Kalman Filters for nonlinear systems:

a comparison of performance.

Tine Lefebvre, Herman Bruyninckx and Joris De Schutter

Katholieke Universiteit Leuven, Dept. of Mechanical Eng., Division P.M.A.,
Celestijnenlaan 300B, B-3001 Heverlee, Belgium

Corresponding author: Tine Lefebvre
T. Lefebvre is visiting Stanford University till September 2004.
Please send all mail regarding this manuscript to
Tine Lefebvre, Stanford University,
Robotics Laboratory, Gates Building 1A,
Stanford CA 94305-9010, USA
Tel.: (+1) 650 380 3530, Fax: (+1) 650 725 1449.
E-mail: Tine.Lefebvre@mech.kuleuven.ac.be

Submitted as a Full Paper
to the International Journal of Control
Ref: E03138, Revised March 2004

Abstract

The Kalman Filter is a well-known recursive state estimator for linear systems. In practise, the algorithm is often used for nonlinear systems by linearising the system's process and measurement models. Different ways of linearising the models lead to different filters. In some applications, these 'Kalman Filter variants' seem to perform well, while for other applications they are useless. When choosing a filter for a new application, the literature gives us little to rely on. This paper tries to bridge the gap between the theoretical derivation of a Kalman Filter variant and its performance in practise when applied to a nonlinear system, by providing an application-independent analysis of the performances of the common Kalman Filter variants.

This paper separates performance evaluation of Kalman Filters into (i) consistency, and (ii) information content of the estimates; and it separates the filter structure into (i) the process update step, and (ii) the measurement update step. This decomposition provides the insights supporting an objective and systematic evaluation of the appropriateness of a particular Kalman Filter variant in a particular application.

1 Introduction

During the last decades, many research areas looked into the matter of on-line state estimation. The uncertainty on the state value varies over time due to the changes in the system state (the process updates) and due to the information in the measurements (the measurement updates). The uncertainty can be represented in different ways, e.g. by intervals or fuzzy sets.

In Bayesian estimation (Bayes 1763, Laplace 1812), a state estimate is represented by a probability density function (pdf). Fast analytical update algorithms require the pdf to be an analytical function of a limited number of time-varying parameters, which is only true for some systems. A well-known example are systems with linear process and measurement models and with additive Gaussian uncertainties. The pdf is then a Gaussian distribution, which is fully determined by its mean vector and covariance matrix. This mean and covariance are updated analytically with the Kalman Filter algorithm (KF, Kalman 1960, Sorenson 1985).

For most nonlinear systems, the pdf cannot be written as an analytical function with time-varying parameters. In order to have a computationally interesting update algorithm, the KF is used as an approximation. This is achieved by linearisation of the process and measurement models of the system. It also means that the true pdf is approximated by a

Gaussian distribution. Different ways of linearisation (different KF variants) lead to different results.

This paper describes (i) how the common KF variants differ in their linearisation of the process and measurement models; (ii) how they take the linearisation errors into account; and (iii) how the quality of their state estimates depends on the previous two choices. The studied algorithms are:

1. The Extended Kalman Filter (EKF, Gelb, Kasper, Nash, Price and Sutherland 1974, Maybeck 1982, Bar-Shalom and Li 1993, Tanizaki 1996);

2. The Iterated Extended Kalman Filter (IEKF, Gelb, Kasper, Nash, Price and Sutherland 1974, Maybeck 1982, Bar-Shalom and Li 1993, Tanizaki 1996);

3. The Linear Regression Kalman Filter (LRKF, Lefebvre, Bruyninckx and Schutter 2002). This filter comprises the Central Difference Filter (CDF, Schei 1997), the first-order Divided Difference Filter (DD1, Nørgaard, Poulsen and Ravn 2000a, b) and the Unscented Kalman Filter (UKF, Uhlmann 1995, Julier and Uhlmann 1996, 2001, Julier, Uhlman and Durrant-Whyte 2000).

The paper gives the following new insights:

1. The quality of the estimates of the KF variants can be expressed by two criteria, i.e. the consistency and the information content of the estimates (defined in section 3). This paper relates the consistency and information content of the estimates to (i) how the linearisation is performed and (ii) how the linearisation errors are taken into account.

2. Although the filters use similar linearisation techniques for the linearisation of the process and measurement models, there can be a substantial difference in their performance for both updates:

   (a) for the linearisation of the process update (section 4), which describes the evolution of the state, the state estimate and its uncertainty are the only available information;

   (b) the measurement update (section 5), on the other hand, describes the fusion of the information in the state estimate with the information in the new measurement.

Hence, in this update also the new measurement is available and can be used to linearise the measurement model.

Therefore, it can be interesting to use different filters for both updates.

3. Two new insights on the performance of specific KF variants are: (i) the IEKF measurement update outperforms the EKF and LRKF updates if the state —or at least the part of it that causes the non-linearity in the measurement model— is instantaneously fully observable (section 5.2); and (ii) for large uncertainties on the state estimate, the LRKF measurement update yields consistent but non-informative state estimates (section 5.3).

These insights are obtained because:

1. This paper describes all filter algorithms as the application of the basic KF algorithm to linearised process and measurement models. The difference between the KF variants is situated in the choice of linearisation and the compensation of the linearisation errors. In previous work this linearisation was not always recognised, e.g. the UKF is originally derived as a filter which does not linearise the models.

2. The analysis clarifies how some filters automatically adapt their compensation for linearisation errors, while other filters have constant (developer-determined) error compensation.

3. Additionally, the paper compares the filter performances separately for process updates and measurement updates instead of their overall performance when they are both combined. In the existing literature, the performances of the KF variants are often compared by interpreting the estimation results for a specific application after executing a large number of process and measurement update steps.

The analysis starts from a general formulation of nonlinear process and measurement models, making the results application independent. The obtained insights are important for all researchers and developers who want to apply a KF variant to a specific application. They lead to a systematic choice of a filter, where previously the choice was mainly made based on success in similar applications or based on trial and error.

Examples of 2D systems are provided. The models are chosen such that they provide a clear graphical demonstration of the discussed effects. For the measurement update, the filters' performances are system dependent, hence, in that case several models are used for illustration.

2 The Kalman Filter algorithm

The (linear) Kalman Filter

The Kalman Filter (KF, Kalman 1960, Sorenson 1985) is a special case of Bayesian filtering theory. It applies to the estimation of a state $\widetilde{\boldsymbol{x}}$ if the state space description of the estimation problem has linear process and measurement equations subject to additive Gaussian uncertainty:

$$\widetilde{\boldsymbol{x}}_k = \widetilde{\mathbf{F}}_{k-1}\widetilde{\boldsymbol{x}}_{k-1} + \widetilde{\boldsymbol{b}}_{k-1} + \widetilde{\mathbf{C}}_{k-1}\widetilde{\boldsymbol{\rho}}_{p,k-1}; \tag{1}$$

$$\widetilde{\boldsymbol{z}}_k = \widetilde{\mathbf{H}}_k\widetilde{\boldsymbol{x}}_k + \widetilde{\boldsymbol{d}}_k + \widetilde{\mathbf{E}}_k\widetilde{\boldsymbol{\rho}}_{m,k}. \tag{2}$$

$\widetilde{\boldsymbol{z}}$ is the measurement vector. The subscripts k and $k-1$ indicate the time step. $\widetilde{\mathbf{F}}$, $\widetilde{\boldsymbol{b}}$, $\widetilde{\mathbf{C}}$, $\widetilde{\mathbf{H}}$, $\widetilde{\boldsymbol{d}}$ and $\widetilde{\mathbf{E}}$ are (possibly nonlinear) functions of the system input. $\widetilde{\boldsymbol{\rho}}_p$ denotes the process uncertainty, being a random vector sequence with zero mean and known covariance matrix $\widetilde{\mathbf{Q}}$. $\widetilde{\boldsymbol{\rho}}_m$ is the measurement uncertainty and is a random vector sequence with zero mean and known covariance matrix $\widetilde{\mathbf{R}}$; $\widetilde{\boldsymbol{\rho}}_p$ and $\widetilde{\boldsymbol{\rho}}_m$ are mutually uncorrelated and uncorrelated between sampling times[1]. Furthermore, assume a Gaussian prior pdf $p(\widetilde{\boldsymbol{x}}_0)$ with mean $\widetilde{\hat{\boldsymbol{x}}}_{0|0}$ and covariance matrix $\widetilde{\mathbf{P}}_{0|0}$.

For this system, the pdfs[2] $p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_{k-1})$ and $p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_k)$ are also Gaussian distributions. The filtering formulas can be expressed as analytical functions calculating the mean vector $\widetilde{\hat{\boldsymbol{x}}}$ and

[1] Correlated uncertainties can be dealt with by augmenting the state vector, this is the original formulation of the KF (Kalman 1960). Expressed in this new state vector, the process and measurement models are of the form (1) and (2) with uncorrelated uncertainties.

