Nonlinear Kalman Filters Explained: A Tutorial on Moment Computations and Sigma Point Methods

MICHAEL ROTH
GUSTAF HENDEBY
FREDRIK GUSTAFSSON

Nonlinear Kalman filters are algorithms that approximately solve the Bayesian filtering problem by employing the measurement update of the linear Kalman filter (KF). Numerous variants have been developed over the past decades, perhaps most importantly the popular sampling based sigma point Kalman filters.

In order to make the vast literature accessible, we present nonlinear KF variants in a common framework that highlights the computation of mean values and covariance matrices as the main challenge. The way in which these moment integrals are approximated distinguishes, for example, the unscented KF from the divided difference KF.

With the KF framework in mind, a moment computation problem is defined and analyzed. It is shown how structural properties can be exploited to simplify its solution. Established moment computation methods, and their basics and extensions, are discussed in an extensive survey. The focus is on the sampling based rules that are used in sigma point KF. More specifically, we present three categories of methods that use sigma-points 1) to represent a distribution (as in the UKF); 2) for numerical integration (as in Gauss-Hermite quadrature); 3) to approximate nonlinear functions (as in interpolation). Prospective benefits and downsides are listed for each of the categories and methods, including accuracy statements. Furthermore, the related KF publications are listed.

The theoretical discussion is complemented with a comparative simulation study on instructive examples.

Manuscript received March 31, 2015; revised June 18, 2015; released for publication January 22, 2016.

Refereeing of this contribution was handled by Jindrich Dunik.

Authors' address: Dept. Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden. (E-mail: {roth,hendeby,fredrik}@isy.liu.se).

1557-6418/16/$17.00

I. INTRODUCTION

Many real world problems have a common underlying structure in which the task is to estimate a latent dynamic state from measurements that are taken at specific time instances. Examples include air traffic control [1], where the state comprises the position and velocities of aircraft that are observed by an airport radar; navigation [2], where the state includes the user position and its derivatives, and the measurements might come from a GPS or an inertial measurement unit; or speech processing [3], where the states might be formant frequencies and the measurements are extracted from an estimate of the spectrum. Stochastic state-space models can be used to describe all of these examples in a realistic and mathematically appealing way. The resulting estimation task appears in the form of a Bayesian state estimation problem [4].

The Bayesian filtering problem that is addressed in this tutorial has an elegant mathematical solution, the Bayesian filtering equations [4]. Unfortunately, the mathematical elegance cannot be translated to algorithms in all but a few special cases. In contrast, the Kalman filter [5, 6] *is* a recursive algorithm that has been developed for optimal filtering in linear models, and has been used in countless applications since the 1960s. Interestingly, the first problem where the KF made an impact was in fact a *nonlinear* navigation problem in the Apollo mission [7]. This immediate development of the extended Kalman filter (EKF) has been followed by many nonlinear KF adaptations [6, 8–11].

The interest in the Bayesian filtering problem has significantly increased in the last two decades, mainly due to advances in computing. In the early 1990s, the particle filter [12, 13] was developed to approximately implement the Bayesian filtering equations using sequential Monte Carlo sampling. At the same time, a sampling based extension of the Kalman filter was suggested [14] to improve upon the EKF, eliminate the need for Jacobian matrices, while still retaining the KF measurement update and its computational complexity. Several related algorithms followed: the unscented Kalman filter (UKF) and the underlying unscented transformation were further refined in [15, 16]; numerical integration filters [17, 18] were developed from an alternative starting point to arrive at a very similar recursion; and interpolation variants [17, 19, 20] introduced sampling for function approximation in the spirit of the EKF. In this tutorial, we address these aforementioned sampling based sigma point Kalman filters [21].

Although the literature on nonlinear Kalman filters appears very diverse, it turns out that all variants can be expressed in an intuitive common framework (Gaussian filters [22], Riccati-free Kalman filters [2, 23]). The computation of mean values and covariance matrices of nonlinearly transformed (Gaussian) random variables

appears as a common challenge. The employed methods to approximate these statistical moments distinguishes, for example, the UKF from the divided difference filters of [20]. By emphasizing this common structure, we hope to simplify the reader's access to the extensive literature.

The occurring moment computation problem has also been explored beyond the filtering context [24, 25]. We provide a detailed analysis that reveals useful structural properties, some of which are exploited in the existing methods. Deeper understanding of these properties can give rise to new methods and, in special cases, even reduce n-fold moment integrals to one-dimensional problems.

The moment computation methods that are used in sigma point Kalman filters avoid analytical approximations in favor of sampling techniques. We thoroughly investigate the most common approaches and group them into three categories in which the sigma points carry different roles. More specifically, we present sigma points that approximate distributions as in the UT; sigma points that are used in numerical integration; and sigma points that are used to interpolate nonlinear functions. Interestingly, many equivalences between methods can be shown. Accuracy statements are given where available and also the related KF publications are listed, including recent developments. With this critical survey we hope to clarify the basis of most nonlinear Kalman filters and provide realistic expectations of what a KF can achieve.

Existing overview papers on nonlinear Kalman filters include [26], which significantly extends the numerical integration perspective on the UKF in [17]; [27], in which the UKF and interpolation filters are analyzed and compared as local derivative-free estimators; [23], which establishes relations between sigma point filters and analytical EKF variants; and the recent survey [28], which assumes the linear regression Kalman filter perspective of [29] and discusses algorithms in similar categories as we do in Sec. IV, albeit with less material on the numerical integration and function approximation perspectives. The present tutorial is based on the thesis [30] and complements the existing literature by highlighting the moment computation problem that is central to all nonlinear Kalman filters. Especially the extensive treatment of the analytical moment expressions and the resulting insights for sigma point methods distinguish our paper from previous work.

The structure of the paper is as follows. The introduction is followed by an overview of Bayesian filtering and the unifying framework for nonlinear Kalman filters in Sec. II. The moment computation problem is defined in Sec. III and analyzed in greater detail. Methods to approximately solve the moment integrals are surveyed in Sec. IV. A simulation study with several experiments is presented in Sec. V and followed by Sec. VI with concluding remarks.

II. BAYESIAN STATE ESTIMATION AND NONLINEAR KALMAN FILTERS

This section provides the required background on state-space models and filtering theory. Most importantly, nonlinear Kalman filters are presented in a unified framework that highlights moment computations as central step in each of the algorithms. This motivates for the detailed treatment of moments in Sec. III and IV.

The discussion is extensive. Readers with a background in Bayesian filtering can advance to the nonlinear KF part in Sec. II-D. Alternative references include [4] as a classic on Bayesian state estimation, and more recent treatments in [2] or [22]. A standard reference on the Kalman filter and linear estimation is [6].

A. Stochastic State-Space Models

Stochastic state-space models permit an intuitive and powerful mathematical framework for describing real world processes that are subject to uncertainty, for example the position and velocities of aircraft that are observed by an airport radar.

We consider the general discrete-time model

$$x_k = f(x_{k-1}, v_{k-1}), \tag{1a}$$

$$y_k = h(x_k, e_k), \tag{1b}$$

with the state $x \in \mathcal{X}$, the measurement $y \in \mathcal{Y}$, and the process and measurement noise $v \in \mathcal{V}$ and $e \in \mathcal{E}$, respectively. The function $f : \mathcal{X} \times \mathcal{V} \to \mathcal{X}$ determines the evolution of x in the state difference equation (1a). Similarly, the function $h : \mathcal{X} \times \mathcal{E} \to \mathcal{Y}$ determines the measurement in the measurement equation (1b). Uncertainties in the state evolution and measurements are modeled by asserting that the initial state x_0 and the noise v_{k-1} and e_k are random variables with known distribution for all $k > 0$. Common assumptions are that v_{k-1} and e_k are white and uncorrelated to x_0 and each other; or that the noise sequences[1] and the initial state are independent and admit the joint probability density function

$$p(x_0, v_{0:k-1}, e_{1:k}) = p(x_0) \prod_{l=1}^{k} p(v_{l-1}) p(e_l). \tag{1c}$$

The functions f and h and all densities in (1c) are assumed to be known, but allowed to be time-varying. That allows for including deterministic input signals in (1) through either the function f or the process noise density $p(v_{k-1})$. An extra time index on the functions and densities is omitted for brevity. Following physical reasoning, it is assumed that all measurements are affected by noise and that $\dim(\mathcal{E}) = \dim(\mathcal{Y})$. In fact, additive measurement noise is most common in (1b). For some cases, e.g., when the system (1) is operating in a feedback control loop, the joint density $p(x_0, v_{0:k-1}, e_{1:k})$ cannot be factored as in (1c) because e_k affects v_k.

[1] We use the short hand notation $v_{0:k-1}$ to denote the sequence $\{v_0, v_1, \ldots, v_{k-1}\}$.

An alternative to specifying (1) is to describe the state and measurement processes in terms of the conditional densities[2]

$$x_k \sim p(x_k \mid x_{k-1}), \tag{2a}$$

$$y_k \sim p(y_k \mid x_k). \tag{2b}$$

The symbol $\sim$ in (2a) denotes that, given x_{k-1}, the state x_k is drawn from a conditional distribution that admits the transition density $p(x_k \mid x_{k-1})$. The density $p(y_k \mid x_k)$ is termed likelihood. For the case of additive independent noise in (1), the transition density and likelihood can be specified in terms of the process and measurement noise densities $p(v_{k-1})$ and $p(e_k)$, respectively. The formulation (2) highlights some conditional independence properties in the independent noise case that are important in the derivation of the Bayesian filtering equations [4]:

$$p(x_k \mid x_{0:k-1}, y_{1:k-1}) = p(x_k \mid x_{k-1}), \tag{3a}$$

$$p(y_k \mid x_{0:k}, y_{1:k-1}) = p(y_k \mid x_k). \tag{3b}$$

B. Bayesian State Estimation

The models (1) and (2) fully characterize the state and measurement sequences in a probabilistic manner. This interpretation of the state as a random variable is the cornerstone of Bayesian state estimation, and facilitates the elegant mathematical treatment in terms of probability density functions. In this paper, we shift our attention to estimation problems in which we try to recover the marginal[3] density $p(x_k \mid y_{1:l})$ of a single state x_k given the model and a sequence of measurements $y_{1:l}$.

The previous section has shown how $x_{0:k}$ and $y_{1:k}$ are generated from x_0, $v_{0:k-1}$, and $e_{1:k}$. As a consequence, the joint density $p(x_{0:k}, y_{1:k})$ is a nonlinear transformation of the density $p(x_0, v_{0:k-1}, e_{1:k})$. Slightly more general, the density $p(x_{0:k}, y_{1:l})$ can be obtained in the same way. The manipulations of probability density functions

$$p(x_k, y_{1:l}) = \int p(x_{0:k}, y_{1:l}) dx_{0:k-1}, \tag{4a}$$

$$p(x_k \mid y_{1:l}) = \frac{p(x_k, y_{1:l})}{p(y_{1:l})}, \tag{4b}$$

$$p(y_{1:l}) = \int p(x_k, y_{1:l}) dx_k, \tag{4c}$$

show how the Bayesian state estimation problem of finding (4b) can, in principle, be reduced to the basic operations of conditioning and marginalization. This holds for any l and k. In fact, the order of marginalization and conditioning can be interchanged. Depending on the relation between k and l, three different Bayesian state estimation problems can be distinguished. The case $k > l$ constitutes a prediction problem, the case $k < l$ gives a smoothing or retrodiction problem, and for $k = l$ a filtering problem is obtained. The focus of this article is on filtering and the inherent one-step-ahead prediction, but similar challenges appear in the smoothing problem [22].

The above view of Bayesian state estimation in terms of transformation, conditioning, and marginalization of probability density functions is intuitive. It includes filtering, prediction, and smoothing, and it covers the case of arbitrary correlation in the noise. Unfortunately, it faces some issues. First, finding $p(x_{0:k}, y_{1:k})$ from $p(x_0, v_{0:k-1}, e_{1:k})$ is a challenge beyond hope for all but the simplest models (1). Second, the use of state and measurement sequences does not immediately yield a recursive filtering algorithm.

The Bayesian filtering equations, in contrast, do yield a recursion for the filtering density $p(x_k \mid y_{1:k})$:

$$p(x_k \mid y_{1:k-1}) = \int p(x_k \mid x_{k-1}) p(x_{k-1} \mid y_{1:k-1}) dx_{k-1}, \tag{5a}$$

$$p(x_k \mid y_{1:k}) = \frac{p(x_k, y_k \mid y_{1:k-1})}{p(y_k \mid y_{1:k-1})} = \frac{p(y_k \mid x_k) p(x_k \mid y_{1:k-1})}{p(y_k \mid y_{1:k-1})}, \tag{5b}$$

$$p(y_k \mid y_{1:k-1}) = \int p(y_k \mid x_k) p(x_k \mid y_{1:k-1}) dx_k. \tag{5c}$$

The derivation of (5) relies on the structural properties (3) of the model (1) and can be found in [4]. Similarities between (5) and (4) are apparent, and so it comes as no surprise that some of the challenges of (4) are inherited. First, the marginalization integrals might not be tractable. Second, the filtering density $p(x_k \mid y_{1:k})$ cannot in general be described by a finite number of parameters. For example, in a linear system that is driven by two-component Gaussian mixture noise v_k, the number of parameters to describe $p(x_k \mid y_{1:k})$ grows exponentially in k [2].

C. Linear Gaussian Models and the Kalman Filter

One of the few exceptions for which the Bayesian filtering equations are tractable is the linear Gaussian case. Here, the state transition and measurement equations of (1) can be written as

$$x_k = F x_{k-1} + G v_{k-1}, \tag{6a}$$

$$y_k = H x_k + e_k. \tag{6b}$$

Moreover, the initial state and process and measurement noise are mutually independent Gaussian with

$$p(x_0, v_{0:k-1}, e_{1:k}) = \mathcal{N}(x_0; \hat{x}_0, P_0) \prod_{l=1}^{k} \mathcal{N}(v_{l-1}; 0, Q) \mathcal{N}(e_l; 0, R). \tag{6c}$$

Again, we allow the model to vary with time, e.g., F or Q, but omit the time indices.

[2] For simplicity, we assume the densities to exist. A more general treatment is possible by working on transition kernels [31] instead.

[3] as opposed to the joint density $p(x_{1:k} \mid y_{1:l})$

By induction, it can be easily shown that the filtering density $p(x_k \mid y_{1:k})$ remains Gaussian for all $k > 0$. As a consequence, the Bayesian filtering equations (5) appear as the update equations for mean values and covariance matrices that are known as the Kalman filter (KF) of [5].

From the filtering density $p(x_{k-1} \mid y_{1:k-1}) = \mathcal{N}(x_{k-1}; \hat{x}_{k-1|k-1}, P_{k-1|k-1})$, the one-step-ahead prediction density $p(x_k \mid y_{1:k-1}) = \mathcal{N}(x_k; \hat{x}_{k|k-1}, P_{k|k-1})$ can be obtained. The updated mean value and covariance matrix

$$\hat{x}_{k|k-1} = F\hat{x}_{k-1|k-1}, \tag{7a}$$

$$P_{k|k-1} = FP_{k-1|k-1}F^{\mathrm{T}} + GQG^{\mathrm{T}} \tag{7b}$$

constitute the KF time update or prediction step.

Slightly more general, the joint prediction density of the state and output is given by

$$p(x_k, y_k \mid y_{1:k-1}) = \mathcal{N}\left(\begin{bmatrix} x_k \\ y_k \end{bmatrix}; \begin{bmatrix} \hat{x}_{k|k-1} \\ \hat{y}_{k|k-1} \end{bmatrix}, \begin{bmatrix} P_{k|k-1} & M_k \\ M_k^{\mathrm{T}} & S_k \end{bmatrix}\right), \tag{8}$$

with the mean and covariance of the predicted output

$$\hat{y}_{k|k-1} = H\hat{x}_{k|k-1}, \tag{9a}$$

$$S_k = HP_{k|k-1}H^{\mathrm{T}} + R, \tag{9b}$$

and the cross-covariance

$$M_k = P_{k|k-1}H^{\mathrm{T}}. \tag{9c}$$

The conditioning step in the Bayesian filtering equations simplifies considerably, because rules for the Gaussian distribution can be applied. The KF measurement update yields $p(x_k \mid y_{1:k}) = \mathcal{N}(x_k; \hat{x}_{k|k}, P_{k|k})$ with the filtering mean and covariance

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + M_k S_k^{-1}(y_k - \hat{y}_{k|k-1}) \tag{10a},$$

$$P_{k|k} = P_{k|k-1} - M_k S_k^{-1} M_k^{\mathrm{T}}. \tag{10b}$$

It is common to introduce the Kalman gain $K_k = M_k S_k^{-1}$ and express the measurement update as

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k(y_k - \hat{y}_{k|k-1}) \tag{11a}$$

$$P_{k|k} = (I - K_k H)P_{k|k-1}(I - K_k H)^{\mathrm{T}} + K_k R K_k^{\mathrm{T}} \tag{11b}$$

$$= P_{k|k-1} - K_k S_k K_k^{\mathrm{T}} \tag{11c}$$

$$= (I - K_k H)P_{k|k-1}, \tag{11d}$$

where we listed several alternative expressions for the filtering covariance [6].

