Detection of Oscillatory Actuator Failures in Passenger Airliners

by

Dylan Els

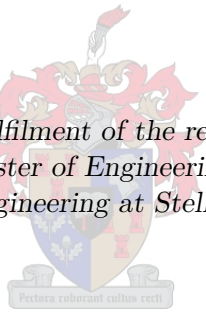

Thesis presented in partial fulfilment of the requirements for the degree of
Master of Engineering
in the Faculty of Engineering at Stellenbosch University

Supervisors:

Dr J.A.A. Engelbrecht Prof H.A. Engelbrecht

Department of Electrical and Electronic Engineering

April 2019

Declaration

1. I have read and understand the Stellenbosch University Policy on Plagiarism and the definitions of plagiarism and self-plagiarism contained in the Policy [Plagiarism: The use of the ideas or material of others without acknowledgement, or the re-use of one's own previously evaluated or published material without acknowledgement or indication thereof (self-plagiarism or text-recycling)].

2. I also understand that direct translations are plagiarism.

3. Accordingly all quotations and contributions from any source whatsoever (including the internet) have been cited fully. I understand that the reproduction of text without quotation marks (even when the source is cited) is plagiarism.

4. I declare that the work contained in this assignment is my own work and that I have not previously (in its entirety or in part) submitted it for grading in this module/assignment or another module/assignment.

	D. Els	***April 2019***
	Initials and Surname	Date

Abstract

This project investigates and develops techniques to detect oscillatory failure cases (OFCs) in aircraft control surface actuators. Oscillatory failures induce additional loads on the structure of the aircraft, requiring additional structural support to withstand these loads, increasing the overall mass of the aircraft. If oscillatory failures can be detected and pacified quickly, then the additional structural support would not be required, and the mass of the aircraft can be reduced, resulting in improved fuel efficiency and aircraft performance.

Oscillatory failure case (OFC) detection is performed by evaluating the difference (residual) between the measured behaviour of the real actuator and the simulated behaviour of a fault-free analytically redundant actuator model running in parallel with the real actuator. An OFC detection system must generate a residual signal using the analytically redundant actuator model, and evaluate the residual signal to determine whether an oscillatory failure is present. The challenge for the residual evaluation stage is to distinguish between the components of the residual signal resulting from modelling uncertainty and sensor noise, and the components resulting from an actual oscillatory failure case. The OFC detection system must detect oscillatory failures within a maximum allowable detection time, but must not produce false alarms.

Five different oscillatory failure detection techniques are investigated and developed, namely oscillation counting, integrated absolute error (IAE), discrete Fourier transform (DFT), multi-window Fourier transform (MWFT), and phase-locked loop (PLL) detection. Oscillation counting is an existing OFC detection technique that was developed by Goupil [1] and is currently in service on the Airbus A380 passenger airliner. The other four techniques are new OFC detection techniques that are developed in this project.

A simulation framework is created to serve as a testbed for the training and testing of the different OFC detection techniques. The simulation framework contains models for the physical actuator, the analytically redundant actuator, the oscillatory failures (both liquid and solid failures), the flight control system, and the aircraft longitudinal dynamics. The simulation models the aircraft's response to an oscillatory failure, since it affects the performance of the OFC detection.

The five OFC detection techniques are trained and rigorously tested using training and testing data generated with the simulation framework. The detection thresholds for each technique are "trained" on fault-free data to determine the lowest detection thresholds that do not produce false alarms. The detection techniques are then tested using testing data to determine the smallest amplitude oscillatory failure that each technique can detect within the specified maximum allowable detection time. The number of false alarms for each technique is also determined.

The results show that DFT, MWFT, and the PLL outperform oscillation counting and IAE by detecting smaller amplitude oscillatory failures and with shorter detection times, with MWFT providing the most promising results. However, oscillation counting and IAE are the most computationally efficient techniques, while DFT, MWFT, and PLL are more computationally expensive. Overall, the multi-window Fourier transform (MWFT) technique is the recommended approach for OFC detection, offering the best detection performance with only a small increase in computational complexity.

Uittreksel

Hierdie projek ondersoek en ontwikkel tegnieke om ossillatoriese faling gevalle (OFGs) in vliegtuig beheeroppervlak aktueerders te bespeur. Ossillatoriese falings induseer bykomende ladings op die struktuur van die vliegtuig, en vereis dus bykomende strukturele ondersteuning om hierdie ladings te weerstaan, wat die algehele massa van die vliegtuig verhoog. Indien ossillatoriese falings bespeur en vinnig gepassifiseer kan word, dan sou die bykomende strukturele ondersteuning nie benodig word nie, en die massa van die vliegtuig sou verminder kon word, wat sou lei tot verbeterde brandstofverbruik en werkverrigting.

Ossillatoriese faling geval (OFG) bespeuring word uitgevoer deur die verskil (residu) te evalueer tussen die gemete gedrag van die werklike aktueerder en die gesimuleerde gedrag van 'n foutvrye analities-oortollige aktueerder model wat in parallel met die werklike aktueerder uitvoer. 'n OFG bespeuringstelsel moet 'n residu sein genereer deur gebruik te maak van die analitiese-oortollige aktueerder model, en moet die residu evalueer om te bepaal of daar 'n ossillatoriese faling teenwoordig is. Die uitdaging vir die residu evaluasie stadium is om te onderskei tussen die komponente van die residu sein wat afkomstig is van model onsekerheid en sensor ruis, en die komponente wat afkomstig is van 'n werklike ossillatoriese faling geval. Die OFG bespeuringstelsel moet ossillatoriese falings bespeur binne 'n maksimum toelaatbare bespeuringstyd, en moet nie vals alarms gee nie.

Vyf verskillende ossillatoriese faling bespeuringstegnieke word ondersoek en ontwikkel, naamlik ossillasie telling, geïntegreerde absolute fout (IAE), diskrete Fourier transform (DFT), multi-venster Fourier transform (MWFT), en fase-sluit lus (PLL) bespeuring. Ossillasie telling is 'n bestaande OFG bespeuring tegniek wat ontwikkel is deur Goupil [1] en word tans gebruik op die Airbus A380 passassiersvliegtuig. Die ander vier tegnieke is nuwe OFG bespeuringstegnieke wat ontwikkel is in hierdie projek.

'n Simulasie raamwerk is geskep om te dien as 'n toetsplatform vir die opleiding en toets van die verskillende OFG bespeuringstegnieke. Die simulasie raamwerk bevat modelle vir die fisiese aktueerder, die analities-oortollige aktueerder, die ossillatoriese falings (beide vloeibare en soliede falings), die vlugbeheerstelsel, en die vliegtuig se longitudinale vlugdinamika. Die simulasie modelleer die vliegtuig se reaksie op die ossillatoriese faling, aangesien dit die prestasie van die OFG bespeuring beïnvloed.

Die vyf OFG bespeuringstegnieke is opgelei en volledig getoets deur gebruik te maak van opleiding en toets data wat genereer is met die simulasie raamwerk. Die bespeuring drempels vir elke tegniek is "opgelei" op foutvrye data om te bepaal wat die laagste bespeuringsdrempel is wat nie vals alarms gee nie. Die bespeuringstegnieke is dan getoets op toets data om te bepaal wat die kleinste amplitude ossillatoriese faling is wat elke tegniek kan bespeur binne die maksimum toelaatbare bespeuringstyd. Die aantal vals alarms vir elke tegniek is ook bepaal.

Die resultate wys dat die DFT, MWFT, en PLL tegnieke oortref die ossillasie telling en IAE tegnieke deur kleiner amplitude ossillatoriese falings te bespeur in korter bespeuringstye, met die MWFT wat die mees belowende resultate lewer. Die ossillasie telling en IAE tegnieke bly egter die mees berekeningsdoeltreffende tegnieke, terwyl die DFT, NWFT, en PLL meer berekeningskoste dra. Algeheel, word die multi-venster Fourier transform (MWFT) tegniek aanbeveel as die voorkeurtegniek, omdat dit die beste bespeuringsprestasie bied met net 'n klein verhoging in berekeningskoste.

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations and Acronyms

ACF	Autocovariance Function
AZCR	Autocovariance Zero Crossing Regularity
ADC	Analogue-to-Digital Converters
CS	Control Surface
DC	Direct Current
DFT	Discrete Fourier Transform
DTFT	Discrete-time Fourier Transform
EFCS	Electronic Flight Control Systems
EMD	Empirical Mode Decomposition
ESPRIT	Estimation of Signal Parameters via Rotational Invariance Techniques
FCC	Flight Control Computer
FDI	Fault Detection and Isolation
FFT	Fast Fourier Transform
IAE	Integrated Absolute Error
LPF	Lowpass Filter
MUSIC	Multiple Signal Classification
MWFT	Multi-window Fourier Transform
OC	Oscillation Counting
OFC	Oscillatory Failure Case
PD	Phase Detector
PDS	Power Density Spectrum
PI	Proportional Integral
PLL	Phase-Locked Loop
QPD	Quadrature Phase Detector
SDFT	Sliding Discrete Fourier Transform
SNR	Signal-to-Noise Ratio
SPRT	Sequential Probability Ratio Test
TBZC	Time between Zero Crossings

VCO	Voltage Controlled Oscillator
ZCR	Zero Crossing Regularity

Symbol Conventions

x	Scalar
$\dot{x}$	First time derivative of x
x_{meas}	Measurement of x
$\hat{x}$	Estimation of x

Constants

π	3.142
e	2.718

General Symbols

A	Amplitude
f	Frequency
N	Window Size
n	Time Step
r	Residual
T	Threshold
T_s	Sampling Period
t	Time
ω	Frequency

Chapter 3

e	Control Error
F_{aero}	Aerodynamic forces
$F_{damping}$	Passive actuator damping force
f_{liq}	Liquid failure case
f_{sol}	Solid failure case
i	Actuator command current
K	Servo-control gain
K_c	Current to speed conversion
K_d	Passive actuator damping coefficient
N_x, N_z	Longitudinal load factor
ΔP	Hydraulic pressure
ΔP_{ref}	Hydraulic pressure at maximum rod speed

p	Actuator position
q	Aircraft pitch rate
S	Actuator piston surface area
u	Actuator position command
v	Actuator rod speed
$\bar{v}$	Aircraft velocity magnitude
v_c	Commanded rod speed
α	Aircraft angle of attack
δ	Control surface deflection
θ	Aircraft pitch angle

Chapter 4

H	Filter Frequency Response
L	Upsampling Factor
r_L	Upsampled Residual
T_c	Counting Threshold
T_E	Energy Threshold

Chapter 5

e	Control Error
IAE_{lim}	Integrated Absolute Error Threshold
n_{lim}	Crossing Count Threshold
t_i	Zero Crossing Time Instant

Chapter 6

k	Frequency-domain Sample
k_f	OFC Frequency-domain Sample
m_z	Zero-padding Factor
n_f	Failure Start Time Step
n_o	Oldest Sample in Window
R	Fourier Transform of Residual
T_{eff}	Effective Threshold
t_d	Failure Detection Time
t_f	Failure Start Time
W	Fourier Transform of Window Function

Chapter 7

A	OFC Amplitude
B	VCO Output Amplitude
c	VCO Gain Factor
F	Loop Filter Transfer Function
v_i	PLL Input Signal
v_o	VCO Output Signal
v_d	QPD Output
v_c	VCO Input
η	Noise
θ_e	Phase Error
θ_{eq}	Steady-state Phase Error
θ_i	Input Phase
θ_o	Output Phase
φ	Quadrature Phase Detector Output
ω_c	VCO Centre Frequency
ω_i	OFC Frequency
ω_o	VCO Output Frequency
$\Delta\omega$	Frequency Difference between v_i and v_c

Acknowledgements

I would like to express my gratitude to the following people for their continued support throughout this project:

- My supervisors, Dr. Japie Engelbrecht and Prof. Herman Engelbrecht, for their guidance, wisdom, and pop culture references.
- Phillipe Goupil, Simone Urbano, and the rest of Airbus, for their intellectual assistance and financial support.
- My mother, for her love, encouragement, and (most importantly) food.
- My office colleagues, Gabriël Roux and Luke Hibbert, for all the jokes, the late nights, the beers, the coffees, the games, and the occasional bout of misery and despair.
- The Bridge players, JC Schoeman, Janneman Gericke, Willem de Jongh, and Clint Lombard for the many hours wasted playing cards.

Chapter 1

Introduction

1.1 Background and Motivation

Aircraft designers and manufacturers have steadily been phasing out mechanical flight control systems in favour of electrical flight control systems (EFCS), or "fly-by-wire" systems. These newer control systems allow for more sophisticated aircraft control, where the actuators of the aircraft control surfaces can be controlled electronically through the use of embedded computers. However, the use of EFCS introduces new failure points that have to be considered [1], especially at flight control surfaces.

Control surface failures can lead to substantial additional loads being applied to the structure of the aircraft, and as a result, structural engineers have to reinforce the aircraft structure to withstand these increased loads. This "local structural load augmentation" adds to the overall weight of the aircraft, resulting in decreased fuel efficiency, flight range, and flight control. A list of these control surfaces is given in Figure 1.1.

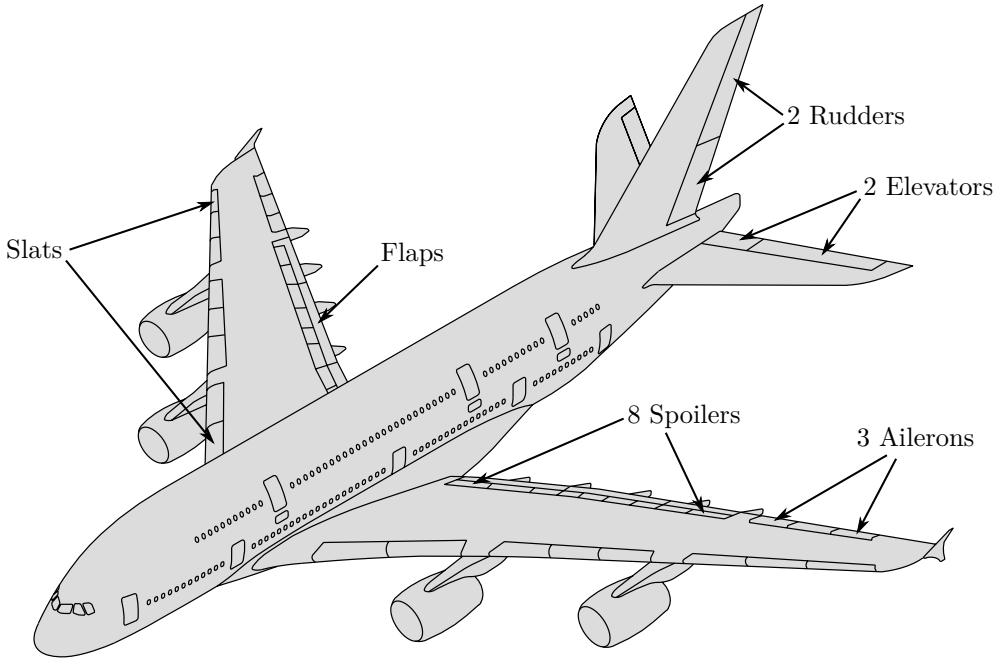

Figure 1.1: A380 Flight Control Surfaces. Image adapted from [2], [3].

One specific failure case that must be addressed is known as the oscillatory failure case (OFC). This occurs when a spurious oscillatory signal propagates through the control loop of a control surface actuator, causing the control surface to oscillate sinusoidally [1]. These failures typically

occur as a result of a single electronic component generating a spurious sinusoidal signal, and result in increased structural loads and overstress, reducing the aircraft's fatigue life [4]. The aircraft structure must therefore be designed to accommodate the additional loads caused by these actuator oscillatory failures.

With the increased economic pressure to improve fuel efficiency, commercial aircraft manufacturers are looking to improve the fault detection schemes implemented on their aircraft. By designing oscillatory failure detection methods to detect, isolate and pacify smaller amplitude oscillatory failures, less structural support could be used on the aircraft, leading to reduced weight and improved flight performance. The smaller the amplitude of the oscillatory failure that can be detected, the greater the potential for mass reduction.

For this project, only OFCs emerging from within the actuator control loop are considered. This limits the potential failure sources to the analogue-to-digital converters (ADC), servo command current, and sensors. Figure 1.2 provides an overview of the actuator system and potential OFC sources.

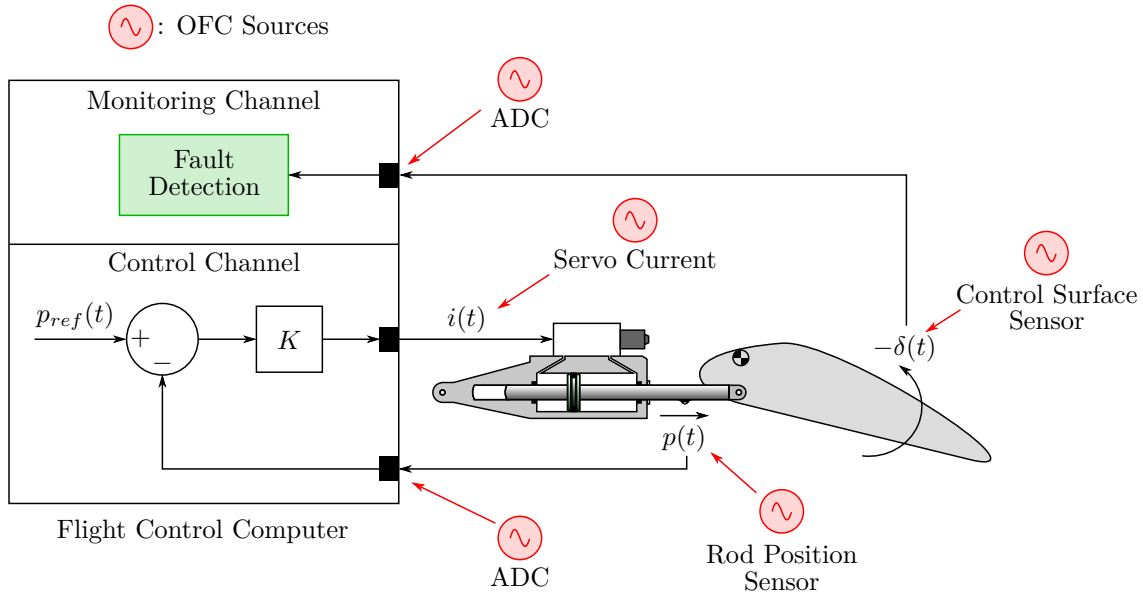

Figure 1.2: Failure Points within the Control Surface, adapted from [1]

Commercial aircraft manufacturers rely on hardware redundancy and automatic reconfiguration for fault-tolerance. Each control surface has two connected actuators, each controlled using independent computers and power sources [1], as illustrated in Figure 1.3. One actuator is in active mode, while the second is in passive, or stand-by mode. When a fault in the active actuator is detected, the active actuator changes to passive mode while the secondary actuator becomes active. This process is known as automatic reconfiguration.

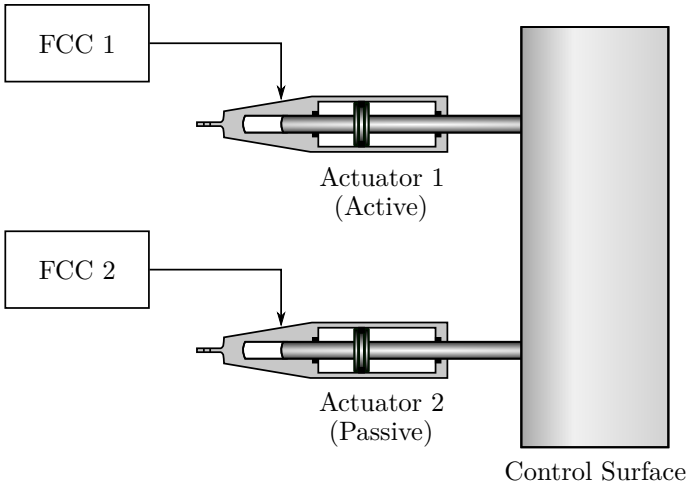

Figure 1.3: Hardware Redundancy [1]

To achieve further redundancy, the flight control computers are divided into two channels: a command channel and a monitoring channel. These channels are shown in Figure 1.2. The command channel controls the actuators and executes pilot commands. The monitoring channel, on the other hand, continuously monitors all components of the flight control system, checking for faults, and issuing reconfiguration commands when faults are detected.

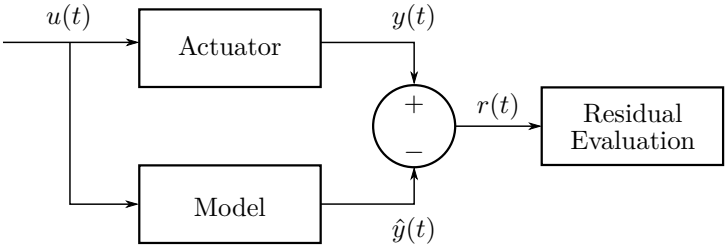

Figure 1.4: Analytical Redundancy

Fault detection systems implemented on the flight control computers utilise model-based approaches, also referred to as analytical redundancy. This involves comparing the measured behaviour of the actuators with some expected fault-free behaviour, which is obtained by simulating the motion of a fault-free actuator in parallel with the real actuator, as shown in Figure 1.4. Such a fault detection scheme can be divided into two stages: residual generation and residual evaluation. Residual generation is the process of determining the difference between the actuator's measured position and its expected position for normal operation, resulting in a signal known as the residual. In fault-free scenarios, this signal consists primarily of modelling errors and sensor noise. When a failure occurs, the residual signal experiences a bias with the characteristics of that type of failure [5]. Residual evaluation is the next stage that involves analysing the residual signal to check if these failure characteristics are present. Overall, this is known as component-level fault detection. Each actuator is monitored independently, using only locally available measurements and signal for fault detection. This also means that fault isolation is not required, since the fault is already localised to a single actuator.

1.2 Research Goal

The overall aim of this project is to investigate and innovate potential signal processing and residual evaluation methods that can detect oscillatory failure cases. The research topic was proposed by

Philippe Goupil of Airbus, who developed the existing oscillation counting technique for oscillatory failure case detection.

The residual evaluation techniques must detect OFCs with unknown frequencies occurring in a known frequency range, must detect as small amplitudes as possible, and must detect them within a maximum allowable detection time, specified as a fixed number of oscillation periods, or cycles. If an OFC of a given amplitude cannot be detected and pacified, then this amplitude must be considered and accommodated in the aircraft structural design. Thus, being able to detect smaller amplitude OFCs can allow for reduced structural support. Furthermore, short-time loads applied to the aircraft when an OFC is induced should be minimised, and thus it is required that OFCs be detected within a given number of oscillation periods, or cycles.

The residual evaluation techniques should avoid raising false alarms, since unnecessary reconfiguration of the actuator hardware removes a layer of hardware robustness. A typical requirement is that any fault detection system should report a false alarm no more than once every 100 000 hours of flight time. This is difficult to confirm without rigorous testing in actual aircraft, but should be considered when evaluating fault detection techniques.

Finally, since the flight control computers employed on commercial airliners are simpler and more robust, the computational resources for fault detection and isolation are limited. The residual evaluation techniques should therefore be computationally efficient.

1.3 System Requirements and Constraints

The following requirements for the OFC detection system were captured in consultation with Philippe Goupil from Airbus:

- The OFC detection system must detect OFCs that have an unknown frequency, but are known to occur in a frequency range from 1 to 10 Hz.
- An OFC must be detected within a maximum allowable detection time of 3 oscillation periods, or cycles.
- The OFC detection system must detect oscillatory failures with as small amplitudes as possible.
- The OFC detection system should not produce false alarms. The false alarm rate for each technique must be determined and compared.
- The OFC detection system should be computationally efficient. The computational complexity of each technique must be quantified and compared.

The following constraint was identified in consultation with Philippe Goupil of Airbus:

- The residual evaluation techniques must use the existing hardware architecture and available sensor measurements described in Figure 1.2. No new sensors may be added.

1.4 Primary Contributions

The following is a list of the primary contributions of this project:

- The oscillation counting approach was implemented as described in literature. However, potential disadvantages as a result of the use of filters were identified and confirmed.
- The integrated absolute error (IAE), an approach used for the detection of oscillating control loops in production plants, was implemented and adapted for oscillatory failure case detection. The approach was modified by implementing an adaptive threshold that would provide better performance given the unique requirements of the OFC detection problem.

- The discrete Fourier transform (DFT) was investigated as a potential technique for oscillatory failure detection. The concept of a frequency-dependent threshold was developed, based on initial poor detection performance when using a fixed threshold value. Following the implementation of the frequency-dependent threshold, a mathematical relationship between the detection time, threshold, and OFC characteristics was derived. This relationship predicts the performance of the DFT approach. The calculation of the DFT was implemented using the sliding-DFT algorithm, which provided improved computational performance over the fast Fourier transform.
- The structure of the sliding-DFT was exploited to allow for an implementation of the DFT that used shorter windows to calculate higher-frequency components. This was termed the multi-window Fourier transform (MWFT), and resulted in shorter detection times and a decrease in computational complexity over the single-window DFT implementation.
- A detection approach based on the use of a phase-locked loop (PLL) and quadrature phase detector was designed. This system was designed to detect oscillatory failures within 0.3 seconds, based on 3 cycles of a 10 Hz failure. However, the PLL struggled to operate at the low OFC frequencies, and this motivated the design of an upsampling and modulation stage that modulated the residual signal to a higher frequency band. A theoretical noise analysis of the PLL was also performed to predict its detection performance.
- A simulation framework was developed in collaboration with Airbus to test and benchmark all the detection methods through extensive Monte Carlo simulations. This framework contained a longitudinal model of an A380 aircraft, with the control surface actuators modelled as non-linear systems. The simulator allowed for the injection of oscillatory failures into the actuator system, and simulated the aircraft's response to the failure, further increasing the challenge of OFC detection.
- The five detection methods, namely oscillation counting, IAE, DFT, MWFT, and the PLL, were rigorously tested using the simulation framework. These methods were compared based on the smallest amplitude oscillatory failure that each method could reliably detect within the specified detection time of 3 cycles. Also compared were the average detection times to determine which method could report the failure in the shortest amount of time. The computational complexity of each technique was quantified and compared in terms of the number of real mathematical operations. Finally the number of false detections reported by each method was compared.

1.5 Thesis Outline

This thesis is organised as follows:

- Chapter 2 provides a literature review of the existing techniques for oscillatory failure detection, as well as more theoretical approaches to sinusoidal signal detection and estimation.
- Chapter 3 introduces the mathematical model and assumptions of the structure of the actuator and of oscillatory failures. Thereafter, the simulation framework that was created to serve as a testbed for the training and testing of the different OFC detection techniques is described.
- Chapter 4 reports on a detailed investigation into the current approach to OFC detection, namely oscillation counting.
- Chapter 5 investigates the integrated absolute error technique, an approach used for oscillatory failure detection in control loops in production plants.
- Chapter 6 presents the design of the OFC detection technique based on the discrete Fourier transform. This chapter covers both the single-window and the multi-window approaches.
- Chapter 7 presents the design of the OFC detection technique that uses a phase-locked loop and quadrature phase detector.

- Chapter 8 verifies and compares the five OFC detection techniques using the simulation framework introduced in Chapter 3.
- Chapter 9 offers some concluding remarks, and recommendations for future work.

The residual signal is usually expressed degrees, since it is based on control surface deflection angle. For simplicity, Chapters 4 to 7 assume the residual signal is expressed in volts. This is to avoid confusion between signal phase and control surface deflection. The results, however, will be expressed as degrees.

Chapter 2

Literature Review

This chapter presents a literature survey of existing techniques for detecting oscillatory failures in aircraft control surface actuators, as well as related techniques used in other applications. The problem of oscillation detection is a well-established field in literature, and has been explored for the application of oscillatory failure case detection in aircraft control surfaces, and oscillations in control loops of large production plants. First, the existing Oscillation Counting technique that is employed on the Airbus A380 aircraft for OFC detection is reviewed, along with improvements that have been proposed for both its residual generation and its residual evaluation stages. Next, a survey is performed of related techniques for detecting oscillations in the control loops of large production plants. Finally, a survey is performed of additional techniques for the general detection of oscillations and sinusoids. Based on the literature survey, energy/integral techniques, frequency domain techniques, and demodulation techniques are identified as promising techniques to be investigated and applied to the OFC detection problem.

2.1 Oscillatory Failure Detection in the A380

A solution to the OFC detection problem has already been devised and implemented on the A380 aircraft, as presented in a paper by Goupil [1]. This approach made use of a non-linear analytically redundant model of an actuator to generate a residual signal. This was done separately for each of the aircraft's actuators. The residual was then used to check for oscillations by defining a positive and negative threshold value, and counting the number of times the residual crossed these thresholds over a certain time period. This method was termed oscillation counting. The residual signal was filtered into two frequency sub-bands before oscillation counting was performed to reduce the effect of noise. Overall, oscillation counting is a very computationally efficient approach to residual evaluation, and has found considerable success in the detection of OFCs on the A380 aircraft. Chapter 4 offers a more detailed description of the system.

2.1.1 Residual Generation Improvements

Many authors have since attempted to improve the oscillation counting system to detect smaller amplitude OFCs. Most of the effort has been to improve the estimates of the flight actuator dynamics to reduce the effects of modelling errors, lowering the noise level of the nominal residual, and consequently allowing for the detection of smaller amplitude OFCs. These approaches thus improve the residual generation stage, and revert back to the oscillation counting approach for residual evaluation.

Alwi and Edwards [6] manipulated the analytical model from Goupil's paper to estimate the true actuator rod speed, which was not available from direct measurements. An adaptive sliding mode Levant differentiator [7] was used to estimate the rod speed. This Levant differentiator allowed for the estimation of the rod speed using the position measurements, without the amplification of noise inherent with differentiation. By obtaining this speed estimate, Alwi and Edwards could reconstruct the oscillatory failure mathematically.

Other residual generation proposals include the use of non-linear Luenberger state observers [8], where the residual signal is fed back into the observer to improve its position estimates. Lavigne,

Zolghadri, Goupil *et al.* [9] proposed the use of extended Kalman filters to estimate the actuator's position, as well as to estimate certain time-varying parameters in the analytical model that are usually assumed to be constant.

2.1.2 Residual Evaluation Techniques

Although there is continued research into the improvement of the residual generation stage, proposals for improved residual evaluation techniques for the application of OFC detection have been relatively sparse.

Varga and Ossmann [10] investigated the detection of many different faults that occurred in flight actuators, including jamming, runaway, oscillatory failures, and loss of efficiency. The presence of a fault was detected using the Narendra signal evaluation scheme [11], which is a time-weighted evaluation of the energy of the residual signal. Once the presence of a fault was confirmed, the type of fault present was identified using different fault identification algorithms for each failure type. For the specific case of OFC detection, the identification algorithm estimated the power spectrum of the residual signal using the discrete Fourier transform (DFT). A recursive implementation [12] of the DFT was proposed here for efficient calculation of the power spectrum. OFCs were flagged if any frequency components crossed some threshold. The authors found that OFCs could be detected within the required detection time, but the method struggled to detect and identify oscillatory failures with higher frequencies.

Goupil returned to the OFC detection problem with a new residual evaluation technique [13], based on Wald's sequential probability ratio test (SPRT) [14]. This test assumes two hypotheses, each with a corresponding probability density function. The first hypothesis assumes that the residual contains no faults, while the second hypothesis assumes a failure exists. The likelihood ratio is the ratio of the probability that the residual matches the failure hypothesis over the probability that the residual contains no failures. If this ratio is above a certain threshold, a failure is inferred. This approach was tested using Gaussian distributions and Laplace distributions, and found that Laplace distributions provided the best results. Computationally, the SPRT was found to be very efficient, with a decrease of about 50 % in computational workload on the A380's flight control computer versus the oscillation counting approach. An additional advantage of this method is that the false alarm probability rate and missed detection probability rate can be used directly to determine relevant thresholds and parameters. The SPRT has similarly been applied to fault detection systems in [15] and [16].

2.2 Detection of Oscillating Control Loops

Oscillatory failures have also been known to affect control loops in large production plants, and fault detection schemes have been developed to address these issues.

2.2.1 Integrated Absolute Error

Hägglund [17] proposed a method to monitor control-loop performance and to detect oscillations in control loops that were caused by valve stiction, load disturbances, or badly tuned controllers. Hägglund made use of the control error signal as the "residual". The proposed detection algorithm made use of the integrated absolute error (IAE), which was calculated by integrating the control error between successive zero crossings. Nominal control-loop operation would typically result in small values of the IAE, and thus failures could be inferred when the IAE exceeded some threshold. Furthermore, oscillatory failures would show periodic threshold crossings. Therefore, a oscillatory failure could be inferred if a certain number of IAE threshold crossings occurred.

Forsman and Stattin [18] attempted to improve the robustness of the IAE oscillation detection scheme by deriving an index that measured the periodicity of the control error signal. This suggested that the ratio between successive IAE calculations, i.e. IAE_{i+1}/IAE_i, should be approximately 1 during an oscillatory failure. Similarly, the time between consecutive zero crossings (TBZC) should remain relatively constant. Forsman and Stattin therefore suggested a metric based on the successive IAE ratios and successive TBZC ratios that measured how "oscillative" the control error

signal was. Additionally, the proposed metric handled positive control errors and negative control errors separately, so that asymmetrical oscillations could be detected.

