

Detection and identification of faults in clock ensembles with the generalized likelihood ratio test

Christian Trainotti^{1,*} , Gabriele Giorgi¹  and Christoph Günther^{1,2}

¹ Institute of Communications and Navigation, German Aerospace Center (DLR), Münchener Straße 20, 82234 Weßling, Germany

² Chair of Communication and Navigation, Technical University of Munich, Theresienstraße 90, 80333 München, Germany

E-mail: christian.trainotti@dlr.de

Received 11 January 2022, revised 5 May 2022

Accepted for publication 7 June 2022

Published 19 July 2022



CrossMark

Abstract

In this paper we propose an approach for performing fault detection and identification in clock ensembles based on the generalized likelihood ratio test. We show that by applying a set of purposefully-designed statistical tests, one can successfully detect faults occurring in a clock of the ensemble, and identify which measurement in the ensemble is most likely to have triggered the detection. We first develop the theoretical framework for the characterization of the detectors and their performance, and validate the derivations via Monte Carlo simulations. Then, we apply the statistical tests to an ensemble of cesium clocks, aiming at detecting and identifying three types of non-nominal behaviors. The faulty conditions are obtained by injecting a pattern of phase steps, a phase and frequency drift, and an oscillatory phase component.

Keywords: GLRT, clock ensemble, integrity, fault detection, timing

(Some figures may appear in colour only in the online journal)

1. Introduction

Clock ensembles play a fundamental role in providing a stable and robust time scale for high-integrity applications. Compared to single clock approaches, clock ensembles are characterized by improved stability and robustness. However, clock faults can still impact the stability and availability of the generated time scale. Thus a clock ensemble must be equipped with a dedicated fault monitoring system performing three steps: fault detection, fault identification, and fault adaptation. Detecting a fault means raising an alarm when a fault occurrence is observed, while the identification step aims at locating where in the system the fault is occurring. Finally, fault

adaptation consists in taking corrective actions to reduce the impact of the fault on the system, for instance by repairing or temporarily removing the faulty unit. Since in this paper we only tackle the first two steps, we will talk of fault detection and identification (FDI).

A FDI algorithm is implemented in a detector. The detector observes a particular quantity (the observable) and scans for deviations from the expected nominal behavior, which are caused by the occurrence of a fault. Different faults act differently on the system, thus we must choose the right observable (or a set thereof) to properly identify different types of expected errors. In this work we use three observables, namely Kalman filter residuals, phase measurements, and dynamic Allan variance (DAVAR) of differential clock measurements. For each of these we devise a dedicated detector based on the generalized likelihood ratio test (GLRT), which accounts for the statistical distribution of the observable in both nominal and faulty conditions. We derive a model-based test and a

* Author to whom any correspondence should be addressed.


 Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Table 1. Selected literature on integrity for clocks and clock ensembles.

Method	Single clock	Clock ensemble
Kalman filter residuals	[4–6, 29]	[3, 23–25]
GLRT on clock measurements	[7–11]	<i>This work</i>
Dynamic Allan variance	[12–16]	<i>This work</i>
Trend analysis—smoothing—recursive filters	[18–20]	[27, 28]
Optimal stopping method	[21]	—
Average, least-squares, standard deviation on sliding windows	[17]	[26]
Interferometry	[22]	—

self-consistency test. In the model-based test the detector compares the current observation to the value predicted by a clock model: any deviation between the two can originate from a fault, but also from a mismodelling of the expected behavior. We apply the model-based test on the entire set of clock measurements (*overall model test*) to perform fault detection, and then singularly on each measurement to perform fault identification (*w-test*) [1, 2]. In the case of homogeneous ensembles, which only contain clocks of the same type, we construct the *self-consistency test*. This test does not require to define clock noise parameters. The availability of multiple clocks of the same type allows the estimation of the nominal clock behavior directly from the observations. Instead of detecting deviations between observed and modelled behaviors, we now detect deviations between one measurement and the average observed behavior of all measurements.

After deriving the different types of statistical tests, we describe the expected performance of each in terms of probability of false alarm and probability of missed detection. Monte Carlo simulations are then used to validate the theoretical derivation. Finally, we perform an experimental assessment of all the tests devised by injecting different types of fault patterns in a laboratory clock ensemble, and we verify whether the tests can successfully detect and identify the faulty signal.

1.1. State of the art

Table 1 shows a selection of literature on fault detection in clocks and clock ensembles. While most of the studies focus on detecting faults on single clock signals, some works specifically target clock ensembles. In [3] the authors propose the usage of a Kalman filter for fault detection in a clock ensemble: by tracking the evolution of the Kalman filter residuals it is possible to detect discrepancies between model and observations. This method is further investigated in [4] and applied in [5, 6]. The use of the GLRT for detecting clock failures is described in [7], and further developed in [8–10]. In [11] the method is applied to the detection of anomalies on a GPS satellite clock. The DAVAR method, detailed in [12–14], makes use of the overlapping Allan variance (OAVAR) computed on sliding windows of given length, and compares the temporal changes in the data to predefined DAVAR values. The detector is triggered if the deviations exceed a statistically significant threshold. The application of this method to clocks onboard satellites of global navigation satellite systems (GNSSs) is

shown in [15, 16]. In [17] the author describes three algorithms for detection of jumps in frequency measurements: the BLKAVG algorithm computes the average of frequency points on non-overlapping moving windows, detecting a jump when the difference in averages among two windows is larger than a given threshold; the SEQAVG algorithm employs an averaging window only once a potential jump is detected, to confirm or reject the detection; finally, the CUSUM algorithm computes the cumulative sum of the differences between each frequency data point and the average value on the entire data set. The slope of the cumulative sum reflects the average behavior of the measurement and can be used to detect frequency jumps. Other techniques for fault detection in clock signals include trend analysis and filtering of the frequency signal [18–20], the optimal stopping method [21], and interferometric analysis [22]. Fault detection in clock ensembles is further discussed in [23–25] using Kalman filters. Instead, the methods in [26–28] use least-squares fitting, standard deviation computed on sliding windows, and infinite impulse response filters.

In this work we integrate and extend some of these approaches for FDI. Expanding from our previous works [30, 31] we provide a theoretical framework for the design of fault detectors in clock ensembles. The techniques shown here can be directly applied to real-time dependable timing system, greatly improving the robustness of the generated time scale. The paper shows how to take advantage of the GLRT for the observation of different quantities, so that the integrity layer of a system can be easily expanded with further detectors using the same approach. Furthermore, the GLRT allows us to predict the performance of the detectors in terms of probability of false alarm (P_{fa}), probability of missed detection (P_{md}), and minimum detectable bias (MDB). This enables tuning the detectors to fulfill the desired integrity requirements, namely which magnitude of fault is detectable and with which probability. By taking as example a precise timing facility running a clock ensemble to provide timing services, these methods would benefit both the provider and its users. Firstly, the provider can guarantee a given integrity level, since the detection performances can be determined beforehand; then, in case of a faulty clock, warnings can be promptly provided to the users, while the fault is located in the system and corrective actions are put in place. An additional advantage of these methods is represented by the self-consistency test, which detects whether one of the clocks diverges from the ensemble's average behavior. Since no *a priori* clock noise model is assumed, this test is not triggered by any influence affecting all clocks

in the same way. The evaluation of clocks on board GNSS satellites is an example of such a benefit: the estimation of the clocks' states is affected by the residuals of the orbit determination process, which determine an increase of the observed OAVAR for sampling intervals around half of the orbital period [32]. A model-based test would in this case trigger a detection since the observed behavior differs from the clock model, even if the misbehavior is not due to the clock itself. Instead, the self-consistency test would not trigger, since the additive bias is present on all clocks.

2. Clock ensemble

The system under study consists of a set of clocks, a measurement device, a Kalman filter and a number of detectors. While we keep the theoretical development as general as possible concerning number and type of clocks, the experimental verification is performed with an ensemble of 5 cesium frequency references. Figure 1 shows a schematic of the ensemble. The measurement device generates differential clock measurements, which are phase and/or frequency differences between two clocks. In our setup the first clock acts as reference against which we measure the remaining units, but different measurement topologies can be used. The clock measurements are fed to the Kalman filter, which first predicts the states of the clocks according to a specified clock model, and then updates the predicted states with the available observations. The Kalman filter is executed on a laboratory computer connected to the measurement device. This computer also runs the FDI algorithm. Furthermore, a dedicated device can inject a desired fault pattern into the signal of a clock, so that we can run experiments for assessing the performance of the detectors. The injected fault is indicated by $\nabla(t)$ in figure 1.

2.1. Clock model and ensemble model

We use a two-state clock model with constant frequency drift and Markov processes, as described in [33–35]. The state vector of each clock contains one state representing the phase deviation and one state for the frequency deviation. When required for more accurate modelling, the state vector of the single clock is extended with one or more additional states, related to additive Markov processes [34]. We combine the states of each clock to obtain the state vector \mathbf{x} of the entire ensemble. In time discrete notation we describe the dynamics of the system and its observation with

$$\mathbf{x}_k = \Phi(\tau_0) \mathbf{x}_{k-1} + \mathbf{d}(\tau_0) + \mathbf{n}_{k-1}, \quad (1)$$

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}_k, \quad (2)$$

where the subscript \bullet_k indicates the quantities at time t_k , $\tau_0 = t_k - t_{k-1}$ is the constant time discretization, Φ the state propagation matrix, \mathbf{d} the drift vector, \mathbf{z} the vector of clock measurements, and \mathbf{H} the measurement matrix. The process noise \mathbf{n}_k and the measurement noise \mathbf{v}_k are assumed Gaussian:

$$\mathbf{n}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}(\tau_0)), \quad (3)$$

$$\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \quad (4)$$

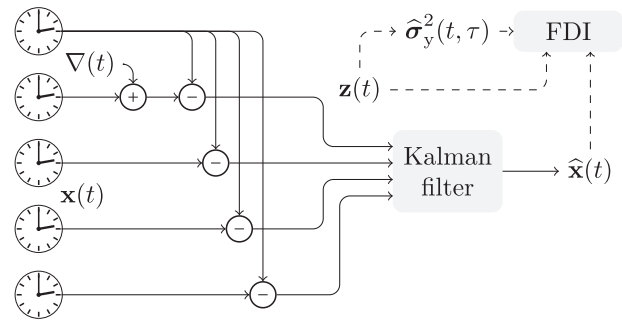


Figure 1. Functional representation of the clock ensemble under study, with measurement topology and processing units: the clocks, whose state vector is $\mathbf{x}(t)$, are measured by a device generating clock measurements $\mathbf{z}(t)$. A Kalman filter computes the estimates $\hat{\mathbf{x}}(t)$ of the clock states. The FDI algorithm uses the clock measurements $\mathbf{z}(t)$, the states estimate $\hat{\mathbf{x}}(t)$, and the dynamic Allan variance (DAVAR) $\hat{\sigma}_y^2(t, \tau)$ of the clock measurements to observe and evaluate the state of the system. For assessing the FDI algorithm, we inject a fault $\nabla(t)$ to the signal of the second clock. The computer executing the algorithm is not depicted.

where \mathbf{Q} is the process noise covariance matrix and \mathbf{R} is the covariance matrix of the measurement noise. The matrices Φ and \mathbf{Q} are block diagonal, where each block is the corresponding matrix for the single clock i , ${}_i\Phi(\tau_0)$ and ${}_i\mathbf{Q}(\tau_0)$, computed with the model parameters of the respective clock ${}_i\sigma_1$ and ${}_i\sigma_2$:

$${}_i\Phi(\tau_0) = \begin{bmatrix} 1 & \tau_0 \\ 0 & 1 \end{bmatrix}, \quad (5)$$

$${}_i\mathbf{Q}(\tau_0) = \begin{bmatrix} {}_i\sigma_1^2\tau_0 + {}_i\sigma_2^2\frac{\tau_0^3}{3} & {}_i\sigma_2^2\frac{\tau_0^2}{2} \\ {}_i\sigma_2^2\frac{\tau_0^2}{2} & {}_i\sigma_2^2\tau_0 \end{bmatrix}. \quad (6)$$

If Markov processes need to be added to model specific behaviors, we expand the matrices accordingly [34]. Denoting the constant frequency drift of the i -th clock as ${}_i d$, we generate the drift vector by stacking the vectors ${}_i\mathbf{d}(\tau_0)$ of the single clocks:

$${}_i\mathbf{d}(\tau_0) = {}_i d \begin{bmatrix} \frac{\tau_0^2}{2} \\ \tau_0 \end{bmatrix}. \quad (7)$$

The measurement matrix \mathbf{H} describes how the clocks are measured with respect to each other. It depends on the measurement topology of the ensemble and on whether differential phase or frequency measurements are available. In our case we measure the phase of each clock with respect to the first clock. Thus we obtain $M = N - 1$ measurements, where N is the number of clocks in the ensemble. Without Markov processes the measurement matrix is

$$\mathbf{H} = \begin{bmatrix} -1 & 0 & 1 & 0 & 0 & 0 & \dots & 0 & 0 \\ -1 & 0 & 0 & 0 & 1 & 0 & \dots & 0 & 0 \\ \vdots & & & & & & & & \\ -1 & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix}. \quad (8)$$

Additionally, we define a second measurement matrix $\overline{\mathbf{H}}$, which maps the clock units to the measurements. It can be obtained from \mathbf{H} by eliminating the null columns. For instance, the matrix corresponding to the aforementioned topology with four clocks and three differential measurements is

$$\overline{\mathbf{H}} = \begin{bmatrix} -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \end{bmatrix}. \quad (9)$$

This matrix represents the measurement topology without dependence on the number of states used to model the system.

2.2. Kalman filter

We run a clock ensembling algorithm based on a Kalman filter, as described in [34, 36]. At each time step, the filter provides an estimate of the clock states by executing the following steps:

$$\widehat{\mathbf{x}}_k^- = \Phi(\tau_0) \widehat{\mathbf{x}}_{k-1} + \mathbf{d}(\tau_0), \quad (10)$$

$$\mathbf{P}_k^- = \Phi(\tau_0) \mathbf{P}'_{k-1} \Phi^\top(\tau_0) + \mathbf{Q}(\tau_0), \quad (11)$$

$$\mathbf{K}_k = \mathbf{P}_k^- \mathbf{H}^\top (\mathbf{H} \mathbf{P}_k^- \mathbf{H}^\top + \mathbf{R})^{-1}, \quad (12)$$

$$\widehat{\mathbf{x}}_k = \widehat{\mathbf{x}}_k^- + \mathbf{K}_k (\mathbf{z}_k - \mathbf{H} \widehat{\mathbf{x}}_k^-), \quad (13)$$

$$\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{H}) \mathbf{P}_k^-, \quad (14)$$

$$\mathbf{P}_{r,k} = \mathbf{S} \mathbf{P}_k \mathbf{S}^\top, \quad (15)$$

$$\mathbf{P}'_k = \frac{\mathbf{P}_{r,k} + \mathbf{P}_{r,k}^\top}{2}. \quad (16)$$

$\widehat{\mathbf{x}}^-$ is the predicted estimate of the state vector \mathbf{x} , with associated error covariance matrix \mathbf{P}^- , \mathbf{K} is the Kalman gain, used to generate the updated state estimate $\widehat{\mathbf{x}}$, with associated error covariance \mathbf{P} . To avoid the unbounded growth of the components in \mathbf{P} , we employ the covariance x -reduction method (15) described in [36]. The reduced covariance matrix \mathbf{P}_r is obtained via the diagonal matrix \mathbf{S} , an identity matrix whose entries corresponding to the phase components in the state vector are set to zero. Finally, (16) ensures that the resulting reduced covariance matrix is diagonal. In (14), \mathbf{I} denotes an identity matrix, whose size is the number of states in \mathbf{x} .