[2] $p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_j)$ denotes the pdf of the state $\widetilde{\boldsymbol{x}}$ at time k, given the measurements $\widetilde{\hat{\mathbf{Z}}}_j = \{\widetilde{\hat{\boldsymbol{z}}}_1, \ldots, \widetilde{\hat{\boldsymbol{z}}}_j\}$ up to time j.

covariance matrix $\widetilde{\mathbf{P}}$ of these pdfs:

$$\begin{aligned}
\widetilde{\hat{\boldsymbol{x}}}_{k|k-1} &= E_{p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_{k-1})}\left[\widetilde{\boldsymbol{x}}_k\right]; \\
&= \widetilde{\mathbf{F}}_{k-1}\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1} + \widetilde{\boldsymbol{b}}_{k-1}; && (3) \\
\widetilde{\mathbf{P}}_{k|k-1} &= E_{p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_{k-1})}\left[\left(\widetilde{\boldsymbol{x}}_k - \widetilde{\hat{\boldsymbol{x}}}_{k|k-1}\right)\left(\widetilde{\boldsymbol{x}}_k - \widetilde{\hat{\boldsymbol{x}}}_{k|k-1}\right)^T\right]; \\
&= \widetilde{\mathbf{F}}_{k-1}\widetilde{\mathbf{P}}_{k-1|k-1}\widetilde{\mathbf{F}}_{k-1}^T + \widetilde{\mathbf{C}}_{k-1}\widetilde{\mathbf{Q}}_{k-1}\widetilde{\mathbf{C}}_{k-1}^T; && (4) \\
\widetilde{\hat{\boldsymbol{x}}}_{k|k} &= E_{p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_{k})}\left[\widetilde{\boldsymbol{x}}_k\right]; \\
&= \widetilde{\hat{\boldsymbol{x}}}_{k|k-1} + \widetilde{\mathbf{K}}_k\widetilde{\boldsymbol{\eta}}_k; && (5) \\
\widetilde{\mathbf{P}}_{k|k} &= E_{p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_{k})}\left[\left(\widetilde{\boldsymbol{x}}_k - \widetilde{\hat{\boldsymbol{x}}}_{k|k}\right)\left(\widetilde{\boldsymbol{x}}_k - \widetilde{\hat{\boldsymbol{x}}}_{k|k}\right)^T\right]; \\
&= \left(\widetilde{\mathbf{I}}_{n\times n} - \widetilde{\mathbf{K}}_k\widetilde{\mathbf{H}}_k\right)\widetilde{\mathbf{P}}_{k|k-1}; && (6)
\end{aligned}$$

where

$$\begin{aligned}
\widetilde{\boldsymbol{\eta}}_k &= \widetilde{\hat{\boldsymbol{z}}}_k - \widetilde{\mathbf{H}}_k\widetilde{\hat{\boldsymbol{x}}}_{k|k-1} - \widetilde{\boldsymbol{d}}_k; && (7) \\
\widetilde{\mathbf{S}}_k &= \widetilde{\mathbf{E}}_k\widetilde{\mathbf{R}}_k\widetilde{\mathbf{E}}_k^T + \widetilde{\mathbf{H}}_k\widetilde{\mathbf{P}}_{k|k-1}\widetilde{\mathbf{H}}_k^T; && (8) \\
\widetilde{\mathbf{K}}_k &= \widetilde{\mathbf{P}}_{k|k-1}\widetilde{\mathbf{H}}_k^T\widetilde{\mathbf{S}}_k^{-1}. && (9)
\end{aligned}$$

$\widetilde{\boldsymbol{\eta}}$ is called the innovation, its covariance is $\widetilde{\mathbf{S}}$. $\widetilde{\mathbf{K}}$ is the Kalman gain. Equations (3)–(4) are referred to as the process update, equations (5)–(9) as the measurement update. $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$ is called the predicted state estimate and $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$ the updated state estimate. If no measurement $\widetilde{\hat{\boldsymbol{z}}}_k$ is available at a certain time step k, then equations (5)–(9) reduce to $\widetilde{\hat{\boldsymbol{x}}}_{k|k} = \widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$ and $\widetilde{\mathbf{P}}_{k|k} = \widetilde{\mathbf{P}}_{k|k-1}$.

Kalman Filters for nonlinear systems

The KF algorithm is often applied to systems with nonlinear process and measurement mod-

els[3],

$$\widetilde{\boldsymbol{x}}_k = \widetilde{\boldsymbol{f}}_{k-1}(\widetilde{\boldsymbol{x}}_{k-1}) + \widetilde{\mathbf{C}}_{k-1}\widetilde{\boldsymbol{\rho}}_{p,k-1}; \tag{10}$$

$$\widetilde{\boldsymbol{z}}_k = \widetilde{\boldsymbol{h}}_k(\widetilde{\boldsymbol{x}}_k) + \widetilde{\mathbf{E}}_k\widetilde{\boldsymbol{\rho}}_{m,k}; \tag{11}$$

by linearisation:

$$\widetilde{\boldsymbol{x}}_k = \widetilde{\mathbf{F}}_{k-1}\widetilde{\boldsymbol{x}}_{k-1} + \widetilde{\boldsymbol{b}}_{k-1} + \widetilde{\boldsymbol{\rho}}^*_{p,k-1} + \widetilde{\mathbf{C}}_{k-1}\widetilde{\boldsymbol{\rho}}_{p,k-1}; \tag{12}$$

$$\widetilde{\boldsymbol{z}}_k = \widetilde{\mathbf{H}}_k\widetilde{\boldsymbol{x}}_k + \widetilde{\boldsymbol{d}}_k + \widetilde{\boldsymbol{\rho}}^*_{m,k} + \widetilde{\mathbf{E}}_k\widetilde{\boldsymbol{\rho}}_{m,k}. \tag{13}$$

The difference between these models and models (1)–(2) is the presence of the terms $\widetilde{\boldsymbol{\rho}}^*_p$ and $\widetilde{\boldsymbol{\rho}}^*_m$ representing the linearisation errors. The additional uncertainty on the linearised models due to these linearisation errors is modelled by the covariance matrices $\widetilde{\mathbf{Q}}^*$ and $\widetilde{\mathbf{R}}^*$. Unfortunately, applying the KF[4] (3)–(9) to systems with nonlinear process and/or measurement models leads to non-optimal estimates and covariance matrices. Different ways of linearising the process and measurement models, i.e. different choices for $\widetilde{\mathbf{F}}$, $\widetilde{\boldsymbol{b}}$, $\widetilde{\mathbf{Q}}^*$, $\widetilde{\mathbf{H}}$, $\widetilde{\boldsymbol{d}}$ and $\widetilde{\mathbf{R}}^*$, yield other results. This paper aims at making an objective comparison of the performances of the commonly used linearisations (KF variants).

3 Consistency and information content of the estimates

The KF variants for nonlinear systems calculate an estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$ and covariance matrix $\widetilde{\mathbf{P}}_{k|i}$ for a pdf which is non-Gaussian. The performance of these KFs depends on how representative the Gaussian pdf with mean $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$ and covariance $\widetilde{\mathbf{P}}_{k|i}$ is for the (unknown) pdf $p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_i)$. Figure 1 shows a non-Gaussian pdf $p(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ and three possible Gaussian approximations $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$, $p_2(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ and $p_3(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$. Intuitively we feel that $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ is a good approximation because 'the same' values of x are probable. Similarly $p_3(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ is not a good approximation because a lot of probable values for x of the original distribution have a probability density

[3]Models which are nonlinear functions of the uncertainties $\widetilde{\boldsymbol{\rho}}_p$ and $\widetilde{\boldsymbol{\rho}}_m$, can be dealt with by augmenting the state vector with the uncertainties. Expressed in this new state vector, the process and measurement models are of the form (10) and (11).

[4]$\widetilde{\mathbf{C}}_{k-1}\widetilde{\boldsymbol{\rho}}_{p,k-1}$ of equation (1) corresponds to $\widetilde{\boldsymbol{\rho}}^*_{p,k-1} + \widetilde{\mathbf{C}}_{k-1}\widetilde{\boldsymbol{\rho}}_{p,k-1}$ of equation (12); hence instead of $\widetilde{\mathbf{C}}_{k-1}\widetilde{\mathbf{Q}}_{k-1}\widetilde{\mathbf{C}}^T_{k-1}$ in equation (4), $\widetilde{\mathbf{Q}}^*_{k-1} + \widetilde{\mathbf{C}}_{k-1}\widetilde{\mathbf{Q}}_{k-1}\widetilde{\mathbf{C}}^T_{k-1}$ is used.
$\widetilde{\mathbf{E}}_k\widetilde{\boldsymbol{\rho}}_{m,k}$ of equation (2) corresponds to $\widetilde{\boldsymbol{\rho}}^*_{m,k} + \widetilde{\mathbf{E}}_k\widetilde{\boldsymbol{\rho}}_{m,k}$ of equation (13); hence instead of $\widetilde{\mathbf{E}}_k\widetilde{\mathbf{R}}_k\widetilde{\mathbf{E}}^T_k$ in equation (9), $\widetilde{\mathbf{R}}^*_k + \widetilde{\mathbf{E}}_k\widetilde{\mathbf{R}}_k\widetilde{\mathbf{E}}^T_k$ is used.

of approximately zero in $p_3(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$. Finally pdf $p_2(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ is 'more uncertain' than $p(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ because a larger domain of x values is uncertain.

[Figure 1 about here.]

These intuitive reflexions are formulated in two criteria: the consistency and the information content of the state estimate. The consistency of the state estimate is a necessary condition for a filter to be acceptable. The information content of the state estimates defines an ordering between all consistent filters.

1. the consistency of the state estimate.