The above presentation of the Kalman filter assumes a linear Gaussian model (6), and in fact implements the Bayesian filtering equations. The derivation differs from Kalman's perspective [5] and does not establish certain optimality properties that also hold for non-Gaussian linear systems with known noise statistics [6]. In that case, the KF is still the best linear unbiased estimator (BLUE) but the propagated mean values and covariance matrices are no longer those of Gaussian random variables.

D. Kalman Filters for Nonlinear Models

The preceding sections have shown two extremes. The Bayesian filtering equations (5) are a conceptual solution to a general filtering problem that typically cannot be implemented. In contrast, the Kalman filter solves a specific problem but gives a simple recursive algorithm for updating mean values and covariance matrices. The following section shows how these complementing realities are combined in nonlinear Kalman filters.

Most real world processes exhibit nonlinearities. Furthermore, the process and measurement noise are not necessarily Gaussian. If the nonlinearities are mild and if the noise is not too far from Gaussian, however, a filter that shares a measurement update of the form (10) is a promising candidate. These filters are henceforth termed nonlinear Kalman filters.

The measurement update in the linear KF (10) can be derived from the Gaussian density in (8). In analogy, all nonlinear KF variants employ a Gaussian density

$$\hat{p}(x_k, y_k \mid y_{1:k-1}) \approx \mathcal{N}\left(\begin{bmatrix} x_k \\ y_k \end{bmatrix}; \begin{bmatrix} \hat{x}_{k|k-1} \\ \hat{y}_{k|k-1} \end{bmatrix}, \begin{bmatrix} P_{k|k-1} & M_k \\ M_k^{\mathrm{T}} & S_k \end{bmatrix}\right), \tag{12}$$

to approximate $p(x_k, y_k \mid y_{1:k-1})$, which can be easily shown to be non-Gaussian in the nonlinear case [22]. The measurement update (10) of the linear KF follows and is thus part of any nonlinear KF. This unifying framework is known under the names assumed density filtering [32] or Gaussian filtering [22]. Another interpretation is that the resulting filters locally approximate the Bayesian filtering equations [10, 27]. The differences between specific KF algorithms lie in the way in which the mean values $\hat{x}_{k|k-1}$ and $\hat{y}_{k|k-1}$, and the covariance matrices $P_{k|k-1}$, M_k, and S_k are computed.

Apart from the algorithmic convenience of the KF measurement update, the approximation in (12) is reasonable because the Gaussian is the maximum entropy distribution for a given mean and covariance [33]. That is, if we use the exact mean and covariance of $p(x_k, y_k \mid y_{1:k-1})$ then the Gaussian density in (12) is the least restrictive choice in some sense. Moreover, the Kullback-Leibler divergence

$$\mathrm{KL}(p\|\hat{p}) = -\int p(x_k, y_k \mid y_{1:k-1}) \times \ln\left(\frac{\hat{p}(x_k, y_k \mid y_{1:k-1})}{p(x_k, y_k \mid y_{1:k-1})}\right) dx_k dy_k \tag{13}$$

is minimized by matching the moments [34] of the exact density and its Gaussian approximation, which gives further motivation for using the exact mean values and covariance matrices in (12).

Next, we present the moment integrals to compute the parameters of (12). First, the one-step-ahead prediction of the state is discussed. The model (1a) shows that x_k depends nonlinearly on x_{k-1} and v_{k-1}. In the Kalman filter context with independent noise we assume that the joint density

$$p(x_{k-1}, v_{k-1} \mid y_{1:k-1}) \approx \mathcal{N}(x_{k-1}; \hat{x}_{k-1|k-1}, P_{k-1|k-1}) \mathcal{N}(v_{k-1}; 0, Q) \quad (14)$$

is Gaussian. The mean value and covariance matrix of the predicted state are then given by the moment integrals

$$\begin{aligned} \hat{x}_{k|k-1} &\approx \mathrm{E}\{x_k \mid y_{1:k-1}\} \\ &= \iint f(x_{k-1}, v_{k-1}) \\ &\quad \times p(x_{k-1}, v_{k-1} \mid y_{1:k-1}) dx_{k-1} dv_{k-1}, \end{aligned} \quad (15a)$$

and

$$\begin{aligned} P_{k|k-1} &\approx \mathrm{cov}\{x_k \mid y_{1:k-1}\} \\ &= \iint (f(x_{k-1}, v_{k-1}) - \mathrm{E}\{x_k \mid y_{1:k-1}\}) \\ &\quad \times (f(x_{k-1}, v_{k-1}) - \mathrm{E}\{x_k \mid y_{1:k-1}\})^{\mathrm{T}} \\ &\quad \times p(x_{k-1}, v_{k-1} \mid y_{1:k-1}) dx_{k-1} dv_{k-1}. \end{aligned} \quad (15b)$$

For linear models, (15) simplifies to the KF time update (7). If the noise v_{k-1} enters additively, it does not influence (15a) and appears in (15b) similar to the Q-dependent term in (7b).

The computation of the remaining parameters $\hat{y}_{k|k-1}$, S_k, and M_k can be approached in two different ways. The first option is to accept that

$$y_k = h(f(x_{k-1}, v_{k-1}), e_k) \quad (16)$$

is a function of x_{k-1} and both the process and the measurement noise. Consequently, the expected values must be carried out with respect to a joint density $p(x_{k-1}, v_{k-1}, e_k \mid y_{1:k-1})$.

Alternatively, and common in Kalman filtering algorithms, an intermediate approximation

$$p(x_k \mid y_{1:k-1}) \approx \mathcal{N}(x_k; \hat{x}_{k|k-1}, P_{k|k-1}) \quad (17)$$

is used in the remaining computations. We present the second option below but note that both approaches have their pros and cons. The choice should depend on the accuracy of utilized moment computation methods and the nonlinearities. The intermediate Gaussian density (17) might not reflect the true prediction density well. On the other hand, the composition of the measurement and state transition functions in (16) might be too complicated to be treated directly. For the linear case, there is no difference between the two.

Combining (17) with the measurement noise density yields

$$p(x_k, e_k \mid y_{1:k-1}) \approx \mathcal{N}(x_k; \hat{x}_{k|k-1}, P_{k|k-1}) \mathcal{N}(e_k; 0, R) \quad (18)$$

in the case of independent Gaussian noise. The remaining moment integrals follow as

$$\begin{aligned} \hat{y}_{k|k-1} &\approx \mathrm{E}\{y_k \mid y_{1:k-1}\} \\ &= \iint h(x_k, e_k) p(x_k, e_k \mid y_{1:k-1}) dx_k de_k, \end{aligned} \quad (19a)$$

$$\begin{aligned} S_k &\approx \mathrm{cov}\{y_k \mid y_{1:k-1}\} \\ &= \iint (h(x_k, e_k) - \mathrm{E}\{y_k \mid y_{1:k-1}\}) \\ &\quad \times (h(x_k, e_k) - \mathrm{E}\{y_k \mid y_{1:k-1}\})^{\mathrm{T}} \\ &\quad \times p(x_k, e_k \mid y_{1:k-1}) dx_k de_k, \end{aligned} \quad (19b)$$

$$\begin{aligned} M_k &\approx \mathrm{cov}\{x_k, y_k \mid y_{1:k-1}\} \\ &= \iint (x_k - \mathrm{E}\{x_k \mid y_{1:k-1}\}) \\ &\quad \times (h(x_k, e_k) - \mathrm{E}\{y_k \mid y_{1:k-1}\})^{\mathrm{T}} \\ &\quad \times p(x_k, e_k \mid y_{1:k-1}) dx_k de_k. \end{aligned} \quad (19c)$$

For the linear case, the above expressions simplify to (9).

The computation of (15) and (19) is the crucial challenge in any nonlinear KF. Due to the involved nonlinearities in the model (1) the integrals are typically intractable. Therefore, approximations are utilized. The list of employed concepts includes linearization in the EKF [6]; truncated Taylor polynomials beyond linearization [8, 9, 35]; statistical linearization and alternative function approximations [11, 36–39]; interpolation approaches in the divided difference filters [17, 19, 20]; Monte Carlo [40] and deterministic numerical integration [17, 18, 41, 42] or a combination [43]; and variants of the unscented transformation [15, 16, 44] in the UKF. The majority of the above algorithms are sampling based and can be categorized as sigma point Kalman filters [21].

With the nonlinear KF application in mind we define and analyze the moment computation problems in Sec. III. The methods to solve it, with reference to the related KF variants, are discussed in Sec. IV.

Due to the approximations, many results of the linear KF are no longer valid. For example, there is no nonlinear equivalent to the stationary KF [6] in which $P_{k|k}$ converges to a constant. Furthermore, the order in which measurements are processed in the case of a (block) diagonal measurement noise covariance R matters, as opposed to the linear case [6]. A recommendation from [45] is to process the most accurate measurement first, and then refine the utilized approximation by, for example, re-linearization in an EKF.

The above discussion used Gaussian noise in (14) and (18) as this results in desirable structural properties of the moment integrals (15) and (19). In the non-Gaussian case the relevant densities can be replaced.

We conclude this section with a warning. The Gaussian distribution has limitations because it cannot model phenomena such as heavy tails, skewness, bounded support, or multimodality. An example is shown in Fig. 1, where a Gamma density is illustrated along a Gaussian with the same mean and covariance. The Gamma density is non-zero only for $x > 0$, which is clearly not the case for the Gaussian. Similarly, the Gaussian is symmetric about its mean value. Accordingly, it is a realistic view that Kalman filters, with their close relation to the Gaussian distribution, cannot solve any arbitrary filtering problem. However, more advanced algorithms such as the interacting multiple model (IMM) filter [1] or marginalized particle filters [46] can solve more complicated cases, and employ nonlinear Kalman filters as building blocks.

III. A DETAILED TREATMENT OF THE MOMENT COMPUTATION PROBLEM

The following section introduces and discusses the moment computation problem that is the central challenge in the nonlinear Kalman filters of Sec. II-D, but also relevant beyond the filtering context [24]. We cover structural properties of the moment integrals and ways to exploit them in greater mathematical detail. Furthermore, a compact solution in terms of the Taylor series is presented.

A. The Moment Computation Problem

Sec. II-D highlighted that the challenges in nonlinear Kalman filtering lie in the computation of mean values and covariance matrices of nonlinearly transformed Gaussian random variables. We here generalize this problem to arbitrary functions of Gaussian random variables and functions, that are not necessarily to be seen in a filtering context. Accordingly, the notation is adjusted: the symbol x now represents a generic random variable[4] rather than a state; $f(x)$ represents a generic function rather than the state transition of (1a).

A problem that is encountered in many applications can be formulated as follows. Let the n-dimensional random variable x have a Gaussian distribution $\mathcal{N}(\hat{x}, P)$ with the known mean vector and the covariance matrix

$$\mathrm{E}\{x\} = \hat{x}, \quad \mathrm{cov}\{x\} = P, \tag{20a}$$

and the probability density function

$$p(x) = \mathcal{N}(x; \hat{x}, P). \tag{20b}$$

Furthermore, let $f : \mathbb{R}^n \to \mathbb{R}^m$ be a known function. How can the random variable

$$z = f(x) \tag{21}$$

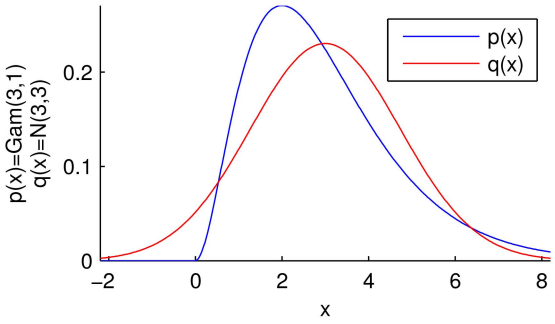

Fig. 1. A Gamma density and a Gaussian with the same mean value and variance.

be characterized in an informative but compact manner?

For special cases, the distribution of z can be obtained via a transformation theorem [47]. In general, however, this is not possible or even desirable. An alternative is to characterize z by its moments. In the Kalman filtering context, the first[5] raw moment and second central moment of z are utilized. Of course, these are better known as the mean value and the covariance matrix and given by the integrals

$$\mathrm{E}\{z\} = \int f(x)p(x)dx, \tag{22a}$$

$$\mathrm{cov}\{z\} = \int (f(x) - \mathrm{E}\{f(x)\}) \times (f(x) - \mathrm{E}\{f(x)\})^{\mathrm{T}} p(x)dx. \tag{22b}$$

Furthermore, cross-covariance matrices of the form

$$\mathrm{cov}\{x,z\} = \int (x - \hat{x})(f(x) - \mathrm{E}\{f(x)\})^{\mathrm{T}} p(x)dx \tag{22c}$$

are required in nonlinear Kalman filters.

An example of a relevant moment computation problem is the conversion between coordinate frames. For example, a Kalman filter might estimate a position in polar coordinates. However, feedback to the user is much more intuitive in Cartesian coordinates, which requires conversion of the polar mean value and its covariance by computing (22a) and (22b). A similar problem is to initialize a Kalman filter with Cartesian state based on the first radar measurements.

The restriction to first and second moments of z is a convenient one, as these are intuitively stored in a vector and a matrix. In principle also higher moments can be considered, although their number grows exponentially with z having m first moments, in the order of m^2 second moments, and so on. Both covariance matrices (22b) and (22c) can be reduced to solving nested problems of the form (22a). Therefore, we concentrate on (22a) in some of the following discussions.

The Gaussian restriction (20b) is chosen because of the Kalman filtering context and because of the structural properties that follow. These are the basis of

[4]The same symbols are used for random variables and their realizations. What is meant should be clear from context.

[5]We avoid the term "order" in connection to moments and reserve it to describe the order of a Taylor series or partial derivatives.

the methods in Sec. IV-C. An extension to arbitrary densities $p(x)$ is immediate.

The moment computation problem (22) requires integration of nonlinear functions over the entire $\mathbb{R}^n$, which is a global problem. However, the integrands are weighted by the probability density function $p(x)$. Hence, only regions in which x has some probability mass contribute to the moment integrals. This is an important aspect to consider when trying to solve (22) by locally approximating f as in Sec. IV-D.

B. Exploiting Structure in the Moment Integrals

The following section shows how the integrals in (22) can be manipulated in order to simplify them and cast them into the structure that some methods of Sec. IV require.

1) Stochastic decoupling:

The following technique is called stochastic decoupling because it can be used to turn expected values with respect to $\mathcal{N}(\hat{x},P)$ into expected values with respect to the standard Gaussian distribution $\mathcal{N}(0,I)$. It employs an $n \times n$ matrix square root $P^{1/2}$ of the covariance P, as presented in (77) of App. A.

Using the rules for substitution in multivariate integrals and $x = P^{1/2}s + \hat{x}$, it is easy to show that

$$\begin{aligned} \mathrm{E}\{z\} &= \int f(x)\mathcal{N}(x;\hat{x},P)dx \\ &= \int f(P^{1/2}s + \hat{x})\mathcal{N}(s;0,I)ds. \end{aligned} \tag{23}$$

The Jacobian determinant of the substitution rule is $\det(P^{1/2}) = \det(P)^{1/2}$ and cancels with the normalization constant of $\mathcal{N}(x;\hat{x},P)$. The quadratic form in the Gaussian density collapsed to $s^{\mathrm{T}}s$.