2.2.2 Autocovariance

Thornhill, Huang and Zhang [19] used the autocovariance function (ACF) of the control error signal to detect oscillations. The ACF was used as a way to filter the control error signal and attenuate noise that could cause spurious zero crossings. It was calculated by computing the inverse Fourier Transform of the power density spectrum (PDS), according to the Wiener-Khintchine Theorem [20]. The required PDS was calculated by simply squaring the discrete Fourier transform (DFT) of the control error signal. The resultant ACF was a sinusoidal signal with the same frequency as the failure, but with greatly attenuated noise. Frequency domain filters were also used to separate different frequency bands before calculating the ACF. These filters were implemented simply by setting the unwanted DFT frequency bins to zero.

With this new filtered signal, three parameters were used to identify the existence of oscillations: the period of the oscillation, the regularity of the oscillation, and the power of the oscillation. The mean and standard deviation of the periods between ACF zero crossings were calculated, and the ratio of the mean over the standard deviation gave an indication of the "regularity" of the oscillation. This regularity, along with the power of the oscillation, was used to decide whether a failure was present.

For this approach, it was recommended to use at least 10 zero crossing intervals to determine the regularity of the signal. The OFC detection problem unfortunately has strict detection time requirements, and this approach is not applicable. The use of the ACF may be able to improve the detectability of small OFCs, but the number of computations required for such a calculation (which requires a DFT, frequency-domain filtering, and an inverse-DFT) is simply too great.

2.2.3 Empirical Mode Decomposition

Srinivasan, Rengaswamy and Miller [21] also performed oscillation detection using the IAE method. The effectiveness of the IAE was improved by using Empirical Mode Decomposition (EMD) [22], a time-domain iterative method of decomposing a signal into several Intrinsic Mode Functions (IMFs) in descending order of frequency. This EMD process was used to obtain the mean of the signal and remove it from the signal. With this zero-mean error signal, the zero crossings were identified. Further filtering was performed by integrating the error signal to remove high frequency noise and spurious zero crossings. The number of spurious crossings were also minimised by clustering crossings that were only a few sample instants apart.

2.2.4 Comparison of Oscillation Detection Methods

Odgaard and Trangbaek [23] performed a study of some of the common methods used for oscillation detection. These methods included:

- Zero crossing regularity (ZCR): The time between zero crossings are compared to the average time between crossings. Small deviations from the average time imply the existence of regular, evenly spaced crossings, and thus an oscillatory failure can be inferred.
- Discrete Fourier transform (DFT): Oscillations are detected by comparing the amplitude of the largest frequency component to the total energy of the low-frequency signal content.
- Autocovariance function (ACF): The ACF of the measured signal is calculated to remove the effects of white noise. The amplitude of the first two peaks of the ACF is found using the ACF's derivative, and these amplitudes are compared.
- Autocovariance zero crossing regularity (ACZR): This approach is very similar to the ZCR, but checks the zero crossings of the autocovariance function instead of the originally measured signal.

These methods were tested by Srinivasan, Rengaswamy and Miller [21] based on whether oscillations in the measurements could be detected reliably, how sensitive each method is to transients, and how robust each method is to changes in window length.

The study found that the ACF method could not reliably detect oscillations in the given test data. The ZCR and AZCR methods could detect oscillations, albeit with significant delays. However, they were dependent on window length and were too responsive to transients. Finally, the DFT approach provided good oscillation detection performance, and could reject transients, but was highly dependent on its window length. The authors, however, did not use any system models for residual generation, and suggested that using a system model could alleviate the effect of transients that they found in their data sets.

2.3 Additional Oscillatory Failure Detection Investigations

The previous two sections have focussed heavily on two applications, namely oscillations in flight actuators and oscillations in control loops. However, there has been a considerable amount of oscillatory failure detection systems outside these two cases.

2.3.1 Vibration Monitoring using Wavelet Transforms

Kar and Mohanty [24] investigated the use of the discrete Wavelet transform in a fault-monitoring system that checked for high-frequency vibration transients, and low-frequency signals emanating from the electric power grid frequency. Due to the high and low frequency nature of the problem, an implementation of the Wavelet transform known as the multi-resolution Fourier transform was found to provide the best results.

2.3.2 Fault Detection using Deterministic Learning

Chen, Wang and Hill [25] tackled the issue of modelling uncertainties in analytically redundant fault detection systems by applying neural networks to learn the unknown system dynamics, as well as the dynamics of oscillatory failures. These neural networks were used to train a bank of estimators with assumptions based on nominal behaviour and oscillatory failure behaviour. Faults were then detected based on the smallest residual principle, where a fault is inferred when the average residual of a fault estimator is smaller than that of a fault-free estimator.

2.4 Sinusoid Detection and Estimation Algorithms

This final section briefly investigates sinusoid detection and estimation techniques.

2.4.1 MUSIC and ESPRIT Algorithms

A tutorial paper by Guo and Bodson [26] approached the problem of adaptive rejection of oscillations in control systems by using frequency estimation techniques such as MUSIC and ESPRIT.

Multiple signal classification (MUSIC) [27] was initially proposed for the estimation of signal parameters measured with antenna arrays, with the most important parameters being the direction of arrival and high-resolution frequencies. These signals often consist of multiple sinusoids at different frequencies, along with white noise. However, the MUSIC algorithm is a computationally expensive algorithm that requires the use of either eigen-decomposition or singular value decomposition.

A more computationally efficient approach to the MUSIC algorithm is the ESPRIT [28] algorithm, where ESPRIT is the acronym for “Estimation of Signal Parameters via Rotational Invariance Techniques”. However, this approach assumes the existence of multiple sensors arranged in matched pairs, making it potentially unsuitable for the application of OFC detection without additional sensors.

2.4.2 Estimation of Sinusoidal Signals

Finally, many other mathematical techniques for the estimation of signal parameters have been suggested. This includes least squares estimators [29], adaptive notch filters [30], amplitude estimators and observers [31], and periodograms [32]. Many of these techniques could be applicable to the OFC detection problem. However, the application of these techniques is left for future work.

2.5 Summary and Conclusions

Oscillation counting and the integrated absolute error (IAE) approaches have been successfully implemented for the purposes of oscillation detection. However, with regards to OFCs in flight actuators, most authors have attempted to improve the residual generation stage, and return to the oscillation counting approach for residual evaluation. Because the residual generation stage has already been thoroughly researched, the focus of this thesis is placed on the residual evaluation techniques.

The oscillation counting approach is currently in service on the A380, and is investigated in Chapter 4 as a baseline against which other evaluation techniques should be compared.

The IAE is another approach that can offer computationally efficient OFC detection, and is further investigated in Chapter 5. However, surveys have shown that many of the improvements made to the IAE approach, such as regularity checks and autocovariance calculations, struggle to provide adequate detection time, and are more applicable to offline processing.

Finally, some potentially interesting gaps have been found in the research into OFC detection. Very few applications of the discrete Fourier transform for OFC detection have been found, and it is therefore investigated further in Chapter 6. Furthermore, signal-processing techniques used in telecommunications, a field built around the idea of signal detection, have not seen many applications in fault detection systems. Chapter 7 investigates the potential application of the phase-locked loop for OFC detection.

Chapter 3

Conceptualisation and Modelling

3.1 Introduction

This chapter provides an overview of the OFC detection testbench. This is a simulation framework that was created for training and testing of different OFC detection techniques. The simulation framework contains models for the physical actuator, the analytically redundant actuator, the oscillatory failures (both liquid and solid failures), the flight control system, and the aircraft longitudinal dynamics. The simulation framework was designed so that the actuator parameters and the oscillatory failure parameters (amplitude, frequency, type, and injection point) can be varied randomly to produce training and testing data that is sufficiently rich.

An overview of the training and testing procedure that will be used to benchmark all the detection methods is given in Section 3.3. Finally, illustrative examples of flight data sets are shown to provide insight into the different types of oscillatory failures. These examples show how the OFCs manifest in the residual, and how this could impact their detectability.

3.2 Simulation Model

A full model of the longitudinal dynamics of an A380 aircraft is used to generate flight data that serves as various test cases to benchmark residual evaluation techniques. This framework was developed in collaboration with Airbus, and assumes the aircraft is in cruise flight at 38 000 ft, flying at a speed of Mach 0.8. The purpose of the model is to describe the aircraft's reaction to an OFC, as well as additional unexpected behaviour that may stress or complicate the detection of oscillatory failures. This provides authentic and representative test cases.

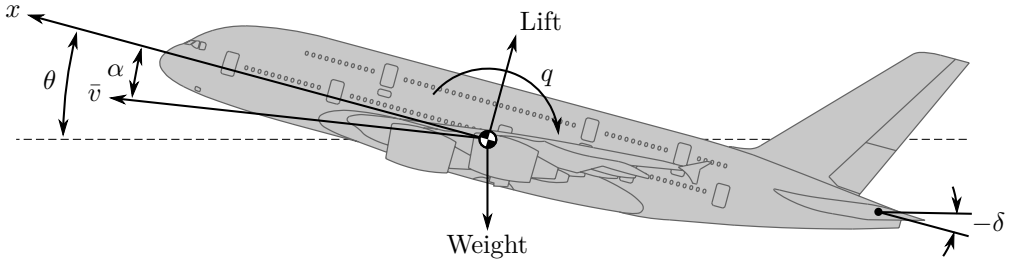

Figure 3.1: Longitudinal Aircraft Parameters. Image adapted from [3] and [33].

Figure 3.1 shows aircraft parameters that are modelled in the simulation. These parameters are described in Table 3.1 [34].

Table 3.1: List of Flight Parameters

Parameter	Description
x, y, z	Aircraft Body Axis
q	Pitch Rate
θ	Pitch Angle
α	Angle of Attack
$\bar{v}$	Velocity Magnitude
δ	Elevator Control Surface Deflection

All these parameters are assumed to be deviations from their trim states. An additional parameter not shown above is the load factor, N_z. This is defined as the ratio of the aircraft's lift over its weight. This serves as the pilot command to the aircraft control system. The trim load factor is 1 g, and the load factor deviation is 0 g for straight and level flight, and rarely exceeds 1 g.

The most important component in the simulation is the control surface actuator, which is assumed to control the deflection of an elevator for longitudinal control. The actuator is essentially modelled twice. The first is the "real" actuator that is affected by aerodynamic forces, parameter variation as a result of manufacturing and environmental differences, and noisy sensors. The second model is the analytical model that is simulated by the flight control computer's monitoring process. These model differences represent modelling inaccuracies that are found in practice, and cause non-zero biases in the fault-free residual that makes fault detection more challenging. Furthermore, the structure of the FCC shown in Figure 1.2 is assumed here, where the measured actuator rod position is used exclusively for control, while the measured control surface deflection is used exclusively for fault monitoring.

The following sections describe the levels of the simulator, starting with a high-level overview of the aircraft model, and leading to the low-level actuator models. All flight controllers are assumed to be continuous-time systems, while the monitoring systems are considered to be discrete-time.

3.2.1 Full Aircraft Model

Figure 3.2 shows the high-level overview of the simulator. It contains a load factor controller, a model of the control surface actuator, an analytically redundant actuator model, and a plant of the longitudinal dynamics of the A380.

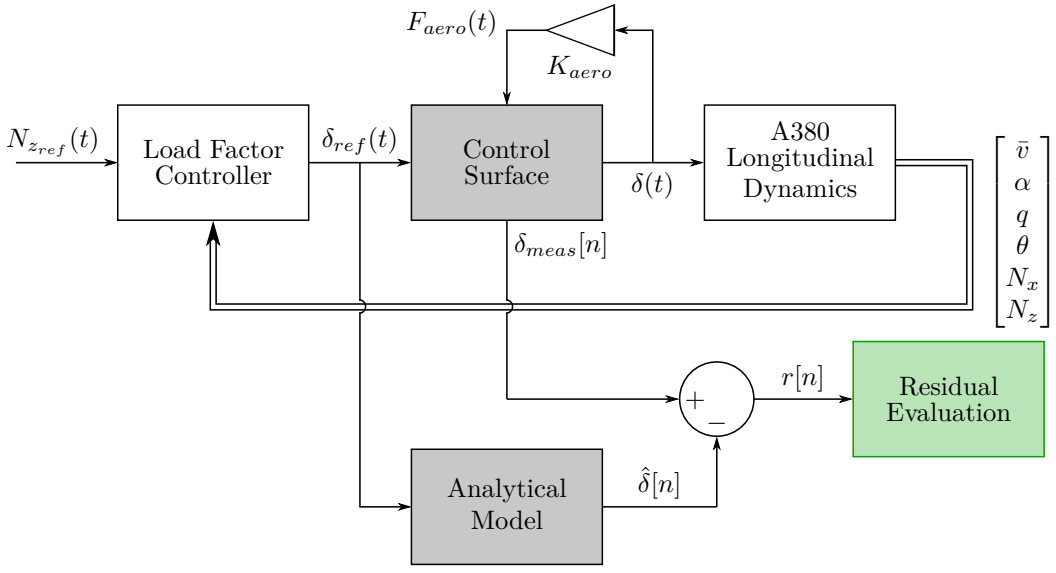

Figure 3.2: Full A380 Longitudinal Model

The model accepts a reference load factor command N_{zref} provided by the pilot. The load factor controller attempts to track the reference load factor command by manipulating the control surfaces. It delivers a control surface deflection command, $\delta_{ref}(t)$, to the control surfaces based on the load factor command and the current aircraft state. This commanded deflection is accepted by the hydraulic actuator control system, which outputs two signals, which are the true control surface deflection $\delta(t)$ and the measured control surface deflection $\delta_{meas}[n]$. The true deflection $\delta(t)$ represents the real deflection of the actuator, and is fed to the A380 longitudinal dynamics. The measured control surface deflection is provided by a sensor that adds sensor noise to the signal, and is used exclusively for fault monitoring.

The analytical model accepts the same reference control surface angle $\delta_{ref}(t)$ as the real control surface system, and provides an estimate of the deflection of the control surface, $\hat{\delta}[n]$. The difference between the measurement, $\delta_{meas}[n]$, and the estimate serves as the residual $r[n]$ for fault detection:

$$r[n] = \delta_{meas}[n] - \hat{\delta}[n] \tag{3.2.1}$$

To generate data for training and testing purposes, the model requires load factor commands that are representative of real pilot commands. Since this model describes the dynamics of straight and level cruise flight, the total reference load factor is chosen not to exceed 0.15 g. This is chosen with the assumption that during cruise flight, the load factor of the aircraft is usually within 10 % to 15 % of the maximum load factor deviation of 1 g. To generate load factor commands that are representative of pilot actions, white noise filtered with a lowpass filter is used as the load factor command N_{zref}. The bandwidth of the filter is chosen such that 99 % of the power of the resulting control surface command $\delta_{ref}(t)$ exists in the 0 to 0.3 Hz range of the power density spectrum. These assumptions are based on sample flight data provided by Airbus.

$F_{aero}(t)$ represents the aerodynamic forces experienced by the control surface. For this model, it is assumed that the magnitude of the aerodynamic force is a linear function of the absolute control surface deflection. When the control surface maintains 0° deflection, F_{aero} applies 0 N of force. When at a non-zero deflection angle, the magnitude of the aerodynamic force is linearly proportional to the deflection, and acts in the opposite direction of the deflection. In other words, F_{aero} always attempts to force the control surface back to 0°. The aerodynamic force therefore opposes movement away from the trim deflection state, and assists movements back towards the trim state. The gain factor K_{aero} is chosen such that the maximum aerodynamic force does not exceed one tenth of the maximum force applied to the control surface by the hydraulic actuator.

3.2.2 Control Surface Model

Figure 3.3 shows the control surface system, which contains a non-linear model of the hydraulic actuator. The system accepts the deflection command $\delta_{ref}(t)$ from the load factor controller, and converts this deflection command to a rod position command $u(t)$. The actuator responds to the command by moving its piston to a new position $p(t)$ which rotates the control surface to a new deflection angle $\delta(t)$.

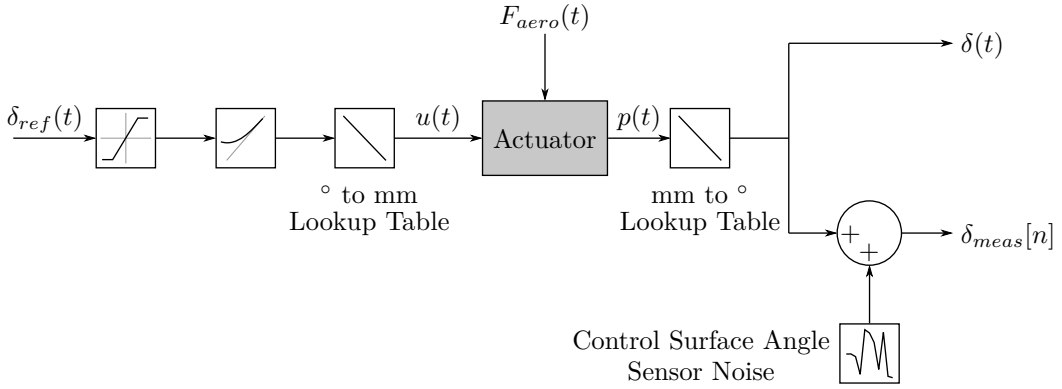

Figure 3.3: Control Surface Model

Note the non-linear saturation and rate limiter blocks at the system input. These limit the commands so that they do not exceed the physical limitations of the control surfaces. Specifically, the control surface deflection angle is limited to $[-30°, 15°]$, and the control surface angular rate is limited to $\pm 30\,°/s$. The lookup tables convert the angle command to a rod position command in mm, and vice versa.

The control surface model has two output signals. The first signal, $\delta(t)$, represents the true deflection of the control surface, and directly affects the longitudinal dynamics of the aircraft. The second signal, $\delta_{meas}(t)$, is the measurement signal from a control surface angle sensor. This measurement is exclusively used for monitoring purposes, as suggested by Figure 1.2, and is corrupted by sensor noise.

3.2.3 Actuator Model

Figure 3.4 shows the model of the control surface actuator. Given the desired reference position $u(t)$ commanded by the flight control computer, the actuator uses a proportional feedback controller to control the linear position of the piston accurately. This position $p(t)$ affects the deflection angle of the control surface $\delta(t)$.

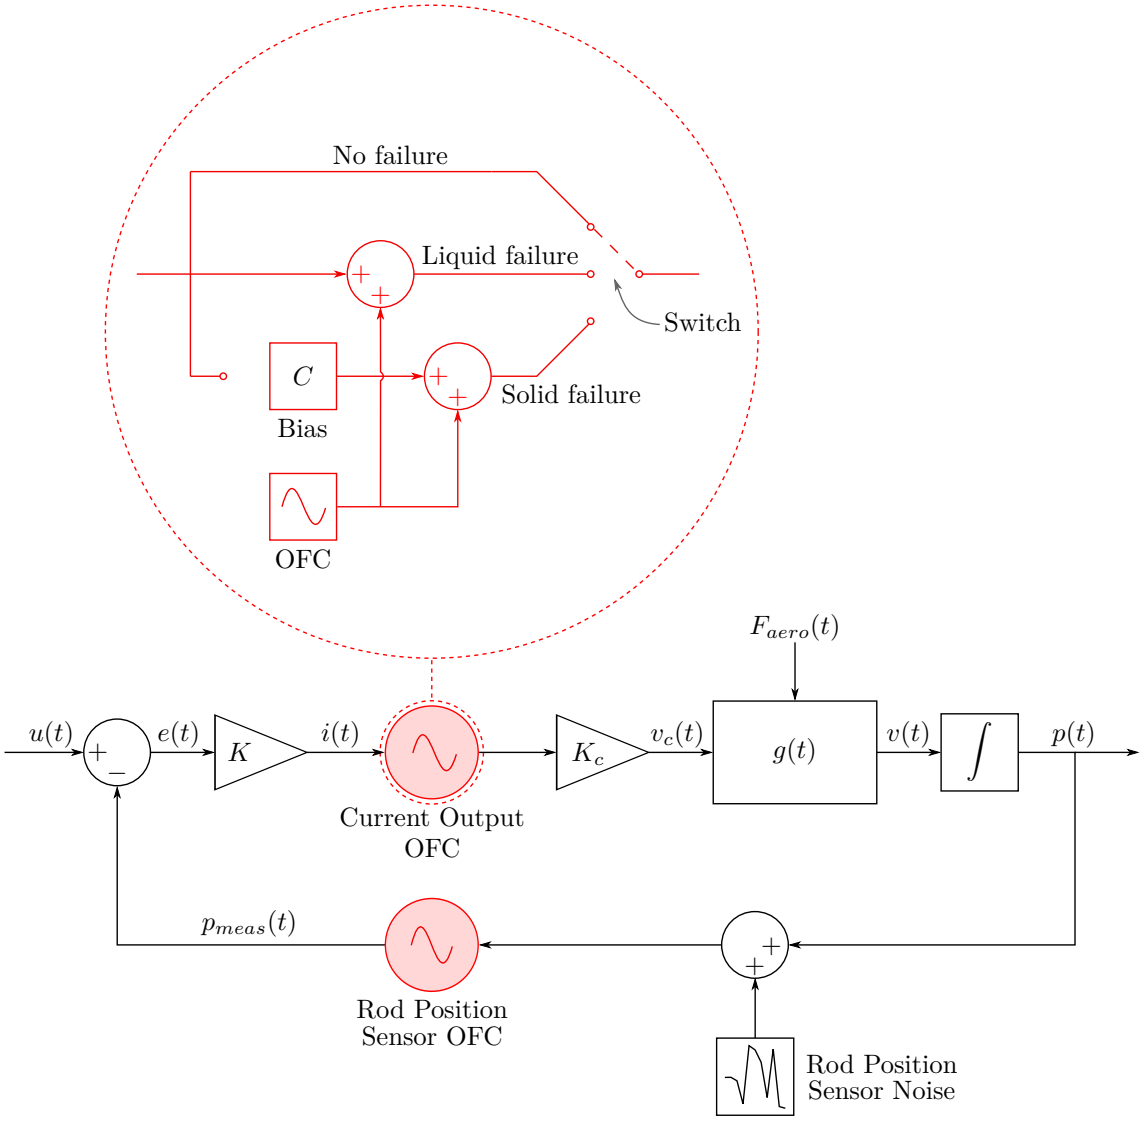

Figure 3.4: Hydraulic Actuator Model

The actuator makes use of a proportional feedback controller for accurate position control. The error signal $e(t)$ is the difference between the actuator reference position $u(t)$ and the measured position $p_{meas}(t)$. The error is converted to a command current $i(t)$ which ultimately controls the velocity of the actuator's piston $v(t)$. The integrator converts this velocity to position. The system includes a position sensor that provides the measurement $p_{meas}(t)$. The sensor adds band-limited white noise to the measurement. The output of the system, $p(t)$, is tapped off before the sensor noise block, and represents the true position of the actuator.

This actuator system is derived from the following mathematical model that calculates the linear actuator rod speed $v(t)$ [1]:

$$v(t) = v_c(t) \cdot \sqrt{\frac{\Delta P(t) - [F_{aero}(t) + F_{damping}(t)]/S}{\Delta P_{ref}}} \tag{3.2.2}$$

This equation relates the speed of the actuator rod to the commanded speed v_c, the hydraulic pressure delivered to the actuator ΔP, and the external forces, F_{aero} and $F_{damping}$, which result from aerodynamic disturbances and the damping of the passive actuator respectively. The following variables and parameters are therefore used in Equation (3.2.2):

- $v_c(t)$ is the commanded rod speed from the actuator controller, measured in mm/s.
- $\Delta P(t)$ is the hydraulic pressure delivered to the actuator, measured in bar.
- ΔP_{ref} is the differential pressure corresponding to the maximum rod speed.
- $F_{aero}(t)$ represents the aerodynamic forces applied on the control surface. It is assumed to have an average value of 0 N.
- $F_{damping}(t)$ represents the damping effect of the passive actuator.
- S is the actuator piston surface area in mm^2.

Equation (3.2.2) can be interpreted as equating the speed of the actuator rod $v(t)$ to the commanded speed $v_c(t)$, weighted by the effects of the hydraulic pressure and the external forces. These non-linear effects are represented in Figure 3.4 as $g(t)$. Increased hydraulic pressure $\Delta P(t)$ increases the weighting, while the external forces counteract the desired motion of the actuator.

The commanded speed is provided by the feedback controller, and is mathematically described as:

$$v_c(t) = K_c \cdot K \cdot (u(t) - p_{meas}(t)) \tag{3.2.3}$$
$$= K_c \cdot i(t) \tag{3.2.4}$$

Here, K is the servo-control gain which generates a servo current command in mA:

$$i(t) = K \cdot (u(t) - p_{meas}(t)) \tag{3.2.5}$$

This current is delivered to the actuator, and translates to rod speed through the conversion factor K_c.

The final parameter of note is $F_{damping}$, the external force caused by the damping effect of the passive actuator. The damping force is a function of the actuator's current velocity:

$$F_{damping} = K_d \cdot v^2(t), \tag{3.2.6}$$

where K_d is the damping coefficient of the passive actuator.

Typical values for the actuator parameters were provided by Airbus and are summarised in Table 3.2.

Table 3.2: List of Flight Actuator Parameter Values

Parameter	Value	Unit
ΔP_{ref}	335	bar
F_{aero}	0	N
S	5800	mm^2
K	0.6	mA/mm
ΔP	$160 - 300$	bar
K_d	$6.8 - 10$	$N/(mm/s)^2$

The values listed for ΔP and K_d represent the range of possible values that these parameters could assume, and these parameters change based on environmental conditions. K_c is not listed above, since it is implemented as a lookup table.

An oscillatory failure case is modelled as a pure sinusoid that is induced within the control loop of the actuator. An OFC can exist at the command current or at the position sensor. These two OFC injection points are indicated in Figure 3.4. Distinction is made between liquid and solid oscillatory failures. Liquid failures are additive faults, where an oscillatory signal is superimposed on the existing signals in the control loop, and the actuator can still respond to reference commands. Solid failures completely replace the existing signals in the loop, and the flight control computer no longer has any control over the actuator. In literature, the location of an oscillatory failure is generalised by assuming that it affects the position command input. Liquid and solid failures are then mathematically modelled as [35]:

$$u(t) = \begin{cases} u_0(t) & \text{for nominal behaviour} \\ u_0(t) + f_{liq}(t) & \text{for liquid failures} \\ u_0(t) + f_{sol}(t) & \text{for solid failures} \end{cases} \tag{3.2.7}$$

where $u_0(t)$ is the fault-free control signal provided by the flight control computer, and the liquid and solid failures are defined as

$$f_{liq}(t) = A\cos(\omega t) \tag{3.2.8}$$

$$f_{sol}(t) = A\cos(\omega t) - u_0(t), \tag{3.2.9}$$

where A and ω are the amplitude and frequency of the oscillatory failure signal. These OFCs are considered to be sinusoids with frequency uniformly distributed over the 1 to 10 Hz frequency range. OFCs with frequencies higher than 10 Hz have little effect on the control surface motion due to the lowpass characteristics of the actuator.

Therefore, these injection points have three failure states:

1. No failure: The command current and position sensor measurements are unchanged.
2. Liquid failure: A sinusoidal signal is added to either the command current signal or the position measurement signal.
3. Solid failure: The command current signal or position measurement signal is completely replaced with a sinusoid at some offset.

The selection of which failure is present is changed by manipulating the switch at some defined failure time t_f. No more than one OFC is enabled at any given time.

3.2.4 Analytically Redundant Model

The analytically redundant model simulated in the FCC's monitoring channel has a similar structure to the real actuator model, but is modelled as a discrete system, and assumes fault-free operation. This is shown in Figures 3.5 and 3.6.

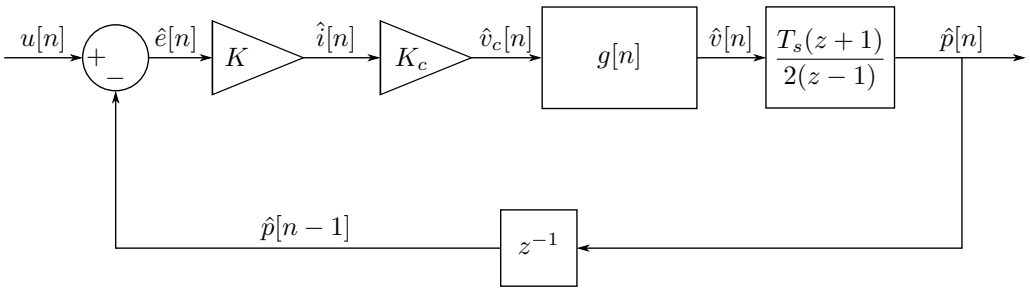

Figure 3.5: Hydraulic Actuator Analytical Model

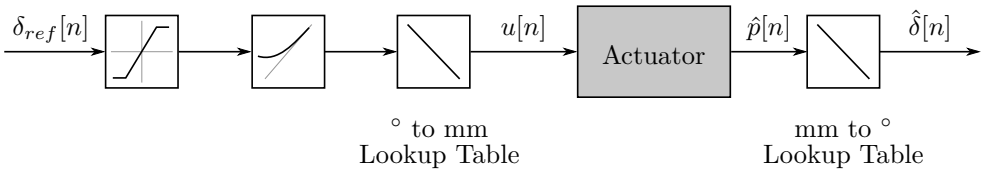

Figure 3.6: Control Surface Analytical Model

Mathematically, the analytically redundant actuator model in Figure 3.5 is expressed as:

$$\hat{v}[n] = K_c K(u[n] - \hat{p}[n-1])\sqrt{\frac{\Delta P - K_d \cdot \hat{v}^2[n]/S}{\Delta P_{ref}}} \tag{3.2.10}$$

Equation (3.2.10) is a discretisation of Equation (3.2.2), and has been simplified by assuming that ΔP is constant, and F_{aero} is zero. The estimated position $\hat{p}[n]$ is calculated from the estimated speed $\hat{v}[n]$ using trapezoidal integration.

The above model is implemented in Simulink. However, this model cannot be used in its current form when implemented on an embedded computer, since $\hat{v}[n]$ appears on both sides of Equation (3.2.10). The equation has to be manipulated such that the rod speed can be computed. After rearranging the equation, the analytically redundant model becomes

$$\hat{v}[n] = K_c K(u[n] - \hat{p}[n-1])\sqrt{\frac{\Delta P}{\Delta P_{ref} + K_d[K_c K(u[n] - \hat{p}[n-1])]^2/S}}. \tag{3.2.11}$$

3.3 Procedure for Training and Testing

With the A380 simulator in place, various residual evaluation strategies can be tested. Each evaluation method only receives the residual signal, $r[n]$. To test all the methods, two large data sets are required, namely a training data set and a testing data set.

The training data set contains no faults and is used to "train" the thresholds for each residual evaluation method. These thresholds represent the envelopes of nominal behaviour in each evaluation method. Essentially, the training process attempts to train each detection method by finding the largest or worst case residual that is not associated with failures. This is done by generating an enormous amount of nominal data, and forcing the various thresholds as low as possible without flagging false positives in any of the training datasets. By setting the thresholds as low as they can possibly be, they are given the best possible chances to detect very small OFC, and this ensures that all detection methods are tested fairly. For additional robustness, the defined thresholds can be scaled by some safety margin, at the obvious expense of minimum detectable amplitudes.

Two types of load factor command signals are used for training data. The first is the randomised load factor commands, limited to 0.15 g, that are representative of real load factor commands

during cruising flight. The second signal type is the chirp signal, which serves to stress the aircraft dynamics across all frequencies between 0 and 10 Hz. It is highly unlikely that a chirp signal will occur during actual flight, but this serves to stress the detection methods thoroughly, since the oscillations of the chirp tend to propagate to the residual signal.

Once the thresholds are trained using the training data, the methods are tested using test data. The test data exclusively contains the filtered white noise load factor commands, and contain OFCs of various amplitudes, frequencies, types (liquid or solid) and injection points (command current or position sensor). All four detection methods are tested with the same test sets, and the results are consolidated into the following performance metrics:

1. Smallest amplitude OFC detected within 3 cycles.
2. Smallest amplitude OFC detected within 6 cycles.
3. Average detection time.
4. Number of false detections.

The first metric listed above is the most important, and represents the experimentally-determined smallest amplitude OFC that is guaranteed to be detected within the required detection time for each detection method. The overall best method is the one with the lowest amplitude. The second metric is the absolute smallest OFC that was detected by each method. This represents the potential performance of each method if the detection time requirements were relaxed. Note that if an OFC is not detected within 6 cycles, it is assumed to have gone undetected. This avoids the results becoming skewed by small OFCs that happen to get detected after, for example, 50 cycles. The average detection time metric simply compares the detection time performance of each method. Finally, the number of false alarms detected by each method provides an indication of the robustness of the method. With sufficiently rigorous training, however, no false alarms should be present. The entire training-testing process is shown in Figure 3.7.

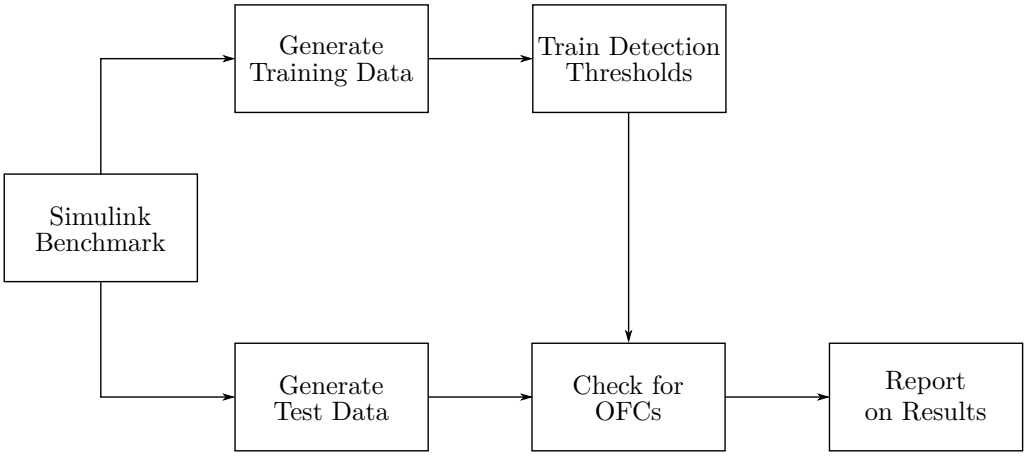

Figure 3.7: Overview of Training and Testing Procedure

Each data set consists of 30 seconds of simulation time. The seeds for all pseudo-random number generators used to represent noise sources, including measurement noise and the load factor command, are changed with each test. Furthermore, the parameters K_d and ΔP of the actual actuator are randomised uniformly between the bounds given in Table 3.2, as these parameters are dependent on external conditions such as fluid temperature. For the analytically redundant model, the parameters K_d and ΔP are set to their nominal values.

