

A computationally efficient global indicator to detect spurious measurement drifts in Kalman filtering

C. Parellier, A. Barrau and S. Bonnabel

Abstract—We consider the problem of detecting additive structured correlated perturbations affecting the measurement outputs of a system whose state is estimated by a Kalman filter. We advocate the time series of the gradients of the log-likelihood with respect to the output measurements as an indicator, notably through its fast Fourier transform (FFT). This provides a novel unifying method to detect structured perturbations, namely small sinusoidal perturbations with unknown frequency, and slowly growing errors, such as a ramp, or more generally any known incipient profile with unknown starting time. The method allows for identification of their parameters too, i.e., frequency, and starting time of the ramp. Thanks to recent results on backpropagation in Kalman filters, and the use of the FFT, the method remains numerically tractable even for large datasets, as demonstrated by simulations.

I. INTRODUCTION

In this paper, we consider a (physical) linear system whose state is estimated at all times by a Kalman filter (KF) based on linear noisy measurements. In this context one may face the problem of a faulty sensor, which would deviate from the values it is supposed to deliver, due to a malfunction.

The field of fault detection and isolation provides too many tools to be listed exhaustively. In a nutshell, within the context of filtering, most approaches revolve around the residuals, also called innovation process, whose values and statistics are inherently computed by the KF. For linear systems with Gaussian noise, the normalized innovation squared (NIS) follows a chi-square distribution which allows for testing incoming measurements under the nominal model [1]. In this way, biases in the instrument may be identified when testing for the mean of the innovation. Hypothesis testing on the residuals has opened up for powerful techniques based on likelihood ratios [16], that provide efficient tools to detect abrupt changes [15], [3].

In the present paper, we propose a novel approach which is as follows. We consider the offline change detection setting, where given a data sample of N measurements, one tries to detect if an output fault has occurred, of which type it is, and the time instant it started to occur. In spite of the terminology, we aim at onboard implementations, where the data consist of the last N measurements. We assume the system is linear

This work is supported by CIFRE Grant 2019/1974 from French Agence Nationale de la Recherche et de la Technologie (ANRT).

C. Parellier and S. Bonnabel are with Mines Paris - PSL, PSL Research University, Centre for Robotics, 60 Bd Saint-Michel, 75006 Paris, France. `firstname.lastname@minesparis.psl.eu`

C. Parellier is with Safran Tech, Groupe Safran, Rue des Jeunes Bois-Chateaufort, 78772, Magny Les Hameaux, France. `firstname.lastname@safrangroup.com`

A. Barrau is with Offroad, 5 Rue Charles de Gaulle, 94140 Alfortville, France.

with Gaussian noises, and a Kalman filter estimates its state. We suppose faults come as additive perturbations of the output measurements of two different types:

- Oscillating fault: a sinusoidal signal of unknown moderate amplitude and unknown frequency corrupts the output measurements;
- Incipient faults: a function (typically a ramp with unknown slope) starts corrupting measurements at unknown time instant and slowly grows.

These types of faults are more difficult to detect than outliers because they do not correspond to an abrupt change in the signal. Additive ramps on the outputs or “slowly-growing errors” are a typical failure of the Global Positioning System (GPS), see [7], [8], [4], [5], with recent further developments in [18], [14], and arise also in process control [10].

Our idea is then to consider the likelihood over a time window of N measurements—which is readily computed from the Kalman filter’s estimates—and to evaluate how a given perturbation profile makes it vary. This constitutes a global indicator in that the log-likelihood is a sum over all measurements. Any profile which yields a large increase in likelihood when subtracted to the measurement signal is then likely to corrupt the measurements.

