

Real-Time Anomaly Detection and Categorization for Satellite Reaction Wheels

Alejandro Penacho Riveiros, Yu Xing, Nicola Bastianello, Karl H. Johansson

Abstract—In this paper, we address the problem of detecting anomalies in the reaction wheel assemblies (RWAs) of a satellite. These anomalies can alert of an impending failure in a RWA, and effective detection would allow to take preventive action. To this end, we propose a novel algorithm that detects and categorizes anomalies in the friction profile of an RWA, where the profile relates spin rate to measured friction torque. The algorithm, developed in a probabilistic framework, runs in real-time and has a tunable false positive rate as a parameter. The performance of the proposed method is thoroughly tested in a number of numerical experiments, with different anomalies of varying severity.

Index Terms—anomaly detection, satellite, reaction wheel assembly, log-likelihood ratio

I. INTRODUCTION

In the last decades, the number of satellites deployed in space has increased rapidly [1], owing to the technological advances that have considerably decreased the cost of an orbital launch. However, repairing satellites in Earth orbit is still quite challenging, as it usually requires a dedicated space mission for the purpose [2]. Hence, satellites are expected to operate reliably throughout their lifetime. Some of the most critical components of a satellite are its *reaction wheel assemblies* (RWAs), consisting of a rotating disk that can exchange angular momentum with the main body of the satellite to control its rotation [3]. RWAs spend most of their lifetimes in motion and thus are especially prone to failures that could disable the satellite, as documented during the missions of Cassini [4], [5], Kepler [6], Dawn [7], and XMM-Newton [8].

To avoid costly repairs, it is desirable to prevent RWAs failures altogether [9]. While the causes of such failures remain uncertain, some anomalous patterns preceding a failure have been observed in the aforementioned missions [10], [11]. Designing robust and trustworthy algorithms to detect such anomalous patterns would enable alerting of ground control to any imminent failure, and allow the ground control to take mitigative measures, such as switching to an alternative attitude control [10], [12], [13]. To this end, we focus on detecting anomalies in RWAs through their *friction profile*, which relates spin rate to the measured friction torque. The friction torque takes into account the contribution of both dry and viscous friction, and pinpointing in which component an anomaly has occurred is of interest.

This work was supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 101070162.

The authors are with the School of Electrical Engineering and Computer Science, and Digital Futures, KTH Royal Institute of Technology, Sweden, { alejpr | yuxing2 | nicolba | kallej }@kth.se.

a) *Literature review*: The nature of the anomaly detection problem can be approached from two different angles: as a space system, or as a rotary machine.

The literature on fault detection for satellite systems, although extensive, relies almost exclusively on data-driven approaches. This can be attributed to the complexity of these systems and the uncertainty regarding the possible faults that can occur. Approaches can be classified into two groups, output-based and prediction-based.

The first group encompasses the approaches based on training a model capable of determining the condition of the satellite based on the telemetry. One possibility is to use supervised training with nominal and faulty datasets, so the model is capable of classifying the telemetry of the satellite [14]. However, the lack of real faulty data from satellites has encouraged the development of methods based on outlier detection, such as Orca [15], [16] and Novelty [17], which can be trained purely on nominal data. The second group consists of the training of a model to generate the expected telemetry of the satellite under nominal conditions. The model is then run in parallel with the satellite and a warning is triggered when there is a significant discrepancy between real and simulated data. The model can be obtained using autoregressive functions [18], neural networks [19] or Markov chains [20], among other options.

Approaching the subject from the field of fault detection on rotary machines, the extensive knowledge of the dynamics of these systems has allowed the use of model-based approaches, such as [21]. Additionally, model-based methods have been successfully used for fault prognosis [22], [23]. Data-driven approaches are also popular for rotary machines [24], with a variety of techniques including Support Vector Machines [25], K-Nearest Neighbor [26] and deep learning [27].

Despite the amount of knowledge collected on space systems and rotary machines, research on fault detection for satellite reaction wheels has not experienced considerable development, with some model-based approaches tackling fault prognosis [28] and fault detection [29]. More aligned with the objectives of this work are data-driven approaches for detecting anomalies in the friction of the reaction wheel, either purely statistical [11] or learning-based [30]. These papers also considered the anomalies to be an early warning of an upcoming failure of the reaction wheel. However, both methodologies were applied retroactively to the satellite case study.

The approach presented here leverages methods in change-point detection, concerned with the detection of sudden

changes in trendlines. Literature on this topic can be classified into Bayesian and non-Bayesian approaches. Bayesian approaches deal with the computation of the probability of a changepoint having taken place. There have been some works mostly on offline algorithms [31], [32], which need the full data time series before starting computations. Online algorithms have also been developed [33]. Although this family of algorithms can keep track of several hypothetical changepoints to obtain accurate estimates, they are usually too complex when the objective is to obtain a simple binary answer.

On non-Bayesian approaches, many popular and computationally inexpensive algorithms have been produced in the last decades, such as CUSUM [34] or the generalized likelihood-ratio [35], [36]. These methods work in real-time, and provide a good baseline to develop a method tuned to the specific characteristics of reaction wheels. Therefore, the proposed algorithm is developed within the non-Bayesian framework.

b) Contributions: The paper has the following contributions:

- 1) We design a novel anomaly detection algorithm for satellites' RWAs. The algorithm detects deviations of the friction profile in real time and pinpoints whether the deviations occur due to changes in either the viscous or dry components.
- 2) The algorithm is developed in a rigorous framework, making use of the log-likelihood ratio. This allows us to select the tolerable false positive rate as a parameter.
- 3) The performance of the algorithm is thoroughly tested on a number of different anomalies, both in viscous and dry friction and of different severities.

II. PROBLEM FORMULATION & ALGORITHM DESIGN

We start this section by formalizing the anomaly detection problem and then discuss our proposed detection algorithm.