3. Faults, observables and hypotheses

In this paper we consider two different classes of faults. The first kind consists of abrupt changes in the observed phase or frequency signals, such as jumps and steps, which present a very fast dynamics. The second type of faults shows a slower behavior and includes phase or frequency drifts as well as oscillatory components in the signals. A drift can arise due to the aging of the clocks, while the oscillation can generate from periodic effects, such as daily temperature fluctuations or orbital-related dynamics. While abrupt faults are promptly reflected into the Kalman filter residuals, drifts cause a step-to-step phase and frequency deviation which might be too small to be detected by testing the filter residuals. Thus, to detect this

type of faults we need a quantity which ‘accumulates’ the error, such as the phase measurements and their DAVAR. Clearly, the required accumulation of a sufficient number of samples introduces a delay between the fault onset and the time of first detection.

In the following sections we introduce the theoretical descriptions of the three observables: Kalman filter residuals, phase measurements, and DAVAR. For each observable we define a *null hypothesis* \mathcal{H}_0 , which describes its statistical distribution when the system operates under nominal conditions. An *alternative hypothesis* \mathcal{H}_A describes instead the observable’s distribution when a fault is occurring in the system. With these hypotheses we then derive a statistical test decision based on the GLRT.

3.1. Kalman filter residuals

The Kalman filter residuals track the deviations between the observed clock measurements \mathbf{z} and the measurements predicted by the model. Recalling (13) and (14), we express the residuals and their covariance as:

$$\boldsymbol{\rho}_k^{(\text{kf})} = \mathbf{z}_k - \mathbf{H} \widehat{\mathbf{x}}_k^- \quad (17)$$

$$\boldsymbol{\Omega}_k^{(\text{kf})} = \mathbf{R} + \mathbf{H} \mathbf{P}_k^- \mathbf{H}^\top. \quad (18)$$

In nominal conditions we expect the residuals to remain small, while their magnitude increases if the observed measurements differ from the modelled behavior. We therefore define the two hypotheses:

$$\begin{cases} \mathcal{H}_0^{(\text{kf})} : \boldsymbol{\rho}_k^{(\text{kf})} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}_k^{(\text{kf})}) \\ \mathcal{H}_A^{(\text{kf})} : \boldsymbol{\rho}_k^{(\text{kf})} \sim \mathcal{N}(\mathbf{C} \nabla_k, \boldsymbol{\Omega}_k^{(\text{kf})}) \end{cases}. \quad (19)$$

Here we introduce the vector of faults ∇ affecting the system, and the matrix \mathbf{C} describing how the faults act on the observables. The size of vector ∇ depends on the shape of matrix \mathbf{C} , whose free design allows the selection of a given fault mode to be tested. For example, setting $\mathbf{C} = \mathbf{I}_M$, an identity matrix of size M , enables us to perform fault detection by assuming a potential fault on all elements of the observable vector, with ∇ becoming a vector of M entries. When assigning $\mathbf{C} = \mathbf{c}_i$ instead, with \mathbf{c}_i a M -elements vector with 1 at the i -th entry and zero elsewhere, we can scan all single measurements $i = 1 \dots M$ to locate the faulty one. In this case ∇ is a scalar.

3.2. Phase measurements

We describe the time-varying distribution of the phase measurements of the clocks with the model equation (1):

$$\mathbf{x}_k \sim \mathcal{N}[\Phi(t_k - t_0) \mathbf{x}(t_0) + \mathbf{d}(t_k - t_0), \mathbf{Q}(t_k - t_0)], \quad (20)$$

where $\mathbf{x}(t_0)$ contains the initial phase and frequency states of the ensemble. Then, from (2) we have

$$\mathbf{z}_k \sim \mathcal{N}(\zeta_k, \boldsymbol{\Omega}_k^{(z)}), \quad (21)$$

where

$$\zeta_k = \mathbf{H} [\Phi(t_k - t_0) \mathbf{x}(t_0) + \mathbf{d}(t_k - t_0)], \quad (22)$$

$$\Omega_k^{(z)} = \mathbf{H}\mathbf{Q}(t_k - t_0) \mathbf{H}^\top + \mathbf{R}. \quad (23)$$

We define the phase residuals as the difference between observed and expected phase measurements,

$$\rho_k^{(z)} = \mathbf{z}_k - \zeta_k, \quad (24)$$

so that the corresponding hypotheses can be defined:

$$\begin{cases} \mathcal{H}_0^{(z)} : \rho_k^{(z)} \sim \mathcal{N}(\mathbf{0}, \Omega_k^{(z)}) \\ \mathcal{H}_A^{(z)} : \rho_k^{(z)} \sim \mathcal{N}(\mathbf{C}\nabla_k, \Omega_k^{(z)}) \end{cases}. \quad (25)$$

These hypotheses have the same form of those derived for the Kalman filter residuals (19). While the latter track the evolution of the phase measurements with respect to the previous time step, the phase measurements residuals describe the overall phase evolution from the initial time instant.

3.3. Dynamic Allan variance

The third observable is the DAVAR of differential phase measurements computed on a sliding window of given length. In [12] the DAVAR value is centered in the middle of the sliding window, whereas in this paper the evaluation time corresponds to the right limit of the window. In this way the DAVAR can be used in a real-time scenario employing the latest available data samples. Thus, the DAVAR at time t_k uses the measurements in the window $[t_k - W, t_k]$, where W is the length of the sliding window. We use the notation $_{ij}\hat{\sigma}_{k,\tau,W}^2$ for the DAVAR of the measurement between clocks i and j , at time instant k , for sampling interval τ and window length W . The expected value $_{ij}\bar{\sigma}_\tau^2$ and the variance $\text{Var}(_{ij}\hat{\sigma}_{k,\tau,W}^2)$ of the DAVAR are assumed constant in time and can be expressed using the clock noise parameters:

$$_{ij}\bar{\sigma}_\tau^2 = 3\frac{ijR}{\tau^2} + \frac{ij\sigma_1^2}{\tau} + \frac{ij\sigma_2^2\tau}{3} + ij\frac{\tau^2}{2}, \quad (26)$$

$$\text{Var}(_{ij}\hat{\sigma}_{k,\tau,W}^2) = f(\tau, W, ijR, ij\sigma_1^2, ij\sigma_2^2, ij\tau). \quad (27)$$

Here, the noise parameters refer to the combined contribution of the two clocks included in each measurement. The function in (27) is explicitly derived in [37] for different noise components. Since an expression of the variance also considering Markov processes is not available yet, we use a simplified clock model without additional processes, whose parameters are reported in section 4.4, expression (73).

The DAVAR at time t_k distributes according to a χ^2 -distribution [38], namely (in the following we drop the ij subscript):

$$p_{\tau,W} \frac{\hat{\sigma}_{k,\tau,W}^2}{\bar{\sigma}_\tau^2} \sim \chi^2(p_{\tau,W}), \quad (28)$$

where

$$p_{\tau,W} = 2 \frac{\bar{\sigma}_\tau^4}{\text{Var}(\hat{\sigma}_{\tau,W}^2)} \quad (29)$$

are the degrees of freedom. By combining (28) and (29) we define the normalized DAVAR

$$\xi_{k,\tau,W} = \frac{2\hat{\sigma}_{k,\tau,W}^2 \bar{\sigma}_\tau^2}{\text{Var}(\hat{\sigma}_{\tau,W}^2)}. \quad (30)$$

We express the distribution of ξ using the gamma distribution

$$\xi_{k,\tau,W} \sim \chi^2(p_{\tau,W}) \equiv \Gamma\left(\frac{p_{\tau,W}}{2}, 2\right), \quad (31)$$

since it allows us to also express the faulty case. If a fault occurs in the system, we expect a change of the measured DAVAR, therefore we modify the theoretical value by the quantity $\nabla_{k,\tau,W}$:

$$\begin{aligned} \xi_{k,\tau,W} &= \frac{2\hat{\sigma}_{k,\tau,W}^2 (\bar{\sigma}_\tau^2 + \nabla_{k,\tau,W})}{\text{Var}(\hat{\sigma}_{\tau,W}^2)} \\ &= \frac{2\hat{\sigma}_{k,\tau,W}^2 \bar{\sigma}_\tau^2}{\text{Var}(\hat{\sigma}_{\tau,W}^2)} \frac{\bar{\sigma}_\tau^2 + \nabla_{k,\tau,W}}{\bar{\sigma}_\tau^2} \end{aligned} \quad (32)$$

$$\sim \Gamma\left(\frac{p_{\tau,W}}{2}, 2 \frac{\bar{\sigma}_\tau^2 + \nabla_{k,\tau,W}}{\bar{\sigma}_\tau^2}\right). \quad (33)$$

The quantity ξ refers to a single measurement. To design an overall model test, we need to express the distribution of a vector of observables ξ containing the observation of all measurements. However, expanding expression (33) to a multivariate gamma distribution is not trivial. In this work we focus on the scalar case, leaving the extension to the multivariate case for future work. It is worth noting that we can derive an overall model test by approximating the gamma distribution (33) to a Gaussian distribution. This holds for a number of degrees of freedom $p_{\tau,W}$ sufficiently large, which occurs when the sampling interval τ is small compared to the window length W . Under this approximation we can follow the same procedure previously derived for tests on the phase measurements: compute the DAVAR residuals as the difference between observed and expected behaviors, and derive the null and alternative hypotheses, which would be in the same form of (19) and (25). However, in this paper we only focus on the w -test derived from the exact distribution (33) of ξ , which is valid for any sampling interval. The corresponding hypotheses read

$$\begin{cases} \mathcal{H}_0^{(\xi)} : \xi_{k,\tau,W} \sim \Gamma\left(\frac{p_{\tau,W}}{2}, 2\right) \\ \mathcal{H}_A^{(\xi)} : \xi_{k,\tau,W} \sim \Gamma\left(\frac{p_{\tau,W}}{2}, 2 \frac{\bar{\sigma}_\tau^2 + \nabla_{k,\tau,W}}{\bar{\sigma}_\tau^2}\right) \end{cases}. \quad (34)$$

4. Fault detection and identification with the GLRT

In this section we derive the statistical tests based on the observables' distributions. In the GLRT framework we devise a model-based test, which can be in the form of an overall model test or a w -test, and a self-consistency test. The derivation of the detectors starts from the definition of the hypotheses

Table 2. Overview of the observables used and of the tests derived in this work.

Quantity	Notation		Model-based test		Self-consistency test	
	Observable	Residual	Distribution	Overall model test	w -test	
Kalman filter residuals	—	$\rho_k^{(kf)}$	\mathcal{N}	✓	✓	x^a
Phase measurements	\mathbf{z}_k	$\rho_k^{(z)}$	\mathcal{N}	✓	✓	✓
Normalized DAVAR	$\xi_{k,\tau,W}$	—	Γ	x^b	✓	x^b

^aSelf-consistency test is not applicable, since Kalman filter residuals already include model information.

^bNot derived in this work.

on the observed distributions. In section 3 we showed that the hypotheses are of the same form for normally distributed observables (Kalman filter residuals and phase measurements), thus we need only one derivation of a test expression which can then be applied to all of them. Additionally, we need a separate derivation of the w -test for gamma distributed observables (normalized DAVAR). Table 2 gives an overview of the observables and the tests derived in this work.

Assume we have a vector of observables with given distribution $\mathbf{y} \sim f(\mathbf{y}|\alpha)$, where the parameter α takes values in the set \mathcal{G} . If we consider the hypotheses

$$\begin{cases} \mathcal{H}_0 : \alpha \in \mathcal{G}_0 \\ \mathcal{H}_A : \alpha \in \mathcal{G} \end{cases}, \quad (35)$$

where $\mathcal{G}_0 \subset \mathcal{G}$, the GLRT for this problem is then defined as

$$\text{Reject } \mathcal{H}_0 \text{ if } \frac{\max_{\alpha \in \mathcal{G}_0} f(\mathbf{y}|\alpha)}{\max_{\alpha \in \mathcal{G}} f(\mathbf{y}|\alpha)} < a. \quad (36)$$

The threshold a defines the critical region \mathcal{K} , i.e. the set of values of \mathbf{y} leading to a rejection of \mathcal{H}_0 .

Since the measurements are affected by noise, the test can deliver a wrong result. We talk of false alarm (or type I error) if \mathcal{H}_0 is rejected when in fact \mathcal{H}_0 is true. The probability of false alarm P_{fa} of the test (or size of the test) is

$$P_{fa} = \int_{\mathcal{K}} f(\mathbf{y}|\mathcal{H}_0) d\mathbf{y}. \quad (37)$$

A missed detection (type II error) occurs instead when \mathcal{H}_0 is wrongly accepted, when in fact \mathcal{H}_A is true. The probability of missed detection P_{md} is

$$P_{md} = 1 - \int_{\mathcal{K}} f(\mathbf{y}|\mathcal{H}_A) d\mathbf{y}. \quad (38)$$

4.1. Model-based GLRT for normally distributed observables

The Kalman filter residuals and the phase measurements residuals follow a normal distribution. Thus, we generalize the hypotheses (19) and (25) by removing all observable-specific annotations:

$$\begin{cases} \mathcal{H}_0 : \rho \sim \mathcal{N}(0, \Omega) \\ \mathcal{H}_A : \rho \sim \mathcal{N}(C\nabla, \Omega) \end{cases}. \quad (39)$$

This allows us to derive a general expression of the test, which is then used for the different observables by using the respective matrices. By following the derivation shown in appendix A, we obtain the test

$$\text{Reject } \mathcal{H}_0 \text{ if } T = \widehat{\nabla}^\top \mathbf{Q}_{\widehat{\nabla}}^{-1} \widehat{\nabla} > k, \quad (40)$$

with

$$\widehat{\nabla} = (C^\top \Omega^{-1} C)^{-1} C^\top \Omega^{-1} \rho, \quad (41)$$

$$\mathbf{Q}_{\widehat{\nabla}} = (C^\top \Omega^{-1} C)^{-1}. \quad (42)$$

If q is the number of entries in $\widehat{\nabla}$, the test T is χ^2 -distributed with q degrees of freedom. The distribution is central under the null hypothesis and non-central under the alternative hypothesis:

$$\begin{cases} \mathcal{H}_0 : T \sim \chi^2(q, 0) \\ \mathcal{H}_A : T \sim \chi^2(q, \lambda) \end{cases}, \quad (43)$$

with non-centrality parameter

$$\lambda = \nabla^\top \mathbf{Q}_{\widehat{\nabla}}^{-1} \nabla. \quad (44)$$

We design two versions of this test by varying the shape of matrix \mathbf{C} . In the limit case of $q = M$ we obtain the overall model test, while setting $q = 1$ yields the w -test.

4.1.1. Overall model test. In the overall model test we have $q = M$ and $\mathbf{C} = \mathbf{I}$, leading to $\widehat{\nabla} = \rho$. The estimate of fault vector $\widehat{\nabla}$ completely absorbs the residuals ρ , and the test reduces to

$$\text{Reject } \mathcal{H}_0 \text{ if } T = \rho^\top \Omega^{-1} \rho > k_M. \quad (45)$$

The test is χ^2 -distributed with M degrees of freedom, and we compute the threshold k_M from the expression

$$P_{fa} = \int_{k_M}^{+\infty} \chi^2(\eta|q = M, \lambda = 0) d\eta. \quad (46)$$

The probability of missed detection for a given value of non-centrality parameter λ is

$$P_{\text{md}} = \int_0^{k_M} \chi^2(\eta|q = M, \lambda) d\eta. \quad (47)$$

We employ the overall model test for the detection step, in which all the entries of the observation vector are tested as a whole.