The reformulation (23) is utilized in several methods throughout Sec. IV. An early reference of its use in filtering is [19].

A related variant is obtained from the substitution rule $x = \sqrt{2}P^{1/2}t + \hat{x}$ and yields

$$\mathrm{E}\{z\} = \int \frac{f(P^{1/2}t + \hat{x})}{\pi^{n/2}} \exp(-t^{\mathrm{T}}t)dt, \tag{24}$$

an integral expression with respect to an exponential weight $\exp(-t^{\mathrm{T}}t)$. The problem of computing (24) is the starting point for Gauss-Hermite quadrature [48] and the numerical integration methods of Sec. IV-C.

2) From Cartesian to spherical coordinates:

The following technique can help to simplify moments of functions that depend on the range only. Furthermore, the separation of moment integrals into spherical and range components has been used to derive integration rules [18, 42, 49].

The change of variables below is a straightforward extension of the Cartesian to polar coordinate change that is taught in undergraduate calculus courses. Consider the relation $s = \Phi(r,\theta_1,\ldots,\theta_{n-1})$ with

$$\begin{aligned} s_1 &= r\sin(\theta_1), \\ s_2 &= r\cos(\theta_1)\sin(\theta_2), \\ s_3 &= r\cos(\theta_1)\cos(\theta_2)\sin(\theta_3), \\ &\vdots \\ s_{n-1} &= r\cos(\theta_1)\cdots\cos(\theta_{n-2})\sin(\theta_{n-1}), \\ s_n &= r\cos(\theta_1)\cdots\cos(\theta_{n-2})\cos(\theta_{n-1}), \end{aligned} \tag{25}$$

where $r \in [0,\infty)$, $\theta_{n-1} \in [-\pi,\pi)$, and $\theta_i \in [-\pi/2,\pi/2)$ for $i = 1,\ldots,n-2$. From (25) follows that $s^{\mathrm{T}}s = r^2$ and thus

$$\mathcal{N}(s;0,I) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{r^2}{2}\right). \tag{26}$$

Using the substitution rule (25) in the moment integral (22a) yields an integral in terms of the range r and the $n-1$ angles θ

$$\begin{aligned} \mathrm{E}\{z\} = \iint f(P^{1/2}\Phi(r,\theta) + \hat{x}) \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{r^2}{2}\right) \\ \times r^{n-1}\cos^{n-2}(\theta_1)\cos^{n-3}(\theta_2)\cdots\cos(\theta_{n-2})dr\,d\theta \end{aligned} \tag{27}$$

over the above mentioned values, e.g., all $r \geq 0$.

A closely related result can be obtained from the fact that any Gaussian variable with the spherically symmetric distribution $\mathcal{N}(0,I)$ can be generated from a random variable ξ that is uniform on the unit sphere [50]. The substitution rule $x = P^{1/2}r\xi + \hat{x}$ yields

$$\mathrm{E}\{z\} = \iint f(P^{1/2}r\xi + \hat{x}) \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{r^2}{2}\right) r^{n-1}dr\,d\xi, \tag{28}$$

where the integration is over all $r \geq 0$ and $\xi^{\mathrm{T}}\xi = 1$.

An interesting case occurs when the function depends on the range only, that is, $f(P^{1/2}\Phi(r,\theta) + \hat{x}) = f(P^{1/2}r\xi + \hat{x}) = \varphi(r)$. Then the spherical integration can be performed analytically with the result

$$\mathrm{E}\{z\} = \frac{2\pi^{n/2}}{\Gamma\left(\frac{n}{2}\right)} \int \varphi(r) \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{r^2}{2}\right) r^{n-1}dr. \tag{29}$$

Alternatively, if $\varphi(r) = \psi(r^2)$ can be written as a function of the squared range, we can arrive at a familiar expression. The substitution rule $r = \sqrt{\rho}$ is a one-to-one relation for $r \geq 0$. The Jacobian determinant is $dr/d\rho = \frac{1}{2}\rho^{-1/2}$, and the integral can be written as

$$\mathrm{E}\{z\} = \int \psi(\rho) \frac{2^{-n/2}}{\Gamma\left(\frac{n}{2}\right)} \rho^{n/2-1} \exp\left(-\frac{\rho}{2}\right) d\rho. \tag{30}$$

The integrand can be recognized as $\psi(\rho)$ times a chi-squared density with n degrees of freedom [49]. In [30], the above technique is used to evaluate the Kullback-Leibler divergence between Gaussian and Student's t

densities, a problem which relates to the filter development of [51].

3) Extension beyond the Gaussian:

The Gaussian distribution is one instance of the wider class of elliptically contoured distributions [49, 50], in which the probability density function $p(x)$ is constant for elliptical regions

$$(x-\hat{x})^{\mathrm{T}}P^{-1}(x-\hat{x}) = c, \tag{31}$$

for $c > 0$. These densities can be written as

$$p(x) = \frac{1}{\sqrt{\det(P)}}h((x-\hat{x})^{\mathrm{T}}P^{-1}(x-\hat{x})), \tag{32}$$

where the function $h : \mathbb{R} \to \mathbb{R}$ is known as density generator [50] and satisfies

$$\int \cdots \int h(s_1^2 + \cdots + s_n^2)ds_1 \cdots ds_n = 1 \tag{33}$$

to ensure that (32) is a valid probability density function. In the Gaussian case, the density generator is

$$h(s^{\mathrm{T}}s) = (2\pi)^{-n/2}\exp\left(-\frac{s^{\mathrm{T}}s}{2}\right) = \mathcal{N}(s;0,I). \tag{34}$$

Both stochastic decoupling and the change to spherical coordinates can be applied to problems in which $p(x)$ is elliptically contoured. A discussion of this is given in [30].

4) Exploiting structure in the function:

In contrast to the preceding paragraphs, the following discussion evolves around structure in $f(x)$ rather than $p(x)$. In fact, the techniques can be applied to arbitrary $p(x)$ and show how the moment computation problem (22) can be broken down systematically.

We consider the case in which x is composed of two random vectors x_1 and x_2, and

$$z = f(x_1, x_2). \tag{35}$$

Following the rules for joint probability density functions, any density $p(x)$ can be factored into

$$p(x) = p(x_1, x_2) = p(x_1 \mid x_2)p(x_2). \tag{36}$$

In the Gaussian case the expression for the marginal density $p(x_2)$ and the conditional density $p(x_1 \mid x_2)$ are well known [49].

Next, we recall to the reader's attention the expressions for the conditional mean

$$\mathrm{E}\{z \mid x_2\} = \int f(x_1, x_2)p(x_1 \mid x_2)dx_1 \tag{37a}$$

and the conditional covariance

$$\begin{aligned}\mathrm{cov}\{z \mid x_2\} = \int (f(x_1, x_2) - \mathrm{E}\{f(x_1, x_2) \mid x_2\})(f(x_1, x_2) \\ - \mathrm{E}\{f(x_1, x_2) \mid x_2\})^{\mathrm{T}}p(x_1 \mid x_2)dx_1.\end{aligned} \tag{37b}$$

These resemble (22a) and (22b) except for the extra conditioning on x_2. A useful consequence of the above is that the moment computation problem (22) can be formulated in terms of the nested expressions

$$\begin{aligned}\mathrm{E}\{z\} &= \mathrm{E}\{\mathrm{E}\{f(x_1, x_2) \mid x_2\}\} \\ &= \int \left(\int f(x_1, x_2)p(x_1 \mid x_2)dx_1 \right) p(x_2)dx_2\end{aligned} \tag{38a}$$

and

$$\begin{aligned}\mathrm{cov}\{z\} = \mathrm{E}\{\mathrm{cov}\{f(x_1, x_2) \mid x_2\}\} \\ + \mathrm{cov}\{\mathrm{E}\{f(x_1, x_2) \mid x_2\}\},\end{aligned} \tag{38b}$$

where the outer expectations average over x_2. The nested covariance expression (38b) is less known than its counterpart (38a) and therefore derived in App. B.

The above expressions are especially useful if $f(x_1, x_2)$ is an affine (or linear) function for a given x_2, i.e.,

$$f(x_1, x_2) = F(x_2)x_1 + g(x_2). \tag{39}$$

Then, the conditional expectations appear as the familiar expressions

$$\mathrm{E}\{z \mid x_2\} = F(x_2)\mathrm{E}\{x_1 \mid x_2\} + g(x_2), \tag{40a}$$

$$\mathrm{cov}\{z \mid x_2\} = F(x_2)\mathrm{cov}\{x_1 \mid x_2\}F(x_2)^{\mathrm{T}}. \tag{40b}$$

The moment computation problem (22) appears as

$$\mathrm{E}\{z\} = \mathrm{E}\{F(x_2)\mathrm{E}\{x_1 \mid x_2\} + g(x_2)\}, \tag{41a}$$

$$\begin{aligned}\mathrm{cov}\{z\} = \mathrm{E}\{F(x_2)\mathrm{cov}\{x_1 \mid x_2\}F(x_2)^{\mathrm{T}}\} \\ + \mathrm{cov}\{F(x_2)\mathrm{E}\{x_1 \mid x_2\} + g(x_2)\}.\end{aligned} \tag{41b}$$

C. The Exact Solution for Differentiable Functions

The section is concluded by a discussion about the exact solution to the moment computation problem (22) for functions that have a convergent Taylor series.

1) Taylor series expansion:

Given that f is differentiable in a neighborhood of $\hat{x}$, it can be expanded as an infinite Taylor series

$$\begin{aligned}f(x) = f(\hat{x}) + \sum_{i=1}^{n} \frac{\partial}{\partial x_i} f(\hat{x})\tilde{x}_i \\ + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} \frac{\partial^2}{\partial x_i \partial x_j} f(\hat{x})\tilde{x}_i\tilde{x}_j \\ + \frac{1}{3!}\sum_{i=1}^{n}\sum_{j=1}^{n}\sum_{k=1}^{n} \frac{\partial^3}{\partial x_i \partial x_j \partial x_k} f(\hat{x})\tilde{x}_i\tilde{x}_j\tilde{x}_k + \ldots.\end{aligned} \tag{42}$$

Here, $\tilde{x} = x - \hat{x}$ denotes the deviation from $\hat{x}$, and $(\partial/\partial x_i)f(\hat{x})$ denotes the partial derivative of f with respect to x_i that is evaluated at $\hat{x}$. For polynomial f of degree d, the dth order Taylor series is exact everywhere. From (20) follows that $\tilde{x} \sim \mathcal{N}(0, P)$.

The infinite series (42) is a polynomial in the components of $\tilde{x}$ with the partial derivatives as coefficients.

Consequently, its expected value can be computed from the moments

$$\mathrm{E}\{\tilde{x}_i\},\mathrm{E}\{\tilde{x}_i\tilde{x}_j\},\mathrm{E}\{\tilde{x}_i\tilde{x}_j\tilde{x}_k\},\ldots \tag{43}$$

for all combinations of $i = 1,\ldots,n$; $j = 1,\ldots,n$; $k = 1,\ldots,n$; and so on. The existence of all moments is guaranteed in the Gaussian case. Moreover, all odd moments are zero due to symmetry, for example the first and third item in (43). The even moments can be computed from the characteristic function [49]. For example, the second and fourth moments,

$$\mathrm{E}\{\tilde{x}_i\tilde{x}_j\} = P_{ij}, \tag{44}$$

$$\mathrm{E}\{\tilde{x}_i\tilde{x}_j\tilde{x}_k\tilde{x}_l\} = P_{ij}P_{kl} + P_{ik}P_{jl} + P_{il}P_{jk}, \tag{45}$$

can be expressed in terms of the scalar entries of the covariance matrix P.

We have established that (22a) can, in principle, be computed if the required partial derivatives and moments are available. For polynomial functions, we can compute the exact mean. For arbitrary differentiable functions, we can compute the moments of their Taylor polynomial of degree d. A compact formula that includes all the moments of a scalar x is given in [52].

The next step towards computing (22b) involves an outer product of the infinite series (42). The monomials of a specific degree are no longer grouped as in (42), which renders the approach impractical. A reformulation of the Taylor series can improve upon this situation as we show below.

2) A compact Taylor series notation:

The moment computation problem (22) is easily solved for affine functions $Ax + b$:

$$\mathrm{E}\{Ax + b\} = A\mathrm{E}\{x\} + b, \tag{46a}$$

$$\mathrm{cov}\{Ax + b\} = A\mathrm{cov}\{x\}A^{\mathrm{T}}, \tag{46b}$$

$$\mathrm{cov}\{x, Ax + b\} = \mathrm{cov}\{x\}A^{\mathrm{T}}. \tag{46c}$$

Motivated by the above simplicity, we now turn the nonlinear problem (22) into an infinite dimensional linear problem.

A related technique that uses the Kronecker product to reshape (42) is described in [53]. Higher moments of the Gaussian distribution and the Kronecker product are explored in [54]. Derivative operators in conjunction with the Kronecker product, as used below, are also treated in [55]. The formulas below have been developed in [30], where a more detailed account can be found. A related discussion with a KF background is given in [56].

The tools that facilitate a compact expression for (42) are the d-fold Kronecker product [53]

$$a^{\otimes d} = a \otimes \cdots \otimes a, \tag{47a}$$

the derivative operator

$$D = \left[\frac{\partial}{\partial x_1},\ldots,\frac{\partial}{\partial x_n}\right] \tag{47b}$$

that acts on each row of a function, and a combination of the two

$$D^{\otimes d} = D \otimes \cdots \otimes D \tag{47c}$$

that gives all partial derivatives of order d when applied to a function f.

Using the above, the infinite series in (42) can be written as the linear relation

$$f(x) = f(\hat{x}) + \mathbf{F}(\hat{x})\tilde{\mathbf{x}}(\tilde{x}) \tag{48a}$$

in the infinite dimensional coefficient matrix and monomial vector

$$\mathbf{F}(\hat{x}) = [Df(\hat{x}), \quad (D^{\otimes 2})f(\hat{x}), \quad (D^{\otimes 3})f(\hat{x}), \quad \ldots] \tag{48b}$$

$$\tilde{\mathbf{x}}(\tilde{x}) = [\tilde{x}^{\mathrm{T}}, \quad \tfrac{1}{2}(\tilde{x}^{\otimes 2})^{\mathrm{T}}, \quad \tfrac{1}{3!}(\tilde{x}^{\otimes 3})^{\mathrm{T}}, \quad \ldots]^{\mathrm{T}} \tag{48c}$$

The expression (48a) mimics $Ax + b$. The randomness has been concentrated in $\tilde{\mathbf{x}}(\tilde{x})$, while $\mathbf{F}(\hat{x})$ is fully deterministic. Because of the structure in f, $\mathbf{F}(\hat{x})$ is typically sparse because not all rows of f depend on all entries in x. The $1/d!$ factors of the degree d terms can be either assigned to (48b) or (48c). The latter option turns out to be more convenient.

Now, the moment computation problem can be formulated as

$$\mathrm{E}\{z\} = f(\hat{x}) + \mathbf{F}(\hat{x})\mathrm{E}\{\tilde{\mathbf{x}}(\tilde{x})\}, \tag{49a}$$

$$\mathrm{cov}\{z\} = \mathbf{F}(\hat{x})\mathrm{cov}\{\tilde{\mathbf{x}}(\tilde{x})\}\mathbf{F}(\hat{x})^{\mathrm{T}}, \tag{49b}$$

$$\mathrm{cov}\{x,z\} = \mathrm{cov}\{x,\tilde{\mathbf{x}}(\tilde{x})\}\mathbf{F}(\hat{x})^{\mathrm{T}}. \tag{49c}$$

Apparently, we have replaced (22) by the simpler problem of finding the moments of $\tilde{\mathbf{x}}(\tilde{x})$. From the symmetry in the distribution of $\tilde{x}$ follows a "checkerboard pattern" in

$$\mathrm{E}\{\tilde{\mathbf{x}}(\tilde{x})\} = \\ [0^{\mathrm{T}}, \quad \tfrac{1}{2}\mathrm{E}\{\tilde{x}^{\otimes 2}\}^{\mathrm{T}}, \quad 0^{\mathrm{T}}, \quad \tfrac{1}{4!}\mathrm{E}\{\tilde{x}^{\otimes 4}\}^{\mathrm{T}}, \quad \ldots]^{\mathrm{T}} \tag{50a}$$

and

$$\mathrm{cov}\{\tilde{\mathbf{x}}(\tilde{x})\} = \\ \begin{bmatrix} P & 0 & \frac{1}{3!}\mathrm{cov}\{\tilde{x},\tilde{x}^{\otimes 3}\} & \\ 0 & \frac{1}{2^2}\mathrm{cov}\{\tilde{x}^{\otimes 2}\} & 0 & \cdots \\ \frac{1}{3!}\mathrm{cov}\{\tilde{x}^{\otimes 3},\tilde{x}\} & 0 & \frac{1}{(3!)^2}\mathrm{cov}\{\tilde{x}^{\otimes 3}\} & \\ & \vdots & & \ddots \end{bmatrix}. \tag{50b}$$

The cross-covariance $\mathrm{cov}\{x,\tilde{\mathbf{x}}(\tilde{x})\}$ is the first row of the matrix in (50b). A remaining challenge is the sheer number of terms in the expressions. However, each entry can be obtained from the moments of the Gaussian distribution [49]. These moment expressions simplify significantly if the stochastic decoupling technique is

applied. Then, however, the inherent sparsity that might be present in $\mathbf{F}(\hat{x})$ typically disappears.