A. Problem formulation

The focus of this work is on detecting anomalies for the RWA of a satellite, of which an example cross-sectional view can be seen in Figure 1. For the purposes of anomaly

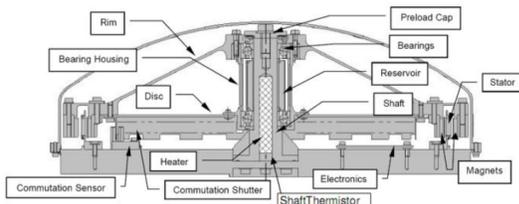


Fig. 1. Cross-sectional view of the XMM-Newton Reaction Wheel Assembly [8]. The main component is the disc, which rotates around the shaft in order to store angular momentum. The disc is accelerated by means of magnets, so the only point of contact between the rotating disc and the body of the satellite is through the bearings between the inner part of the disc and the shaft. This point of contact produces a small amount of friction torque that can be estimated from telemetry. Along with the torque, telemetry provides an accurate measurement of the disc's angular speed.

detection, the behavior of the RWA will be characterized by its *friction profile* [37]. The friction profile relates the total friction torque experienced at the RWA bearing (f , in millinewton-meter) with the spin rate (ω , in radians per second), both quantities that can be obtained from telemetry. There are two factors contributing to the friction profile: the *dry* and *viscous* friction torques. The former is a constant torque that does not depend on the spin rate (the intercept of the lines in Figure 2), while the latter is typically an increasing function of the spin rate (the slope of the lines in Figure 2) [38].

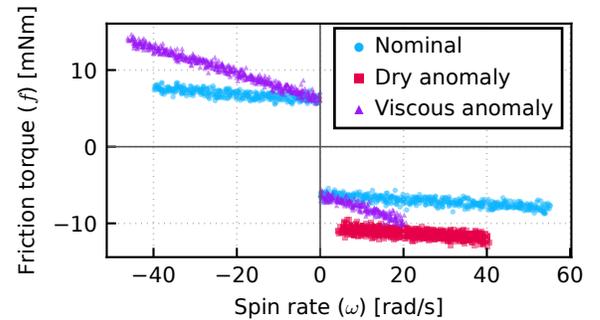


Fig. 2. Examples of different friction profiles. The intercept of a line represents the dry friction, hence a change in the intercept is categorized as a dry friction anomaly. The slope of a line instead represents the effect of viscous friction, with a change in slope signifying a viscous friction anomaly.

Let $f(\omega)$ denote the friction profile, that is, the total friction torque experienced at spin rate ω . Under nominal conditions this function does not change. However, when changes in the friction predicted by the nominal $f(\omega)$ deviate from the data measured at the RWA, an anomaly occurs, which may indicate an impending failure. The objective of the following sections is to design an algorithm that can detect such changes from nominal behavior. Additionally, the algorithm should pinpoint whether the friction profile changes are caused by a change in viscous or dry friction. Characterizing the specific nature of the anomaly can indeed lead to a better assessment of the RWA health [11], [30]. In summary, we are interested in solving the following problem.

Problem: Design an algorithm to detect anomalies in the behavior of an RWA by analyzing its friction profile. Additionally, the algorithm should categorize the anomalies as dry or viscous friction anomalies.

B. Preliminaries

The first step is to select a model for the friction profile of the reaction wheel, based on which we will design the proposed algorithm in section II-C. We choose a well-accepted linear model that accounts for dry and viscous friction and includes random Gaussian noise in order to account for measurement errors [38]. Denoting by $\omega_\ell \in \mathbb{R}$ the spin rate at time $\ell \in \mathbb{N}$, the total friction torque is

modeled as

$$\begin{aligned} f_\ell &= -f_\ell^c \text{sign}(\omega_\ell) - f_\ell^v \omega_\ell + \varepsilon_\ell \\ &= h_\ell x_\ell + \varepsilon_\ell, \end{aligned} \quad (1)$$

where f_ℓ^c and f_ℓ^v are the dry and viscous friction parameters, respectively, ε_ℓ is random Gaussian noise with zero mean, and $x_\ell = [f_\ell^c \ f_\ell^v]^T$, $h_\ell = -[\text{sign}(\omega_\ell) \ \omega_\ell]$. The friction parameters are scaled so that the noise has unit variance,¹ and they may change over time due to the occurrence of an anomaly.

In the following, we define an anomaly as a change between $x_{\ell-1}$ and x_ℓ , due to either an increase in dry friction (the first component), viscous friction (the second), or both. Examples of the effect of both types of anomalies are shown in Figure 2, and can be recognized by the greater (in absolute value) intercept in the y-axis for the dry friction anomaly, and the greater slope for the viscous anomaly.

The simplicity of model (1) will allow for the design of an efficient detection algorithm which may be deployed *on-board the satellite* itself and in *real time*, while still achieving effective detection, as shown in section III.

C. Algorithm design

With model (1) in place, we can now develop the proposed algorithm. The idea is to estimate *online* the current value of the friction parameters x_k , and notify of an anomaly whenever significant changes occur. In the following we introduce the estimation procedure – based on *least squares* – and the anomaly detection trigger – based on the *generalized likelihood ratio*.

a) State estimation: Consider a subset of data points sampled in the window of time $\mathcal{S} = [k_1, k_2]$. Owing to the linearity of model (1) w.r.t. the friction coefficients, and the assumption of zero mean Gaussian noise, the optimal state estimation can be carried out by solving the least squares problem

$$\hat{x}_\mathcal{S} = \arg \min_x \sum_{\ell \in \mathcal{S}} [y_\ell - h_\ell x]^2 = \Pi_\mathcal{S}^{-1} q_\mathcal{S}, \quad (2)$$

where we denote $\Pi_\mathcal{S} = \sum_{\ell \in \mathcal{S}} h_\ell^T h_\ell$ and $q_\mathcal{S} = \sum_{\ell \in \mathcal{S}} h_\ell^T f_\ell$. Hereafter, $\hat{x}_\mathcal{S}$ denotes the least squares estimate over the data collected during $\mathcal{S} = [\ell_1, \ell_2]$. The estimate can be updated *recursively*. Indeed, updating $q_\mathcal{S}$ as new points are added or removed from \mathcal{S} is straightforward, and updating $\Pi_\mathcal{S}^{-1}$ can be made inexpensive by using the Sherman-Morrison formula:

$$\Pi_{\mathcal{S} \cup \{\ell\}}^{-1} = [\Pi_\mathcal{S} + h_\ell^T h_\ell]^{-1} = \Pi_\mathcal{S}^{-1} - \frac{\Pi_\mathcal{S}^{-1} h_\ell^T h_\ell \Pi_\mathcal{S}^{-1}}{1 + h_\ell \Pi_\mathcal{S}^{-1} h_\ell^T}. \quad (3)$$

This allows us to quickly update the state estimate as new data points are obtained, which is critical for the computational performance of the algorithm.

