

Kalman filters for real-time magnetic island phase tracking

D.P. Borgers^c, M. Lauret^{a,b,*}, M.R. de Baar^{a,b}, the TCV team, the TEXTOR team

^aFOM Institute DIFFER - Dutch Institute for Fundamental Energy Research, Association EURATOM-FOM, Trilateral Euregio Cluster - P.O. Box 1207 Nieuwegein, The Netherlands

^bControl Systems Technology, Department of Mechanical Engineering - Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

^cHybrid and Networked Systems, Department of Mechanical Engineering - Eindhoven University of Technology

Abstract

For control of neoclassical tearing modes (NTMs) and the resulting rotating magnetic islands in tokamak plasmas, the frequency and phase of the magnetic islands need to be accurately tracked in real-time. In previous experiments on TEXTOR, this was achieved using a phase-locked loop (PLL). For ASDEX Upgrade however, the desired frequency range in which the islands are to be tracked (5 Hz – 20 kHz) is much larger than is possible with a PLL. In this contribution an extended Kalman filter (EKF) and an unscented Kalman filter (UKF) are proposed for real-time frequency, phase and amplitude tracking of sinusoidal signals using noisy measurements. Compared to PLLs, the EKF and UKF are able to track magnetic islands in a much larger frequency range. The filters are applied on synthetic data and on experimental data from the TEXTOR and TCV tokamaks, from which we conclude that the UKF can be useful for real-time control of magnetic islands on ASDEX Upgrade.

Keywords: Extended Kalman Filter, Unscented Kalman Filter, magnetic islands, phase tracking, real-time, rotating tearing modes control

1. Introduction

Tokamaks are toroidal devices in which a hot plasma is confined by means a magnetic field. External coils generate the dominant toroidal magnetic field B_ϕ . A toroidal electric current I_p in the plasma adds a smaller poloidal component B_θ . Ideally, the resulting magnetic field has helical field lines that lie on nested surfaces of constant flux.

On some of these surfaces, magnetic islands can develop, caused by the so-called neoclassical tearing modes (NTMs). Due to magnetic reconnection, the topology of nested flux surfaces changes, which degrades the performance and stability of the reactor.

The magnetic axis of the island is referred to as the O-point. The O-point is a local minimum in the current density. In reactor-relevant tokamaks such as ITER, real-time control of the width of these modes is required [1]. The control scheme is based on local non-inductive current drive using

high power mm-wave beams. The beam frequency is a multiple of the electron cyclotron frequency, and the deposition of the beam is a localized, resonant process. As the island is rotating and the deposition location is fixed, modulation of the beam power is required to assure deposition in the islands O-point. This implies that the phase of the island rotation needs to be identified in real-time, with good accuracy.

The island rotation can be inferred from magnetic field measurements outside the plasma or from local electron temperature T_e measurements at the radial location of the island. The repetitive passing of the O-point leads to a repetitive variation of T_e and the magnetic field. The poloidal phase of the island needs to be inferred from these measurements, and made available to a real-time control system. Given the variation of the signal's amplitude and the amount of measurement noise, simple threshold methods cannot be applied. Measurements of the frequency content of the signals and concomitant phase determination in principle require sampling over a number of oscillations. In situations with

*M.Lauret@tue.nl

rapidly varying frequencies this leads to significant errors in the phase measurement. In situations with locking islands this leads to significant latency.

A phase-locked loop (PLL) can be used to estimate frequency and phase of a measured periodic signal in real-time by locking the oscillator of the PLL with the measured signal [2]. Successful O-point tracking and modulation control using a PLL have been achieved in TEXTOR [3]. However, the frequency band for which PLLs work is inherently limited, and the method can only be applied in systems in which the mode frequency is limited within a well known band. On ASDEX Upgrade, a tokamak characterized by strong torque from neutral beams and by strong island-wall interaction, the mode frequency band is from 100 Hz to 10 kHz. This range is much wider than the frequency band of PLLs, which necessitates the development of a viable alternative method, capable of tracking periodic signals in a frequency range of at least 100 Hz to 10 kHz.

In this contribution two real-time methods (an extended and an unscented Kalman filter), for accurate and low latency tracking of oscillating signals with variable frequency, amplitude and phase are discussed. The frequency band for which these methods works is, in principle, only bounded by the Nyquist frequency (half the sample frequency).

The paper is organised as follows. In Section 2, the signal model is discussed. In Section 3, theory and principles of Kalman filters are explained. In Section 4, the methods are tested on synthetic data, and in Section 5, the methods are tested on off-line data of magnetic field and electron temperature measurements.

2. Signal model

The Kalman filters that we will propose in Section 3 estimate the states of a system, based on its measured output and a state-space model of the system.

Because of the repetitive passing of the O-point of the magnetic islands, the magnetic field measurements or electron temperature measurements are roughly sinusoidal signals. Therefore we do not model fusion processes, but we simply derive a state-space model that describes a cosine. Of course, this model can be used to track many other sinusoidal signals, such as, for example, Alfvén waves. The filters might also be able to track a sawtooth using a different signal model, such as a

sawtooth model as described in [4] or [5]. Many other signals (including nonperiodical signals) can be tracked using an appropriate state-space model.

We assume that the measured signal can be described by

$$y(t) = A(t) \cos(\theta(t)) + v(t). \quad (1)$$

Here y is the measured signal, A the amplitude of the sinusoid, θ its phase in radians and v the measurement noise.

In discrete time $t_k = kt_s$, where t_s is the constant sample time and $k \in \mathbb{Z}$, the signal is written as

$$y_k = A_k \cos(\theta_k) + v_k. \quad (2)$$

To track this signal using a Kalman filter, a state-space model is needed with states that describe the evolution of the amplitude A_k , phase θ_k and frequency $f_k = (\theta_{k+1} - \theta_k)/(2\pi t_s)$. Here the following nonlinear discrete-time model is used, that at time t_k has states

$$x_k = \begin{pmatrix} x_{1,k} \\ x_{2,k} \\ x_{3,k} \end{pmatrix} = \begin{pmatrix} A_k \cos(\theta_k) \\ A_k \sin(\theta_k) \\ 2\pi f_k t_s \end{pmatrix}. \quad (3)$$

Assuming that the amplitude and frequency are almost constant in the small time-interval $(t_k, t_{k+1}]$, state x_1 at time t_{k+1} becomes

$$\begin{aligned} x_{1,k+1} &= A_{k+1} \cos(\theta_{k+1}) \\ &= A_k \cos(\theta_k + 2\pi f_k t_s) + w_{1,k} \\ &= x_{1,k} \cos(x_{3,k}) - x_{2,k} \sin(x_{3,k}) + w_{1,k}, \end{aligned} \quad (4)$$

where the unmodelled variations in amplitude and frequency are included in the process noise $w_{1,k}$.