4.1.2. w -test. The w -test is derived by setting $q = 1$. The matrix \mathbf{C} becomes a column vector \mathbf{c}_i , where the i th entry is equal to one, and zero elsewhere. By iteratively applying the test to all measurements in the ensemble (varying i from 1 to M), we can sweep all the measurements and locate the one which is most probably faulty. This process is also known as data-snooping. The test becomes

$$\text{Reject } \mathcal{H}_0 \text{ if } T = \frac{\widehat{\nabla}^2}{\sigma_{\widehat{\nabla}}^2} > k_1, \quad (48)$$

where $\widehat{\nabla} = \frac{\mathbf{c}_i^\top \boldsymbol{\Omega}^{-1} \boldsymbol{\rho}}{\mathbf{c}_i^\top \boldsymbol{\Omega}^{-1} \mathbf{c}_i}$ and $\sigma_{\widehat{\nabla}}^2 = \mathbf{c}_i^\top \boldsymbol{\Omega}^{-1} \mathbf{c}_i$ are now scalar. The test is χ^2 -distributed with 1 degree of freedom, central under the null hypothesis, non-central under the alternative hypothesis. We compute the threshold k_1 from the expression

$$P_{\text{fa}} = \int_{k_1}^{+\infty} \chi^2(\eta|q = 1, \lambda = 0) d\eta, \quad (49)$$

while the non-centrality parameter $\lambda = \nabla^2 \mathbf{c}_i^\top \boldsymbol{\Omega}^{-1} \mathbf{c}_i$ is related to the P_{md} :

$$P_{\text{md}} = \int_0^{k_1} \chi^2(\eta|q = 1, \lambda) d\eta. \quad (50)$$

The MDB represents the minimum magnitude of fault that can be detected by the w -test with a given value of P_{md} :

$$|\nabla_{\text{MDB}}| = \sqrt{\frac{\lambda}{\mathbf{c}_i^\top \boldsymbol{\Omega}^{-1} \mathbf{c}_i}}. \quad (51)$$

4.2. Self-consistency test for normally distributed observables

The model-based test requires to know *a priori* how the clocks behave in their environment to determine the parameters of the clock model. This is problematic if the clocks cannot be directly controlled or accessed, for instance for clocks onboard satellites: in this case the application of a clock model not reproducing the actual nominal conditions of the system can lead to mismodelling, missed detection, and in general poor performance of the integrity monitoring techniques. This motivates the self-consistency test, in which we only set the structure of the ensemble and the distribution function of the measurements, but the model parameters are left free and are estimated from the observations [1]. In the case of homogeneous ensembles, where the clocks are of the same type, the test estimates the expected value and variance of the measurements, as well as assessing the occurrence of a fault in a measurement. While the model-based tests are computed from the residuals, defined as the difference between the obser-

vation and the modelled behavior, the self-consistency test is computed directly from the observations. We develop the self-consistency test for the phase measurements (21), while the derivation for gamma distributed observables is left for future work.

In an homogeneous ensemble, we expect all clock measurements to exhibit the same mean value and variance, thus we rewrite expression (21) as

$$\mathbf{z}_k \sim \mathcal{N}(\zeta_k \mathbf{u}, v_k^2 \boldsymbol{\Psi}), \quad (52)$$

where \mathbf{u} is a M -sized vector of ones, ζ_k the magnitude of the expected value, $\boldsymbol{\Psi}$ dictates the shape of the covariance matrix, and v_k^2 controls the magnitude of the covariance matrix. The test directly estimates the scalars ζ_k and v_k^2 from the observations \mathbf{z}_k . We set the matrix $\boldsymbol{\Psi}$ according to the measurement topology of the ensemble:

$$\boldsymbol{\Psi} = \frac{\overline{\mathbf{H}} \overline{\mathbf{H}}^\top}{|\overline{\mathbf{H}} \overline{\mathbf{H}}^\top|} = \frac{1}{M} \overline{\mathbf{H}} \overline{\mathbf{H}}^\top. \quad (53)$$

By dropping the subscript for time dependence, the hypotheses read

$$\begin{cases} \mathcal{H}_0^{(s)} : \mathbf{z} \sim \mathcal{N}(\zeta \mathbf{u}, v^2 \boldsymbol{\Psi}) \\ \mathcal{H}_A^{(s)} : \mathbf{z} \sim \mathcal{N}(\zeta \mathbf{u} + \mathbf{C} \nabla, v^2 \boldsymbol{\Psi}) \end{cases}. \quad (54)$$

The derivation in appendix B leads to the self-consistency test

$$\text{Reject } \mathcal{H}_0^{(s)} \text{ if } T = \frac{\widehat{\nabla}^\top \mathbf{R}_{\widehat{\nabla}}^{-1} \widehat{\nabla}}{q \widehat{v}_A^2} > k_s, \quad (55)$$

with

$$\mathbf{R}_{\widehat{\nabla}} = (\mathbf{C}^\top \boldsymbol{\Psi}^{-1} \mathbf{P}_{\mathbf{u}}^\perp \mathbf{C})^{-1}, \quad (56)$$

$$\widehat{\nabla} = \mathbf{R}_{\widehat{\nabla}} \mathbf{C}^\top \boldsymbol{\Psi}^{-1} \mathbf{P}_{\mathbf{u}}^\perp \mathbf{z}, \quad (57)$$

$$\widehat{v}_A^2 = (M - q - 1)^{-1} \mathbf{z}^\top \boldsymbol{\Psi}^{-1} \mathbf{P}_{\mathbf{C}}^\perp \mathbf{P}_{\mathbf{u}}^\perp \mathbf{z}, \quad (58)$$

where q is the number of columns in \mathbf{C} , and we use the following idempotent matrices:

$$\mathbf{P}_{\mathbf{u}}^\perp = \mathbf{I} - \mathbf{u} (\mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \boldsymbol{\Psi}^{-1}, \quad (59)$$

$$\mathbf{P}_{\mathbf{C}}^\perp = \mathbf{I} - \mathbf{C} (\mathbf{C}^\top \boldsymbol{\Psi}^{-1} \mathbf{C})^{-1} \mathbf{C}^\top \boldsymbol{\Psi}^{-1}, \quad (60)$$

$$\underline{\mathbf{P}}_{\mathbf{u}}^\perp = \mathbf{I} - \mathbf{u} (\mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{P}_{\mathbf{C}}^\perp \mathbf{u})^{-1} \mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{P}_{\mathbf{C}}^\perp. \quad (61)$$

The test (55) is F -distributed

$$\begin{cases} \mathcal{H}_0^{(s)} : T \sim F(q, M - 1 - q, 0) \\ \mathcal{H}_A^{(s)} : T \sim F(q, M - 1 - q, \lambda) \end{cases}, \quad (62)$$

with non-centrality parameter

$$\lambda = \nabla^\top \mathbf{R}_{\widehat{\nabla}}^{-1} \nabla v^{-2}. \quad (63)$$

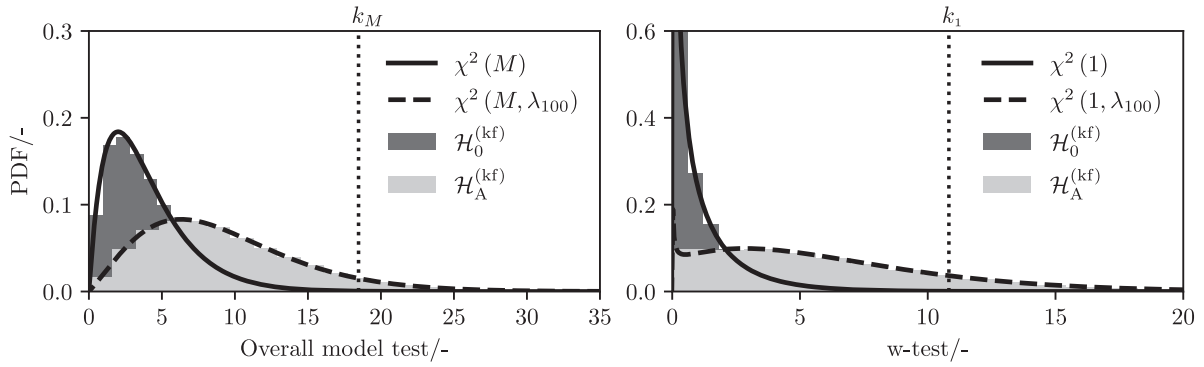


Figure 2. Distribution at $t_k = 100$ s of the overall model test and of the w -test applied to the Kalman filter residuals in the nominal case $\mathcal{H}_0^{(kf)}$ and in the faulty case $\mathcal{H}_A^{(kf)}$ ($\nabla = 10^{-11}$ s added on the first residual at every time step). The histograms show the distributions of the test from Monte Carlo simulations, while the lines show the theoretical distributions.

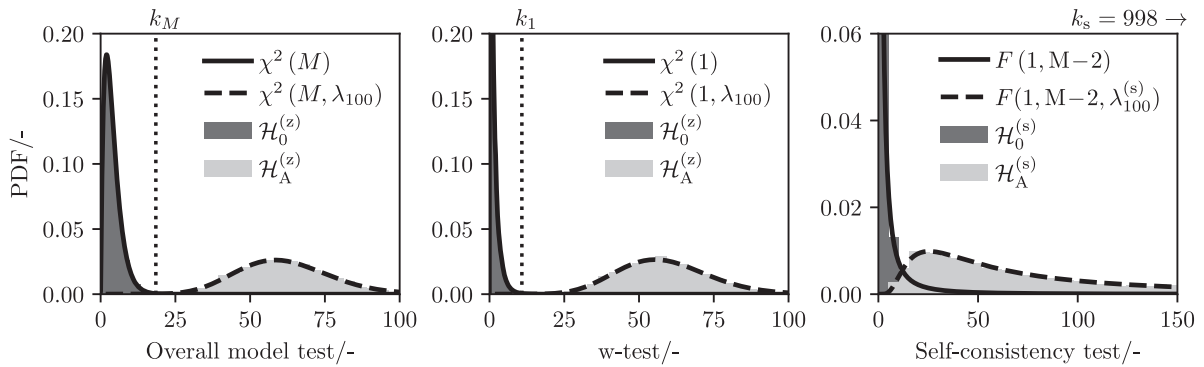


Figure 3. Left: distribution of the overall model test applied to the phase residuals. Center: distribution of the w -test applied to the phase residuals. Right: distribution of the self-consistency test applied to the phase measurements (threshold out of scale). All distributions are provided in the nominal and in the faulty case for $t_k = 500$ s, where a drift ∇ with slope $\frac{d\nabla}{dt} = 2.5 \times 10^{-12}$ s/s is added on the first measurement. The histograms show the distributions of the test from Monte Carlo simulations, while the lines show the theoretical distributions.




Figure 4. Evolution of the MDB for the w -test ($|\nabla_{MDB}|$) and self-consistency test ($|\nabla_{MDB}^{(s)}|$) on the phase measurements, as function of time. The MDB grows in time since the magnitude of the observables' covariance matrix increases. The dashed lines represent the magnitude of the injected fault ∇ . In this particular case, the self-consistency test (right plot) would not detect the slowly increasing drift.

The test threshold k_s can be found by inverting the expression for the P_{fa} :

$$P_{fa} = \int_{k_s}^{+\infty} F(\eta|q, M-1-q, \lambda) d\eta. \quad (64)$$

In the following we apply the self-consistency test in a manner similar to the w -test, thus we set $\mathbf{C} = \mathbf{c}_i$, $q = 1$, and we sweep the different measurements by changing the position

i of the non-zero entry in the vector \mathbf{c}_i . The MDB for the self-consistency test is

$$|\nabla_{MDB}^{(s)}| = \sqrt{\frac{\lambda}{R_{\hat{\nabla}}^{-1}} v^2}, \quad (65)$$

where in this case $R_{\hat{\nabla}}$ is a scalar.




Figure 5. Expected (P_{md}) and observed (\hat{P}_{md}) probability of missed detection for the tests applied to phase measurements. The theoretical values are found using the corresponding integrals defining the P_{md} , while we compute the observed values by counting the number of triggerings of the given test during the Monte Carlo simulation. As the injected fault increases faster than then the MDB (see figure 4), the P_{md} decreases over time.

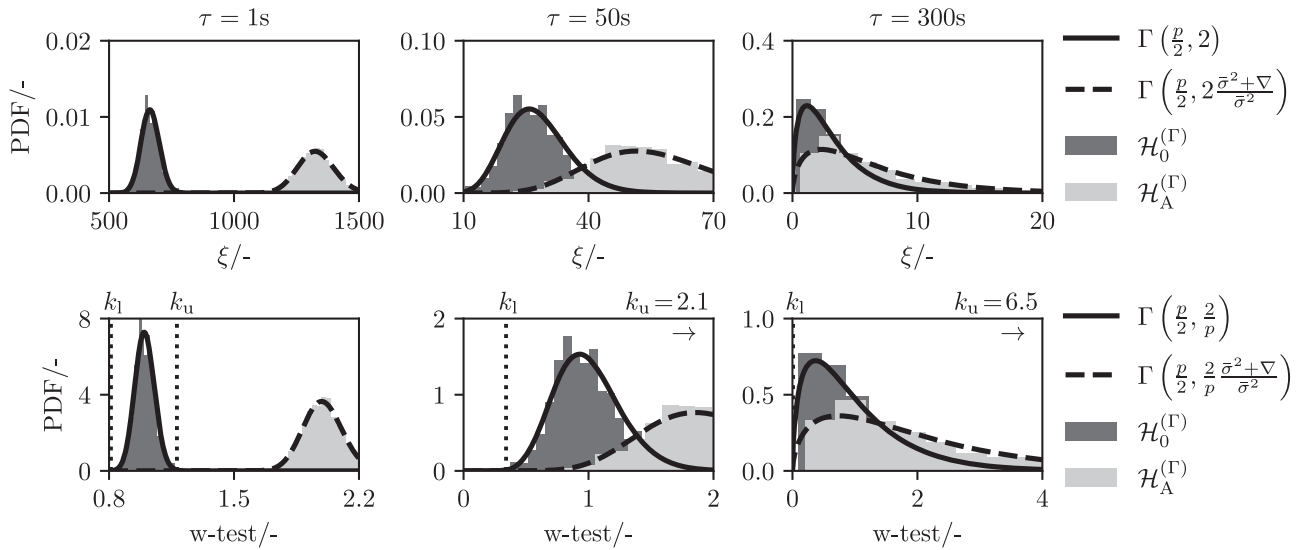


Figure 6. The upper plots show the distribution of the normalized DAVAR ξ for different sampling intervals, for a sliding window set to $W = 10^3$ s. The lower plots show the distributions of the corresponding w -test and the test thresholds. The histograms show the distributions resulting from Monte Carlo simulations, while the lines show the theoretical distributions. Please note the different scaling of the axes.

4.3. w -test for gamma-distributed observables

As previously shown, the normalized DAVAR $\xi_{k,\tau,W}$ follows a gamma distribution, for which a dedicated test is here developed. In this work we focus on the scalar version, leading to the w -test. Dropping all subscripts, the hypotheses are

$$\begin{cases} \mathcal{H}_0^{(\Gamma)} : \xi \sim \Gamma\left(\frac{p}{2}, 2\right) \\ \mathcal{H}_A^{(\Gamma)} : \xi \sim \Gamma\left(\frac{p}{2}, 2\frac{\bar{\sigma}^2 + \nabla}{\bar{\sigma}^2}\right) \end{cases} \quad (66)$$

The derivation in appendix C leads to

$$T = \frac{\hat{\sigma}^2}{\bar{\sigma}^2}, \quad (67)$$

which distributes according to a gamma distribution

$$\begin{cases} \mathcal{H}_0^{(\Gamma)} : T \sim \Gamma\left(\frac{p}{2}, \frac{2}{p}\right) \\ \mathcal{H}_A^{(\Gamma)} : T \sim \Gamma\left(\frac{p}{2}, \frac{2}{p} \frac{\bar{\sigma}^2 + \nabla}{\bar{\sigma}^2}\right) \end{cases} \quad (68)$$

Thus, the w -test on the DAVAR becomes

$$\text{Reject } \mathcal{H}_0^{(\Gamma)} \text{ if } T < k_l \text{ or } T > k_u, \quad (69)$$

where k_l is the lower threshold and k_u is the upper threshold, which can be computed from the expression of the P_{fa} :

$$P_{fa} = 1 - \int_{k_l}^{k_u} \Gamma\left(\eta \left| \frac{p}{2}, \frac{2}{p} \right.\right) d\eta. \quad (70)$$




Figure 7. Expected (P_{md}) and observed (\hat{P}_{md}) probability of missed detection for the w -test on the DAVAR, for different sampling intervals and fixed length of the sliding window $W = 10^3$ s.