It should be noted that the results obtained from the Monte Carlo training-testing procedure are influenced by the choices for the nominal values and ranges of the actuator parameters, and by the accuracy of the aircraft model used in the simulation model. The absolute performance of all the

OFC detection techniques will therefore vary based on the uncertainty in the actuator modelling, the amount of sensor noise, and the modelled response of the aircraft. The results should therefore be used to compare the relative performance of the different OFC detection techniques, and not as an indication of the absolute performance of the techniques.

3.4 Example Data Sets

This section investigates the different failure cases that could occur, how the control surface reacts to each failure, and how their respective residual signals could potentially look. For illustrative purposes, 30 seconds of flight data are shown for each case, with the oscillatory failure occurring at 15 seconds.

3.4.1 Nominal Flight Data Set

In Figure 3.8, the default, nominal flight case is shown where no failure is present, and filtered white noise is used as the load factor commands.

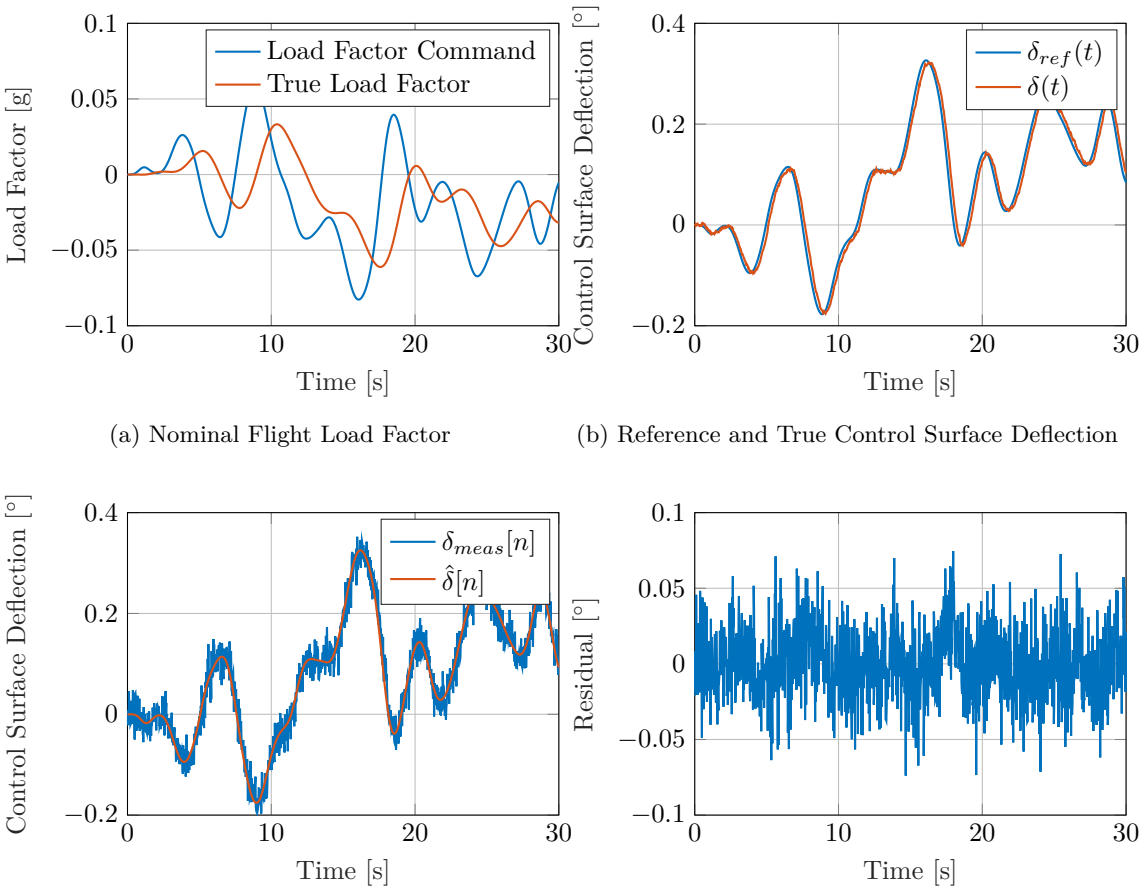

(a) Nominal Flight Load Factor

(b) Reference and True Control Surface Deflection

(c) Measured and Estimated Deflection

(d) Residual

Figure 3.8: Nominal Flight Data

Figure 3.8a shows the commanded load factor and the true load factor. Figure 3.8b shows the commanded control surface deflection, and the true deflection. Figure 3.8c shows the deflection as measured by the noisy sensor on the control surface, along with the estimated deflection provided by the analytical model. These two signals are used to generate the residual, pictured in Figure 3.8d. Here, the residual primarily consists of measurement noise from the control surface sensor, but also

contains some additional low frequency content due to the simplifications and assumptions made in the analytical model.

3.4.2 Nominal Training Data Set with Chirp Signal

Figure 3.9 shows a data set where the load factor command is a chirp signal of amplitude 0.15 g that sweeps the frequency between 0 and 10 Hz. This data set is an example of a no-fault scenario that stresses the OFC detection methods in terms of their robustness towards false alarms. Standard practice employed by Airbus is to use such a training case to set the thresholds of the fault detection methods to ensure robustness, although at the expense of a higher minimum detectable OFC amplitude.

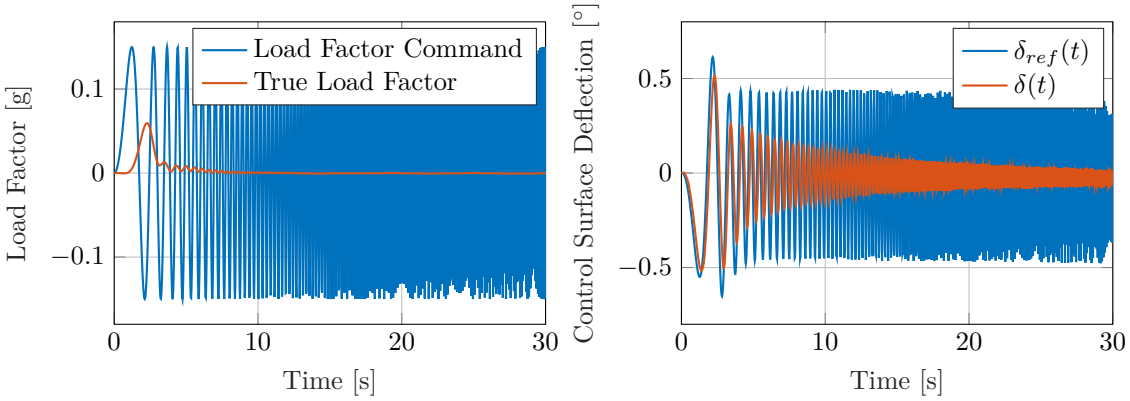

(a) Chirp Flight Load Factor

(b) Reference and True Control Surface Deflection

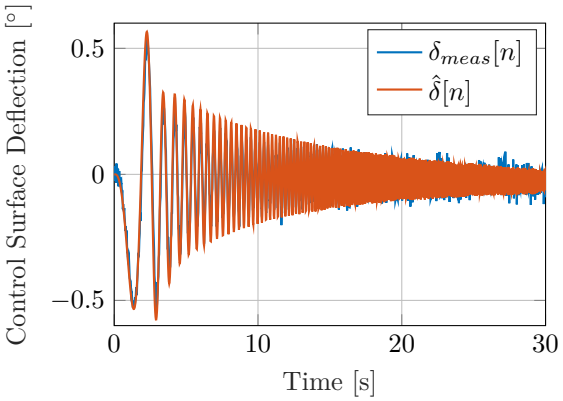

(c) Measured and Estimated Deflection

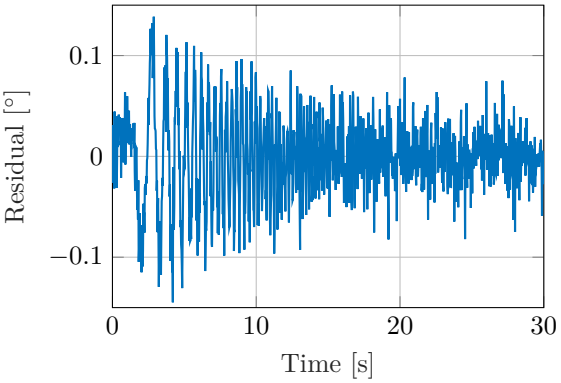

(d) Residual

Figure 3.9: Nominal Flight Data with Chirp Load Factor Command

Figure 3.9a shows the commanded load factor, which is a chirp signal, and the true load factor response. The true load factor struggles to follow the given command due to the rejection of high-frequency signals by the aircraft's longitudinal model. As can be seen in Figures 3.9b and 3.9c, the amplitude of the oscillation at the control surface decreases as frequency increases, showing that the flight control system and actuator dynamics tend to reject high frequencies. The resultant residual in Figure 3.9d shows significant oscillation over the first 10 seconds, which could quite likely cause false alarms. Once again, the non-zero residual signal results from the differences between the true actuator and the simulated actuator due to modelling uncertainty, and the sensor noise. The modelling uncertainty originates from the fact that the random parameter values for the true actuator differ from the nominal parameter values used for the simulated actuator.

3.4.3 Command Current Liquid OFC

Figure 3.10 shows the first example of an oscillatory failure case, located at the current command $i(t)$ of the actuator. Here, the induced OFC has an amplitude of 1 mA, at a frequency of 1 Hz.

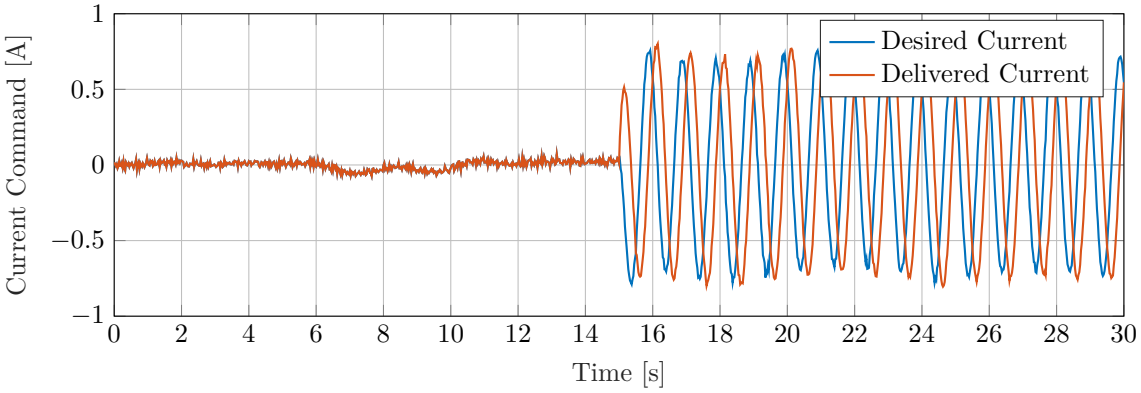

(a) Desired and Delivered Actuator Current

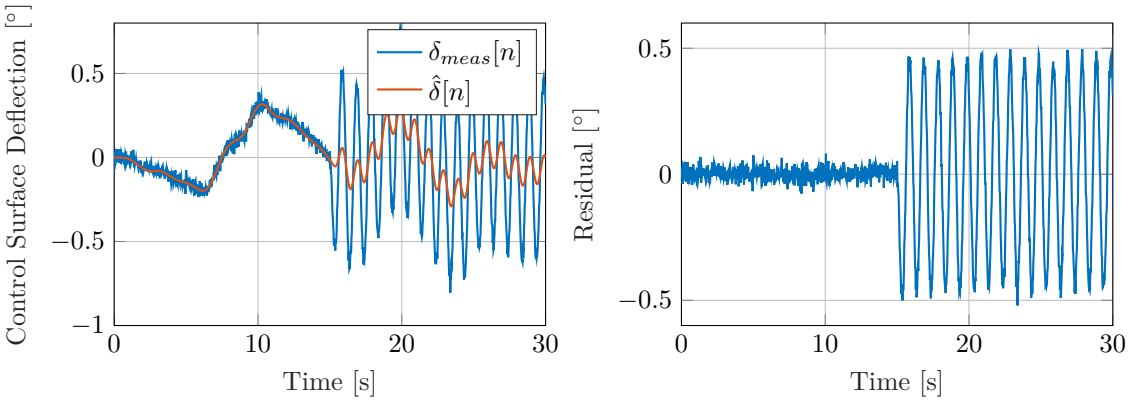

(b) Measured and Estimated Deflection

(c) Residual

Figure 3.10: Liquid OFC at the Actuator Command Current

Figure 3.10a compares the desired command current computed by the flight control computer, before the OFC injection point, to the current delivered to the actuator after the OFC has been applied. The desired current shows that the actuator control system attempts to reduce the effect of the OFC by delivering opposing commands. This reduces the amplitude of the OFC from 1 mA to about 0.75 mA.

The corrupted command current controls the velocity of the actuator's rod, causing it, and consequently the control surface, to oscillate. Since this is a liquid failure, the oscillations are added to the current command. The control surface still responds to the load factor command, as seen in Figure 3.10b. The resulting residual signal consists primarily of a component originating from the oscillatory failure. It is also noted that the 0.75 mA amplitude translates to an amplitude of about 0.5° at the control surface.

3.4.4 Rod Sensor Liquid OFC

Figure 3.11 shows an example of a 1 mm liquid OFC with a frequency of 1 Hz at the actuator's rod position sensor.

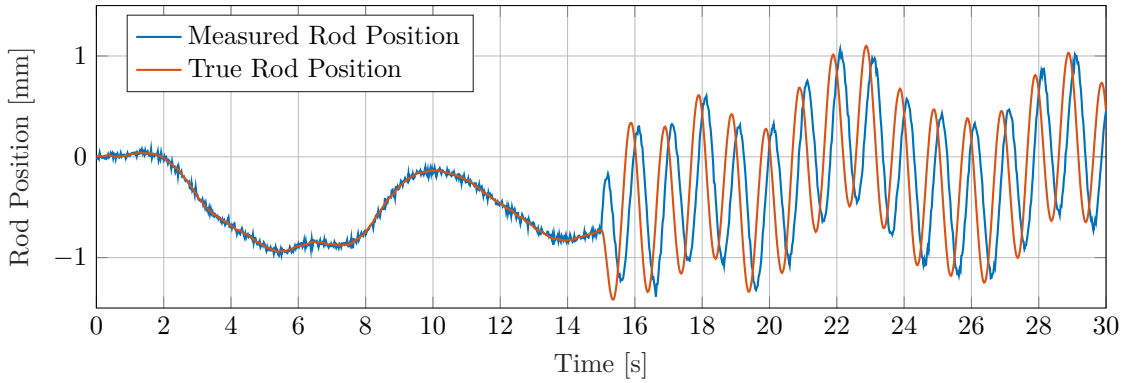

(a) Measured and True Rod Position

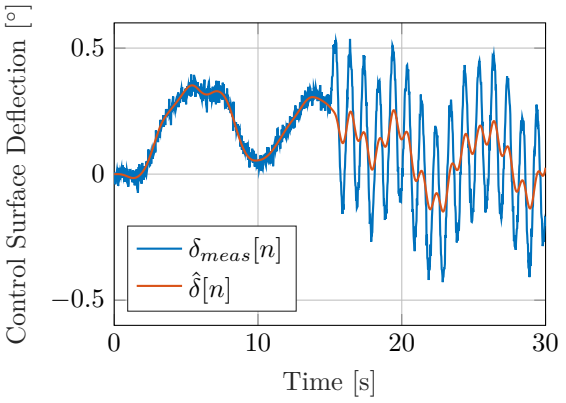

(b) Measured and Estimated Deflection

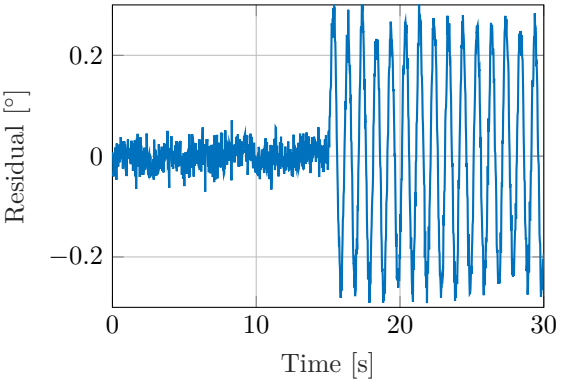

(c) Residual

Figure 3.11: Liquid Failure at the Rod Position Sensor

Similarly to the liquid failure at the command current, the control system attempts to reject this unwanted oscillation. In this case, however, the OFC at the sensor does not directly affect the control surface deflection. Instead, the actuator control system receives false measurements, and attempts to correct this by providing opposing commands to the actuator, causing the rod to oscillate. This reaction to the faulty measurements causes the control surface to oscillate.

Once again, it can be seen that the control surface still reacts to load factor commands, and the oscillations are superimposed on the desired deflection.

3.4.5 Command Current Solid Failure

Shown here is the first example of a solid failure. This is an OFC of 1 mA at the command current, with a frequency of 1 Hz.

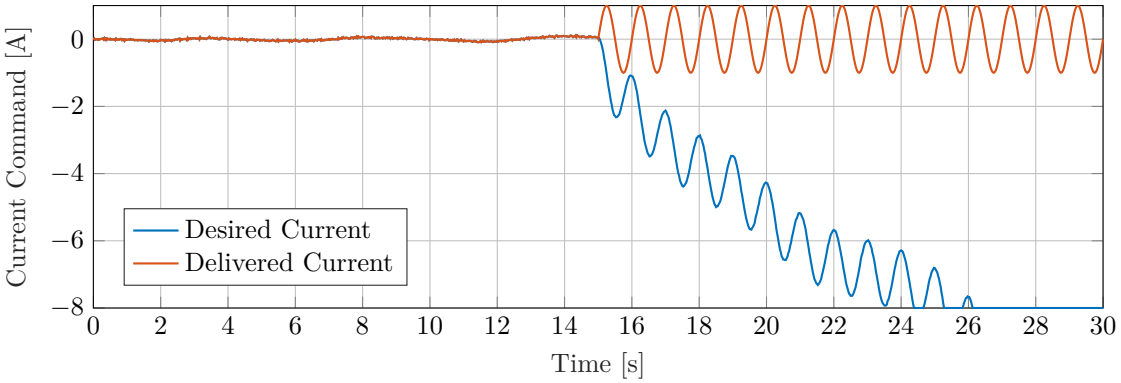

(a) Desired and Delivered Actuator Current

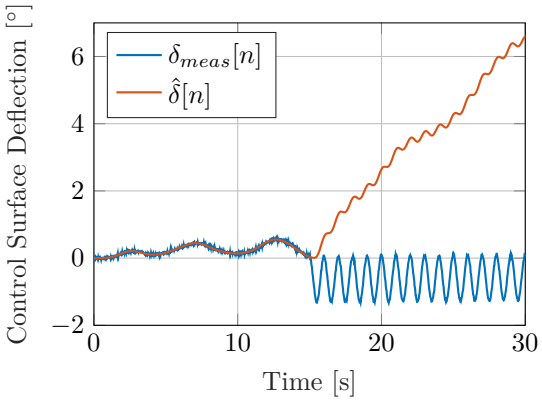

(b) Measured and Estimated CS Deflection

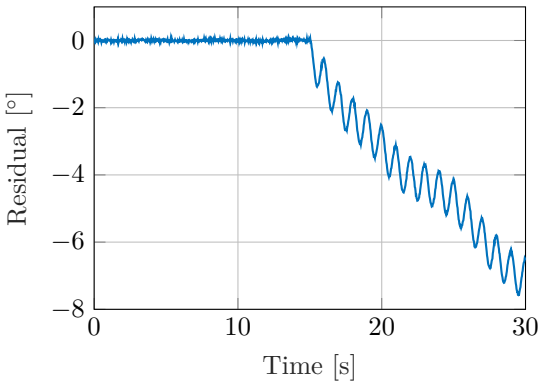

(c) Residual

Figure 3.12: Solid Failure at the Commanded Current

The failure is shown in Figure 3.12a. The OFC has no bias offset, and thus the oscillation occurs around 0 mA. The failure completely ignores the desired command current, and replaces it with a pure sinusoid. As a result, the actuator does not react to any commands given to it, and the control surface begins to oscillate in place, as shown in Figure 3.12b. To compensate for this, the aircraft commands ever increasing control surface deflections, which are completely ignored. The result is that the residual shows a large offset as a result of the subtraction of the estimated deflection from the measured deflection.

However, this solid failure case is the only solid failure case that is considered to be detectable by any OFC detection methods. Other solid failure cases are shown in the next section.

3.4.6 Current Command Solid Failure with Non-zero Bias

Shown here is a solid failure at the command current, with a non-zero offset. The results shown here are also applicable to *any* solid failure that occurs at the rod position sensor.

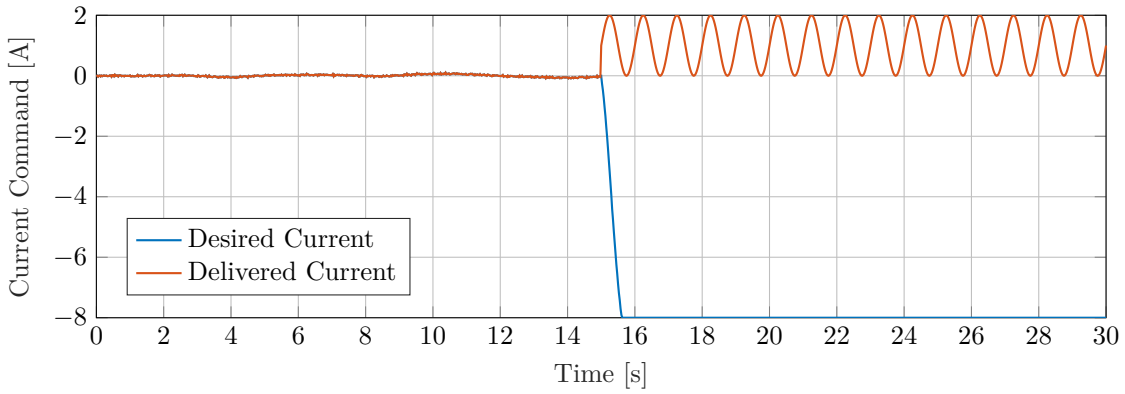

(a) Desired and Delivered Actuator Current

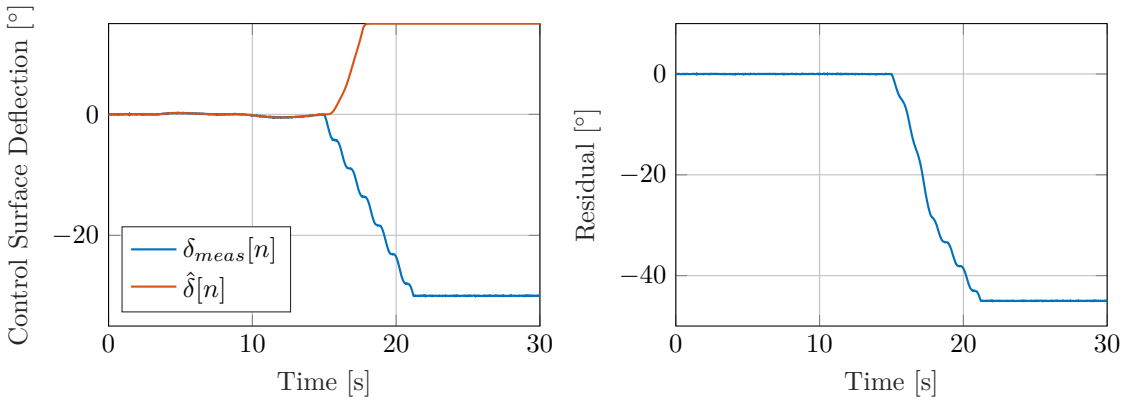

(b) Measured and Estimated Deflection

(c) Residual

Figure 3.13: Solid Failure at the Commanded Current with Non-zero Offset

The result of this failure is that there is permanently a large offset in the control error, causing the rod position to move in the direction of the offset, until its position saturates. This causes the control surface to diverge until it reaches its physical deflection limit at $-30°$.

Runaway detection would be more suitable for the detection of such a failure than oscillation detection.

3.5 Summary and Contributions

A simulation framework was created in collaboration with Airbus for the purpose of training and testing different OFC detection techniques. The simulation framework contains models for the physical actuator, the analytically redundant actuator, the oscillatory failures (both liquid and solid failures), the flight control system, and the aircraft longitudinal dynamics. The simulation framework was designed so that the actuator parameters and the oscillatory failure parameters (amplitude, frequency, type, and injection point) can be varied randomly to produce training and testing data that is sufficiently rich.

Six examples of simulated flight data have been illustrated, providing an overview of the differences between nominal flight data, training data, and failure case data. Nominal and fault-free training data sets will be used to determine suitable thresholds for each detection technique, and

test data sets with liquid and solid OFCs present will be used to test and evaluate the OFC detection techniques in terms of smallest amplitude OFC detected, detection time, and false alarms. However, from the examples, it was discovered that most solid failure cases would not be classified as oscillatory failures, as they instead cause actuator runaway.

The following four chapters will present the investigation and theoretical development of the four OFC detection techniques considered in this thesis, namely oscillation counting, integrated absolute error (IAE), discrete Fourier transform (DFT), and phase-locked loop (PLL).

Chapter 4

Time-Domain Residual Thresholding

4.1 Introduction

This chapter provides a detailed overview of the oscillation counting method, along with some additional observations that have not been thoroughly addressed in literature.

Oscillation counting is an existing OFC detection technique that was developed by Goupil [1] and is currently in service on the Airbus A380 passenger airliner. Oscillation counting is a time-domain residual evaluation technique, and is fundamentally a limit-checking approach to fault detection. This approach has found considerable success in the detection of oscillatory failures on the A380, and is considered to be the baseline for comparing other residual evaluation techniques. This oscillation counting technique was implemented in Matlab for this project, using the available information from literature, so that it could be trained and tested in the simulation framework, and so that its performance could be compared against the other OFC detection techniques developed in this project.

Specific concerns and design considerations were highlighted while developing the Matlab implementation, and the potential effect of these observations are discussed in Section 4.3.

4.2 Background Theory

4.2.1 Overview

The oscillation counting approach attempts to count the number of oscillations that occur in a sliding window of the residual. This is achieved by applying a positive and negative threshold to the residual data, and counting the number of successive and alternating crossings of the residual over some given threshold T. If the number of crossings that occur within the sliding window exceed some value T_c, then an oscillatory failure is inferred. This is depicted in Figure 4.1. Crossings are counted when the residual crosses the positive threshold with a rising edge, or the negative threshold with a falling edge.

The size of the sliding window, N, is expressed in number of samples, and is chosen to span over at least three cycles of the lowest expected OFC frequency. For example, if the lowest expected frequency is 1 Hz, and the sampling rate is 40 Hz, the length of the window is selected to be 3 seconds, or 120 samples. The use of this sliding window ensures that old threshold crossings are discarded. This helps to improve the robustness of the system, since only recent crossings are considered, while much older threshold crossings have no effect on the detection. Therefore, the crossing count is accumulated if a crossing is detected on the most recent residual sample, and the count is decremented when a crossing occurs at the oldest sample of the window.

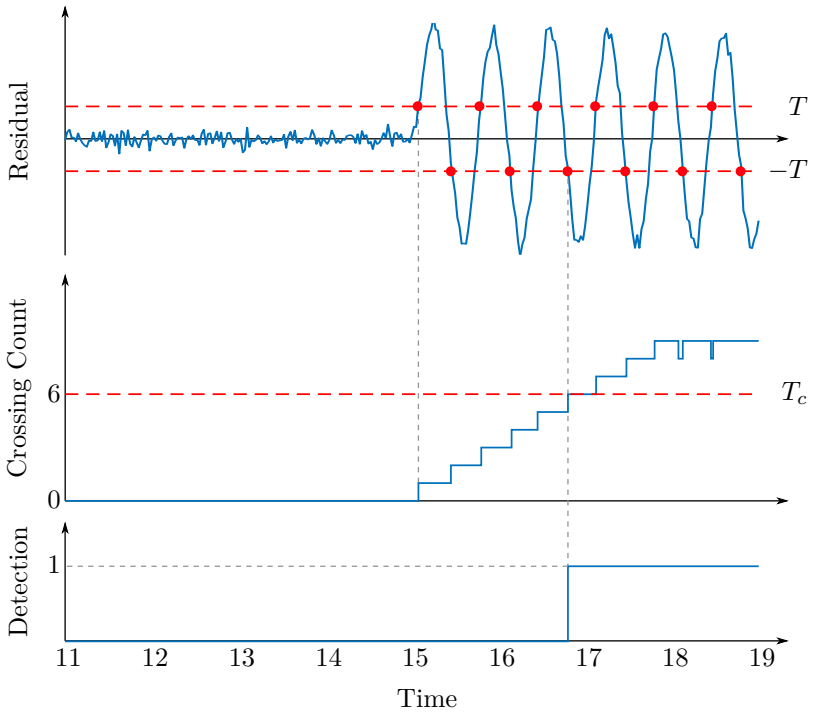

Figure 4.1: Overview of Oscillation Counting, adapted from [35]

Oscillation counting offers a robust approach to OFC detection, since checking for the existence of multiple crossings ensures that occasional and unexpected crossings that result from noise or modelling errors are ignored. The threshold for the number of crossings is directly related to the detection time requirements. For a 3 cycle detection time, T_c is chosen to be 6. In theory, this approach guarantees detection time, since for each period of the OFC, the first crossing occurs between 0° and 90° of the sinusoid's phase, and the second crossing occurs between 180° and 270°, as illustrated in Figure 4.2. Provided the amplitude exceeds the threshold, the first six threshold crossings will always be counted before the end of the third OFC cycle. In this way, detection time is independent of frequency and amplitude.

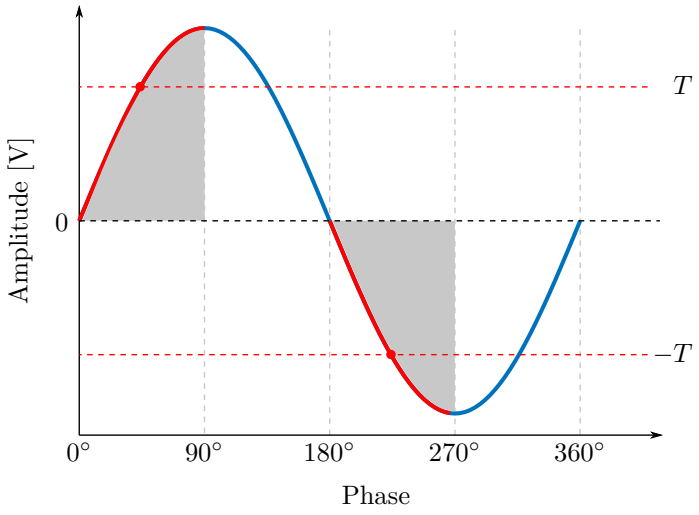

Figure 4.2: Threshold Crossings over Phase

A diagram of the entire oscillation counting technique is provided in Figure 4.3. Here, it can be seen that the residual is first filtered before oscillation counting is performed. Additionally, the technique handles the checks for liquid failures and solid failures separately, where the solid failure check offsets the thresholds by the estimated control surface deflection. These elements are discussed further in the proceeding sections.

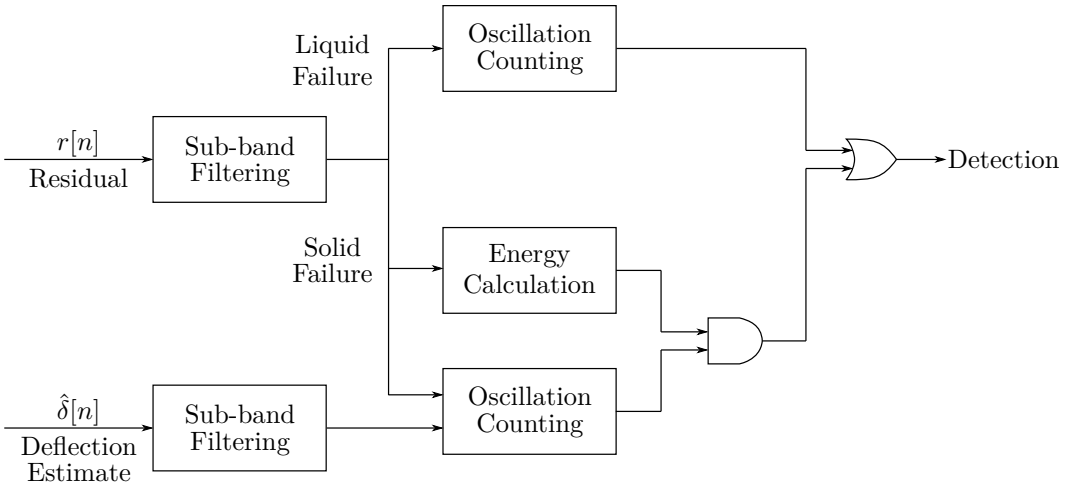

Figure 4.3: Oscillation Counting OFC Detection System

4.2.2 Residual Filtering

Robustness can be further improved with suitable filtering of the residual signal. Filtering assists in attenuating noise, allowing the thresholds to be lowered without fear of false alarms. Filtering can also be used to remove any DC components of the residual, ensuring that the mean of the signal remains at 0 V, and thereby guaranteeing threshold crossings.