A similar derivation can be used to find $x_{2,k+1}$. The result is the nonlinear state-space model for x_{k+1} together with output y_k

$$\begin{aligned} x_{k+1} &= \begin{pmatrix} x_{1,k} \cos(x_{3,k}) - x_{2,k} \sin(x_{3,k}) \\ x_{1,k} \sin(x_{3,k}) + x_{2,k} \cos(x_{3,k}) \\ x_{3,k} \end{pmatrix} + w_k \\ &= f(x_k) + w_k, \\ y_k &= x_{1,k} + v_k. \end{aligned} \quad (5)$$

The measured signal to be estimated has time-varying frequency, phase and amplitude. These variances, as well as model errors, are accounted for in the process noise w_k .

The resulting model (5) is observable [6] for $x_1, x_2 > 0$, $0 < x_3 < \pi$, which implies that

the states of the system can be inferred from the measured output in the frequency range $A > 0$, $0 < f < f_s/2$, where f_s is the sample frequency.

As shown in [7], multiple sinusoids can be tracked simultaneously by expanding the signal model to include multiple sinusoids with different frequencies and amplitudes.

In [8] an extra state $x_4 = \alpha$ was proposed, where α is the rate of change of the amplitude, such that $A_{k+1} = \alpha_k A_k$. Something similar can be done for the frequency tracking using $f_{k+1} = f_k + \Delta_{f,k}$, where $\Delta_{f,k}$ is the estimated frequency increase.

Depending on the signal to be estimated, these extensions might improve the performance of the filter.

Since no knowledge of the fusion processes is used, the model is very general and can be used to describe any sinusoidal signal. If a model of the physical process that produces the periodic signal is available, this can also be included into the signal model, which will have a positive effect on the stability and performance of the filter.

3. Kalman filters

For the problem of tracking the island phase, frequency and amplitude, an observer is needed that reliably estimates the states of the system (and thereby the frequency, phase and amplitude), and is insensitive to measurement noise. The observers that are investigated here are the extended Kalman filter (EKF) [9] and unscented Kalman filter (UKF) [10], which are both extensions of the Kalman filter to nonlinear systems.

A general description of a time-invariant discrete-time dynamical system is given by

$$\begin{aligned} x_{k+1} &= f(x_k, u_k, w_k), \\ y_k &= h(x_k, v_k), \end{aligned} \quad (6)$$

where at time t_k , x_k is the state vector, u_k the input, w_k the process noise, y_k the output and v_k the measurement noise, and the functions f and h are independent of time (time-invariant). In general, x_k , u_k , w_k , y_k and v_k are vectors.

In this article it is assumed that w and v are

zero-mean noise (7), (8) and are uncorrelated (11):

$$\mathbb{E}\{w_k\} = 0, \quad (7)$$

$$\mathbb{E}\{v_k\} = 0, \quad (8)$$

$$\mathbb{E}\{w_k w_k^\top\} = Q_k, \quad (9)$$

$$\mathbb{E}\{v_k v_k^\top\} = R_k, \quad (10)$$

$$\mathbb{E}\{w_k v_k^\top\} = 0. \quad (11)$$

Here, $\mathbb{E}\{\cdot\}$ denotes the expectation of (\cdot) .

For the model presented in Section 2, y_k and v_k are scalars and the functions f and h are nonlinear and time-invariant.

The Kalman filter was first described in [11] and can be used to estimate the states of a linear time-invariant (LTI) discrete-time systems, described by the state-space model of Eq. (6), with functions f and h linear in their parameters.

The Kalman filter is an iterative filter and consists of two steps per iteration. First the model of the plant is used to predict the *a priori* state and output estimates \hat{x}_k^- and \hat{y}_k^- , and second the *a priori* state estimate is corrected based on the measured output of the plant:

$$\hat{x}_k = \hat{x}_k^- + K_k(y_k - \hat{y}_k^-). \quad (12)$$

The Kalman gain K_k is based on the noise covariance matrices Q_k and R_k , and on P_k , which is an iterative estimate of $\mathbb{E}\{ee^\top\}$, where $e = x - \hat{x}$ is the state estimate error. For the complete algorithm of the Kalman filter, see the EKF (Algorithm 3.1), which reduces straightforwardly to the Kalman filter when the system is linear.

For LTI systems with zero-mean noise, the Kalman filter is the optimal linear filter in the least-squares sense, meaning that no other linear filter can achieve a smaller error e in a least-squares sense than the Kalman filter. Since the Kalman filter only assumes zero-mean noise, the measurement noise need not be white, but can also be colored. A nonlinear filter can only outperform the Kalman filter when the random processes w and v are non-Gaussian [11].

3.1. Extended Kalman filter

For nonlinear systems the Kalman filter cannot be used, which led to the development of the extended Kalman filter (EKF) [9]. The main idea behind the EKF is that at each iteration the nonlinear system is linearized around the current state estimate.

The nonlinear system is used to compute the *a priori* estimates, and the linearized system is used to compute the Kalman gain. The complete algorithm is shown in Algorithm 3.1.

Algorithm 3.1 Complete algorithm of the extended Kalman filter for nonlinear differentiable systems.

- 1: Initialize \hat{x} and P with \hat{x}_0 and $P_0 \neq 0$.
- 2: Assign the expected noise covariances to Q and R .
- 3: **for** $k = 1, 2, 3, \dots$ **do**
- 4: Linearize the system dynamics (Eq. (6))

$$F_k = \frac{\partial f(\hat{x}_{k-1}, u_{k-1}, 0)}{\partial x_{k-1}}$$

$$H_k = \frac{\partial h(\hat{x}_{k-1}, 0)}{\partial x_{k-1}}$$

- 5: Predict the state and the estimate error covariance

$$\hat{x}_k^- = f(\hat{x}_{k-1}, u_k, 0)$$

$$P_k^- = F_k P_{k-1} F_k^\top + Q$$

- 6: Compute the Kalman gain

$$K_k = P_k^- H_k^\top (H_k P_k^- H_k^\top + R)^{-1}$$

- 7: Update the state and covariance estimate using the measurement

$$\hat{x}_k = \hat{x}_k^- + K_k (y_k - h(\hat{x}_k^-, 0))$$

$$P_k = (I - K_k H_k) P_k^-$$

- 8: **end for**
-

For nonlinear systems, convergence and non-biasedness of the EKF are not guaranteed. So, even if the filter converges, it might converge to incorrect state estimates. The extended Kalman filter is nonetheless often used because of the simplicity of the algorithm and the fact that the EKF often works well in practise.

3.2. Unscented Kalman filter

The linearization step in the EKF algorithm introduces errors that can deteriorate the performance and stability of the filter. A more accurate nonlinear filter is the unscented Kalman filter of [10], which evaluates the nonlinear function at a certain set of $2N_x + 1$ points (the *sigma points*)

around the previous state estimate \hat{x}_{k-1} , where N_x is the length of the state vector x . All these points are propagated through the nonlinear system, and based on the weighted mean and covariance of these transformed sigma points, the *a priori* state estimate \hat{x}_k^- and the Kalman gain K_k are found without any linearization.