The P_{md} can be found from:

$$P_{md} = \int_{k_1}^{k_u} \Gamma \left(\eta \left| \frac{p}{2}, \frac{2}{p} \frac{\bar{\sigma}^2}{\bar{\sigma}^2 + \nabla_{MDB}} \right. \right) d\eta. \quad (71)$$

The test (69) is double-sided. One can argue that a fault in the system can only cause a degradation of the observed stability and thus an increase in DAVAR, making the lower side of the test pointless. However, this test detects differences between observation and model, therefore we can use the lower side to detect mismodelling. If we select a model predicting an DAVAR higher than the real value, the test will fall below the lower threshold, and the triggering of the test can be accordingly marked.

4.4. Monte Carlo analysis

We run a set of Monte Carlo simulations of an ensemble comprising 5 cesium references to verify the theoretical distributions of the devised tests, and to compare the observed values of P_{fa} and P_{md} to the expected ones. This is only possible in simulation, since we need a large number of runs to evaluate the statistical behavior. Two simulations are carried out: one for the nominal case, and one for the faulty case, where a fault is injected in the first measurement. The selected model parameters of the cesium standards are

$$\begin{aligned} {}_i\sigma_1^2 &= 9 \times 10^{-24} \text{ s}, & {}_i\sigma_2^2 &= 0, & {}_i d &= 0, \\ {}_i U_M &= 10^{-23} \text{ s}, & {}_i R_M &= 0.57 \text{ s}^{-1}, \end{aligned} \quad (72)$$

where ${}_i U_M$ and ${}_i R_M$ parametrize the Markov process. As discussed in section 3.3, we also need a model without Markov processes to obtain a prediction model for the DAVAR, thus we choose

$${}_i\sigma_1^2 = 4.5 \times 10^{-23} \text{ s}, \quad {}_i\sigma_2^2 = 0, \quad {}_i d = 0, \quad {}_i U_M = 0. \quad (73)$$

In both cases we assume a value for the measurement noise variance of $\mathbf{R} = \mathbf{I}_M \times 10^{-25} \text{ s}^2$. These values are obtained by a laboratory characterization of the cesium frequency references used in this work. Long term measurements of the clocks are taken and their DAVAR computed, which is then fitted to




Figure 8. The solid gray line represents the nominal OAVAR $\bar{\sigma}^2(\tau)$, the dotted line shows the faulty OAVAR, which is shifted from the nominal by the injected fault ∇ , and the solid black lines represent the MDB ∇_{MDB} for the DAVAR test, using $W = 10^3$ s.

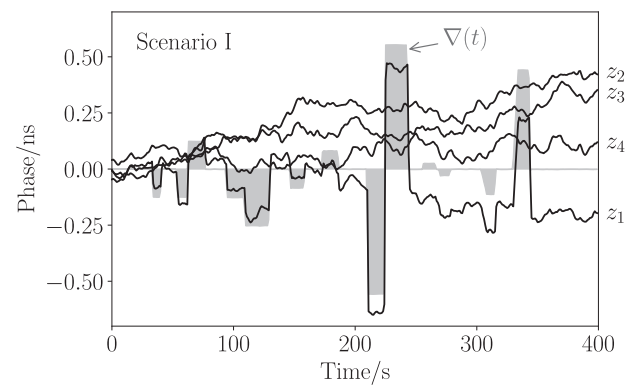


Figure 9. Phase measurements $\mathbf{z}(t)$ in the first scenario. The shaded area shows the injected pattern of phase steps $\nabla(t)$.

derive the model parameters. We set the probability of false alarm to

$$P_{fa} = 10^{-3}, \quad (74)$$

leading to the threshold value for the overall model test $k_M = 18.5$, for the w -test $k_1 = 10.8$, and for the self-consistency test $k_s = 998.5$. The thresholds k_1 and k_u for the DAVAR test depend on the length of sliding window W and on the sampling interval. A representation is given in the following experimental results. In the different simulations, the P_{md} can be computed using the threshold value and the magnitude of the injected drift.

4.4.1. Kalman filter residuals test. We run the Kalman filter in the nominal case and in the faulty case, where we inject a fault of magnitude $\nabla = 10^{-11} \text{ s}$ on the first residual, at all time steps. Thus we have $\nabla = \nabla \mathbf{c}_1$, where $\mathbf{c}_1 = [1 \ 0 \ \dots \ 0]^T$ is a column vector of M entries. We then check the distributions of the overall model test and of the w -test under the two hypotheses. In the faulty case we expect a time dependent non-centrality parameter

$$\lambda_k = \nabla^2 \mathbf{c}_1^T \left(\mathbf{\Omega}_k^{(kf)} \right)^{-1} \mathbf{c}_1. \quad (75)$$

The matrix $\mathbf{\Omega}_k^{(kf)}$ depends on time, but it converges after some iterations of the Kalman filter, as long as the measurements do not show abrupt changes. For $t_k = 100 \text{ s}$ we obtain $\lambda_{100} = 5.2$.

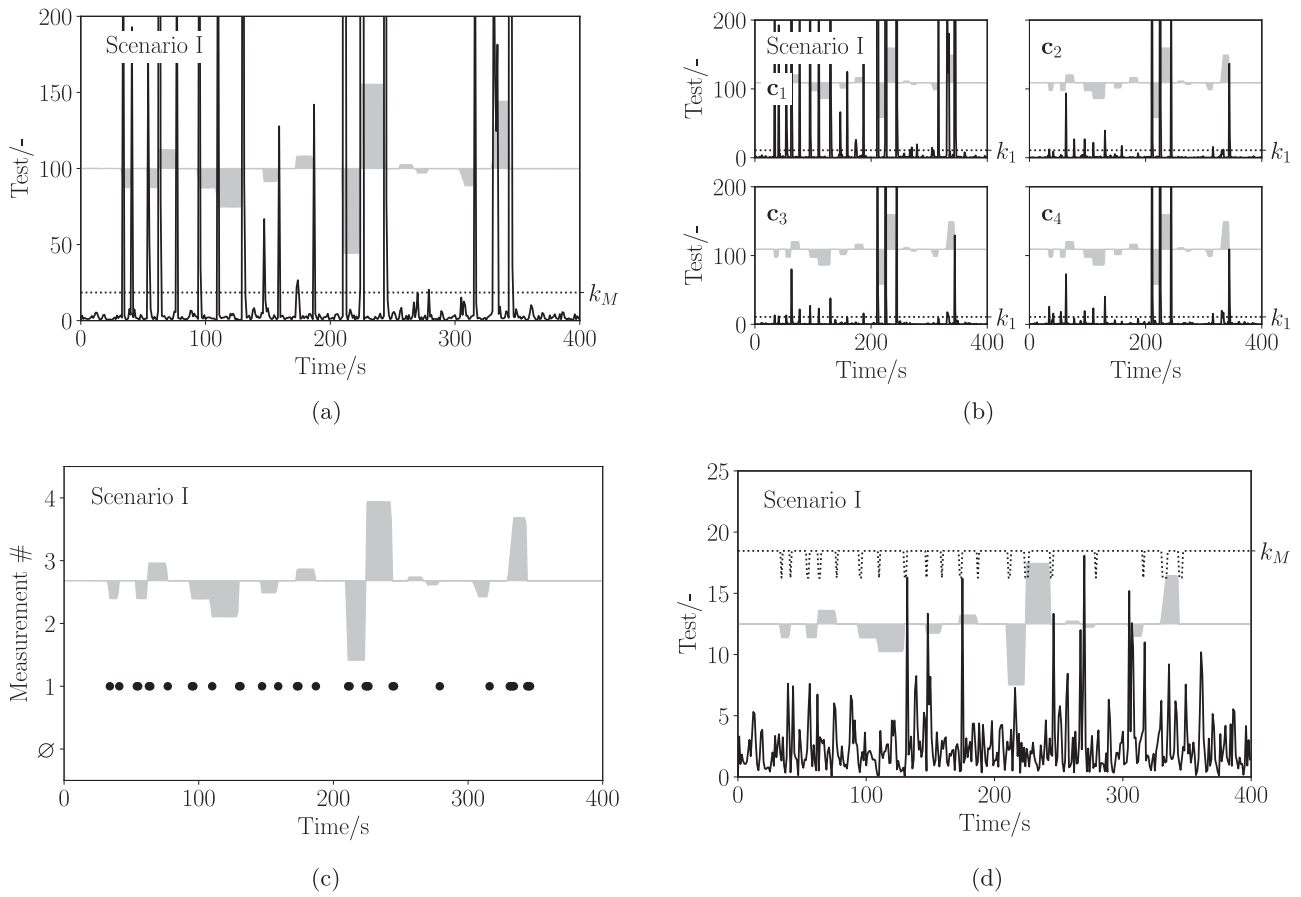


Figure 10. Scenario I: tests on the Kalman filter residuals in the presence of phase steps. The dotted line is the test threshold, the shaded curve shows the injected steps. (a) shows the overall model test, (b) the w -test selecting the single measurements c_1 to c_4 , (c) shows which measurement is most probably deemed as faulty (the symbol \emptyset is used for unsuccessful identification, when a fault is detected but not identified), while (d) shows the reduced test, when the faulty measurement is excluded from the observation vector. Removing an entry from the vector requires a recomputation of the threshold using a smaller value of degrees of freedom q , leading to the dips seen in the plot.




Figure 11. Scenario I: phase measurement residuals overall model test in the presence of phase steps. The dotted line is the test threshold, the shaded curve shows the injected steps.

From expressions (47) and (50) we expect a P_{md} for the overall model test of $P_{md}^{(om)}(\lambda_{100}) = 0.94$, while for the w -test of $P_{md}^{(wt)}(\lambda_{100}) = 0.84$.

Figure 2 shows the distributions of the overall model test and of the w -test in the nominal case $\mathcal{H}_0^{(kf)}$ and faulty case $\mathcal{H}_A^{(kf)}$, as well as the theoretical expected distributions. By counting the number of triggerings in the nominal case based on 10^5 simulations, we compute $\hat{P}_{fa} = 9.6 \times 10^{-4}$, which is in good

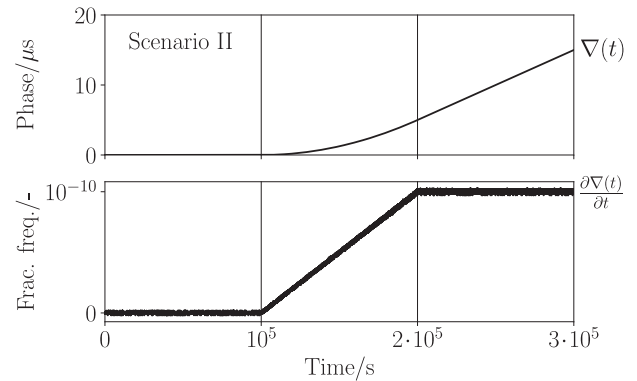


Figure 12. Scenario II: injected fault in terms of phase and fractional frequency differences. Three parts can be identified: initially no fault is added, then the injected phase grows quadratically, finally reaching a linear behavior.

agreement with the expected value of $P_{fa} = 10^{-3}$. If instead we count the times the test was not triggered in the faulty case, we obtain $\hat{P}_{md}^{(om)} = 0.94$ and $\hat{P}_{md}^{(wt)} = 0.84$ for the w -test. These values are in agreement with the theoretical derivation of the tests on the Kalman filter residuals.

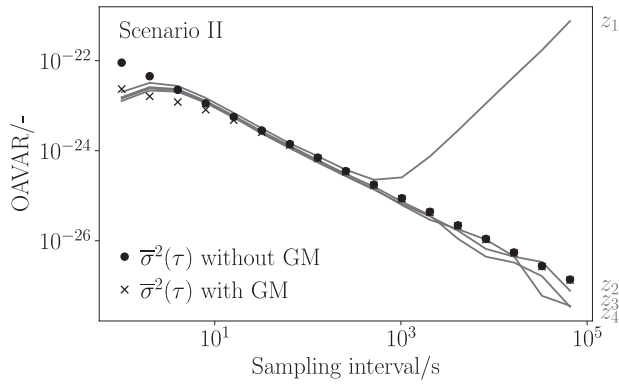


Figure 13. Measured OAVAR $\hat{\sigma}^2$ of the four channels computed on the entire batch of experimental data, and theoretical expected OAVAR $\bar{\sigma}^2$ for the two clock models, with and without Markov processes. The drift injected to the first measurement z_1 is visible at long sampling intervals.

4.4.2. Phase measurements test. In this case we inject a fault with linearly increasing magnitude with slope $\frac{d\nabla}{dt} = 2.5 \times 10^{-12}$ s/s in the first entry of the clock measurements vector \mathbf{z} . While the distribution of test in the null hypothesis is constant in time, in the faulty case the distribution is time dependent, since the injected fault is growing. Thus we verify the distribution for $t_k = 500$ s, leading to $\lambda_{500} = 57.2$ for the overall model test and w -test, and $\lambda_{500}^{(s)} = 53.6$ for the self-consistency test. Figure 3 shows the theoretical and observed distributions of the three tests in the nominal and faulty cases for $t_k = 500$ s, while in figure 4 the magnitude of the injected fault is compared to the MDB for the w -test and self-consistency test. Since the entries of the matrix $\Omega_k^{(z)}$ used to compute the non-centrality parameter are increasing, the MDB grows in time. A slowly increasing drift can be detected only when its magnitude becomes larger than the MDB. In figure 5 we plot a comparison of the expected and observed probabilities of missed detection for the overall model test, w -test and self-consistency test as function of time. The theoretical values are computed with (47) for the overall model test, (50) for the w -test and (71) for the self-consistency test. At the beginning, the injected fault is smaller than the MDB and remains undetected. The detection occurs later when the fault magnitude exceeds the MDB, which is reflected in the decrease in P_{md} .

4.4.3. DAVAR test. To verify the distributions of the tests on the DAVAR, we simulated an ensemble in the nominal case using the clock parameters (73), while in the faulty case we increased the first parameter of the second clock by a factor 3, so that ${}_2\sigma_1^2 = 1.35 \times 10^{-22}$ s. In this way the first entry in the measurement vector shows a fault, whose magnitude $\nabla(\tau)$ is computed by subtracting the theoretical DAVAR values of the first and second clocks. We employ a sliding window of 1000 s, and we compute the test for different values of sampling interval τ . The upper plots in figure 6 show the distribution of the normalized DAVAR (66) for an increasing sampling interval τ , while the lower plots show the distribution of the corresponding w -test developed using the gamma distribution (68).

Figure 7 shows the observed and expected P_{md} for the DAVAR w -test at increasing values of τ , while figure 8 compares the magnitude of the injected fault $\nabla(\tau)$ to the value of the MDB (∇_{MDB}). The error can be captured at shorter sampling intervals, while in the long term the statistical uncertainty in the determination of the DAVAR increases, thus weakening the power of detection.