The problem with this idea is its computation complexity, which has made it prohibitively costly so far. Indeed, it may prove intractable to recompute the entire likelihood for each candidate profile (for instance each frequency of an injected sinusoid) and see how it varies, especially if we target onboard applications. We advocate herein this can be done much more efficiently, leveraging a number of tools. First, we consider the perturbation to be small, and compute a first-order expansion of the negative log-likelihood w.r.t. a candidate perturbation function $\delta y_1, \dots, \delta y_N$:

$$\mathcal{L}(y_1 - \delta y_1, \dots, y_N - \delta y_N) \approx \mathcal{L}(y) - \sum_{k=1}^N \frac{\partial \mathcal{L}}{\partial y_k} \delta y_k.$$

This implies computing the gradients w.r.t. all the measurements denoted y_1, y_2, \dots , that is, $(\frac{\partial \mathcal{L}}{\partial y_k})_{1 \leq k \leq N}$, which is generally very costly. Our very recent work on backpropagation in Kalman filters [12], though, has rendered it possible in very limited running time. Then, we use the fast Fourier transform (FFT) of the obtained signal (the gradients), and its properties to efficiently evaluate the variation in the likelihood corresponding to a given profile (that is, the righthmost term above), which yields a fast way of finding the parameters we seek, whether a frequency or a time when an incipient fault starts corrupting the signal.

The paper is organized as follows. In Section II we recall the Kalman filter equations and provide formulas for the sensitivity of the likelihood with respect to measured outputs. We then specify the types of faults our method is able to detect. In Section III we detail the method, and discuss its situation with respect to other—more classical—methods and its advantages in terms of numerical complexity. Section IV illustrates the interest through numerical simulations.

II. SENSITIVITY EQUATIONS AND CONSIDERED TYPES OF FAULTS

In this section we recall a few facts on the Kalman filter and the associated sensitivity equations, and then describe the types of faults we intend to address.

A. Kalman filter equations in discrete time

Consider a linear system defined by the state-transition matrix F_n and the measurement model matrix H_n ,

$$x_n = F_n x_{n-1} + w_n \quad (1)$$

$$y_n = H_n x_n + v_n, \quad (2)$$

corrupted by centered white Gaussian noises denoted by w_n and v_n , with respectively covariance matrices Q and R . We suppose the initial information on the state comes in the following form $x_0 \sim \mathcal{N}(\mu, P_0)$ with μ a known vector. At a given time n , the goal is to estimate with maximum accuracy the state x_n , given current and past observations y_n, y_{n-1}, \dots, y_1 . This can be achieved by computing the statistics $p(x_n | y_n, y_{n-1}, \dots, y_1)$ which enables computation of state estimates, and the uncertainty they convey. It turns out the latter is a Gaussian density, whose mean $\hat{x}_{n|n}$ and covariance $P_{n|n}$ can be obtained recursively through the KF's equations. They consist of a prediction step in which variables are evolved through the noise-free dynamical model

$$\hat{x}_{n|n-1} = F_n \hat{x}_{n-1|n-1} + B_n u_n, \quad P_{n|n-1} = F_n P_{n-1|n-1} F_n^T + Q, \quad (3)$$

where $B_n u_n$ is a known input. It is followed by an update step in which state is corrected in the light of the latest measurement

$$S_n = H_n P_{n|n-1} H_n^T + R, \quad K_n = P_{n|n-1} H_n^T S_n^{-1}, \quad (4)$$

$$P_{n|n} = (I - K_n H_n) P_{n|n-1}, \quad z_n = y_n - H_n \hat{x}_{n|n-1}, \quad (5)$$

$$\hat{x}_{n|n} = \hat{x}_{n|n-1} + K_n z_n, \quad (6)$$

where K_n is the (optimal) Kalman gain. The innovation $z_n = y_n - H_n \hat{x}_{n|n-1}$ is the prediction error between the measurement y_n and its predicted value $H_n \hat{x}_{n|n-1}$. It is a centered Gaussian random variable with covariance S_n .

B. Likelihood function

A well-known calculation [1] shows that up to an additive constant, the negative log-likelihood of the measurements over a time window, that is, $\mathcal{L} := -\log p(y_1, y_2, \dots, y_N)$, writes

$$\mathcal{L} = \sum_{n=1}^N \log \det S_n + z_n^T S_n^{-1} z_n. \quad (7)$$

This is also known as the energy function [13]. In the sequel we will refer to it as the ‘‘inconsistency’’, as we found it easier

to think about it as a measure of inconsistency in the present context. The weighted prediction error terms are easily interpreted: the commensurateness of a prediction error $y - Hx$ is to be evaluated against the associated covariance S_n (a large discrepancy is to be expected if the associated covariance is large). This results in a normalized innovation squared (NIS) $\frac{1}{N} \sum_{n=1}^N z_n^T S_n^{-1} z_n$ which follows a chi-square distribution with $N \times \dim z$ degrees of freedom under the nominal model [1].

