{x}_k \mid \mathbf{Z}^k)$ from the set of received measurements $\mathbf{Z}^k = \{\mathbf{Z}_1, \ldots, \mathbf{Z}_k\}$ where each $\mathbf{Z}_k$ contains m_k measurement vectors: $\mathbf{Z}_k = \{\zeta_k(i)\}_{i=1}^{m_k}$.

The PDAF and the PDAFAI employ a key approximation, namely that the posterior pdf can be collapsed into a single Gaussian which then is propagated to the next time step:

$$p(\mathbf{x}_k \mid \mathbf{Z}^{k-1}) \approx \mathcal{N}(\mathbf{x}_k; \hat{\mathbf{x}}_{k|k-1}, \mathbf{P}_{k|k-1}). \tag{1}$$

The previous data are in other words summarized by the vector $\hat{\mathbf{x}}_{k|k-1}$ and the matrix $\mathbf{P}_{k|k-1}$, which both are treated as known and nonrandom parameters of the prior pdf during the next estimation cycle [1, p. 129].

A. Kinematics

The kinematic state $\mathbf{x}_k$ will typically contain position and velocity, and possibly accelerations, heading, maneuver strengths, etc. The kinematic transition prior is assumed Gaussian and linear:

$$p(\mathbf{x}_k \mid \mathbf{x}_{k-1}) = \mathcal{N}(\mathbf{x}_k; \mathbf{F}\mathbf{x}_{k-1}, \mathbf{Q}). \tag{2}$$

In principle, changes in the mean target power d_k should also be modeled by the transition prior. However, it is very difficult to estimate a time-varying target power accurately for fluctuating targets [18]. Therefore the target power state is treated as a constant parameter $d_k = d_{k-1} = d$.

B. Measurement Model

The measurement model of the single-target tracking problem comprises not only the mapping from kinematic state to the true target-originating measurement, but also models for clutter and for cardinalities of both target and clutter measurements. The measurement model used here is fairly conventional, with the only exception that we consider a more heavy-tailed background distribution than conventionally used. This section is somewhat cursorial. We do not list all the independence assumptions commonly invoked, but refer the reader to [2] or [19] for more comprehensive lists of assumptions.

Each measurement vector $\zeta_k(i)$ in $\mathbf{Z}_k$ can be parameterized into a kinematic part $\mathbf{z}_k(i)$ and, when applicable, an amplitude part $a_k(i)$,

$$\zeta_k(i) = [\mathbf{z}_k^{\mathrm{T}}(i), a_k(i)]^{\mathrm{T}}. \tag{3}$$

The kinematic part $\mathbf{z}_k(i)$ has for clutter measurements a uniform distribution

$$p_0^z(\mathbf{z}_k(i) \mid \mathbf{x}_k) = p_0^z(\mathbf{z}_k(i) \mid \mathbf{x}_k) = \frac{1}{V_k} \tag{4}$$

where V_k is the volume of the surveillance region, i.e., the validation gate. For target-originating measurements, the kinematic component is related to the state by the Gaussian-linear model

$$p_1^z(\mathbf{z}_k(i) \mid \mathbf{x}_k) = \mathcal{N}(\mathbf{z}_k(i); \mathbf{H}\mathbf{x}_k, \mathbf{R}). \tag{5}$$

The amplitude $a_k(i)$ has the pdf $p_1^a(a_k(i) \mid d, \mathbf{q})$ if $\zeta_k(i)$ originates from the target, otherwise it has the pdf $p_0^a(a_k(i) \mid \mathbf{q})$, where $\mathbf{q}$ contains parameters describing the background noise. In this paper the target return is modeled by the Swerling I-model, while the Rayleigh distribution and also the K-distribution are considered as models for the clutter return. The superscript a indicates that these are pdfs of thresholded measurements, and thus only the upper tails of the unthresholded pdfs $p_1(a_k(i) \mid d, \mathbf{q})$ and $p_0(a_k(i) \mid \mathbf{q})$, suitably normalized.

The clutter-only amplitude likelihood is for the conventional case of Rayleigh distributed background noise given by

$$\begin{aligned} p_0^a(a_k(i) \mid \eta) &= \frac{1}{P_{\mathrm{FA}}} p_0(a_k(i) \mid \eta) \\ &= \frac{1}{P_{\mathrm{FA}}} \frac{a}{\eta} \exp\left(\frac{-(a)^2}{2\eta}\right) \end{aligned} \tag{6}$$

where η is the background strength. The amplitude model (6) is commonly encountered in the tracking literature, although most frequently normalized to $\eta = 1$ [1, p. 249, 2].

As a more flexible and realistic alternative this paper suggests the K-distribution [14], which is more heavy-tailed than the Rayleigh distribution. The K-distribution is given by its scale parameter b and its shape parameter ν, where the latter parameter controls the degree of heavy-tailedness. While the Rayleigh distribution is regained for $\nu \to \infty$, very heavy-tailed clutter is characterized by values of ν between 0 and 1. The clutter-only amplitude likelihood for K-distributed clutter is

$$\begin{aligned} p_0^a(a_k(i) \mid \nu, b) &= \frac{1}{P_{\mathrm{FA}}} p_0(a_k(i) \mid \nu, b) \\ &= \frac{1}{P_{\mathrm{FA}}} \frac{4a^\nu}{\sqrt{b}^{\nu+1}\Gamma(\nu)} K_{\nu-1}\left(\frac{2a}{\sqrt{b}}\right). \end{aligned} \tag{7}$$

The target model considered in this paper is the popular Swerling I model, which is equivalent to the Swerling II model as long as the radar or sonar scan is comprised of a single pulse or ping only. In Rayleigh distributed background noise, the clutter plus target amplitude likelihood is

$$\begin{aligned} p_1^a(a_k(i) \mid d, \eta) &= \frac{1}{P_{\mathrm{D}}} p_1(a_k(i) \mid d, \eta) \\ &= \frac{1}{P_{\mathrm{D}}} \frac{a}{d+\eta} \exp\left(\frac{-(a)^2}{2(d+\eta)}\right). \end{aligned} \tag{8}$$

The normalized version of (8) is also found in [1, p. 249, 2].

For a Swerling I target in K-distributed background noise, the clutter plus target amplitude likelihood is given by marginalization over the "texture" η [14, p. 113]:

$$\begin{aligned} p_1^a(a \mid d, \nu, b) &= \frac{1}{P_{\mathrm{D}}} p_1(a \mid d, \nu, b) \\ &= \frac{1}{P_{\mathrm{D}}} \frac{a}{b^\nu \Gamma(\nu)} \int_0^\infty \frac{\eta^{\nu-1}}{\eta + d} \\ &\quad \cdot \exp\left(-\frac{\eta}{b} - \frac{a^2}{2(\eta + d)}\right) d\eta. \end{aligned} \tag{9}$$

Notice that the scalar η refers both to deterministic background strength in the Rayleigh case, and to random texture in the K-distribution case. This is done because η plays the role of a "local" Rayleigh parameter in the compound formulation of the K-distribution [15, pp. 75–76].

Finally, in this paper it is assumed that false alarms are distributed according to a Poisson process with parameter λV_k,

$$\mu(m_k) = \mathrm{Poisson}(m_k; \lambda V_k) = e^{-\lambda V_k} \frac{(\lambda V_k)^{m_k}}{m_k!} \tag{10}$$

and that the true target originating measurement is received with a probability $P_{\mathrm{D}}P_{\mathrm{G}}$, where P_{D} is the probability of detection and P_{G} is the gate probability. These are standard assumptions, which are regularly invoked in standard references such as [1, 2, 19].