5. Hardware setup and measurements

In our laboratory we run an ensemble of 5 cesium frequency references (Symmetricon 5071A), structured as in figure 1. With a SpectraDynamics high resolution offset generator we inject a fault of desired magnitude in the signal of the second cesium. Finally, a computer collects the measurements generated by a K+K FXE80 counter, and runs the ensembling and FDI algorithms in Matlab. We run three scenarios with different faults: in the first we inject a series of phase steps; the second scenario involves a phase and frequency drift; finally, a periodic phase component is injected in the third case. Here we set the integrity requirements as $P_{fa} = 10^{-3}$, $P_{md} = 10^{-6}$.

5.1. Scenario I: injection of phase steps

Figure 9 shows the phase measurements in the first scenario along with the injected phase steps. The phase pattern includes steps of different size and sign. For this scenario we only compute the tests on the Kalman filter residuals and on the phase measurements, but not on the DAVAR.

5.1.1. Tests on the Kalman filter residuals. Applying the overall model test on the Kalman filter residuals generates the results shown in figure 10. All the injected steps cause a spike in the test value, whose level depends on the size of the respective step. Only the smallest steps around $t_k = 270$ s do not trigger the test, but this is expected as their size is comparable with the MDB $|\nabla_{MDB}^{(kf)}| = 52.1$ ps. For locating the faulty measurement once the overall model test is triggered, we perform the w -test shown in figure 10(b). Although the test on the first channel (using \mathbf{c}_1) shows the highest values, the tests on the other channels trigger as well. This is due to the correlation between the differential measurements using a common clock, also encoded in the non-diagonal structure of the matrix Ω_k . Thus, we cannot unambiguously identify the faulty measurement yet. To perform identification we apply the following strategy:

- (a) Find which vector \mathbf{c}_j generates the highest test value;
- (b) Create a reduced observables vector ρ_R by removing the entry j corresponding to \mathbf{c}_j ;
- (c) Perform the overall model test on the reduced observables ρ_R :
 1. If the test is triggered, ρ_R still contains the faulty measurement. Restart from step (a) to remove one more entry from the reduced vector ρ_R ;
 2. If the test is *not* triggered, the faulty measurement was correctly removed: measurement j is considered as faulty;

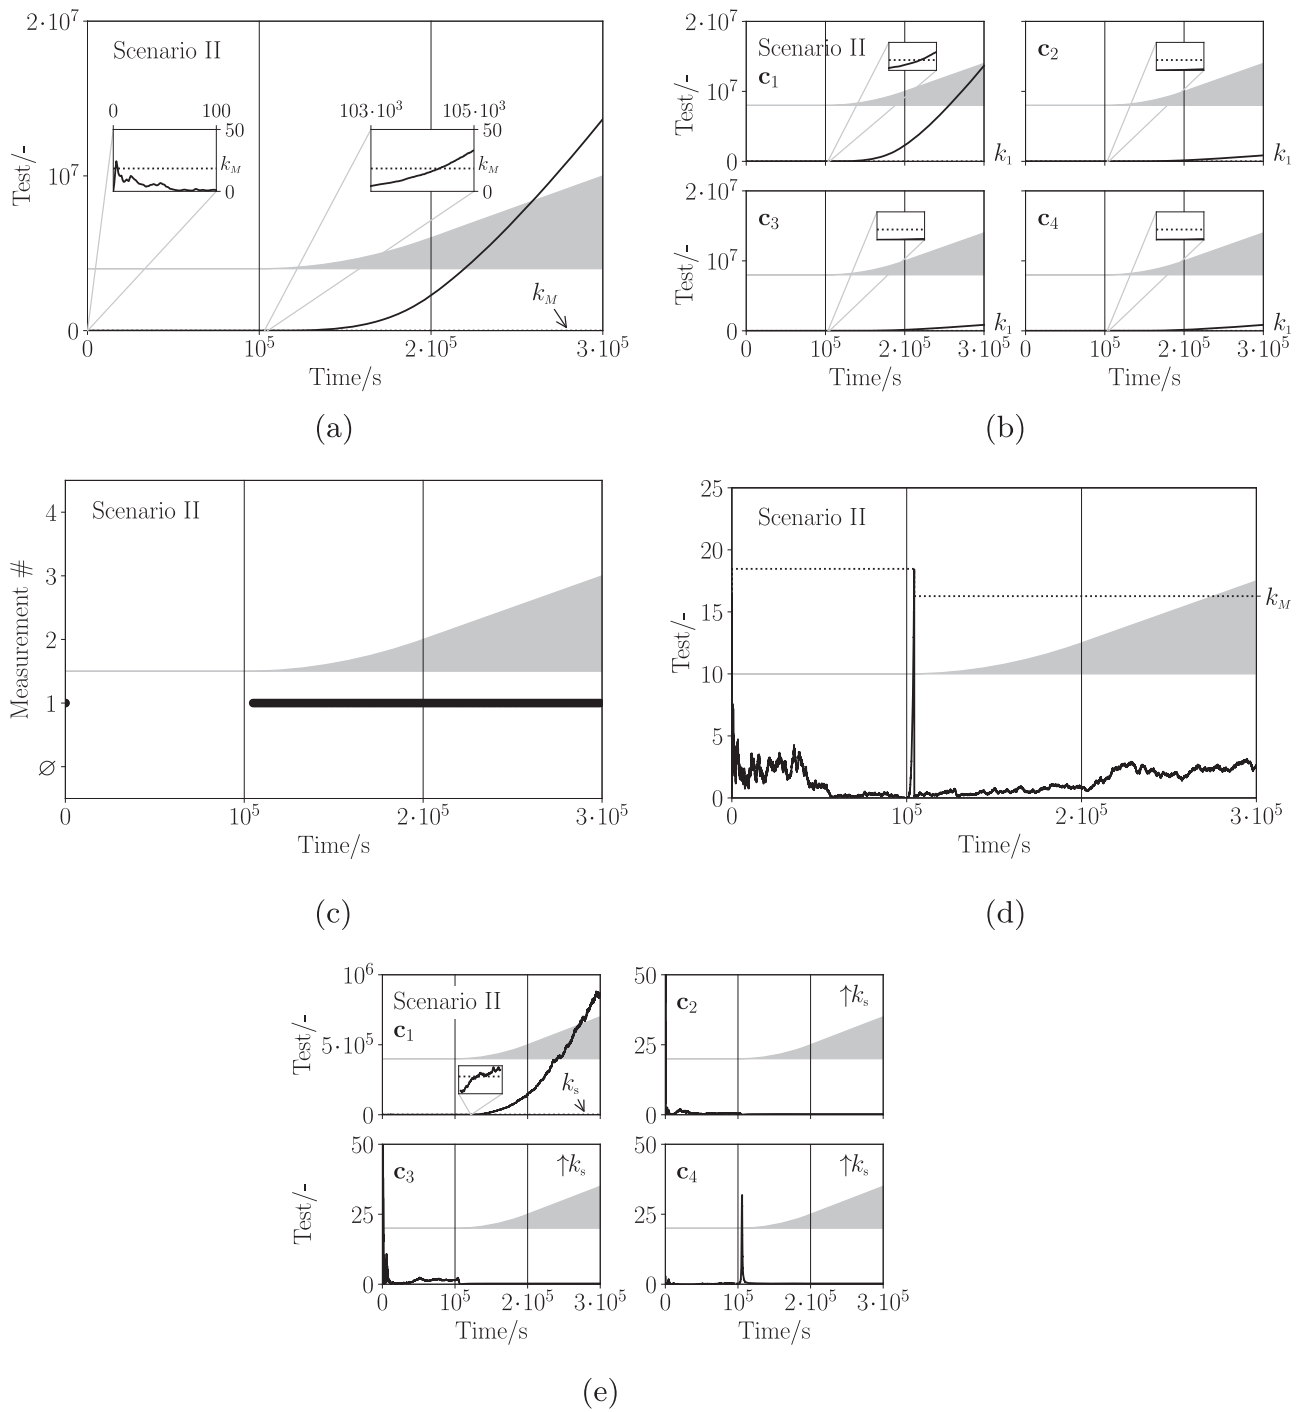


Figure 14. Scenario II: tests on the phase measurements in the presence of phase drift. The dotted line is the test threshold, the shaded curve shows the injected drift. (a) shows the overall model test, (b) the w -test on single measurements (all inset plots show the range 1.03×10^5 s to 1.05×10^5 s), (c) shows which measurement is most probably deemed as faulty (the symbol \emptyset is used for unsuccessful identification), while (d) shows the reduced test, when the faulty measurement is removed. Finally, (e) shows the self-consistency test on the phase measurements. Note the different scales of the y-axes in (e): the test threshold $k_s = 998$ is out of scale in all plot except that for c_1 .

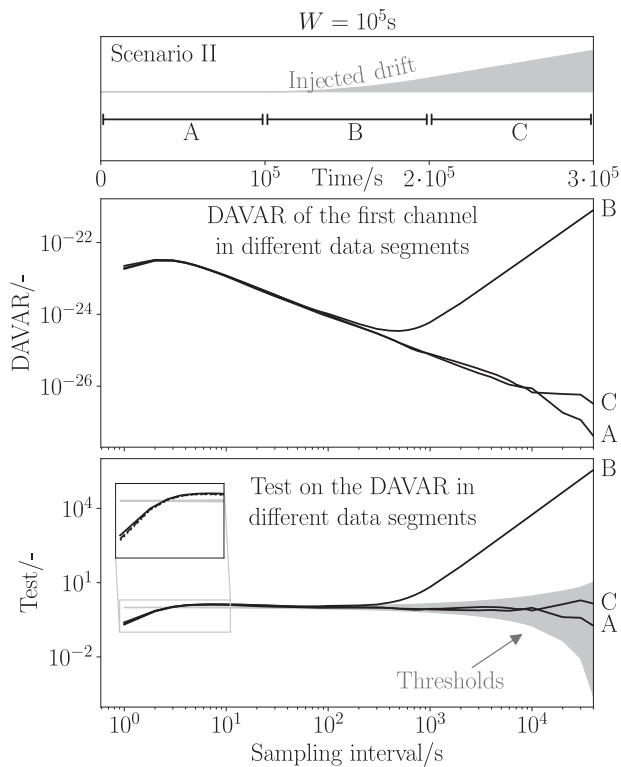


Figure 15. Scenario II: measured DAVAR $\hat{\sigma}_{k,r}^2$ of the faulty channel and corresponding test, computed on three data segments (A, B, and C) starting at different times. The upper plot shows the alignment of the data segments with the injected drift (shaded area): the first window is computed before the drift onset, the second during the part with quadratic drift, while the third in the part with linear drift. The central plot shows the DAVAR computed in the three segments. As expected, only the quadratic drift causes an increase in the observed values. The lower plot shows the test computed on the three data segments. The shaded area represents the interval between the lower and upper test thresholds, i.e. the region where $\mathcal{H}_0^{(r)}$ is accepted.

- (d) Repeat the process until a faulty measurement is identified. If all entries are removed from the observables vector ρ , the algorithm cannot identify a faulty measurement, and the identification process fails.

The result of this process represents the measurement which is most probably faulty, shown in figure 10(c). Here, we use the symbol \emptyset when the fault identification process is unsuccessful. As expected, the first measurement is identified as faulty in all the detected steps. Figure 10(d) shows the overall model test computed on the reduced vector of observables. The reduced test is never triggered, meaning that the algorithm successfully identifies the faulty measurement. Removing entries in the vector of observables requires recomputing the test threshold with a smaller number of degrees of freedom, leading to the dips shown in the threshold value.

5.1.2. Tests on the phase measurements. Although this scenario is not designed to assess the phase measurements tests, we show in figure 11 the performance of the overall model test. As expected, only the largest steps trigger the test, since

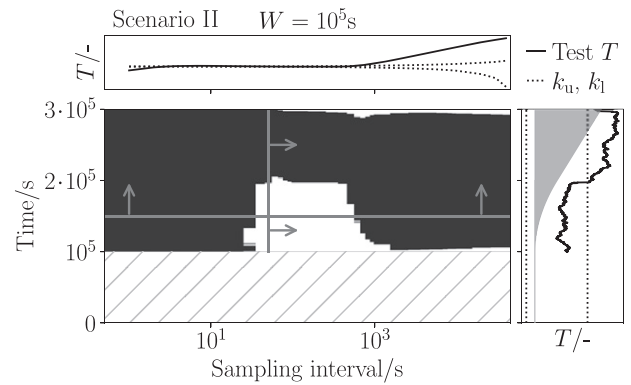


Figure 16. Result of applying the DAVAR test to the first channel in the second scenario, for a sliding window of length $W = 10^5$ s. The black shading represents the points in time and sampling interval where the test is triggered, while the hatched area covers missing data. The upper and right plots show the projections at a selected time instant and at a selected sampling interval. The dotted lines show the test thresholds.

the magnitude of the covariance matrix (and therefore the magnitude of the MDB) increases over time.

5.2. Scenario II: injection of a drift

In the second experiment we inject a phase and frequency drift, shown in figure 12. It consists of three parts: until $t_k = 10^5$ s no drift is added, while the frequency drift increases linearly for $10^5 \text{ s} < t_k < 2 \times 10^5 \text{ s}$, reaching a maximum value of 1 mHz, which is then kept constant until the end of the experiment. Correspondingly, the phase drift shows first a quadratic increase and finally a linear growth. Figure 13 shows the OAVAR computed on the entire measurement length and the OAVAR of the clock models: the dotted points describe the model without Markov process, while the crossed points correspond to the model with Markov process.

5.2.1. Tests on the phase measurements. Figure 14 shows the results of the test applied to the phase measurement residuals. The injected drift is successfully detected about 4320 s after its onset at $t_k = 10^5$ s. The overall model test in figure 14(a) shows a small increase at the beginning of the experiment. This is due to the nature of this test, which is more sensitive to small phase deviations during the initialization of the detection process. The w -test (figure 14(b)) correctly locates the drift on the first measurement, as seen in figure 14(c). The inset plots show that the w -tests computed using \mathbf{c}_2 , \mathbf{c}_3 , and \mathbf{c}_4 grow more slowly than with \mathbf{c}_1 , and are triggered on average 5000 s later. Figure 14(d) shows the overall model test on the reduced vector of phase residuals. The self-consistency test results in the values in figure 14(e). This test identifies the drift with a delay of 21 945 s using \mathbf{c}_1 . Although the detection with this test occurs with a significant delay, one should consider that the injected drift is relatively small (100 ps/s), thus a longer time is needed for the fault to exceed the detection threshold.

5.2.2. Test on the DAVAR. We expect the test on the DAVAR to detect the fault at long sampling intervals, since the drift




Figure 17. Scenario II: analysis of the detection delay for the DAVAR w -test. We define the detection delay as the difference between the start of drift injection at $t_k = 10^5$ s and the time of first detection. (a) shows the DAVAR w -test for $W = 10^4$ s on the first measurement as function of time. The different lines show the test for different sampling intervals. (b) shows the detection delay on the first measurement as function of the sampling interval, for different lengths of the sliding window W .