C. Sensitivity with respect to measurements

It is customary to differentiate the log-likelihood \mathcal{L} with respect to the parameter matrices of the Kalman filter Q and R , see [9]. It is less frequent to study its sensibility with respect to measurements y_n 's. In this paper, we propose to use the log-likelihood \mathcal{L} for fault detection through its sensibility with respect to measurements. The corresponding equations are easily derived through the chain rule of calculus.

Proposition 1: For each $1 \leq k \leq N$ the gradient writes

$$\frac{\partial \mathcal{L}}{\partial y_k} = 2S_k^{-1} z_k - 2 \sum_{n=k+1}^N K_k^T (\Pi_{i=k+1}^{n-1} F_i^T (I - K_i H_i)^T) F_n^T H_n^T S_n^{-1} z_n, \quad (8)$$

where by convention $\sum_{N+1}^N = 0$ and $\Pi_{j+1}^j = I$ (identity matrix).

Proof: To use the chain rule of calculus it is easier to denote by $\frac{\partial a}{\partial b}$ the Jacobian matrix of vector a with respect to vector b , which makes gradients row vectors. In this formalism what we seek to prove is the transposed expression

$$\frac{\partial \mathcal{L}}{\partial y_k} = 2z_k^T S_k^{-1} - 2 \sum_{n=k+1}^N z_n^T S_n^{-1} H_n F_n (\Pi_{i=k+1}^{n-1} (I - K_i H_i) F_i) K_k,$$

where the product is now in reverse order, i.e., $\Pi_{i=k+1}^{n-1} (I - K_i H_i) F_i := (I - K_{n-1} H_{n-1}) F_{n-1} \dots (I - K_{k+1} H_{k+1}) F_{k+1}$. Writing $l_{n|n-1} := z_n^T S_n^{-1} z_n = (y_n - H_n \hat{x}_{n|n-1})^T S_n^{-1} (y_n - H_n \hat{x}_{n|n-1})$ and using standard calculus the Jacobians write

$$\frac{\partial l_{n|n-1}}{\partial y_n} = 2z_n^T S_n^{-1}, \quad \frac{\partial l_{n|n-1}}{\partial \hat{x}_{n|n-1}} = -2z_n^T S_n^{-1} H_n. \quad (9)$$

We have $\frac{\partial z_n}{\partial y_n} = I$, and a simple recursion proves that

$$\frac{\partial \hat{x}_{n|n-1}}{\partial y_k} = F_n (I - K_{n-1} H_{n-1}) F_{n-1} \dots (I - K_{k+1} H_{k+1}) F_{k+1} K_k,$$

which describes how y_k affects subsequent state estimates $n \geq k + 1$. The desired formula follows from the chain rule. ■

Expression (8) may look cumbersome, but we will see it may be computed quite efficiently in the sequel.

D. Considered Types of Failures

Throughout the paper, we suppose the measurements to be corrupted by an additive perturbation, that is, what we actually measure is $(y + \delta y)_{1 \leq n \leq N}$ with $(\delta y)_{1 \leq n \leq N}$ a perturbation function. The considered errors are not independent, in that $(\delta y)_{1 \leq n \leq N}$ is a (deterministic) structured known function of time n , parameterized by quantities to be identified.

1) *Oscillating fault (sinusoidal error)*: This type of failure adds a periodic disturbance that oscillates at a certain frequency, that is, is of the sinusoidal type. It is more difficult to detect it than a punctual defect since it describes a specific signature, but locally the perturbation may remain a small deviation. The challenge is then to be able to detect that such a perturbation is occurring, and to identify its frequency.

2) *Incipient fault (ramp error)*: This failure is characterized by a deviation from healthy measurements that increases linearly over time. As ramps act slowly, they do not correspond to abrupt changes making them more difficult to detect than instantaneous degradations, see [10].




Fig. 1. Illustration of an outlier GPS measurement in case of a navigating vehicle with an inertial measurement unit (IMU) and a GPS receiver.




Fig. 2. Illustration of a GPS failure of the ramp type, i.e., a slowly growing error (SGE), in case of a navigating vehicle with an inertial measurement unit (IMU) and a GPS receiver.