¹This can be done by computing the scaling factor from historical data. While this simplifies the discussion of the proposed algorithm, a version that does not rely on rescaling the friction parameters can also be designed.

b) Detection metric: Consider k as the candidate time for an anomaly; we define the two subsets of data points $\mathcal{P}_k = [0, k-1]$ and $\mathcal{T}_k = [0, k+w-1]$, where $w \in \mathbb{N}$ is the window size, a parameter of the algorithm that will be discussed later. \mathcal{P}_k represents the past data points collected before time k , while \mathcal{T}_k also includes $w-1$ data points collected after k . The metric we use to trigger an anomaly detection at k , the log-likelihood ratio (LLR), is defined as follows [39]:

$$\text{LLR}_k = \log \left[\frac{p(f_{\mathcal{P}_k} \mid x = \hat{x}_{\mathcal{T}_k})}{p(f_{\mathcal{P}_k} \mid x = \hat{x}_{\mathcal{P}_k})} \right]. \quad (4)$$

The metric measures how much the optimal state estimate on the points in \mathcal{P}_k is degraded when we include the points in $[k, k+w-1]$. If there is no change of state, we can expect $\hat{x}_{\mathcal{T}_k} \approx \hat{x}_{\mathcal{P}_k}$, and the ratio will be close to 1, resulting in an LLR close to zero. If there is a change of state, however, the estimate $\hat{x}_{\mathcal{T}_k}$ will perform much worse on the data in \mathcal{P}_k and yield a non-zero value of LLR_k . Recalling the notation of state estimate (2) over a window of time and taking advantage of the Gaussian distribution of the noise, the LLR at time k can be rewritten as:

$$\begin{aligned} \text{LLR}_k &= \sum_{\ell \in \mathcal{P}_k} [f_\ell - h_\ell \hat{x}_{\mathcal{T}_k}]^2 - \sum_{\ell \in \mathcal{P}_k} [f_\ell - h_\ell \hat{x}_{\mathcal{P}_k}]^2 \\ &= \sum_{\ell \in \mathcal{P}_k} [h_\ell \hat{x}_{\mathcal{T}_k}]^2 - [h_\ell \hat{x}_{\mathcal{P}_k}]^2 - 2f_\ell h_\ell [\hat{x}_{\mathcal{T}_k} - \hat{x}_{\mathcal{P}_k}] \\ &= \hat{x}_{\mathcal{T}_k}^T \Pi_{\mathcal{P}} \hat{x}_{\mathcal{T}_k} - \hat{x}_{\mathcal{P}_k}^T \Pi_{\mathcal{P}} \hat{x}_{\mathcal{P}_k} - 2q_{\mathcal{P}_k} [\hat{x}_{\mathcal{T}_k} - \hat{x}_{\mathcal{P}_k}] \\ &= [\hat{x}_{\mathcal{T}_k} + \hat{x}_{\mathcal{P}_k}]^T \Pi_{\mathcal{P}} - 2q_{\mathcal{P}_k}^T [\hat{x}_{\mathcal{T}_k} - \hat{x}_{\mathcal{P}_k}] \\ &= [\hat{x}_{\mathcal{T}_k} - \hat{x}_{\mathcal{P}_k}]^T \Pi_{\mathcal{P}} [\hat{x}_{\mathcal{T}_k} - \hat{x}_{\mathcal{P}_k}] \end{aligned} \quad (5)$$

As discussed above, all quantities involved can be computed recursively, and thus the LLR can be obtained quickly at every step, allowing us to track this metric online.

The following Lemma 1 characterizes the distribution of the LLR in the presence or not of an anomaly, which will allow us to design the proposed algorithm in the next section.

Lemma 1: Consider a signal $\{f_\ell\}_{\ell=0}^{k+w}$ generated according to (1), with the state begin x_0 for $\ell \in [0, k-1]$ and $x_0 + \Delta x$ for $\ell \in [k, k+w-1]$, $\Delta x \geq 0$. Then the log-likelihood ratio is distributed as

$$\text{LLR}_k \sim X_1^2 + X_2^2, \quad (6)$$

where $X_i \sim \mathcal{N}(\mu_i, \lambda_i)$, $i = 1, 2$, corresponding to a generalized χ^2 distribution. The values of λ_i are the eigenvalues of $I - \Pi_{\mathcal{P}_k}^{-\frac{1}{2}} \Pi_{\mathcal{T}_k}^{-1} \Pi_{\mathcal{P}_k}^{-\frac{1}{2}}$, while the values of μ_i are the components of $U^T \Pi_{\mathcal{P}_k}^{-\frac{1}{2}} [I - \Pi_{\mathcal{P}_k}^{-1} \Pi_{\mathcal{T}_k}]$, with U consisting of left-singular vectors in the singular value decomposition of $[I - \Pi_{\mathcal{P}_k}^{-\frac{1}{2}} \Pi_{\mathcal{T}_k}^{-1} \Pi_{\mathcal{P}_k}^{-\frac{1}{2}}]^{\frac{1}{2}}$. For the nominal case ($\Delta x = 0$), the normal distributions have zero means, and LLR_k is a sum of two χ^2 distributions with 1 degree of freedom.

Computing the c.d.f. of the distribution functions presented in the lemmas is not trivial, but there are approximate methods for this task [40].

c) *The algorithm:* We are now ready to describe Algorithm 1. As mentioned above, the main idea behind the algorithm is to estimate the state online (computing $\hat{x}_{\mathcal{T}_k}$ and $\hat{x}_{\mathcal{P}_k}$), and then evaluating the LLR for the candidate change point k (see lines 6–9). Relying on Lemma 1, we can compute the probability that under nominal conditions an equal or higher value of the LLR is attained (line 11 and Algorithm 2):

$$p_k = P \left(\text{LLR}_k > \sum_{i=1}^2 N(0, \lambda_i)^2 \right). \quad (7)$$

This distribution depends on two parameters, λ_1 and λ_2 , and it is thus possible to generate a lookup table for the cumulative distribution function, which yields p_k given LLR_k . If the probability p_k is below a predefined threshold r , we declare that an anomaly has occurred (lines 18–22). We remark that r represents the *false positive rate* that we choose to allow. Notice that in practice, due to the measurements being noisy, the LLR in the presence of an anomaly may not present a strictly increasing trend (resp., the probability p_k may not be strictly decreasing). For this reason, we accept k as a change point only if p_k attained the minimum value over the window $[n, n + k_{\text{verif}}]$ (see lines 18–22). Finally, once an anomaly is declared, the algorithm is rebooted taking the time of the anomaly as the new starting point (see lines 20–21). We can then obtain an estimation of the friction profile by using Least Squares on the points before the changepoint.

d) *Parameter selection:* The proposed Algorithm 1 has three tunable parameters. The first is the false positive rate $r \in \mathbb{R}$, which represents the probability of incorrectly detecting an anomaly at time k . A small value for r is of course desirable, but considering that the measurements are noisy, reducing r too much could actually lead to failure in detecting an anomaly. The false positive rate can be tuned according to Lemma 1.