Beyond moment computation, the above results are useful for assessing the nonlinearity of f in the neighborhood of an expansion point $\hat{x}$. If the error between the degree d Taylor approximation of f and f is small for carefully chosen x, then f is well represented by a polynomial. In Sec. V-B, Taylor polynomials up to degree 5 are used to assess a range rate measurement in terms of its nonlinearity.

3) Quadratic functions as special case:

We here investigate the special case of a specific $z = f(x)$ with m rows that are given by

$$z_l = a_l + b_l^{\mathrm{T}} x + x^{\mathrm{T}} C_l x. \tag{51}$$

The a_l are scalars, the b_l are n-vectors, and the C_l are $n \times n$-matrices. With the Jacobian and Hessian matrices

$$J_l(x) = b_l^{\mathrm{T}} + x^{\mathrm{T}}(C_l + C_l^{\mathrm{T}}), \quad H_l = (C_l + C_l^{\mathrm{T}}) \tag{52}$$

that contain the first and second order partial derivatives, respectively, (51) can be written as

$$z_l = f_l(\hat{x}) + J_l(\hat{x})(x - \hat{x}) + \tfrac{1}{2}(x - \hat{x})^{\mathrm{T}} H_l (x - \hat{x}). \tag{53}$$

Because f is a polynomial of degree 2, (53) is exact everywhere.

The moments of z can be derived using the tools from the previous subsection together with (45). The mean and covariance are determined by [1]

$$\mathrm{E}\{z_l\} = f_l(\hat{x}) + \tfrac{1}{2}\mathrm{tr}(H_l P), \tag{54a}$$

$$\mathrm{cov}\{z_l, z_k\} = J_l(\hat{x}) P J_k(\hat{x})^{\mathrm{T}} + \tfrac{1}{2}\mathrm{tr}(H_l P H_k P), \tag{54b}$$

$$\mathrm{cov}\{x, z\} = P J(\hat{x})^{\mathrm{T}}, \tag{54c}$$

with $k = 1,\ldots,m$, and $l = 1,\ldots,m$. These expressions can be implemented entirely sampling based as shown in Sec. IV-D.

IV. SIGMA POINT METHODS FOR SOLVING THE MOMENT COMPUTATION PROBLEM

In this section we discuss methods for solving the moment computation problem (22) of Sec. III. The discussion is focused on approximate techniques that are all based on sampling. Apart from a brief coverage of Monte Carlo integration, deterministic methods are presented. The interpretation of the sigma point differs considerably in the unscented transformation, numerical integration, and interpolation approaches.

For convenience, we introduce the abbreviations

$$\hat{z} \approx \mathrm{E}\{z\}, \quad S \approx \mathrm{cov}\{z\}, \quad M \approx \mathrm{cov}\{x, z\} \tag{55}$$

for the moment approximations of the transformed random variable $z = f(x)$.

A. Monte Carlo Integration

This section describes Monte Carlo integration in its most basic form. It is the only *stochastic* moment computation method that we discuss in detail. For a more thorough treatment of Monte Carlo methods in general the reader is referred to [57] and the dedicated chapters in [34, 58].

The basis of Monte Carlo integration is sampling. Hence, it is applicable to the moment computation problem (22) with arbitrary distributions of the random variable x, as long as they allow sampling. Specifically, a large number N of realizations $\{x^{(i)}\}_{i=1}^{N}$ is generated, passed through the function f to yield $\{z^{(i)} = f(x^{(i)})\}_{i=1}^{N}$, and then used to approximate $\mathrm{E}\{z\}$ in (22a) by the sample average

$$\hat{z} = \frac{1}{N}\sum_{i=1}^{N} z^{(i)}. \tag{56a}$$

In a similar manner, the covariance matrices $\mathrm{cov}\{z\}$ and $\mathrm{cov}\{x, z\}$ of (22b) and (22c) are approximated by

$$S = \frac{1}{N-1}\sum_{i=1}^{N} (z^{(i)} - \hat{z})(z^{(i)} - \hat{z})^{\mathrm{T}}, \tag{56b}$$

$$M = \frac{1}{N-1}\sum_{i=1}^{N} (x^{(i)} - \hat{x})(z^{(i)} - \hat{z})^{\mathrm{T}}. \tag{56c}$$

Justification for the above expressions is given by the law of large numbers which states that the averages in (56) indeed converge to the true moments as N tends to infinity. Another important point is that the accuracy of (56a), expressed in terms of its covariance

$$\mathrm{cov}\{\hat{z}\} = \frac{1}{N}\mathrm{cov}\{z\}, \tag{57}$$

depends solely on the true covariance, not the dimension n of x [58]. As a result some problems can be solved with very few samples N, although n is big, whereas others cannot be solved by Monte Carlo integration at all. The standard deviation of a scalar $\hat{z}$ decays as $1/\sqrt{N}$ which is rather slow.

Monte Carlo methods impose no restriction on the functional form of f. If the true moments exist, then (56) will work for sufficiently large N. Due to the random sampling, there is always some variation in the results of (56). These variations decrease with increased sample size N, and can be reduced further by employing variance reduction techniques [57]. If it is difficult to generate samples of x directly, the concept of importance sampling can be used, which actually builds the basis for the particle filter [12, 13]. Even more advanced sampling schemes are Markov chain Monte Carlo methods [57].

The idea to use Monte Carlo integration in a nonlinear Kalman filter is mentioned in [23, 40]. It could be argued that the required computational complexity to run a Monte Carlo Kalman filter contradicts the idea of maintaining only mean values and covariance matrices. Interestingly, a related KF variant beyond the scope of this paper, the ensemble Kalman filter [59], uses Monte Carlo sampling to mimic the KF algorithm in

high-dimensional state spaces that prohibit the storage of $n \times n$ covariance matrices.

B. The Sigma Point Approximation of a Distribution and the Unscented Transformation

The ideas of the following paragraphs were introduced in the filtering context in [14], together with the term sigma points. Later, the somewhat mysterious name *unscented* transformation (UT) was given to the moment computation method, and the resulting filter became known as the unscented Kalman filter (UKF) [16].

As starting point, we can assume the moment computation problem (22) for a discrete random variable x that takes values from the set $\{x^{(i)}\}_{i=1}^{N}$ with the probabilities $\Pr\{x^{(i)}\} = w^{(i)}$. The integrals in (22) then appear as finite sums. The resulting simplicity therefore justifies the idea to approximate the continuous density $p(x)$ of (22) with a discrete point mass function or Dirac mixture density [60].

The UT [15] implements the above idea by systematically selecting N sigma points and weights $\{x^{(i)}, w^{(i)}\}_{i=1}^{N}$. As in MC integration, the function f is evaluated at every sigma point to yield $\{z^{(i)}\}_{i=1}^{N}$. The moment integrals (22) are approximated by

$$\hat{z} = \sum_{i=1}^{N} w^{(i)} z^{(i)}, \tag{58a}$$

$$S = \sum_{i=1}^{N} w^{(i)} (z^{(i)} - \hat{z})(z^{(i)} - \hat{z})^{\mathrm{T}}, \tag{58b}$$

$$M = \sum_{i=1}^{N} w^{(i)} (x^{(i)} - \hat{x})(z^{(i)} - \hat{z})^{\mathrm{T}}. \tag{58c}$$

The expressions resemble a weighted version of (56). However, the number of sigma points is typically much smaller than in Monte Carlo sampling with $N = 2n + 1$ in the most common variant [14].

1) Sigma point and weight selection:

The problem of approximating a probability density function by a number of representative samples goes beyond the moment computation context. For example, the quantization problem [33, 61] in information theory is closely related. For computational reasons, the number of points should be small. In order to approximate the continuous density well, however, more points would be valuable.

The UT sigma points are chosen deterministically such that the mean and covariance of x are preserved in the weighted sample mean and covariance

$$\sum_{i=1}^{N} w^{(i)} x^{(i)} = \hat{x}, \tag{59a}$$

$$\sum_{i=1}^{N} w^{(i)} (x^{(i)} - \hat{x})(x^{(i)} - \hat{x})^{\mathrm{T}} = P. \tag{59b}$$

The condition (59) guarantees that (58) is correct for affine functions. Also constraints on higher moments can, in principle, be included [16]. For any fixed number of points $N > n$, infinitely many combinations of weights and samples satisfy (59).

The original sampling scheme [15] utilizes the columns $\{u_i\}_{i=1}^{n}$ (78) of the $n \times n$ matrix square root $P^{1/2}$ (77) (App. A) to generate $N = 2n + 1$ sigma points and weights

$$x^{(0)} = \hat{x}, \tag{60a}$$

$$x^{(\pm i)} = \hat{x} \pm \sqrt{(n + \kappa)} u_i, \tag{60b}$$

$$w^{(0)} = \frac{\kappa}{n + \kappa}, \tag{60c}$$

$$w^{(\pm i)} = \frac{1}{2(n + \kappa)}, \tag{60d}$$

where $i = 1, \ldots, n$. The signed superscripts $(\pm i)$ underline the symmetry in the sigma points. Accordingly, the summation in (58) is replaced by a sum from $-n$ to n. The alternative expression for (58a)

$$\hat{z} = w^{(0)} z^{(0)} + \sum_{i=1}^{n} w^{(\pm i)} (z^{(+i)} + z^{(-i)}) \tag{61}$$

establishes a relation to the divided difference methods in Sec. IV-D.

The authors of [15] suggest that κ can be chosen as any number such that $\kappa \neq -n$ and recommend $\kappa = 3 - n$. The square root in (60b), however, suggests that $\kappa > -n$. Furthermore, the obtained negative $w^{(i)}$ for $0 > \kappa > -n$ compromise the interpretation as point mass function and can lead to indefinite S in (58b). For $\kappa = 0$, the UT is identical to a cubature rule (Sec. IV-C) that is used in [18].

For large n, the sigma points in (60) are located far from the mean $\hat{x}$, which might yield degraded results for functions f that change severely over the x-space. The scaled UT [62] addresses this by moving the samples towards $\hat{x}$ without violating (59). The procedure has been condensed into a selection scheme similar to (60) by [21].

In the selection scheme, the parameter κ is replaced by $\lambda = \alpha^2(n + \kappa) - n$ with $0 < \alpha \leq 1$. The samples and weights for (58a) or (61) are

$$x^{(0)} = \hat{x}, \tag{62a}$$

$$\begin{aligned} x^{(\pm i)} &= \hat{x} \pm \sqrt{n + \lambda} u_i \\ &= \hat{x} \pm \alpha \sqrt{n + \kappa} u_i, \end{aligned} \tag{62b}$$

$$w^{(0)} = \frac{\lambda}{n + \lambda}, \tag{62c}$$

$$w^{(\pm i)} = \frac{1}{2(n + \lambda)}, \tag{62d}$$

with $i = 1, \ldots, n$. The covariance (58b) is computed with the weight $w_c^{(0)} = (\lambda/(n + \lambda)) + 1 - \alpha^2 + \beta$ for the $z^{(0)}$ term instead. Here, the scalar parameter β is a correction

term [16] which is recommended to be 2 for Gaussian x. With some effort, the motivation for the above choice can be extracted from [16] or [21]. The latter recommends that the choice of α should depend on f as well as $\mathrm{cov}\{x\} = P$, whereas the choice of κ is secondary.

Alternative sigma point generation schemes beyond moment matching were investigated in [60]. Instead of a simple generation rule, the points are chosen to minimize a cost function. The procedure is more costly but allows the user to choose an arbitrary number of sigma points N. In the spirit of the stochastic decoupling technique (Sec. III-B) the authors of [39] employ the method of [60] to generate a sigma point approximation to $\mathcal{N}(0, I)$ offline. These carefully selected samples are then transformed to represent $x \sim \mathcal{N}(\hat{x}, P)$.

2) Accuracy considerations:

The UT is built upon the idea that a good approximation of the distribution of x should also give a good approximation of the distribution of z. Its implementation is simple and requires hardly more than evaluating f for all sigma points. Albeit the intuition, accuracy statements are difficult to make. After all, how the characteristics of the weighted sigma points in x carry over to z is highly dependent on the function f.

In an attempt to show the UT accuracy, the authors of [15] expand the covariance (22b) as infinite Taylor series and compare the terms of low degree to the solutions provided by linearization and UT. It is argued that the former gives only the first term in the infinite series, whereas the UT does not truncate. So, more terms of low degree in the infinite Taylor expansion of (22b) are matched by UT. This observation has lead to the easily misinterpreted statement that the UT is "correct up to the second order" [15, App. II], but it is not clear whether this "order" refers to degree of the Taylor polynomial of (22b) or to the moments of the transformed samples. Certainly, the statement has led to misunderstandings.

To be clear, the UT does not give the correct first and second[6] moments (22a) and (22b) for arbitrary nonlinearities. Furthermore, the UT does not give the moments of a degree two Taylor approximation of f (Sec. III-C). This is discussed and demonstrated for a simple example in [23], and further clarified in the simulation study of polynomial function in Sec. V-A. Another extensive discussion of the statements in [15] is given in [21]. Although the UT often works well, concise accuracy statements are difficult to make. The numerical integration perspective in Sec. IV-C alleviates this to some extent.

3) Unscented and related Kalman filters:

The first UKF was suggested in [14] but the ideas took several years before they appeared in journal format [15]. A complete account of the developers is given in [16]. An interpretation of the UKF as performing statistical linearization is discussed in [63]. The UKF as one member of the sigma point KF class is discussed in [21], including its use for parameter estimation and an account of the accuracy statements in [15]. The relations between the UKF and divided difference filters [17, 20] are investigated in [27], and the authors of [23] discuss asymptotic relations to analytical EKF. An extensive account of the numerical integration perspective of UKF is given in [26]. The authors of [44] investigate the scaling parameter in UKF variants with the sample set (60) and devise an adaptive scaling method for κ in an online filter. The recent smart sampling Kalman filter [39] makes use of the alternative sampling method [60] and allows the user to select an arbitrary number of sigma points N. References to comparative simulation studies are listed in Sec. V.

C. Sigma Points and Numerical Integration

The discussion below results in methods that appear very similar to the UT and, in fact, share the formulas (58).

As starting point, we consider the problem of computing the integral of an exponentially weighted function $g(t)$

$$\int \cdots \int g(t) \exp(-t^{\mathrm{T}} t) dt_1 \ldots dt_n. \tag{63}$$

The relation to the moment computation problem (22) with Gaussian x is established using the stochastic decoupling technique (Sec. III-B). Specifically, the conversion of (22a) to (63) is given in (24).

The problem of computing (63) has been addressed by mathematicians for centuries. The suggested solutions, for example Gauss-Hermite quadrature rules [48], use formulas that are similar to (58a):

$$\int g(t) \exp(-t^{\mathrm{T}} t) dt \approx \sum_{i=1}^{N} w^{(i)} g(t^{(i)}). \tag{64}$$

The classical numerical integration literature [64, 65] provides rules to select the N weights $w^{(i)}$ and sigma points $t^{(i)}$ such that (64) is exact for polynomial functions of degree[7] d. The number of required points N increases as d increases. The built-in accuracy statement is an advantage over the UT. However, by showing the equivalence between an integration rule and a UT variant, the accuracy statements can be transferred to the latter.