   A state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$ with covariance matrix $\widetilde{\mathbf{P}}_{k|i}$ is called consistent if

   $$E_{p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_i)}\left[\left(\widetilde{\boldsymbol{x}}_k - \widetilde{\hat{\boldsymbol{x}}}_{k|i}\right)\left(\widetilde{\boldsymbol{x}}_k - \widetilde{\hat{\boldsymbol{x}}}_{k|i}\right)^T\right] \leq \widetilde{\mathbf{P}}_{k|i}. \tag{14}$$

   For consistent results, the matrix $\widetilde{\mathbf{P}}_{k|i}$ is equal to or larger than the expected squared deviation with respect to the estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$ under the (unknown) distribution $p(\widetilde{\boldsymbol{x}}_k|\widetilde{\hat{\mathbf{Z}}}_i)$. The mean and covariance of pdfs $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ and $p_2(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ in figure 1 obey equation (14). Pdf $p_3(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$, on the other hand, is inconsistent.

   Inconsistency of the calculated state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$ and covariance matrix $\widetilde{\mathbf{P}}_{k|i}$ ('divergence' of the filter) is the most encountered problem with the KF variants. In this case, $\widetilde{\mathbf{P}}_{k|i}$ is too small and does no longer represent a reliable measure for the uncertainty on $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$. Even more, once an inconsistent state estimate is met, the subsequent state estimates are also inconsistent. This is because the filter believes the inconsistent state estimate to be more accurate than it is in reality and hence, it attaches too much weight to this state estimate when processing new measurements.

   Testing for inconsistency is done by consistency tests such as tests on the sum of a number of Normalised Innovation Squared values (SNIS; Willsky 1976, Bar-Shalom and Li 1993).

2. the information content of the state estimate.

   The calculated covariance matrix $\widetilde{\mathbf{P}}_{k|i}$ indicates how uncertain the state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$ is: a large covariance matrix indicates an inaccurate (and little useful) state estimate; the smaller the covariance matrix, the larger the information content of the state estimate.

E.g. both pdfs $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ and $p_2(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ of figure 1 are consistent with $p(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$, however, $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ has a smaller variance, hence it is more informative than $p_2(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ (a smaller domain of x values is probable). The most informative, consistent approximation is the Gaussian with the same mean and covariance as the original distribution, i.e. $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ for the example.

There is a trade-off between consistent and informative state estimates: inconsistency can be avoided by making $\widetilde{\mathbf{P}}_{k|i}$ artificially larger, see equation (14). However, making $\widetilde{\mathbf{P}}_{k|i}$ too large, i.e. larger than necessary for consistency, corresponds to loosing information about the actual accuracy of the state estimate.

The different KF variants linearise the process and the measurement models in the uncertainty region around the state estimate. Consistent estimates are obtained by adding process and measurement uncertainty on the linearised models to compensate for the linearisation errors. In order for the estimates to be informative (i) the linearisation errors need to be as small as possible; and (ii) the extra uncertainty on the linearised models should not be larger than necessary to compensate for these errors. The following sections describe how the Extended Kalman Filter, the Iterated Extended Kalman Filter and the Linear Regression Kalman Filter differ in their linearisation of the process and measurement models; how they take the linearisation errors into account; and how the quality[5] of the state estimates, expressed in terms of consistency and information content, depends on these two choices.

4 Nonlinear process models

This section contains a comparison between the process updates of the different KF variants when dealing with a nonlinear process model (10) with linearisation (12). The KF variants differ by their choice of $\widetilde{\mathbf{F}}_{k-1}$, $\widetilde{\boldsymbol{b}}_{k-1}$ and $\widetilde{\mathbf{Q}}^*_{k-1}$. After linearisation, they all use process update equations[4,p. 7] (3)–(4) to update the state estimate and its uncertainty.

Section 4.1 describes the linearisation of the process model by the EKF and IEKF, section 4.2 by the LRKF. The formulas are summarised in table 1. Section 4.4 presents some examples.

[5]The more nonlinear the behaviour of the process or measurement model in the uncertainty region around the state estimate, the more pronounced the difference in quality performance (consistency and information content of the state estimates) between the KF variants.

[Table 1 about here.]

4.1 The (Iterated) Extended Kalman Filter

The EKF and the IEKF linearise[6] the process model by a first-order Taylor series around the updated state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}$:

$$\widetilde{\mathbf{F}}_{k-1} = \left.\frac{\partial \widetilde{\boldsymbol{f}}_{k-1}}{\partial \widetilde{\boldsymbol{x}}}\right|_{\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}}; \tag{15}$$

$$\widetilde{\boldsymbol{b}}_{k-1} = \widetilde{\boldsymbol{f}}_{k-1}(\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}) - \widetilde{\mathbf{F}}_{k-1}\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}. \tag{16}$$

The basic (I)EKF algorithms do not take the linearisation errors into account (n is the dimension of the state vector $\widetilde{\boldsymbol{x}}$):

$$\widetilde{\mathbf{Q}}^{*}_{k-1} \equiv \widetilde{\mathbf{0}}_{n\times n}. \tag{17}$$

This leads to inconsistent state estimates when these errors can not be neglected.

4.2 The Linear Regression Kalman Filter

The Linear Regression Kalman Filter (LRKF) uses the function values of r regression points $\widetilde{\boldsymbol{\mathcal{X}}}^{j}_{k-1|k-1}$ in state space to model the behaviour of the process function in the uncertainty region around the updated state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}$. The regression points are chosen such that their mean and covariance matrix equal the state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}$ and its covariance matrix $\widetilde{\mathbf{P}}_{k-1|k-1}$. The CDF, DD1 and UKF filters correspond to specific choices. The function values of the regression points are

$$\widetilde{\boldsymbol{\mathcal{X}}}^{j}_{k|k-1} = \widetilde{\boldsymbol{f}}_{k-1}(\widetilde{\boldsymbol{\mathcal{X}}}^{j}_{k-1|k-1}). \tag{18}$$

The LRKF algorithm uses a linearised process function (12) where $\widetilde{\mathbf{F}}_{k-1}$, $\widetilde{\boldsymbol{b}}_{k-1}$ and $\widetilde{\mathbf{Q}}^{*}_{k-1}$ are obtained by statistical linear regression through the $\left(\widetilde{\boldsymbol{\mathcal{X}}}^{j}_{k-1|k-1}, \widetilde{\boldsymbol{\mathcal{X}}}^{j}_{k|k-1}\right)$ points, $j = 1, \ldots, r$; i.e. the deviations $\widetilde{\boldsymbol{e}}_{j}$ between the function values in $\widetilde{\boldsymbol{\mathcal{X}}}^{j}_{k-1|k-1}$ for the nonlinear and the

[6] The EKF and IEKF only differ in their measurement update (sections 5.1–5.2).

linearised function are minimised in least-squares sense:

$$\widetilde{\boldsymbol{e}}_j = \widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1} - \left(\widetilde{\mathbf{F}}\widetilde{\boldsymbol{\mathcal{X}}}^j_{k-1|k-1} + \widetilde{\boldsymbol{b}}\right); \tag{19}$$

$$(\widetilde{\mathbf{F}}_{k-1}, \widetilde{\boldsymbol{b}}_{k-1}) = \arg\min_{(\widetilde{\mathbf{F}},\widetilde{\boldsymbol{b}})} \sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j^T \widetilde{\boldsymbol{e}}_j. \tag{20}$$

The sample covariance of the deviations $\widetilde{\boldsymbol{e}}_j$

$$\widetilde{\mathbf{Q}}^*_{k-1} = \frac{1}{r}\sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j \widetilde{\boldsymbol{e}}_j^T \tag{21}$$

gives an idea of the magnitude of the linearisation errors in the uncertainty region around $\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}$.

Intuitively we feel that when enough[7] regression points are taken, the state estimates of the LRKF process update are consistent and informative. They are consistent because $\widetilde{\mathbf{Q}}^*_{k-1}$ gives a well founded approximation of the linearisation errors (equation (21)). They are informative because the linearised model is a good approximation of the process model in the uncertainty region around $\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}$ (equations (19)–(20)).

4.3 Extra process uncertainty

In all of the presented filters, the user can decide to add extra process uncertainty $\widetilde{\mathbf{Q}}^*_{k-1}$ (or to multiply the calculated covariance matrix $\widetilde{\mathbf{P}}_{k|k-1}$ by a fading factor larger than 1, Bar-Shalom and Li 1993). This is useful if the basic filter algorithm is not consistent. E.g. this is the case for the (I)EKF or for an LRKF with a number of regression points too limited to capture the nonlinear behaviour of the process model in the uncertainty region around $\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}$.

For a particular problem, values for $\widetilde{\mathbf{Q}}^*_{k-1}$ that result in consistent and informative state estimates are obtained by off-line tuning or on-line parameter learning (adaptive filtering, Mehra 1972). In many practical cases consistency is assured by taking the added uncertainty too large, e.g. by taking a constant $\widetilde{\mathbf{Q}}^*$ over time which compensates for decreasing linearisation errors. This, however, results in less informative estimates.

[7] This depends on the nonlinearity of the model in the uncertainty region around the state estimate. A possible approach is to increase the number of regression points until the resulting linearisation (with error covariance) does not change any more. Of course, because the true pdf is unknown, it is not possible to guarantee that the iteration has converged to a set of regression points representative for the model behaviour in the uncertainty region.