C. The Joint Measurement pdf

The additional information received between time step $k-1$ and time step k is of a random nature as specified by the measurement pdfs elaborated in Section IIB. These pdfs can be combined into a joint measurement pdf, which we decompose as follows using the definition of conditional probability:

$$p(\mathbf{Z}_k \mid \mathbf{Z}^{k-1}) = p(\mathbf{Z}_k, m_k \mid \mathbf{Z}^{k-1}) = p(\mathbf{Z}_k \mid m_k, \mathbf{Z}^{k-1}) P\{m_k\}.$$

The cardinality m_k of $\mathbf{Z}_k$ is distributed according to

$$\begin{aligned} P\{m_k\} &= P_{\mathrm{D}}P_{\mathrm{G}}\mu(m_k - 1) + (1 - P_{\mathrm{D}}P_{\mathrm{G}})\mu(m_k) \\ &= e^{-\lambda V_k} \frac{(\lambda V_k)^{m_k - 1}}{m_k!} (P_{\mathrm{D}}P_{\mathrm{G}} m_k + (1 - P_{\mathrm{D}}P_{\mathrm{G}})\lambda V_k). \end{aligned}$$

The measurement pdf $p(\mathbf{Z}_k \mid m_k, \mathbf{Z}^{k-1})$ as conditioned on the cardinality m_k as well as the

previous data $\mathbf{Z}^{k-1}$ is a mixture over the mutually exclusive and exhaustive hypotheses $\theta_k(i)$:

$\theta_k(0)$ No measurement originates from the target

$\theta_k(1)$ Measurement 1 originates from the target

$\vdots$

$\theta_k(m_k)$ Measurement m_k originates from the target.

For $i = 0$ (i.e., no target detection) the prior event probability is

$$\gamma_0 = P\{\theta_k(0) \mid m_k\} = \frac{(1 - P_{\mathrm{D}}P_{\mathrm{G}})\lambda V_k}{P_{\mathrm{D}}P_{\mathrm{G}}m_k + (1 - P_{\mathrm{D}}P_{\mathrm{G}})\lambda V_k} \tag{11}$$

while for all other i it is

$$\begin{aligned}\gamma_1 &= P\{\theta_k(1) \mid m_k\} = \cdots = P\{\theta_k(m_k) \mid m_k\} \\ &= \frac{P_{\mathrm{D}}P_{\mathrm{G}}}{P_{\mathrm{D}}P_{\mathrm{G}}m_k + (1 - P_{\mathrm{D}}P_{\mathrm{G}})\lambda V_k}.\end{aligned} \tag{12}$$

For each $\theta_k(i)$ the corresponding hypothesis-conditioned measurement pdf can be written

$$\begin{aligned}p(\mathbf{Z}_k \mid \theta_k(0), m_k, \mathbf{Z}^{k-1}) &= \frac{1}{V_k^{m_k}} \prod_{j=1}^{m_k} p_0^a(a_k(j)) \\ p(\mathbf{Z}_k \mid \theta_k(i), m_k, \mathbf{Z}^{k-1}) &= \mathcal{N}(\boldsymbol{\nu}_k(i); \mathbf{0}, \mathbf{S}_k) p_1^a(a_k(i)) \\ &\quad \cdot \frac{1}{V_k^{m_k - 1}} \prod_{j \neq i}^{m_k} p_0^a(a_k(j))\end{aligned} \tag{13}$$

where we have introduced the innovation

$$\boldsymbol{\nu}_k(i) = \mathbf{z}_k(i) - \mathbf{H}\hat{\mathbf{x}}_{k|k-1} \triangleq \mathbf{z}_k(i) - \hat{\mathbf{z}}_{k|k-1} \tag{14}$$

and its corresponding covariance

$$\mathbf{S}_k = \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^{\mathrm{T}} + \mathbf{R}. \tag{15}$$

Notice that the innovations depend on the prediction $\hat{\mathbf{x}}_{k|k-1}$, and not on the state $\mathbf{x}_k$.

Combining all this yields

$$\begin{aligned}p(\mathbf{Z}_k \mid m_k, \mathbf{Z}^{k-1}) &= \sum_{i=0}^{m_k} p(\mathbf{Z}_k \mid \theta_k(i), m_k, \mathbf{Z}^{k-1}) P\{\theta_k(i) \mid m_k\} \\ &= \frac{\gamma_1 c_M g^M}{V_k^{m_k} P_{\mathrm{D}} \sqrt{2\pi}^M} \prod_{j=1}^{m_k} p_0^a(a_k(j)) \\ &\quad \cdot \left(\mathrm{b} + \sum_{i=1}^{m_k} e_k(i) l_k^a(i) \right).\end{aligned} \tag{16}$$

The third line of (16) has introduced the short-hand notations

$$e_k(i) = \exp(-\tfrac{1}{2}\boldsymbol{\nu}_k(i)^{\mathrm{T}}\mathbf{S}_k^{-1}\boldsymbol{\nu}_k(i)) \tag{17}$$

and

$$\mathrm{b} = \left(\frac{2\pi}{g}\right)^{M/2} \lambda V_k \frac{1 - P_{\mathrm{D}}P_{\mathrm{G}}}{c_M P_{\mathrm{D}}} \tag{18}$$

where the volume of the validation gate is $V_k = c_M g^M \sqrt{|\mathbf{S}_k|}$, and the constant $c_M = \pi^{M/2}/\Gamma((M/2) + 1)$ is the volume of a unit-radius M-dimensional sphere.

For target tracking without AI (as in the PDAF), the amplitude likelihood ratio $l^a(i)$ is just unity. For a Swerling I target in Rayleigh background noise it is

$$\begin{aligned}l^a(i) &= \frac{P_{\mathrm{FA}}}{P_{\mathrm{D}}} \frac{p_1(a \mid d, \eta)}{p_0(a \mid d, \eta)} \\ &= \frac{P_{\mathrm{FA}}}{P_{\mathrm{D}}} \cdot \frac{\eta}{\eta + d} \cdot \exp\left(\frac{a^2 d}{2\eta(\eta + d)}\right).\end{aligned} \tag{19}$$

For a Swerling I target in K-distributed background noise it is

$$\begin{aligned}l^a(i) &= \frac{P_{\mathrm{FA}}}{P_{\mathrm{D}}} \frac{p_1(a \mid d, \nu, b)}{p_0(a \mid d, \nu, b)} \\ &= \frac{P_{\mathrm{FA}}}{P_{\mathrm{D}}} \cdot \frac{(a\sqrt{b})^{1-\nu}}{4K_{\nu-1}\left(\frac{2a}{\sqrt{b}}\right)} \\ &\quad \cdot \int_0^\infty \frac{\eta^{\nu-1}}{\eta + d} \exp\left(-\frac{\eta}{b} - \frac{a^2}{2(\eta + d)}\right) d\eta.\end{aligned} \tag{20}$$