Goupil made use of two 4^{th} order Cauer, or Elliptic bandpass filters to divide the residual into two spectral sub-bands. Dividing the residual into sub-bands allows the system to have different detection requirements for different frequencies. For example, the amplitude and detection time requirements could be more strict for lower frequencies than for higher frequencies.

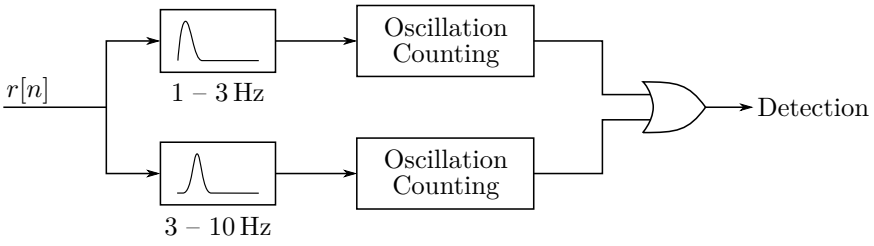

Figure 4.4: Oscillation Counting

In this investigation, the filter bandwidths are chosen to be 1 to 3 Hz and 3 to 10 Hz respectively. These cutoff frequencies are based on the logarithmic scale, where 3 Hz is exactly 3 times 1 Hz, and 10 Hz is approximately 3 times 3 Hz. Elliptic filters are chosen due to their superior frequency drop-off compared to other filters, providing maximum noise attenuation with minimal filter order, and subsequently minimal computational cost. Additionally, two windows with different sizes are used for the oscillation counting, namely 120 samples for 1 to 3 Hz, and 40 samples for 3 to 10 Hz.

4.2.3 Distinction between Liquid and Solid Failures

Goupil defines two potential oscillatory failure cases: liquid and solid failures [1]. These are expressed mathematically in Equations (3.2.7) and (3.2.8) in Chapter 3. Once the estimated actuator is subtracted from the measured position, the residual is expressed as

$$r[n] \approx \begin{cases} 0 & \text{for nominal behaviour} \\ A\cos\left(2\pi f n T_s\right) & \text{for liquid failures} \\ A\cos\left(2\pi f n T_s\right) - \hat{\delta}[n] & \text{for solid failures} \end{cases} \tag{4.2.1}$$

The residual for liquid failures contains the pure sinusoid centred at 0 V, and oscillation counting can be performed around zero. The solid failure, on the other hand, is offset by the negative estimated deflection, $\hat{\delta}[n]$. This means that counting oscillations around 0 V may not be viable. Instead, oscillations have to be counted around the opposite of the estimated deflection, $-\hat{\delta}[n]$. Oscillation counting is performed with the positive and negative thresholds offset by the negative estimation. Figure 4.5 shows an exaggeration example of the residual during a solid failure case, along with the negated estimated control surface deflection. Note that because the residual is filtered into different sub-bands, the estimated deflection must also be filtered before being used to offset the threshold.

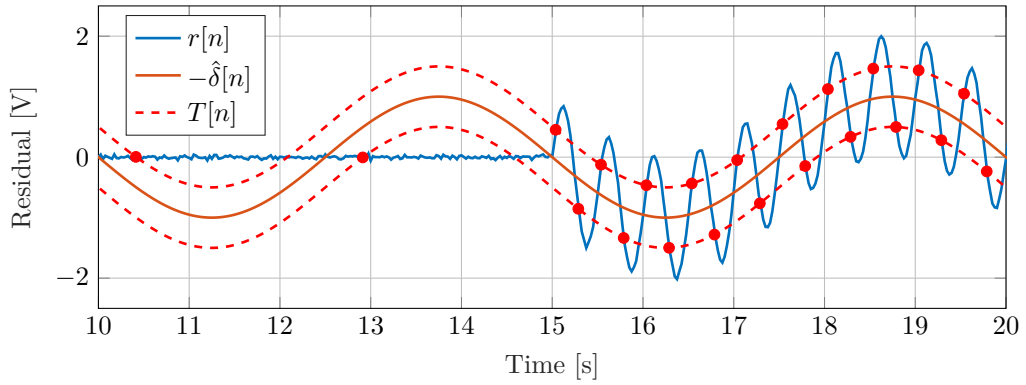

Figure 4.5: Residual and Estimated Control Surface Deflection During a Solid Failure Case

However, in the process of counting oscillations around the opposite of the deflection estimate, a new false alarm risk manifests. In a fault-free case where the control surface receives a sinusoidal command input, the offset thresholds regularly cross the fault-free residual, as seen in Figure 4.6.

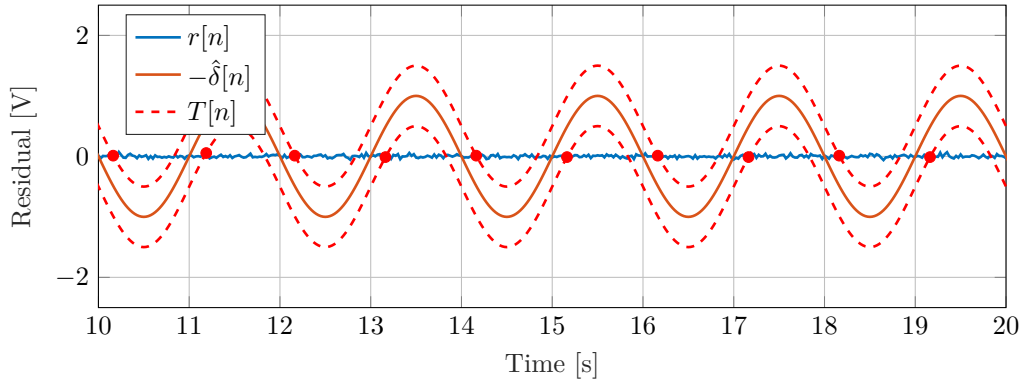

Figure 4.6: Potential Case for False Detections During Solid Counting

With no failure present, the residual is approximately zero, but the oscillation counting threshold is now offset by the sinusoidal deflection estimate, causing false threshold crossings. To counter

this, an additional robustness check is added to the system that checks whether the energy of the residual is above some threshold or not, as shown in Equation (4.2.2).

$$\sum_{k=0}^{N-1} r^2(k) > T_E \tag{4.2.2}$$

If the energy of the residual within the sliding window is greater than the energy in a fault-free case, and threshold crossings are present in the solid failure check, then a solid failure is inferred. An example of the residual energy is shown in Figure 4.7.

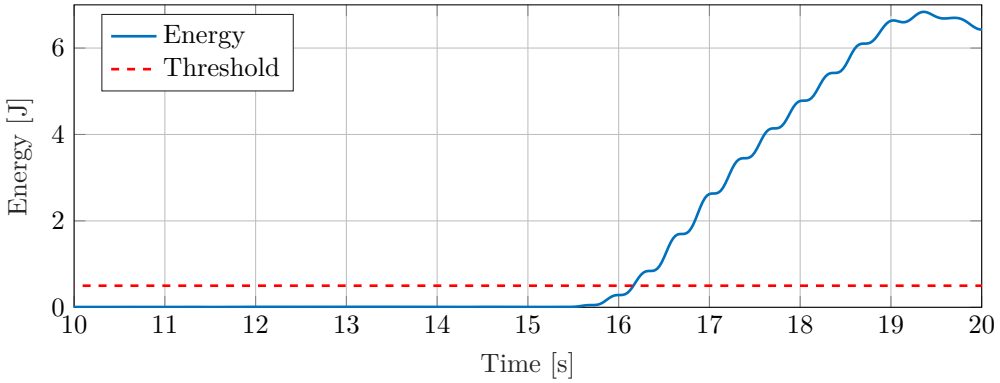

Figure 4.7: Residual Energy during Solid Failure

In summary, the oscillation counting method involves filtering the residual signal into two sub-bands, and counting the number of threshold crossings occur over a set period. Liquid and solid failures are handled separately, and for additional robustness, the energy in the residual is calculated and compared to a nominal energy value.

4.3 Critical Analysis, Insights, and Design Considerations

For benchmarking purposes, the oscillation counting technique described above was implemented in Matlab, and during the implementation, important observations and insights emerged that are not covered in existing literature. Concerns arose regarding the effects of the sub-band filters, sampling rate, and the implications of counting "successive and alternating" crossings. These concerns and design considerations are discussed in this section.

4.3.1 Counting Alternating Crossings

An important consideration in the detection of OFCs with the oscillation counting approach is the counting of "alternate" threshold crossings. In other words, if a positive threshold crossing occurred, the next crossing that would be counted would have to be at the negative threshold, and vice versa. This could greatly reduce the number of incorrect detections as a result of noisy threshold crossings, improving the robustness of the system.

This led to the first critical observation of the oscillation counting technique, which was that the alternating crossing check often caused late detections. If a spurious positive threshold crossing occurred due to noise, followed by a positive threshold crossing due to a newly induced OFC, this first OFC threshold crossing would not be counted, as illustrated in Figure 4.8. This adds a half-cycle delay to the OFC detection.

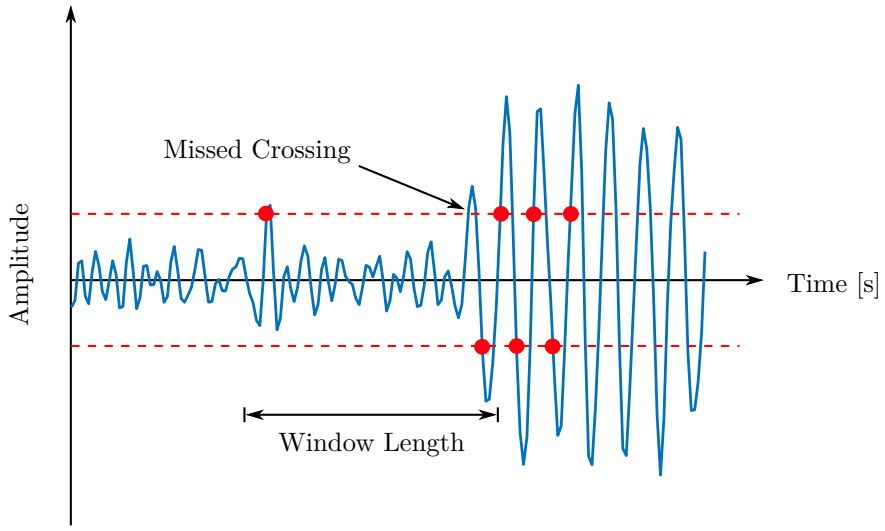

Figure 4.8: Potential for Missed Threshold Crossing

In the above figure, a spurious residual spike causes a threshold crossing before an OFC begins. This crossing falls within one window length of the start of the OFC. When the OFC starts, the first, positive threshold crossing is not counted. The negative crossing is then the first OFC crossing detected. Before the total number of detected crossings reaches 6 however, the first spurious crossing leaves the window, and the crossing counter is decremented. Now, one additional crossing is required before detection, resulting in a late detection.

Overall, this problem may appear very infrequently, but should be taken into account when a very low threshold is chosen. If detection time is a strict requirement, the counting of only alternating threshold crossings may not be recommended, especially if the detection of small amplitude OFCs is desired. Instead, count all threshold crossings that occur within the sliding window, irrespective of order of appearance.

To remain consistent with literature, the implementation used in the final results will make use of the alternating crossing check.

4.3.2 Filtering Effects

Elliptic filters are used to divide the residual into sub-bands. As stated, this reduces the amount of noise in the residual signal, and allows for different window sizes to be used for the detection of different frequencies. Unfortunately, the use of bandpass filters introduces delays and transients that may affect the detectability of OFCs within the maximum detection time. Figure 4.9 shows the effect of the $1 - 3\,\text{Hz}$ filter on a $1\,\text{Hz}$ OFC with an amplitude of $1\,\text{V}$, in a case with no noise or modelling errors.

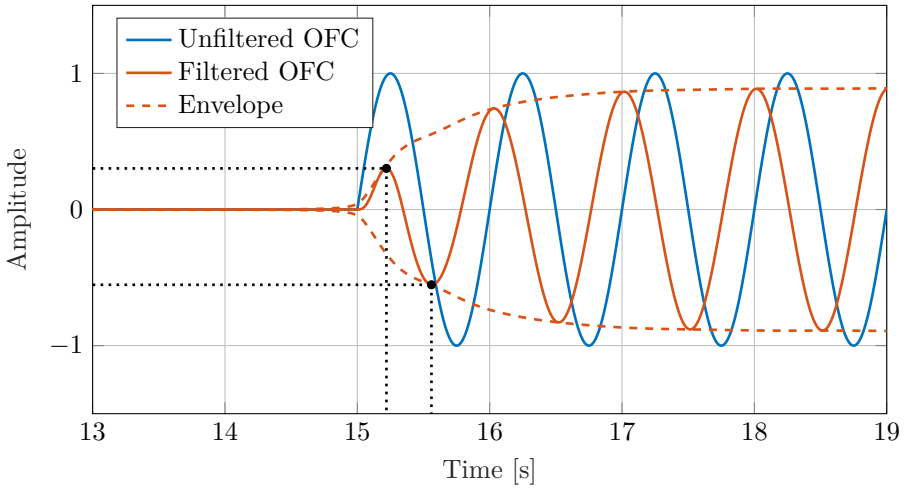

Figure 4.9: Effect of Bandpass Filter on 1 Hz OFC

Here, it is shown that the filter causes the OFC's amplitude to rise gradually to its maximum value, which is visualised by the envelope plot. As a result, the first two peaks of the filtered OFC, indicated in the figure above, have magnitudes of 0.3 V and 0.55 V respectively. Because of this attenuation, these first two peaks may not cross the threshold, resulting in a delay in the detection time by at least one cycle. This effect is shown in Figure 4.10.

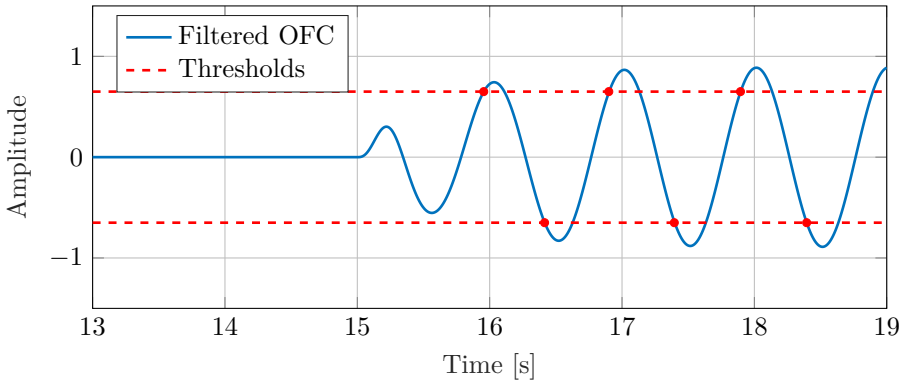

Figure 4.10: Missed Threshold Crossings due to Bandpass Filter

Similar effects can be seen for higher frequency OFCs, such as the 3 Hz OFC present in Figure 4.11, which shows the first two peaks having values of 0.36 V and 0.84 V. The potential result of this is that for an OFC to have guaranteed detection time, the first peak must be higher than the threshold. Based on the 0.3 V of the first peak in the 1 Hz case, it can be assumed that the amplitude of the OFC must be larger than the threshold by a factor of 3.3 for guaranteed detection time.

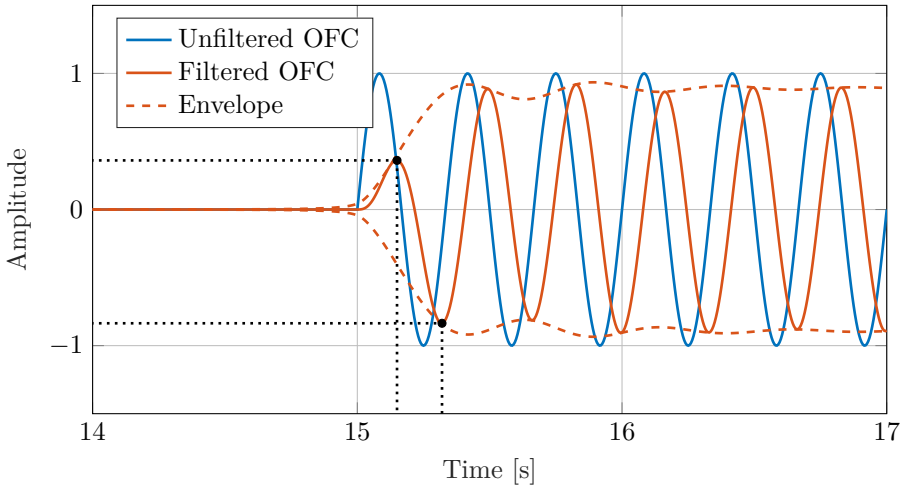

Figure 4.11: Effect of Bandpass Filter on 3 Hz OFC

This "rise time" is not the only side-effect of the filters that could be detrimental to OFC detection. Two more effects of the filters are observed in Figures 4.9 and 4.11. Firstly, the amplitude of the filtered OFCs is smaller than that of the unfiltered OFCs. This is due to the passband ripple of the filter frequency response, shown in Figure 4.12. The amount of passband ripple can be adjusted, with smaller ripple resulting in a flatter cutoff slope. For these tests, the passband ripple was chosen as 1 dB. This results in a minimum attenuation factor of 0.89 over the passband. This can further reduce the method's ability to detect small amplitude OFCs.

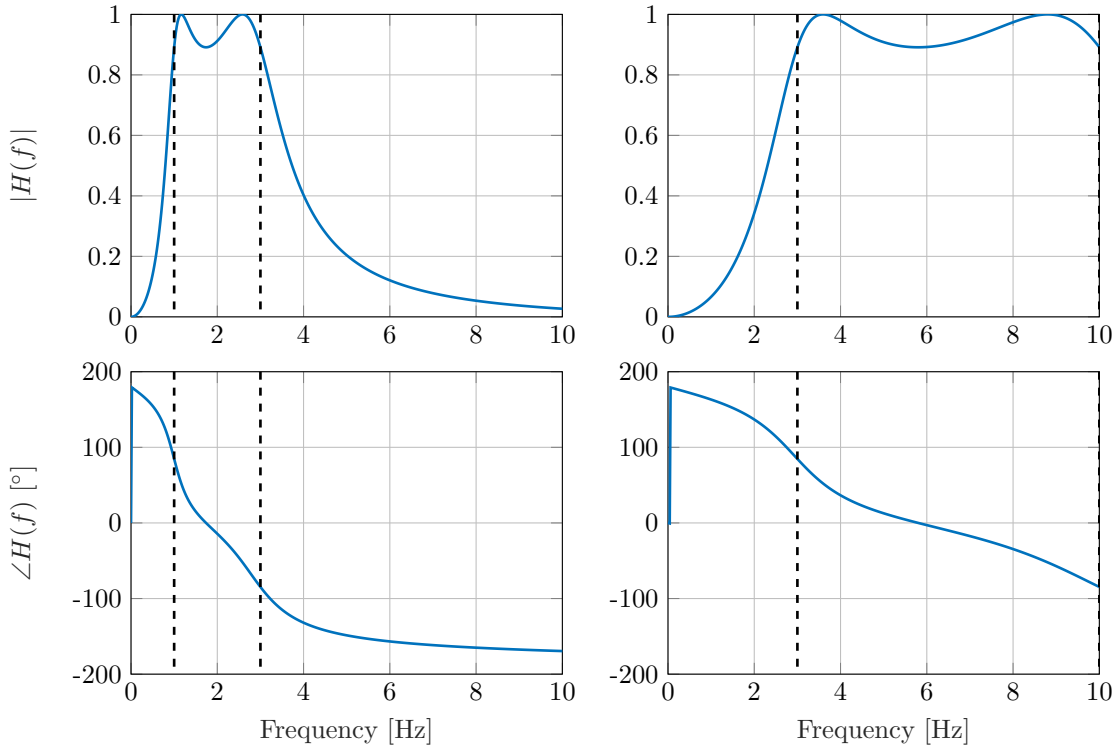

Figure 4.12: Frequency Response of 1 – 3 Hz (Left) and 3 – 10 Hz (Right) Bandpass Filters

Secondly, the filters induce phase delays in the residual, which can further influence detection time. The 1 Hz OFC in Figure 4.9 shows a phase lead, while the 3 Hz OFC in Figure 4.11 shows a phase

lag. In the case of the phase lag, the threshold crossings occur later than they would without filtering. Figure 4.12 shows that at 3 Hz and at 10 Hz, the phase delay reaches a maximum of 90° for these 4^{th} filters. However, the potential phase lag tends to increase with higher order filters. This lag effect, coupled with the rise time of the filter could result in significant delays before detection. Fortunately, by inspecting Figure 4.2 further, it can be seen that an OFC should be detected 90° before the end of the last period, suggesting that the 90° phase lag may be allowable.

In summary, with the combined effects of the effective rise time of the filter and the possible 90° phase lag, oscillation counting could have a potential 1.5 cycle delay before detecting a small amplitude OFC, and the filters can cause slight attenuation in the OFC, degrading the minimum detectable amplitude performance. As a rule of thumb, only an OFC that is 3.3 times larger than the threshold should be considered to have guaranteed detection time, ignoring the potential effects of noise.

4.3.3 Upsampling

One final potential pitfall for the oscillation counting is the use of low sampling frequencies. For example, at a sampling rate of 40 Hz, a full cycle of a 10 Hz OFC is only described by four samples. This means that the samples may not show the full amplitude of the OFC. As shown in Figure 4.13, the amplitude of a sampled 10 Hz sinusoid can potentially reduce to 70 % of its true amplitude.

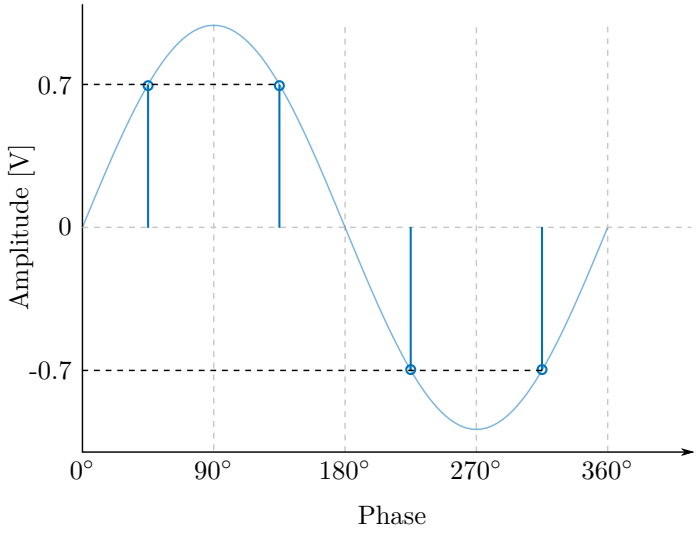

Figure 4.13: Effects of Low Sampling Frequency on High Frequency OFCs

For the purposes of oscillation counting, it is recommended to use a sampling rate that is at least ten times higher than the highest OFC frequency to ensure that there are at least 10 samples in a single OFC cycle. For example, if the highest expected OFC frequency is 10 Hz, then a 100 Hz sampling rate is recommended. However, if the sampling frequency is lower than desired, upsampling can be used to increase the sampling frequency of the residual signal. The sampling frequency of a signal can be increased by a factor of L by inserting $L-1$ zeros between each sample of the original signal. This is mathematically expressed as:

$$r_L[n] = \begin{cases} L \cdot r[n/L], & n = 0, \pm L, \pm 2L, \ldots \\ 0, & \text{otherwise} \end{cases} \tag{4.3.1}$$

The sub-band filters can then remove the unwanted high-frequency components that are introduced by the upsampling. For the case of a 40 Hz sampling frequency, an upsampling factor of 3 is used to increase the sampling frequency to 120 Hz. Figures 4.14 and 4.15 show the potential improvements of upsampling a 10 Hz sinusoidal signal that was sampled at 40 Hz.

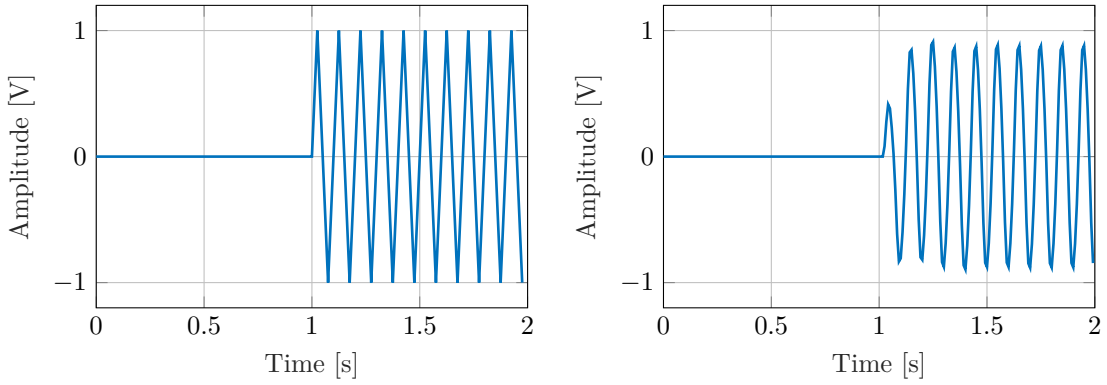

Figure 4.14: Upsampling Results for 10 Hz Sinusoid

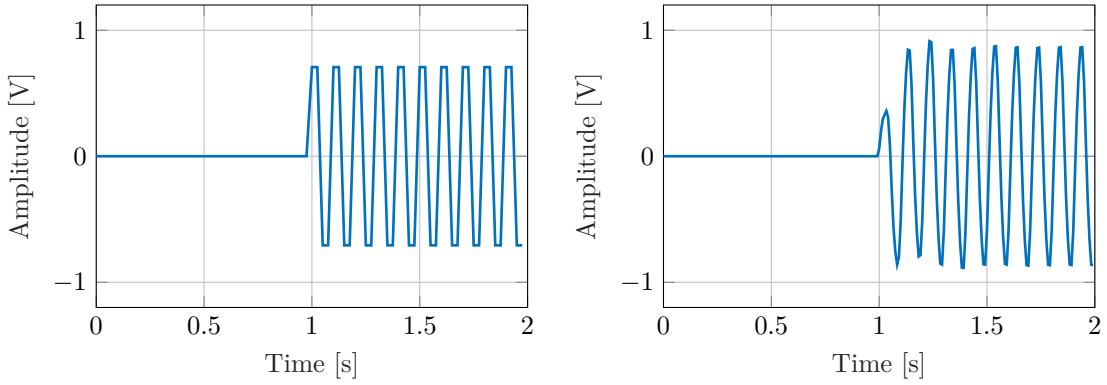

Figure 4.15: Upsampling Results for 10 Hz Sinusoid with 45° Phase Shift

4.4 Threshold Training

Threshold training is performed on training data using a binary search algorithm to find the smallest threshold that would not result in a false alarm in the training data.

An upper bound and a lower bound for the threshold are assumed. These are chosen as 0 V and 30 V, based on the maximum allowable deflection of the control surface of 30°. The oscillation counting threshold is set to the average, or middle of these two bounds. This threshold is then used to check fault-free training data for OFCs. If an OFC is (incorrectly) detected, the threshold is considered to be too low, and the initial lower bound is moved up to this middle value. The new threshold is recalculated as the average between these new bounds. If the OFC check does not result in a false alarm, the threshold can potentially be lowered. The upper bound is then moved down to the middle value, and once again the new threshold is the average of the new bounds. This process is repeated until a desired accuracy is achieved.

The iterative search ensures that the threshold is set to as low a value as possible that still avoids false alarms. It results in a threshold that is set close to the noise level, which allows for occasional noisy threshold crossings without fear of false alarms. Similar approaches to threshold selection is followed for each detection method to ensure fair comparisons in the final results.

4.5 Summary

A detailed investigation and review of the current implementation of OFC detection used on the A380 aircraft is reported in this chapter. The oscillation counting residual evaluation scheme accepts a residual signal, filters the signal into frequency sub-bands, and counts successive and alternating threshold crossings. Once six crossings have been detected, a failure is inferred. This

method required the use of separate checks for liquid and solid failures, increasing the computational complexity of the system.

The oscillation counting approach was implemented in Matlab according to literature. During implementation, potential causes for late detections were identified. The sub-band filters resulted in transient effects and phase shifting that could cause late detections. The check for alternating threshold crossings could occasionally result in the first peak of an oscillatory failure being ignored by the detection algorithm. Finally, the use of an upsampling stage was recommended in cases where the residual signal is sampled at a low sampling frequency.

A binary search algorithm was implemented to train the oscillation counting threshold. This allowed the threshold to be set as low as possible without triggering false alarms, while still allowing occasional threshold crossings during nominal flight operation.

4.6 Contributions

The following contributions were made in this chapter:

- A potential cause for late detections was identified in the oscillation counting technique, based on the counting of only alternating threshold crossings. The proposed solution for this is to count all detection crossings rather than specifically alternating crossings.
- The potentially negative effects of filtering on detection time were identified, specifically the rise time and the phase shift. Based on the behaviour of the filters' rise time, it was proposed that the smallest amplitude oscillation that could be detected within 3 cycles could be approximated by applying a scaling factor of 3.3 to the chosen threshold. Additionally, it was concluded that the effects of the filters' phase shift should not cause late detections provided the bandpass filters are of 4^{th} order.
- The potentially negative effects of low sampling rate were identified, and an upsampling stage was proposed to increase the sampling rate of the residual signal.
- A binary search algorithm was implemented as a way to determine the lowest threshold that would not cause false alarms.

Chapter 5

Energy-Based Techniques

5.1 Introduction

This chapter investigates residual evaluation schemes that rely on calculating the energy of the residual signal. The primary focus is placed on the integrated absolute error (IAE) approach, which was initially proposed for online detection of oscillation in control loops of processing plants. This approach has garnered a lot of attention in literature, and many improvements have been suggested. However, many of these improvements were suggested for offline fault detection, isolation and diagnosis, where detection time and computational complexity is not as critical.

Section 5.2 provides background theory to the IAE detection scheme. Section 5.3 then explains how the approach is implemented and adapted for the OFC detection problem.

5.2 Background Theory

Hägglund [17] formulated the integrated absolute error (IAE) for the online detection of oscillations in control loops. The IAE was calculated by integrating the control error between successive zero crossings. This is mathematically described as

$$\text{IAE} = \int_{t_{i-1}}^{t_i} |e(t)|dt, \tag{5.2.1}$$

where t_{i-1} and t_i are two zero crossing time instances. Nominal control-loop operation typically results in small IAE values, and failures can be inferred when the IAE exceeds some threshold IAE_{lim}. This is illustrated in Figure 5.1. Furthermore, oscillatory failures would show periodic threshold crossings. Therefore, a oscillatory failure can be inferred if a certain number of threshold crossings, defined as n_{lim}, occur over some window of time.

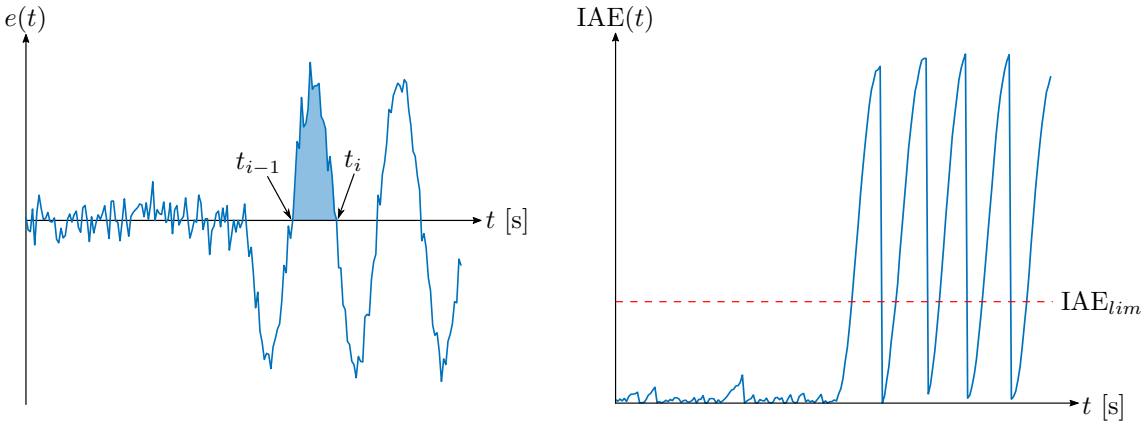

Figure 5.1: Illustration of the Integrated Absolute Error (IAE)

Hägglund states that the threshold is chosen based on the desired amplitude that should be detected. For a pure sinusoidal signal with amplitude A and frequency ω, the threshold should be

$$\text{IAE}_{lim} \leq \int_0^{\pi/\omega} |A \sin(\omega t)| dt = \frac{2A}{\omega}. \tag{5.2.2}$$

Here, the value ω in the threshold is chosen to be a constant based on the parameters of the PI(D) controller used in the control loop. However, Equation (5.2.2) also clearly shows that the calculated value of the IAE is dependent on frequency, and high frequency failures will not manifest with the same magnitude as low frequency failures, and may therefore go undetected. Hägglund illustrates this as a detection domain on an amplitude-frequency axis, as shown in Figure 5.2. The blue region shows amplitude-frequency combinations that can be detected, while the red regions show amplitude-frequency combinations that would likely go undetected. Hägglund acknowledges that the detection method requires a sufficiently high amplitude with a sufficiently low frequency for successful detections.