These types of errors occur frequently when using a GPS. They are known in the related literature as slowly growing errors (SGE), see e.g., [4], [7], [8], [5], [18], [14]. This is illustrated on Figure 2. They should be opposed to standard outliers, as illustrated on 1. They also occur in a variety of other domains [10]. The challenge then, in terms of fault identification, is to be able to localize precisely the time when the ramp is injected.

III. PROPOSED METHOD

The method proposed in the present paper for failure detection is based on the sensitivity of the likelihood with respect to the measurements. The negative log-likelihood \mathcal{L} (or inconsistency) is a function of all measurements. Suppose the system was fed with spurious measurements $y_k = \tilde{y}_k + \delta y_k$ for $1 \leq k \leq N$, where \tilde{y}_k is the output corresponding to the system without defect. The likelihood satisfies

$$\mathcal{L}(y_1 - \delta y_1, \dots, y_N - \delta y_N) = \mathcal{L}(y) - \sum_{k=1}^N \frac{\partial \mathcal{L}}{\partial y_k} \delta y_k + \text{h.o.t.} \quad (10)$$

where h.o.t. stands for “higher order terms” in the perturbation function δy , and where we recall we are dealing with scalar products between the gradients and the δy_k . The inconsistency associated with the healthy measurements $\tilde{y}_1, \dots, \tilde{y}_N$ is expected to be significantly smaller.

Our method builds upon (10), which indicates how inconsistency varies when subtracting the candidate error signal δy from the measured output. This may be schematically written

$$\mathcal{L}_{\text{measured}} - \mathcal{L}_{\text{healthy}} = \mathcal{L}(y) - \mathcal{L}(\tilde{y}) = \Delta \mathcal{L} \approx \frac{\partial \mathcal{L}}{\partial y} \cdot \delta y \quad (11)$$

where “ \cdot ” denotes a scalar product¹. If $\frac{\partial \mathcal{L}}{\partial y} \cdot \delta y$ is high, it is likely to affect the measurements, as subtracting $\delta y_1, \dots, \delta y_N$ from the measurements makes the inconsistency (7) drop.

The rationale is as follows. Contrary to standard innovation tests, we calculate the global impact of a failure profile on the likelihood. Ultimately, it may prove a better indicator capable of identifying a broader spectrum of failures. Moreover, structured failures may be better detected in that the perturbation may be small (and thus the signal to noise ratio low) and yet its effect may build up in the sum in the right-hand-side of (10).

A. Fourier transform

When we turn to errors being related to each other over time, in the sense that $(\delta y_k)_{1 \leq k \leq N}$ is a structured function of time, it may prove useful to resort to the discrete Fourier transform (DFT) of the functions. To simplify, let us assume that y_k is a scalar output, and let us denote by $\frac{\partial \mathcal{L}}{\partial y}$ the sequence $(\frac{\partial \mathcal{L}}{\partial y_k})_{1 \leq k \leq N}$. From (10), we see we are interested in the quantity $\sum_{k=1}^N \frac{\partial \mathcal{L}}{\partial y_k} \delta y_k$. Using Parseval’s theorem we have

$$\sum_{k=1}^N \frac{\partial \mathcal{L}}{\partial y_k} \delta y_k = \frac{1}{N} \sum_{k=1}^N \mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\})_k \cdot \mathcal{F}(\{\delta y\})_k^*, \quad (12)$$

where $\mathcal{F}(\cdot)$ denotes the DFT. We now leverage this formula.

B. Application to sinusoidal error detection

Assume the signal is corrupted by a sine wave, that is,

$$\delta y_k = \epsilon \sin(kl \frac{2\pi}{N}) \quad (13)$$

for $1 \leq j \leq N$, where ϵ is a small scalar. The number of cycles l represents the frequency, which is the quantity we are seeking to identify. It corresponds to an oscillation period of $T = N/l$. The DFT of this sine wave writes:

$$\mathcal{F}(\{\delta y\})_k = \begin{cases} -\epsilon \frac{N}{2} l & \text{if } k = l \\ \epsilon \frac{N}{2} l & \text{if } k = N - l \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

Recalling (12) and (11), This leads to

$$\Delta \mathcal{L} = \frac{\epsilon}{2} (i \mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\})_l - i \mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\})_{N-l}). \quad (15)$$