III. PROBABILISTIC DATA ASSOCIATION

The PDAF and the PDAFAI express the state estimate at time k as a weighted average of the prediction $\hat{\mathbf{x}}_{k|k-1}$ and state estimates conditioned on the latest measurements $\mathbf{z}_k(i)$. This leads to the following Kalman filter-like equations for prediction and measurement update of the state estimate $\hat{\mathbf{x}}_{k|k}$ and its associated covariance $\mathbf{P}_{k|k}$:

$$\begin{aligned}\hat{\mathbf{x}}_{k|k-1} &= \mathbf{F}\hat{\mathbf{x}}_{k-1|k-1} \\ \mathbf{P}_{k|k-1} &= \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^{\mathrm{T}} + \mathbf{Q} \\ \hat{\mathbf{x}}_{k|k} &= \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k \sum_{i=1}^{m_k} \beta_k(i)\boldsymbol{\nu}_k(i) \\ \mathbf{P}_{k|k} &= \mathbf{P}_{k|k-1} - (1 - \beta_k(0))\mathbf{K}_k\mathbf{S}_k\mathbf{K}_k^{\mathrm{T}} + \tilde{\mathbf{P}}_k\end{aligned} \tag{21}$$

where

$$\begin{aligned}\mathbf{K}_k &= \mathbf{P}_{k|k-1}\mathbf{H}^{\mathrm{T}}\mathbf{S}_k^{-1} \\ \mathbf{S}_k &= \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^{\mathrm{T}} + \mathbf{R}_k \\ \tilde{\mathbf{P}}_k &= \mathbf{K}_k \left[\sum_{i=1}^{m_k} \beta_k(i)\boldsymbol{\nu}_k(i)\boldsymbol{\nu}_k(i)^{\mathrm{T}} - \boldsymbol{\nu}_k\boldsymbol{\nu}_k^{\mathrm{T}}\right] \mathbf{K}_k^{\mathrm{T}} \\ \boldsymbol{\nu}_k &= \sum_{i=1}^{m_k} \beta_k(i)\boldsymbol{\nu}_k(i).\end{aligned} \tag{22}$$

The index i ranges over all m_k extracted and validated measurements inside a validation gate $\mathcal{G}$, defined by

$$\mathcal{G}: (\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_{k|k-1})^{\mathrm{T}}(g^2\mathbf{S}_k)^{-1}(\mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_{k|k-1}) < 1 \quad (23)$$

where the scalar g is called the gate size.

The association probabilities $\beta_k(i) = P\{\theta_k(i) \mid \mathbf{Z}^k\}$ in (21) are under the Poisson assumption (10) given by

$$\beta_k(0) = \frac{\mathrm{b}}{\mathrm{b} + \sum_{j=1}^{m_k} l_k^a(j)e_k(j)} \quad (24)$$

$$\beta_k(i) = \frac{l_k^a(i)e_k(i)}{\mathrm{b} + \sum_{j=1}^{m_k} l_k^a(j)e_k(j)} \quad (25)$$

with $e_k(i)$ and b as in (17) and (18), respectively. The reader is referred to [1] or [2] for further details on the PDAF(AI).

IV. THE MODIFIED RICCATI EQUATION

The data association problem solved by the PDAF(AI) is nonlinear due to the measurement origin uncertainty. The covariance $\mathbf{P}_{k|k}$ corresponding to the state estimate $\hat{\mathbf{x}}_{k|k}$ does therefore depend on the data $\mathbf{Z}_k$. Consequently, it is impossible to predict the performance of the PDAF(AI) in the same way as the conventional Riccati equation predicts the performance of the Kalman filter. A similar recursion can nevertheless be carried out if $\mathbf{P}_{k|k}$ is replaced by its conditional expectation

$$\mathbf{P}_k = E[\mathbf{P}_{k|k} \mid \mathbf{Z}^{k-1}] \triangleq E[(\hat{\mathbf{x}}_{k|k} - \mathbf{x}_k)(\hat{\mathbf{x}}_{k|k} - \mathbf{x}_k)^{\mathrm{T}} \mid \mathbf{Z}^{k-1}]. \quad (26)$$

Under some reasonably relaxed conditions outlined in [4] (notably gate probability $P_{\mathrm{G}} \approx 1$), the recursion of $\mathbf{P}_k$ can be written

$$\begin{aligned} \mathbf{P}_{k|k-1} &= \mathbf{F}\mathbf{P}_{k-1}\mathbf{F}^{\mathrm{T}} + \mathbf{Q} \\ \mathbf{P}_k &= \mathbf{P}_{k|k-1} - q_2(\mathbf{S}_k; P_{\mathrm{D}}, P_{\mathrm{FA}})\mathbf{K}_k\mathbf{S}_k\mathbf{K}_k^{\mathrm{T}} \\ \mathbf{S}_k &= \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^{\mathrm{T}} + \mathbf{R}. \end{aligned} \quad (27)$$

This recursion, known as the MRE, differs from the conventional Riccati equation due to the information reduction factor q_2, which is given by

$$q_2 = \sum_{m=1}^{\infty} \eta_{2m}(\mathbf{S}_k)P\{m\} \quad (28)$$

$$\eta_{2m}(\mathbf{S}_k)\mathbf{S}_k = E\left[\sum_{i=1}^{m} \beta(i)^2 \boldsymbol{\nu}_k(i)\boldsymbol{\nu}_k(i)^{\mathrm{T}} \mid m, \mathbf{Z}^{k-1}\right]. \quad (29)$$

It can be shown that the right hand side of (29) is proportional to $\mathbf{S}_k$. The proportionality constant $\eta_{2m}(\mathbf{S}_k)$ is found by averaging over the pdf $p(\mathbf{Z}_k \mid m_k, \mathbf{Z}^{k-1})$, or equivalently over $p(\boldsymbol{\nu}_k(1), \ldots, \boldsymbol{\nu}_k(m_k), a_k(1), \ldots, a_k(m_k) \mid m_k, \mathbf{Z}^{k-1})$:

$$\begin{aligned} \eta_{2m} = {} & \frac{m\gamma_1 c_M}{P_{\mathrm{G}}\sqrt{2\pi}^M}\left(\frac{M}{g^M}\right)^{m-1} \int_0^g \cdots \int_0^g \int_t^\infty \cdots \int_t^\infty \\ & \cdot p_1^a(a_k(j)) \left[\prod_{j=2}^{m} p_0^a(a_k(j))\right] \\ & \cdot \varrho_1^{M+1}\left[\prod_{j=2}^{m} \varrho_j^{M-1}\right] \exp\left(-\frac{1}{2}\varrho_1^2\right) \\ & \cdot \frac{l^a(1)\exp\left(-\frac{1}{2}\varrho_1^2\right)}{\mathrm{b} + \sum_{i=1}^{m} l^a(i)\exp\left(-\frac{1}{2}\varrho_i^2\right)} da(m)\ldots d\varrho_1. \end{aligned} \quad (30)$$

Notice that q_2 depends on $\mathbf{S}_k$ through γ_1 and b which, as can be seen from (12) and (18), depend on λ and the gate volume $V_k = c_M g^M \sqrt{|\mathbf{S}_k|}$. The only way to evaluate (30) is by importance sampling. This can be done by averaging the fraction on the last line over distributions proportional to the preceding factors as explained in the Appendix.

The derivation of the main results (27), (29), and (30) is entirely analogous to the derivation of the conventional MRE in [17]. The only difference is that amplitudes with corresponding pdfs are included in this paper. For a full derivation of the MRE with AI the reader is referred to [15, ch. 5].

V. RESULTS

In this section the performance of the PDAFAI in K-distributed clutter is investigated by means of both simulations and by evaluation of the MRE.