4.4 Examples

The different process updates are illustrated by a simple 2D nonlinear process model ($\widetilde{\boldsymbol{x}}(i)$ denotes the ith element of $\widetilde{\boldsymbol{x}}$):

$$\begin{cases} \widetilde{\boldsymbol{x}}_k(1) &= (\widetilde{\boldsymbol{x}}_{k-1}(1))^2\,; \\ \widetilde{\boldsymbol{x}}_k(2) &= \widetilde{\boldsymbol{x}}_{k-1}(1) + 3\;\widetilde{\boldsymbol{x}}_{k-1}(2); \end{cases} \tag{22}$$

with no process uncertainty: $\widetilde{\boldsymbol{\rho}}_{p,k-1} \equiv \widetilde{\boldsymbol{0}}_{2\times 1}$. $\widetilde{\boldsymbol{x}}_k(1)$ depends nonlinearly on $\widetilde{\boldsymbol{x}}_{k-1}$. The process update of $\widetilde{\boldsymbol{x}}_k(2)$ is linear. The updated state estimate and its uncertainty at time step $k-1$ are:

$$\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1} = \begin{bmatrix} 10 \\ 15 \end{bmatrix}; \qquad \widetilde{\mathbf{P}}_{k-1|k-1} = \begin{bmatrix} 36 & 0 \\ 0 & 3600 \end{bmatrix}. \tag{23}$$

Monte Carlo simulation The mean value and the covariance of the true pdf $p(x_k|\widetilde{\hat{\mathbf{Z}}}_{k-1})$ are calculated with a (computationally expensive) Monte Carlo simulation based on a Gaussian pdf $p(x_{k-1}|\widetilde{\hat{\mathbf{Z}}}_{k-1})$ with mean $\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}$ and covariance matrix $\widetilde{\mathbf{P}}_{k-1|k-1}$. The results of this computation are used to illustrate the (in)consistency and information content of the state estimates of the different KF variants. The mean and covariance matrix of the $p(x_k|\widetilde{\hat{\mathbf{Z}}}_{k-1})$ calculated by the Monte Carlo algorithm are:

$$\widetilde{\hat{\boldsymbol{x}}}_{k|k-1} = \begin{bmatrix} 136 \\ 55 \end{bmatrix}; \qquad \widetilde{\mathbf{P}}_{k|k-1} = \begin{bmatrix} 16994 & 721 \\ 721 & 32436 \end{bmatrix}. \tag{24}$$

(I)EKF

[Figure 2 about here.]

Figure 2 shows the updated and (I)EKF predicted state estimates and their uncertainty ellipses[8]. The dotted line is the uncertainty ellipse of the distribution obtained by Monte

[8]The uncertainty ellipsoid

$$(\widetilde{\boldsymbol{x}}_k - \widetilde{\hat{\boldsymbol{x}}}_{k|i})^T \widetilde{\mathbf{P}}_{k|i}^{-1} (\widetilde{\boldsymbol{x}}_k - \widetilde{\hat{\boldsymbol{x}}}_{k|i}) = 1 \tag{25}$$

is a graphical representation of the uncertainty on the state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$. Starting from the point $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$, the distance to the ellipse in a direction is a measure for the uncertainty on $\widetilde{\hat{\boldsymbol{x}}}_{k|i}$ in that direction.

Carlo simulation. The IEKF state prediction and its covariance matrix are:

$$\widetilde{\hat{\boldsymbol{x}}}_{k|k-1} = \begin{bmatrix} 100 \\ 55 \end{bmatrix}; \qquad \widetilde{\mathbf{P}}_{k|k-1} = \begin{bmatrix} 14400 & 720 \\ 720 & 32436 \end{bmatrix}. \tag{26}$$

Due to the neglected linearisation errors, the state estimate is inconsistent: the covariance $\widetilde{\mathbf{P}}_{k|k-1}$ is smaller than the covariance calculated by the Monte Carlo simulation for the first state component $\widetilde{\boldsymbol{x}}(1)$ which had a nonlinear process update. For consistent results this covariance should even be larger because the IEKF estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$ differs from the mean of the pdf $p(x_k|\widetilde{\hat{\mathbf{Z}}}_{k-1})$, calculated by the Monte Carlo simulation.

LRKF

[Figure 3 about here.]

Figure 3 shows the $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k-1|k-1}$ points (top figure) and the $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$ points (bottom figure), the updated state estimate (top) and the predicted state estimate (bottom) and their uncertainty ellipses for the LRKF. The $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k-1|k-1}$ points are chosen with the UKF algorithm of Julier and Uhlmann (1996) where $\kappa = 3 - n = 1$. This corresponds to choosing 6 regression points, including 2 times the point $\widetilde{\hat{\boldsymbol{x}}}_{k-1|k-1}$. The uncertainty ellipse obtained by Monte Carlo simulation coincides with the final uncertainty ellipse of the LRKF predicted state estimate (bottom figure). This indicates consistent and informative results. The LRKF predicted state estimate and its covariance matrix are

$$\widetilde{\hat{\boldsymbol{x}}}_{k|k-1} = \begin{bmatrix} 136 \\ 55 \end{bmatrix}; \qquad \widetilde{\mathbf{P}}_{k|k-1} = \begin{bmatrix} 16992 & 720 \\ 720 & 32436 \end{bmatrix}. \tag{27}$$

4.5 Conclusion: the process update

The LRKF performs better than the (I)EKF when dealing with nonlinear process functions:

1. the LRKF linearises the function based on its behaviour in the uncertainty region around the updated state estimate. The (I)EKF on the other hand only uses the function evaluation and its Jacobian in this state estimate.

2. the LRKF deals with linearisation errors in a theoretically founded way (provided that enough regression points are chosen). The (I)EKF on the other hand needs trial and error

for each particular example to obtain good values for the covariance matrix representing the linearisation errors.

3. unlike the (I)EKF, the LRKF does not need the function Jacobian. This is an advantage where this Jacobian is difficult to obtain or non-existing (e.g. for discontinuous process functions).

5 Nonlinear measurement models

The previous section contains a comparison between the (I)EKF and LRKF process updates; this section focuses on their measurement updates for a nonlinear measurement model (11) with linearisation (13). The EKF, IEKF and LRKF choose $\widetilde{\mathbf{H}}_k$, $\widetilde{\boldsymbol{d}}_k$ and $\widetilde{\boldsymbol{R}}_k^*$ in a different way. After linearisation they use the KF update equations[4,p. 7] (5)–(9).

The linearisation of the measurement model by the IEKF (section 5.2) takes the measurement into account; the EKF (section 5.1) and LRKF (section 5.3) linearise the measurement model based only on the predicted state estimate and its uncertainty. For the latter filters, the linearisation errors ($\widetilde{\mathbf{R}}_k^*$) are larger, especially when the measurement function is quite nonlinear in the uncertainty region around the predicted state estimate. A large uncertainty on the linearised measurement model $\widetilde{\mathbf{R}}_k^* + \widetilde{\mathbf{E}}_k \widetilde{\mathbf{R}}_k \widetilde{\mathbf{E}}_k^T$ (due to a large uncertainty on the state estimate) results in throwing away the greater part of the information of the possibly very accurate measurement. The different linearisation formulas are summarised in table 2. Section 5.5 presents some examples.

[Table 2 about here.]

5.1 The Extended Kalman Filter

The EKF linearises the measurement model around the predicted state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$:

$$\widetilde{\mathbf{H}}_k = \left. \frac{\partial \widetilde{\boldsymbol{h}}_k}{\partial \widetilde{\boldsymbol{x}}} \right|_{\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}} ; \tag{28}$$

$$\widetilde{\boldsymbol{d}}_k = \widetilde{\boldsymbol{h}}_k(\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}) - \widetilde{\mathbf{H}}_k \widetilde{\hat{\boldsymbol{x}}}_{k|k-1}. \tag{29}$$

The basic EKF algorithm does not take the linearisation errors into account:

$$\widetilde{\mathbf{R}}_k^* \equiv \widetilde{\boldsymbol{0}}_{m\times m}; \tag{30}$$

where m is the dimension of the measurement vector $\widetilde{\boldsymbol{z}}_k$. If the measurement model is nonlinear in the uncertainty region around the predicted state estimate, the linearisation errors are not negligible. This means that the linearised measurement model does not reflect the relation between the true state value and the measurement. I.e. the true state value is 'far from'[9] the linearised measurement model. After processing the measured value, given the linear measurement model and the measurement uncertainty, the state is believed to be in a region which does not include the true state estimate, i.e. the updated state estimate is inconsistent.

5.2 The Iterated Extended Kalman Filter

The EKF of the previous section, linearises the measurement model around the predicted state estimate. The IEKF tries to do better by linearising the measurement model around the updated state estimate:

$$\widetilde{\mathbf{H}}_k = \left.\frac{\partial \widetilde{\boldsymbol{h}}_k}{\partial \widetilde{\boldsymbol{x}}}\right|_{\widetilde{\hat{\boldsymbol{x}}}_{k|k}}; \tag{31}$$

$$\widetilde{\boldsymbol{d}}_k = \widetilde{\boldsymbol{h}}_k(\widetilde{\hat{\boldsymbol{x}}}_{k|k}) - \widetilde{\mathbf{H}}_k \widetilde{\hat{\boldsymbol{x}}}_{k|k}. \tag{32}$$

This is achieved by iteration: the filter first linearises the model around a value $\widetilde{\hat{\boldsymbol{x}}}_{k|k}^0$ (often taken equal to the predicted state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$) and calculates the updated state estimate. Then, the filter linearises the model around the newly obtained estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k}^1$ and calculates a new updated state estimate (based on $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$, $\widetilde{\mathbf{P}}_{k|k-1}$ and the new linearised model). This process is iterated until a state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k}^i$ is found for which $\widetilde{\hat{\boldsymbol{x}}}_{k|k}^i$ is close to $\widetilde{\hat{\boldsymbol{x}}}_{k|k}^{i-1}$. The state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$ and uncertainty $\widetilde{\mathbf{P}}_{k|k}$ are calculated starting from the state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$ with its uncertainty $\widetilde{\mathbf{P}}_{k|k-1}$ and the measurement model linearised around $\widetilde{\hat{\boldsymbol{x}}}_{k|k}^i$.