As a result, if a sinusoidal wave with number of cycles l is corrupting the measurements, we shall see two peaks in the DFT of the computed function. More precisely, owing to

¹There are two different scalar products, not to be confused. Each $\frac{\partial \mathcal{L}}{\partial y_k} \delta y_k$ is a scalar product between vectors of dimension $\dim y$. We omit the dot as it can be written as a row times column matrix multiplication. The \cdot is to be understood as a scalar product between N -dimensional elements, i.e., the sequence of gradients $(\frac{\partial \mathcal{L}}{\partial y_k})_{1 \leq k \leq N}$ and that of errors $(y_k)_{1 \leq k \leq N}$.

the properties of the DFT, it is in fact sufficient to evaluate the Fourier transform on half the interval, leading to the following mathematical result.

Proposition 2: Assume the signal to be corrupted by an additive sine wave (13) with period $T = N/l$. Then the maximum of the drop in inconsistency in the sense of (11) corresponds to

$$\max_{1 \leq l \leq N/2} -\epsilon \Im(\mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\}_l)), \quad (16)$$

where \Im denotes the imaginary part.

Proof: As $\frac{\partial \mathcal{L}}{\partial y}$ is a real function, we have $\mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\}_k) = \mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\}_{N-k})^*$, so $\Delta \mathcal{L} = \epsilon \Re(i \mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\}_l)) = -\epsilon \Im(\mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\}_l))$. ■ It is thus sufficient to compute the DFT on half the interval for the present purpose, and to seek a peak.

C. Application to ramp error detection

Suppose the signal is corrupted by a ramp that starts at some time instant m and of the following form

$$\delta y_k^{(m)} = \begin{cases} 0 & \text{if } k < m \\ \epsilon(k-m) & \text{if } k \geq m \end{cases} \quad (17)$$

and that one seeks to identify the starting point m .

It turns out this can be efficiently detected too, using once again the DFT. Indeed, in the present case the inconsistency drop (11) due to the ramp writes

$$\Delta \mathcal{L} = \sum_{k=1}^N \frac{\partial \mathcal{L}}{\partial y_k} \delta y_k^{(m)} = \sum_{k=1}^N \frac{\partial \mathcal{L}}{\partial y_k} x_{m-k}$$

where x is a fixed sequence defined by

$$x_l = \begin{cases} -\epsilon l & \text{if } l \leq 0 \\ 0 & \text{if } l > 0. \end{cases}$$

Prolonging the sequence $(\frac{\partial \mathcal{L}}{\partial y_k})_{1 \leq k \leq N}$ with null values outside the interval $\{1, 2, \dots, N\}$, the sum of interest may be written in turn as a convolution product

$$\Delta \mathcal{L} = [(\frac{\partial \mathcal{L}}{\partial y}) * x]_m. \quad (18)$$

Applying the convolution theorem for discrete sequences we get our final result for ramp detection.

Proposition 3: The drop in inconsistency related to a ramp starting at time m is given by the value

$$\Delta \mathcal{L} = \mathcal{F}^{-1}(\{\mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\}) \mathcal{F}(\{x\})\})_m \quad (19)$$

Proof: By the convolution theorem $\Delta \mathcal{L} = \sum_{k=1}^N \frac{\partial \mathcal{L}}{\partial y_k} \delta y_k^{(m)} = [(\frac{\partial \mathcal{L}}{\partial y}) * x]_m = \mathcal{F}^{-1}(\{\mathcal{F}(\{\frac{\partial \mathcal{L}}{\partial y}\}) \mathcal{F}(\{x\})\})_m$. ■

As a result, a large value of function (19) at m indicates a likely time at which the ramp has started corrupting the measurements. It thus remains to evaluate and compare the function at different m .

We see that the method may more generally be applied to detect the time when any (fixed) known profile x is added to the signal, beyond a ramp, e.g., a truncated ramp [10], or step functions. Note that, the slope of the ramp does not play a role for detection, since it acts linearly on the scalar product.

D. Numerical complexity

We now show that the whole pipeline necessitates a computation cost of $O(N \log N)$ operations, and can thus be implemented even if the number of measurements N is large.

a) *Efficient computation of all gradients:* The first step of our method consists in computing all the gradients of the inconsistency $(\frac{\partial \mathcal{L}}{\partial y_k})_{1 \leq k \leq N}$. From (8), we see the overall computation time for obtaining all gradients is at least quadratic in N .