A. System Setup

We evaluate the MRE for a linear constant velocity model whose kinematic state is

$$\mathbf{x}_k = [x_k, \dot{x}_k, y_k, \dot{y}_k]^{\mathrm{T}}.$$

The matrices in (2) and (5) are given by

$$\mathbf{F} = \begin{bmatrix} 1 & T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (31)$$

$$\mathbf{Q} = \begin{bmatrix} \frac{\sigma_v^2}{3}T^3 & \frac{\sigma_v^2}{2}T^2 & 0 & 0 \\ \frac{\sigma_v^2}{2}T^2 & \sigma_v^2 T & 0 & 0 \\ 0 & 0 & \frac{\sigma_v^2}{3}T^3 & \frac{\sigma_v^2}{2}T^2 \\ 0 & 0 & \frac{\sigma_v^2}{2}T^2 & \sigma_v^2 T \end{bmatrix} \quad (32)$$

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \quad (33)$$

TABLE I
Filter Model Parameters

Parameter	Value	Specification
g	6	Gate size
P_{FA}	$2 \cdot 10^{-4}$, $2 \cdot 10^{-2}$	False alarm rates
σ_r^2	$1/12$ m^2	Measurement noise
σ_v^2	$(0.0125$ m/s$^2)^2$	Process noise
T	1 s	Sampling time
c_1	5	Lower track-loss threshold
c_2	15	Upper track-loss threshold
SNR	15 dB	Signal-to-noise ratio
N	300 s	Duration of scenario

The measurement noise matrix $\mathbf{R}$ is in practical applications typically specified in polar coordinates. However, in order to make it meaningful to talk about the steady-state output of the MRE, linearity must be ensured. To keep things simple we use

$$\mathbf{R} = \sigma_r^2 \mathbf{I} \tag{34}$$

where $\mathbf{I}$ is the 2×2 identity matrix. Values for the constants appearing in (31)–(34) can be found in Table I.

B. Performance Analysis using Simulations

In order to validate the MRE, its output should be compared with simulation results. This was done for the conventional MRE in [17], and a similar validation is presented for the AI case in this paper.

In addition to verifying the MRE, this section is also written with the purpose of complementing the simulation results reported in [16]. In that paper realistic Monte-Carlo simulations were used to illustrate the gains that could be expected from using AI in practical scenarios with heavy-tailed clutter. The realistic approach required that the trackers could not know the model parameters a priori; instead they had to be estimated, causing a corresponding performance loss.

The simulations in this paper are carried out according to what we instead may describe as a clean approach. In this approach, the trackers are given perfect knowledge about all model parameters such as P_{FA}, P_{D}, ν, b, and so on. Thus there is no mismatch between filter and simulation models.

The output from the simulations is summarized by two measures: the position error and the rate of track-loss. Caution is required when attempting to estimate the position root mean square error (RMSE) from empirical data. Occurrences of track-loss will, if the number of samples is high enough, cause severe outliers in the error pdf which make it impossible to obtain a well-defined value for the RMSE. In [17] this problem was circumvented by using very few (more precisely 10) Monte-Carlo runs. In this paper we take a different approach: instead of estimating the RMSE empirically we store the entire error pdf as represented by a histogram.

TABLE II
Lost Tracks out of 10000 for $\nu = 1$

		PDAF	PDAFAI
SNR = 10	$P_{\mathrm{FA}} = 10^{-4}$	4367	3789
	$P_{\mathrm{FA}} = 10^{-2}$	3361	834
	$P_{\mathrm{FA}} = 10^{-1}$	9549	721
SNR = 15	$P_{\mathrm{FA}} = 10^{-4}$	67	22
	$P_{\mathrm{FA}} = 10^{-2}$	482	12
	$P_{\mathrm{FA}} = 10^{-1}$	6727	19

TABLE III
Lost Tracks out of 10000 for $\nu = 8$

		PDAF	PDAFAI
SNR = 10	$P_{\mathrm{FA}} = 10^{-4}$	181	74
	$P_{\mathrm{FA}} = 10^{-2}$	1401	51
	$P_{\mathrm{FA}} = 10^{-1}$	9327	64
SNR = 15	$P_{\mathrm{FA}} = 10^{-4}$	3	0
	$P_{\mathrm{FA}} = 10^{-2}$	275	1
	$P_{\mathrm{FA}} = 10^{-1}$	6611	4

In the same way as was done in [16] and [20], this paper treats track-loss as a two-stage process. Denote the true simulated target state by $\mathbf{x}_k^{\mathrm{s}} = [x_k^{\mathrm{s}}, \dot{x}_k^{\mathrm{s}}, y_k^{\mathrm{s}}, \dot{y}_k^{\mathrm{s}}]^{\mathrm{T}}$. A track is then considered tentatively lost at time k if the position error $\sqrt{(x_{k|k} - x_k^{\mathrm{s}})^2 + (y_{k|k} - y_k^{\mathrm{s}})^2}$ exceeds a threshold c_1. If the error later goes below c_1 the lost label is removed. On the other hand, if the error never manages to go below c_1 again, we consider it lost at time k. If the error exceeds the higher threshold $c_2 > c_1$ we immediately consider it lost at time k, irrespectively of whether the error later goes below c_1.

Track-loss results can be seen in Tables II and III. The simulation results were generated using 10000 Monte-Carlo runs for each scenario. The numerical recipe of [16] was used to evaluate the amplitude likelihood for the PDAFAI. The results indicate, as one would expect, that the PDAFAI always performs better than the PDAF. Notice that the PDAFAI offers significant improvements irrespective of the false alarm rate.

The track-loss rate of the PDAF does in general attain its lowest values for low false alarm rates. The PDAFAI can on the other hand often benefit from higher false alarm rates. This is appears to be especially so when the clutter is strongly heavy-tailed ($\nu = 1$), and when the target is weak (10 dB). In such cases, the detection probability P_{D} will necessarily be very low unless a high false alarm rate is used, so it is hardly surprising that a higher false alarm rate yields improved performance. Of course the increased false alarm rate also increases the risk of having the tracker misled by clutter measurements. However,

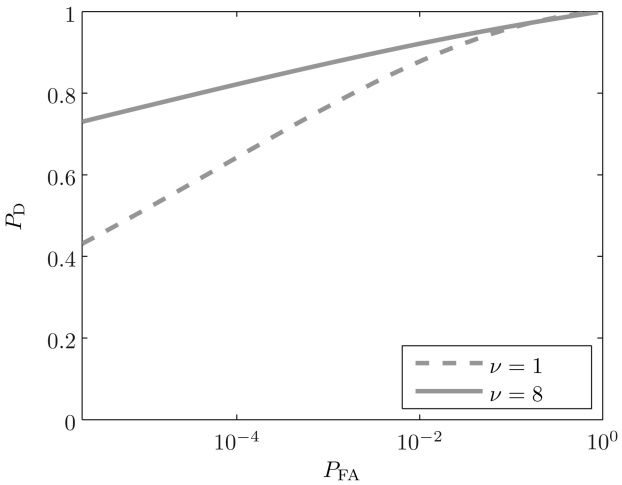

Fig. 1. ROC curves in K-distributed background noise for target with SNR = 15 dB.

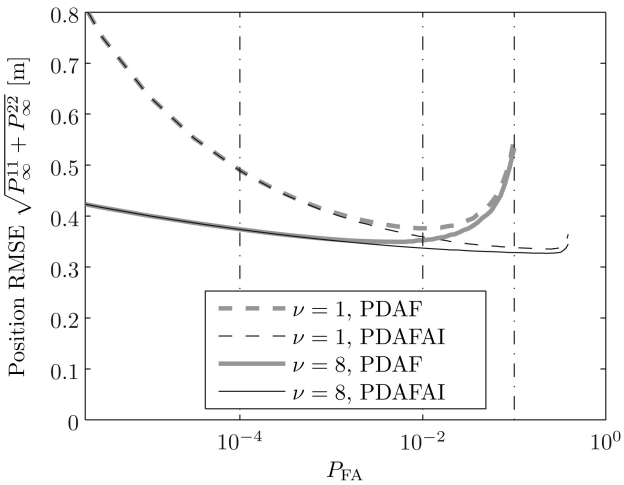

Fig. 2. Evaluation of MRE for target with SNR = 15 dB. Vertical lines indicate false alarm rates used in Figs. 3, 4, and 5.

this risk is much smaller for the PDAFAI than for the PDAF, since the PDAFAI is able to discard most of these clutter measurements by noting that their amplitudes are lower than what would be expected for measurements from the target.