Like the EKF algorithm, the basic IEKF algorithm does not take the linearisation errors into account:

$$\widetilde{\mathbf{R}}_k^* \equiv \widetilde{\boldsymbol{0}}_{m\times m}. \tag{33}$$

[9] 'Far from' (and 'close to') must be understood as: the deviation of the true state with respect to the linearised measurement model is not justified (is justified) by the measurement uncertainty $\widetilde{\mathbf{R}}_k^* + \widetilde{\mathbf{E}}_k \widetilde{\mathbf{R}}_k \widetilde{\mathbf{E}}_k^T$.

If the measurement model is nonlinear in the uncertainty region around the updated state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$, state estimates will be inconsistent. In case of a measurement model that instantaneously fully observes the state (or at least the part of the state that causes the nonlinearities in the measurement model), the linearisation errors will be small[10] in the uncertainty region around $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$. The true state estimate is then 'close to' the linearised measurement function and the updated state estimate is consistent. The result is also informative because no uncertainty due to linearisation errors needs to be added.

5.3 The Linear Regression Kalman Filter

The LRKF evaluates the measurement function in r regression points $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$ in the uncertainty region around the predicted state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$. The $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$ are chosen such that their mean and covariance matrix are equal to the predicted state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$ and its covariance $\widetilde{\mathbf{P}}_{k|k-1}$. The CDF, DD1 and UKF filters correspond to specific choices. The function values of the regression points through the nonlinear function are

$$\widetilde{\boldsymbol{\mathcal{Z}}}^j_k = \widetilde{\boldsymbol{h}}_k(\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}). \tag{34}$$

The LRKF algorithm uses a linearised measurement function (13) where $\widetilde{\mathbf{H}}_k$, $\widetilde{\boldsymbol{d}}_k$ and $\widetilde{\mathbf{R}}^*_k$ are obtained by statistical linear regression through the points $\left(\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}, \widetilde{\boldsymbol{\mathcal{Z}}}^j_k\right)$, $j = 1, \ldots, r$. The statistical linear regression is such that the deviations $\widetilde{\boldsymbol{e}}_j$ between the nonlinear and the linearised function in the regression points are minimised in least-squares sense:

$$\widetilde{\boldsymbol{e}}_j = \widetilde{\boldsymbol{\mathcal{Z}}}^j_k - \left(\widetilde{\mathbf{H}}\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1} + \widetilde{\boldsymbol{d}}\right); \tag{35}$$

$$(\widetilde{\mathbf{H}}_k, \widetilde{\boldsymbol{d}}_k) = \arg\min_{(\widetilde{\mathbf{H}}, \widetilde{\boldsymbol{d}})} \sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j^T \widetilde{\boldsymbol{e}}_j. \tag{36}$$

The sample covariance matrix of the deviations $\widetilde{\boldsymbol{e}}_j$ gives an idea of the magnitude of the linearisation errors:

$$\widetilde{\mathbf{R}}^*_k = \frac{1}{r} \sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j \widetilde{\boldsymbol{e}}_j^T. \tag{37}$$

[10] This assumes that the iterations lead to an accurate $\widetilde{\hat{\boldsymbol{x}}}^i_{k|k}$. The linearisations are started around a freely chosen $\widetilde{\hat{\boldsymbol{x}}}^0_{k|k}$. In order to assure quick and correct iteration, (part of) this value can be chosen based on the measurement information if this information is more accurate than the predicted state estimate.

Intuitively we feel that when enough[7, p. 11] regression points ($\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$, $\widetilde{\boldsymbol{\mathcal{Z}}}^j_k$) are taken the state estimates of the LRKF measurement update are consistent because $\widetilde{\mathbf{R}}^*_k$ gives a well founded approximation of the linearisation errors (equation (37)). However, if the measurement model is highly nonlinear in the uncertainty region around $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$, the ($\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$, $\widetilde{\boldsymbol{\mathcal{Z}}}^j_k$) points deviate substantially from a hyperplane. This results in a large $\widetilde{\mathbf{R}}^*_k$ and non-informative updated state estimates (see the example in section 5.5).

5.4 Extra measurement uncertainty

In order to make the state estimates consistent, the user can tune an inconsistent filter by adding extra measurement uncertainty $\widetilde{\mathbf{R}}^*_k$.

Only off-line tuning or on-line parameter learning can lead to a good value for $\widetilde{\mathbf{R}}^*_k$ for a particular problem. In many practical cases consistency is assured by choosing the added uncertainty too large, e.g. by taking a constant $\widetilde{\mathbf{R}}^*$ over time which compensates for decreasing linearisation errors. This reduces the information content of the results.

5.5 Examples

First example

The comparison between the different measurement updates is illustrated with the measurement function $z_k = h_1(\widetilde{\boldsymbol{x}}_k) + \rho_{m,k}$;

$$h_1(\widetilde{\boldsymbol{x}}_k) = (\widetilde{\boldsymbol{x}}_k(1))^2 + (\widetilde{\boldsymbol{x}}_k(2))^2 \,. \tag{38}$$

$\widetilde{\boldsymbol{x}}_k = \begin{bmatrix} 15 \\ 20 \end{bmatrix}$ is the true value and $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1} = \begin{bmatrix} 10 \\ 15 \end{bmatrix}$ is the predicted state estimate with covariance matrix $\widetilde{\mathbf{P}}_{k|k-1} = \begin{bmatrix} 36 & 0 \\ 0 & 3600 \end{bmatrix}$. The processed measurement is $\hat{z}_k = 630$ and the measurement covariance is $R_k = 400$.

Second example

To illustrate the consistency of the state estimate of an IEKF when the measurement observes

the state completely, a second example is used. The measurement function is

$$\widetilde{\boldsymbol{z}}_k = \widetilde{\boldsymbol{h}}(\widetilde{\boldsymbol{x}}_k) + \widetilde{\boldsymbol{\rho}}_{m,k} = \begin{bmatrix} h_1(\widetilde{\boldsymbol{x}}_k) + \widetilde{\boldsymbol{\rho}}_{m,k}(1) \\ h_2(\widetilde{\boldsymbol{x}}_k) + \widetilde{\boldsymbol{\rho}}_{m,k}(2) \end{bmatrix}; \tag{39}$$

with

$$\begin{cases} h_1(\widetilde{\boldsymbol{x}}_k) &= (\widetilde{\boldsymbol{x}}_k(1))^2 + (\widetilde{\boldsymbol{x}}_k(2))^2; \\ h_2(\widetilde{\boldsymbol{x}}_k) &= 3\,(\widetilde{\boldsymbol{x}}_k(2))^2 / \widetilde{\boldsymbol{x}}_k(1); \end{cases} \tag{40}$$

$\widetilde{\boldsymbol{x}}_k = \begin{bmatrix} 15 \\ 20 \end{bmatrix}$ is the true value and $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1} = \begin{bmatrix} 10 \\ 15 \end{bmatrix}$ is the predicted state estimate with covariance matrix $\widetilde{\mathbf{P}}_{k|k-1} = \begin{bmatrix} 36 & 0 \\ 0 & 3600 \end{bmatrix}$. The processed measurement and the measurement covariance matrix are:

$$\widetilde{\hat{\boldsymbol{z}}}_k = \begin{bmatrix} 630 \\ 85 \end{bmatrix}; \qquad \widetilde{\mathbf{R}}_k = \begin{bmatrix} 400 & 0 \\ 0 & 400 \end{bmatrix}. \tag{41}$$

In all figures, the true state value $\widetilde{\boldsymbol{x}}_k$ is plotted; if this value is 'far' outside the uncertainty ellipse of a state estimate, the corresponding estimate is inconsistent. Because the measurement is accurate and the initial estimate is not, the uncertainty on the state estimate should drop considerably when the measurement is processed. The updated state estimate is not informative if this is not the case.

EKF

[Figure 4 about here.]

Figure 4 shows the state estimates, uncertainty ellipses and measurement functions for the EKF applied to the first example (equation (38)). The true measurement function is nonlinear. $\widetilde{\boldsymbol{x}}_k$ is the true value of the state, and is 'close to' this function. The linearisation around the uncertain predicted state estimate is not a good approximation of the function around the true state value: the true state value is 'far from' the linearised measurement function. The

resulting updated state estimate

$$\widetilde{\hat{\boldsymbol{x}}}_{k|k} = \begin{bmatrix} 10 \\ 25 \end{bmatrix}; \qquad \widetilde{\mathbf{P}}_{k|k} = \begin{bmatrix} 36 & -24 \\ -24 & 16 \end{bmatrix}; \tag{42}$$

is inconsistent.

IEKF

[Figure 5 about here.]