However, leveraging our very recent results on back-propagation through Kalman filters [12], it turns out one may compute all the gradients linearly in the number of measurements N . Indeed, (8) satisfies a simple recursion, so that intermediary computations steps are reused in the previous gradient, along the lines of dynamic programming. It is proved in [12] indeed that the gradients satisfy the following backward recursion

$$u_N = 0_{1 \times d} \quad (20)$$

$$\frac{\partial \mathcal{L}}{\partial y_n} = K_n^T u_n + 2S_n^{-1} z_n \quad n \leq N \quad (21)$$

$$u_{n-1} = F_n^T [(I - K_n H_n)^T u_n - 2H_n^T S_n^{-1} z_n] \quad (22)$$

It can be checked that this coincides with (8) indeed, but it makes the computation cost of $(\frac{\partial \mathcal{L}}{\partial y_k})_{1 \leq k \leq N}$ of order $O(N)$ as it amounts to implementing a recursive backward calculation over the time window. Those equations thus appear as an enabler for efficient implementation.

b) *Efficient computation of the DFT:* Once the sequence of gradients has been obtained, we may use the FFT to compute its DFT, leading to a $O(N \log N)$ computation cost.

E. Interest of the method

The method bears a superficial resemblance to the local approach [2], based on a Taylor expansion of the log-likelihood, but the latter performs statistical tests related to the asymptotic behavior. More recently, [10] proposed methods to detect ramps, based on a likelihood ratio between nominal model and ramp starting at time m , for all possible m . In the latter and more generally in the fault detection literature, one tests one or multiple hypotheses for a change in a parameter, which yields a large computation cost when the number of hypotheses increases. Let us compare our approach to the existing for both considered types of faults.

a) *Oscillating fault:* A standard approach in this case would be to compute the Fourier transform (DFT) of the output signal. However, this shall not prove efficient for frequency recovery, as the signal to noise ratio may be weak, especially if the amplitude of the perturbation is small. By contrast, our approach allows for a much clearer identification since it assesses the direct effect of additive sinusoids on the likelihood with all possible frequencies: by subtracting the correct sinusoid the likelihood dramatically changes, and this shows directly on the FFT of the gradients.

b) *Incipient fault*: Detecting ramps is a challenge as they do not come as abrupt changes. The innovation does not allow for detection as the fault is present in all measurements, and starts slowly. A classical route would consist in using the multiple model adaptive estimation (MMAE) [11], where in essence we test multiple candidate models, each corresponding to a different starting time in the interval. This leads to a computationally very demanding solution as many Kalman filters must be run in parallel. To lower computation demands, one could test the interacting multiple models (IMM) instead [1], [17] but it is suboptimal and testing is efficiency and tuning it on the given problem is beyond the scope of this paper. In any case, our point is that we do not have to test many different models with combinatorial explosion: indicators are obtained with linear complexity in the number of measurements data. Another approach [6] to such soft failures in navigation consists in comparing the Kalman filter estimates with the prior (dead-reckoning) model and perform a chi-square test, but this does not readily yield the starting point of occurrence.

IV. NUMERICAL EXPERIMENTS

We now illustrate on a simple example the relevance of the proposed indicators.

A. Setting

We prove the efficiency of our approach on synthetic data from a particle moving on the real line, and whose acceleration is controlled. The dynamics are

$$p_n = p_{n-1} + \Delta t v_{n-1} + w_n^p \quad (23)$$

$$v_n = v_{n-1} + \Delta t a_{n-1} + w_n^v \quad (24)$$

where $\Delta t = 1$ is the integration step, $a_n \in \mathbb{R}$ the acceleration (control input) and w_n is the process noise which is assumed to be zero mean Gaussian white noise with diagonal covariance matrix Q . The measurements y_n 's consist of the noisy position:

$$y_n = p_n + v_n$$

where v_n is the measurement noise which is assumed to be a centered Gaussian white noise with variance R .

We let² $Q = \text{diag}(1e^{-8}, 1e^{-4})$ and $R = 1.5$. The state is estimated by a Kalman filter (see Section II-A) with input $u_n = a_n$ and position measurement y_n .