C. Performance Analysis using the MRE with AI

The output of the MRE is most conveniently summarized by the stationary position RMSE as predicted by the MRE:

$$e(P_{FA}, P_D) = \sqrt{P_\infty^{11} + P_\infty^{22}}. \tag{35}$$

For a given scenario, this quantity generates a surface over the P_{FA}-P_D-plane known as the tracking operating characteristic (TOC). Since the ROC curve $P_D(P_{FA})$ is given by the scenario, it makes sense to plot the TOC as a function of P_{FA} along the ROC curve to investigate what would be the optimal false alarm rate for the given scenario.

In Fig. 1 we have plotted the ROC curves for a 15 dB target both in significantly heavy-tailed noise (dashed curve) and in not so heavy-tailed noise (solid curve). It can clearly be seen how the more heavy-tailed case is more challenging due to lower detection probability for the same false alarm rate. This is most noticeable for low false alarm rates, while the difference appears negligible for very high false alarm rates.

Corresponding TOC curves have been plotted in Fig. 2 for both the conventional MRE corresponding to the PDAF, and for the AI-based MRE corresponding to the PDAFAI. This figure tells us that there exists a unique false alarm rate which will minimize the expected RMSE of the PDAF. For not so heavy-tailed clutter, the expected RMSE will hardly increase even if a significantly lower false alarm rate than the "optimal" one is used. Taking the results of Table III into account, we may therefore

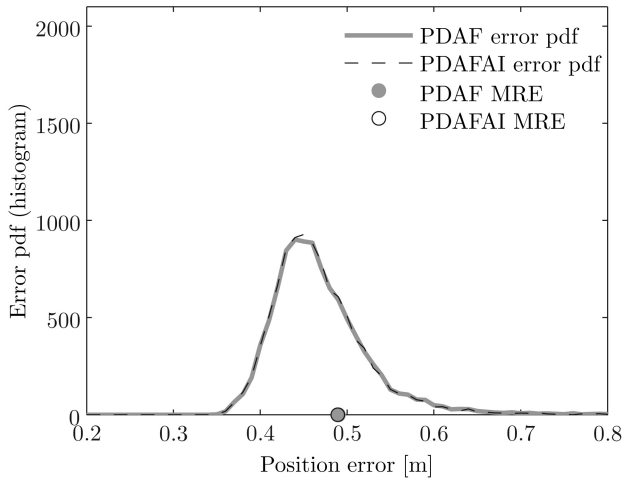

Fig. 3. Error pdf for $\nu = 1$ and $P_{FA} = 10^{-4}$.

in general recommend the usage of low false alarm rates for the PDAF in near-Rayleigh clutter. In more heavy-tailed clutter one may on the other hand prefer a somewhat high false alarm rate (around 10^{-2}) if low RMSE is considered more important than a low track-loss rate. In either case, too high false alarm rates (roughly $\geq 10^{-1}$) cause the MRE to diverge, indicating that the PDAF cannot cope with such large amounts of clutter.

The most striking observation to make from Fig. 2 is that the PDAFAI is able to beat this divergence for much higher false alarm rates (roughly up to $P_{FA} = 0.5$). On the other hand, for low false alarm rates the MRE predicts only marginal improvements due to AI. As far as the MRE is concerned, the main purpose of AI is therefore to make it possible to use false alarm rates so high that target detection can be expected.

Validation results are presented in Figs. 3, 4, and 5, where error statistics of the PDAF and the PDAFAI have been collected. The histograms were

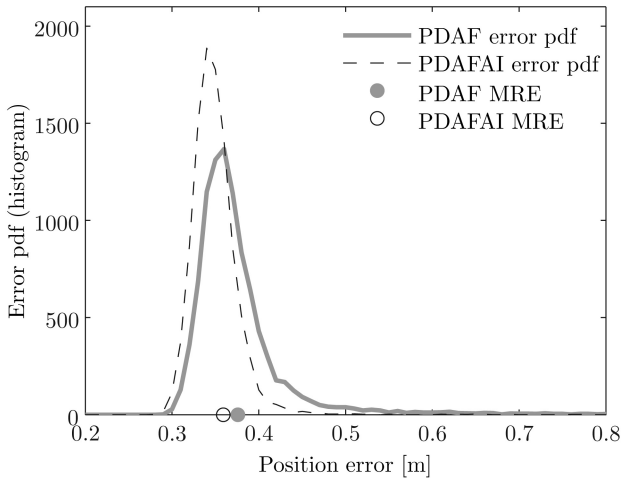

Fig. 4. Error pdf for $\nu = 1$ and $P_{\mathrm{FA}} = 10^{-2}$.

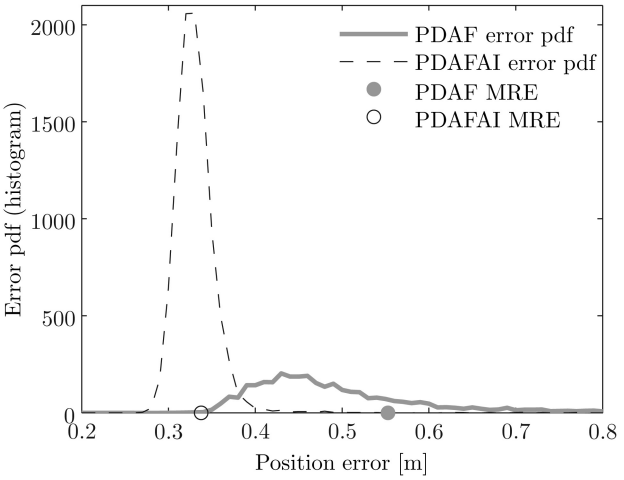

Fig. 5. Error pdf for $\nu = 1$ and $P_{\mathrm{FA}} = 10^{-1}$.

generated using 10000 Monte-Carlo runs as explained in Section VB.

All these three figures reveal a slightly positive (i.e., pessimistic) bias of the MRE as compared with the modes of the error pdfs. On the other hand, the error pdfs are also slightly skewed due to upper tails. This is especially visible for the PDAF in the high-P_{FA} cases treated in Figs. 4 and 5. We may therefore conclude that the MRE does a reasonably good job in predicting the expected error of both the PDAF and the PDAFAI.

One should treat the analysis provided by the MRE with some caution. First, it is very optimistic. The error graphs in [16] showed a slightly worse performance, even after removal of bad tracks. The obvious explanation for this is that in the real world (or in realistic simulations as in [16]), several parameters are not known a priori. The resulting estimation errors cause an inevitable performance loss which is difficult to account for in a simple equation such as the MRE.

Second, the MRE deals only with the expected RMSE. As illustrated in Figs. 3 and 4 this corresponds to the first-order moment of the error pdf. Track-loss is on the other hand an outlier phenomena which can only be adequately addressed if the upper tail of the error pdf is studied. Tables II and III reveal that even when the MRE predicts the same performance from the PDAFAI as from the PDAF, the PDAFAI is actually vastly superior to the PDAF as measured by such rare but undesirable occurrences.

Several criteria must be taken into account when choosing a nominal false alarm rate. Two criteria that we have not discussed here are the computational power of the tracking system and track management. Increasing the number of false alarms clearly increases the computational burden, and it may therefore in practice be impossible to use a higher false alarm rate than, say, 10^{-4}. Lowering the false alarm rate will in particular reduce the number of false preliminary tracks [1, p. 104], which in many tracking systems consume a huge percentage of the computational power.