The second parameter is the window size $w \in \mathbb{N}$, which determines how many data points collected after k are employed to assess whether an anomaly has taken place at time k . There is a trade-off for the choice of w : on the one hand, a longer window yields better estimates and thus improves anomaly detection, lowering the false positive rate. On the other hand, the algorithm should run in real-time, and longer windows result in longer detection delays. Additionally, a longer window may include multiple anomalies, thus degrading the effectiveness of the algorithm.

The last parameter is the verification window $k_{\text{verif}} \in \mathbb{N}$, which represents the number of steps that we wait before declaring a changepoint. Indeed, as mentioned above, due to noise in the measurements the evolution of p_k may not be strictly decreasing towards an anomaly. Therefore, k_{verif} is used to ensure that an anomaly is not triggered by a local minimum of p_k .

III. NUMERICAL RESULTS & DISCUSSION

In this section, we present numerical results testing the proposed algorithm on different anomalies and discuss the

Algorithm 1 Anomaly Detector

Require: window size w , verification window k_{verif} , false positive rate r

- 1: $k \leftarrow w$
- 2: $\mathcal{P}, \mathcal{T} \leftarrow [0, w - 1], [0, 2w - 1]$ \triangleright sets of old and all data
- 3: $k_{\min}, p_{\min} \leftarrow 0, 1$ \triangleright time and p of worst step
- 4: **loop**
- 5: *// Collect new data-point*
- 6: $k \leftarrow k + 1$
- 7: $\mathcal{P}, \mathcal{T} \leftarrow \mathcal{P} \cup \{k - 1\}, \mathcal{T} \cup \{k + w - 1\}$
- 8: *// Compute LLR at this step*
- 9: $\text{LLR}_k \leftarrow \sum_{\ell \in \mathcal{P}} [(f_\ell - h_\ell \hat{x}_{\mathcal{P}})^2 - (f_\ell - h_\ell \hat{x}_{\mathcal{T}})^2]$
- 10: *// Compute the probability that the LLR has this value under nominal conditions*
- 11: $p_k \leftarrow \text{compute_nominal_probability}(\text{LLR}_k, \mathcal{P}, \mathcal{T})$
- 12: *// check if n is a candidate change point*
- 13: **if** $p_k < p_{\min}$ **then**
- 14: $p_{\min} \leftarrow p_k$
- 15: $k_{\min} \leftarrow k$
- 16: **end if**
- 17: *// declare an anomaly for the current candidate*
- 18: **if** $k > (k_{\min} + k_{\text{verif}})$ and $p_{n, \min} \leq r$ **then**
- 19: $k \leftarrow k + 1$
- 20: $\mathcal{P} \leftarrow [k_{\min}, k_{\min} + w - 1]$
- 21: $\mathcal{T} \leftarrow [k_{\min} + w, k_{\min} + 2w]$
- 22: **end if**
- 23: **end loop**

Algorithm 2 compute_nominal_probability

Require: current LLR, set of past points \mathcal{P} , set of total points \mathcal{T}

- 1: $\Pi_{\mathcal{T}} \leftarrow \sum_{k \in \mathcal{T}} h_k^T h_k$
- 2: $\Pi_{\mathcal{P}} \leftarrow \sum_{k \in \mathcal{P}} h_k^T h_k$
- 3: $\lambda \leftarrow \text{eig}(I - \Pi_{\mathcal{P}}^{1/2} \Pi_{\mathcal{T}}^{-1} \Pi_{\mathcal{P}}^{1/2})$
- 4: *// LLR follows $\lambda_1 \varepsilon_1^2 + \lambda_2 \varepsilon_2^2$ under nominal conditions*
- 5: $p_k \leftarrow p(\lambda_1 \varepsilon_1^2 + \lambda_2 \varepsilon_2^2 > \text{LLR})$

role of Algorithm 1’s tunable parameters. For each simulation, we select a spin rate profile and construct a friction torque profile, including the effect of noisy measurements. Throughout the section, we will fix $k_{\text{verif}} = w/2$.

A. Anomaly detection precision

We start by evaluating how precisely the proposed algorithm pinpoints the time when the anomaly occurs. To do so, we consider a dry friction increase of either 0.3 or 0.5 at time step $k = 1500$. The telemetry is generated according to (1), with the spin rate profile given by $\omega_k = 20 - 10 \cos(\pi k / 1200)$. We apply the algorithm on 200 simulation runs, with different realizations of the random noise, and choosing false positive rate $r = 10^{-5}$ and window size $w = 500$. The results are reported in Figure 3, with dry friction change of 0.3 (left) and 0.5 (right). The upper plots show the evolution of p_k in logarithmic scale, that

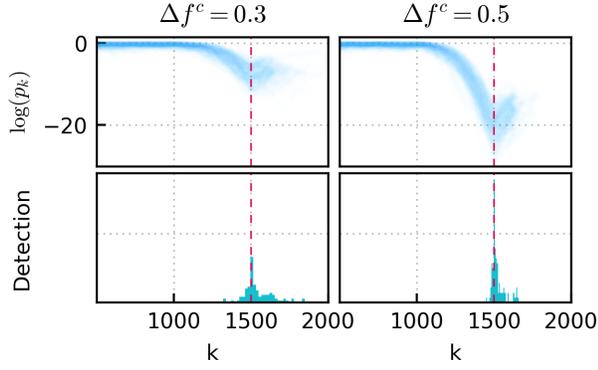


Fig. 3. Results of 200 simulations with a change of dry friction of 0.3 (left) or 0.5 (right) at step $k = 1500$. The upper plots show the evolution of the LLR, while the lower ones show the distribution of the step at which the anomaly was declared by the algorithm. In 4.5% of the simulations, no anomaly was declared.

is, the estimated probability that an anomaly has occurred. The lower plots show the distribution of the step at which the anomaly was declared. We can see that on average the algorithm accurately detects the anomaly at the correct time, with some variability as a consequence of the measurement noise.