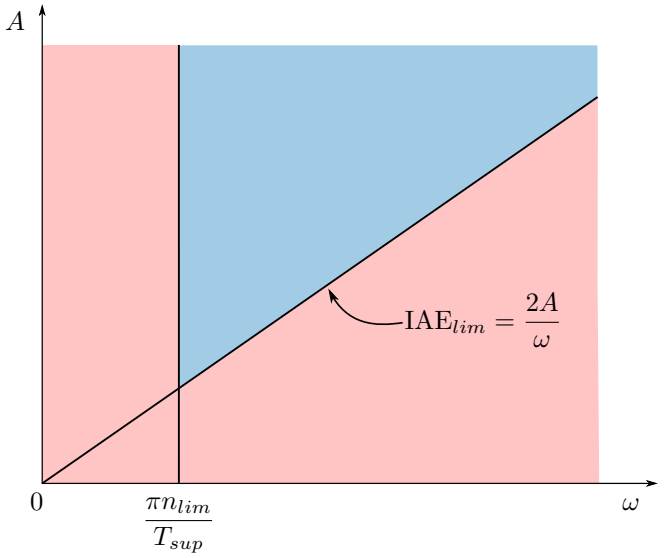

Figure 5.2: Viable Values of Amplitude and Frequency for IAE Detection

In general, the IAE residual evaluation technique has successfully been used for the detection of oscillations in closed-loop control systems, and offers a computationally efficient approach to the OFC detection problem. This is a time-domain technique, similar to the oscillation counting approach proposed by Goupil. It is therefore worth investigating if the IAE has any advantage over the oscillation counting approach, or vice versa.

5.3 Applying the IAE to OFC Detection

This section describes how the IAE technique is adapted for OFC detection in this project, specifically highlighting a new approach to IAE thresholding.

5.3.1 Overview of Adapted IAE Method

The IAE is used here as a residual evaluation technique, and is applied to the residual signal rather than to the control error signal. The IAE is therefore defined as:

$$\text{IAE} = \int_{t_{i-1}}^{t_i} |r(t)| dt, \tag{5.3.1}$$

and the detection threshold from Equation (5.2.2) is

$$\text{IAE}_{lim} = \frac{2A}{\omega}, \tag{5.3.2}$$

where A is the amplitude of the smallest OFC that can be detected, and must be chosen beforehand. An unfortunate shortcoming of this technique is that its threshold is dependent on frequency. A 10 Hz OFC would have to be 10 times larger than a 1 Hz OFC for them both to cross the same threshold.

An alternative approach to choosing the threshold is to assume that the time between the two zero crossings is half the period of an OFC, and thus the frequency of the signal can easily be estimated using this time. In other words, the threshold IAE_{lim} can be calculated using this frequency estimation:

$$\text{IAE}_{lim} = \frac{2A \cdot 2(t_i - t_{i-1})}{2\pi} \tag{5.3.3}$$

This offers a threshold that will adapt to the time between zero crossings, and will set low thresholds for high frequencies, and higher thresholds for low frequencies.

5.3.2 Implementation

In the discrete domain, the IAE calculation and threshold calculation simplifies to

$$\text{IAE}[n_i] = \sum_{n=n_{i-1}}^{n_i-1} r[n] \quad \text{and} \quad \text{IAE}_{lim}[n_i] = \frac{2A(n_i - n_{i-1})}{\pi}. \tag{5.3.4}$$

where n_{i-1} and n_i represent two successive zero crossing samples. The IAE can be implemented in a simple fashion, summarised as follows: At each sample time, check for a sign change between the current and previous residual sample. This indicates a zero crossing occurrence. If no sign change is present, then increment an integration time variable, and add the absolute value of the most recent residual sample to a running total. If a zero crossing has occurred, the threshold is calculated based on the integration time, and this is compared to the running total. If the running total is greater than the threshold, a "load disturbance" is flagged. Once the check is performed, the integration time and running total are both reset to zero.

The implementation of the IAE shares similar characteristics to the oscillation counting implementation. The residual is evaluated over a window of samples to forget old threshold crossings, and a failure is flagged if 6 threshold crossings occur over this window. Sub-band filters are also used here for noise attenuation, and ultimately to ensure fair comparisons between the two methods. Figure 5.3 shows an example of an OFC detection using the IAE.

The red stem plots shown above represents the threshold which is recalculated at every zero crossing. A load disturbance is detected if the IAE is greater than the threshold. The OFC is detected after 6 threshold crossings. It can be seen in the IAE graph that the magnitude of the threshold is dependent on the time between zero crossings. The threshold is set very low for short durations, and higher for long durations between zero crossings.

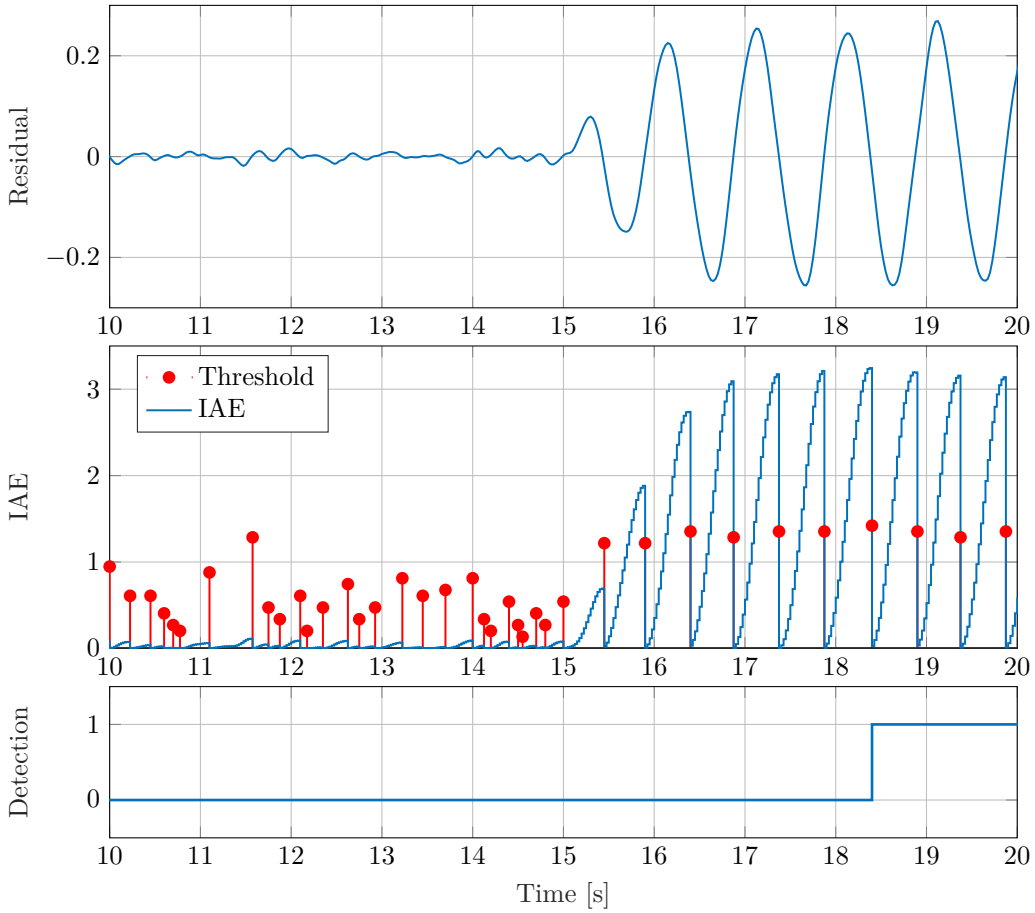

Figure 5.3: Detection of an OFC with the IAE Method

There are a few observations in Figure 5.3 that should be noted. Firstly, the integration over the first half-period does not cross the threshold. This is due to the rise time of the filters, as discussed in Section 4.3.2, and results in a late detection. Secondly, the threshold crossing occurs at the very end of a half-cycle, unlike in the oscillation counting approach. This suggests that the IAE is more prone to late detections than oscillation counting.

Additionally, the IAE may also require a unique solution for solid failure detection. One possible solution is to check where the residual crosses the filtered estimated position, rather than zero crossings. This is similar to the oscillation counting approach of offsetting the thresholds with the filtered estimated position. However, a similar false alarm risk would exist here as well.

5.3.3 Threshold Training

The threshold that must be tuned for the IAE is A, the amplitude of the OFC. This is the value used in Equation (5.3.4) to calculate the adaptive threshold. The same search algorithm that was used to determine the threshold for oscillation counting (described in Section 4.4) is used to find the smallest amplitude that can be detected while still ensuring zero false alarms.

5.4 Summary and Contributions

The integrated absolute error (IAE) offers another time-domain solution to the OFC detection problem, and can be implemented with a very low computational cost. This method integrates the residual between zero crossings, and compares the integrated value to a threshold.

The IAE was implemented in Matlab, based on existing literature. However, the IAE calculation suffered from a scaling factor that was dependent on the frequency of the OFC. A new adaptive

threshold was implemented to mitigate this effect, where the threshold was scaled by the time between successive zero crossings.

Unfortunately, the IAE suffers from the same filtering effects as the oscillation counting approach, and because threshold crossings are only observed at the end of each OFC cycle, the IAE will likely suffer from late detections.

Chapter 6

Frequency-Domain Residual Thresholding

6.1 Introduction

A popular method of determining the frequency content of a signal is the Fourier transform. However, given the real-time requirements and embedded nature of the OFC detection problem, the use of the discrete Fourier transform (DFT), and even its optimised implementation, the fast Fourier transform (FFT), have been avoided. These algorithms are considered to be too computationally expensive. Nevertheless, this chapter investigates the application of the DFT, its potentially superior detection performance over other proposed methods, and any performance optimisations that could make its implementation feasible. This may offer an upper bound of detection performance to which other detection methods should strive, and could even be applicable in future, more powerful embedded monitoring systems.

This chapter begins with a brief definition of the discrete Fourier transform in Section 6.2, and follows up with an overview of the proposed detection method in Section 6.3. Section 6.4 provides a mathematical proof that this approach can guarantee detection time. Sections 6.5 and 6.6 offer solutions to improvements in frequency resolution and computational optimisations. Finally, Section 6.7 offers a more experimental approach to the use of the Fourier transform that makes use of varying window sizes.

6.2 Background Theory

The two main concepts that are used throughout this chapter are the discrete-time Fourier transform (DTFT), and the discrete Fourier transform (DFT). Both of these transforms decompose time signals into frequency-domain representations of their frequency content.

The DTFT, $X(\omega)$, of a discrete-time signal $x[n]$ is defined as [36]:

$$X(\omega) \triangleq \sum_{n=0}^{N-1} x[n]e^{-j\omega n}, \quad \omega \in [0, 2\pi] \text{ rad/sample} \tag{6.2.1}$$

The DFT, $X[k]$, is a discretisation of the DTFT:

$$X[k] \triangleq \sum_{n=0}^{N-1} x[n]e^{-j2\pi kn/N}, \quad k = 0,\ 1,\ ...,\ N-1 \tag{6.2.2}$$

where

$$\begin{aligned} x[n] &\triangleq \text{input signal at time step } n \\ n &\triangleq \text{sample number} \\ k &\triangleq \text{frequency sample number} \\ N &\triangleq \text{total number of samples} \end{aligned}$$

6.3 Overview of Detection Algorithm

The frequency-domain residual thresholding technique for OFC detection relies on the use of the DFT to analyse the residual signal. The DFT provides a breakdown of the frequency content of the residual, and the magnitude of these frequency components are used to identify OFCs. In the event of an OFC, the spurious sinusoid would manifest in the frequency domain as a spike at one particular frequency. By representing the residual data in the frequency domain, it becomes much easier to extract and verify the existence of any abnormal frequency components from the noise and modelling errors. This naturally leads to a very simple detection method, which is to check whether any frequency component exceeds some predefined threshold value.

This detection technique is designed to use the short-time Fourier transform (STFT), where, at each time step, the most recent N samples of the residual are used to calculate its DFT. This window of N samples is referred to as a rolling window, and is used to provide an estimate of the instantaneous frequency content. The process is shown in Figure 6.1. Detection time was discovered to be highly dependent on window size, with shorter windows offering shorter detection times.

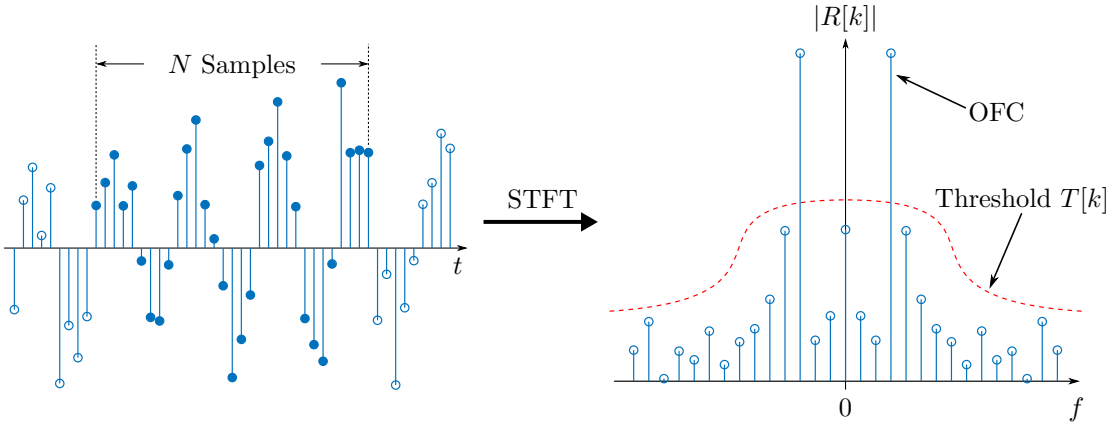

(a) Residual $r[n]$ (b) Frequency Components

Figure 6.1: Overview of Frequency Domain Thresholding

Initially, the DFT detection technique was designed to use a single fixed threshold value T for all frequency components. However, this "fixed" threshold showed poor detection performance, and exhibited late detections for high-frequency failures. The investigation into thresholding the DFT soon led to the concept of a "frequency-dependent threshold", $T[k]$, where the threshold is varied across frequency. Further investigation into this frequency-dependent threshold led to a mathematical proof that the use of a fixed threshold to cover all frequencies would be insufficient for ensuring consistent detection times, as will be shown in Section 6.4. Therefore, the threshold is instead chosen to be unique for each frequency component based on the expected frequency content of the residual for nominal flight.

The threshold training process relies on large amounts of training data, similar to that of the oscillation counting method. Essentially, the detector is provided with large amounts of fault-free residual data, and the maximum value of each frequency component of the residual is determined and stored. The result is a range of values over frequency that specify the maximum magnitude of the frequency components that existed in the fault-free datasets. Any deviation over any of these values suggests a fault. More detail of this process is given in Section 6.4.

6.4 Detection Time and Thresholding Analysis

As specified in the project requirements, this detection method must be able to confirm the presence of an OFC within 3 cycles of its onset. This requirement has implications on the size of the rolling window and the magnitude of the frequency-dependent threshold. A shorter window allows

the DFT to update the instantaneous frequency estimate more quickly, reducing detection time. This naturally has the further advantage of requiring fewer computations every time step, since fewer frequency components are calculated for shorter windows, based on Equation (6.2.2). The disadvantage of using short windows is that the resolution of the DFT is inversely proportional to N, the size of the window. Such large quantization errors may result in OFCs not being detected. This is explained further in Section 6.5. It is also preferable to have at least one full period of the lowest expected OFC frequency in the window, to improve the chances that the OFC will be correctly differentiated from noise.

The next section considers the transient behaviour of the frequency components that are calculated using the short-time Fourier transform applied to a rolling window. The transient behaviour of the frequency components must be considered, since it affects the detection time.

6.4.1 Transient Behaviour of Frequency Components

The transient behaviour of the frequency components that are calculated using the STFT is affected by the length of the rolling window that is used. The length of the window is chosen based on the lowest frequency that needs to be detected, namely 1 Hz. A 3 second window with a sampling rate F_s of 40 Hz results in a window size of 120 samples, with a frequency resolution of $0.\dot{3}$ Hz. This rolling window would include 3 cycles of the lowest OFC frequency, but 30 cycles of the highest expected frequency (10 Hz), raising concerns about whether a high-frequency OFC would indeed be detected in time.

Fortunately, a mathematical expression for the detection time of an OFC can be obtained. To simplify the derivation process, we assume an ideal monitoring system where modelling errors in the residual $r[n]$ during nominal, fault-free operation are negligible. An OFC induced into the system manifests as a pure sinusoid in the residual signal, as follows:

$$r[n] = \begin{cases} 0 & \text{for nominal behaviour} \\ A\cos\left(2\pi f n T_s\right) & \text{for liquid failures} \\ A\cos\left(2\pi f n T_s\right) - \hat{\delta}[n] & \text{for solid failures} \end{cases} \tag{6.4.1}$$

where T_s is the sampling period, A is the amplitude of the OFC, f is the frequency of the OFC, and $\hat{\delta}[n]$ is the residual offset in solid failure cases. $\hat{\delta}[n]$ is considered to be a low-frequency signal, and is ignored for this analysis.

The OFC manifests at some time t_f. The residual can therefore be expressed as:

$$r[n] = \begin{cases} 0 & \text{for } nT_s < t_f \\ A\cos\left(2\pi f n T_s\right) & \text{for } nT_s \geq t_f \end{cases} \tag{6.4.2}$$

Since the system measurements occur at discrete time intervals, the failure will begin to manifest in the residual at sample n_f.

$$n_{f-1}T_s < t_f \leq n_f T_s \tag{6.4.3}$$

An example of this residual is shown in Figure 6.2, along with an illustration of the rolling window.

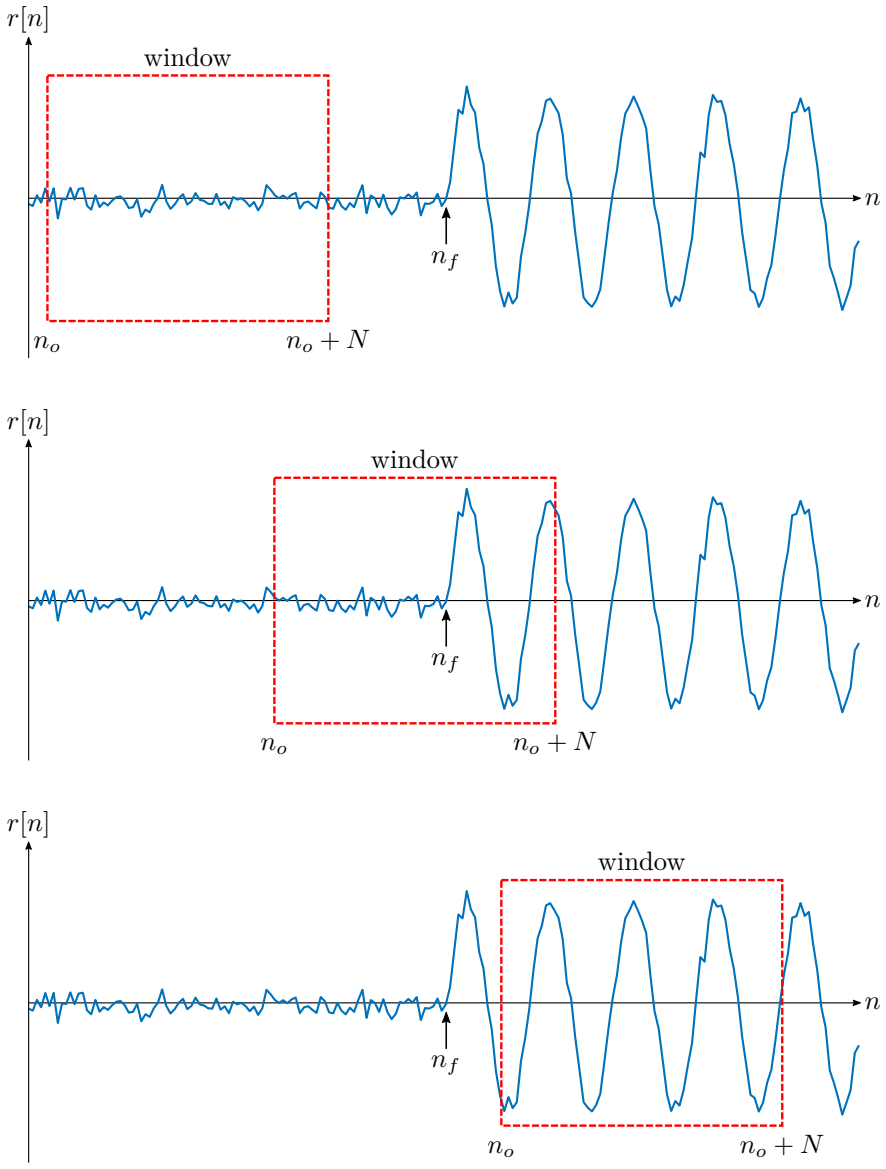

Figure 6.2: Rolling Window

The DFT, $R[k]$, of the window of residual $r[n]$ is defined as:

$$R[k] = \sum_{n=n_o}^{n_o+N-1} r[n]e^{-j2\pi k(n-n_o)/N}, \quad k = 0,\ 1,\ ...,\ N-1 \tag{6.4.4}$$

In this equation, it is clear that the DFT is calculated over a window of N samples, where n_o denotes the first sample of the window, and the current time step is

$$n = n_o + N - 1. \tag{6.4.5}$$

Assume a case where sample n_f falls within the window, as in the second plot of Figure 6.2. Expanding Equation (6.4.4) with the residual expression given in Equation (6.4.2) gives:

$$R[k] = \sum_{n=n_o}^{n_f-1} r[n]e^{-j2\pi k(n-n_o)/N} + \sum_{n=n_f}^{n_o+N-1} r[n]e^{-j2\pi k(n-n_o)/N}. \tag{6.4.6}$$

With $r[n]$ equal to zero in the first term, the equation simplifies to

$$\begin{aligned} R[k] &= \sum_{n=n_f}^{n_o+N-1} A\cos(\omega n T_s) e^{-j2\pi k(n-n_o)/N} \\ &= \sum_{n=n_f}^{n_o+N-1} \frac{A}{2}\left(e^{j2\pi f n T_s} + e^{-j2\pi f n T_s}\right) e^{-j2\pi kn/N} e^{j2\pi k n_o/N} \\ &= \frac{A}{2} e^{j2\pi k n_o/N}\left[\sum_{n=n_f}^{n_o+N-1} e^{(j2\pi n)(fT_s-k/N)} + \sum_{n=n_f}^{n_o+N-1} e^{(-j2\pi n)(fT_s+k/N)}\right]. \end{aligned} \tag{6.4.7}$$

The two terms above represent the positive and negative frequency components. These two frequency components contain the same information, so it is decided to ignore the effects of the negative frequency component from this point forth. The positive component coincides with $f/F_s = k/N$, where both f/F_s and k/N represent the same discrete frequency on the unit circle, measured in cycles per sample. Note that k here also represents the integer number of complete periods of that frequency over N samples. The sample k that coincides with the OFC frequency f will from here be represented as k_o. The above equation becomes

$$R[k_o] = \frac{A}{2} e^{j2\pi k n_o/N} \sum_{n=n_f}^{n_o+N-1} e^0. \tag{6.4.8}$$

Since only the magnitude information of the Fourier transform is useful for the OFC detection, the absolute value of $R[k_o]$ further simplifies to

$$|R[k_o]| = \frac{A}{2} \sum_{n=n_f}^{n_o+N-1} 1 = \frac{A\,(n_o + N - n_f)}{2}. \tag{6.4.9}$$

The current time step is the last sample of the window, $n = n_o + N - 1$. Substituting this into the above equation gives the time-domain expression of the k^{th} element of the DFT:

$$|R[k_o, n]| = \frac{A\,(n - n_f + 1)}{2} \tag{6.4.10}$$

Finally, the frequency component of the OFC, normalised by the length of the window N, can be represented as a piece-wise function:

$$|R[k_o, n]| = \begin{cases} 0 & \text{for } n < n_f \\ \dfrac{A\,(n - n_f + 1)}{2N} & \text{for } n_f \le n < n_f + N \\ \dfrac{A}{2} & \text{for } n \ge n_f + N \end{cases} \tag{6.4.11}$$

In the above equation, A, n_f, and N are all constants. This proves that the magnitude of the OFC's frequency component in the DFT linearly increases until it reaches its maximum value of $A/2$, exactly N time steps after the failure has manifested. This is clear proof that the size of N influences the detection time, but also suggests that the OFC amplitude A, and threshold T, have some effect as well. Larger amplitudes increase the gradient of Equation (6.4.11), resulting in earlier threshold crossings. Figure 6.3 illustrates this behaviour for a 2 Hz OFC.

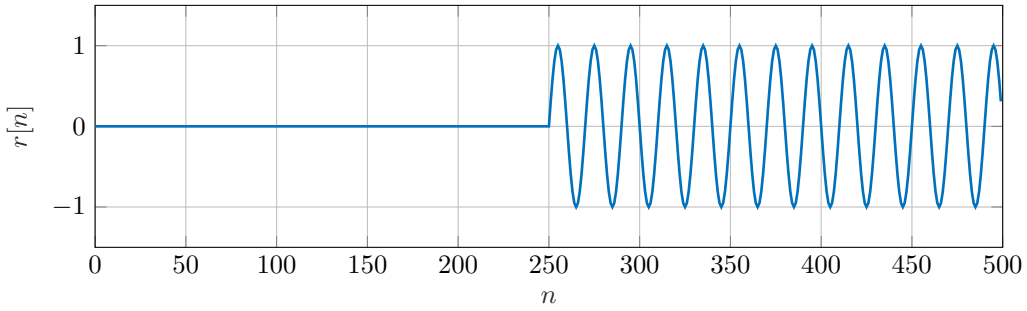

(a) Ideal Residual with 2 Hz OFC

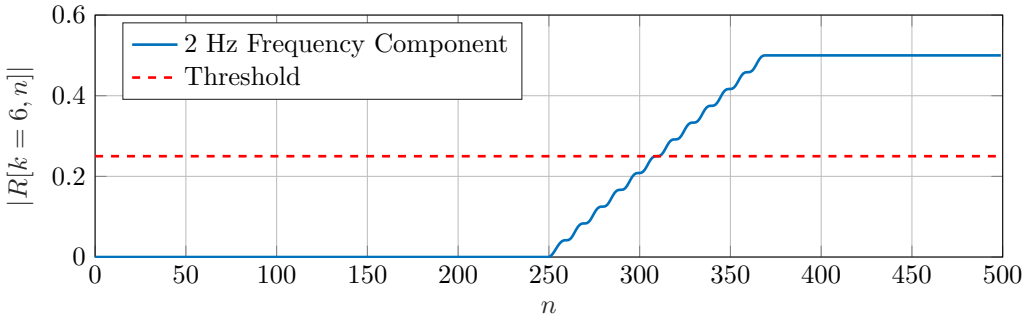

(b) 2 Hz Frequency Component over Time

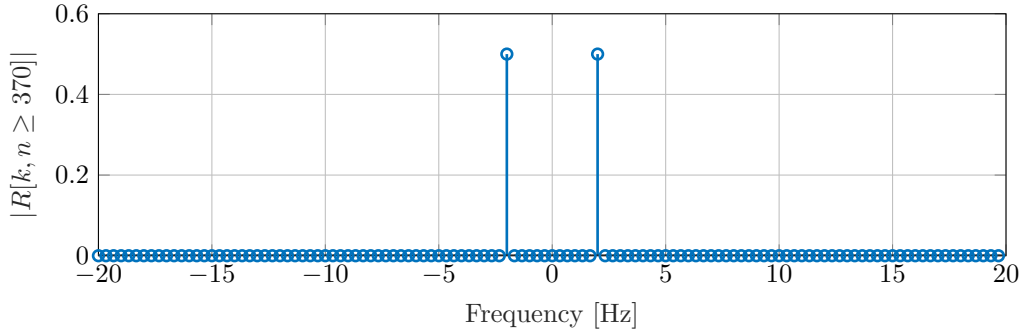

(c) Discrete Fourier Transform of $r[n]$ for $n \geq 370$

Figure 6.3: Illustration of OFC Detection with $f = 2\,\text{Hz}$, $F_s = 40\,\text{Hz}$, and $N = 120$

Figure 6.3b shows the magnitude of the 2 Hz DFT component over time. If the amplitude of the OFC is increased, or the threshold is lowered, the threshold crossing would occur more quickly. Note that the increase over $n \in [250, 369]$ is not linear due to spectral leakage, which occurs in situations where a non-integer number of periods of the OFC exist across the window.

6.4.2 Guaranteeing Detection Time

In order to guarantee the detection time of all OFCs in terms of cycles, the threshold must become a function of frequency. This is because Equation (6.4.11) is independent of frequency. If a single value is used as the threshold for all frequency components, higher frequency OFCs will cross the threshold after more cycles than lower frequency OFCs. In other words, higher frequencies must have lower detection thresholds. This is a valid strategy, since the dynamics of the actuator and the aircraft tend to reject high frequency control inputs.

Assume that an OFC causes a frequency component to cross a frequency-dependent threshold $T[k]$ at sample n_d. At this point,

$$T[k] < R[k, n_d] = \frac{A(n_d - n_f + 1)}{2N}. \tag{6.4.12}$$

The detection time of this OFC in cycles is

$$c = T_s \cdot f \cdot (n_d - n_f), \tag{6.4.13}$$

where c can at most be

$$c \leq T_s \cdot f \cdot (N - 1). \tag{6.4.14}$$

The required threshold to ensure a consistent detection time, c, for all frequencies f, is given as

$$T(f) < \frac{A(cF_s + f)}{2fN}. \tag{6.4.15}$$

Equation (6.4.15) provides the required threshold for consistent detection time c over all frequencies, and is illustrated in Figure 6.4. For comparative purposes, an example of a fixed, frequency-independent threshold is plotted as well.

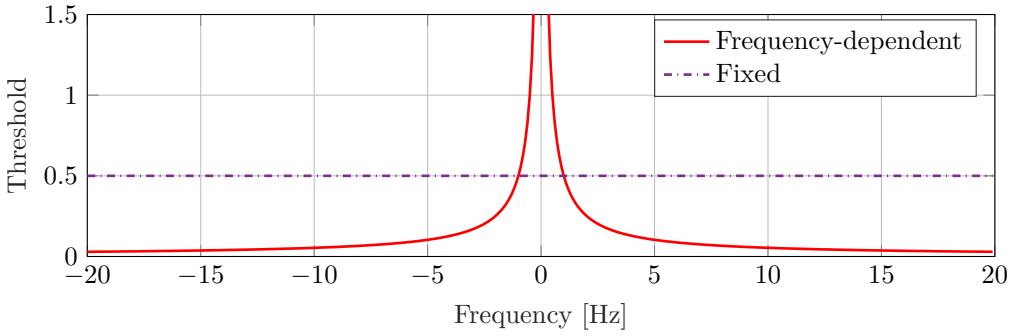

Figure 6.4: Frequency Dependent Threshold

The above analysis assumes that no noise or modelling inaccuracies exist in the system. A noisy residual could assist or hinder the detection by containing in-phase frequency components that reinforce the OFC, or out-of-phase components that attenuate the OFC.

6.4.3 Thresholding

Of course, the threshold that is set must not be low enough to cause false alarms. Real flight data must be used to determine the expected frequency content of normal flight operation. A large quantity of nominal flight data is used to train the threshold by running the OFC detection algorithm on this training data to find the maximum magnitude of each frequency component. This provides a frequency-domain bound for nominal flight situations. Any deviation over this bound can be considered a failure. This threshold may not hold strictly to Equation (6.4.15), but the equation can be used to determine the lowest amplitude OFC that will be reliably detected within 3 cycles. The evolution of the threshold during training is illustrated in Figure 6.5.