However, rigorous methods for track management (i.e., initiation and termination) will not function properly if the detection probability is too low. An extreme example is the probability hypothesis density (PHD) filter, which may run into difficulties due to a single misdetection [21]. The situation is less severe, but still troublesome, for the integrated probabilistic data association (IPDA). According to [19], the IPDA should terminate a confirmed track after as few as 3 misdetections. For $P_{\mathrm{D}} \leq 0.8$ this will happen quite frequently. Further advances in track management must therefore either use AI to a larger extent, or exhibit better robustness to low detection probabilities.

VI. CONCLUDING REMARKS

In this paper the MRE has been extended to the case of AI. Using this approach, the performances of the PDAF and the PDAFAI have been predicted in K-distributed clutter.

To the best of our knowledge, this paper has presented the first systematic approach to the determination of nominal false alarm rates in heavy-tailed clutter. In particular, the analysis shows that one may consider setting the false alarm rate quite high when the PDAFAI is used. The MRE also tells us that the improvement of the PDAFAI over the PDAF should be most noticeable when the false alarm rate is high.

Comparison with simulation results has validated the performance prediction offered by the MRE, but also called attention to shortcomings of this approach. The MRE is only able to predict the performance of tracks on target. It provides no information regarding the rate of track-loss, unless it is so high that the MRE

diverges. Therefore, the MRE is unable to shed light on the improvements of the PDAFAI over the PDAF for low false alarm rates.

This last point illustrates that caution should be exercised when analyzing performance by continuous measures such as the RMSE. The outlier nature of track-loss events should be addressed in future research, for example through the hybrid averaging technique of [6]. The research reported here is an important step in that direction.

The shortcomings of the MRE will also be shared by other RMSE-based performance measures such as the PCRLB. The recent surge in research on the PCRLB and other continuous performance measures should therefore be accompanied with critical investigations of whether conclusions obtained by means of these measures also hold when track-loss is used as the primary criterion. It would be desirable, although probably quite challenging, to obtain results similar to those in [6] for other tracking methods than the PDAF(AI).

APPENDIX

The information reduction factor $q_2(P_{\mathrm{D}},P_{\mathrm{FA}},\lambda V_k)$ is a function that depends on both the detection probability P_{D}, the false alarm rate P_{FA}, and the clutter intensity as represented by λV_k. The volume V_k is given by the innovation covariance $\mathbf{S}_k$, which again is given by the output of the MRE at the previous time step $k-1$. The three quantities P_{D}, P_{FA}, and λV are intertwined in the scalar quantities η_{2m}, whose integrals cannot be expressed in closed form.

The most efficient way of propagating the MRE is by constructing an interpolation table. In the original paper [4] the information reduction factor was evaluated over a discrete grid in the $(P_{\mathrm{D}},\lambda V_k)$-plane, which in the absence of AI is equivalent to the $(P_{\mathrm{D}},P_{\mathrm{FA}})$-plane. When amplitudes are included this approach becomes too simple, since P_{FA} is no longer equivalent to λV_k. Instead a 3-dimensional interpolation table must be constructed over P_{D}, P_{FA} as well as λV.

The integral in (30) can only be evaluated numerically. This integral is, especially for higher values of m_k, so high-dimensional that grid-based techniques will run into difficulties. Instead, importance sampling should be used. We have organized the terms in (30) in order to demonstrate how the sampling should be carried out. The amplitudes $a_1,\ldots,a_m$ are sampled using the truncated densities on the first line as explained in Subsection A, while the radial variables $\varrho_1,\ldots,\varrho_m$ are sampled from densities proportional to the quantities on the second line as explained in Subsection B. The

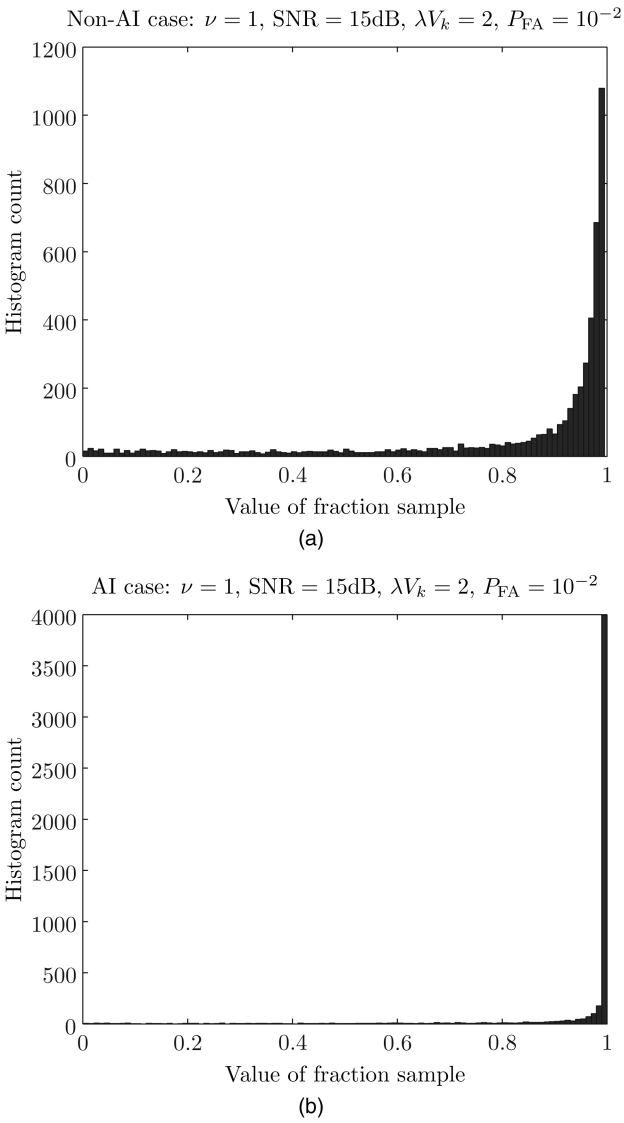

Fig. 6. Distribution of fraction in (38) when variables are drawn as proposed in Subsections A and B. AI can be seen to move probability mass towards one, thereby increasing the kurtosis. Clutter Poisson rate $\lambda V = 2$ indicates a rather tough tracking scenario where AI is necessary to avoid track-loss. For lower values of λV both histograms look more similar to the rightmost one.

ratio on the third line is then averaged over these samples (cf. Fig. 6). Thus, assuming that we have N samples $(a_1^{(l)})_{l=1}^N,\ldots,(\varrho_m^{(l)})_{l=1}^N$, we approximate η_{2m} by an average on the form

$$\eta_{2m} \approx \frac{C_m}{N}\sum_{l=1}^{N}\frac{l(a_1^{(l)}\mid d,\nu,b)\exp\left(-\frac{1}{2}(\varrho_1^{(l)})^2\right)}{\mathrm{b}+\sum_{i=1}^{m} l(a_i^{(l)}\mid d,\nu,b)\exp\left(-\frac{1}{2}(\varrho_i^{(l)})^2\right)}. \tag{36}$$

where C_m is a constant. This constant is comprised of the constant in (30) together with proportionality constants linking the functions of ρ_i in (30) to their corresponding sampling densities for $i = 1,\ldots,m$. By inserting these proportionality constants as obtained in

Subsection B we find

$$C_m = \frac{m\gamma_1 c_M}{P_G\sqrt{2\pi}^M}\left(\frac{M}{g^M}\right)^{m-1} \cdot 2^{M/2}\Gamma\left(\frac{M}{2}+1\right)\cdot\left(\frac{g^M}{M}\right)^{m-1}. \quad (37)$$

It follows that

$$\eta_{2m} \approx \frac{m\gamma_1}{P_G N}\sum_{l=1}^{N}\frac{l(a_1^{(l)} \mid d,\nu,b)\exp\left(-\frac{1}{2}(\varrho_1^{(l)})^2\right)}{\mathrm{b}+\sum_{i=1}^{m} l(a_i^{(l)} \mid d,\nu,b)\exp\left(-\frac{1}{2}(\varrho_i^{(l)})^2\right)}. \quad (38)$$

A. Sampling of Amplitude Components

When amplitudes are ignored there is simply no need to sample the amplitude components, and we may proceed as if the amplitude related terms were all unity. For Rayleigh-distributed clutter the amplitudes can be sampled according to the strategy devised in [12, p. 166].