The complete algorithm of the unscented Kalman filter is shown in Algorithm 3.2, where the following $2N_x + 1$ sigma points are used:

$$x^{(i)} = \begin{cases} \hat{x}_k, & i = 0 \\ \hat{x}_k + \left(\sqrt{\frac{N_x}{1-W_m^{(0)}} P_k} \right)_i, & 0 < i \leq N_x \\ \hat{x}_k - \left(\sqrt{\frac{N_x}{1-W_m^{(0)}} P_k} \right)_{i-N_x}, & N_x < i \leq 2N_x \end{cases} \quad (13)$$

Using certain weights, the weighted mean and covariance of the sigma points is equal to the true mean and covariance of \hat{x}_k . For the calculation of the mean, these weights are

$$W_m^{(i)} = \begin{cases} \frac{\alpha^2(N_x + \kappa) - N_x}{\alpha^2(N_x + \kappa)}, & i = 0 \\ \frac{1 - W_m^{(0)}}{2N_x}, & 0 < i \leq 2N_x \end{cases}, \quad (14)$$

and for the calculation of the covariance they are

$$W_c^{(i)} = \begin{cases} W_m^{(0)} + (1 + \beta - \alpha^2), & i = 0 \\ W_m^{(i)}, & 0 < i \leq 2N_x \end{cases} \quad (15)$$

Here $10^{-4} < \alpha < 1$ is a parameter that determines the spread of the sigma points, $i \in \{0, \dots, 2N_x\}$, and κ and β are secondary parameters used to incorporate prior knowledge of the probability distribution of the state. For Gaussian distributions $\kappa = 3 - N_x$ and $\beta = 2$ [10, 12].

For the tracking problem considered in this article, the exact value of α appeared to be of little influence, so the arbitrary value $\alpha = 10^{-3}$ is used.

The square root in (13) should be computed using a numerically stable and efficient method. Here it is calculated using the Cholesky decomposition [10, 13], which requires that P_k is positive definite. However, P_k is a covariance matrix, so it might be only positive semidefinite. A solution is to calculate the Cholesky decomposition of $P_k + \epsilon I$, where $\epsilon > 0$ is a small parameter.

For nonlinear systems the UKF is more accurate than the EKF, the tradeoff is a higher computational cost [10].

For both filters it can be derived that when $R \gg Q$, the filter will rely more on the model than

Algorithm 3.2 Complete algorithm of the unscented Kalman filter.

- 1: Initialize \hat{x} and P .
- 2: Assign α , β and κ .
- 3: Assign the expected noise covariances to Q and R .
- 4: **for** $k = 1, 2, 3, \dots$ **do**
- 5: Calculate sigma points $x_{k-1}^{(i)}$ using Eq. (13)
- 6: Propagate sigma points through the system dynamics (Eq. (6))

$$\begin{aligned} x_k^{(i)} &= f(x_{k-1}^{(i)}) \\ y_k^{(i)} &= h(x_k^{(i)}) \end{aligned}$$

- 7: Calculate the *a priori* state and output estimate using $W_m^{(i)}$ as in Eq. (14)

$$\hat{x}_k^- = \sum_{i=0}^{2N_x} W_m^{(i)} x_k^{(i)} \quad (16)$$

$$\hat{y}_k^- = \sum_{i=0}^{2N_x} W_m^{(i)} y_k^{(i)} \quad (17)$$

- 8: Calculate the *a priori* covariance estimates using $W_c^{(i)}$ as in Eq. (15)

$$\begin{aligned} P_k^- &= Q + \sum_{i=0}^{2N_x} W_c^{(i)} (x_k^{(i)} - \hat{x}_k^-)(x_k^{(i)} - \hat{x}_k^-)^\top, \\ P_{xy,k}^- &= \sum_{i=0}^{2N_x} W_c^{(i)} (x_k^{(i)} - \hat{x}_k^-)(y_k^{(i)} - \hat{y}_k^-)^\top, \\ P_{yy,k}^- &= R + \sum_{i=0}^{2N_x} W_c^{(i)} (y_k^{(i)} - \hat{y}_k^-)(y_k^{(i)} - \hat{y}_k^-)^\top \end{aligned} \quad (18)$$

- 9: Calculate the Kalman gain

$$K_k = P_{xy,k}^- \left(P_{yy,k}^- \right)^{-1}$$

- 10: Update the estimated state and its covariance

$$\hat{x}_k = \hat{x}_k^- + K_k (y_k - \hat{y}_k^-) \quad (19)$$

$$P_k = P_k^- - K_k P_{yy,k}^- K_k \quad (20)$$

- 11: **end for**
-

the filter will be low. When $Q \gg R$, the filter will rely more on the measurements than on the model, which leads to fast convergence, but also noisy estimates. The stability of the filters depends on the accuracy of the initial conditions and covariance matrices R and Q , but also on how well the model fits the measured data. For instance, tracking a sawtooth using the sinusoidal signal model presented here turned out to be impossible, while tracking a pure sinusoid succeeded almost without fail.

Although quantitative results on the convergence speed and stability of the filters would be useful, these are hard to find because of the nonlinearity of the problem, and the dependence on the signal that is to be tracked.

3.3. Covariance matrices R and Q

The Kalman gain depends heavily on the process noise covariance matrix Q , and the measurement noise covariance matrix R , therefore accurate values for Q and R are essential for good performance and stability of the filter. In many cases, R and Q are constant matrices and are tuned using a trial and error procedure.

Here, the process noise w_k includes the changes in the frequency, phase and amplitude of the dominant harmonic in the measured periodic signal. As the level of fluctuation of these variables is unknown and varies over time, an algorithm that provides an adaptive estimate \hat{Q}_k of the time-varying matrix Q_k is desired.

All unmodelled dynamics that are visible in the measurement (such as upper harmonics) are included in the measurement noise v_k . Assuming that these unmodelled dynamics are roughly the same in each tokamak shot, a reasonable value for R can be found by tuning.

A number of approaches to the adaptive estimation of Q can be taken, including Bayesian estimation, maximum likelihood estimation, correlation methods and covariance-matching techniques [14]. Bayesian estimation and maximum likelihood estimation methods yield reliable estimates, but are computationally difficult. Correlation methods are mainly applicable for systems where Q is time-invariant, and the computations become very difficult when the system is nonlinear. Here we use a covariance-matching technique, which is easy to implement and works reasonably in practise, but the convergence cannot be guaranteed. Other esti-

on the measurements, so the convergence speed of

mation techniques will possibly yield better results and are an interesting subject of future research.

Here a covariance-matching technique is presented for the estimation of Q , which is based on the one presented in [12]. Applying this adaptive estimation of Q in the EKF or UKF yields the adaptive extended Kalman filter (AEKF) or adaptive unscented Kalman filter (AUKF).

If it is assumed that a good estimate \hat{x}_{k-1} of x_{k-1} is available, then

$$f(x_{k-1}) - f(\hat{x}_{k-1}) \approx 0, \quad (21)$$

and

$$x_k - \hat{x}_k^- = f(x_{k-1}) - f(\hat{x}_{k-1}) + w_k \approx w_k. \quad (22)$$

Assuming that the correction step provides a good estimate \hat{x}_k of x_k , then $x_k - \hat{x}_k \approx 0$. Given that the correction term $z_k = \hat{x}_k - \hat{x}_k^- = K_k(y_k - h(\hat{x}_k^-))$, it follows that $z_k \approx w_k$, so z_k can be used to estimate Q_k .