For the moment computation problem (22) a rule of degree d gives the correct mean value (22a) for polynomial f of degree d. The covariance (22b), however, requires integration of outer products of f. Therefore, an integration rule of degree d computes (22b) correctly only if f has at most degree $d/2$. Furthermore, the cross-covariance (22c) is computed correctly if f has at most

[6] As stated earlier, we reserve "order" to describe the order of a Taylor series of partial derivative to avoid confusion.

[7] For a vector or matrix function g, the degree is the highest appearing monomial degree. For example, $g(t) = [t_1^3 t_2^2 - 3, t_2]^{\mathrm{T}}$ has degree 5.

degree $d-1$. The above considerations are confirmed by the simulation results in Sec. V-A.

Integration rules can be grouped into product and nonproduct rules [64] that are further explained below.

1) Product rules:

Product rules compute the n-fold integral (63) by applying integration rules to one dimension at a time. For example, the integration with respect to t_1 can be approximated by an M_1-point rule. The result is a sum of M_1 integrals with respect to the remaining components $t_{2:n}$

$$\sum_{i_1=1}^{M_1} w_1^{(i_1)} \int g(t_1^{(i_1)}, t_{2:n}) \exp\left(-\sum_{j=2}^{n} t_j^2\right) dt_{2:n}. \tag{65}$$

Continuing in this way requires $N = \prod_{i=1}^{n} M_i$ points and weights in total, which renders the method unfeasible for all but small n. The overall degree of such an n-dimensional product rule is determined by the smallest M_i. The Kronecker product can be used to compute the overall weights and samples [66].

Gauss-Hermite quadrature rules [48] can be used to generate the M_i weights and points for the one-dimensional integrals from the Hermite polynomials. An M-point Gauss-Hermite rule integrates polynomials of degree $d = 2M-1$ exactly.

In a Kalman filter context, product rules were investigated first in [17] and later in [26, 41]. As pointed out, the required number of points is problematic for larger n. A technique to divide the filtering problem into smaller chunks that can be approached by product rules is given in [67]. The simulations in Sec. V include 3-point and 5-point (per dimension) quadrature rules.

2) Nonproduct rules:

In contrast to product rules, nonproduct rules are tailored formulas to integrate monomials[8] of at most degree d exactly. The number of required point scales much better with n.

An example from the filtering literature is the rule of degree 3 that is used in the cubature Kalman filter (CKF) of [18]. It requires $N = 2n$ weights and sigma points that are given by

$$t^{(\pm i)} = \pm\sqrt{\frac{n}{2}} e_i, \tag{66a}$$

$$w^{(\pm i)} = \frac{\pi^{n/2}}{2n}, \tag{66b}$$

with $i = 1,\ldots,n$ and the ith column e_i of the identity matrix. The selection scheme (66) can be found in as early references as [64, 65]. Using the variable substitution (24) of Sec. III-B the $t^{(i)}$ can be mapped to

$$x^{(\pm i)} = \hat{x} \pm \sqrt{2} P^{1/2} t^{(\pm i)} = \hat{x} \pm \sqrt{n} u_i. \tag{67a}$$

Furthermore, the weighted transformed samples can be written as

$$w^{(\pm i)} g(t^{(\pm i)}) = \frac{\pi^{n/2}}{2n} \frac{1}{\pi^{n/2}} f(x^{(\pm i)}) = \frac{1}{2n} z^{(\pm i)}. \tag{67b}$$

The above (67) discloses that the degree 3 rule corresponds to the UT (60) with $\kappa = 0$. Therefore, $\kappa = 0$ can be regarded as well founded parameter choice in the UT.

The authors of [18] show how to use a change to spherical coordinates (Sec. III-B) and subsequent separation of the integral (63) into a range and a spherical component to arrive at the above expressions. The authors of [26] assume a numerical integration perspective on the UT to arrive at a similar expression.

An extensive numerical integration source is [64], where the above degree-3 rule with $2n$ points, a degree-5 rule with $2n^2+1$ points, and a degree-7 rule with $(4n^3+8n+3)/3$ points are provided. Related Kalman filtering advances are described in [26], including an assessment of the stability factor that relates to numerical accuracy [65]. The recent high-degree cubature Kalman filters of [42] provide arbitrary degree rules developed from the integral representation (28). The number of required points grows polynomially with n, which yields an improvement over the exponentially growing number of points in Gauss-Hermite quadrature. Also, [42] highlights the degree-5 rule of [64] in their discussion and KF simulations.

3) Randomized integration rules:

We conclude this section with a numerical integration approach that combines Monte Carlo integration with deterministic integration rules.

In (28) of Sec. III-B it is shown how the integral (63) can be separated into a range and a spherical part. This variable change is exploited in [68] in a Monte Carlo scheme, and implemented within the Kalman filtering context by [69] as randomized UKF variant and [43] with the numerical integration perspective. The suggested stochastic integration filter [43] uses randomized rules for both the radial and spherical integration, and refines the results in an iterative scheme. Thereby, less sigma points are required compared to a direct Monte Carlo approach. The required number of points N depends on the user parameters, which are set to determine the degree of the radial and spherical integration rules. The approach is further discussed and successfully tested in [70].

An immediate variation of the above is to use a deterministic integration rule for the spherical and Monte Carlo integration for the range component. Such derived rules can be viewed as randomized deterministic integration rules.

D. Sigma Points for Function Approximation

This section summarizes how sigma points can be used to approximate the function f in the moment computation problem (22). Such local approximations

[8] and hence polynomials.

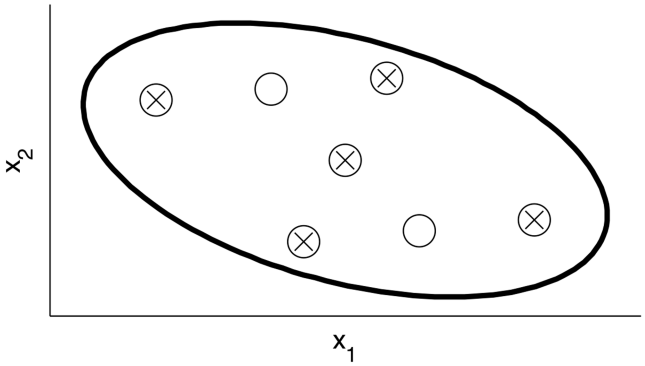

Fig. 2. The extended sigma points of (70) (circles) and the original samples of (68) (crosses) for $n = 2$. The ellipse indicates a contour for which $\mathcal{N}(x;\hat{x},P)$ is constant. The required matrix square root is computed using the singular value decomposition of P.

are in the spirit of linearization or Taylor approximation of f, but with divided differences rather than analytical derivatives.

1) Divided difference methods:

The solutions of the moment computation problem (22) are given by (46) in the case of linear f, and by (54) for quadratic f (51). If we consider a nonlinear f, we might want to mimic these solutions using a local Taylor expansion of f. The required partial derivatives are then evaluated at the expansion point only, and do not take into account the uncertainty in x. If, instead, the local approximation is carried out by an interpolation approach, the interpolation points can be chosen to also reflect the uncertainty in x.

A linearization based on the above idea was first suggested in [19] in the Kalman filtering context, and extended independently to include quadratic approximations by [20] and [17]. Both motivate their development from an interpolation point of view that replaces analytical derivatives with divided differences. The method generates sigma points similar to the UT

$$x^{(0)} = \hat{x}, \tag{68a}$$

$$x^{(\pm i)} = \hat{x} \pm \sqrt{\gamma} u_i, \tag{68b}$$

where $i = 1,\ldots,n$. The u_i are the columns (78) of the square root $P^{1/2}$ (77) (see App. A) and act as perturbation points in the divided differences. The parameter γ determines the length of the interpolation interval.

The sigma points are transformed to yield $z^{(0)}$ and all $z^{(\pm i)}$ and processed as

$$\hat{z} = \frac{\gamma - n}{\gamma} z^{(0)} + \frac{1}{2\gamma} \sum_{i=1}^{n} (z^{(i)} + z^{(-i)}), \tag{69a}$$

$$S = \frac{1}{4\gamma} \sum_{i=1}^{n} (z^{(i)} - z^{(-i)})(z^{(i)} - z^{(-i)})^{\mathrm{T}} + \frac{\gamma - 1}{4\gamma^2} \sum_{i=1}^{n} (z^{(i)} - 2z^{(0)} + z^{(-i)})(z^{(i)} - 2z^{(0)} + z^{(-i)})^{\mathrm{T}}, \tag{69b}$$

$$M = \frac{1}{2\sqrt{\gamma}} \sum_{i=1}^{n} u_i (z^{(i)} - z^{(-i)})^{\mathrm{T}}. \tag{69c}$$

The expressions sacrifice some inconvenient terms for the sake of a faster algorithm. Further details on the omitted terms are provided in [71]. The second order divided difference method from [20] is illustrated in (69). A first order divided difference method [19] is obtained by retaining only the first sum in (69b), and setting $\hat{z} = z^{(0)}$ similar to the Taylor based linearization.

The expression (69a) has the same functional form as the UT mean in (61). A detailed discussion on the relations between divided difference methods and the UT is given in [27].

2) Sigma Point Implementation for the Moments of a Quadratic Function:

The exact solution of the moment computation problem for quadratic f is given by (54). The divided difference solution (69) is an approximation to this in general. If f is indeed quadratic then (69a) and (69c) are exact, but (69b) is not. We here show how an exact S can be computed entirely from sigma points. The material is based on [23] and the extensions in [30, 35].

The sigma points (68) are not sufficient to make S exact. Therefore, the set is extended to include n^2 points

$$x^{(0)} = \hat{x}, \tag{70a}$$

$$x^{(\pm ij)} = \hat{x} \pm \frac{\sqrt{\gamma}}{2}(u_i + u_j), \tag{70b}$$

with $i = 1,\ldots,n$ and $j = 1,\ldots,n$. By construction, the points of (68) are contained in (70): $x^{(\pm ii)} = x^{(\pm i)}$. In addition, $n(n-1)$ unique points are generated. The increased number of points N corresponds to that of a degree-5 integration rule [42, 64]. Fig. 2 illustrates the extended as well as the points of (68) for $n = 2$.

The exact S for quadratic f is given by

$$S = \frac{1}{4\gamma} \sum_{i=1}^{n} (z^{(i)} - z^{(-i)})(z^{(i)} - z^{(-i)})^{\mathrm{T}} + \frac{1}{8\gamma^2} \sum_{i=1}^{n} \sum_{j=1}^{n} (4z^{(ij)} + 4z^{(-ij)} - z^{(i)} - z^{(-i)} - z^{(j)} - z^{(-j)} - 4z^{(0)}) (4z^{(ij)} + 4z^{(-ij)} - z^{(i)} - z^{(-i)} - z^{(j)} - z^{(-j)} - 4z^{(0)})^{\mathrm{T}}. \tag{71}$$

The derivation of the above expression is provided in App. C.

In the Kalman filter context the above leads to a sample implementation of the second order EKF [35].

V. SIMULATION EXPERIMENTS

In this section we provide simulations to illustrate the performance of the moment computation methods of Sec. IV. Insights about differences and advantages are highlighted. Furthermore, it is shown that the methods do not always work as out-of-the-box solutions to any moment computation problem (22).

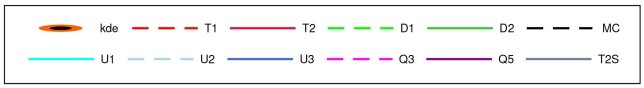

Fig. 3. The legend for all subsequent figures.

TABLE I
List of evaluated methods, their abbreviation, and a reference motivating the selected parameters.

Abbrev.	Method	Parameters
MC	Monte Carlo	$N = 50000$
T1	1st order Taylor [23]	—
T2	2nd order Taylor [23]	—
D1	1st order div. diff. [20]	$\gamma = n$
D2	2nd order div. diff. [20]	$\gamma = n$
U1	Unscented [14]	$\alpha = 1$; $\beta = 0$; $\kappa = 3 - n$
U2	Unscented [21]	$\alpha = 10^{-3}$; $\beta = 2$; $\kappa = 0$
U3	Cubature [18]	$\alpha = 1$; $\beta = 0$; $\kappa = 0$
Q3	3-point Gauss-Hermite [17, 48]	—
Q5	5-point Gauss-Hermite [17, 48]	—
T2S	2nd order Taylor, sigma points [35]	$\gamma = 0.1$

The examples include polynomial functions as well as a radar target tracking problem that occurs in air traffic control. The displayed methods are chosen in favor of a compact description that aims at both new and experienced researchers in the field. Hence, simple rules are favored over the more advanced methods that are mentioned in Sec. IV. The reader is referred to [23, 28, 30, 42] for other comparative simulation studies of moment computation methods in the spirit of our discussion below, and to [21, 26, 27, 29, 70] for comparative studies of the related filtering algorithms.

The investigated moment computation methods are first and second order Taylor approximation using analytical derivatives (T1 and T2, Sec. III-C); first and second order divided difference approximations (D1 and D2, Sec. IV-D.1) and the related second order Taylor approximation using sigma points (T2S, Sec. IV-D.2); Monte Carlo integration (MC, Sec. IV-A); three variations of the unscented transformation with different parameter settings including a degree-3 cubature rule (U1–U3, Sec. IV-B); and 3-point and 5-point (per dimension) Gauss-Hermite quadrature rules (Q3 and Q5, Sec. IV-C) which are more accurate than UT at the expense of more sigma points (3^n and 5^n). Table I summarizes the abbreviations used and the parameters according to Sec. IV. The plot legend for all subsequent illustrations is given in Fig. 3.

A. Polynomial functions

First, two quadratic functions are used to study the moment computation problem (22) with Gaussian input. In this case Sec. III-C describes the exact analytical solution T2 that can be compared to the approximate methods. Second, the composition of the two functions is studied, both in a two stage and single stage approach.

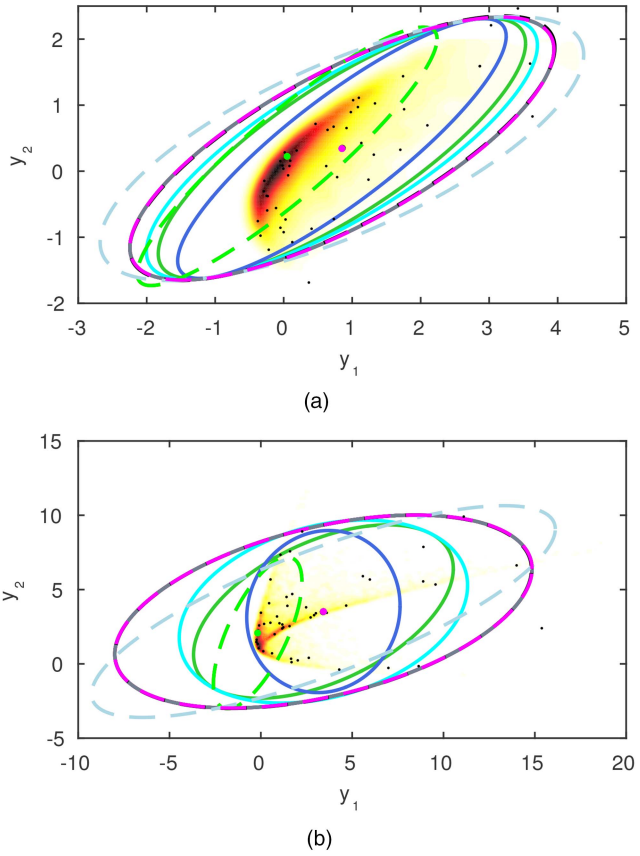

Fig. 4. The result obtained when using different methods to transforming x. The legend for this and subsequent figures is given in Figure 3. (a) $y = f^1(x)$. (b) $y = f^2(x)$.