Sampling from the truncated K-distribution pdfs (7) and (9) is done using inverse transform sampling, implemented by means of interpolation tables. Denote the survival function corresponding to the truncated density $p^a(a \mid d,\nu,b)$ by $S^a(a \mid d,\nu,b)$:

$$S^a(a \mid d,\nu,b) = \int_a^\infty p^a(u \mid d,\nu,b)du. \quad (39)$$

Inverse transform sampling treats the value of survival function itself (or equivalently and more commonly the cumulative distribution function (cdf)) as a random variable which is drawn according to

$$u = S^a(a \mid d,\nu,b) \sim \mathrm{Uniform}(u;[0,1]). \quad (40)$$

Samples of a distributed according to $p^a(a)$ can then be obtained by inverting the survival function.

Alternatively, we may instead draw the negative logarithm of the survival function, which in accordance with (40) must be exponentially distributed:

$$x = f(a) = -\ln S^a(a \mid d,\nu,b) \sim \mathrm{Exponential}(x;1). \quad (41)$$

Amplitude samples are then drawn according to

$$a = f^{-1}(x) \sim p^a(a \mid d,\nu,b). \quad (42)$$

The sample scheme given by (41) is preferred over the sample scheme given by (40) due to numerical benefits. Since no closed-form expression is known for $S^{-1}(\cdot)$, this mapping is most conveniently implemented using linear interpolation. The mapping $f^{-1}(\cdot)$ is, as illustrated in Fig. 7, easily approximated by piecewise linear segments.

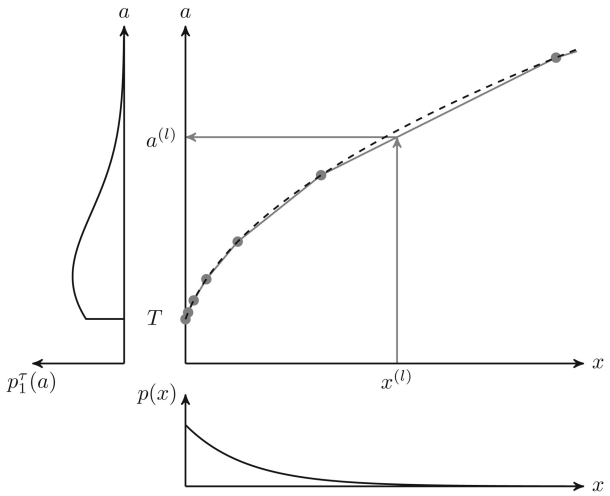

Fig. 7. Sampling technique for amplitude variables. Function $f^{-1}(x)$ (dashed curve) defines mapping from exponential random variable x to truncated K-Swerling I random variable a. Function is approximated by linear interpolation as depicted by solid curve.

An interpolation table for the implementation of $a = f^{-1}(x)$ consists of control points $(x^{(p)},a^{(p)})_{p=1}^P$. We first decide where the control point domain values $a^{(p)}$ should be placed, and thereafter calculate the corresponding codomain values $x^{(p)}$. As argued in [16], the interpolation grid $[a_i^{(1)},\ldots,a_i^{(P)}]$ should have a variable resolution reflecting the fact that the curvature of $p^a(a \mid d,\nu,b)$ decreases as $a \to \infty$. This is obtained by requiring the intervals $\Delta a^{(p)} = a^{(p)} - a^{(p-1)}$ of this grid to be geometrically increasing:

$$\Delta a_i^{(p)} = AB^p. \quad (43)$$

While the lowermost point $a^{(0)}$ of the grid should be the threshold t, its uppermost point $a^{(P)}$ can be any reasonably large value, say 100. The constant A is also a tuning parameter, for which 0.25 has been decided to be an appropriate value. The constant B is then determined by solving

$$a^{(P)} = t + A\frac{1-B^P}{1-B}. \quad (44)$$

The amplitude values of the interpolation grid are then given by

$$t + \left[0, \Delta a^{(1)}, \Delta a^{(1)} + \Delta a^{(2)}, \ldots, \sum_{p=1}^{P}\Delta a^{(p)}\right]. \quad (45)$$

The corresponding control points $x^{(p)}$ are evaluated using (41).

The survival functions $S_1^a(a \mid d,\nu,b)$ and $S_0^a(a \mid \nu,b)$ corresponding to $p_1^a(a_1)$ and $p_0^a(a_1)$ are truncated versions of the survival functions $S_1(a \mid d,\nu,b)$ and $S_0(a \mid \nu,b)$ corresponding to (9) and (7), respectively. Mathematically they can be written

$$S_1^a(a \mid d,\nu,b) = \frac{S_1(a \mid d,\nu,b)}{S_1(t \mid d,\nu,b)} \quad (46)$$

and

$$S_0^a(a \mid \nu, b) = \frac{S_0(a \mid \nu, b)}{S_0(t \mid \nu, b)} \tag{47}$$

where

$$S_1(a \mid d, \nu, b) = \frac{(2/b)^\nu}{\Gamma(\nu)} \int_0^\infty \eta^{\nu-1} \exp\left(-\frac{2}{b}\eta - \frac{a^2}{2(\eta + d)}\right) d\eta \tag{48}$$

and

$$S_0(a; \nu, b) = \frac{2a^\nu}{\sqrt{b}^\nu \Gamma(\nu)} K_\nu\left(\frac{2a}{\sqrt{b}}\right). \tag{49}$$

The integral in the expression for $S_1(a \mid d, \nu, b)$ must be evaluated numerically, for example using the scheme described in [16, Appendix B]. To summarize, we draw samples of a_1 using $S_1^a(a \mid d, \nu, b)$, while samples of all other a_j are drawn using $S_0^a(a \mid \nu, b)$.

The sampling technique is illustrated in Fig. 7, whose major plot shows the mapping $a = f^{-1}(x)$ for a 15 dB target embedded in K-distributed noise with $\nu = 1$. The leftmost plot shows the pdf of target plus clutter, while the bottom plot shows the exponential distribution used to draw $x^{(l)}$. It is illustrated how the linear interpolation (blue curve) provides a very good approximation of the exact mapping.

B. Sampling of Radial Components

The radial variable ϱ_1 is to be drawn from a pdf proportional to the function

$$f(\varrho_1) = \varrho_1^{M+1} e^{-\varrho_1^2/2}. \tag{50}$$

This is a chi-density

$$p(\varrho_1) = \chi(\varrho_1; M+2) = \frac{x^{M+1} e^{-x^2/2}}{2^{M/2}\Gamma\left(\frac{M}{2} + 1\right)}. \tag{51}$$

The sampling density and the actual function to be integrated differ by a proportionality constant $2^{M/2}\Gamma((M/2) + 1)$. In other words,

$$f(\varrho_1) = 2^{M/2}\Gamma\left(\frac{M}{2} + 1\right) p(\varrho_1). \tag{52}$$

The middle factor in (37) follows from this.