The covariance of the correction term z_k equals

$$\begin{aligned} \mathbb{E}\{z_k z_k^\top\} &= \mathbb{E}\{K_k(h(x_k) - h(\hat{x}_k^-) + v_k) \\ &\quad (h(x_k) - h(\hat{x}_k^-) + v)^\top K_k^\top\} \\ &= \mathbb{E}\{K_k(h(x_k) - h(\hat{x}_k^-))(h(x_k) \\ &\quad - h(\hat{x}_k^-))^\top K_k^\top + K_k R K_k^\top\}. \end{aligned} \quad (23)$$

Since the measurement noise should not influence the influence Q , we remove the term $K_k R K_k^\top$ to find our estimate \hat{Q} :

$$\hat{Q}_k = \mathbb{E}\{z_k z_k^\top - K_k R K_k^\top\}. \quad (25)$$

Because Q can be time-varying, the estimate \hat{Q}_k at time t_k should rely on recent data. Here a learning algorithm is used, in which old data is ‘forgotten’ at rate γ :

$$\hat{Q}_k \approx \hat{Q}_{k-1} + \gamma \left((z_k z_k^\top - K_k R K_k^\top) - \hat{Q}_{k-1} \right). \quad (26)$$

The learning rate γ should be small enough to smooth out the effect of the term $z_k z_k^\top - K_k R K_k^\top$, which can vary dramatically from one time step to the next, but large enough to provide good tracking behaviour. Here, $\gamma = t_s$ is used, so that \hat{Q}_k depends mostly on data from the past time unit.

If the filter provides good estimates, then $\mathbb{E}\{e_y e_y^\top\} \approx \mathbb{E}\{v v^\top\} = R$, where $e_{y,k} = y_k - \hat{y}_k$.

If in some time interval, $\mathbb{E}\{e_y e_y^\top\} \ll R$, then in that interval the filter is tracking the noise, implying that \hat{Q} is too large. If in some time interval, $\mathbb{E}\{e_y e_y^\top\} \gg R$, then the filter fails to track the output (and thus the states), implying that \hat{Q} is too small. Thus, any divergence of the filter can be detected by periodically checking whether

$$a_1 R \leq \frac{1}{N} \sum_{i=k-N}^k e_{y,i} e_{y,i}^\top \leq a_2 R, \quad (27)$$

where N , a_1 and a_2 are tuning parameters. If either inequality is not satisfied, \hat{Q} should be adapted. It should be noted that if both inequalities are satisfied, this does not guarantee that the filter is tracking correctly. For the tests performed in this article, we used $N = 2000$, $a_1 = 1/3$ and $a_2 = 3$, and \hat{Q}_k is reset to

$$\hat{Q}_k = \hat{Q}_0 \frac{\|\hat{Q}_k\|}{\|\hat{Q}_0\|} \cdot 20 \quad (28)$$

when \hat{Q}_k is too small, or

$$\hat{Q}_k = \hat{Q}_0 \frac{\|\hat{Q}_k\|}{\|\hat{Q}_0\|} \cdot \frac{1}{20} \quad (29)$$

when \hat{Q}_k is too large. Here $\|\cdot\|$ denotes the Euclidean matrix norm.

This method is based on the strong assumption that the filter is tracking accurately, so the formal validity of this method can be questioned. The estimate can be improved by using prior knowledge of Q . Since $\mathbb{E}\{x_1 x_1\} = \mathbb{E}\{x_2 x_2\}$ and $\mathbb{E}\{x_1 x_2\} = \mathbb{E}\{x_2 x_1\} = 0$, we assume that $\mathbb{E}\{w_1 w_1\} = \mathbb{E}\{w_2 w_2\}$ and $\mathbb{E}\{w_1 w_2\} = \mathbb{E}\{w_2 w_1\} = 0$. By also assuming that w_3 is independent of w_1 and w_2 we find that

$$Q_k = \mathbb{E}\{w_k w_k^\top\} = \begin{bmatrix} q_{1,k} & 0 & 0 \\ 0 & q_{1,k} & 0 \\ 0 & 0 & q_{3,k} \end{bmatrix}, \quad (30)$$

where $q_{1,k} = \mathbb{E}\{w_{1,k} w_{1,k}^\top\}$ and $q_{3,k} = \mathbb{E}\{w_{3,k} w_{3,k}^\top\}$.

4. Tracking of synthetic data

In this section the performance of the AEKF and AUKF is evaluated by means of simulation. Synthetic data created with Matlab will be used, so that all errors are exactly known and thus also the exact performance of the filters.

Tests are performed using signals containing

- a step in the frequency to simulate very fast changes in the frequency;
- a monotonously increasing frequency (chirp signal) to test tracking behaviour;
- a frequency varying in a large range to test the operating range.

All signals are contaminated with white Gaussian noise to simulate measurement noise. The amount of measurement noise is defined by the signal-to-noise ratio (SNR), where the SNR is the ratio of the root-mean-square (RMS) of the true signal and the RMS of the measurement noise. For an interval $[0, T]$, the SNR is

$$\text{SNR} = \frac{\sqrt{\frac{1}{T} \int_{t=0}^T y(t)^2 dt}}{\sqrt{\frac{1}{T} \int_{t=0}^T v(t)^2 dt}}, \quad (31)$$

usually expressed in decibel (dB), using $20 \log_{10}(\text{SNR})$.

To judge the quality of the frequency estimate, we define the relative frequency error

$$e_f = \frac{|f - \hat{f}|}{f}. \quad (32)$$

If the frequency estimate is accurate, then $e_f \ll 1$. To judge the quality of the phase estimate the phase error we look at $\hat{\theta} - \theta$, and to show the evolution of \hat{Q}_k , its Euclidean matrix norm $\|\hat{Q}_k\|_2$ is used. As for the model used here, \hat{Q}_k is a diagonal matrix for all time steps k , the Euclidean norm is equal to the largest element of \hat{Q}_k .

4.1. Signal with step in frequency

The AEKF and AUKF filters are challenged with the tracking of the signal

$$y(t) = \cos(\theta) + v(t), \quad (33)$$

where

$$\frac{1}{2\pi} \frac{d}{dt} \theta = f(t) = \begin{cases} 40 \text{ Hz} & \text{for } 0 \leq t < 2 \\ 60 \text{ Hz} & \text{for } 2 \leq t < 4 \end{cases}, \quad (34)$$

and $v(t)$ is white Gaussian noise, such that the signal-to-noise ratio is 30 dB. The frequency of the signal is shown in Fig. 4.1.

We use a sample time of $t_s = 10^{-3}$ s and initial conditions $f_0 = 44$ Hz, $\theta_0 = 0$, $A_0 = 1$, $P_0 = Q_0 = t_s \cdot \text{diag}([1 \ 1 \ 0.01])$. It has been found by

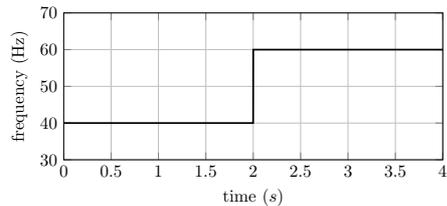


Figure 1: The frequency of reference signal (33).

trial and error that this value for P_0 and Q_0 works well for most signals, so this is used throughout this article,

The relative frequency error and the phase error of both filters are shown in Fig. 2.