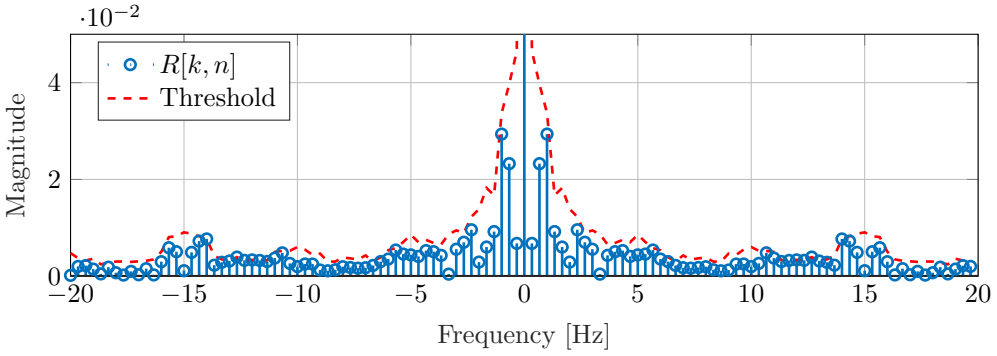

(a) Threshold Trained after 1 Minute of Flight (Training) Data

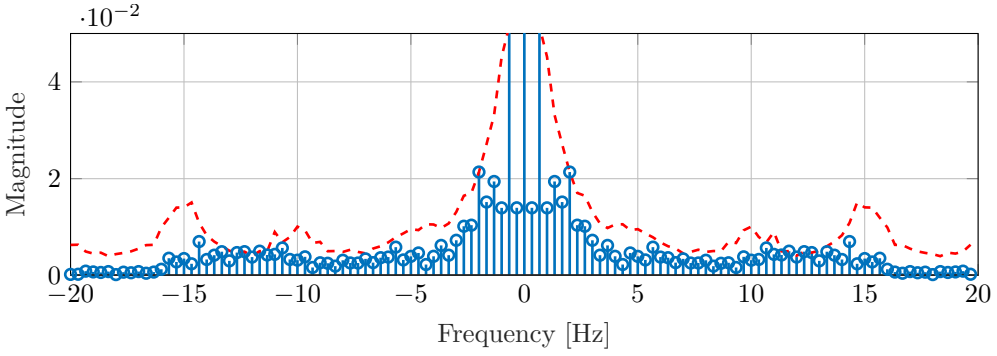

(b) Threshold Trained after 15 Minutes of Flight (Training) Data

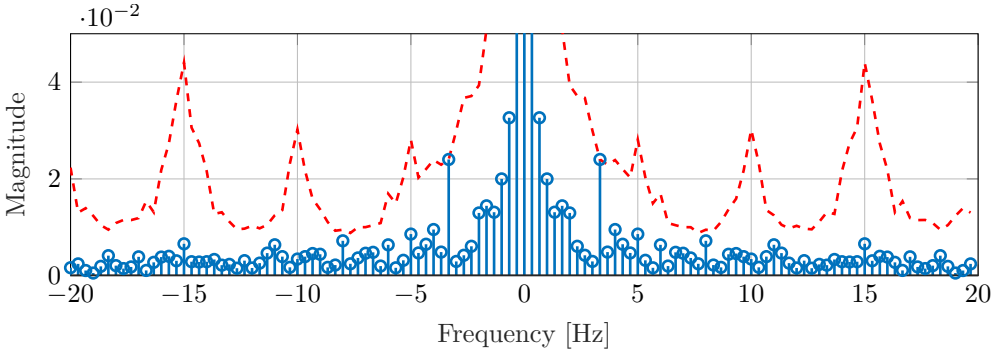

(c) Threshold Trained after 2 Hours of Flight (Training) Data

Figure 6.5: Illustration of Threshold Training

With the threshold trained using the rigorous process above, Equation (6.4.15) can then be used to determine the smallest amplitude OFC that can reliably be detected within 3 cycles. First, the unwanted frequency components are discarded from the threshold $T[k]$. These unwanted frequencies include all frequency components outside of the 1 – 10 Hz range, and all negative frequencies. Then, the effective threshold is the smallest amplitude for each frequency that can be reliably detected within c cycles. This can be calculated as

$$T_{eff}[k] = \frac{2fNT[k]}{cF_s + f}. \tag{6.4.16}$$

Figure 6.6 shows an example of the final threshold after training. For illustrative purposes, the threshold is doubled to show the smallest amplitude that would cross the threshold (since the magnitude of an OFC would appear in the DFT as $A/2$). Also shown is the effective threshold T_{eff}, which shows the smallest amplitudes over frequency required for 3 cycle detection time. This illustrates that the DFT can detect very small amplitude OFCs, but to guarantee detection time, the OFC amplitude has to be significantly larger than the threshold.

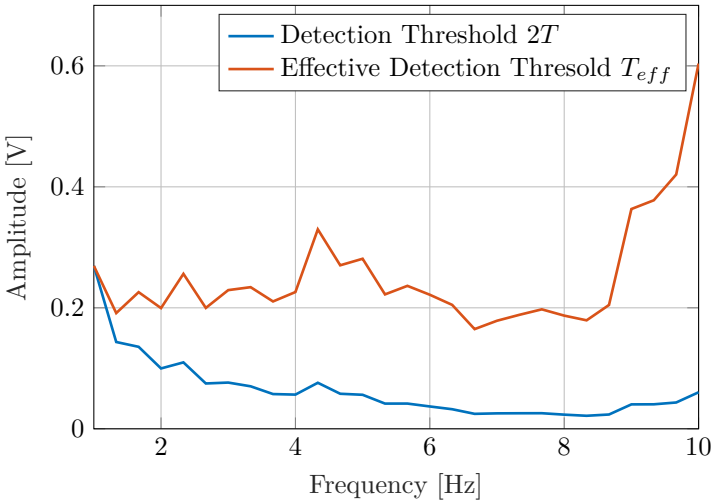

Figure 6.6: Frequency Dependent Threshold

The smallest amplitude that can be reliably detected within c cycles, irrespective of frequency, can be specified as

$$A_{min} = \max\left(T_{eff}[k]\right). \tag{6.4.17}$$

A_{min} thus represents a theoretical smallest OFC that can be detected within 3 cycles at any frequency, and is a representation of the worst case, noise notwithstanding. For all cases where $A \geq A_{min}$, the OFC will be detected within 3 cycles. Furthermore, much smaller amplitude OFCs can eventually be detected, although these may be detected later than allowed.

6.5 Achieving Sub-Sample Accuracy

Due to the way the DTFT is discretised, the DFT is defined only for discrete frequencies at $\omega_k = 2\pi k F_s/N$, and has a frequency resolution of F_s/N Hz. However, an OFC can manifest at any frequency over the continuous frequency range. This means that it is very likely that the frequency of an OFC will fall between two computed frequencies, ω_{k-1} and ω_k. Unfortunately, this may result in a missed detection, and is thus a motivation for frequency-domain sub-sample interpolation. This is most easily achieved using zero padding.

To illustrate the possibility of a missed detection, consider an OFC

$$r[n] = A\cos\left(\omega_i n T_s\right)$$

that falls between two frequency samples. Due to the rectangular window of the residual, the OFC manifests as a sinc function in the DTFT:

$$R(\omega) = \frac{1}{2}[A\cdot W(\omega - \omega_i) + A\cdot W(\omega + \omega_i)], \tag{6.5.1}$$

where $W(\omega)$, the DTFT of the rectangular window function, is defined as [36]

$$W(\omega) = \frac{\sin\left(\omega N/2F_s\right)}{N\sin\left(\omega/2F_s\right)} e^{-j\omega(N-1)/2F_s}. \tag{6.5.2}$$

Evaluating the DFT at the discrete frequencies ω_k gives

$$R(\omega_k) = \frac{1}{2}[A\cdot W(\omega_k - \omega_i) + A\cdot W(\omega_k + \omega_i)]. \tag{6.5.3}$$

The DFT and DTFT of a pure sinusoid is shown in Figure 6.7, along with the actual component of the OFC which lies between two frequency samples. It clearly shows that the spectral components

of the DFT do not reach the actual magnitude of the OFC. An extreme case occurs when ω_i falls exactly halfway between ω_{k-1} and ω_k.

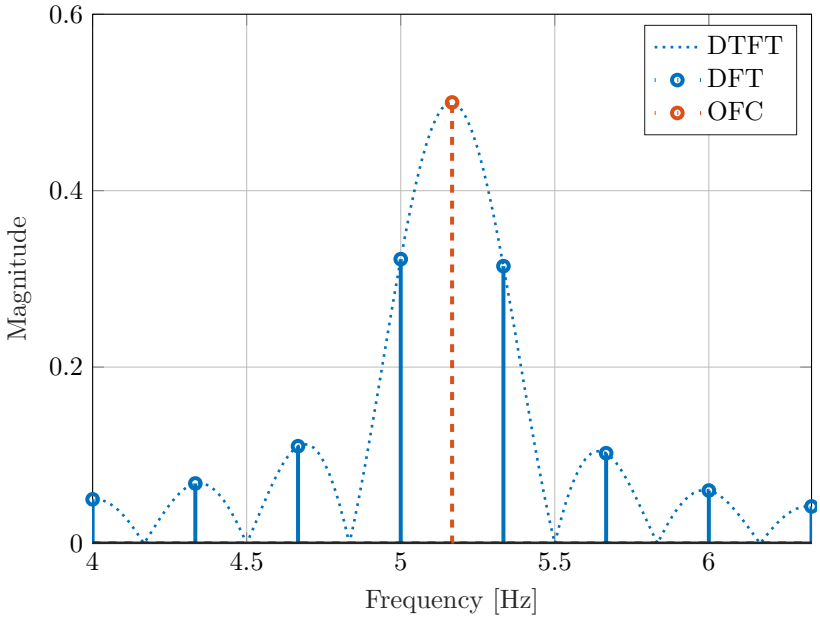

Figure 6.7: Spectral Resolution

Considering only the positive frequency component:

$$|A\cdot W(\omega_k-\omega_i)| = \left|\frac{A\sin\left((\omega_k-\omega_i)N/2F_s\right)}{N\sin\left((\omega_k-\omega_i)/2F_s\right)}\right| \tag{6.5.4}$$

Evaluating for $\omega_k-\omega_i = 2\pi(F_s/2N)$, where ω_i is exactly halfway between ω_{k-1} and ω_k, gives

$$|A\cdot W(\omega_k-\omega_i)| = \frac{A}{N\sin\left(\pi/2N\right)} \approx \frac{2A}{\pi}. \tag{6.5.5}$$

This result, which makes use of the small-angle approximation, shows that the effective amplitude of the OFC as calculated by the DFT could be reduced to 63.7 % of the actual amplitude.

Zero padding can bring out these critical spectral details that are hidden by the low spectral resolution. Zero padding increases the length of the window by some integer factor m_z by appending $(m_z-1)N$ zeros to the end of the residual window. The effect of this is to increase the number of samples in the window, thereby resulting in an increase in the number of samples in the frequency domain. These frequency-domain samples become more closely spaced, improving the resolution of the DFT. This unfortunately also increases the number of calculations that must be performed to calculate the DFT.

With zero padding, important hidden features emerge. An OFC that manifests at a frequency exactly halfway between two frequency samples may appear at only 63.7 % of its amplitude in the DFT, but with a zero-padding factor of 5, for example, this value increases to 98.4 %. In other words, there is a maximum potential amplitude drop of 1.6 %. A required zero-padding factor can be determined by specifying an allowable amplitude drop:

$$\%\text{ amplitude drop} \approx \left(1-\frac{\sin\left(\pi/2m_z\right)}{\pi/2m_z}\right)\times 100 \tag{6.5.6}$$

Figure 6.8 shows the result of zero padding.

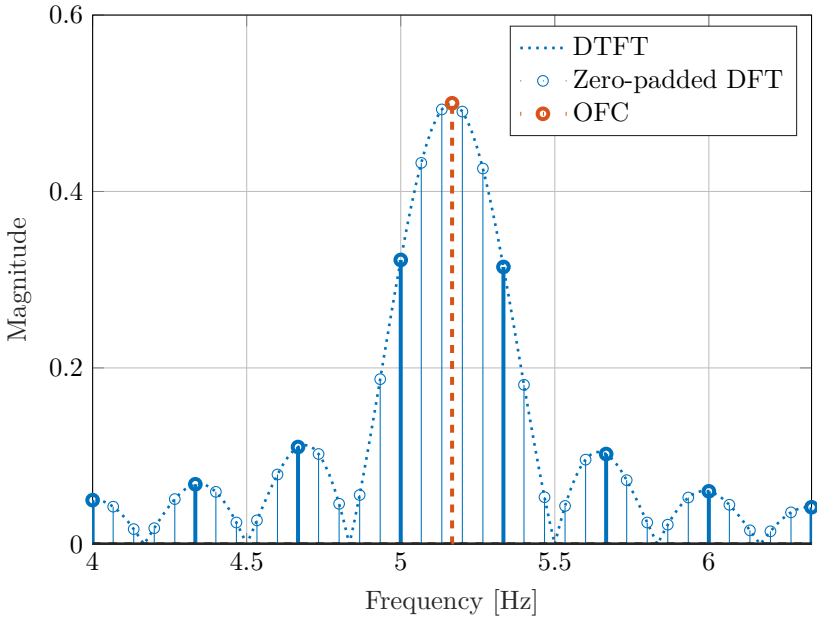

Figure 6.8: Spectral Resolution with Zero-padding Factor of 5

The above result shows that the effects of frequency-domain resolution can successfully be mitigated through the use of zero padding.

At this point, the DFT approach to OFC detection can be summarised as follows:

1. At each time step, obtain the most recent N samples of the residual.
2. Append $(m_z - 1)N$ zeros to the end of the residual window.
3. Calculate the $(m_z N)$-point DFT of the residual window.
4. For each frequency component between 1 and 10 Hz, compare the magnitude of the DFT to a frequency-dependent threshold.

However, this approach is clearly more computationally expensive than the simple threshold crossing checks executed by the oscillation counting method. For feasible implementation, it is necessary to investigate more efficient DFT calculation algorithms and other performance improvements that can be realised.

6.6 Sliding Discrete Fourier Transform (SDFT)

Sections 6.4 and 6.5 offer a compelling argument for the use of Fourier-domain techniques, but unfortunately the use of the DFT comes with significant computational complexity. Calculating the DFT can be efficiently performed using the fast Fourier transform (FFT) algorithm, but this may still not offer sufficient performance. Fortunately, there exists a more computationally efficient approach to calculating the DFT using a rolling window, namely the sliding discrete Fourier transform (SDFT). The SDFT, illustrated in Figure 6.9, is defined as [37]:

$$R[k,n] = e^{j2\pi k/N}(R[k,n-1] - r[n-N] + r[n]) \tag{6.6.1}$$

The derivation of this equation can be found in Appendix A. With normalisation, the SDFT is

$$R[k,n] = e^{j2\pi k/N}\left(R[k,n-1] - \frac{r[n-N]}{N} + \frac{r[n]}{N}\right). \tag{6.6.2}$$

The SDFT is a recursive approach to computing the k^{th} frequency component, or bin, and uses the result of the previous time step to update the current time step. This process requires fewer overall operations than the FFT. Furthermore, since each frequency bin k does not rely on any other frequency samples, $R[k, n]$ can be calculated for a select few values of k. This means that only the 1 to 10 Hz frequency components need to be computed, and the calculation of the rest, including the negative frequency components, can be ignored. The FFT, on the other hand, is forced to calculate all components, irrespective of which frequencies are of interest and which are not.

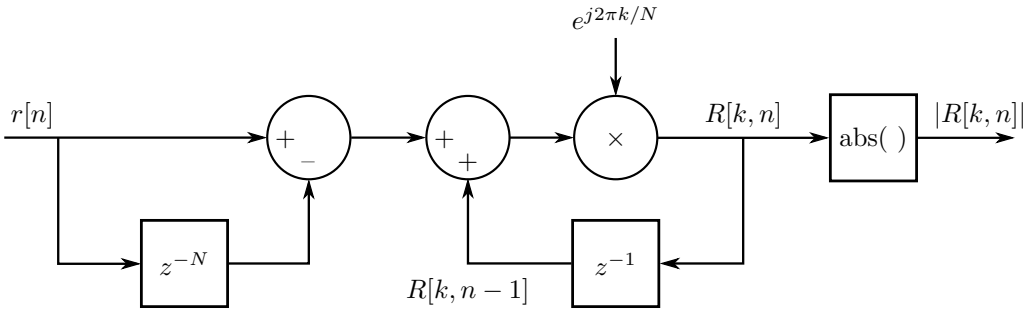

Figure 6.9: Block Diagram of SDFT

The computational complexity of the DFT, FFT, and SDFT is summarised in Table 6.1 [36] in terms of the number of complex operations that must be performed. Here, it is shown that the SDFT requires significantly fewer operations than either the FFT or the DFT.

Table 6.1: Computational Complexity of Discrete Fourier Transform Algorithms

	Complex Multiplications	Complex Additions
DFT	N^2	$N(N-1)$
Radix-2 FFT	$(N/2)\log_2 N$	$N\log_2 N$
SDFT	N	$2N$

Certain performance improvements can further reduce the computational burden of the SDFT. The complex value $\exp(j2\pi k/N)$ remains constant for a particular k, and can be computed beforehand and stored as a lookup table. Another optimisation can be found by avoiding the calculation of frequency bins that are not of interest. With $N = 120$, $F_s = 40$ Hz, and $f \in [1, 10]$, the total number of frequency bins that have to be computed are

$$\text{no. bins} = \frac{N(f_{max} - f_{min})}{F_s} + 1 = 28. \tag{6.6.3}$$

This is a substantial improvement over the FFT's 120 bins.

The SDFT can also be modified to calculate the DFT of zero-padded data. The DFT of a zero-padded window is

$$R[k] = \sum_{n=0}^{N-1} r[n] e^{-j2\pi kn/m_z N}. \tag{6.6.4}$$

By following a similar derivation process as given in Appendix A, the algorithm simplifies to

$$R[k, n] = e^{j2\pi k/m_z N}\left(R[k, n-1] - r[n-N] + r[n]e^{-j2\pi k/m_z}\right). \tag{6.6.5}$$

With normalisation,

$$R[k,n] = e^{j2\pi k/m_z N}\left(R[k,n-1] - \frac{r[n-N]}{N} + \frac{r[n]}{N}e^{-j2\pi k/m_z}\right). \tag{6.6.6}$$

Once again, zero padding comes with an increased computational cost. With a zero-padding factor of 5, the number of frequency bins that must be computed increase from 28 to 136, and adds an additional complex multiplication.

Two final points regarding the use of the SDFT must be acknowledged. Firstly, initialisation of the SDFT is extremely important. This is due to the recursive nature of the algorithm. One approach to initialisation is to calculate the FFT of the first N samples, and continue using the SDFT with this FFT result. Another approach is to initialise $R[k,0]$ as all zeros, and to prepend the residual data with N zeros.

The second, more critical point regarding the SDFT is that the SDFT equation contains a pole on the unit circle, and is therefore marginally stable. This can unfortunately lead to instability if any computational errors arise. One approach to combating this instability is to add a damping factor r_d:

$$R[k,n] = r_d \cdot e^{j2\pi k/N}(R[k,n-1] - r[n-N] + r[n]), \quad 0 << r_d < 1 \tag{6.6.7}$$

Alternative, more accurate techniques have been suggested by Duda [38] and Douglas and Soh [39]. However, all these methods come with some form of approximation error, and increased computational complexity. They may also complicate any attempt to achieve sub-sample accuracy as described in Section 6.5.

6.7 Fourier Analysis using Multiple Window Lengths

According to Equation (6.4.15), the frequency-domain threshold has to be a function of frequency, and as seen in Figure 6.4, the threshold can get quite small at higher frequencies. However, Equation (6.4.15) also suggests that instead of forcing the threshold T to adapt to f, the window length N can be varied based on f. This means that higher frequency components could rather be calculated using shorter windows than the low frequency components. Equation (6.4.15) therefore motivates an investigation into the possibility of varying both the threshold and the window size over frequency, rather than just the threshold. This idea is based on the fact that for high-frequency OFCs, more than 3 cycles can fit in the window. Thus using shorter windows could also lead to faster detection times.

This leads to the concept of the multi-window Fourier transform (MWFT), an implementation of the SDFT where shorter windows are used to calculate the magnitude of higher frequency bins. Furthermore, the resolution of the SDFT becomes a function of frequency. By using shorter windows, the resolution, or gap between bins increases. This approach is made possible by the SDFT algorithm, where each frequency bin is calculated independently, and different values of N can be chosen for each bin.

As an example, consider a case where 4 different window sizes are used: 120, 80, 40, and 20 samples. Figure 6.10 illustrates the MWFT of a signal, along with zero padding.

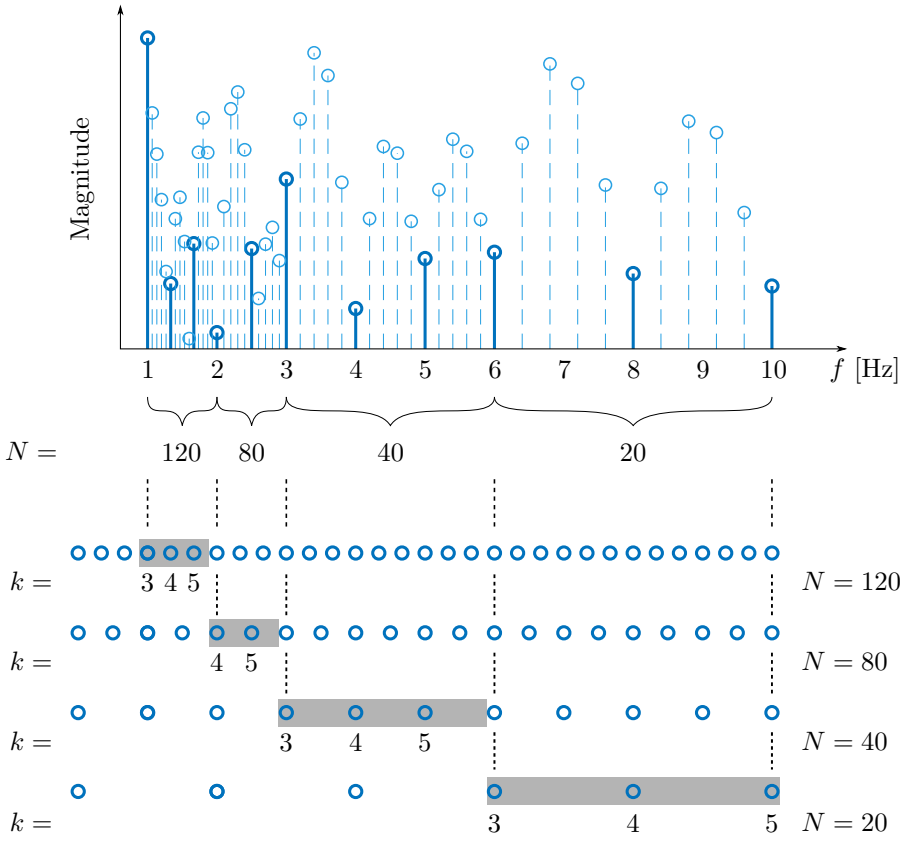

Figure 6.10: Frequency Components of MWFT

Figure 6.10 specifies which window sizes are used for which frequency components, and which values of k to use to calculate the MWFT. The challenge of choosing window sizes is to ensure smooth transitions between the end of one frequency sub-band and the start of another, especially when zero padding is involved. Here, it was chosen that the last frequency sample of one sub-band falls on the same frequency as the first frequency sample of the next sub-band. For example, the $k = 6$ sample of the $N = 120$ window falls on the same frequency as the $k = 4$ sample of the $N = 80$ window. One final consideration is that at least 3 cycles should fit in the window. Recalling that k represents the number of complete cycles in the window, it is clear in Figure 6.10 that there are always at least 3 cycles in the chosen windows.

To illustrate the effect of the MWFT on oscillatory residual, Figure 6.11 compares 4 separate DFTs to the MWFT. A 5 Hz OFC exists in the residual, and its translation to the frequency domain using the 4 different window sizes are shown in the left-hand plot. The MWFT, shown on the right, attempts to merge the results of these 4 DFTs into a single signal that can be thresholded.

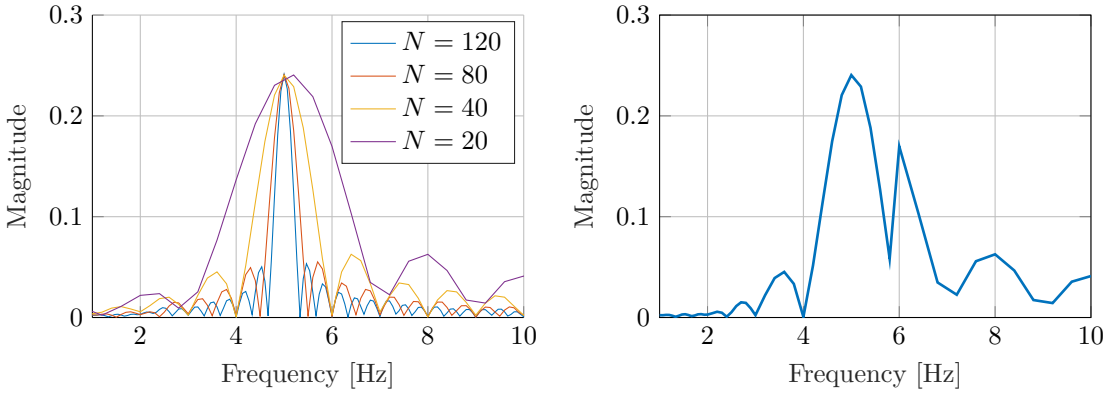

Figure 6.11: 4 DFTs (Left) and the corresponding MWFT (Right), with 5 Hz OFC

Figure 6.11 illustrates an additional effect, which is that for many frequencies the DFT of the $N = 20$ window appears to be higher than that of the $N = 120$ window. To observe the effect of the MWFT on detectable amplitudes, the threshold of the MWFT is compared to that of a single window DFT with $N = 120$. This is shown in Figure 6.12.

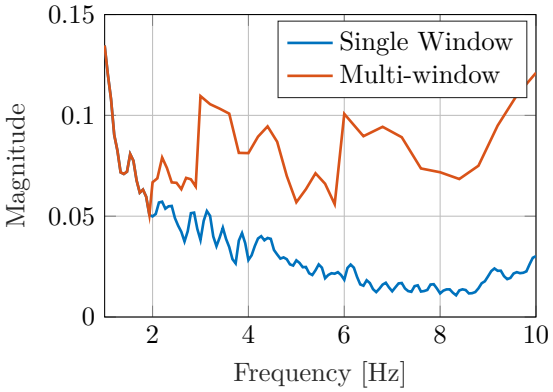

Figure 6.12: Trained Thresholds for Single and Multi-windowed Fourier Transforms, with $m_z = 5$

The trained threshold of the MWFT appears to be higher overall than the single-window approach, primarily due to the reduced spectral resolution. This suggests that the overall threshold increase could cause the MWFT to miss small OFCs that would otherwise be caught by the standard DFT. However, the effective threshold can be determined with Equation (6.4.15) to find the smallest OFCs that can be detected within 3 cycles. This is shown in Figure 6.13, and it is clear that the MWFT shows no loss in minimum detectable amplitude, and sometimes even shows slight improvement over the single-window Fourier transform.

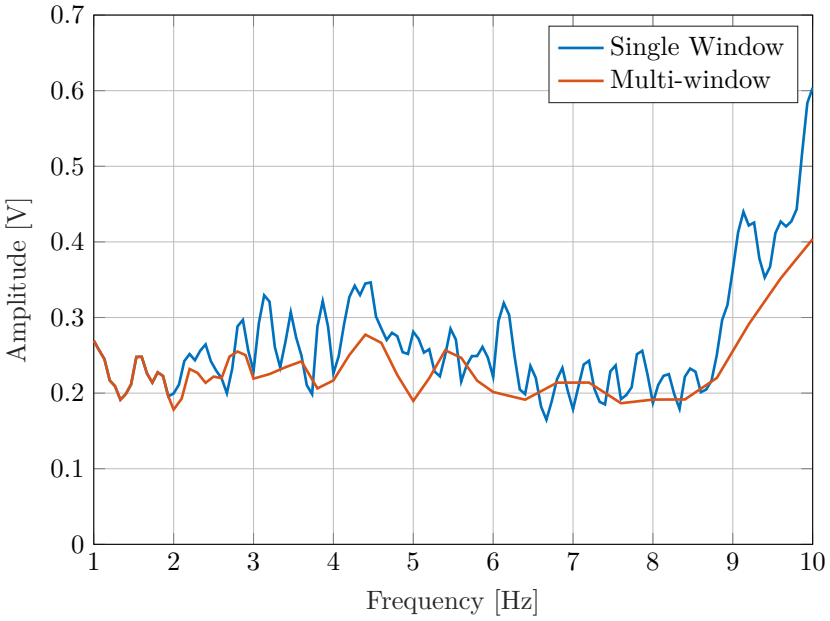

Figure 6.13: Effective Threshold for Single and Multi-window Fourier Transforms

The use of the MWFT approach can also improve the computational performance of the SDFT, as fewer bins have to be computed. With the selected windows, the number of components reduces to 51 with zero padding, or 11 without. This is in comparison to the 136 with zero padding or 28 bins without for the traditional SDFT, as stated in Section 6.6.

Overall, the MWFT offers an alternate application of the DFT to the OFC detection problem with major reductions in computational cost and memory requirements without compromising detection time. Ultimately, this design should be considered as experimental, and further investigation into its propensity towards false alarms and missed detections should be conducted. However, it is more thoroughly tested in Chapter 8.

6.8 Summary and Contributions

In summary, the following contributions were made in this chapter:

- The discrete Fourier transform was implemented as an approach to OFC detection. Initial attempts made use of Matlab's FFT function, and used a single, fixed threshold value for all frequency components.
- Nominal residual signals were noted to have less energy in the higher frequencies than in the lower frequencies. This led to the concept of the frequency-dependent threshold, which used a unique threshold value for each frequency component. This threshold was trained by determining the maximum value of each frequency component over a large amount of fault-free training data.
- Following the advent of the frequency-dependent threshold, a mathematical proof was derived that linked the detection time of the approach to the amplitude and frequency of the OFC, and the chosen threshold. This further motivated the use of a frequency-dependent threshold, and provided a way to gauge the expected performance of the DFT.
- A shortcoming based on the resolution of the DFT was identified, where OFCs with frequencies that fall between the DFT's frequency samples would not exhibit their full magnitudes. Zero padding was discussed and proposed as a solution.

- The sliding-DFT was identified as a computationally efficient approach to the DFT calculation that would be less computationally expensive than the FFT. It was subsequently implemented, and its result was confirmed to be identical to Matlab's FFT function. The SDFT was also confirmed to allow for zero padding.
- The structure of the SDFT was exploited to develop an experimental multi-window approach for the estimation of the residual's spectral content. This multi-window Fourier transform allowed for the use of smaller windows to calculate higher-frequency components, leading to faster detection times for high-frequency OFCs, and further improved the computational efficiency.

Frequency-domain techniques offer an attractive means of OFC detection by extracting the frequency content of the residual signal. The sliding discrete Fourier transform is a computationally inexpensive approach to estimating the short-time frequency content of a signal, and can potentially be implemented in the embedded monitoring systems on the Airbus A380. However, if the SDFT cannot be used for stability reasons, the fast Fourier transform algorithm can be implemented instead, although with an increase in computational cost.

The use of the discrete Fourier transform has significant advantages over time-domain techniques. Firstly, no filters are required. This means that no phase delays or transient effects exist to influence the detection time. Detection time can be guaranteed fairly accurately. Furthermore, OFCs will almost always be detected before 3 cycles have passed. More severe OFCs tend to cross the threshold sooner, based on Equation (6.4.15). However, this approach is admittedly more computationally expensive than oscillation counting. The robustness of the detection system is also highly dependent on the threshold training.

Chapter 7

Demodulation-Based Detection Technique

7.1 Introduction

The oscillatory failure case (OFC) detection problem can be viewed simply as the detection of sinusoids in the presence of noise. This problem has already been thoroughly investigated and its solution has been applied in the field of telecommunications. These techniques may offer novel solutions to the OFC detection problem.

This chapter investigates the use of a phase-locked loop (PLL) for OFC detection. A PLL attempts to track the phase and frequency of a sinusoidal signal, and can provide an indication of the dominant frequency component of a signal. As such, they are typically used for carrier recovery, clock synchronisation, and phase and frequency modulation [40]. Overall, the pedigree of the PLL provides strong evidence that this is an avenue worth investigating.

This chapter will introduce the reader to the basic theory of the PLL, as well as potential pitfalls presented by certain non-linear effects. Section 7.2 provides essential background theory on phase-locked loops. Section 7.3 details the design of the PLL-based OFC detection algorithm. Section 7.4 presents an analysis of the theoretical signal-to-noise ratio performance of the PLL design. Section 7.5 describes how the detection threshold of the PLL is trained. Finally, Section 7.6 considers alternate designs that are noteworthy, but ultimately not used.

7.2 Phase-Locked Loop Theory

7.2.1 Overview

A phase-locked loop consists of four basic components [40]:

- A phase detector (PD), which determines the phase difference between two sinusoidal input signals.
- A voltage-controlled oscillator (VCO), which outputs an oscillating signal whose frequency is controlled by an input voltage.
- A loop filter, which removes high-frequency components from the phase detector's output, and acts as a controller that can be designed to some required specifications using control theory.
- A feedback loop, which feeds the output of the VCO to the phase detector.

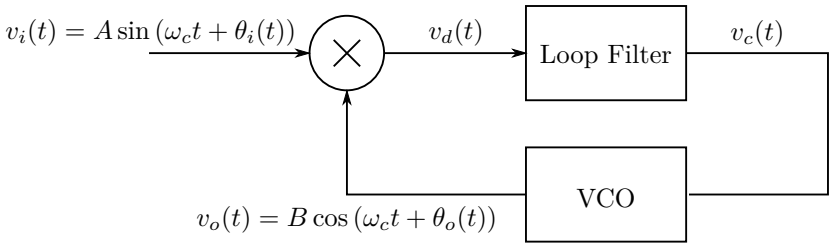

Figure 7.1: Overview of Phase-locked Loop Operation

The PLL system is illustrated in Figure 7.1. The phase detector, which in this case is a simple multiplier, determines the phase error between the input signal and the VCO's output signal. This phase error is filtered by the loop filter, and then used to control the VCO to reduce the phase error (and consequently the frequency error) to some constant value, or to zero. This allows the PLL to "lock" onto the frequency of the input signal.

The PLL is a non-linear system, in which both the phase detector and the VCO are non-linear devices. However, for most design purposes a linear model of the system is assumed, with some non-linear analysis performed to ensure that the system operates as expected.

7.2.2 Linear Modelling

The derivation of the linear PLL model is provided in the works of Gardner [41] and Lathi [20], and is investigated and summarised here.