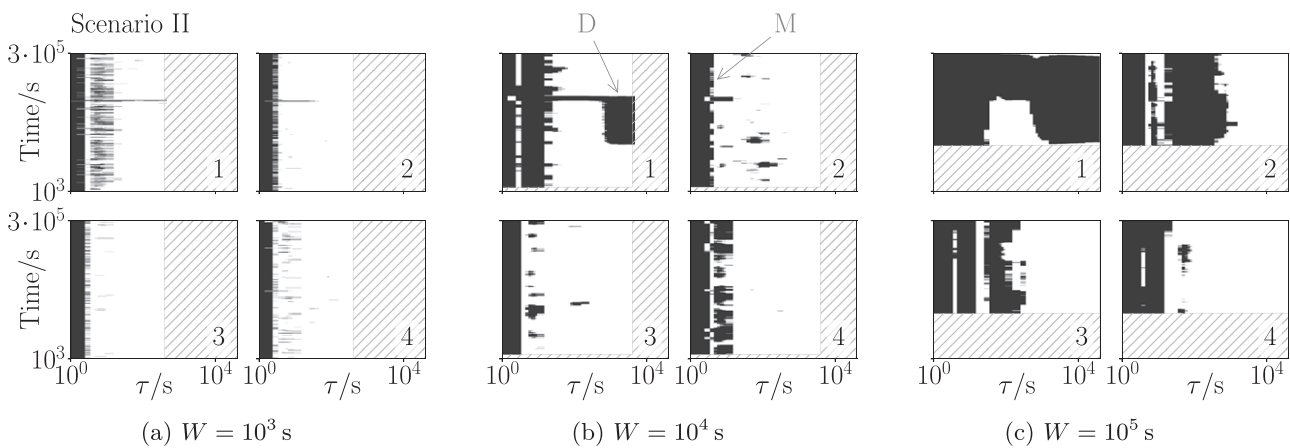


Figure 18. Scenario II: DAVAR w -test based on the gamma distribution, for different lengths of sliding window W , for the measurement channels 1 to 4. The black areas indicate fault detection, white means no detection. The hatched areas cover missing data: shorter windows provide values only at shorter sampling times, but longer windows can be evaluated only at later times. The area ‘D’ shows drift detection, while the area ‘M’ corresponds to triggering due to mismodelling.

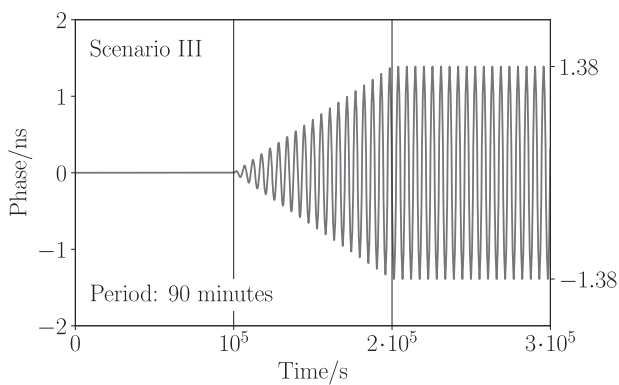


Figure 19. Scenario III: injected fault phase difference $\nabla(t)$. We can see three parts: initially no fault is added, then the amplitude grows linearly, finally reaching an oscillation of constant amplitude.

increases the measured DAVAR in the region above $\tau = 10^2$ s (see figure 13). However, the successful detection of the drift depends on the length of the sliding window and on the current time step. The drift can be seen in the DAVAR only if the sliding window is long enough, in this case at least $W > 10^3$ s.

Figure 15 shows the DAVAR and the w -test (67) for three windows of length $W = 10^5$ s starting at different time instants: before the drift onset, during the quadratic phase increase, and during the linear phase drift. The shaded area in the lower test plot represents the area where the null-hypothesis is accepted, delimited by the lower threshold k_l and the upper threshold k_u . The DAVAR computed on the window during the quadratic drift clearly shows that the fault can be detected in the region above $\tau = 200$ s. However, the fault remains undetected in the last data segment, since the DAVAR is invariant to linear phase drifts. In the lower plot of figure 15 we note that the test is triggered also for short sampling intervals. This triggering is due to mismodelling: in figure 13 a mismatch can be noted between the measured DAVAR and the values predicted by both clock models. In summary, we expect the DAVAR test to trigger due to the drift for $\tau > 200$ s, $W > 10^3$ s, and 10^5 s $< t_k < 2 \times 10^5$ s. The rejections due to mismodelling should occur for all window lengths, all times, and $\tau < 10$ s.

The value of the DAVAR test is function of time step t_k , sampling interval τ , length of the sliding window W , and measurement channel. To show these dependencies, we introduce




Figure 20. Scenario III: tests on the phase measurements in the presence of an oscillation in phase. The dotted line is the test threshold, the shaded curve shows the injected oscillation. (a) shows the overall model test, and (b) shows the self-consistency test on the phase measurements. As expected, neither test can successfully detect the injected oscillation.

the compact visualization shown in figure 16. Here we plot the results of the w -test on the DAVAR for the first channel, for one value of sliding window length W . This plot is obtained by reworking the three-dimensional surface of the DAVAR as function of time and sampling interval, corresponding to the waterfall plots in [13]. Since we are mainly interested in knowing whether the test value lies within the thresholds (accept $\mathcal{H}_0^{(T)}$) or outside (reject $\mathcal{H}_0^{(T)}$), we can forgo plotting the test value on the z -axis, and we use instead a color code for showing test acceptance or rejection. Thus, we reduce the 3D surface to a plot on the $(t_k - \tau)$ plane: the black pixels show the pairs (t_k, τ) where the test is triggered, while we leave a white pixel when accepting $\mathcal{H}_0^{(T)}$. More information can be added to the visualization: for example, by cutting the three-dimensional DAVAR curve at a given time and draw a projection onto the $(T - \tau)$ plane, we can visualize the test value as function of the sampling interval, along with the test thresholds (upper plot). Conversely, if we cut the surface at a given sampling interval, we project onto the $(T - t_k)$ plane on the right side the evolution of the test value as function of time (the shaded area represents the injected drift). The hatched area covers points where the test value is not available. The first value is available at $t_k = W$, since we must wait to collect enough points for computing the DAVAR on the window $[0, W]$. In

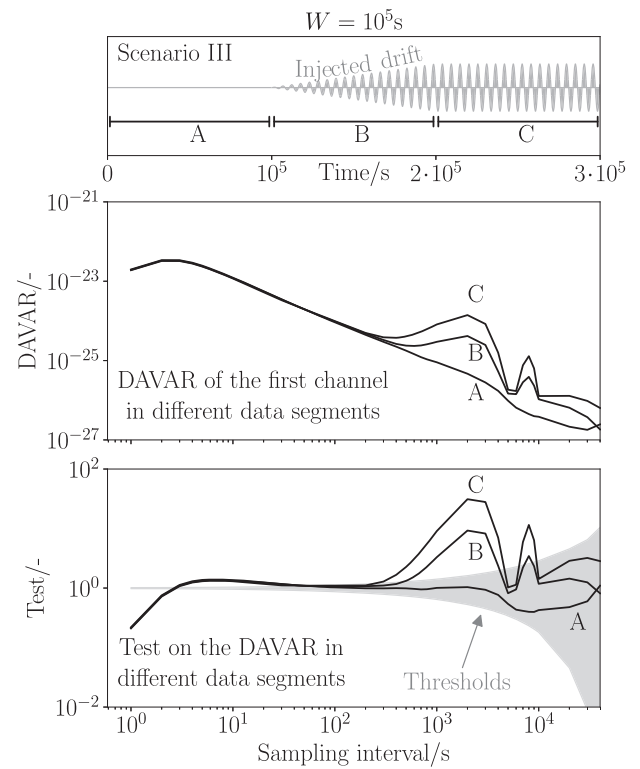


Figure 21. Scenario III: measured DAVAR $\hat{\sigma}^2$ of the faulty channel and corresponding test, computed on three data segments (A, B, and C) starting at different times. The upper plot shows the alignment of the segments with the injected drift (shaded area): the first window is computed before the drift onset, the second during the part with increasing amplitude, while the third in the part with constant amplitude. The central plot shows the DAVAR in the three batches, and the expected effects of the injected oscillation. The lower plot shows the test computed on the three data segments. The shaded area represents the interval between the lower and upper test thresholds.

this example we see that the test detects the mismodelling at short values of τ , and the drift at longer sampling intervals. The detection delay depends on the window length and on the sampling interval. In figure 17(a) the w -test on the first measurement computed for $W = 10^4$ s is plotted as function of time: the different lines show the test computed for different sampling intervals. Higher values of τ result in higher test values and the detection tends to occur earlier. This is visualized also in figure 17(b), where we plot the detection delay as function of the sampling interval for the first measurement. The minima occur for both window lengths at $\tau = 2 \times 10^3$ s, with detection delays of about 3000 s.

Figure 18 shows the result of the w -test on the DAVAR (68) based on the gamma distribution, for the 4 measurement channels z_i , and for different lengths of sliding window. We note how shorter sliding windows provide results at earlier time instants, but only for shorter sampling intervals. Conversely, longer windows yield values of DAVAR at longer sampling intervals, but only at later times. As expected, the test detects the drift only on the first measurement z_1 , when using sufficiently long sliding windows (area ‘D’). The drift

Both in the second and third scenarios we evaluate the test on the DAVAR, which successfully detects the injected drift and oscillation. However, the χ^2 -distribution of the DAVAR complicates the derivation of the GLRT. In this paper we derived a scalar version in section 4.3: although this test performs well in detecting drifts, it is applied to the single measurements and thus it does not take advantage of the structure of the clock ensemble. In future work we want to extend the statistical description of the DAVAR of the clock ensemble using a multivariate gamma distribution, and thus develop an exact overall model test on the DAVAR. The approximation to a Gaussian distribution allows us to develop an overall model test and a self-consistency test on the DAVAR. Although the approximation works well for short sampling intervals, this assumption degrades the test results at long term, which is exactly the region of interest when detecting drifts. Thus, longer sliding windows must be employed, which in turn delays the detection.

The detection of the mismodelling in the DAVAR test at short sampling intervals leads us to an interesting discussion point. On the one hand, the model-based tests clearly require a precise clock model to provide efficient fault detection without triggering too many false alarms. On the other hand, we can consider this as an opportunity to warn the user that wrong clock models may be in use. Our examples showed exactly this case: since the DAVAR tests trigger for short sampling intervals at all the times, on all the channels, and in different scenarios, we conclude that our clock model is imprecise in that region, which is confirmed by observing the difference between modelled and observed DAVAR in figure 13. Clearly, the detector cannot autonomously discern between faulty behavior or wrong model, but an experienced user shall be able to determine which one is the case.

An help in this direction comes from the self-consistency test. Since the test compares different measurements between each other, this method is not influenced by mismodelling. Therefore, by comparing the results of w -test and self-consistency test, we are able to determine which triggerings are due to faulty clock behaviors and which ones are due to mismodelling or other unmodelled effects. These effects include for instance daily temperature fluctuations, orbital effects, and in general everything equally affecting all the measurements and not predicted by the clock model. However, the self-consistency test works only on homogeneous clock ensembles, containing only clocks of the same type. Furthermore, the estimation of the expected value and of the magnitude of the covariance matrix requires at least three measurements. Finally, the derivation of the self-consistency test for χ^2 -distributed observables is not straightforward and requires further analyses.

The different tests developed in this work complement each other, and could be bundled together in an operational scenario to form an ensemble of powerful fault detectors.

Acknowledgments

This work was partially supported by the DLR Galileo Competence Center in the framework of the project ‘RPTF—Robust Precise Time Facility’.

Appendix A. Derivation of the GLRT for normally distributed observables

The ratio in (36) becomes

$$\text{Reject } \mathcal{H}_0 \text{ if } \frac{\max_{\nabla=0} f(\rho|\nabla)}{\max_{\nabla \neq 0} f(\rho|\nabla)} = \frac{f(\rho|0)}{f(\rho|\hat{\nabla})} < a, \quad (\text{A.1})$$

where the probability density function (PDF) is

$$f(\rho|\nabla) = (2\pi)^{-\frac{M}{2}} |\Omega|^{-\frac{1}{2}} e^{-\frac{1}{2}(\rho - C\nabla)^\top \Omega^{-1}(\rho - C\nabla)}, \quad (\text{A.2})$$

and

$$\hat{\nabla} = \arg \max_{\nabla \neq 0} f(\rho|\nabla). \quad (\text{A.3})$$

To find $\hat{\nabla}$ we take the derivative of the logarithm of (A.2) and set it to zero. This operation does not influence the position of the maximum. Thus, since Ω^{-1} is symmetric, it results

$$\begin{aligned} 0 &= -\frac{1}{2} \frac{\partial}{\partial \nabla} [(\rho - C\nabla)^\top \Omega^{-1}(\rho - C\nabla)]_{|\hat{\nabla}} \\ &= C^\top \Omega^{-1}(\rho - C\hat{\nabla}), \end{aligned} \quad (\text{A.4})$$

from which

$$\begin{aligned} \hat{\nabla} &= (C^\top \Omega^{-1} C)^{-1} C^\top \Omega^{-1} \rho \\ &= \mathbf{Q}_{\hat{\nabla}\hat{\nabla}} C^\top \Omega^{-1} \rho, \end{aligned} \quad (\text{A.5})$$

where we defined the matrix $\mathbf{Q}_{\hat{\nabla}\hat{\nabla}} = (C^\top \Omega^{-1} C)^{-1}$. By substituting (A.5) in (A.1), the ratio becomes

$$\exp\left(\frac{1}{2} \hat{\nabla}^\top \mathbf{Q}_{\hat{\nabla}\hat{\nabla}}^{-1} \hat{\nabla} - \rho^\top \Omega^{-1} C \hat{\nabla}\right). \quad (\text{A.6})$$

From (A.5) we have

$$\mathbf{Q}_{\hat{\nabla}\hat{\nabla}}^{-1} \hat{\nabla} = C^\top \Omega^{-1} \rho = (\rho^\top \Omega^{-1} C)^\top, \quad (\text{A.7})$$

which we can substitute in the second term of (A.6):

$$\rho^\top \Omega^{-1} C \hat{\nabla} = \hat{\nabla}^\top \mathbf{Q}_{\hat{\nabla}\hat{\nabla}}^{-1} \hat{\nabla}. \quad (\text{A.8})$$

Thus, the test (A.1) becomes

$$\text{Reject } \mathcal{H}_0 \text{ if } \exp\left(-\frac{1}{2} \hat{\nabla}^\top \mathbf{Q}_{\hat{\nabla}\hat{\nabla}}^{-1} \hat{\nabla}\right) < a, \quad (\text{A.9})$$

which can be equivalently rewritten by taking the logarithm of both sides, leading to

$$\text{Reject } \mathcal{H}_0 \text{ if } T = \hat{\nabla}^\top \mathbf{Q}_{\hat{\nabla}\hat{\nabla}}^{-1} \hat{\nabla} > 2 \log(a) = k. \quad (\text{A.10})$$

Appendix B. Derivation of the self-consistency test for normally distributed observables

The hypotheses (54) can be joined as

$$\mathbf{z} \sim \mathcal{N}(\mathbf{A}\boldsymbol{\eta}, v^2\boldsymbol{\Psi}), \tag{B.1}$$

with

$$\begin{cases} \mathcal{H}_0^{(s)} : \mathbf{A} = \mathbf{u}, \boldsymbol{\eta} = \zeta \\ \mathcal{H}_A^{(s)} : \mathbf{A} = [\mathbf{u} \ \mathbf{C}], \boldsymbol{\eta} = \begin{bmatrix} \zeta \\ \boldsymbol{\nabla} \end{bmatrix}, \end{cases} \tag{B.2}$$

and PDF

$$f(\mathbf{z}|\boldsymbol{\eta}, v^2) = \frac{(2\pi v^2)^{-\frac{M}{2}}}{\sqrt{|\boldsymbol{\Psi}|}} e^{-\frac{1}{2v^2}(\mathbf{z}-\mathbf{A}\boldsymbol{\eta})^\top \boldsymbol{\Psi}^{-1}(\mathbf{z}-\mathbf{A}\boldsymbol{\eta})}. \tag{B.3}$$