The signal is tracked accurately, without any bias. When the step in the frequency occurs, the AEKF locks onto the new frequency within 50 samples (0.05 s) and the AUKF within 100 samples (0.1 s). In steady-state, the phase error is no more than four degrees.

Fig. 3 shows $\|\hat{Q}_k\|_2$ as a function of time. In steady state the norm converges, which means that also \hat{Q}_k converges entrywise. When at time $t = 2$ s the step in the frequency occurs, the estimate \hat{Q}_k is increased immediately according to 28, showing that the estimate adapts to the current situation. At $t = 3$ s, \hat{Q}_k is considered too large and decreased using 29, which explains why e_f also decreases at $t = 3$ s.

At a signal-to-noise ratio of 30 dB the filters succeeded in accurately tracking the phase and frequency, using almost any initial state estimate \hat{x}_0 , even insensible ones. When the signal-to-noise ratio is decreased, the range of initial conditions that leads to a tracking filter becomes smaller. It was observed that the unscented Kalman filter achieves tracking for a larger range of initial conditions than the extended Kalman filter, especially for low SNR.

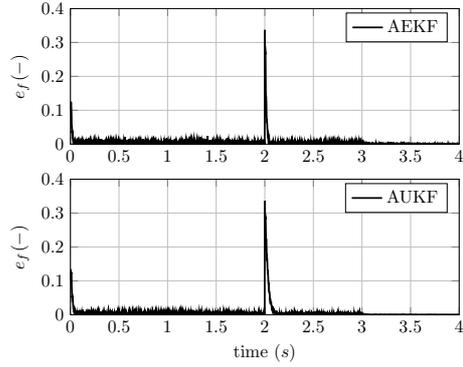
4.2. Chirp signal

To test the tracking behaviour in transient situations a test signal is used with monotonically increasing frequency:

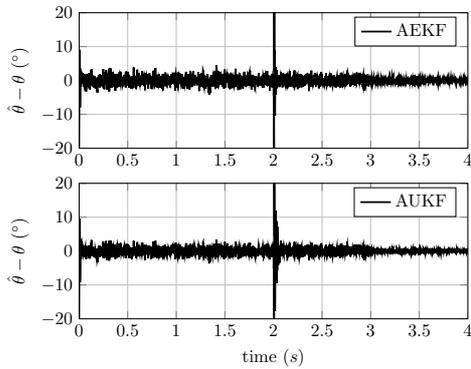
$$y(t) = \cos(\theta) + v(t), \quad (35)$$

where

$$\frac{1}{2\pi} \frac{d}{dt} \theta = f(t) = \begin{cases} 40 \text{ Hz} & \text{for } 0 \leq t < 1 \\ 40t \text{ Hz} & \text{for } 1 \leq t < 3 \\ 120 \text{ Hz} & \text{for } 3 \leq t \end{cases}, \quad (36)$$



(a) relative frequency error



(b) phase error

Figure 2: Tracking signal (33), showing the relative frequency error (a) and the estimation error of the phase (b).

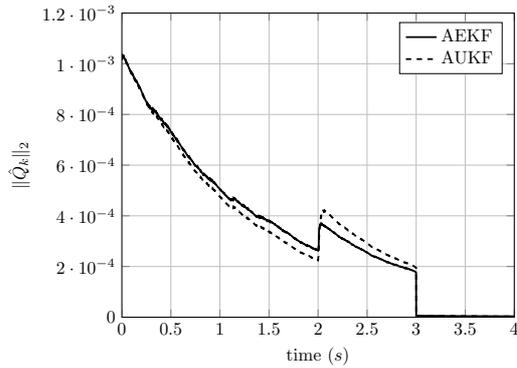


Figure 3: The Euclidean norm of \hat{Q}_k as a function of time, while tracking signal (33). The converging norm implies that the elements of \hat{Q}_k also converge.

and $v(t)$ is again white Gaussian noise, such that the SNR is 30 dB.

The frequency of the signal is shown in Fig. 4.

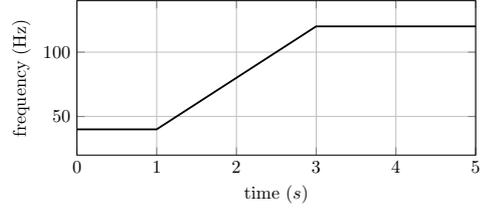
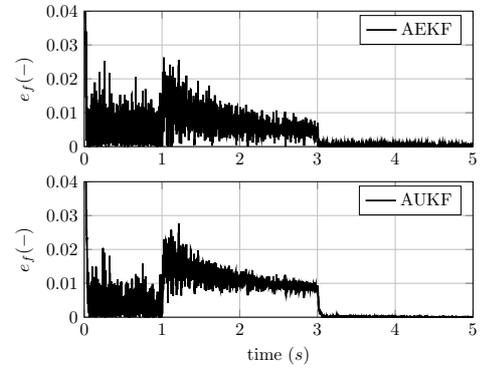
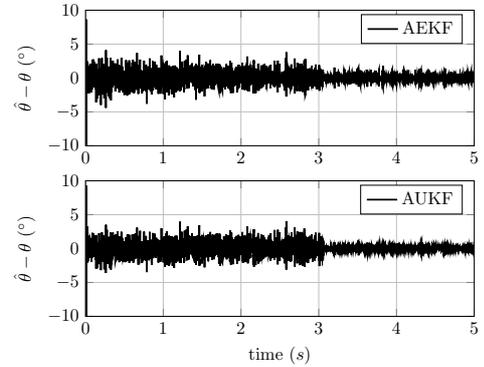


Figure 4: The frequency of reference signal (35).

The used sample time and initial conditions are the same as in the previous test, the results are shown in Fig. 5.



(a) relative frequency error



(b) phase error

Figure 5: Tracking signal (35), with signal-to-noise ratio of 30 dB, showing the relative frequency error (a) and the estimation error of the phase (b).

In steady state the estimates are nonbiased, but when the frequency is increasing, the frequency estimate is both biased and lagging behind. The lag is caused by the fact that the filter can only react to errors and not anticipate them. The bias is in agree-

ment with the claim in [15] that the estimates will be biased if the assumed noise characteristics are incorrect. The model assumes zero-mean noise, but for $1 \leq t < 3$ the linearly increasing frequency leads to $\mathbb{E}\{w_3\} = \mathbb{E}\{\Delta_f\} = 40$ Hz/s, so w_3 is nonzero mean noise. Note that the phase estimate error, see Fig. 5(b) is again within four degrees, and is nonbiased.

The frequency bias in the transient case can be removed by including the extra state $x_4 = 2\pi\Delta_f t_s$ into the model. However, one should only add states when absolutely necessary, because adding states increases the complexity of the tracking/estimation problem and can make the filter less stable.

In Fig. 5(a) it can be seen that the effect of the noise on the estimated frequency is decreasing over time. This is caused by the converging of \hat{Q}_k , which can be seen in Fig. 6. Again, at $t = 3$ s, \hat{Q}_k is considered too large and decreased using 29.