The two two-dimensional quadratic functions f^1 and f^2 are specified as in (51) with the parameters

$$\begin{bmatrix} a_1^1 \\ a_2^1 \end{bmatrix} = \begin{bmatrix} -0.2 \\ 0 \end{bmatrix}, \quad [b_1^1 \quad b_2^1] = \begin{bmatrix} 0.1 & 0.2 \\ 0.2 & 0.3 \end{bmatrix},$$
$$C_1^1 = \begin{bmatrix} 0.15 & 0 \\ 0 & 0.05 \end{bmatrix}, \quad C_2^1 = \begin{bmatrix} 0.025 & 0.009 \\ 0.003 & 0.005 \end{bmatrix}, \tag{72a}$$

$$\begin{bmatrix} a_1^2 \\ a_2^2 \end{bmatrix} = \begin{bmatrix} 0 \\ 1.2 \end{bmatrix}, \quad [b_1^2 \quad b_2^2] = \begin{bmatrix} -0.8 & 0.4 \\ 0.8 & 0.4 \end{bmatrix},$$
$$C_1^2 = \begin{bmatrix} 0.625 & -0.303 \\ -0.303 & 0.275 \end{bmatrix}, \quad C_2^2 = \begin{bmatrix} 0.08 & 0 \\ 0 & 0.28 \end{bmatrix}, \tag{72b}$$

where the superscript is used to distinguish between the two functions. The input is assumed Gaussian $x \sim \mathcal{N}([0,1]^{\mathrm{T}}, 4I)$.

1) Quadratic Transformation:

The results from applying the different methods in Table I to compute the mean and covariance of $f^1(x)$ and $f^2(x)$ are presented in Fig. 4. Illustrated are the ellipses that contain 95 percent of the probability mass of a Gaussian with the computed mean and covariance. We notice that f^1 is less nonlinear than f^2, which

can be observed from the visualization of the correct target distribution which has been obtained as a kernel density estimate from MC samples. Both T2 and T2S by construction provide the exact mean and covariance. With f^1 and f^2 being polynomials of degree 2, the correct mean is recovered by all integration rules of degree 2 and above. This includes the UT variants U1, U2, U3 which can be seen as integration rules of degree 3 [26], but also the algebraically equivalent D2. The 3-point and 5-point (per dimension) quadrature rules with degrees $2 \cdot 3 - 1 = 5$ and $2 \cdot 5 - 1 = 9$, respectively, also manage to capture the correct covariance which requires an integration rule of degree 4 for f^1 and f^2. In contrast to the UT variants ($N = 2n = 4$ or $N = 2n + 1 = 5$), Q3 and Q5 use more sigma points ($N = 3^n = 9$ and $N = 5^n = 25$). Still, with $n = 2$ the computational demands remain low. The MC method with 50 000 samples is very close to the correct solution. Only T1 and D1 with their underlying linearity assumption fail to capture the correct mean and severely underestimate the covariance.

The methods perform mostly well for the quadratic functions, especially for the mildly nonlinear f^1 the approximations are similar, including T1 and D1. For f^2 the differences are more pronounced and the degree 3 rules U1, U3, and D2 do not reflect the true uncertainty ellipse well.

In summary, T2, T2S, Q3, and Q5 are exact for quadratic functions. The UT variants fail to capture the correct covariance with U2 slightly better for the chosen function f^2. D1 and T1 can perform very poorly, even though the nonlinearity is mild in the sense that it is only "one degree above linear." Other experiments with randomly selected quadratic functions support the above findings [30].

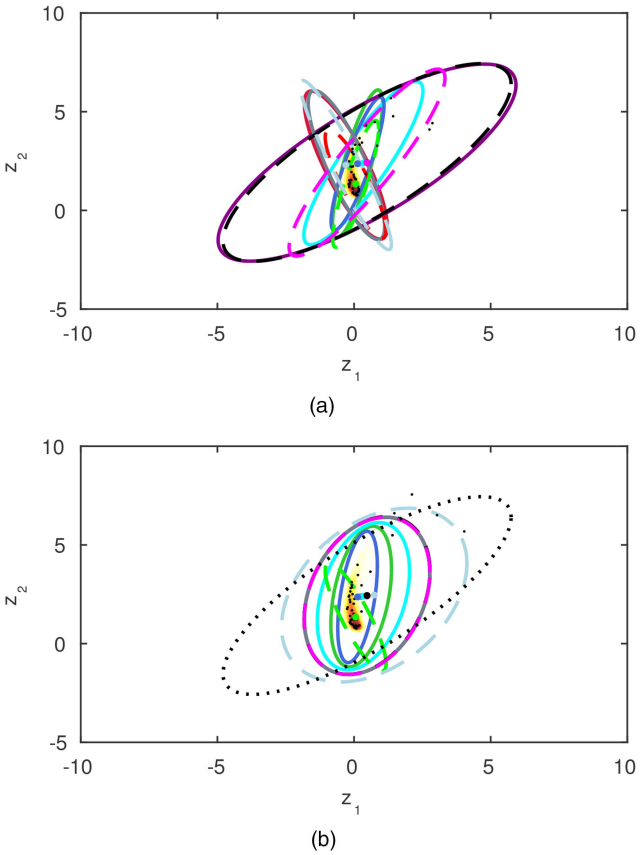

Fig. 5. The result obtained when applying the different approximations to a quartic function. Note that the MC solution is only correct in (a) as an intermediate approximation is performed in (b). Therefore the composed MC estimate has been added as dotted black line in the latter case. (a) Direct function evaluation, $z = f^3(x)$, corresponding to keeping the same sigma points. (b) Sequential function evaluation, $y = f^1(x)$ and $z = f^2(y)$, with intermediate sigma point regeneration.

2) Quartic Transformation:

In a second simulation a quartic function f^3 is constructed by composing f^2 and f^1, $f^3(x) = f^2(f^1(x))$. This allows us to study the selected methods on a higher degree polynomial than covered by most methods in Table I, as well as touching on the interesting questions: *"When considering a series of transformations, is it preferable to keep the same sigma points throughout the approximation, or should new sigma points be selected in each step?"* This has implications to the filtering problem where the time propagation step and measurement prediction are performed in series (Sec. II-D) and for cases where a function can be decomposed into simpler ones.

Fig. 5 shows the result obtained when applying the methods in Table I to the quartic function f^3. The difference between the result from applying the methods directly to the quartic function (Fig. 5(a)) and from performing the approximation in two steps (Fig. 5(b)) is striking. In the former case, only the Q5 rule with its degree 9 is accurate enough to capture the true mean and covariance. The MC method also recovers the true mean and covariance, albeit with some small sampling errors. With degree 5, the rule Q3 still provides the correct mean whereas the remaining methods fail to do so. In the later case, none of the methods (including the MC method) can be expected to recover the true mean and covariance, as a result of the intermediate step where the non-Gaussian distribution of $f^1(x)$ is approximated as a Gaussian, which results in information loss.

The two step computations result in comparable approximations in the sense that all methods underestimate the true covariance (as provided by MC and Q5 in Fig. 5(a) and as dotted line in Fig. 5(b)) and rotate its principal axes. T1 and D1 assert the smallest covariance and put the mean close to the most likely value of the resulting distribution. The similarity among the remaining approximations can be explained with the comparable Gaussian approximations of $f^1(x)$ that all methods provide (Fig. 4(a)). The lack of performance can be explained by the fact that all methods ignore the skewness in the distribution of $f^1(x)$. Note that MC, T2, T2S, Q3, and Q5 are all identical, as they all perfectly capture the two quadratic transformations, but that none of them are correct due to the intermediate approximation.

The result obtained from applying the approximations directly to f^3 yields less comparable results. D1, U1, and D2 provide roughly the same estimates, slightly worse than the sequential results. U2, T2, and T2S seriously underestimates the uncertainty and provide results far from the correct Q5 estimate. Q3 obtains the correct mean but underestimates the covariance. T1 and D1 no longer agree due to the differences between the numerical and analytical derivatives for the degree-4 polynomial.

Overall, we conclude that of the studied methods only Q5 and MC are suitable for the quartic function. Interestingly, there seems to be a gain from performing a two stage procedure where new sigma points are selected when using methods of insufficient accuracy, albeit the intermediate Gaussian approximation. Still, directly computing the moments of $f^3(x)$ with a rule of sufficient degree such as Q5 should be preferred if computationally feasible.

B. Tracking Application

In this second set of examples we consider moment computations problems that occur in target tracking, say, for air traffic control. In our set-up the tracked aircraft are described with a simple state-space model with the state vector

$$x = [\mathsf{x}, \quad \mathsf{y}, \quad \mathsf{s}, \quad \mathsf{h}]^{\mathrm{T}}, \tag{73}$$

where x and y make up the position in Cartesian coordinates [m], and s and h are the speed [m/s] and heading [rad], respectively.

A radar provides measurements $y = h(x) + e$ of the range, r, the bearing, ϕ, and the range rate, $\dot{\mathsf{r}}$, with

$$h(x) = \begin{bmatrix} \mathsf{r} \\ \phi \\ \dot{\mathsf{r}} \end{bmatrix} = \begin{bmatrix} \sqrt{\mathsf{x}^2 + \mathsf{y}^2} \\ \mathrm{atan2}(\mathsf{y}, \mathsf{x}) \\ \frac{\mathsf{s}}{\mathsf{r}}(\mathsf{x}\cos(\mathsf{h}) + \mathsf{y}\sin(\mathsf{h})) \end{bmatrix}, \tag{74}$$

where atan2 is the quadrant compensated arctangent function.

The dynamic model for how the state evolves in time is given by $x_k = f(x_{k-1}, v_{k-1})$, where

$$\begin{bmatrix} \mathsf{x}_k \\ \mathsf{y}_k \\ \mathsf{s}_k \\ \mathsf{h}_k \end{bmatrix} = \begin{bmatrix} \mathsf{x}_{k-1} + \cos(\mathsf{h}_{k-1})(\mathsf{s}_{k-1} + \Delta\mathsf{s}_{k-1}) \\ \mathsf{y}_{k-1} + \sin(\mathsf{h}_{k-1})(\mathsf{s}_{k-1} + \Delta\mathsf{s}_{k-1}) \\ \mathsf{s}_{k-1} + \Delta\mathsf{s}_{k-1} \\ \mathsf{h}_{k-1} + \Delta\mathsf{h}_{k-1} \end{bmatrix} \tag{75}$$

and $v = [\Delta\mathsf{s}, \Delta\mathsf{h}]^{\mathrm{T}}$ is process noise. Such models can be used to describe highly maneuvering targets, for example a flying robot, but also slowly maneuvering aircraft.

One time step and measurement will be studied, with

$$\begin{bmatrix} \mathsf{x}_{k-1} \\ \mathsf{y}_{k-1} \\ \mathsf{s}_{k-1} + \Delta\mathsf{s}_{k-1} \\ \mathsf{h}_{k-1} + \Delta\mathsf{h}_{k-1} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 2500 \\ 4330 \\ 250 \\ -\frac{2\pi}{3} \end{bmatrix}, \begin{bmatrix} 2500 & 0 & 0 & 0 \\ 0 & 2500 & 0 & 0 \\ 0 & 0 & 25 & 0 \\ 0 & 0 & 0 & \frac{900\pi^2}{180^2} \end{bmatrix} \right) \tag{76}$$

in the simulations below. This represents one step in a filter solution to track the target, but retains the focus on moment computation problems. For simplicity, the speed and heading uncertainty for the current time step is assumed to be included in the state already. This has no impact on the end result as they always appear together in (75). The position and speed accuracy has been exaggerated and the heading error is large to make the example more illustrative. Hence, the observed effects can be expected to be less pronounced in practice.

1) Range and Bearing Measurements:

The range and bearing measurements obtained by a radar are mildly nonlinear functions of the position, as shown in the heatmaps of Fig. 6. The state estimate is illustrated with an uncertainty ellipse that contains 95 percent of the probability mass. For the relevant positions, the heatmap appears similar to a linear function. The approximate linearity is confirmed by computing the moments of Taylor polynomials of different degree, as suggested in Sec. III-C. There is hardly any change for degrees above one.

Following the above considerations, it is not surprising to see that all the methods perform equally well. The result is illustrated in Fig. 7, where it is virtually impossible to tell the different methods apart.

2) Range Rate Measurement:

Adding range rate to the measurement vector provides important additional information about the target motion, but does at the same time add a more complex and nonlinear function that depends on all the state components. At this distance, the range rate is almost constant with respect to changes in the position, whereas Fig. 8 illustrates that the range rate is nonlinear (approximately sinusoidal) in the heading angle.

Again, the methods in Table I are used to obtain the mean and variance of the range rate. The result is given in Fig. 9. The methods T2, T2S, D2, U1, U2, U3, Q3, and Q5 are almost perfect matches with the MC method. T1 and D1 provide virtually identical results considerably underestimating the variance, and with a significant bias coinciding with the most likely values of the distribution. Thus, a reasonable choice of method would be U3 which uses the least sigma points.

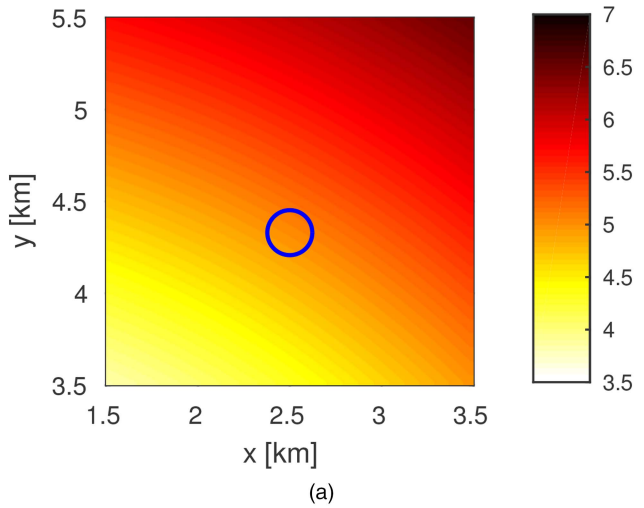

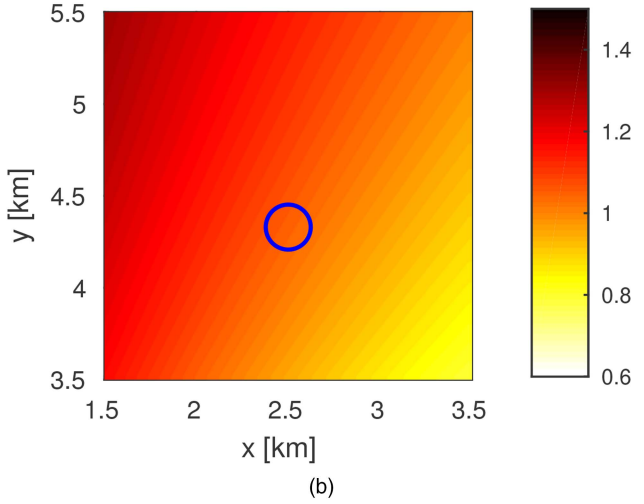

Fig. 6. Illustration of the severity of the radar measurement nonlinearity. Note, the range and bearing ((a) and (b)) are independent from the speed and heading. The blue ellipse represents the state uncertainty. (a) Range as function of target position. (b) Bearing as function of target position.

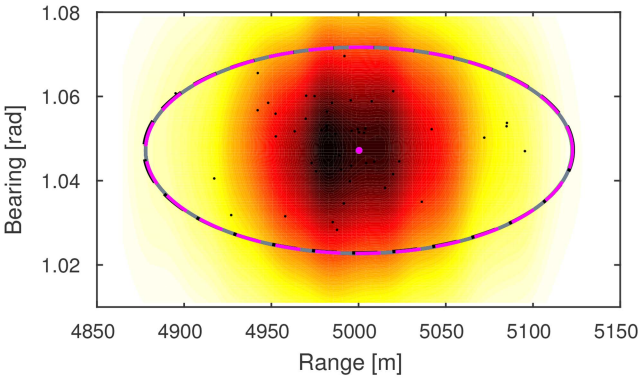

Fig. 7. Result obtained when applying the measurement function to x_{k-1} and visualizing the range and bearing components.

Table II shows the moments of different Taylor polynomials of the range rate that have been computed using the compact representation of Sec. III-C. It can be seen that the mean value changes for every two polynomial degrees added to the approximation. One

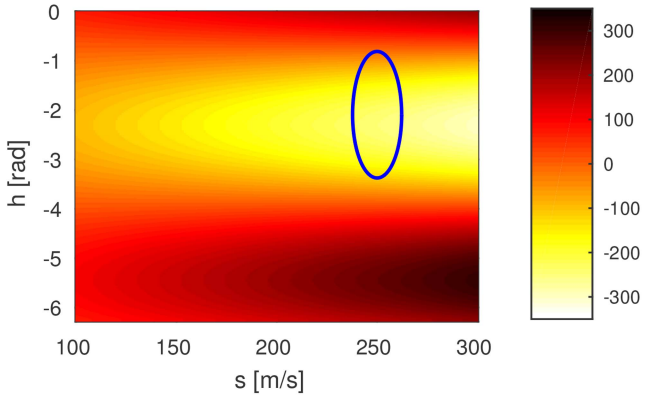

Fig. 8. Range rate as a function of speed and heading of the target when x and y have their mean values. The blue ellipse illustrates the state uncertainty.