Additionally, we notice that the larger the change in f^c is, the more precise the detection. To further investigate this trend, we evaluate the empirical detection probability for different values of the dry friction change Δf^c . The detection probability is computed as the percentage of the 100 runs in which the algorithm successfully detected the anomaly. We use the spin rate profile $\omega_k = 20 + 5 \cos(\pi k/1000)$, with $w = 500$ and three different values of r . In Figure 4 we

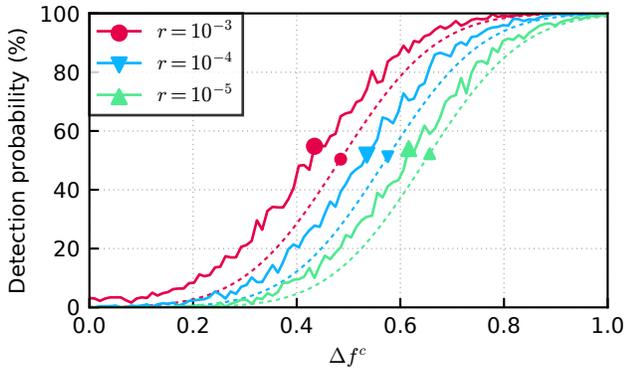


Fig. 4. Probability of detecting an increase of dry friction depending on its magnitude, for different false positive rates r . Solid lines represent the simulation results, dashed lines the theoretical distribution characterized in Lemma 1.

plot the empirical detection probability (solid lines) and the theoretical one (dashed lines) as characterized by Lemma 1. These results show that indeed the larger Δf^c is, the more accurately the algorithm can detect the anomaly. Moreover, they show how a smaller value of the false positive rate r

yields a more conservative detection with a smaller detection probability. Section III-C will provide further details on the effect of r .

B. False positives

We turn now to evaluating how often the proposed algorithm yields a false positive detection. In particular, we consider a set of 100 runs of the algorithm on telemetry *without* any anomaly, with spin rate profile $\omega_k = 20 + 5 \cos(k/5000)^2$. In Figure 5 we plot the average run length

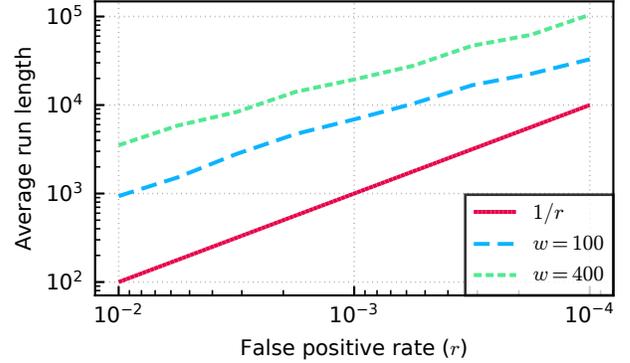


Fig. 5. Average run length until a false positive detection, for different values of the parameter r and window size w . The lower bound $1/r$ represents the average run length in the case of independent terms in $\{\text{LLR}_k\}_{k \in \mathbb{N}}$. Notice that r decreases when moving to the right of the x -axis.

before the algorithm returned a false positive, for different values of r (in decreasing order). We compare the empirical run length with $1/r$, which represents the average run length if $\{\text{LLR}_k\}_{k \in \mathbb{N}}$ were a sequence of independent r.v.s. We can see that the smaller r is, the fewer false positives are returned. Moreover, since the terms in $\{\text{LLR}_k\}_{k \in \mathbb{N}}$ are not actually independent, the average run in practice is larger than $1/r$.

C. The algorithm's parameters

In the previous sections, we have seen how different choices of the parameters r and w in Algorithm 1 affect its performance. In this section, we further explore the effect of these parameters on the detection probability. We run 100 simulations with spin rate profile $\omega_k = 100 - k/1250$ and an increase in dry friction of $\Delta f^c = 0.3$. We use as default values $r = 10^{-4}$ and $w = 500$. The empirical (solid lines) and theoretical (dashed lines) detection probabilities (as characterized by Lemma 1) are presented in Figures 6 and 7, for different values of w and r , respectively. From Figure 6 we observe how a larger window improves detection, as the state estimates with more data are more precise. And, as observed in Figure 4 as well, larger values result in a less conservative and thus more accurate detection algorithm. We further remark that the theoretical result of Lemma 1 provides a good lower bound to the empirical detection probability.

²This spin rate profile differs from that of the previous section in order to test the algorithm in diverse scenarios.

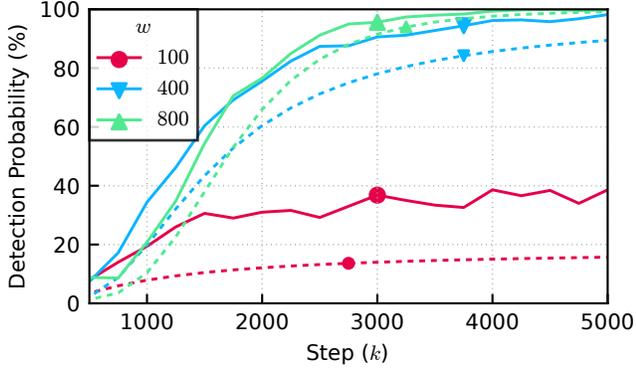


Fig. 6. Probability of detecting a $\Delta f^c = 0.3$ change in dry friction over time, for different window sizes. Solid lines represent the empirical probability, dashed lines represent the theoretical bound.

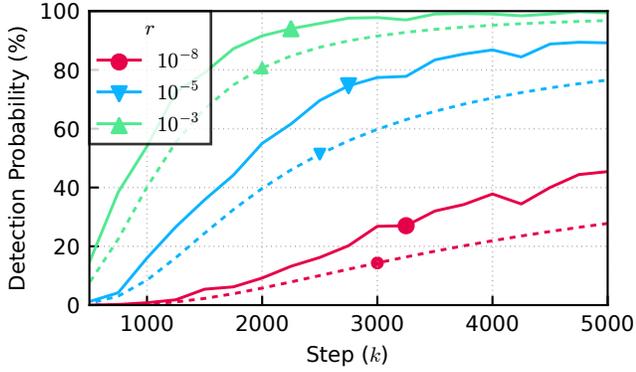


Fig. 7. Probability of detecting a $\Delta f^c = 0.3$ change in dry friction over time, for different false positive rates. Solid lines represent the empirical probability, dashed lines represent the theoretical bound.