The true system's trajectory and the KF's estimates are displayed on Figure 3. The trajectory is obtained by injecting a large acceleration a_n roughly every 20 seconds.

B. Results: Ramp error detection

The measured position is then corrupted by an additive incipient ramp (SGE) starting at $m = 39$ and defined by:

$$\delta y_k = \begin{cases} 0 & \text{if } k < 39 \\ 0.1(k-39) + 0.1 & \text{if } k \geq 39 \end{cases} \quad (25)$$

This is shown on Figure 4.

²The magnitudes of the noises were chosen having in mind inertial navigation with an IMU and a GPS as a motivation. In commercial airliners, the IMUs are highly accurate and come with very low process noise.




Fig. 3. True (simulated) position and velocity over time, and their estimates obtained by the Kalman filter, in the absence of fault.

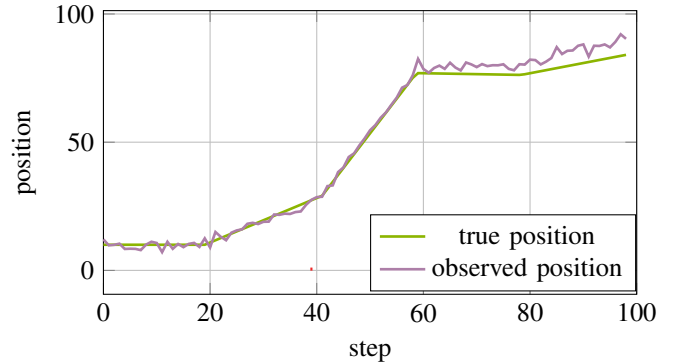


Fig. 4. True position and measured position corrupted by white noise plus a spurious ramp starting at $m = 39$ and with slope $\epsilon = 0.1$.

Figure 5 shows the result of the detection of the ramp, plotting the inconsistency drop given by (18) or alternatively (19). The theory is confirmed, as the maximum of the inconsistency drop is clearly seen at the starting time of the ramp. This is remarkable as we can see from Figure 4 that the signal-to-noise ratio is very low. When compared to the normalized innovation squared (NIS), we see the proposed indicator is much more relevant. The plot of the NIS confirms that no dramatic change is visible in the innovation process.




Fig. 5. Comparison of our proposed indicator and the classic NIS to detect a ramp fault starting at $m = 39$.




Fig. 6. True position and measured position corrupted by white noise plus a spurious sinusoid $t \mapsto 2.5\sin(\frac{t}{2})$ i.e. of frequency $\frac{1}{4\pi}$ and amplitude 2.5.




Fig. 7. Comparison of our indicator with the Fourier transform of the innovation process when the signal is corrupted by white noise plus a sinusoidal defect of frequency $\frac{1}{4\pi} \approx 0.08$.

C. Results: Sinusoidal error

Figure 6 shows the measured position in the presence of measurement noise plus a sine wave $t \mapsto 2.5\sin(\frac{t}{2})$. This corresponds to a frequency $f = \frac{1}{4\pi}$ and as we have $N = 100$ measurements, it corresponds to a number of cycles $l = N/T = 100/(4\pi) \approx 8$.

Figure 7 shows the result of the detection of the sinusoid, plotting the (normalized) quantity $\mathfrak{I}(\mathcal{F}(\frac{\partial \mathcal{L}}{\partial y}))_l$ that is opposite to the inconsistency drop, see (16). The theory is confirmed, as the maximum (of the opposite signal) is clearly seen at the actual frequency of the sinusoid. This is remarkable, here again, as we can see from Figure 6 the signal-to-noise ratio seems very degraded. It is instructive to see the Fourier transform of the innovation process does not reveal any relevant information regarding the frequency that corrupts the measured signal.

V. CONCLUSION

In this paper we have proposed a method to detect structured faults in the measured outputs of a dynamical system. It relies on the use of a Kalman filter, and a Taylor expansion of the likelihood of the measurements, and its efficient implementation is enabled by the recent results of [12]. This provides a global indicator that allows for detection of structured profiles along with their parameters in spite of an unfavorable signal-to-noise ratio, and without the need to know their magnitude. Contrary to multiple hypothesis testing, the method is very efficient computationally and

requires (almost) linear runtime in the data size, enabling onboard implementations.