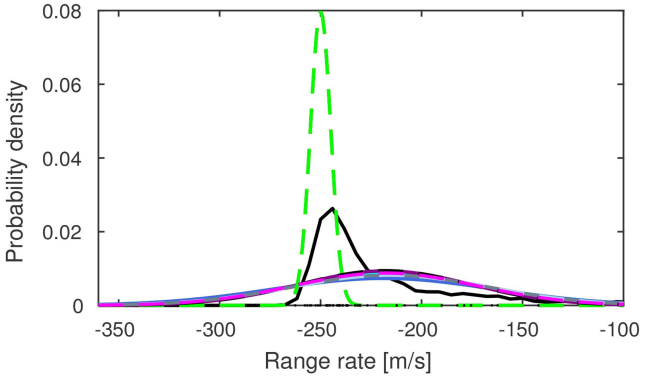

Fig. 9. Illustration of transformed range rate using the different approximations. The black solid line here visualize the true underlying function obtained with MC simulations.

TABLE II
Mean and variance matrix of the range rate based on Taylor expansions of different degrees.

Degree	Mean	Covariance
1	−250	25
2	−215.7	2375.5
3	−215.7	2370.1
4	−218	1784.4
5	−218	1784.5

interpretation of the result is that the range rate is difficult to approximate by a polynomial of low degree for the given input uncertainty.

3) Dynamic Model and Measurement:

In the linear KF it is common to alternate between time updates and measurement updates. As discussed in Sec. II-D, in nonlinear KF the situation is not as clear because the user can either combine the state transition and measurement functions to jointly compute the moments of the predicted state and output; or introduce an intermediate Gaussian approximation for the predicted state and then compute the moments of the predicted output. Also, see the related example in Sec. V-A.2.

Here, the two approaches are compared. The result is presented in Fig. 10. First, notice that the state pre-

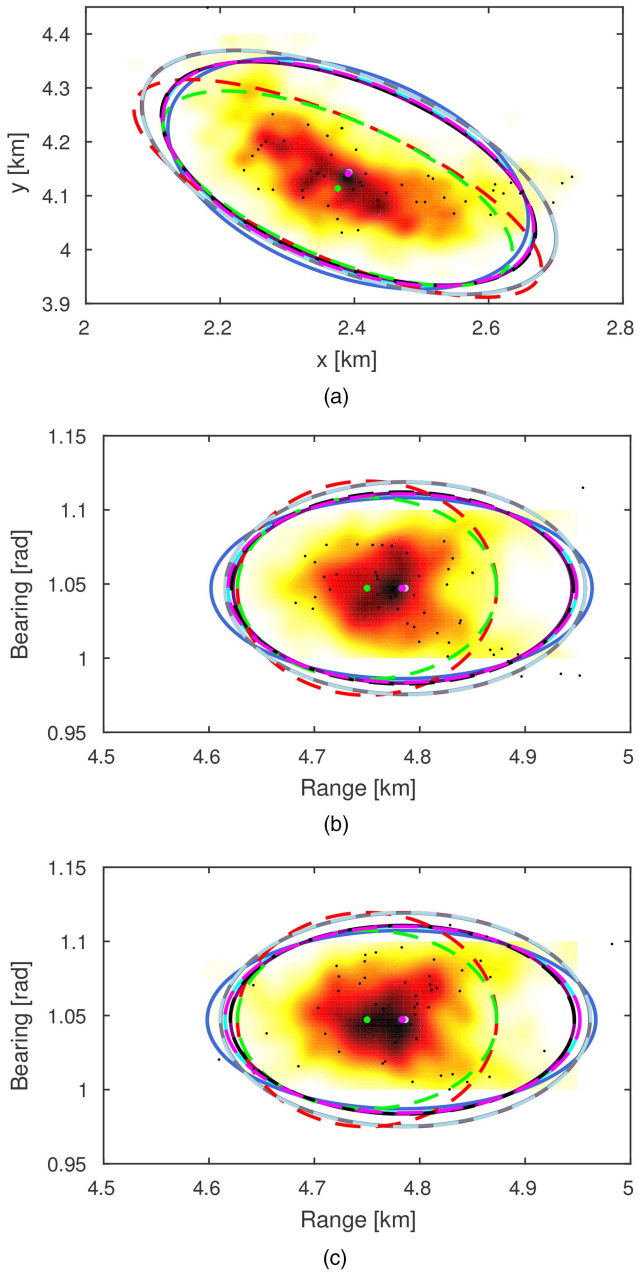

Fig. 10. Illustration of separated and composed transformation from x_{k-1} to y_k. Note that the MC solution is only correct in (c) as an intermediate approximation is performed in (b). Therefore the composed MC estimate has been added as dotted black line in the latter case. (a) Time propagation, $x_k = f(x_{k-1})$. (b) Sequential time propagation and measurement prediction, $y_k = h(x_k)$. (c) Combined time propagation and measurement prediction, $y_k = h(f(x_{k-1}))$.

diction results in a slightly banana shaped distribution (Fig. 10(a)), as a result of the uncertain heading but certain speed. Hence, the state transition is only mildly nonlinear. All methods, except for T1 and D1, produce mean and covariance estimates close to the true values. Again, T1 and D1 stand out by underestimating the covariance.

Comparing the result from performing the transformation in two steps (Fig. 10(b)) and in a single step (Fig. 10(c)) it is hard to tell the two approaches apart. In both cases the true distribution is well approximated by Q3, Q5, T2, T2S, D2, U1, U2, and U3, whereas T1 and D1 provide worse results with bias and too small covariances.

A conclusion from this experiment, supported by the experiences in Sec. V-A.2, is that it is not obvious whether to keep the sigma points between the two steps or to generate new ones. With the almost linear range and bearing measurements it made no difference. Furthermore, our results confirm the widespread view that T1 and D1 tend to introduce biases and to underestimate the covariance due to their underlying linearity assumptions.

VI. CONCLUDING REMARKS

In this tutorial we have discussed the nonlinear Kalman filter solution to the Bayesian filtering problem, with a focus on sigma point filters. It was shown that the central challenge in all Kalman filter algorithms is the computation of the mean values and the covariance matrices of nonlinearly transformed random variables. These moments are then used in the measurement update. This unifying view facilitates an easier access to the filtering literature for researchers who are new to the field, and provides additional (perhaps subtle) insights for the more experienced audience.

The underlying moment computation problem was discussed in greater detail including its structural properties, common reformulations, and even a solution in terms of the Taylor series. It is hoped that the presentation will facilitate a greater understanding of existing methods as well as catalyze new methods with a solid theoretical foundation.

The presentation of sigma point based moment computation methods which are motivated from the different concepts of density approximation (UT), numerical integration, and interpolation has shown significant similarities between the approaches. Noticing the interrelations is beneficial, for example in the case of the degree of integration rules and the accuracy statements that follow for the UT.

The simulations have shown that for mildly nonlinear functions all approaches perform well. It was seen that methods that are based on analytical Taylor expansion work well if the function is well approximated by polynomials of the assumed degree. However, if the degree of the approximation is too low, then the performance can deteriorate. The investigated UT variants (including the cubature setup) showed solid performance throughout the examples, but not without failure on some. Similar statements can be made for the divided difference approaches. Both failed for one of the investigated quadratic functions. Two more accurate Gauss-Hermite integration rules were shown to perform well, but at the expense of more sigma points which prohibits their use for high-dimensional problems.

We conclude with the insight that for moment computation problems, it is important to understand the addressed nonlinearities, to understand the type of problems that can be accurately handled, and to understand the limitations of sigma point methods. This is especially crucial in nonlinear Kalman filters which are based on the chosen moment computation methods.

APPENDIX

A. About Matrix Square Roots

Any positive definite $n \times n$ matrix, in particular the covariance P of $x \sim \mathcal{N}(\hat{x}, P)$, can be factored as

$$P^{1/2} P^{\mathrm{T}/2} = P \tag{77}$$

with an $n \times n$ matrix square root $P^{1/2}$.

The square root is not uniquely defined because any $P^{1/2}U$ with orthogonal U is also a matrix square root for P.

Matrix square roots can be computed using different algorithms [72]. The Cholesky factorization is a fast method for computing a triangular $P^{1/2}$. The singular value decomposition enjoys desirable numerical properties but is computationally more demanding.

The ith column of $P^{1/2}$ is given by

$$u_i = P^{1/2} e_i, \tag{78}$$

where e_i is the ith basis vector in the standard basis of $\mathbb{R}^n$.

From the columns of $P^{1/2}$

$$P = \sum_{i=1}^{n} u_i u_i^{\mathrm{T}} \tag{79}$$

can be re-constructed.

B. Nested Covariance Computation for Partitioned x

Expression (38b) of Sec. III-B can be derived as follows:

$$\begin{aligned}
\operatorname{cov}\{z\} &= \mathrm{E}\{zz^{\mathrm{T}}\} - \mathrm{E}\{z\}\mathrm{E}\{z\}^{\mathrm{T}} \\
&= \mathrm{E}\{\mathrm{E}\{zz^{\mathrm{T}} \mid x_2\}\} \\
&\quad - \mathrm{E}\{\mathrm{E}\{z \mid x_2\}\}\mathrm{E}\{\mathrm{E}\{z \mid x_2\}\}^{\mathrm{T}} \\
&= \mathrm{E}\{\operatorname{cov}\{z \mid x_2\}\} \\
&\quad + \mathrm{E}\{\mathrm{E}\{z \mid x_2\}\mathrm{E}\{z \mid x_2\}^{\mathrm{T}}\} \\
&\quad - \mathrm{E}\{\mathrm{E}\{z \mid x_2\}\}\mathrm{E}\{\mathrm{E}\{z \mid x_2\}\}^{\mathrm{T}} \\
&= \mathrm{E}\{\operatorname{cov}\{z \mid x_2\}\} + \operatorname{cov}\{\mathrm{E}\{z \mid x_2\}\}.
\end{aligned}$$

C. Sigma Point Implementation of the Exact Covariance for Quadratic f

The function f is quadratic. Each row f_k can be expressed by its Taylor expansion (53) with the Jacobian $J_k(\hat{x})$ and the Hessian matrix H_k. All J_k compose the Jacobian matrix J. The dependence on $\hat{x}$ is omitted for convenience. For the following discussion we discuss the first two rows f_1 and f_2.

First, we use the sigma points in (68) to derive the Jacobian term of (54b). Using (53), the transformed points can be written as

$$z_1^{(\pm i)} = f_1(\hat{x}) \pm \sqrt{\gamma} J_1 u_i + \frac{\gamma}{2} u_i^{\mathrm{T}} H_1 u_i. \tag{80}$$

For quadratic f this is not an approximation but exact for any γ. After pairwise processing and rearranging, we can isolate the Jacobian

$$z^{(i)} - z^{(-i)} = 2\sqrt{\gamma} J u_i, \tag{81a}$$

$$J u_i = \frac{1}{2\sqrt{\gamma}} (z^{(i)} - z^{(-i)}). \tag{81b}$$

A summation over all Ju_i together with the result (79) gives

$$\begin{aligned}
JPJ^{\mathrm{T}} &= \sum_{i=1}^{n} J u_i u_i^{\mathrm{T}} J^{\mathrm{T}} \\
&= \frac{1}{4\gamma} \sum_{i=1}^{n} (z^{(i)} - z^{(-i)})(z^{(i)} - z^{(-i)})^{\mathrm{T}},
\end{aligned} \tag{82}$$

the desired first term of (54b).

Next, we address the trace term of (54b) which for f_1 and f_2 can be written as

$$\begin{aligned}
\operatorname{tr}(H_1 P H_2 P) &= \sum_{i=1}^{n} \sum_{j=1}^{n} \operatorname{tr}(H_1 u_i u_i^{\mathrm{T}} H_2 u_j u_j^{\mathrm{T}}) \\
&= \sum_{i=1}^{n} \sum_{j=1}^{n} u_j^{\mathrm{T}} H_1 u_i u_i^{\mathrm{T}} H_2 u_j,
\end{aligned} \tag{83}$$

a sum of n^2 terms. One term in the sum is given by

$$u_j^{\mathrm{T}} H_1 u_i u_i^{\mathrm{T}} H_2 u_j. \tag{84}$$

For $i = j$, we can extract (84) from (80)

$$z_1^{(i)} + z_1^{(-i)} = 2 f_1(\hat{x}) + \gamma u_i^{\mathrm{T}} H_1 u_i, \tag{85a}$$

$$u_i^{\mathrm{T}} H_1 u_i = \frac{1}{\gamma} (z_1^{(i)} - 2 z_1^{(0)} + z_1^{(-i)}) \tag{85b}$$

and

$$\begin{aligned}
&u_i^{\mathrm{T}} H_1 u_i u_i^{\mathrm{T}} H_2 u_i \\
&\quad = \frac{1}{\gamma^2} (z_1^{(i)} - 2 z_1^{(0)} + z_1^{(-i)})(z_2^{(i)} - 2 z_2^{(0)} + z_2^{(-i)}).
\end{aligned} \tag{85c}$$

That is as far as we can get with the original sigma points (68). From here we start working with the extended sigma points (70) which are passed through f to yield

$$\begin{aligned}
z_1^{(\pm ij)} &= f_1(\hat{x}) \pm \frac{\sqrt{\gamma}}{2} J_1 (u_i + u_j) \\
&\quad + \frac{\gamma}{8} (u_i + u_j)^{\mathrm{T}} H_1 (u_i + u_j).
\end{aligned} \tag{86}$$

Again, we can isolate the Hessian term by a combination of transformed samples similar to (85b). The resulting expression

$$\begin{aligned} z_1^{(ij)} - 2z_1^{(0)} + z_1^{(-ij)} \\ &= \frac{\gamma}{4}(u_i + u_j)^{\mathrm{T}} H_1 (u_i + u_j) \\ &= \frac{\gamma}{4}(u_i^{\mathrm{T}} H_1 u_i + u_j^{\mathrm{T}} H_1 u_j + 2u_j^{\mathrm{T}} H_1 u_i) \end{aligned} \tag{87}$$

involves a mixed j,i-term that we require for computing (84) and also homogeneous terms (85c). We can now isolate the mixed term $2u_j^{\mathrm{T}} H_1 u_i$ of (87) by subtracting expressions of the form (85b):

$$\begin{aligned} 2u_j^{\mathrm{T}} H_1 u_i = \frac{4}{\gamma}(z_1^{(ij)} - 2z_1^{(0)} + z_1^{(-ij)}) \\ - \frac{1}{\gamma}(z_1^{(i)} - 2z_1^{(0)} + z_1^{(-i)}) \\ - \frac{1}{\gamma}(z_1^{(j)} - 2z_1^{(0)} + z_1^{(-j)}). \end{aligned} \tag{88}$$

Simplification yields

$$\begin{aligned} u_j^{\mathrm{T}} H_1 u_i = \frac{1}{2\gamma}(4z_1^{(ij)} + 4z_1^{(-ij)} \\ - z_1^{(i)} - z_1^{(-i)} - z_1^{(j)} - z_1^{(-j)} - 4z_1^{(0)}). \end{aligned} \tag{89}$$

Next, (84) can be written as product

$$\begin{aligned} u_j^{\mathrm{T}} H_1 u_i u_i^{\mathrm{T}} H_2 u_j = \frac{1}{4\gamma^2} \\ \times (4z_1^{(ij)} + 4z_1^{(-ij)} - z_1^{(i)} - z_1^{(-i)} - z_1^{(j)} - z_1^{(-j)} - 4z_1^{(0)}) \\ \times (4z_2^{(ij)} + 4z_2^{(-ij)} - z_2^{(i)} - z_2^{(-i)} - z_2^{(j)} - z_2^{(-j)} - 4z_2^{(0)}). \end{aligned} \tag{90}$$

Summation over i and j from 1 to n yields the trace term of (83) and hence the second term of (54b).