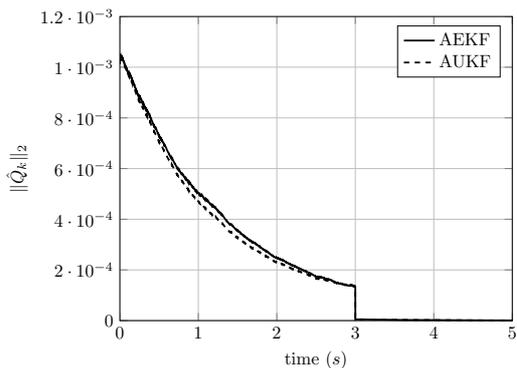


Figure 6: The Euclidean norm of \hat{Q}_k as a function of time, while tracking signal (35). The converging norm implies that the elements of \hat{Q}_k also converge.

4.3. Frequency range test

For ASDEX Upgrade, a filter is requested that is able to track harmonic signals in the range 5 Hz–20 kHz. It is assumed that on ASDEX Upgrade a sample frequency of 100 kHz will be used. Because the system is observable for $0 < f < f_s/2$, $A > 0$, tracking should be possible in the range 0 Hz–50kHz. However, at frequencies near 0 or $f_s/2$, the system is almost unobservable, which makes accurate tracking difficult. The observability at low frequencies can be improved by decreasing the sample frequency, but this also decreases the maximum detectable frequency.

Here we show that the AUKF can track fast changing frequencies from 0 Hz to 50kHz by applying the AUKF to the following signal:

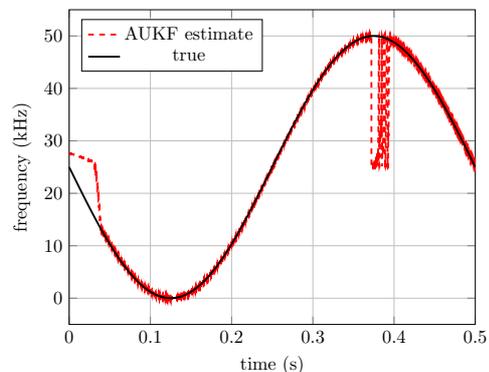
$$y(t) = \cos(\theta) + v(t), \quad (37)$$

where

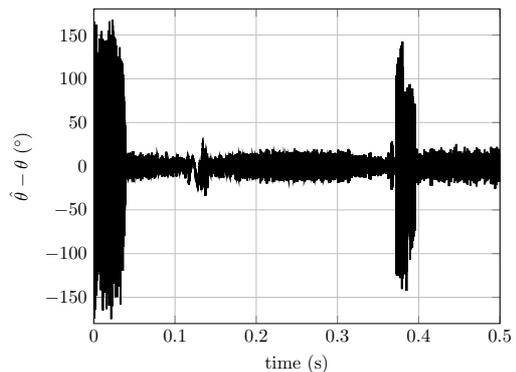
$$\frac{1}{2\pi} \frac{d}{dt} \theta = f(t) = 25000 + 25000 \sin(4\pi t) \text{ Hz}, \quad (38)$$

and $v(t)$ is white Gaussian noise, such that the SNR is 10 dB. The frequency is shown in Fig. 7(a).

The true frequency and the AUKF estimate are shown in Fig. 7(a). In this figure we can see that near the Nyquist frequency, the AUKF loses its lock onto the signal, but quickly regains it when the frequency decreases.



(a) signal frequency



(b) phase error

Figure 7: The results of the frequency range test using the AUKF and a signal-to-noise ratio of 10 dB, showing the frequency of the signal to be tracked and the AUKF estimate (a) and the phase error $\hat{\theta} - \theta$ (b).

The phase error of the AUKF estimate is shown in Fig. 7(b). Once the AUKF has locked onto the

signal, the phase error is typically less than 20 degrees. Only near 0 Hz and 50 kHz does the phase error exceed 20 degrees, which is a result of the observability issues. It can also be seen that faster frequency changes lead to larger phase errors.

It can be concluded that a Kalman filter can track a sinusoid in the range 0 Hz–50 kHz when the sample frequency is 100 kHz. However, near 0 Hz and 50 kHz tracking becomes very difficult and large errors will occur.

5. Tracking of NTMs in experiments

In this section the Kalman filters are tested on real data from the Tokamak à Configuration Variable (TCV), at the EPFL, Switzerland, and from the Tokamak Experiment for Technology Oriented Research [16] (TEXTOR), located in Jülich, Germany. The TCV data are Mirnov coil measurements, whereas the TEXTOR data are line-of-sight ECE measurements [17].

5.1. TCV Mirnov coil data

The filters are tested on the Mirnov coil data of TCV shot #40605. In Fig. 8 a section of the data is shown, together with the AUKF estimate. In Fig. 9 the spectrogram of the data is shown, together with the frequency estimates of a discrete-time phase-locked loop (DPLL), used by the TCV team, and the AUKF.

The AEKF only succeeds in tracking when a model is used that includes the second harmonic. The AUKF is able to accurately track the signal using a model with only three states.

The data is measured at the sample frequency $f_s = 250$ kHz, the AUKF uses $R = 80$, and initial conditions $A_0 = 30$, $f_0 = 4$ kHz, $\theta_0 = 0$ and $P_0 = \hat{Q}_0 = t_s \cdot \text{diag}([1 \ 1 \ 0.01])$.

From Figs. 8 and 9 it can be seen that the signal is contaminated mostly with white noise and a second harmonic with low energy. When regarding the second harmonic also as noise, the SNR is about 7dB in the interval [0.4s, 2.0s]. The AUKF performs comparable to the DPLL, but the DPLL does lock onto the signal slightly earlier than the AUKF, while using the same initial frequency. However, it should be noted that the DPLL contains a band-pass filter that removes all components of the measured signal outside the frequency band in which the magnetic islands are expected. The AUKF does not contain such a band-pass filter, so the input signal of the

AUKF contains much more noise and other disturbances than the band-pass filtered signal, and is therefore harder to track. However, because the AUKF does not contain such a band-pass filter it is able to track signals in a much broader frequency range than the PLL.

5.2. TEXTOR ECE data

Next, the filters are tested using channel 4 of the ECE data of TEXTOR shots #107904 and #108029. The data is filtered with a high-pass filter with a cutoff frequency of 8Hz, to remove the DC offset from the signals. For both shots, the AUKF succeeds in tracking the ECE signal, but the AEKF does not.

5.2.1. Shot #107904

In Figs. 10(a) and 10(b) two sections of the channel 4, shot #107904 are shown, together with the output estimate using the AUKF. The spectrogram of the complete shot, together with the estimated frequency of the AUKF is shown in Fig. 10(c).

The data is measured at the sample frequency $f_s = 100$ kHz, the AUKF uses $R = 0.03$, and initial conditions $A_0 = 0.1$, $f_0 = 2$ kHz, $\theta_0 = 0$ and $P_0 = \hat{Q}_0 = t_s \cdot \text{diag}([1 \ 1 \ 0.01])$.