D. Realistic scenario

In the previous sections, the results were derived by employing a periodical spin rate profile. In this section, we conclude by testing the proposed algorithm on a more realistic spin rate profile, which is depicted at the top of Figure 8. Such a profile makes the friction torque estimation, and thus the anomaly detection, more challenging. As we can see in the middle plot of the figure, the proposed algorithm still provides an accurate estimate of the friction. The lower figure then depicts the evolution of the log-likelihood ratio over time, with crosses denoting when an anomaly has been declared. The results were derived with $r = 10^{-4}$ and $w = 400$. We notice how the algorithm successfully detects most anomalies even in the presence of this more realistic spin rate profile. The run corresponds to 3 hours 32 minutes of data, which corresponds to 88000 datapoints. To process them all, the algorithm takes around 5.6 seconds using Python 3.11 on an Intel i7-1365U. Proportionally, it takes 1 second to process 38 minutes of data.

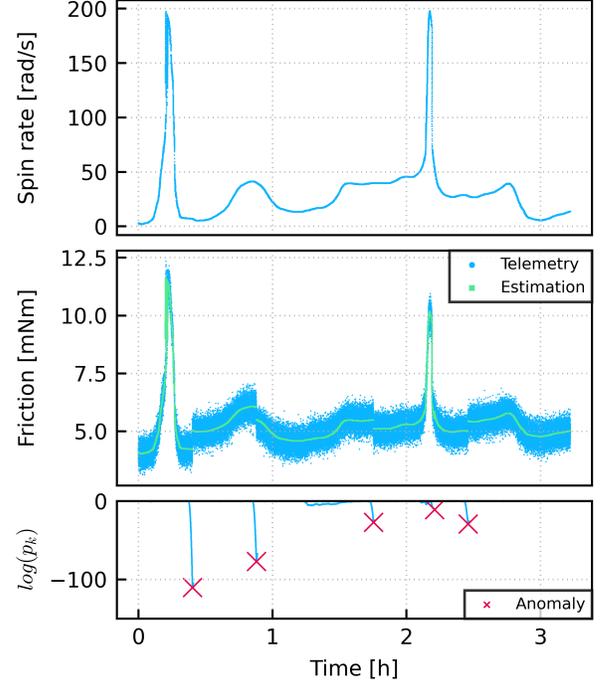


Fig. 8. State tracking for simulated satellite data

APPENDIX I PROOF OF LEMMA 1

We start by assuming that an anomaly occurs at time k , that is, the state of the RWA evolves according to

$$x_\ell = \begin{cases} x_0 & \text{if } \ell < k \\ x_0 + \Delta x & \text{otherwise} \end{cases} \quad (8)$$

for some Δx and k .

Letting $\mathcal{P} = [0, k - 1]$ and $\mathcal{T} = [0, k + w - 1]$, we define the following vectors and matrices:

$$\begin{aligned} H_{\mathcal{T}} &= [h_1^T \ \cdots \ h_{k+w-1}^T]^T && \in \mathbb{R}^{(k+w-1) \times 2} \\ F_{\mathcal{T}} &= [f_1^T \ \cdots \ f_{k+w-1}^T]^T && \in \mathbb{R}^{(k+w-1)} \\ \varepsilon_{\mathcal{T}} &= [\varepsilon_1 \ \cdots \ \varepsilon_{k+w-1}]^T && \in \mathbb{R}^{(k+w-1)} \\ H_{\mathcal{P}} &= [h_1^T \ \cdots \ h_{k-1}^T \ 0 \ \cdots \ 0]^T && \in \mathbb{R}^{(k+w-1) \times 2} \\ F_{\mathcal{P}} &= [f_1 \ \cdots \ f_{k-1} \ 0 \ \cdots \ 0]^T && \in \mathbb{R}^{(k+w-1)}. \end{aligned}$$

With this notation in place, we can then characterize the state estimates as a function of the noise ε as follows:

$$\begin{aligned} \hat{x}_{\mathcal{P}} &= \Pi_{\mathcal{P}}^{-1} q_{\mathcal{P}} = \Pi_{\mathcal{P}}^{-1} H_{\mathcal{P}}^T F_{\mathcal{P}} \\ &= \Pi_{\mathcal{P}}^{-1} H_{\mathcal{P}}^T [H_{\mathcal{P}} x_0 + \varepsilon_{\mathcal{T}}] \\ &= x_0 + \Pi_{\mathcal{P}}^{-1} H_{\mathcal{P}}^T \varepsilon_{\mathcal{T}} \end{aligned} \quad (9)$$

$$\begin{aligned} \hat{x}_{\mathcal{T}} &= \Pi_{\mathcal{T}}^{-1} q_{\mathcal{T}} = \Pi_{\mathcal{T}}^{-1} H_{\mathcal{T}}^T F_{\mathcal{T}} \\ &= \Pi_{\mathcal{T}}^{-1} H_{\mathcal{T}}^T [H_{\mathcal{T}} x_0 + [H_{\mathcal{T}} - H_{\mathcal{P}}] \Delta x + \varepsilon_{\mathcal{T}}] \\ &= x_0 + [I - \Pi_{\mathcal{T}}^{-1} \Pi_{\mathcal{P}}] \Delta x + \Pi_{\mathcal{T}}^{-1} H_{\mathcal{T}}^T \varepsilon_{\mathcal{T}}. \end{aligned} \quad (10)$$

As a consequence, the log-likelihood ratio is given by

$$\text{LLR}_k = [\hat{x}_{\mathcal{T}} - \hat{x}_{\mathcal{P}}]^T \Pi_{\mathcal{P}} [\hat{x}_{\mathcal{T}} - \hat{x}_{\mathcal{P}}] = z^T z$$

where

$$\begin{aligned} z &= \Pi_{\mathcal{P}}^{1/2} [\hat{x}_{\mathcal{T}} - \hat{x}_{\mathcal{P}}] \\ &= \Pi_{\mathcal{P}}^{1/2} [\Pi_{\mathcal{T}} H_{\mathcal{T}}^T \varepsilon_{\mathcal{T}} + [I - \Pi_{\mathcal{T}}^{-1}] \Delta x - \Pi_{\mathcal{P}} H_{\mathcal{P}}^T \varepsilon_{\mathcal{T}}] \\ &= \Pi_{\mathcal{P}}^{1/2} [\Pi_{\mathcal{T}}^{-1} H_{\mathcal{T}}^T - \Pi_{\mathcal{P}}^{-1} H_{\mathcal{P}}^T] \varepsilon_{\mathcal{T}} \\ &\quad + \Pi_{\mathcal{P}}^{1/2} [I - \Pi_{\mathcal{T}}^{-1} \Pi_{\mathcal{P}}] \Delta x \\ &= R \varepsilon_{\mathcal{T}} + Q \Delta x, \end{aligned}$$

$$\text{with } R = \Pi_{\mathcal{P}}^{1/2} [\Pi_{\mathcal{T}}^{-1} H_{\mathcal{T}}^T - \Pi_{\mathcal{P}}^{-1} H_{\mathcal{P}}^T] \text{ and } Q = \Pi_{\mathcal{P}}^{1/2} [I - \Pi_{\mathcal{T}}^{-1} \Pi_{\mathcal{P}}].$$