REFERENCES

[1] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan
Estimation with Applications to Tracking and Navigation.
Wiley-Interscience, Jun. 2001.

[2] F. Gustafsson
Statistical Sensor Fusion.
Studentlitteratur AB, Mar. 2010.

[3] Y. Zheng and M. Hasegawa-Johnson
Particle filtering approach to Bayesian formant tracking,
in *IEEE Workshop on Statistical Signal Processing*, Oct. 2003, pp. 601–604.

[4] A. H. Jazwinski
Stochastic Processes and Filtering Theory.
Academic Press, Mar. 1970.

[5] R. E. Kalman
A new approach to linear filtering and prediction problems,
Journal of basic Engineering, vol. 82, no. 1, pp. 35–45, Mar. 1960.

[6] B. D. O. Anderson and J. B. Moore
Optimal Filtering.
Prentice Hall, Jun. 1979.

[7] M. Grewal and A. Andrews
Applications of Kalman filtering in aerospace 1960 to the present,
IEEE Control Systems Magazine, vol. 30, no. 3, pp. 69–78, 2010.

[8] M. Athans, R. Wishner, and A. Bertolini
Suboptimal state estimation for continuous-time nonlinear systems from discrete noisy measurements,
IEEE Transactions on Automatic Control, vol. 13, no. 5, pp. 504–514, 1968.

[9] H. W. Sorenson and A. R. Stubberud
Recursive filtering for systems with small but non-negligible non-linearities,
International Journal of Control, vol. 7, no. 3, pp. 271–280, 1968.

[10] H. W. Sorenson
On the development of practical nonlinear filters,
Information Sciences, vol. 7, pp. 253–270, 1974.

[11] A. Gelb
Applied Optimal Estimation.
MIT Press, May 1974.

[12] N. J. Gordon, D. J. Salmond, and A. F. M. Smith
Novel approach to non-linear/non-Gaussian Bayesian state estimation,
IEE Proceedings-F Radar and Signal Processing, vol. 140, no. 2, pp. 107–113, Apr. 1993.

[13] F. Gustafsson
Particle filter theory and practice with positioning applications,
IEEE Aerospace and Electronics Systems Magazine, vol. 25, no. 7, pp. 53–82, Jul. 2010.

[14] S. Julier, J. Uhlmann, and H. Durrant-Whyte
A new approach for filtering nonlinear systems,
in *Proceedings of the American Control Conference*, vol. 3, 1995, pp. 1628–1632.

[15] S. Julier, J. Uhlmann, and H. F. Durrant-Whyte
A new method for the nonlinear transformation of means and covariances in filters and estimators,
IEEE Transactions on Automatic Control, vol. 45, no. 3, pp. 477–482, Mar. 2000.

[16] S. J. Julier and J. K. Uhlmann
Unscented filtering and nonlinear estimation,
Proceedings of the IEEE, vol. 92, no. 3, pp. 401–422, Mar. 2004.

[17] K. Ito and K. Xiong
Gaussian filters for nonlinear filtering problems,
IEEE Transactions on Automatic Control, vol. 45, no. 5, pp. 910–927, May 2000.

[18] I. Arasaratnam and S. Haykin
Cubature Kalman filters,
IEEE Transactions on Automatic Control, vol. 54, no. 6, pp. 1254–1269, Jun. 2009.

[19] T. S. Schei
A finite-difference method for linearization in nonlinear estimation algorithms,
Automatica, vol. 33, no. 11, pp. 2053–2058, Nov. 1997.

[20] M. Nørgaard, N. K. Poulsen, and O. Ravn
New developments in state estimation for nonlinear systems,
Automatica, vol. 36, no. 11, pp. 1627–1638, Nov. 2000.

[21] R. van der Merwe
"Sigma-point Kalman filters for probabilistic inference in dynamic state-space models,"
Ph.D. dissertation, OGI School of Science & Engineering, Oregon Health & Science University, Portland, OR, USA, Apr. 2004.

[22] S. Särkkä
Bayesian Filtering and Smoothing.
New York: Cambridge University Press, Oct. 2013.

[23] F. Gustafsson and G. Hendeby
Some relations between extended and unscented Kalman filters,
IEEE Transactions on Signal Processing, vol. 60, no. 2, pp. 545–555, Feb. 2012.

[24] M. Evans and T. Swartz
Methods for approximating integrals in statistics with special emphasis on Bayesian integration problems,
Statistical Science, vol. 10, no. 3, pp. 254–272, Aug. 1995.

[25] ———
Approximating Integrals via Monte Carlo and Deterministic Methods, 1st ed.
New York: Oxford University Press, May 2000.

[26] Y. Wu, D. Hu, M. Wu, and X. Hu
A numerical-integration perspective on Gaussian filters,
IEEE Transactions on Signal Processing, vol. 54, no. 8, pp. 2910–2921, Aug. 2006.

[27] M. Šimandl and J. Duník
Derivative-free estimation methods: New results and performance analysis,
Automatica, vol. 45, no. 7, pp. 1749–1757, Jul. 2009.

[28] P. Stano, Z. Lendek, J. Braaksma, R. Babuska, C. de Keizer, and A. den Dekker
Parametric Bayesian filters for nonlinear stochastic dynamical systems: A survey,
IEEE Transactions on Cybernetics, vol. 43, no. 6, pp. 1607–1624, Dec. 2013.

[29] T. Lefebvre, H. Bruyninckx, and J. De Schutter
Kalman filters for non-linear systems: a comparison of performance,
International Journal of Control, vol. 77, no. 7, pp. 639–653, 2004.

[30] M. Roth
"Kalman filters for nonlinear systems and heavy-tailed noise,"
Licentiate thesis, Linköping University, Linköping, Sweden, Sep. 2013. [Online]. Available: http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-97544.

[31] O. Cappé, E. Moulines, and T. Ryden
Inference in Hidden Markov Models.
New York, NY: Springer, Dec. 2010.

[32] P. S. Maybeck
Stochastic Models, Estimation, and Control: Volume 2.
Academic Press, Jun. 1982.

[33] T. M. Cover and J. A. Thomas
Elements of information theory, 2nd ed.
Hoboken, N.J.: Wiley-Interscience, 2006.

[34] C. M. Bishop
Pattern Recognition and Machine Learning.
Springer, Aug. 2006.

[35] M. Roth and F. Gustafsson
An efficient implementation of the second order extended Kalman filter,
in *Proceedings of the 14th International Conference on Information Fusion (FUSION)*, Jul. 2011.

[36] J. Sarmavuori and S. Särkkä
Fourier-Hermite Kalman filter,
IEEE Transactions on Automatic Control, vol. 57, no. 6, pp. 1511–1515, 2012.

[37] F. Sandblom and L. Svensson
Moment estimation using a marginalized transform,
IEEE Transactions on Signal Processing, vol. 60, no. 12, pp. 6138–6150, 2012.

[38] M. F. Huber
Chebyshev polynomial Kalman filter,
Digital Signal Processing, vol. 23, no. 5, pp. 1620–1629, Sep. 2013.

[39] J. Steinbring and U. D. Hanebeck
LRKF revisited—the smart sampling Kalman filter (S2KF),
Journal of Advances in Information Fusion, vol. 9, no. 2, pp. 106–123, Dec. 2014.

[40] H. Tanizaki
Nonlinear Filters: Estimation and Applications, 2nd ed.
Springer, Aug. 1996.

[41] I. Arasaratnam, S. Haykin, and R. J. Elliott
Discrete-time nonlinear filtering algorithms using Gauss-Hermite quadrature,
Proceedings of the IEEE, vol. 95, no. 5, pp. 953–977, 2007.

[42] B. Jia, M. Xin, and Y. Cheng
High-degree cubature Kalman filter,
Automatica, vol. 49, no. 2, pp. 510–518, Feb. 2013.

[43] J. Dunik, O. Straka, and M. Simandl
Stochastic integration filter,
IEEE Transactions on Automatic Control, vol. 58, no. 6, pp. 1561–1566, 2013.

[44] J. Dunik, M. Simandl, and O. Straka
Unscented Kalman filter: Aspects and adaptive setting of scaling parameter,
IEEE Transactions on Automatic Control, vol. 57, no. 9, pp. 2411–2416, Sep. 2012.

[45] S. Blackman and R. Popoli
Design and Analysis of Modern Tracking Systems.
Artech House, Aug. 1999.

[46] G. Hendeby, R. Karlsson, and F. Gustafsson
The Rao-Blackwellized particle filter: A filter bank implementation,
EURASIP Journal on Advances in Signal Processing, vol. 2010, no. 1, Dec. 2010.

[47] A. Gut
An Intermediate Course in Probability, 2nd ed.
Springer, Jun. 2009.

[48] M. Abramowitz and I. A. Stegun
Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables.
Dover publications, 1965, vol. 55.

[49] T. W. Anderson
An Introduction to Multivariate Statistical Analysis, 3rd ed.
Wiley-Interscience, 2003.

[50] K.-T. Fang, S. Kotz, and K. W. Ng
Symmetric Multivariate and Related Distributions.
Chapman and Hall/CRC, Nov. 1989.

[51] M. Roth, E. Özkan, and F. Gustafsson
A Student's t filter for heavy tailed process and measurement noise,
in *38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Vancouver, Canada, May 2013.

[52] S. Saha, P. K. Mandal, Y. Boers, H. Driessen, and A. Bagchi
Gaussian proposal density using moment matching in SMC methods,
Statistics and Computing, vol. 19, no. 2, pp. 203–208, Jun. 2009.

[53] T. Kollo
Advanced Multivariate Statistics with Matrices.
Dordrecht, The Netherlands: Springer, Aug. 2005.

[54] B. Holmquist
Moments and cumulants of the multivariate normal distribution,
Stochastic Analysis and Applications, vol. 6, no. 3, pp. 273–278, Jan. 1988.

[55] J. E. Chacón and T. Duong
Multivariate plug-in bandwidth selection with unconstrained pilot bandwidth matrices,
TEST, vol. 19, no. 2, pp. 375–398, Aug. 2010.

[56] D. Manolakis
Kronecker product based second order approximation of mean value and covariance matrix in nonlinear transformations,
IEEE Signal Processing Letters, vol. 18, no. 1, pp. 43–46, Jan. 2011.

[57] C. P. Robert and G. Casella
Monte Carlo Statistical Methods, 2nd ed., ser. Springer Texts in Statistics.
Springer, 2004.

[58] D. J. C. MacKay
Information Theory, Inference and Learning Algorithms, 1st ed.
Cambridge, UK: Cambridge University Press, Oct. 2003.

[59] M. Roth, C. Fritsche, G. Hendeby, and F. Gustafsson
The ensemble Kalman filter and its relations to other nonlinear filters,
in *European Signal Processing Conference 2015 (EUSIPCO 2015)*, Nice, France, Aug. 2015.

[60] U. D. Hanebeck, M. F. Huber, and V. Klumpp
Dirac mixture approximation of multivariate Gaussian densities,
in *Proceedings of the 48th IEEE Conference on Decision and Control*, 2009, pp. 3851–3858.

[61] R. Gray and D. Neuhoff
Quantization,
IEEE Transactions on Information Theory, vol. 44, no. 6, pp. 2325–2383, 1998.

[62] S. J. Julier
The scaled unscented transformation,
in *Proceedings of the American Control Conference*, vol. 6, 2002, pp. 4555–4559.

[63] T. Lefebvre, H. Bruyninckx, and J. De Schutter
Comment on "A new method for the nonlinear transformation of means and covariances in filters and estimators" [with authors' reply],
IEEE Transactions on Automatic Control, vol. 47, no. 8, pp. 1406–1409, Aug. 2002.

[64] A. H. Stroud
Approximate Calculation of Multiple Integrals.
Prentice-Hall, 1971.

[65] A. H. Stroud and D. Secrest
Approximate integration formulas for certain spherically symmetric regions,
Mathematics of Computation, vol. 17, no. 82, pp. 105–135, Apr. 1963.

[66] C. F. Van Loan
The ubiquitous Kronecker product,
Journal of Computational and Applied Mathematics, vol. 123, no. 1–2, pp. 85–100, Nov. 2000.

[67] P. Closas, C. Fernandez-Prades, and J. Vila-Valls
Multiple quadrature Kalman filtering,
IEEE Transactions on Signal Processing, vol. 60, no. 12, pp. 6125–6137, 2012.

[68] A. Genz and J. Monahan
Stochastic integration rules for infinite regions,
SIAM Journal on Scientific Computing, vol. 19, no. 2, pp. 426–439, Mar. 1998.

[69] O. Straka, J. Dunik, and M. Simandl
Randomized unscented Kalman filter in target tracking,
in *15th International Conference on Information Fusion*, Jul. 2012, pp. 503–510.

[70] J. Dunik, O. Straka, M. Simandl, and E. Blasch
Random-point-based filters: analysis and comparison in target tracking,
IEEE Transactions on Aerospace and Electronic Systems, vol. 51, no. 2, pp. 1403–1421, Apr. 2015.

[71] M. Nørgaard, N. K. Poulsen, and O. Ravn
"Advances in derivative-free state estimation for nonlinear systems,"
Department of Mathematical Modelling, DTU, Lyngby, Denmark, Tech. Rep. IMM-REP-1998-15, Apr. 2000.

[72] G. H. Golub and C. F. Van Loan
Matrix Computations, 3rd ed.
Baltimore: Johns Hopkins University Press, Oct. 1996.

Michael Roth is a Ph.D. student in the Division of Automatic Control, Department of Electrical Engineering, Linköping University, since 2010. He received his Master's degree (Diplom-Ingenieur) in Electrical Engineering from TU Berlin in 2010. His main research is on algorithms for filtering and smoothing in non-linear, non-Gaussian, and high-dimensional state space models. Moreover, he has been involved in several projects about system identification, control, and signal processing for medical applications.

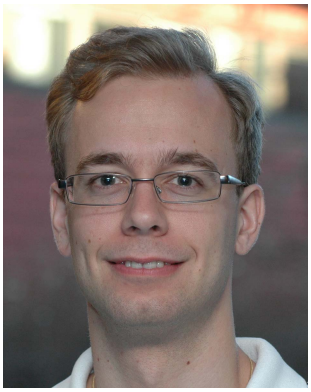

Gustaf Hendeby is Associate Professor in the Division of Automatic Control, Department of Electrical Engineering, Linköping University. He received his M.Sc. in Applied Physics and Electrical Engineering in 2002 and his Ph.D. in Automatic Control in 2008, both from Linköping University. He worked as Senior Researcher at the German Research Center for Artificial Intelligence (DFKI) 2009–2011, and Senior Scientist at Swedish Defense Research Agency (FOI) and held an adjunct Associate Professor position at Linköping University 2011–2015. Dr. Hendeby's main research interests are stochastic signal processing and sensor fusion with applications to nonlinear problems, target tracking, and simultaneous localization and mapping (SLAM). He has experience of both theoretical analysis as well as implementation aspects.

Fredrik Gustafsson is professor in Sensor Informatics at the Department of Electrical Engineering, Linköping University, since 2005. He received the M.Sc. degree in electrical engineering in 1988 and the Ph.D. degree in Automatic Control in 1992, both from Linköping University. During 1992–1999 he held various positions in automatic control, and 1999–2005 he had a professorship in Communication Systems. His research interests are in stochastic signal processing, adaptive filtering and change detection, with applications to communication, vehicular, airborne, and audio systems. He is a co-founder of the companies NIRA Dynamics (automotive safety systems), Softube (audio effects) and SenionLab (indoor positioning systems).

He was an associate editor for *IEEE Transactions of Signal Processing* 2000–2006, *IEEE Transactions on Aerospace and Electronic Systems* 2010–2012, and EURASIP *Journal on Applied Signal Processing* 2007–2012. He was awarded the Arnberg prize by the Royal Swedish Academy of Science (KVA) 2004, elected member of the Royal Academy of Engineering Sciences (IVA) 2007, and elevated to IEEE Fellow 2011. He was awarded the Harry Rowe Mimno Award 2011 for the tutorial "Particle Filter Theory and Practice with Positioning Applications," which was published in the *AESS Magazine* in July 2010, and was co-author of "Smoothed state estimates under abrupt changes using sum-of-norms regularization" that received the Automatica paper prize in 2014.