The remaining radial variables ϱ_j should be drawn from densities proportional to the monomials ϱ_j^{M-1}. This is done by drawing

$$\varrho_j = g u^{1/M} \qquad \text{where} \quad u \sim \text{Uniform}([0,1]) \tag{53}$$

from which it follows that

$$p(\varrho_j) = \frac{M}{g^M} \varrho^{M-1}. \tag{54}$$

It follows that the sampling density and the function to be integrated differ by a proportionality constant M/g^M. Thus the third factor in (37) must be included when ϱ_j is drawn using (53) for $j = 2, \ldots, m$.

REFERENCES

[1] Bar-Shalom, Y. and Li, X. R.
Multitarget-Multisensor Tracking: Principles and Techniques.
Storrs, CT: YBS Publishing, 1995.

[2] Kirubarajan, T. and Bar-Shalom, Y.
Probabilistic data association techniques for target tracking in clutter.
Proceedings of the IEEE, **92**, 3 (Mar. 2004), 536–556.

[3] Lerro, D. and Bar-Shalom, Y.
Interacting multiple model tracking with target amplitude feature.
IEEE Transactions on Aerospace and Electronic Systems, **29**, 2 (Apr. 1993), 494–509.

[4] Fortmann, T., et al.
Detection thresholds for tracking in clutter—A connection between estimation and signal processing.
IEEE Transactions on Automatic Control, **30**, 3 (Mar. 1985), 221–229.

[5] Bar-Shalom, Y., Li, X. R., and Kirubarajan, T.
Estimation with Application to Tracking and Navigation.
Hoboken, NJ: Wiley, 2001.

[6] Li, X. R. and Bar-Shalom, Y.
Detection threshold selection for tracking performance optimization.
IEEE Transactions on Aerospace and Electronic Systems, **30**, 3 (July 1994), 742–749.

[7] Tichavsky, P., Muravchik, C. H., and Nehorai, A.
Posterior Cramer-Rao bounds for discrete-time nonlinear filtering.
IEEE Transactions on Signal Processing, **46**, 5 (1998), 1386–1396.

[8] Ristic, B., Arulampalam, S., and Gordon, N.
Beyond the Kalman Filter: Particle Filters for Tracking Applications.
Norwood, MA: Artech House, 2004.

[9] Zhang, X. and Willett, P.
Cramer-Rao bound for discrete-time linear filtering with measurement origin uncertainties.
In *Proceedings of The Workshop on Estimation, Tracking and Fusion: A Tribute to Yaakov Bar-Shalom*, Monterey, CA, May 2001, 546–561.

[10] Hernandez, M. L., Farina, A., and Ristic, B.
PCRLB for tracking in cluttered environments: Measurement sequence conditioning approach.
IEEE Transactions on Aerospace and Electronic Systems, **42**, 2 (Apr. 2006), 680–704.

[11] Boers, Y. and Driessen, J. N.
Bounds for target tracking accuracy with probability of detection smaller than one.
In *Proceedings of the 10th International Conference on Information Fusion*, 2007.

[12] Blanding, W. R.
Multitarget tracking using maximum likelihood techniques.
Ph.D. dissertation, University of Connecticut, 2007.

[13] Abraham, D.
Choosing a non-Rayleigh reverberation model.
In *Proceedings of Oceans 1999 MTS/IEEE, Riding the Crest into the 21st Century*, vol. 1, Seattle, WA, Sept. 1999, 284–288.

[14] Ward, K. D., Tough, R. J. A., and Watts, S.
Sea Clutter: Scattering, the K Distribution and Radar Performance (IET series on radar, sonar and navigation), vol. 20.
London, UK: The Institution of Engineering and Technology, 2006.

[15] Brekke, E.
Clutter mitigation for target tracking.
Ph.D. dissertation, Norwegian University of Science and Technology (NTNU), Trondheim, Norway, June 2010.

[16] Brekke, E., Hallingstad, O., and Glattetre, J.
Tracking small targets in heavy-tailed clutter using amplitude information.
IEEE Journal of Oceanic Engineering, **35**, 2 (May 2010), 314–329.

[17] Fortmann, T., et al.
Detection thresholds for tracking in clutter—A connection between estimation and signal processing.
Bolt Beranek and Newman Inc., Technical Report 5495, 1983.

[18] Rutten, M. G., Gordon, N. J., and Maskell, S.
Recursive track-before-detect with target amplitude fluctuations.
IEE Proceedings: Radar, Sonar and Navigation, **152** (Oct. 2005), 345–352.

[19] Li, N. and Li, X. R.
Target perceivability and its applications.
IEEE Transactions on Signal Processing, **49**, 11 (2001), 2588–2604.

[20] Brekke, E., Hallingstad, O., and Glattetre, J.
Target tracking in state dependent wake clutter.
In *Proceedings of the IEEE Oceans Conference*, Sydney, NSW, Australia, May 2010.

[21] Erdinc, O., Willett, P., and Bar-Shalom, Y.
Probability hypothesis density filter for multitarget multisensor tracking.
In *Proceedings of the 8th International Conference on Information Fusion*, Philadelphia, PA, 2005, 146–153.

Edmund Brekke (S'10—M'11) was born on October 23, 1980 in Trondheim, Norway. He received the M.S. degree in industrial mathematics in 2005, and the Ph.D. degree in engineering cybernetics in 2010, both from the Norwegian University of Science and Technology (NTNU), Trondheim, Norway.

Dr. Brekke is currently working as a research fellow at the Acoustic Research Laboratory (ARL) at the National University of Singapore (NUS). His research interests include target tracking, detection theory, seismic signal processing and autonomous navigation.

Oddvar Hallingstad (M'89) was born on October 17, 1945 in Hol, Norway. He received the M.E. degree from the Norwegian Technical University (NTH), Trondheim, Norway in 1971 and his Ph.D. in 1978 also from NTH.

He joined the Norwegian Defence Research Establishment (FFI), Kjeller, Norway in 1977 where he worked on the inertial navigation system for the new costal vessels and the inertial navigation system for the Penguin air to sea missile. Since 1989 he has been at the Graduate University Center at Kjeller lecturing estimation theory and being an advisor for master and Ph.D. students. He is also an adjunct professor at The Norwegian University of Science and Technology (NTNU) and a part-time employee at FFI.

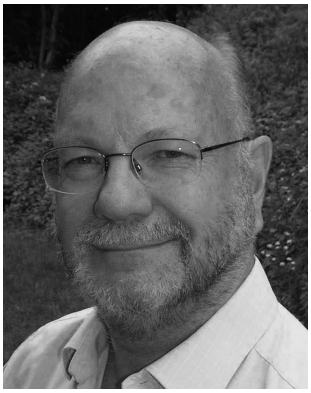

John Glattetre (M'04) was born on November 7. 1944 in Brooklyn, NY and moved to Norway in 1957. He received the B.E. from Oslo Technical College in 1971, the M.E. degree from the Norwegian Technical University (NTH), Trondheim, in 1977 and his Ph.D. from NTH in 1986.

He was an authorized aircraft radio mechanic in Scandinavian Airlines System (SAS). He performed flight calibration of aeronautical navigation systems for the Norwegian Directorate of Civil Aviation. After a short period with micro electronics, he joined the Norwegian Defence Research Establishment (FFI) as a research scientist and worked with sound propagation and signal processing in underwater acoustics. He presently does R&D for Kongsberg Maritime AS, Horten, Norway, where his emphasis is on active sonar, sound propagation, design and signal processing including tracking. He has been coadvisor for 2 Ph.D. students, a task which is successfully concluded.

Dr. Glattetre is a Member of the Acoustical Society of America's Technical Committee on Underwater Acoustics.