Assume the input to the PLL is a sinusoidal signal with frequency ω_i and initial phase offset ψ:

$$v_i(t) = A\sin(\omega_i t + \psi) \tag{7.2.1}$$

The VCO generates a sinusoidal signal described as

$$v_o(t) = B\cos(\omega_c t + \theta_o(t)). \tag{7.2.2}$$

The frequency ω_c is known as the free running, or centre frequency of the VCO, and is associated with a 0 V input to the VCO. The instantaneous frequency of the VCO's output is

$$\omega_o = \omega_c + \dot{\theta}_o(t) \tag{7.2.3}$$

where $\dot{\theta}_o(t)$ is the frequency deviation from the centre frequency,

$$\dot{\theta}_o(t) = cv_c(t) \tag{7.2.4}$$

where, $v_c(t)$ is a reference voltage provided by the loop filter and c is a conversion factor from V to rad/s.

The phase of the input signal $v_i(t)$ can be expressed as a function of the centre frequency of the VCO:

$$\begin{aligned}\omega_i t + \psi &= \omega_c t + (\omega_i - \omega_c)t + \psi \\ &= \omega_c t + \theta_i(t)\end{aligned} \tag{7.2.5}$$

It is also useful to define the frequency deviation between the incoming signal and the VCO's free-running frequency:

$$\Delta\omega = \omega_i - \omega_c \tag{7.2.6}$$

This value is ultimately the offset that has to be applied to the VCO to match its output frequency with the input frequency. To obtain this value and control the VCO, the phase error

$$\theta_e(t) = \theta_i(t) - \theta_o(t) \tag{7.2.7}$$

is desired. This can easily be obtained through simple multiplication:

$$\begin{aligned} v_d(t) &= A\sin\left(\omega_c t + \theta_i(t)\right) \times B\cos\left(\omega_c t + \theta_o(t)\right) \\ &= \frac{1}{2}AB\left[\sin\left(\theta_i(t) - \theta_o(t)\right) + \sin\left(2\omega_c + \theta_i(t) + \theta_o(t)\right)\right] \end{aligned} \tag{7.2.8}$$

This multiplication is known as a sinusoidal phase detector, and its output contains two terms: the first is the sine of the phase error $\theta_e(t)$, while the second is a frequency component at twice the VCO centre frequency. This second term is referred to as the double frequency term. Since only the phase error is desired, a lowpass loop filter $F(s)$ is needed to reject this double frequency term. If it is assumed that the loop filter completely removes this high-frequency component, the input voltage to the VCO, $v_c(t)$, can be expressed as

$$v_c(t) = f(t) * \frac{1}{2}AB\sin\theta_e(t). \tag{7.2.9}$$

Common practice at this point is to assume that small angle approximations are applicable, and to let $\sin\theta_e(t) \approx \theta_e(t)$. This linearisation simplifies the process of designing the loop filter for some desired transient response. Taking Equation (7.2.9) to the Laplace domain:

$$V_c(s) = \frac{1}{2}ABF(s)\Theta_e(s) \tag{7.2.10}$$

Similarly, Equation (7.2.4) can also be taken to the Laplace domain, and with a little rearranging, the transfer function of the VCO is obtained.

$$G(s) = \frac{\Theta_o(s)}{V_c(s)} = \frac{c}{s} \tag{7.2.11}$$

Thus, the linearised VCO is quite simply an integrator and a gain factor. Figure 7.2 illustrates this linearised model.

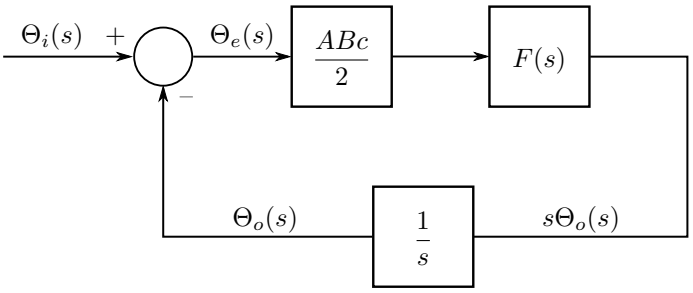

Figure 7.2: Linear Phase-locked Loop Model

The linear closed-loop transfer function of the PLL is

$$\frac{\Theta_o(s)}{\Theta_i(s)} = \frac{\frac{1}{2}ABcF(s)}{s + \frac{1}{2}ABcF(s)} \tag{7.2.12}$$

It is important to note that the closed-loop specifications of the PLL is dependent on the amplitude of both the input signal, A, and of the VCO, B. While the VCO's amplitude will be constant, the amplitude of the residual may vary drastically, changing the dynamics of the system. This must therefore me taken into account during the design stage. Bandpass limiters are an effective means to remove a signal's amplitude information while still maintaining phase information, and Section 7.5 provides more information about the input limiter used in the final design.

7.2.3 Steady-state Analysis

As a quick interlude, a brief explanation of system type is given here. A type 1 closed-loop system contains a single free integrator in the control loop, and can therefore track step inputs with zero steady-state error, and ramp inputs with finite steady-state error. Similarly, a type 2 system has two free integrators in the control loop, and can track step and ramp inputs with zero steady-state error.

Steady-state analysis of the PLL is performed with the final value theorem:

$$\theta_{eq} = \lim_{t\to\infty} \theta_e(t) = \lim_{s\to 0} s\Theta_e(s) \tag{7.2.13}$$

where θ_{eq} represents the steady-state phase error, or equilibrium point of θ_e. Given the closed-loop transfer function of the PLL in equation 7.2.12, the Laplace transform of the phase error can be expressed as

$$\Theta_e(s) = \frac{s}{s + \frac{1}{2}ABcF(s)}\Theta_i(s) \tag{7.2.14}$$

Since the VCO is essentially a free integrator, the system is at least of type 1. Thus, for a phase step $\Delta\theta$ in the input, the phase error, and therefore the phase detector output, will strive to zero.

A step in frequency, $\Delta\omega$, translates to a ramp change in the input phase. Such an input will result in a constant phase error in a type 1 system. This steady-state error is proportional to the change in frequency, as shown using the final value theorem:

$$\theta_{eq} = \lim_{s\to 0}\left[s\frac{s}{s + \frac{1}{2}ABcF(s)}\frac{\Delta\omega}{s^2}\right] = \frac{\Delta\omega}{\frac{1}{2}ABcF(0)}. \tag{7.2.15}$$

7.2.4 Lock Detection

It is sometimes necessary to determine whether the PLL has detected and locked onto a signal, or is just reacting to noise. Indeed, this is the essence of the OFC detection problem. This can be achieved through the use of a quadrature phase detector (QPD) [41], [42].

The QPD does this by multiplying the PLL's input signal with a 90° phase shifted version of the VCO's output. This process is similar to that of the sinusoidal phase detector, and is described as

$$\begin{aligned}\varphi(t) &= A\sin\left(\omega_c t + \theta_i(t)\right) \times B\sin\left(\omega_c t + \theta_o(t)\right) \\ &= \frac{1}{2}AB\left[\cos\left(\theta_i(t) - \theta_o(t)\right) + \cos\left(2\omega_c t + \theta_i(t) + \theta_o(t)\right)\right].\end{aligned} \tag{7.2.16}$$

Once the PLL has locked onto the incoming signal, the error θ_e in the first term approaches a constant value, while the double frequency term above can be removed with a lowpass filter. Due to the cosine in the first term, the QPD output approaches a maximum DC value of $AB/2$ as θ_e approaches zero. By selecting the VCO's output amplitude B, and choosing a minimum amplitude A that should be detectable, a threshold value can be chosen.

7.2.5 Non-linear Analysis

While the linearised system described above is sufficiently accurate for closed-loop feedback design, the inherent non-linearities of the system must be evaluated to ensure that the system works as expected, and that non-linear effects will not interfere with the system's detection performance. Two non-linear PLL concepts are introduced here, namely the hold-in range and lock-in range. These non-linearities stem from the sinusoidal characteristics of the phase detector, and limit the range of frequencies the PLL is able to track. Leonov, Kuznetsov, Yuldashev *et al.* proposed mathematically rigorous definitions for these entities. For this analysis, the double frequency is assumed to be completely removed. Figure 7.3 illustrates the assumed non-linear model, where the output of the phase detector is now considered to be the sine of the phase error, and small-angle approximation is no longer assumed.

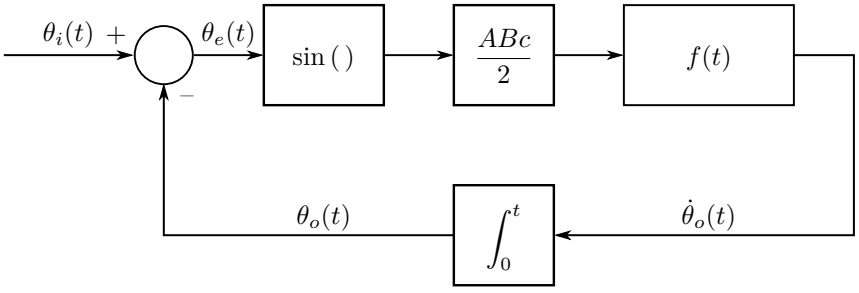

Figure 7.3: Non-linear Phase-locked Loop Model

Non-linear analysis of PLLs is often performed through simulation and phase-plane plots [40], and these methods are employed here. Alternative methods that have been used include the Lyapunov Redesign [44] and the Circle/Popov Criteria [45].

7.2.5.1 Hold-in Range

The hold-in range $\Delta\omega_H$ is the frequency range for which the PLL is able to maintain phase tracking. More specifically for sinusoidal phase detectors, the frequency range is determined by calculating the frequency offset $\Delta\omega$ that causes the phase error to exceed the range of linear analysis [40].

Assume that $v_i(t)$, the input to a type 1 PLL, experiences a sudden step in frequency $\Delta\omega$. The PLL reacts by forcing the VCO input voltage $v_c(t)$ to some constant value that is directly proportional to $\Delta\omega$.

$$\lim_{t\to\infty} v_c(t) = \frac{\Delta\omega}{c} \tag{7.2.17}$$

where $\Delta\omega$ is the frequency deviation between the PLL's input and the VCO's centre frequency, according to Equation (7.2.6)

From Equation (7.2.10), and taking the sine of the phase error into account, the steady state value of $v_c(t)$ can also be expressed as

$$\lim_{t\to\infty} v_c(t) = \frac{1}{2}ABF(0)\sin\theta_{eq}. \tag{7.2.18}$$

Thus, the sine of the steady-state phase error is

$$\sin\theta_{eq} = \frac{\Delta\omega}{\frac{1}{2}ABcF(0)}. \tag{7.2.19}$$

However, since $|\sin\theta_{eq}| \leq 1$, the frequency deviation $\Delta\omega$ the PLL can track is limited to

$$|\Delta\omega| \leq \frac{1}{2}ABcF(0), \tag{7.2.20}$$

where A is the amplitude of the input signal, B is the amplitude of the VCO output signal, c is the gain of the VCO, and $F(0)$ is the steady-state gain of the loop filter. This results shows that the tracking of a large change in frequency in $v_i(t)$ may require a $v_c(t)$ value that is too large for the phase detector's limited output. If the input frequency exceeds this limit, the PLL will lose lock and begin to show a "ringing" effect in $v_c(t)$, and subsequently in the frequency of $v_o(t)$. Fortunately, this limit can be adjusted by changing the gain values in the system, such as the gain of the loop filter or of the VCO, but it may provide a design challenge for systems requiring large bandwidths.

Note that this effect is limited to type 1 PLLs. If the loop filter contains an integrator, the DC gain is infinite, and the PLL can theoretically lock onto any frequency. However, if the frequency step is too large, the error will slip cycles for a while until lock is achieved.

7.2.5.2 Lock-In Range and Cycle Slipping

Type 1 loops are limited by their hold-in range, and will lose lock permanently while the frequency deviation $\Delta\omega$ exceeds this hold-in range. Type 2 loops with infinite DC gain are not limited by this hold-in range, but instead can temporarily lose lock due to non-linear transient effects known as cycle skipping or cycle slipping.

The conditions under which the system will slip cycles is dependent both on the phase error and frequency error. The effect is best visualised using phase portraits. An example is given in Figure 7.4.

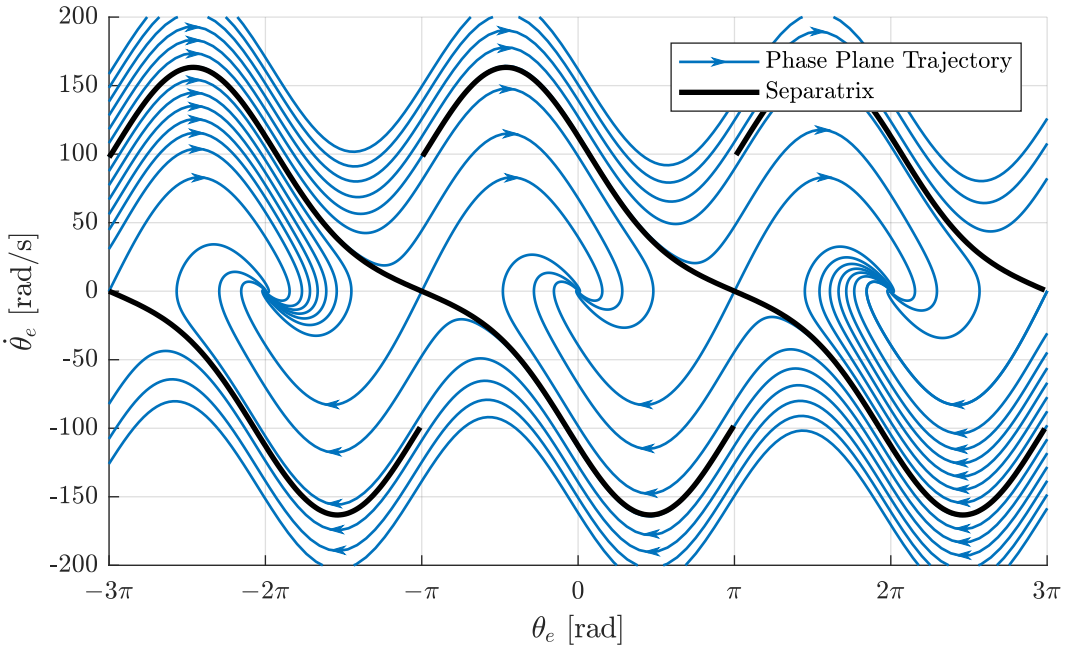

Figure 7.4: Phase Plane Trajectories of a Type 2 PLL, Adapted from [41]

Figure 7.4 illustrates that the trajectories terminate at equilibrium points θ_{eq} that exist at $\theta_e = 2\pi k$ for $k \in \mathbb{Z}$. This shows that for any initial phase or frequency error, the system responds by reducing the phase error to zero. Unstable saddle points exist at $\theta_e = \pi + 2\pi k$. The separatrices, marked in thick black lines above, designate the initial conditions for which the PLL will instantaneously lock onto the nearest equilibrium point.

If an initial condition on the phase plane lies between the upper and lower separatrices on some 2π interval, its trajectory will strive towards the equilibrium point at the centre of that interval. If an initial condition lies outside the separatices, the PLL will slip cycles until the trajectory eventually crosses a separatrix and reaches an equilibrium point at a different 2π interval [41]. As such, cycle slipping can be formally defined as:

$$\left| \lim_{t\to\infty} \theta_e(t) - \theta_e(0) \right| \geq \pi \tag{7.2.21}$$

Alternate definitions for cycle slipping and lock-in range have been suggested in literature. Leonov, Kuznetsov, Yuldashev *et al.* [43] defines cycle slipping as

$$\sup_{t>0} |\theta_e(0) - \theta_e(t)| > 2\pi, \tag{7.2.22}$$

while Hsieh and Hung [46] offers an approximation of the lock-in range $\Delta\omega_L$ as

$$\Delta\omega_L \approx \pm\frac{1}{2} ABcF(\infty) \tag{7.2.23}$$

where $F(\infty)$ represents the frequency response of the loop filter at infinite frequency.

Cycle slipping occurs because the sinusoidal phase detector cannot output a large enough value for the loop to lock instantly, and consequently the integrator stage in the loop filter has to build up to match the frequency offset. This causes θ_e to exceed 2π. As θ_e increases, the phase detector output tends to oscillate, or ring, until the loop locks. The result of this is an increase in lock time, which is considered to be unacceptable for the application of OFC detection.

Figure 7.5 illustrates the effect of cycle slipping on the PLL, and the resultant increase in convergence time. It's clear in Figure 7.5a that θ_e quickly settles to 0 rad when the input frequency is within the lock-in range. Figure 7.5c shows that when a frequency outside of the lock-in range is introduced, θ_e has to build up to a large value for the PLL to achieve lock, and eventually θ_e settles to 4π rad. The effects on the VCO input $v_c(t)$ are shown in Figures 7.5b and 7.5d. When cycle slipping occurs, v_c rings before settling to its steady-state value, resulting in a doubling of the settling time.

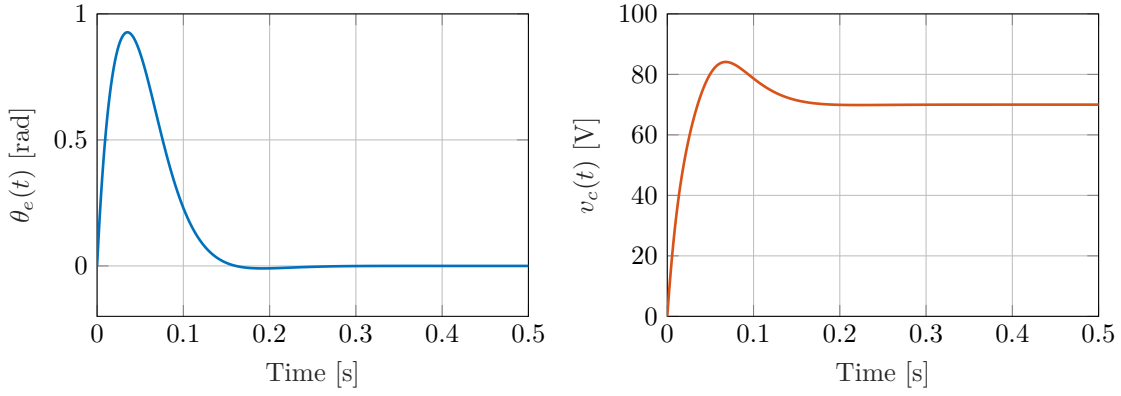

(a) Phase Error without Cycle Slipping

(b) VCO Input without Cycle Slipping

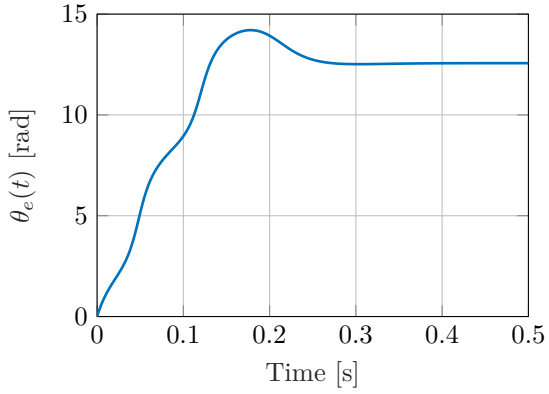

(c) Phase Error with Cycle Slipping

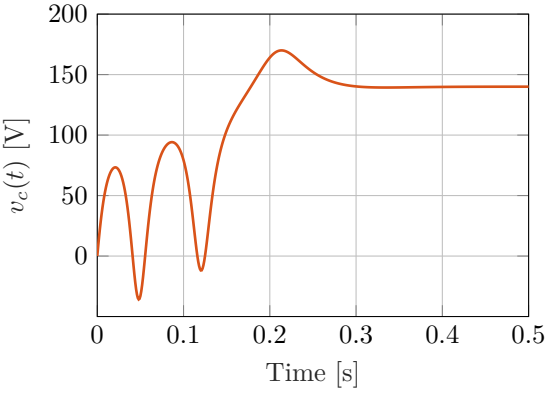

(d) VCO Input with Cycle Slipping

Figure 7.5: Comparison of PLL Operation without (Top) and with (Bottom) Cycle Slipping

Using simulation and phase plane analysis, the lock-in range $\Delta\omega_L$ can be determined. This process is more thoroughly explored in Section 7.3.2.

7.3 Phase-Locked Loop OFC Detection Design

With all the PLL theory established, the design of the phase-locked loop detection method can be described. The initial design of the PLL detection method was based on the concept of frequency demodulation, and attempted to detect OFCs by monitoring the VCO's input voltage, v_c. In theory, a large offset in v_c would indicate the presence of a dominant frequency component at a

frequency "far" from the VCO's centre frequency. OFCs would be detected by checking for large offsets in v_c. If v_c remained approximately zero, no failure was present, or a failure was present with a frequency approximately equal to the centre frequency. Therefore, the approach would use multiple PLLs with overlapping frequency bands to cover the whole 1 to 10 Hz spectrum.

This approach was ultimately flawed. During fault-free operation, v_c would continuously attempt to track phase and frequency changes in the noisy residual signal, often resulting in large transients that were guaranteed to cause false alarms. As an alternative, a new approach using the lock detector was proposed, which eventually led to the final design of the detection system. Figure 7.6 shows the architecture of the proposed PLL detection method.

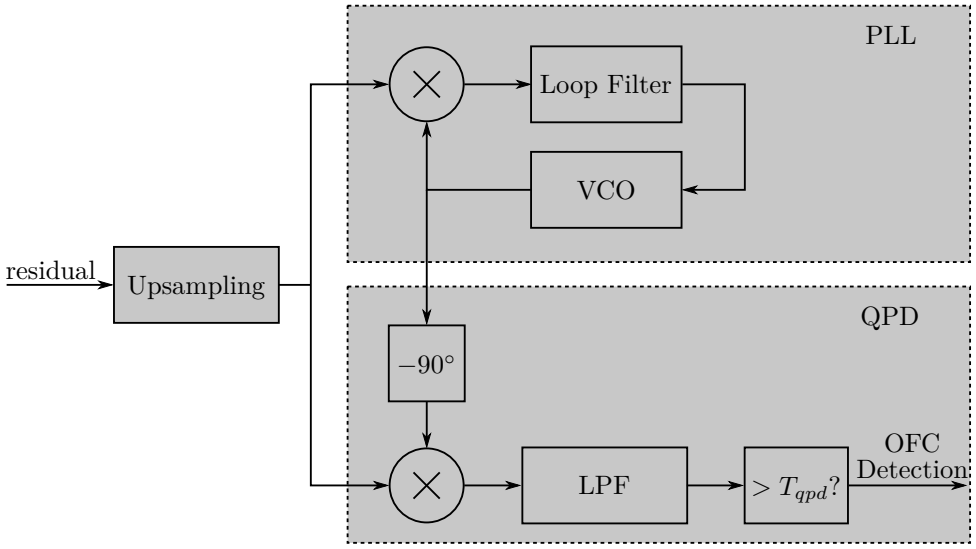

Figure 7.6: Overview of Phase Locked Loop Detection Algorithm

The proposed approach consists of three main stages. The first stage involves upsampling and modulating the residual signal to a higher frequency band. The PLL is the second stage, and attempts to track any OFC signal that may be present. The final stage, the QPD, checks if the PLL has locked onto some dominant frequency, and provides the confirmation of the existence of an OFC. These stages are explained in more detail in this section. Alternate approaches and additional design considerations are summarised in Section 7.6.

7.3.1 Upsampling

The fact that OFC frequencies of interest are in the 1 to 10 Hz range complicates the design of the PLL. The frequencies to be detected span an entire decade on a logarithmic scale, and the PLL is required to provide detection confirmation within as little as 0.3 s (for a 10 Hz case), requiring a fast converging loop filter. As a result, the loop filter must have a low cutoff frequency for accurate frequency tracking and removal of the double frequency term, but also a fast settling time. These requirements are in direct conflict with each other.

Furthermore, the effectiveness of the QPD is largely influenced by the amplitude of the residual. However, with a sampling rate of 40 Hz, a 10 Hz OFC consists of only 4 samples per cycle, and amplitude information is not fully conveyed through the multiplier.

Noting that PLLs are mostly used in high-frequency applications, where the frequencies of interest are much higher than the loop filter cutoff frequency [40], many of the issues listed above may be negated if the frequencies of the residual were simply "moved" to a higher frequency band. This would place the centre frequency considerably higher than the loop filter's cutoff frequency, and would reduce the effect of the logarithmic scaling issues.

Upsampling, a method of interpolating new samples between the existing samples, offers a solution to increase both the sampling frequency of the residual, as well as to modulate the frequency content of the residual to higher frequencies. The proposed upsampling method is illustrated in Figures 7.7 and 7.8, and transforms a signal, $r[n]$, with a low sampling frequency to a signal $\hat{r}[n]$ with a higher sampling frequency.

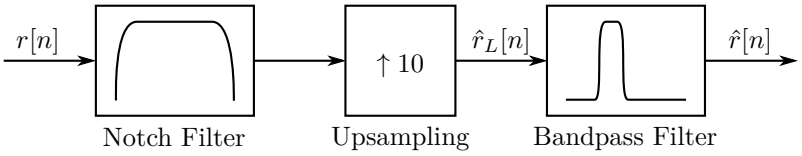

Figure 7.7: Upsampling Stage Block Diagram

The residual signal is upsampled by a factor of $L = 10$, which involves inserting zeros between each sample of the original signal [36] resulting in $\hat{r}_L$.

$$\hat{r}_L[n] = \begin{cases} L \cdot r[n/L], & n = 0, \pm L, \pm 2L, ... \\ 0, & \text{otherwise} \end{cases} \tag{7.3.1}$$

In the frequency domain, this upsampling results in a periodic repetition of the original frequency content of the signal. A bandpass filter with a frequency response of $H_{BP}(f)$ is then used to extract the desired frequencies. This allows single-sideband (SSB) modulation to be performed.

$$v_i[n] = \hat{r}[n] = \hat{r}_L[n] * h_{BP}[n] \tag{7.3.2}$$

where $h_{BP}[n]$ is the discrete-time impulse response of the bandpass filter.

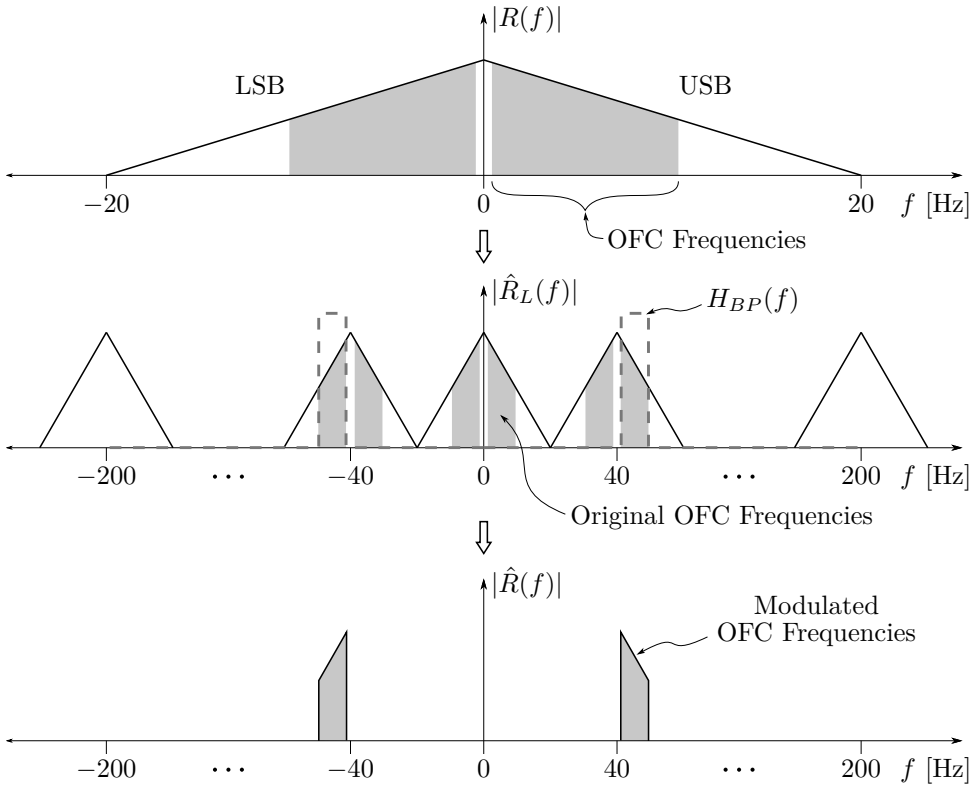

Figure 7.8: Frequency-domain Upsampling and Modulation Illustration

Thus, after upsampling the 1 Hz to 10 Hz frequency components of the original residual signal, the bandpass filter can extract the same information from the 41 Hz to 50 Hz band. A single PLL with a centre frequency of 45.5 Hz can then check for the presence of OFCs.

However, this approach comes with a number of concerns. Firstly, any DC content in the original residual becomes a 40 Hz frequency component that the PLL may attempt to track. Any significant DC offset could therefore result in a false detection. To prevent this, a notch filter is placed before the upsampling stage to remove any DC components. Figure 7.9 shows the output for a step input. The 40 Hz component is present for about 1 second before the notch filter reaches a steady state. This transient may still result in false alarms, but decreasing the settling time of the filter would require a larger notch width, potentially attenuating important frequency components and resulting in missed detections.

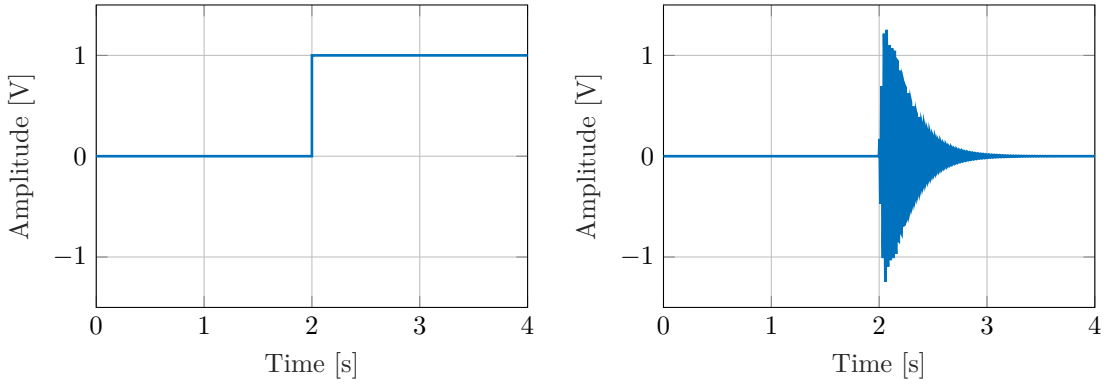

(a) Step Input (b) Upsampled Signal

Figure 7.9: Upsampling Results for Step Input

Secondly, the bandpass filter is not perfect, and cannot completely suppress all unwanted frequencies. This allows frequency content from the lower sideband to leak through, causing low-frequency oscillation in the amplitude of the upsampled residual. In general however, this only affects low-frequency OFCs, which have less strict detection time requirements. An example of this effect is shown in Figures 7.10 and 7.11, where a 2 Hz input results in two frequency components, at 38 Hz and 42 Hz.

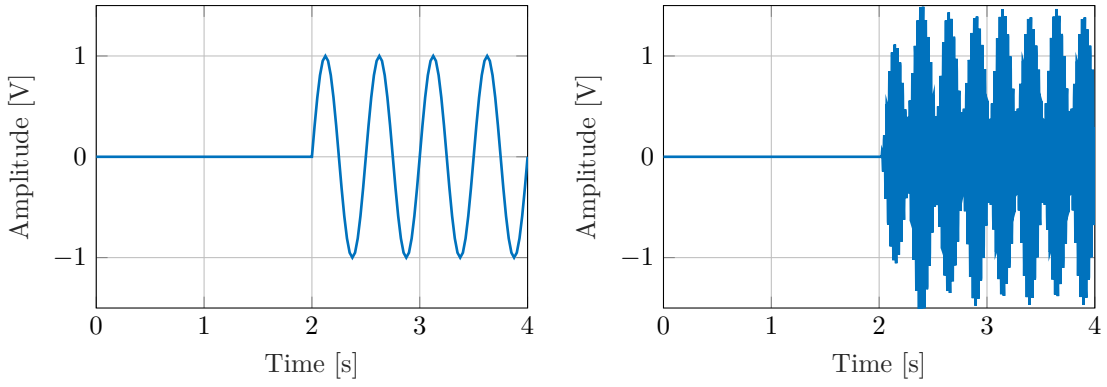

(a) 2 Hz Sinusoidal Input (b) Upsampled Signal

Figure 7.10: Upsampling Results for 2 Hz Input

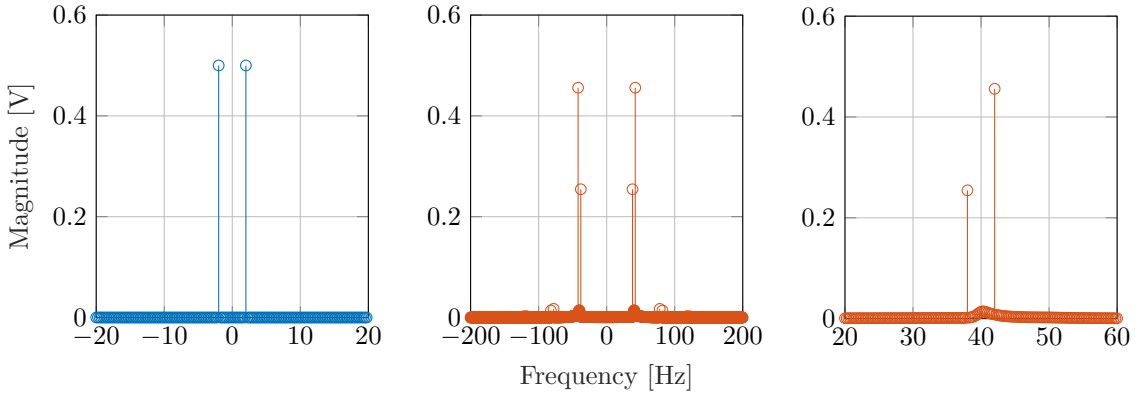

(a) Frequencies of Input (b) Frequencies of Output (c) 38 and 42 Hz Components

Figure 7.11: Frequency Content of Original and Upsampled Signals

Finally, the number of computations per second increases by a factor of 10, making the system much more computationally expensive.