From $t = 2$ s to $t = 2.7$ s, there is a large disturbance on the signal, caused by ECRH scattering [18] that is only partially filtered out. This disturbance is not white, and in phase with the rotation of the islands, so the measured signal can no longer be approximated by a sinusoid. This can be solved by changing the signal model, but here it is solved by interpreting the disturbance as measurement noise, which means that the measurement noise covariance is increased during this period to $R = 0.20$. As seen in Fig. 10(a), during this period the filter cannot produce a nice fit on the ECE signal, but it does produce a nice sinusoid that is as close as possible to the measured signal. If the measurement noise covariance is not increased when the measurement is disturbed by ECRH scattering, the AUKF produces a better fit on the ECE signal by oscillating the estimated frequency. Then the output estimate is more accurate, but the frequency and phase estimates are less accurate.

From $t = 1$ s to $t = 5$ s the filter can accurately track the signal, even when the frequency suddenly drops at around $t = 3.2$ s. At the beginning and the end of the shot, the signal power is too low for the filter to achieve tracking.

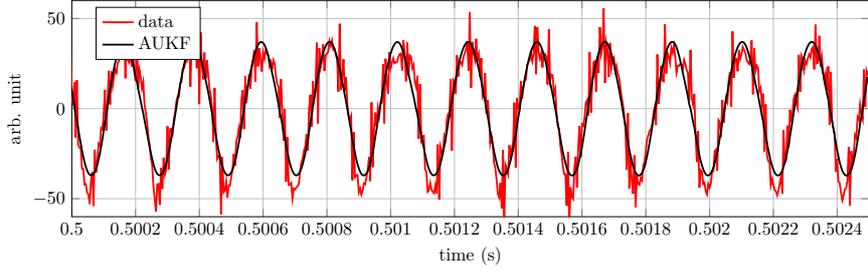


Figure 8: A segment of the Mirnov coil measurement from the TCV plant and the corresponding output estimate of the AUKF. Data is courtesy of the TCV team.

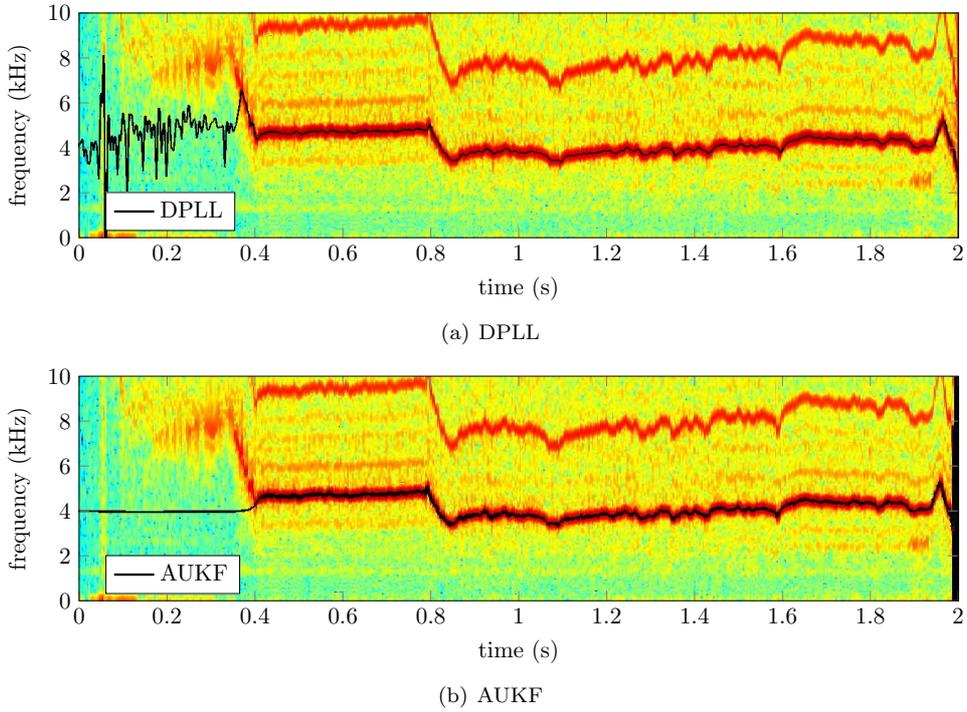


Figure 9: The spectrogram of the TCV Mirnov coil data, together with (a) the frequency estimate using a discrete phase-locked loop, and (b) the AUKF frequency estimate. Data is courtesy of the TCV team.

5.2.2. Shot #108029

The spectrogram of shot #108029 is shown in Fig. 11, together with the estimated frequency of the AUKF. The data is measured at the sample frequency $f_s = 100$ kHz, the AUKF uses $R = 0.02$, and initial conditions $A_0 = 0.2$, $f_0 = 1$ kHz, $\theta_0 = 0$ and $P_0 = \hat{Q}_0 = t_s \cdot \text{diag}([1 \ 1 \ 0.01])$. Again the AUKF achieves good tracking, apart from a few occasions where the AUKF momentarily loses the lock on the signal because of disturbances, but in each case the lock is quickly regained. Just like in

shot #107904, when the frequency suddenly drops almost like a step function, the AUKF estimate is lagging behind, but only for a duration of less than 0.05 s (less than 50 periods). Note that the very fast frequency variations between $t = 5$ s and $t = 6$ s are being tracked accurately.

Before $t = 2$ s the signal consists of sawtooth measurement, which cannot be tracked with a sinusoidal signal model. After $t = 6$ s the signal power is too low for the filter to achieve tracking.

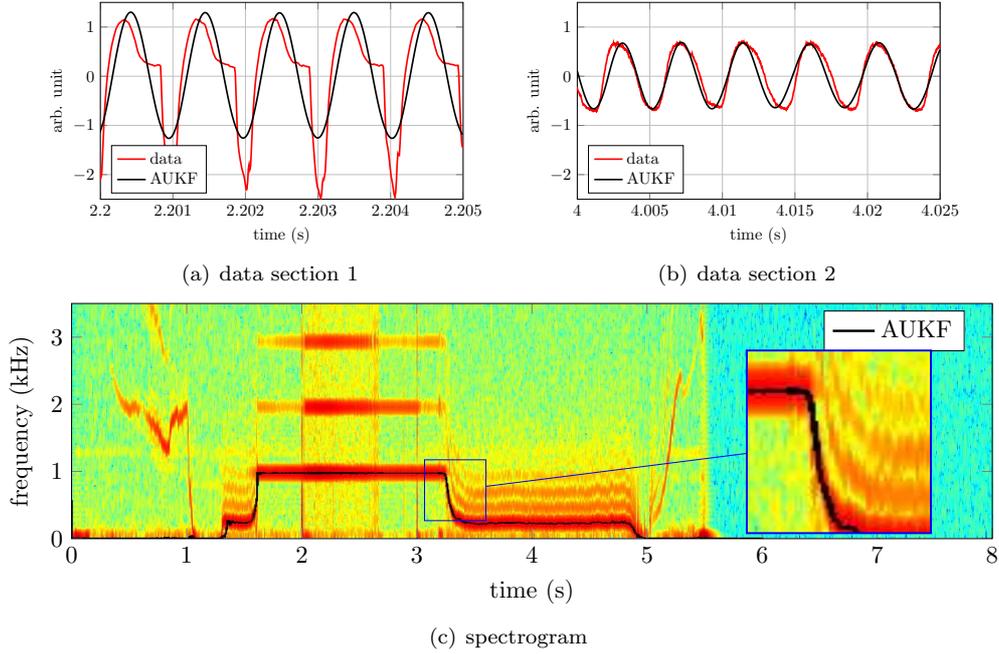


Figure 10: Two sections of TEXTOR shot #107904, channel 4 (a,b), together with the AUKF estimate, and the spectrogram of the complete shot (c), together with the AUKF frequency estimate. Data is courtesy of the TEXTOR team.