Figure 5 shows the measurement function, the linearised measurement function around the point $\widetilde{\hat{\boldsymbol{x}}}^{i}_{k|k}$, the true state value $\widetilde{\boldsymbol{x}}_k$ and the state estimates for the IEKF applied to the first example (equation (38)). The measurement model does not fully observe the state. This results in an uncertain updated state estimate $\widetilde{\hat{\boldsymbol{x}}}^{i}_{k|k}$ around which the filter linearises the measurement function. As was the case for the EKF, the linearisation errors are not negligible and the true value is 'far from' the linearised measurement function. The updated state estimate

$$\widetilde{\hat{\boldsymbol{x}}}_{k|k} = \begin{bmatrix} 10 \\ 23 \end{bmatrix}; \qquad \widetilde{\mathbf{P}}_{k|k} = \begin{bmatrix} 36 & -16 \\ -16 & 7.0 \end{bmatrix}; \tag{43}$$

is inconsistent.

[Figure 6 about here.]

If however the measurement model fully observes the state, the IEKF updated state estimate is accurately known; hence, the linearisation errors are small and the true state value is 'close to' the linearised measurement function. In this case, the updated state estimate is consistent. Figure 6 shows the measurement function, the linearised measurement function, the true state value $\widetilde{\boldsymbol{x}}_k$, the state estimates and the uncertainty ellipses for the IEKF applied to the second example (equations (39)–(40)). The updated state estimate and covariance matrix

$$\widetilde{\hat{\boldsymbol{x}}}_{k|k} = \begin{bmatrix} 14 \\ 21 \end{bmatrix}; \qquad \widetilde{\mathbf{P}}_{k|k} = \begin{bmatrix} 2.6 & -1.7 \\ -1.7 & 1.3 \end{bmatrix}; \tag{44}$$

are consistent and informative due to the small, ignored, linearisation errors.

LRKF

[Figure 7 about here.]

[Figure 8 about here.]

[Figure 9 about here.]

A LRKF is run on the first example (equation (38)). The $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$ points are chosen with the UKF algorithm with $\kappa = 3 - n = 1$ (Julier and Uhlmann 1996). This corresponds to choosing 6 regression points, including 2 times the point $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$. Figures 7 and 8 show the nonlinear measurement function, the $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$-points and the linearisation. The predicted state estimate is uncertain, hence the $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$-points are widespread. R^*_k is large ($R^*_k = 2.6 \times 10^7$) due to the large deviations between the ($\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$, $\widetilde{\boldsymbol{\mathcal{Z}}}^j_k$) points and the linearised measurement function (see figure 8). The updated state estimate and its covariance matrix are

$$\widetilde{\hat{\boldsymbol{x}}}_{k|k} = \begin{bmatrix} 10 \\ 2.6 \end{bmatrix}; \qquad \widetilde{\mathbf{P}}_{k|k} = \begin{bmatrix} 36 & 0 \\ 0 & 3600 \end{bmatrix}. \tag{45}$$

Figure 9 shows the $\widetilde{\boldsymbol{\mathcal{X}}}^j_{k|k-1}$ points, the measurement function, the LRKF linearised measurement function, the true state value $\widetilde{\boldsymbol{x}}_k$, the state estimates and the uncertainty ellipses. The updated state estimate is consistent, however, it can hardly be called an improvement over the previous state estimate ($\widetilde{\mathbf{P}}_{k|k} \approx \widetilde{\mathbf{P}}_{k|k-1}$). The information in the measurement is neglected due to the high 'measurement uncertainty' $\widetilde{\mathbf{R}}^*_k + \widetilde{\mathbf{E}}_k \widetilde{\mathbf{R}}_k \widetilde{\mathbf{E}}^T_k$ on the linearised function.

Note that some kind of Iterated LRKF (similar to the Iterated EKF) would not solve this problem: the updated state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$ and its covariance matrix $\widetilde{\mathbf{P}}_{k|k}$ are more or less the same as the predicted state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$ and its covariance matrix $\widetilde{\mathbf{P}}_{k|k-1}$. Hence, the regression points and the linearisation would approximately be the same after iteration.

5.6 Conclusion: the measurement update

Measurements which fully observe the part of the state that makes the model nonlinear, are best processed by the IEKF. In this case (and assuming that the algorithm iterates to a good linearisation point[10, p. 16]), the IEKF linearisation errors are negligible.

In the other cases, none of the presented filters outperforms the others. A filter should be chosen for each specific application: the LRKF makes an estimate of its linearisation errors ($\widetilde{\mathbf{R}}_k^*$), the EKF and IEKF on the other hand require off-line tuning or on-line parameter learning of $\widetilde{\mathbf{R}}_k^*$ to yield consistent state estimates. Because the IEKF additionally takes the measurement into account when linearising the measurement model, its linearisation errors are smaller than those of the EKF and LRKF. This means that once a well-tuned IEKF is available, the state estimates it returns can be far more informative than those of the LRKF or a well-tuned EKF.

Finally, note that the LRKF does not use the Jacobian of the measurement function, which makes it possible to process discontinuous measurement functions.

6 Conclusions

This paper gives insight in the advantages and drawbacks of the Extended Kalman Filter (EKF), the Iterated Extended Kalman Filter (IEKF) and the Linear Regression Kalman Filter (LRKF). These insights are a result of the distinct analysis approach taken in this paper:

1. The paper describes all filter algorithms as the application of the basic KF algorithm to linearised process and measurement models. The difference between the KF variants is situated in the choice of linearisation and the compensation of the linearisation errors. In previous work this linearisation was not always recognised, e.g. the UKF is originally derived as a filter which does not linearise the models.

2. The analysis clarifies how some filters automatically adapt their compensation for linearisation errors, while other filters have constant (developer-determined) error compensation.

3. The quality of the state estimates is expressed by two criteria: the consistency and the information content of the estimates. This paper relates the consistency and information content of the estimates to (i) how the linearisation is performed and (ii) how the linearisation errors are taken into account. The understanding of the linearisation processes allows us to make a well-founded choice of filter for a specific application.

4. The performance of the different filters is compared for the process and measurement

updates separately, because a good performance for one of these updates does not necessarily mean a good performance for the other update. This makes it interesting in some cases to use different filters for both updates.

For process updates the LRKF performs better than the other mentioned KF variants because (i) the LRKF linearises the process model based on its behaviour in the uncertainty region around the updated state estimate. The (I)EKF on the other hand only uses the function evaluation and its Jacobian in this state estimate; and (ii) the LRKF deals with linearisation errors in a theoretically founded way, provided that enough regression points are chosen. The (I)EKF on the other hand needs trial and error for each particular application to obtain a good covariance matrix representing the linearisation errors.

The IEKF is the best way to handle nonlinear measurement models that fully observe the part of the state that makes the measurement model nonlinear. In the other cases, none of the presented filters outperforms the others: the LRKF makes an estimation of the linearisation errors, the EKF and IEKF on the other hand require extensive off-line tuning or on-line parameter learning in order to yield consistent state estimates. However, unlike the EKF and LRKF, the IEKF additionally uses the measurement value in order to linearise the measurement model. Hence, its linearisation errors are smaller and once a well-tuned IEKF is available, the state estimates it returns can be far more informative than those of the LRKF or a well-tuned EKF.

The insights described in this paper are important for all researchers and developers who want to apply a KF variant to a specific application. They lead to a systematic choice of a filter, where previously the choice was mainly made based on success in similar applications or based on trial and error. Further work should report on practical applications using these insights and an effort should be made to analyse and include future KF algorithms in this framework.

Acknowledgements

T. Lefebvre is a Postdoctoral Fellow of the Fund for Scientific Research–Flanders (F.W.O.) in Belgium. Financial support by the Belgian Programme on Inter-University Attraction Poles initiated by the Belgian State—Prime Minister's Office—Science Policy Programme (IUAP), and by K.U.Leuven's Concerted Research Action GOA/99/04 are gratefully acknowledged.

References

Bar-Shalom, Y., and Li, X., 1993, Estimation and tracking: principles, techniques and software (Boston, London: Artech House).

Bayes, T., 1763, An essay towards solving a problem in the doctrine of chances. Philosophical Transactions of the Royal Society of London, **53**, 370–418. Reprinted in Biometrika, **45**, 293–315, 1958.

Gelb, A., Kasper, J., Nash, R., Price, C., and Sutherland, A., 1974, Applied optimal estimation (Cambridge, MA: MIT Press).

Julier, S., Uhlman, J., and Durrant-Whyte, H., 2000, A new method for the transformation of means and covariances in filters and estimators. IEEE Transactions on Automatic Control, **45**(3), 477–482.

Julier, S., and Uhlmann, J., 1996, A general method for approximating nonlinear transformations of probability distributions. Technical report, Robotics Research Group, Department of Engineering Science, University of Oxford.

Julier, S., and Uhlmann, J., 2001, Data fusion in non-linear systems. In D. Hall and J. Llinas (eds.), Handbook of Multisensor Data Fusion (CRC Press), Chapter 13.

Kalman, R. E., 1960, A new approach to linear filtering and prediction problems. Transactions of the ASME, Journal of Basic Engineering, **82**, 34–45.

Laplace, P.-S., 1812, Théorie Analytique des Probabilités (Paris: Courcier Imprimeur).

Lefebvre, T., Bruyninckx, H., and De Schutter, J., 2002, Comment on "A new method for the nonlinear transformation of means and covariances in filters and estimators". IEEE Transactions on Automatic Control, **47**(8), 1406–1408.

MAYBECK, P., 1982, Stochastic models, estimation, and control, volume 2 (Academic Press).