The maximum likelihood estimates are obtained by finding the maxima of the PDF:

$$\begin{aligned} \hat{\boldsymbol{\eta}} &= \arg \max_{\boldsymbol{\eta}} f(\mathbf{z}|\boldsymbol{\eta}, v^2) \\ &= (\mathbf{A}^\top \boldsymbol{\Psi}^{-1} \mathbf{A})^{-1} \mathbf{A}^\top \boldsymbol{\Psi}^{-1} \mathbf{z}, \end{aligned} \tag{B.4}$$

and

$$\begin{aligned} \hat{v}^2 &= \arg \max_{v^2} f(\mathbf{z}|\boldsymbol{\eta}, v^2) \\ &= M^{-1}(\mathbf{z} - \mathbf{A}\hat{\boldsymbol{\eta}})^\top \boldsymbol{\Psi}^{-1}(\mathbf{z} - \mathbf{A}\hat{\boldsymbol{\eta}}) \\ &= M^{-1} \hat{\mathbf{e}}^\top \boldsymbol{\Psi}^{-1} \hat{\mathbf{e}}, \end{aligned} \tag{B.5}$$

where we introduced the estimation error vector

$$\begin{aligned} \hat{\mathbf{e}} &= \mathbf{z} - \mathbf{A}\hat{\boldsymbol{\eta}} \\ &= [\mathbf{I}_M - \mathbf{A}(\mathbf{A}^\top \boldsymbol{\Psi}^{-1} \mathbf{A})^{-1} \mathbf{A}^\top \boldsymbol{\Psi}^{-1}] \mathbf{z} \\ &= (\mathbf{I}_M - \mathbf{P}_A) \mathbf{z}, \end{aligned} \tag{B.6}$$

and $\mathbf{P}_A = \mathbf{A}(\mathbf{A}^\top \boldsymbol{\Psi}^{-1} \mathbf{A})^{-1} \mathbf{A}^\top \boldsymbol{\Psi}^{-1}$ is the idempotent projection matrix associated with \mathbf{A} . Under the two hypotheses we have

$$\hat{\mathbf{e}}_0 = \mathbf{z} - \mathbf{u}\hat{\zeta}_0, \tag{B.7}$$

$$\hat{\mathbf{e}}_A = \mathbf{z} - \mathbf{u}\hat{\zeta}_A - \mathbf{C}\hat{\boldsymbol{\nabla}}. \tag{B.8}$$

The estimate (B.4) is unbiased

$$\mathcal{E}(\boldsymbol{\eta}) = (\mathbf{A}^\top \boldsymbol{\Psi}^{-1} \mathbf{A})^{-1} \mathbf{A}^\top \boldsymbol{\Psi}^{-1} \mathcal{E}(\mathbf{z}) = \boldsymbol{\eta}, \tag{B.9}$$

since $\mathcal{E}(\mathbf{z}) = \mathbf{A}\boldsymbol{\eta}$, where $\mathcal{E}(\bullet)$ is the expectation operator. However, (B.5) is biased:

$$\begin{aligned} \mathcal{E}(v^2) &= M^{-1} \mathcal{E}(\hat{\mathbf{e}}^\top \boldsymbol{\Psi}^{-1} \hat{\mathbf{e}}) \\ &= M^{-1} \text{tr}[\mathcal{E}(\hat{\mathbf{e}}\hat{\mathbf{e}}^\top) \boldsymbol{\Psi}^{-1}]. \end{aligned} \tag{B.10}$$

We observe that the covariance of $\hat{\mathbf{e}}$ is

$$\mathbf{Q}_{\hat{\mathbf{e}}\hat{\mathbf{e}}} = \mathcal{E}(\hat{\mathbf{e}}\hat{\mathbf{e}}^\top) - \mathcal{E}(\hat{\mathbf{e}})\mathcal{E}(\hat{\mathbf{e}}) = \mathcal{E}(\hat{\mathbf{e}}\hat{\mathbf{e}}^\top), \tag{B.11}$$

since

$$\mathcal{E}(\hat{\mathbf{e}}) = \mathcal{E}(\mathbf{z}) - \mathbf{A}\boldsymbol{\eta} = \mathbf{0}. \tag{B.12}$$

By introducing (B.6) and exploiting the properties of idempotent matrices we obtain

$$\begin{aligned} \mathbf{Q}_{\hat{\mathbf{e}}\hat{\mathbf{e}}} &= (\mathbf{I}_M - \mathbf{P}_A) \mathcal{E}(\mathbf{z}\mathbf{z}^\top) (\mathbf{I}_M - \mathbf{P}_A)^\top \\ &= v^2 (\mathbf{I}_M - \mathbf{P}_A) \boldsymbol{\Psi}. \end{aligned} \tag{B.13}$$

Substitution of (B.11) and (B.13) in (B.10) yields

$$\begin{aligned} \mathcal{E}(v^2) &= v^2 M^{-1} \text{tr}(\mathbf{I}_M - \mathbf{P}_A) \\ &= v^2 M^{-1} (M - n), \end{aligned} \tag{B.14}$$

where n is the number of column in \mathbf{A} : $n = 1$ under the null hypothesis and $n = q + 1$ under the alternative hypothesis. The estimate (B.14) is biased, so we can use a corrected estimator:

$$\hat{v}^2 = \hat{v}^2 \frac{M}{M - n}. \tag{B.15}$$

The estimates (B.4) and (B.5) can be now rewritten for the null hypothesis \mathcal{H}_0

$$\begin{aligned} \hat{\zeta}_0 &= (\mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{z} \\ \hat{v}_0^2 &= (M - 1)^{-1} \hat{\mathbf{e}}_0^\top \boldsymbol{\Psi}^{-1} \hat{\mathbf{e}}_0, \end{aligned} \tag{B.16}$$

and for the alternative hypothesis \mathcal{H}_A

$$\begin{aligned} \begin{bmatrix} \hat{\zeta}_A \\ \hat{\boldsymbol{\nabla}} \end{bmatrix} &= \begin{bmatrix} \mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{u} & \mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{C} \\ \mathbf{C}^\top \boldsymbol{\Psi}^{-1} \mathbf{u} & \mathbf{C}^\top \boldsymbol{\Psi}^{-1} \mathbf{C} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{u}^\top \boldsymbol{\Psi}^{-1} \\ \mathbf{C}^\top \boldsymbol{\Psi}^{-1} \end{bmatrix} \mathbf{z} \\ \hat{v}_A^2 &= (M - q - 1)^{-1} \hat{\mathbf{e}}_A^\top \boldsymbol{\Psi}^{-1} \hat{\mathbf{e}}_A. \end{aligned} \tag{B.17}$$

Now that the maximum likelihood estimates are available, we can build the test. The GLRT becomes

$$\begin{aligned} \frac{f(\mathbf{z}|\hat{\zeta}_0, \hat{v}_0^2)}{f(\mathbf{z}|\hat{\zeta}_A, \hat{\boldsymbol{\nabla}}, \hat{v}_A^2)} &= \frac{\hat{v}_A^M \exp\left(-\frac{1}{2\hat{v}_0^2} \hat{\mathbf{e}}_0^\top \boldsymbol{\Psi}^{-1} \hat{\mathbf{e}}_0\right)}{\hat{v}_0^M \exp\left(-\frac{1}{2\hat{v}_A^2} \hat{\mathbf{e}}_A^\top \boldsymbol{\Psi}^{-1} \hat{\mathbf{e}}_A\right)} \\ &= \left(\frac{\hat{v}_A^2}{\hat{v}_0^2}\right)^{\frac{M}{2}} e^{-\frac{q}{2}}. \end{aligned} \tag{B.18}$$

The test becomes

$$\text{Reject } \mathcal{H}_0^{(s)} \text{ if } \frac{\hat{v}_0^2}{\hat{v}_A^2} > \left(\alpha e^{\frac{q}{2}}\right)^{-\frac{2}{M}} = \bar{\alpha}, \tag{B.19}$$

or equivalently:

$$\text{Reject } \mathcal{H}_0^{(s)} \text{ if } \frac{M - q - 1}{M - 1} \frac{\hat{\mathbf{e}}_0^\top \boldsymbol{\Psi}^{-1} \hat{\mathbf{e}}_0}{\hat{\mathbf{e}}_A^\top \boldsymbol{\Psi}^{-1} \hat{\mathbf{e}}_A} > \bar{\alpha}. \tag{B.20}$$

Please note that the term $M - q - 1$ is always positive since the maximum value of q is $M - 2$. We want now to rewrite the test as function of $\hat{\boldsymbol{\nabla}}$. By multiplying (B.17) by a matrix \mathbf{B}

$$\mathbf{B} \begin{bmatrix} \mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{u} & \mathbf{u}^\top \boldsymbol{\Psi}^{-1} \mathbf{C} \\ \mathbf{C}^\top \boldsymbol{\Psi}^{-1} \mathbf{u} & \mathbf{C}^\top \boldsymbol{\Psi}^{-1} \mathbf{C} \end{bmatrix} \begin{bmatrix} \hat{\zeta}_A \\ \hat{\boldsymbol{\nabla}} \end{bmatrix} = \mathbf{B} \begin{bmatrix} \mathbf{u}^\top \boldsymbol{\Psi}^{-1} \\ \mathbf{C}^\top \boldsymbol{\Psi}^{-1} \end{bmatrix} \mathbf{z}, \tag{B.21}$$

with

$$\mathbf{B} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{C}^\top \Psi^{-1} \mathbf{u} (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} & \mathbf{I} \end{bmatrix}, \quad (\text{B.22})$$

we obtain

$$\begin{bmatrix} \mathbf{u}^\top \Psi^{-1} \mathbf{u} & \mathbf{u}^\top \Psi^{-1} \mathbf{C} \\ \mathbf{0} & \mathbf{C}^\top \Psi^{-1} \mathbf{P}_u^\perp \mathbf{C} \end{bmatrix} \begin{bmatrix} \hat{\zeta}_A \\ \hat{\nabla} \end{bmatrix} = \begin{bmatrix} \mathbf{u}^\top \Psi^{-1} \\ \mathbf{C}^\top \Psi^{-1} \mathbf{P}_u^\perp \end{bmatrix} \mathbf{z}, \quad (\text{B.23})$$

where

$$\mathbf{P}_u^\perp = \mathbf{I} - \mathbf{P}_u = \mathbf{I} - \mathbf{u} (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1}. \quad (\text{B.24})$$

From the second equation in (B.23) we obtain an estimate of the fault magnitude

$$\begin{aligned} \hat{\nabla} &= (\mathbf{C}^\top \Psi^{-1} \mathbf{P}_u^\perp \mathbf{C})^{-1} \mathbf{C}^\top \Psi^{-1} \mathbf{P}_u^\perp \mathbf{z} \\ &= \mathbf{R}_{\hat{\nabla}} \mathbf{C}^\top \Psi^{-1} \mathbf{P}_u^\perp \mathbf{z}, \end{aligned} \quad (\text{B.25})$$

where we defined the matrix

$$\mathbf{R}_{\hat{\nabla}} = (\mathbf{C}^\top \Psi^{-1} \mathbf{P}_u^\perp \mathbf{C})^{-1}. \quad (\text{B.26})$$

We can rewrite the estimation error vectors under the two hypotheses

$$\begin{cases} \mathcal{H}_0^{(s)} : \hat{\mathbf{e}}_0 = \mathbf{P}_u^\perp \mathbf{z} \\ \mathcal{H}_A^{(s)} : \hat{\mathbf{e}}_A = \mathbf{P}_C^\perp \mathbf{P}_u^\perp \mathbf{z} \end{cases}, \quad (\text{B.27})$$

where we used following idempotent matrices

$$\mathbf{P}_C^\perp = \mathbf{I} - \mathbf{C} (\mathbf{C}^\top \Psi^{-1} \mathbf{C})^{-1} \mathbf{C}^\top \Psi^{-1}, \quad (\text{B.28})$$

$$\mathbf{P}_u^\perp = \mathbf{I} - \mathbf{u} (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1}. \quad (\text{B.29})$$

From (B.16) we have

$$(\mathbf{u}^\top \Psi^{-1} \mathbf{u}) \hat{\zeta}_0 = \mathbf{u}^\top \Psi^{-1} \mathbf{z}, \quad (\text{B.30})$$

which can be inserted in the first equation of (B.23):

$$\hat{\zeta}_A - \hat{\zeta}_0 = -(\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1} \mathbf{C} \hat{\nabla}. \quad (\text{B.31})$$

With this, we can express the difference of the residuals:

$$\begin{aligned} \hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A &= \mathbf{u} (\hat{\zeta}_A - \hat{\zeta}_0) + \mathbf{C} \hat{\nabla} \\ &= [\mathbf{I} - \mathbf{u} (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1}] \mathbf{C} \hat{\nabla} \\ &= \mathbf{P}_u^\perp \mathbf{C} \hat{\nabla}. \end{aligned} \quad (\text{B.32})$$

The ratio in (B.20) can be reworked as

$$\frac{\hat{\mathbf{e}}_0^\top \Psi^{-1} \hat{\mathbf{e}}_0}{\hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} = \frac{\hat{\mathbf{e}}_0^\top \Psi^{-1} \hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A}{\hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} + 1. \quad (\text{B.33})$$

The numerator is equal to

$$(\hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A)^\top \Psi^{-1} (\hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A) + 2(\hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A)^\top \Psi^{-1} \hat{\mathbf{e}}_A, \quad (\text{B.34})$$

but the term $(\hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A)^\top \Psi^{-1} \hat{\mathbf{e}}_A$ is zero. In fact, by using the property of symmetry and idempotence of \mathbf{P}_u^\perp and $\mathbf{P}_C^\perp \mathbf{P}_u^\perp$ we have

$$\begin{aligned} (\hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A)^\top \Psi^{-1} \hat{\mathbf{e}}_A &= \mathbf{z}^\top (\mathbf{P}_u^\perp - \mathbf{P}_C^\perp \mathbf{P}_u^\perp) \Psi^{-1} \mathbf{P}_C^\perp \mathbf{P}_u^\perp \mathbf{z} \\ &= \mathbf{z}^\top (\mathbf{P}_u^\perp - \mathbf{I}) \mathbf{P}_C^\perp \mathbf{P}_u^\perp \Psi^{-1} \mathbf{z} \\ &= \mathbf{z}^\top \mathbf{P}_C^\perp \mathbf{P}_u^\perp \mathbf{P}_u \Psi^{-1} \mathbf{z}. \end{aligned} \quad (\text{B.35})$$

We can show that the product $\mathbf{P}_u^\perp \mathbf{P}_u$ is zero

$$\begin{aligned} \mathbf{P}_u^\perp \mathbf{P}_u &= (\mathbf{I} - \mathbf{u} (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1}) \mathbf{u} \\ &\quad \cdot (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1} \\ &= (\mathbf{u} - \mathbf{u} (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1} \mathbf{u}) (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1} \\ &= (\mathbf{u} - \mathbf{u}) (\mathbf{u}^\top \Psi^{-1} \mathbf{u})^{-1} \mathbf{u}^\top \Psi^{-1} \\ &= 0. \end{aligned} \quad (\text{B.36})$$

Thus, the ratio becomes

$$\frac{\hat{\mathbf{e}}_0^\top \Psi^{-1} \hat{\mathbf{e}}_0}{\hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} = \frac{(\hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A)^\top \Psi^{-1} (\hat{\mathbf{e}}_0 - \hat{\mathbf{e}}_A)}{\hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} + 1. \quad (\text{B.37})$$

By inserting (B.32) we obtain

$$\begin{aligned} \frac{\hat{\mathbf{e}}_0^\top \Psi^{-1} \hat{\mathbf{e}}_0}{\hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} &= \frac{\hat{\nabla}^\top \mathbf{C}^\top \Psi^{-1} \mathbf{P}_u^\perp \mathbf{C} \hat{\nabla}}{\hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} + 1 \\ &= \frac{\hat{\nabla}^\top \mathbf{R}_{\hat{\nabla}}^{-1} \hat{\nabla}}{\hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} + 1. \end{aligned} \quad (\text{B.38})$$

Thus the test (B.20) can be rewritten as

$$\text{Reject } \mathcal{H}_0^{(s)} \text{ if } \frac{\hat{\nabla}^\top \mathbf{R}_{\hat{\nabla}}^{-1} \hat{\nabla}}{\hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} > \beta, \quad (\text{B.39})$$

with

$$\beta = \frac{M - 1}{M - q - 1} \bar{\alpha} - 1. \quad (\text{B.40})$$

The numerator in (B.39) distributes according to a χ^2 -distribution with q degrees of freedom, central under $\mathcal{H}_0^{(s)}$ and non-central under $\mathcal{H}_A^{(s)}$. The non-centrality parameter is $\lambda = \hat{\nabla}^\top \mathbf{R}_{\hat{\nabla}}^{-1} \hat{\nabla}$. The denominator is also χ^2 -distributed, central under both hypotheses and with $M - q - 1$ degrees of freedom.