We take now the singular value decomposition $R = U \Sigma V^T$, where $U \in \mathbb{R}^{2 \times 2}$, $V \in \mathbb{R}^{(k+w) \times (k+w)}$ and $\Sigma \in \mathbb{R}^{2 \times (k+w)}$, which yields:

$$\begin{aligned} z &= U \Sigma V^T \varepsilon_{\mathcal{T}} + U \tilde{Q} \Delta x \\ &= U \tilde{\Sigma} E V^T \varepsilon_{\mathcal{T}} + U \tilde{Q} \Delta x \\ &= U [\tilde{\Sigma} \eta + \tilde{Q} \Delta x] \end{aligned} \quad (11)$$

where $\tilde{Q} = UQ$, $\tilde{\Sigma}$ is a square diagonal matrix with the non-zeros entries of Σ (which we call σ_1 and σ_2) and $E = [e_1 \ e_2]^T$, e_i being a vector where the i -th entry is 1 and the other are 0. We have then defined $\eta = EV^T \varepsilon_{\mathcal{T}}$, whose entries are independent and normally distributed since both E and V^T are orthonormal. We can then rewrite the distribution of the LLR_k as:

$$\begin{aligned} \text{LLR}_k &= z^T z = [\tilde{\Sigma} \eta + \tilde{Q} \Delta x]^T U^T U [\tilde{\Sigma} \eta + \tilde{Q} \Delta x] \\ &= \left| \tilde{\Sigma} \eta + \tilde{Q} \Delta x \right|^2 \\ &= \left[\sigma_1 \eta_1 + e_1^T \tilde{Q} \Delta x \right]^2 + \left[\sigma_2 \eta_2 + e_2^T \tilde{Q} \Delta x \right]^2 \end{aligned} \quad (12)$$

This last expression corresponds to the sum of two squared normally distributed variables. In order to obtain the non-zero entries of Σ , we note that $RR^T = U \Sigma^2 U^T = I - \Pi_{\mathcal{P}} \Pi_{\mathcal{T}}^{-1} \Pi_{\mathcal{P}}$, so they have the same singular values. Since the right matrix is symmetric, these coincide with its eigenvalues λ_1 and λ_2 , which can be easily computed. Then, we obtain the singular values of Σ as $\sigma_1 = \sqrt{\lambda_1}$ and $\sigma_2 = \sqrt{\lambda_2}$.

This result holds for any Δx , including $\Delta x = 0$, thus characterizing the distribution of LLR both in the presence or not of an anomaly. \square

REFERENCES

- [1] J. C. McDowell, *General Catalog of Artificial Space Objects*, Jul. 2023.
- [2] *The Hubble Program - Servicing Missions - SM1*.
- [3] J. Narkiewicz, M. Sochacki, and B. Zakrzewski, "Generic Model of a Satellite Attitude Control System," *International Journal of Aerospace Engineering*, vol. 2020, e5352019, Jul. 2020.
- [4] T. S. Brown, "The Cassini Reaction Wheels: Drag and Spin-Rate Trends from an Aging Interplanetary Spacecraft at Saturn," in *AIAA Guidance, Navigation, and Control Conference*, ser. AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, Jan. 2016.
- [5] A. Y. Lee and E. K. Wang, "In-Flight Performance of Cassini Reaction Wheel Bearing Drag in 1997–2013," *Journal of Spacecraft and Rockets*, vol. 52, no. 2, pp. 470–480, 2015.
- [6] K. A. Larson, K. M. McCalmont, C. A. Peterson, and S. E. Ross, "Kepler Mission Operations Response to Wheel Anomalies," in *SpaceOps 2014 Conference*, American Institute of Aeronautics and Astronautics.
- [7] M. D. Rayman and R. A. Mase, "Dawn's operations in cruise from Vesta to Ceres," *Acta Astronautica*, vol. 103, pp. 113–118, Oct. 2014.
- [8] M. Pantaleoni, P. Chapman, T. Godard, *et al.*, "Curing XMM-Newton's reaction wheel cage instability: The in-flight re-lubrication experience," May 2014.
- [9] K.-C. Liu, P. Maghami, and C. Blaurock, "Reaction Wheel Disturbance Modeling, Jitter Analysis, and Validation Tests for Solar Dynamics Observatory," in *AIAA Guidance, Navigation and Control Conference and Exhibit*, ser. Guidance, Navigation, and Control and Co-located Conferences, American Institute of Aeronautics and Astronautics, Aug. 2008.
- [10] F. M. Ekin, "Solar Dynamics Observatory Reaction Wheel Bearing Friction Increase: Detection, Analysis, and Impacts," in *SpaceOps 2014 Conference*, ser. SpaceOps Conferences, American Institute of Aeronautics and Astronautics, May 2014.
- [11] J. Kampmeier, R. Larsen, L. F. Migliorini, and K. A. Larson, "Reaction Wheel Performance Characterization Using the Kepler Spacecraft as a Case Study," in *2018 SpaceOps Conference*, ser. SpaceOps Conferences, American Institute of Aeronautics and Astronautics, May 2018.
- [12] D. Sahnou, J. Kruk, T. Ake, *et al.*, "Operations with the new FUSE observatory: Three-axis control with one reaction wheel," *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 6266, Jun. 2006.
- [13] M. J. Dube, J. Fisher, S. Loewenthal, and P. Ward, "Recovery and Operational Best Practices for Reaction Wheel Bearings," NASA Johnson Space Center, 2020.
- [14] S. Gao, Q. Han, N. Zhou, *et al.*, "Dynamic and wear characteristics of self-lubricating bearing cage: Effects of cage pocket shape," *Nonlinear Dynamics*, vol. 110, no. 1, pp. 177–200, Sep. 2022.
- [15] S. D. Bay and M. Schwabacher, "Mining distance-based outliers in near linear time with randomization and a simple pruning rule," in *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, ser. KDD '03, New York, NY, USA: Association for Computing Machinery, Aug. 2003, pp. 29–38.
- [16] D. Iverson, "Data Mining Applications for Space Mission Operations System Health Monitoring," in *SpaceOps 2008 Conference*, ser. SpaceOps Conferences, American Institute of Aeronautics and Astronautics, May 2008.
- [17] S. Fustes, G. Picart, J.-Y. Tourneret, L. Chaari, A. Ferrari, and C. Richard, "Improving Spacecraft Health Monitoring with Automatic Anomaly Detection Techniques," in *SpaceOps 2016 Conference*, ser. SpaceOps Conferences, American Institute of Aeronautics and Astronautics, May 2016.
- [18] R. Fujimaki, T. Yairi, and K. Machida, "An Anomaly Detection Method for Spacecraft Using Relevance Vector Learning," in *Advances in Knowledge Discovery and Data Mining*, T. B. Ho, D. Cheung, and H. Liu, Eds., ser. Lecture Notes in Computer Science, Berlin, Heidelberg: Springer, 2005, pp. 785–790.
- [19] S. K. Ibrahim, A. Ahmed, M. A. E. Zeidan, and I. E. Ziedan, "Machine Learning Methods for Spacecraft Telemetry Mining," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 4, pp. 1816–1827, Aug. 2019.
- [20] J. Pang, D. Liu, Y. Peng, and X. Peng, "Collective Anomalies Detection for Sensing Series of Spacecraft Telemetry with the Fusion of Probability Prediction and Markov Chain Model," *Sensors*, vol. 19, no. 3, p. 722, Jan. 2019.
- [21] K. Loparo, M. Adams, W. Lin, M. Abdel-Magied, and N. Afshari, "Fault detection and diagnosis of rotating machinery," *IEEE Transactions on Industrial Electronics*, vol. 47, no. 5, pp. 1005–1014, Oct. 2000.
- [22] M. Rocchi, F. Mosciaro, F. Grottesi, *et al.*, "Fault prognosis for rotating electrical machines monitoring using recursive least square," in *2014 6th European Embedded Design in Education and Research Conference (EDERC)*, Sep. 2014, pp. 269–273.