MEHRA, R., 1972, Approaches to adaptive filtering. IEEE Transactions on Automatic Control, **17**(5), 693–698.

NØRGAARD, M., POULSEN, N., and RAVN, O., 2000a, Advances in derivative-free state estimation for nonlinear systems. Technical Report IMM-REP-1998-15 (revised edition), Technical University of Denmark, Denmark.

NØRGAARD, M., POULSEN, N., and RAVN, O., 2000b, New developments in state estimations for nonlinear systems. Automatica, **36**(11), 1627–1638.

SCHEI, T., 1997, A finite-difference method for linearisation in nonlinear estimation algorithms. Automatica, **33**(11), 2053–2058.

SORENSON, H. W., 1985, Kalman filtering: theory and application (New York, NY: IEEE Press).

TANIZAKI, H., 1996, Nonlinear filters: estimation and applications — second, revised and enlarged edition (Berlin-Heidelberg: Springer-Verlag).

UHLMANN, J., 1995, Dynamic map building and localization: new theoretical foundations. PhD thesis, University of Oxford, UK.

WILLSKY, A., 1976, A survey of design methods for failure detection in dynamic systems. Automatica, **12**, 601–611.

List of Tables

List of Figures

List of Figures

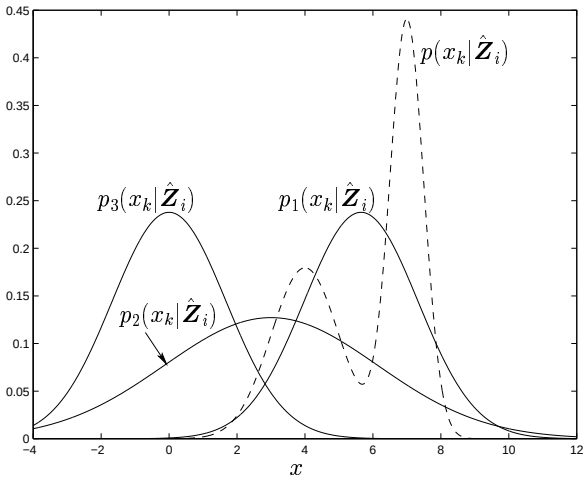

Figure 1: Non-Gaussian pdf $p(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ with three Gaussian approximations $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$, $p_2(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ and $p_3(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$. $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ and $p_2(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ are consistent, $p_3(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ is inconsistent. $p_1(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$ is more informative than $p_2(x_k|\widetilde{\hat{\mathbf{Z}}}_i)$.

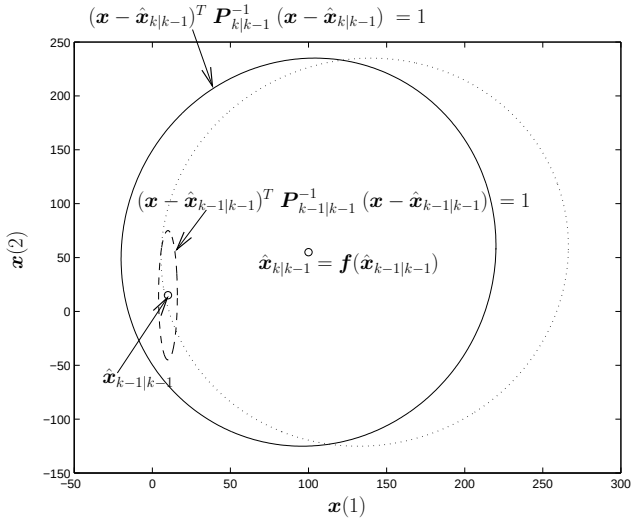

Figure 2: Nonlinear process model. Uncertainty ellipses for the updated state estimate at $k-1$ (dashed line), for the (I)EKF predicted state estimate (full line) and the Monte Carlo uncertainty ellipse (dotted line). The predicted state estimate is inconsistent due to the neglected linearisation errors: the uncertainty ellipse of the IEKF predicted estimate is shifted with respect to the Monte Carlo uncertainty ellipse and is somewhat smaller.

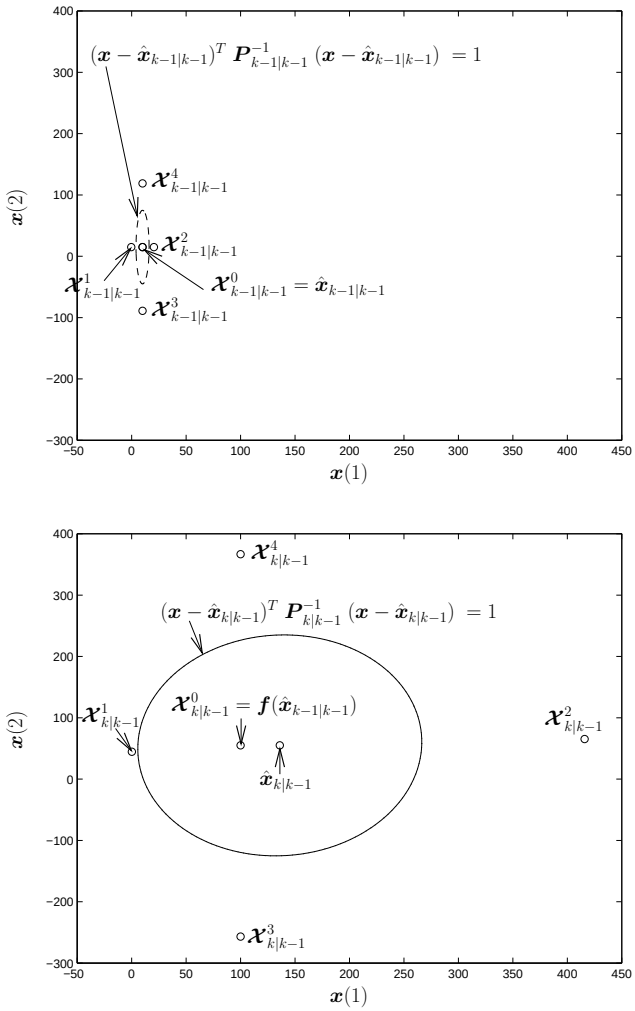

Figure 3: Nonlinear process model. Uncertainty ellipses for the updated state estimate at $k-1$ (dashed line, top figure), for the LRKF predicted state estimate (full line, bottom figure), and Monte Carlo uncertainty ellipse (dotted line which coincides with the full line, bottom figure). The LRKF predicted state estimate is consistent and informative: its uncertainty ellipse coincides with the Monte Carlo uncertainty ellipse.

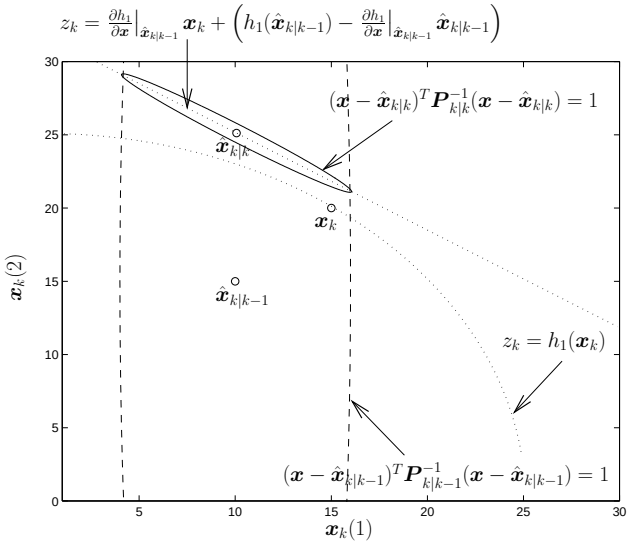

Figure 4: Nonlinear measurement model $z = h_1(\widetilde{\boldsymbol{x}}_k)$ and EKF linearisation around $\widetilde{\hat{\boldsymbol{x}}}_{k|k-1}$ (dotted lines). The true state $\widetilde{\boldsymbol{x}}_k$ is 'far from' this linearisation and the obtained state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$ (uncertainty ellipse in full line) is inconsistent.

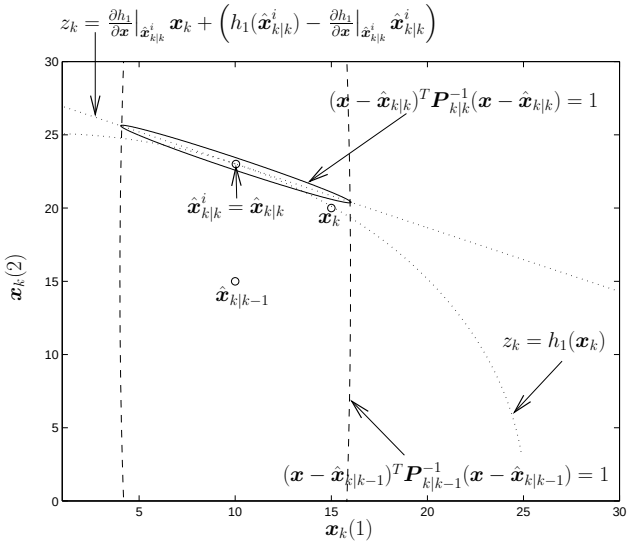

Figure 5: Nonlinear measurement model $z = h_1(\widetilde{\boldsymbol{x}}_k)$ that does not observe the full state, and its IEKF linearisation around $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$ (dotted lines). The true state $\widetilde{\boldsymbol{x}}_k$ is 'far from' this linearisation, leading to an inconsistent state estimate $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$ (uncertainty ellipse in full line).