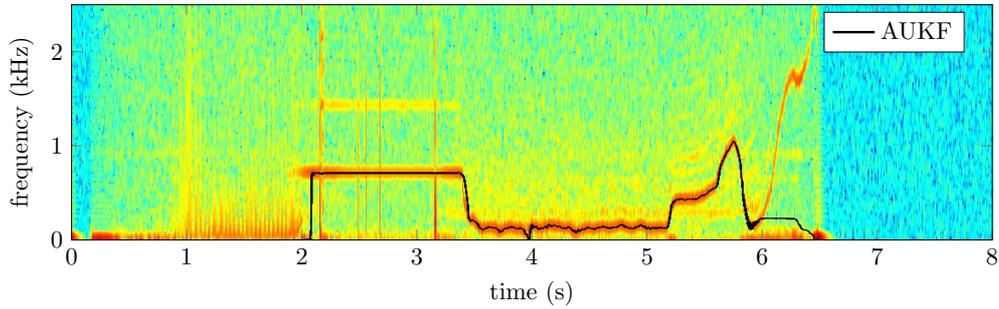


Figure 11: Spectrogram of channel 4 of the ECE data of TEXTOR shot #108029, together with the frequency estimate of the AUKF. Data is courtesy of the TEXTOR team.

6. Conclusion

An Unscented and an Extended Kalman Filter (UKF and EKF) have been defined for real-time frequency, amplitude and phase estimation using noisy measurements. These filters have been applied to synthetic and experimental rotating tearing mode data from TEXTOR and TCV.

The analysis shows that the UKF shows good tracking over a wide range of mode rotational frequencies, and is robust to disturbances (like ECRH scattering) and fast frequency changes. Subject to

the applied noise, the UKF can track modes from 0 Hz to 50 kHz, a much larger range than the frequency band of a PLL. The quality of the estimates of course improves when the signal-to-noise ratio increases.

Given the wide frequency band of this method, it could prove relevant for real-time mode analysis and NTM control on ASDEX Upgrade. The implementations of the EKF and UKF as discussed here are not optimised for real-time feedback control applications yet, as the covariance matrices Q and R have to be optimized for every machine and sce-

nario. We also note that the methodology works for general periodic signals, so it can potentially also be used for other periodic signals in the plasma, like Alfvén waves and perhaps sawtooth signals.

7. Acknowledgements

The work in this paper has been performed within the framework of the NWO-RFBR Centre of Excellence (grant 047.018.002) on Fusion Physics and Technology. This work was supported by NWO, ITER-NL, the European Communities under the contract of the association Euratom/FOM was carried out within the framework of the European Fusion Programme. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

References

- [1] R. J. L. Haye, Neoclassical tearing modes and their control, *Physics of Plasmas* 13 (2006) 055501.
- [2] J. Crawford, *Advanced phase-lock techniques*, Artech House, Boston, 2008.
- [3] B. A. Hennen, E. Westerhof, P. W. J. M. Nuij, J. W. Oosterbeek, M. R. de Baar, W. A. Bongers, A. Bürger, D. J. Thoen, M. Steinbuch, the TEXTOR Team, Real-time control of tearing modes using a line-of-sight electron cyclotron emission diagnostic, *Plasma Physics and Controlled Fusion* 52 (2010) 104006.
- [4] M. Bora, D. Sarmah, Sawtooth disruptions and limit cycle oscillations, *Communications in Nonlinear Science and Numerical Simulation* 13 (2006) 26.
- [5] D. Constantinescu, O. Dumbrajs, V. Igochine, K. Lackner, R. Meyer-Spasche, H. A. U. T. Zohm, A low-dimensional model system for quasi-periodic plasma perturbations, *Physics of Plasmas* 18 (2011) 062307.
- [6] H. Nijmeijer, Observability of autonomous discrete time non-linear systems: a geometric approach, *International Journal of Control* 36 (1982) 867–874.
- [7] D. Alves, R. Coelho, An adaptive real-time multi-tone estimator and frequency tracker for non-stationary signals, *Real Time Conference (RT), 2010 17th IEEE-NPSS* (2010) 1–7.
- [8] H. Hajimolahoseini, M. Taban, H. Soltanian-Zadeh, Extended Kalman filter frequency tracker for nonstationary harmonic signals, *Measurement* 45 (2012) 126–132.
- [9] G. Welch, G. Bishop, *An Introduction to the Kalman Filter*, ACM Press, Addison-Wesley, 2001.
- [10] S. Julier, J. Uhlmann, Unscented filtering and nonlinear estimation, *Proceedings of the IEEE* 92 (2004) 401–422.
- [11] R. Kalman, A new approach to linear filtering and prediction problems, *Transactions of the ASME—Journal of Basic Engineering* 82 (1960) 35–45.
- [12] P. Dash, S. Hasan, B. Panigrahi, Adaptive complex unscented Kalman filter for frequency estimation of time-varying signals, *Science, Measurement Technology, IET* 4 (2010) 93–103.
- [13] M. Heath, *Scientific computing: An introductory survey*, McGraw-Hill Higher Education, 2nd edition, 1996.
- [14] R. Mehra, Approaches to adaptive filtering, *Automatic Control, IEEE Transactions on* 17 (1972) 693–698.
- [15] L. Ljung, Asymptotic behavior of the extended Kalman filter as a parameter estimator for linear systems, *Automatic Control, IEEE Transactions on* 24 (1979) 36–50.
- [16] K. H. Finken, et al., The dynamic ergodic divertor in the textor tokamak: plasma response to dynamic helical magnetic field perturbations, *Plasma Physics and Controlled Fusion* 46 (2004) B143.
- [17] J. Oosterbeek, A. Bürger, E. Westerhof, M. De Baar, M. Van Den Berg, W. Bongers, M. Graswinckel, B. Hennen, O. Kruijt, J. Thoen, R. Heidinger, S. Korsholm, F. Leipold, S. Nielsen, A line-of-sight electron cyclotron emission receiver for electron cyclotron resonance heating feedback control of tearing modes, *Review of Scientific Instruments* 79 (2008). Cited By (since 1996) 19.
- [18] E. Westerhof, S. K. Nielsen, J. W. Oosterbeek, M. Salewski, M. R. De Baar, W. A. Bongers, A. Bürger, B. A. Hennen, S. B. Korsholm, F. Leipold, D. Moseev, M. Stejner, D. J. Thoen, Strong scattering of high power millimeter waves in tokamak plasmas with tearing modes, *Phys. Rev. Lett.* 103 (2009) 125001.