Although the method may be run onboard, it operates on the fixed sequence of the N past measurements. In the future, it would be worth investigating whether an online counterpart may be derived. Another advantage of the method is that it is not reserved for scalar outputs, and suits multidimensional outputs. In this regard, it would be interesting to test it on more complex datasets obtained through real experiments, notably when navigating with a GPS. Finally, it would be nice to adapt it to the context where the additive perturbation depends on more parameters, for instance a ramp that starts at unknown time m and saturates at unknown time $m' > m$.

REFERENCES

- [1] Bar-Shalom, Y., Li, X.R., Kirubarajan, T.: Estimation with applications to tracking and navigation: theory algorithms and software. John Wiley & Sons (2001)
- [2] Basseville, M.: On-board component fault detection and isolation using the statistical local approach. *Automatica* **34**(11), 1391–1415 (1998)
- [3] Basseville, M., Nikiforov, I.V., et al.: Detection of abrupt changes: theory and application, vol. 104. prentice Hall Englewood Cliffs (1993)
- [4] Bhatti, U.I., Ochieng, W.Y., Feng, S.: Integrity of an integrated GPS/INS system in the presence of slowly growing errors. part i: A critical review. *Gps Solutions* **11**, 173–181 (2007)
- [5] Bhatti, U.I., Ochieng, W.Y., Feng, S.: Performance of rate detector algorithms for an integrated gps/ins system in the presence of slowly growing error. *GPS solutions* **16**, 293–301 (2012)
- [6] Brumback, B., Srinath, M.: A chi-square test for fault-detection in kalman filters. *IEEE Transactions on Automatic Control* **32**(6), 552–554 (1987)
- [7] Diesel, J., King, J.: Integration of navigation systems for fault detection, exclusion, and integrity determination-without waas. In: Proceedings of the 1995 National Technical Meeting of The Institute of Navigation. pp. 683–692 (1995)
- [8] Feng, S., Ochieng, W.: A difference test method for early detection of slowly growing errors in GNSS positioning. *The Journal of Navigation* **60**(3), 427–442 (2007)
- [9] Gupta, N., Mehra, R.: Computational aspects of maximum likelihood estimation and reduction in sensitivity function calculations. *IEEE transactions on automatic control* **19**(6), 774–783 (1974)
- [10] Kiasi, F., Prakash, J., Shah, S.: Detection and diagnosis of incipient faults in sensors of an LTI system using a modified GLR-based approach. *Journal of Process Control* **33**, 77–89 (2015)
- [11] Magill, D.: Optimal adaptive estimation of sampled stochastic processes. *IEEE Transactions on Automatic Control* **10**(4), 434–439 (1965)
- [12] Parellier, C., Barrau, A., Bonnabel, S.: Speeding up backpropagation of gradients through the Kalman filter via closed-form expressions. *IEEE Transactions on Automatic Control* (2023). <https://doi.org/10.1109/TAC.2023.3297879>
- [13] Särkkä, S.: Bayesian filtering and smoothing. No. 3, Cambridge university press (2013)
- [14] Sun, A., Zhang, Q., Yu, Z., Meng, X., Liu, X., Zhang, Y., Xie, Y.: A novel slow-growing gross error detection method for gnss/accelerometer integrated deformation monitoring based on state domain consistency theory. *Remote Sensing* **14**(19), 4758 (2022)
- [15] Willsky, A., Jones, H.: A generalized likelihood ratio approach to the detection and estimation of jumps in linear systems. *IEEE Transactions on Automatic Control* **21**(1), 108–112 (1976). <https://doi.org/10.1109/TAC.1976.1101146>
- [16] Willsky, A.S.: A survey of design methods for failure detection in dynamic systems. *Automatica* **12**(6), 601–611 (1976)
- [17] Zhang, Y., Li, X.R.: Detection and diagnosis of sensor and actuator failures using IMM estimator. *IEEE Transactions on aerospace and electronic systems* **34**(4), 1293–1313 (1998)
- [18] Zhao, X., Gao, C., Pang, C., Zhang, C., Wang, Y.: A double-threshold test method for soft faults assisted by bp neural network. *Control and Decision* **35**(6), 1384–1390 (2020)