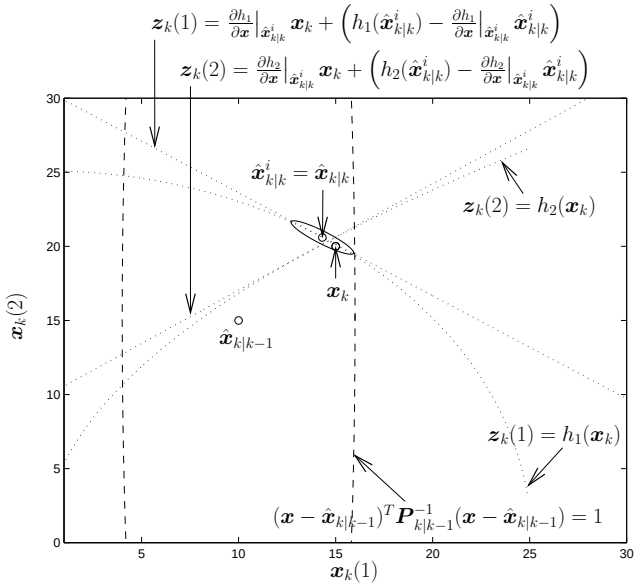

Figure 6: Nonlinear measurement model $\widetilde{\boldsymbol{z}} = \widetilde{\boldsymbol{h}}(\widetilde{\boldsymbol{x}}_k)$ that fully observes the state, and its IEKF linearisation around $\widetilde{\hat{\boldsymbol{x}}}_{k|k}$ (dotted lines). The true state $\widetilde{\boldsymbol{x}}_k$ is 'close to' this linearisation, leading to a consistent state estimate (uncertainty ellipse in full line).

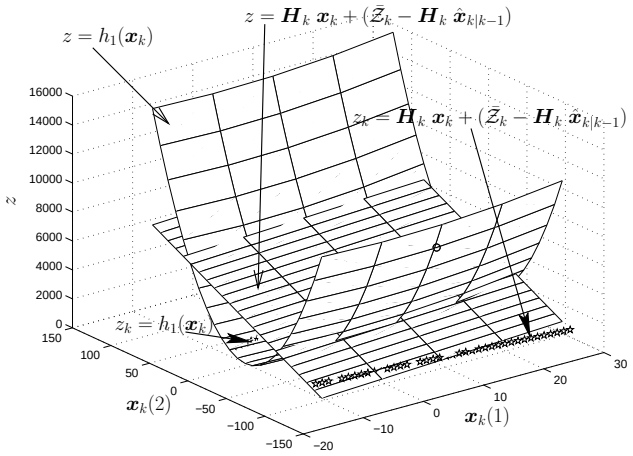

Figure 7: Nonlinear measurement model $z = h_1(\widetilde{\boldsymbol{x}})$ and LRKF linearisation. The linearisation errors are large.

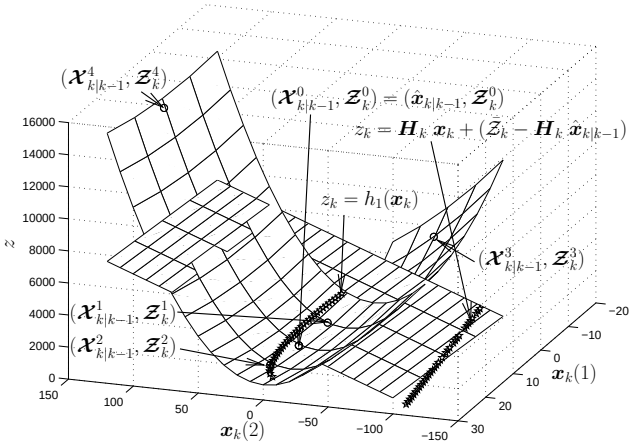

Figure 8: Figure 7 seen from another angle.

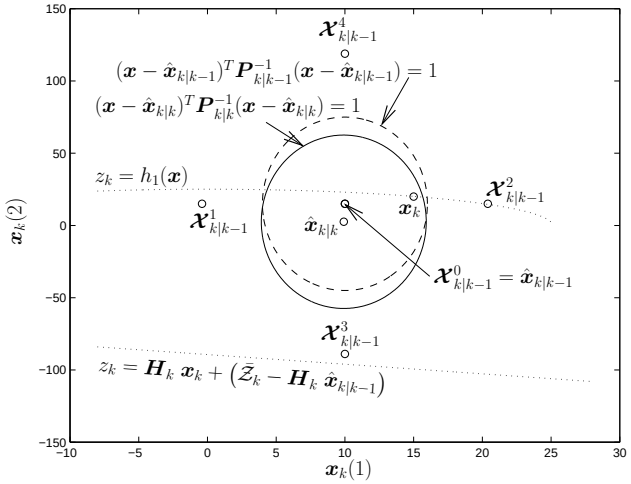

Figure 9: Nonlinear measurement model $z = h_1(\widetilde{\boldsymbol{x}})$ and LRKF linearisation (dotted lines). The large linearisation errors result in a large measurement uncertainty $\widetilde{\mathbf{R}}^*_k + \widetilde{\mathbf{E}}_k\widetilde{\mathbf{R}}_k\widetilde{\mathbf{E}}^T_k$. The updated state estimate (uncertainty ellipse in full line) is consistent but non-informative.

List of Tables

	$\widetilde{\mathbf{F}}_{k-1}$	$\widetilde{\boldsymbol{b}}_{k-1}$	$\widetilde{\mathbf{Q}}^*_{k-1}$
EKF	$\left.\frac{\partial \widetilde{\boldsymbol{f}}_{k-1}}{\partial \widetilde{\boldsymbol{x}}}\right\vert_{\widetilde{\hat{\boldsymbol{x}}}_{k-1\vert k-1}}$	$\widetilde{\boldsymbol{f}}_{k-1}(\widetilde{\hat{\boldsymbol{x}}}_{k-1\vert k-1}) - \widetilde{\mathbf{F}}_{k-1}\widetilde{\hat{\boldsymbol{x}}}_{k-1\vert k-1}$	$\widetilde{\mathbf{0}}_{n\times n}$
IEKF	$\left.\frac{\partial \widetilde{\boldsymbol{f}}_{k-1}}{\partial \widetilde{\boldsymbol{x}}}\right\vert_{\widetilde{\hat{\boldsymbol{x}}}_{k-1\vert k-1}}$	$\widetilde{\boldsymbol{f}}_{k-1}(\widetilde{\hat{\boldsymbol{x}}}_{k-1\vert k-1}) - \widetilde{\mathbf{F}}_{k-1}\widetilde{\hat{\boldsymbol{x}}}_{k-1\vert k-1}$	$\widetilde{\mathbf{0}}_{n\times n}$
LRKF	$\arg\min_{(\widetilde{\mathbf{F}},\widetilde{b})} \sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j^T \widetilde{\boldsymbol{e}}_j$	$\arg\min_{(\widetilde{\mathbf{F}},\widetilde{b})} \sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j^T \widetilde{\boldsymbol{e}}_j$	$\frac{1}{r}\sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j \widetilde{\boldsymbol{e}}_j^T$

Table 1: Summary of the linearisation of the process model by the Extended Kalman Filter (EKF), the Iterated Extended Kalman Filter (IEKF) and the Linear Regression Kalman Filter (LRKF).

	$\widetilde{\mathbf{H}}_k$	$\widetilde{\boldsymbol{d}}_k$	$\widetilde{\mathbf{R}}_k^*$
EKF	$\frac{\partial \widetilde{\boldsymbol{h}}_k}{\partial \widetilde{\boldsymbol{x}}}\Big\vert_{\widetilde{\hat{\boldsymbol{x}}}_{k\vert k-1}}$	$\widetilde{\boldsymbol{h}}_k(\widetilde{\hat{\boldsymbol{x}}}_{k\vert k-1}) - \widetilde{\mathbf{H}}_k \widetilde{\hat{\boldsymbol{x}}}_{k\vert k-1}$	$\widetilde{\mathbf{0}}_{m\times m}$
IEKF	$\frac{\partial \widetilde{\boldsymbol{h}}_k}{\partial \widetilde{\boldsymbol{x}}}\Big\vert_{\widetilde{\hat{\boldsymbol{x}}}_{k\vert k}}$	$\widetilde{\boldsymbol{h}}_k(\widetilde{\hat{\boldsymbol{x}}}_{k\vert k}) - \widetilde{\mathbf{H}}_k \widetilde{\hat{\boldsymbol{x}}}_{k\vert k}$	$\widetilde{\mathbf{0}}_{m\times m}$
LRKF	$\min_{(\widetilde{\mathbf{H}},\widetilde{\boldsymbol{d}})} \sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j^T \widetilde{\boldsymbol{e}}_j$	$\min_{(\widetilde{\mathbf{H}},\widetilde{\boldsymbol{d}})} \sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j^T \widetilde{\boldsymbol{e}}_j$	$\frac{1}{r}\sum_{j=1}^{r} \widetilde{\boldsymbol{e}}_j \widetilde{\boldsymbol{e}}_j^T$

Table 2: Summary of the linearisation of the measurement model by the Extended Kalman Filter (EKF), the Iterated Extended Kalman Filter (IEKF) and the Linear Regression Kalman Filter (LRKF).