If (B.39) is rewritten as

$$\text{Reject } \mathcal{H}_0^{(s)} \text{ if } \frac{q^{-1} \hat{\nabla}^\top \mathbf{R}_{\hat{\nabla}}^{-1} \hat{\nabla}}{(M - q - 1)^{-1} \hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A} > k_s, \quad (\text{B.41})$$

with

$$k_s = \frac{M - q - 1}{q} \beta, \quad (\text{B.42})$$

we can express the distribution of the ratio (B.41) with a F -distribution, which corresponds to the ratio of two independent χ^2 -distributed variables:

$$F(n_1, n_2, \lambda) = \frac{n_1^{-1} \cdot \chi^2(n_1, \lambda)}{n_2^{-1} \cdot \chi^2(n_2, 0)}. \tag{B.43}$$

From (B.17) we note that the denominator in (B.41) equals to \hat{v}_A^2 , so that the self-consistency test finally becomes

$$\text{Reject } \mathcal{H}_0^{(s)} \text{ if } T = \frac{\hat{\nabla}^\top \mathbf{R}_{\hat{\nabla}}^{-1} \hat{\nabla}}{q \cdot \hat{v}_A^2} > k_s, \tag{B.44}$$

which distributes according to

$$T \sim F(q, M - q - 1, \lambda). \tag{B.45}$$

The terms in (B.44) can be computed with (B.25) and

$$\begin{aligned} \hat{v}_A^2 &= (M - q - 1)^{-1} \hat{\mathbf{e}}_A^\top \Psi^{-1} \hat{\mathbf{e}}_A \\ &= (M - q - 1)^{-1} \mathbf{z}^\top \Psi^{-1} \mathbf{P}_C^\perp \mathbf{P}_u^\perp \mathbf{z}. \end{aligned} \tag{B.46}$$

Appendix C. Derivation of the w -test for gamma-distributed observables

To obtain an expression for the w -test for gamma-distributed observables in section 4.3, we need to find the maxima in the GLRT:

$$\text{Reject } \mathcal{H}_0^{(\Gamma)} \text{ if } \bar{T} = \frac{\max_{\nabla=0} f(\xi|\nabla)}{\max_{\nabla \neq 0} f(\xi|\nabla)} < a, \tag{C.1}$$

where

$$\begin{aligned} f(\xi|\nabla) &= \frac{1}{2} \left(\frac{\hat{\sigma}^2 \bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} \right)^{\frac{p}{2}-1} \frac{1}{\Gamma(\frac{p}{2})} \left(\frac{\bar{\sigma}^2}{\bar{\sigma}^2 + \nabla} \right)^{\frac{p}{2}} \\ &\cdot \exp \left(-\frac{\hat{\sigma}^2 \bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} \frac{\bar{\sigma}^2}{\bar{\sigma}^2 + \nabla} \right). \end{aligned} \tag{C.2}$$

The numerator is trivial:

$$\begin{aligned} \max_{\nabla=0} f(\xi|\nabla) &= f(\xi|0) \\ &= \frac{1}{2} \left(\frac{\hat{\sigma}^2 \bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} \right)^{\frac{p}{2}-1} \frac{1}{\Gamma(\frac{p}{2})} \exp \left(-\frac{\hat{\sigma}^2 \bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} \right). \end{aligned} \tag{C.3}$$

For the denominator we have

$$\max_{\nabla \neq 0} f(\xi|\nabla) = f(\xi|\hat{\nabla}), \tag{C.4}$$

with

$$\hat{\nabla} = \arg \max_{\nabla} f(\xi|\nabla). \tag{C.5}$$

To find $\hat{\nabla}$ we take the logarithm of the PDF (C.2), since this does not influence the location of the maximum:

$$L_\Gamma(\xi|\nabla) = \log [f(\xi|\nabla)]$$

$$\begin{aligned} &= \left(\frac{p}{2} - 1 \right) \log \left(\frac{\hat{\sigma}^2 \bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} \right) - \log \left[\Gamma \left(\frac{p}{2} \right) \right] \\ &+ \log \left(\frac{\bar{\sigma}^2}{\bar{\sigma}^2 + \nabla} \right)^{\frac{p}{2}} - \frac{\hat{\sigma}^2 \bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} \frac{\bar{\sigma}^2}{\bar{\sigma}^2 + \nabla}. \end{aligned} \tag{C.6}$$

By taking the derivative we obtain

$$\frac{\partial}{\partial \nabla} L_\Gamma = -\frac{p}{2} \frac{1}{\bar{\sigma}^2 + \nabla} + \frac{\hat{\sigma}^2 \bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} \frac{1}{(\bar{\sigma}^2 + \nabla)^2}. \tag{C.7}$$

If we impose $\frac{\partial}{\partial \nabla} L_\Gamma|_{\hat{\nabla}} = 0$, we can compute $\hat{\nabla}$:

$$\hat{\nabla} = \frac{2}{p} \frac{\hat{\sigma}^2 \bar{\sigma}^4}{\text{Var}(\hat{\sigma}^2)} = \hat{\sigma}^2 - \bar{\sigma}^2, \tag{C.8}$$

where we substituted the expression for the degrees of freedom

$$p = \frac{2\bar{\sigma}^4}{\text{Var}(\hat{\sigma}^2)}. \tag{C.9}$$

This point is indeed a maximum, since the second derivative computed for $\nabla = \hat{\nabla}$ is negative:

$$\frac{\partial^2}{\partial \nabla^2} L_\Gamma|_{\hat{\nabla}} = -\frac{\bar{\sigma}^4}{\hat{\sigma}^4 \text{Var}(\hat{\sigma}^2)} < 0. \tag{C.10}$$

Thus, the second maximum becomes

$$f(\xi|\hat{\nabla}) = \frac{1}{2\Gamma(\frac{p}{2})} \left(\frac{\hat{\sigma}^2 \bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} \right)^{\frac{p}{2}-1} \cdot \exp \left(-\frac{\bar{\sigma}^4}{\text{Var}(\hat{\sigma}^2)} \right), \tag{C.11}$$

so that the ratio (C.1) yields

$$\begin{aligned} \bar{T} &= \left(\frac{\hat{\sigma}^2}{\bar{\sigma}^2} \right)^{\frac{p}{2}} \exp \left[-\frac{\bar{\sigma}^2}{\text{Var}(\hat{\sigma}^2)} (\hat{\sigma}^2 - \bar{\sigma}^2) \right] \\ &= \left(\frac{\hat{\sigma}^2}{\bar{\sigma}^2} \right)^{\frac{p}{2}} \exp \left[-\frac{p}{2} \left(\frac{\hat{\sigma}^2}{\bar{\sigma}^2} - 1 \right) \right]. \end{aligned} \tag{C.12}$$

The test \bar{T} is a bell shaped function of the ratio $\frac{\hat{\sigma}^2}{\bar{\sigma}^2}$, as shown in figure C1, where we plot the function (C.12) for some arbitrary values of degrees of freedom p and threshold a .

We observe that

$$\bar{T} < a \Leftrightarrow \frac{\hat{\sigma}^2}{\bar{\sigma}^2} < k_l \vee \frac{\hat{\sigma}^2}{\bar{\sigma}^2} > k_u, \tag{C.13}$$

thus the test (C.1) can be equivalently rewritten as

$$\text{Reject } \mathcal{H}_0^{(\Gamma)} \text{ if } T < k_l \text{ or } T > k_u, \tag{C.14}$$

where now the test value is $T = \frac{\hat{\sigma}^2}{\bar{\sigma}^2}$.

The distribution of the test (C.14) can be derived from (31):

$$\xi = p \frac{\hat{\sigma}^2}{\bar{\sigma}^2} \sim \Gamma \left(\frac{p}{2}, 2 \frac{\bar{\sigma}^2 + \nabla}{\bar{\sigma}^2} \right). \tag{C.15}$$




Figure C1. Qualitative example of the function $\bar{T}\left(\frac{\hat{\sigma}^2}{\sigma^2}\right)$, for some arbitrary values of degrees of freedom p and threshold a . Setting $\bar{T} < a$ is equivalent to $\frac{\hat{\sigma}^2}{\sigma^2} < k_l \vee \frac{\hat{\sigma}^2}{\sigma^2} > k_u$.

Thus by applying the scaling property of the gamma distribution we finally obtain

$$T = \frac{1}{p} \left(p \frac{\hat{\sigma}^2}{\sigma^2} \right) \sim \Gamma \left(\frac{p}{2}, \frac{2}{p} \frac{\sigma^2 + \nabla}{\sigma^2} \right). \quad (\text{C.16})$$

ORCID iDs

Christian Trainotti  <https://orcid.org/0000-0001-5176-6100>
Gabriele Giorgi  <https://orcid.org/0000-0003-0552-9790>

References

- [1] Teunissen P J G 2009 *Testing Theory: An Introduction (Mathematical Geodesy and Positioning)* (VSSD: Delft Academic Press) <http://vssd.nl/hlf/a030.htm>
- [2] Teunissen P J G 2017 *Batch and Recursive Model Validation* (Berlin: Springer) pp 687–720
- [3] Jones R H and Tryon P V 1983 *J. Res. Natl. Bur. Stan.* **88** 17–24
- [4] Galleani L and Tavella P 2011 Using the Kalman filter to detect frequency jumps in atomic clocks *2011 Joint Conf. IEEE Int. Frequency Control and the European Frequency and Time Forum (FCS) Proc.* (IEEE) pp 1–3
- [5] Huang X, Gong H and Ou G 2014 *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **61** 747–55
- [6] Khare A, Arora R, Banik A and Mehta S D 2016 *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **63** 326–35
- [7] Nunzi E and Carbone P 2007 *Proc. SPIE* **6603** 237–45
- [8] Nunzi E and Carbone P 2008 *Metrologia* **45** S103–7
- [9] Nunzi E, Carbone P and Tavella P 2008 Fault detection in atomic clock frequency standards affected by mean and variance changes and by an additive periodic component: the GLRT approach *2008 IEEE Instrumentation and Measurement Technology Conf.* (IEEE) pp 1594–7
- [10] Nunzi E and Saltanocchi G 2009 A simplified detector for atomic clock anomalies detection *2009 IEEE Workshop on Environmental, Energy, and Structural Monitoring Systems* (IEEE) pp 15–8
- [11] Nunzi E and Saltanocchi G 2010 Real-time detection of anomalies for atomic clock in space by means of the GLRT *EFTF-2010 24th European Frequency and Time Forum* (IEEE) pp 1–6
- [12] Galleani L and Tavella P 2003 The characterization of clock behavior with the dynamic Allan variance *Proc. 2003 IEEE Int. Frequency Control Symp. and PDA Exhibition Jointly with the 17th European Frequency and Time Forum* pp 239–44
- [13] Galleani L and Tavella P 2009 *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **56** 450–64
- [14] Galleani L and Tavella P 2016 *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **63** 624–35
- [15] Sesia I, Galleani L and Tavella P 2011 *IEEE Trans. Aerosp. Electron. Syst.* **47** 884–95
- [16] Rodríguez-Pérez I, García-Serrano C, Catalán C C, García A M, Tavella P, Galleani L and Amarillo F 2011 *Adv. Space Res.* **47** 197–212
- [17] Riley W J 2008 *Metrologia* **45** S154–61
- [18] Galleani L and Tavella P 2010 An algorithm for the detection of the frequency jumps in space clocks *Proc. 42nd Annual Precise Time and Time Interval Systems and Applications Meeting* (Reston, Virginia) pp 503–8
- [19] Galleani L and Tavella P 2017 *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **64** 475–85
- [20] Formichella V and Tavella P 2018 A recursive clock anomalies detector with double exponential smoothing *2018 European Frequency and Time Forum (EFTF)* (IEEE) pp 245–8
- [21] Zucca C, Tavella P and Peskir G 2016 *Metrologia* **53** S89–95
- [22] Chan Y C, Johnson W A, Karuza S K, Young A M and Camparo J C 2010 *IEEE Trans. Instrum. Meas.* **59** 330–4
- [23] Süß M 2014 Robust composite clock using modified Kalman filtering *Dissertation* University of Stuttgart <https://elib.dlr.de/122015/>
- [24] Gödel M and Furthner J 2017 Robust ensemble time onboard a satellite *Proc. 48th Annual Precise Time and Time Interval Systems and Applications Meeting ION* (Monterey, California: ION) (Monterey, US 30 January–1 February 2017) pp 26–43
- [25] Zenzinger A, Bartusch T, Kuehl C, Fischer S and Shrestha A 2012 Failure detection and correction for clock ensemble in space *2012 6th ESA Workshop on Satellite Navigation Technologies (Navitec 2012) & European Workshop on GNSS Signals and Signal Processing* (IEEE) pp 1–8
- [26] Wang Q and Rochat P 2009 An anomaly clock detection algorithm for a robust clock ensemble *Proc. 41st Annual Precise Time and Time Interval Systems and Applications Meeting* pp 121–30
- [27] Wang Q, Droz F and Rochat P 2015 Robust clock ensemble for time and frequency reference system *2015 Joint Conf. IEEE Int. Frequency Control Symp. & the European Frequency and Time Forum* (IEEE) pp 374–8
- [28] Wang Q and Rochat P 2019 Algorithms development and verification for next generation on-board clock monitoring and control unit *2019 Joint Conf. IEEE Int. Frequency Control Symp. European Frequency and Time Forum (EFTF/IFC)* (IEEE) pp 1–5
- [29] Viðarsson L, Pullen S, Green G and Enge P 2001 Satellite autonomous integrity monitoring and its role in enhancing GPS user performance *Proc. 14th Int. Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS 2001)* (Salt Lake City, Utah) pp 690–702
- [30] Trainotti C and Giorgi G 2020 Detection and identification of phase and frequency drifts in clock ensembles *Proc. 51st Annual Precise Time and Time Interval Systems and Applications Meeting* (San Diego, CA, USA) <https://ion.org/ptti/abstracts.cfm?paperID=8062>
- [31] Trainotti C, Giorgi G and Furthner J 2019 Detection and identification of faults in clock ensembles *Proc. 32nd Int. Technical Meeting of the Satellite* (Division of the Insti-

- tute of Navigation) pp 3342–58 <https://ion.org/publications/abstract.cfm?articleID=17104>
- [32] Beard R and Senior K 2017 *Clocks* (Berlin: Springer) pp 121–64
- [33] Zucca C and Tavella P 2005 *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **52** 289–96
- [34] Davis J A, Greenhall C A and Stacey P W 2005 *Metrologia* **42** 1–10
- [35] Trainotti C, Schmidt T and Furthner J 2019 Comparison of clock models in view of clock composition, clock steering and measurement fitting *Proc. 50th Annual Precise Time and Time Interval Systems and Applications Meeting (PTTI 2019)* (Reston, Virginia) pp 265–83 <https://ion.org/publications/abstract.cfm?articleID=16756>
- [36] Greenhall C A 2006 *Metrologia* **43** S311–21
- [37] Galleani L 2011 *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **58** 1550–8
- [38] Howe D A, Allan D U and Barnes J A 1981 Properties of signal sources and measurement methods *35th Annual Frequency Control Symp.* (IEEE) <https://ieeexplore.ieee.org/document/1537481>