- [23] A. Rahimi, K. D. Kumar, and H. Alighanbari, "Failure Prognosis for Satellite Reaction Wheels Using Kalman Filter and Particle Filter," *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 3, pp. 585–588, Mar. 2020.
- [24] M. Hamadache, J. H. Jung, J. Park, and B. D. Youn, "A comprehensive review of artificial intelligence-based approaches for rolling element bearing PHM: Shallow and deep learning," *JMST Advances*, vol. 1, no. 1, pp. 125–151, Jun. 2019.
- [25] V. Sugumaran and K. I. Ramachandran, "Effect of number of features on classification of roller bearing faults using SVM and PSVM," *Expert Systems with Applications*, vol. 38, no. 4, pp. 4088–4096, Apr. 2011.
- [26] C. K. Mechefske and J. Mathew, "Fault detection and diagnosis in low speed rolling element bearings Part II: The use of nearest neighbour classification," *Mechanical Systems and Signal Processing*, vol. 6, no. 4, pp. 309–316, Jul. 1992.
- [27] M. Jalayer, C. Orsenigo, and C. Vercellis, "Fault detection and diagnosis for rotating machinery: A model based on convolutional LSTM, Fast Fourier and continuous wavelet transforms," *Computers in Industry*, vol. 125, p. 103378, Feb. 2021.
- [28] M. S. Islam and A. Rahimi, "A Three-Stage Data-Driven Approach for Determining Reaction Wheels' Remaining Useful Life Using Long Short-Term Memory," *Electronics*, vol. 10, no. 19, p. 2432, Jan. 2021.
- [29] B. Bellali, A. Hazzab, I. K. Bousserhane, and D. Lefebvre, "Parameter Estimation for Fault Diagnosis in Nonlinear Systems by ANFIS," *Procedia Engineering*, 2012 International Workshop on Information and Electronics Engineering, vol. 29, pp. 2016–2021, Jan. 2012.
- [30] K. Naik, A. Holmgren, and J. Kenworthy, "Using Machine Learning to Automatically Detect Anomalous Spacecraft Behavior from Telemetry Data," in *2020 IEEE Aerospace Conference*, Mar. 2020, pp. 1–14.
- [31] P. Fearnhead, "Exact and efficient Bayesian inference for multiple changepoint problems," *Statistics and Computing*, vol. 16, no. 2, pp. 203–213, Jun. 2006.
- [32] O. Seidou, J. J. Asselin, and T. B. M. J. Ouarda, "Bayesian multivariate linear regression with application to change point models in hydrometeorological variables," *Water Resources Research*, vol. 43, no. 8, 2007.
- [33] R. P. Adams and D. J. C. MacKay, *Bayesian Online Changepoint Detection*, Oct. 2007.
- [34] G. Romano, I. Eckley, P. Fearnhead, and G. Rigai, *Fast Online Changepoint Detection via Functional Pruning CUSUM statistics*, Jul. 2022.
- [35] D. L. Solomon, "A Note on the Non-equivalence of the Neyman-Pearson and Generalized Likelihood Ratio Tests for Testing a Simple Null versus a Simple Alternative Hypothesis," *The American Statistician*, vol. 29, no. 2, pp. 101–102, May 1975.
- [36] M. Perlman and L. WU, "The Emperor's new tests," *Statistical Science*, vol. 14, pp. 355–369, 1999.
- [37] J. M. Hacker, J. Ying, and P. C. Lai, "Reaction Wheel Friction Analysis Methodology and On-orbit Experience (Invited)," in *AIAA/AAS Astrodynamics Specialist Conference*, American Institute of Aeronautics and Astronautics.
- [38] H. Olsson, K. J. Åström, C. Canudas de Wit, M. Gäfvert, and P. Lischinsky, "Friction Models and Friction Compensation," *European Journal of Control*, vol. 4, no. 3, pp. 176–195, Jan. 1998.
- [39] D. A. Abraham, "Chapter 11 - Signal Processing," in *Applied Underwater Acoustics*, T. H. Neighbors and D. Bradley, Eds., Elsevier, Jan. 2017, pp. 743–807.
- [40] R. B. Davies, "The Distribution of a Linear Combination of Chi-Squared Random Variables," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 29, no. 3, pp. 323–333, 1980.