Figures 7.12 and 7.13 show the successful application of the upsampling process to a 10 Hz signal, where a simple 45° phase shift can result in a 30% reduction in amplitude, due to the 10 Hz period being described with only four samples. (This was previously shown in Section 4.3.3 for oscillation counting.) For both cases, the upsampling and modulation manages to perform successful interpolation, and consistently provides a 50 Hz signal with the same amplitude. This shows that the upsampling stage is both viable and necessary for the successful implementation of the PLL detection method.

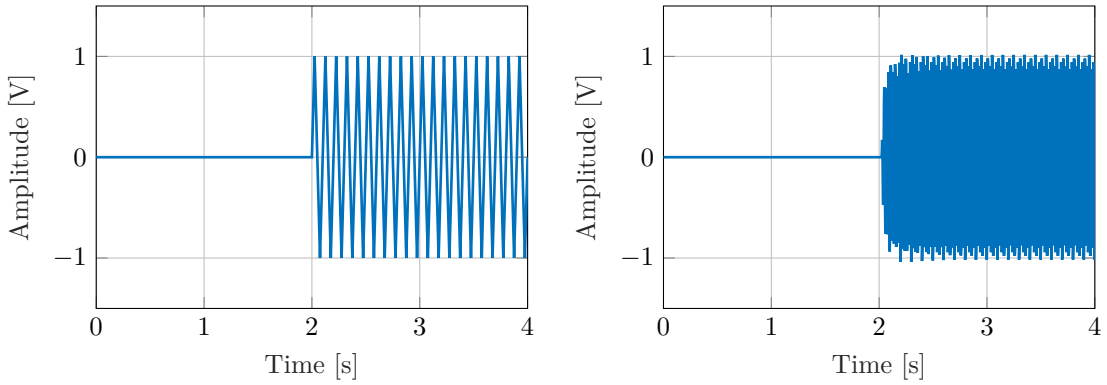

(a) 10 Hz Sinusoidal Input (b) Upsampled Signal

Figure 7.12: Upsampling Results for 10 Hz Input

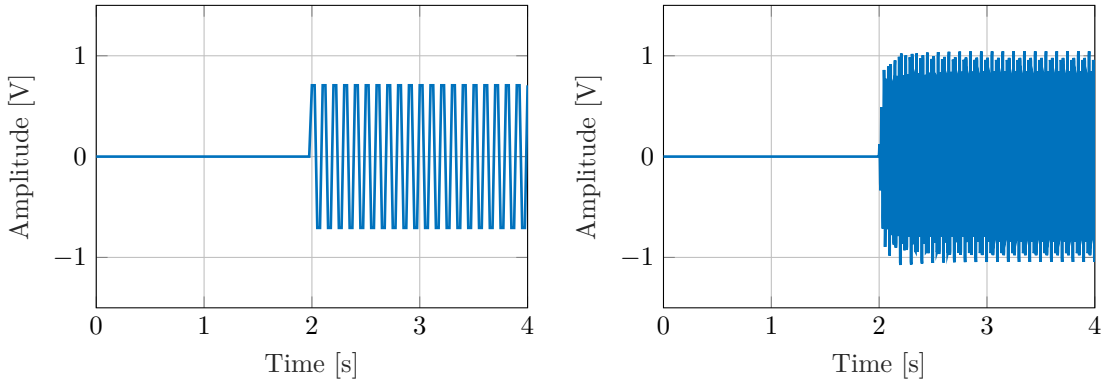

(a) 10 Hz Sinusoidal Input

(b) Upsampled Signal

Figure 7.13: Upsampling Results for 10 Hz Input with 45° Phase Shift

7.3.2 Loop filter

The main aspect of the design of the PLL is its loop filter. The loop filter of a PLL serves to attenuate the double-frequency term of the phase detector's output, and acts as a controller to ensure the following closed-loop specifications:

- 2% settling time of 0.15 s (half the 0.3 s detection time requirement),
- Zero steady-state error for frequency steps,
- PLL hold-in or lock-in range of at least 4.5 Hz.

The settling time requirement is to ensure that the PLL locks onto the incoming frequency quickly, so that the input frequency and output frequency converge quickly, resulting in the DC component of the QPD output as shown in Equation (7.2.16). From the same equation, it is clear that the output amplitude of the QPD will be at a maximum when the phase error is 0°, and the QPD output should ideally be consistent for all input frequencies. This therefore requires a type 2 closed-loop system. The effect of loop order and input frequency on the QPD output can be seen in Figure 7.14. Finally, the hold-in or lock-in range has to ensure that the PLL can lock onto any frequency between 41 and 50 Hz without losing lock or slipping cycles.

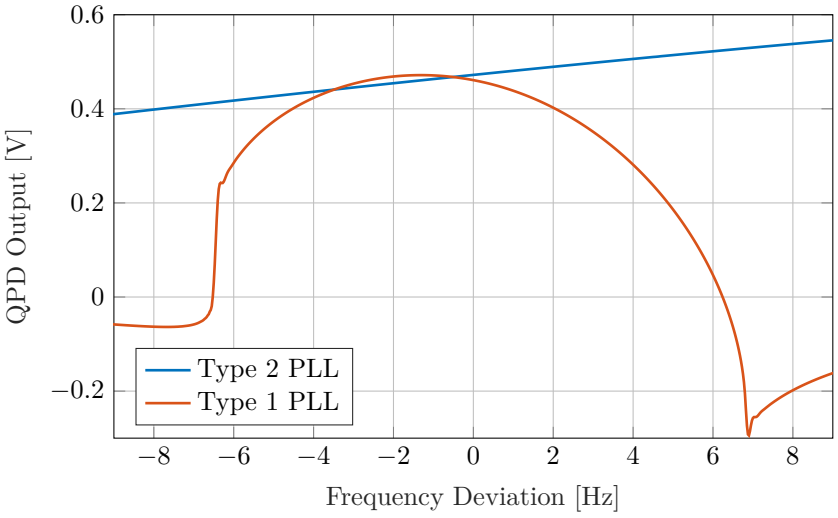

Figure 7.14: QPD Output versus Input Frequency for Type 1 and Type 2 PLL

Figure 7.15 illustrates the operation of a PLL with two different loop filter architectures by introducing a frequency sweep $\Delta\omega(t)$ and showing the resultant the VCO control signal $v_c[n]$. In both figures, it clear that $v_c[n]$ manages to track the frequency sweep $\Delta\omega(t)$ accurately. The type 1 PLL, with a lag controller of the form $K/(\tau s + 1)$, has a hold-in frequency of about 7 Hz, and manages to eliminate most of the double-frequency term. However, its QPD output is a function of $\Delta\omega$. In contrast, the type 2 PLL, which uses a PI controller, offers little attenuation of the high frequency content, but makes the QPD output independent of $\Delta\omega$, as seen in Figure 7.14.

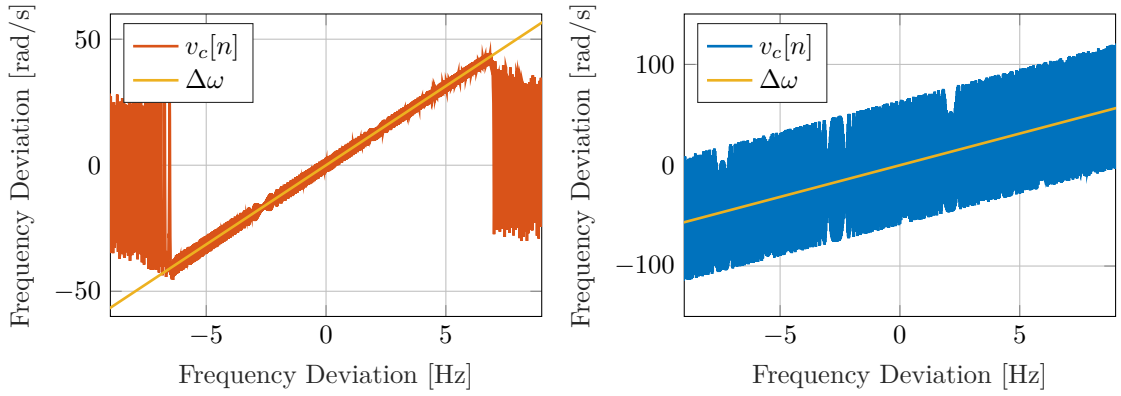

(a) Type 1 PLL

(b) Type 2 PLL

Figure 7.15: VCO Control Signal given a Frequency Sweep Input, with $ABc = 1$.

The listed requirements motivates the use of a proportional plus integral (PI) controller. Thus,

$$F(s) = K\left(\tau + \frac{1}{s}\right) \tag{7.3.3}$$

where K is the gain of the controller, and τ determines the location of the controller's zero. The step response of the system with the final PI controller design is shown in Figure 7.16, which shows the 2 % settling time of 0.15 s. Figure 7.17 shows the root locus design of the controller.

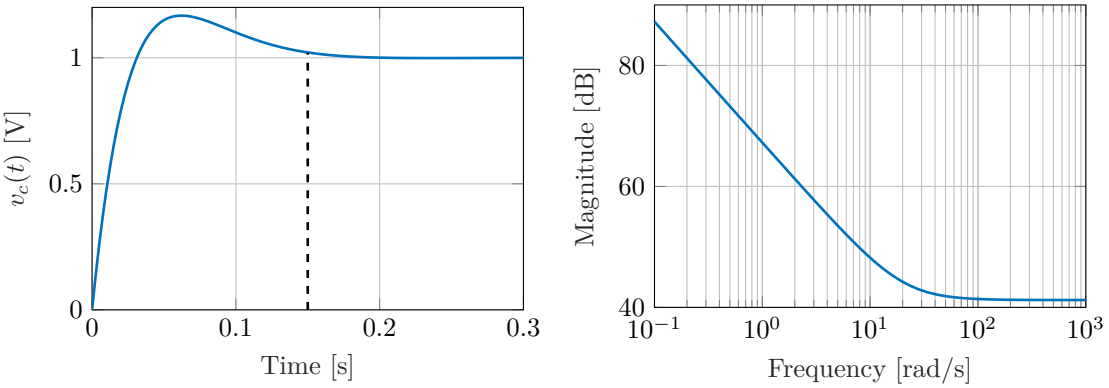

Figure 7.16: Closed-Loop Step Response (Left) and Frequency Response (Right) of PI Controller

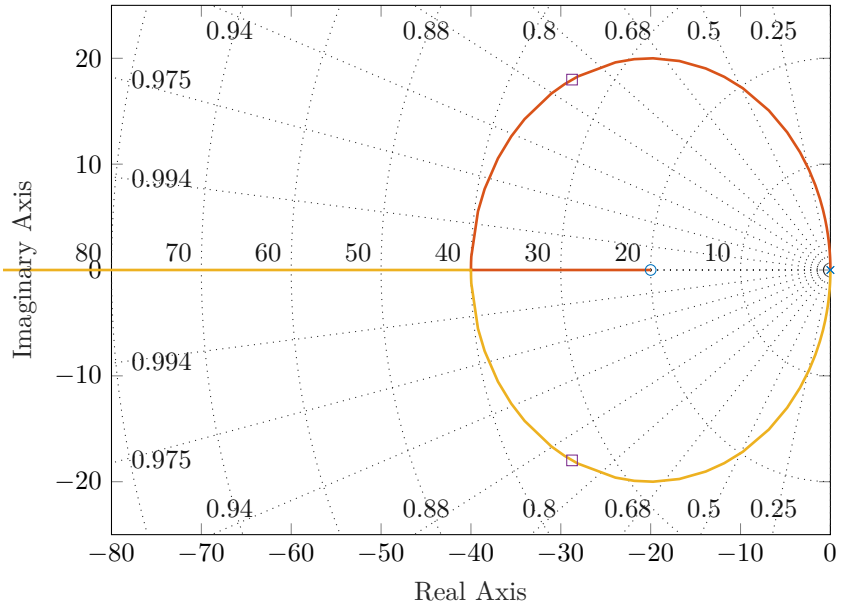

Figure 7.17: PLL Root Locus

The desired closed-loop poles (indicated as squares) are initially calculated from the settling time specifications, assuming an optimal damping ratio. The value τ is chosen to place the zero and to force the root locus to cross the desired closed-loop poles. Finally, the gain is chosen so that the closed-loop poles are placed at the desired locations on the root locus plot. These values are then further tuned and refined to ensure the desired settling time specifications, while also taking into account the potential for cycle slipping.

While the PI controller can easily be designed for the desired transient response, its values K and τ in Equation (7.3.3) affect the lock-in range of the PLL, as suggested in Equation (7.2.23). It is essential that cycle slipping be avoided as far as possible. To determine the lock-in range, a simplified Simulink model of the system, based on Figure 7.3, was iteratively run for different initial conditions for $\Delta\theta$ and $\Delta\omega$. Figure 7.18 shows the resulting lock-in ranges as defined by Equations (7.2.21) and (7.2.22).

Alternatively, a time-domain differential equation can be derived and solved using a solver such as ODE45. The differential equations for a type 2 PLL with a PI controller was derived by Groenendaal and Braun [47]. An simplified but intuitive understanding of how K and τ affect the lock-in range $\Delta\omega_L$ can also be gained from Equation (7.2.23) and Equation (7.3.3), where the lock-in range is (approximately) proportional to K and τ:

$$\Delta\omega_L \approx \pm\frac{1}{2}ABcF(\infty) = \pm\frac{1}{2}ABcK\tau \tag{7.3.4}$$

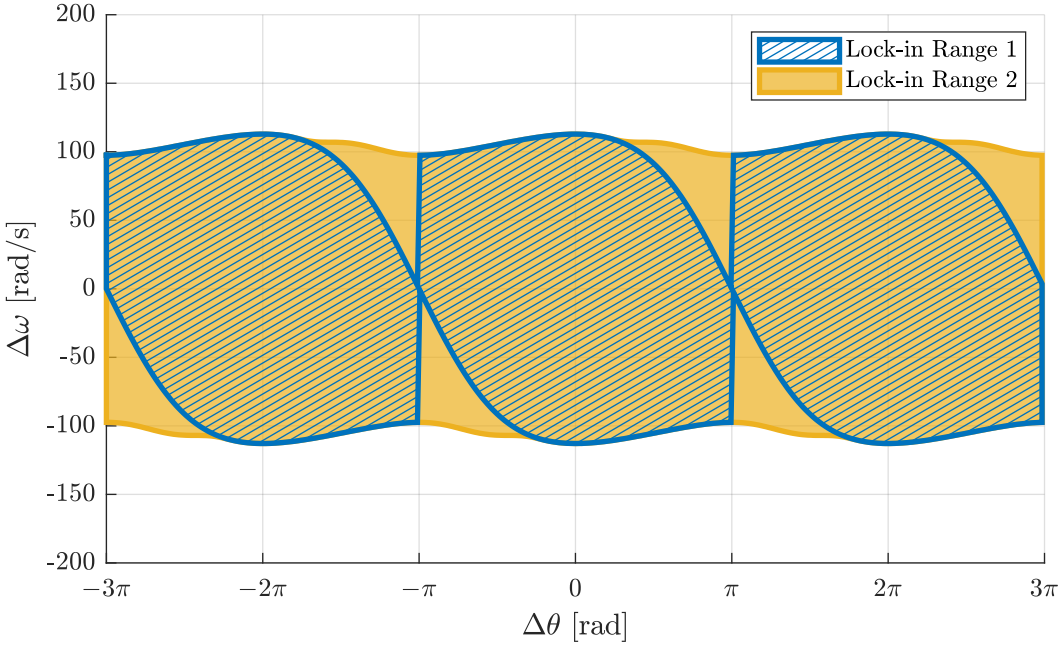

Figure 7.18: Lock-in Regions of PLL

Figure 7.18 shows the lock-in phase space regions, for a PLL with $K = 2300$, $\tau = 0.05$, and $ABc = 1$. The blue hatched area (Lock-in Range 1) shows the initial conditions of the input signal for which the PLL will not slip cycles, according to Equation (7.2.21). Note that this region is strongly affected by $\Delta\omega$, and lowers to 0 rad/s at the saddle points $\pm\pi$. Thus, for this definition, the lock-in range is effectively 0 rad/s, with only the centre frequency being within the range. The yellow shaded area (Lock-in Range 2) in contrast shows the lock-in region defined by Equation (7.2.22). Additionally, according to the approximation in Equation (7.3.4), the lock-in range $\Delta\omega_L$ is 115 rad/s, which agrees with Lock-in Range 2.

However, the above definitions do not take the initial conditions of the loop filter into account. Leonov, Kuznetsov, Yuldashev *et al.* [43] proposed the following process for determining the lock-in domain:

"*We have to increase the frequency deviation* $|\Delta\omega|$ *step by step and at each step, after the loop achieves a locked state, to change* $\Delta\omega = \tilde{\omega}$ *abruptly to* $\Delta\omega = -\tilde{\omega}$ *and to check if the loop can achieve a new locked state without cycle slipping. If so, then the considered value* $\Delta\omega$ *belongs to the lock-in domain.*"

Since the VCO frequency will be centred at 45.5 Hz, with a frequency range of interest of 9 Hz, The lock-in range should be at least 4.5 Hz, or 28.3 rad/s, as the input bandpass filter will remove all but these frequencies from the upsampled residual. This required range is shown in Figure 7.19, along with the final effective lock-in range. Since Lock-in Range 2 exceeds the required lock-in range by a reasonable margin, the PLL should avoid cycle slipping.

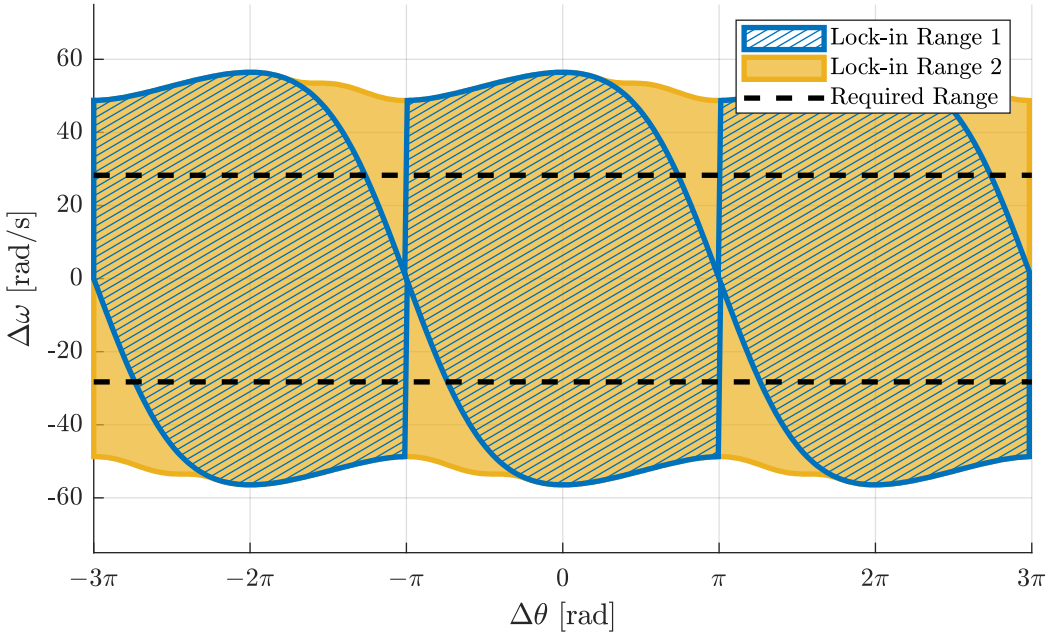

Figure 7.19: Lock-in Regions of PLL

In summary, the PI controller for the PLL was successfully designed for fast settling time, zero phase error for all input frequencies, and a lock-in range large enough to cover the OFC frequency band. The designed PI controller is implemented in software as a difference equation, converted from the continuous domain using the Bilinear Transform.

7.3.3 Lock Detector

The lock detector, or QPD, provides an indication of the presence of an OFC. It outputs a DC signal that provides an indication of the amplitude of a sinusoidal signal. Thresholding this signal can provide a simple means of checking for the presence of OFCs. When B, the amplitude of the VCO output is 1 V, the DC output of the QPD given by Equation (7.2.16) will be $A/2$, half the amplitude of the OFC. An overview of the lock detector is provided in Figure 7.20.

The operation of the QPD requires that the output of the VCO be phase shifted by 90°. This is implemented using differentiation. However, to maintain the amplitude of the signal, it is normalised with the VCO's centre frequency. Assuming the phase of $v_o(t)$ is zero, the phase shift operation is described as

$$\frac{1}{\omega_c}\frac{d}{dt}v_o(t) = \frac{1}{\omega_c}\frac{d}{dt}\cos(\omega_o t) = \frac{-\omega_o}{\omega_c}\sin(\omega_o t) \approx -\sin(\omega_o t). \tag{7.3.5}$$

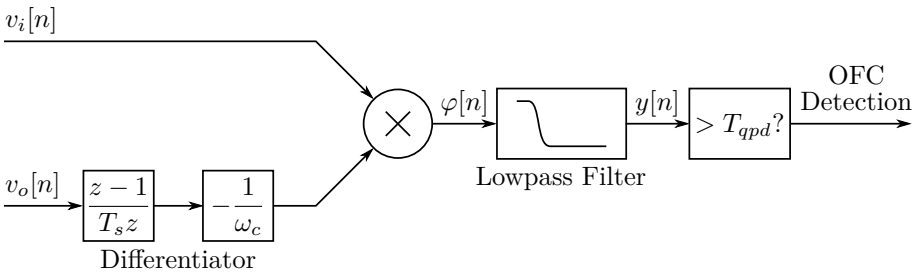

Figure 7.20: Quadrature Phase Detector

Finally, a lowpass filter ensures that the double frequency term in Equation (7.2.16) is removed. A 4^{th} order Butterworth filter is used, as its frequency response for DC components is 0 dB. The cutoff frequency is chosen such that the rise time for a frequency step is within the 0.3 s requirement.

To ensure the detection time of the system is within specification, the rise time of the filter output has to be less than 0.3 s for any input signal. The simplified Simulink model was used to calculate θ_e for initial conditions that require the most amount of time to converge to an equilibrium point, without cycle slipping. These worst cases were used to determine the longest theoretical rise time of the lowpass filter, and this time had to be less than 0.3 s. The rise time of the lowpass filter can be increased by increasing its bandwidth, but this results in a noisier output signal, and thus higher minimum detectable signal-to-noise ratio (SNR). Noise can be further reduced by increasing the order of the filter, but in the interest of minimising computations, a 4^{th} was determined to be adequate for reasonable SNR detectability.

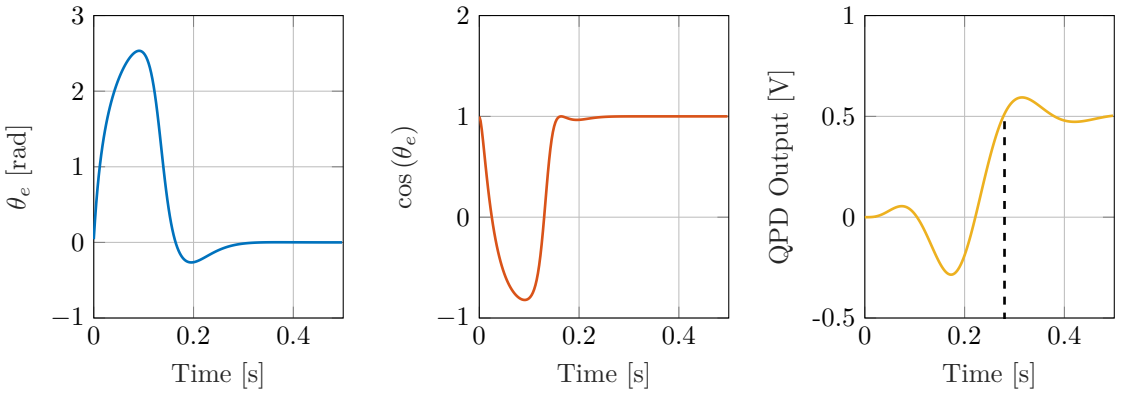

Figure 7.21: Step Response of QPD Output

Figure 7.21 illustrates an extreme example where the closed-loop system does not converge to an equilibrium point within the specified 0.15 s. However, the output of the QPD still manages to rise to 0.5 V in less than 0.3 s. This is achieved with a 4^{th} order Butterworth filter with a cutoff frequency of 5 Hz.

A second lowpass filter with a narrower bandwidth can be appended to the QPD. This could serve as an additional check that could detect even smaller amplitude OFCs. This detector would not offer fast detection times, but may offer a secondary check that could alert the monitoring system of the existence of an OFC, albeit a few seconds late. The structure of the aircraft would have to be robust against these smaller OFCs anyway, but this may offer additional information for diagnosis, or for the prevention of structural fatigue.

Figures 7.22 and 7.23 show the response of the PLL detection method for a 1 Hz and 10 Hz OFC respectively. The red horizontal lines denote the detection thresholds, T_{qpd}, while the black vertical lines indicate the OFC start times. The use of a 5 Hz cutoff frequency for the output filter guarantees that the detection time of the system falls within specification. Also shown is the output of a filter with a cutoff frequency of 0.5 Hz, which shows potential for using a much lower threshold to detect smaller OFCs, but with increased detection time.

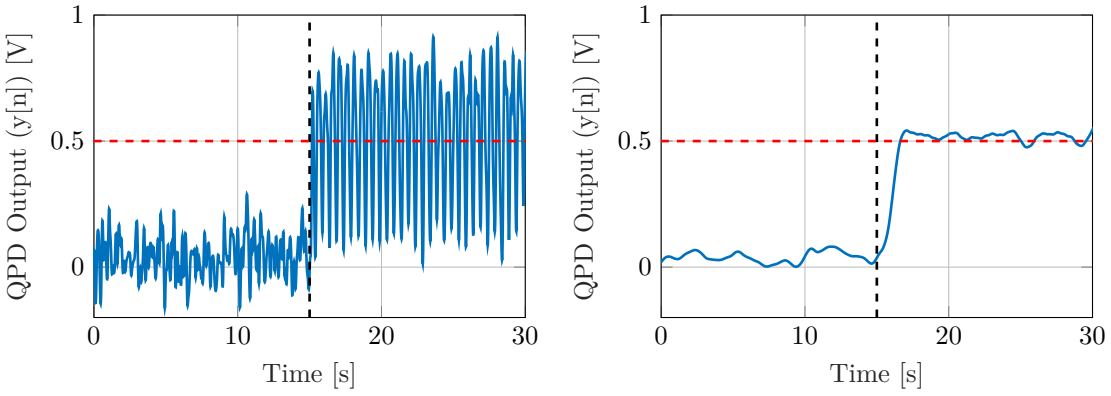

(a) 5 Hz Cutoff Frequency (b) 0.5 Hz Cutoff Frequency

Figure 7.22: Detection of 1 Hz OFC

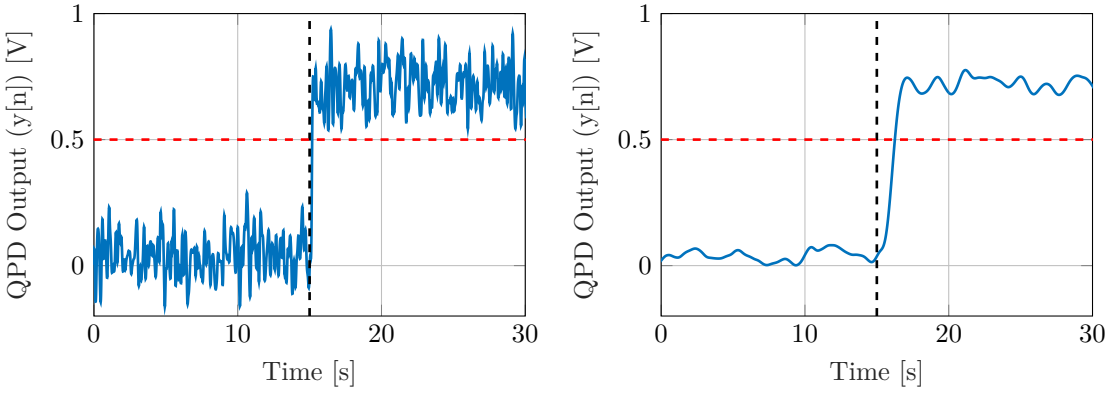

(a) 5 Hz Cutoff Frequency (b) 0.5 Hz Cutoff Frequency

Figure 7.23: Detection of 10 Hz OFC

7.4 Noise Analysis and Theoretical SNR Detection Threshold

The performance of the PLL detection method can be analysed by determining a theoretical minimum signal-to-noise ratio (SNR) for which an OFC can reliably be detected, without the risk of false alarms. The goal of this analysis is to determine the maximum noise power allowable at the QPD output that will not result in false alarms, to define a relationship between the noise power of the residual and the noise power of the QPD output, and subsequently to determine the minimum SNR.

SNR is defined as

$$\text{SNR} = 10 \log \frac{P_s}{P_\eta}, \tag{7.4.1}$$

where P_s is the power of the OFC signal, and P_η is the power of the noise. Assume the noise, $\eta[n]$, of the residual is band-limited white noise, and is wide-sense stationary. For zero-mean wide-sense stationary processes, it can be shown that the average power is equal to the variance of the process [48]. The average power of the random process is the integral of the power density spectrum (PDS) over frequency.

$$P_y = {\sigma_y}^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} S_y(\omega) d\omega \tag{7.4.2}$$

To simplify the analysis, assume the threshold at the QPD output is be 0.5 V. This is the threshold required for a 1 V OFC to be detected, as suggested by Equation (7.2.16). The permissible false

alarm rate is 1 false alarm over 100 000 h of flight time. Thus, for each sample of the residual, the probability of the QPD output crossing the 0.5 V threshold must be less than 1 in $1.44 \cdot 10^{10}$. This requirement implies that the variance of the noise at the QPD output must be less than $6.07 \cdot 10^{-3}$ W. This value can be used to determine the maximum variance at the PLL input $v_i[n]$, as illustrated below.

For discrete-time processes, the PDS of a signal is defined by the discrete-time Fourier transform (DTFT) of its autocorrelation sequence:

$$S_\eta(\omega) = \sum_{k=-\infty}^{\infty} R_\eta[k] e^{-jk\omega} \tag{7.4.3}$$

for $\omega \in [-\pi, \pi)$ rad/sample. $S_\eta(\omega)$ represents the PDS of white noise, which is a flat spectrum with a constant magnitude of $\mathcal{N}$ W/Hz. The average power of the white noise signal is thus

$$P_\eta = \frac{1}{2\pi} \int_{-\pi}^{\pi} S_\eta(\omega) d\omega \tag{7.4.4}$$

If no OFC is present, $r[n] = \eta[n]$, and

$$S_r(\omega) = S_\eta(\omega) = \frac{\mathcal{N}}{2}. \tag{7.4.5}$$

During the upsampling stage ($\eta[n] \rightarrow \hat{\eta}_L[n]$), the spectrum of the noise, band-limited to 20 Hz, becomes spread over 200 Hz. Additionally, as seen in Figure 7.7, a notch filter and bandpass filter are present at the PLL's input. The combined effect of these filters, with an effective frequency response of $|H_I(\omega)|^2$, results in the PLL's input:

$$v_i[n] = \hat{\eta}[n] = \hat{\eta}_L * h_I[n] \tag{7.4.6}$$

The PSD of $v_i[n]$ is

$$S_{\hat{\eta}} = S_\eta |H_I(\omega)|^2. \tag{7.4.7}$$

For lock detection, $v_i[n]$ is multiplied by $v_o[n]$, the VCO's output. This multiplication is expressed as

$$\varphi[n] = \hat{\eta}[n] \cos(\omega_o n + \Theta), \tag{7.4.8}$$

for which the autocorrelation function is [48]

$$R_\varphi[k] = \frac{R_{\hat{\eta}}}{2} \cos(\omega_o k). \tag{7.4.9}$$

$S_\varphi(\omega)$ is the DTFT of $R_\varphi[k]$:

$$S_\varphi(\omega) = \frac{1}{4}[S_{\hat{\eta}}(\omega + \omega_o) + S_{\hat{\eta}}(\omega - \omega_o)] \tag{7.4.10}$$

Finally, $\varphi[n]$ is passed through the QPD's lowpass filter with frequency response of $|H_O(\omega)|^2$. Thus, the PSD of the QPD's output is defined as

$$S_y(\omega) = S_\varphi(\omega) |H_O(\omega)|^2, \tag{7.4.11}$$

and the total average output power is

$$P_y[n] = \frac{1}{2\pi} \int_{-\pi}^{\pi} S_y(\omega) d\omega \tag{7.4.12}$$

The complete equation for the variance of $y[n]$ is

$${\sigma_y}^2 = \frac{\mathcal{N}}{16\pi} \int_{-\pi}^{\pi} |H_O(\omega)|^2 \left[|H_I(\omega + \omega_o)|^2 + |H_I(\omega - \omega_o)|^2 \right] d\omega. \tag{7.4.13}$$

Since the desired ${\sigma_y}^2$ is already known, $\mathcal{N}$ and ${\sigma_\eta}^2$ can be determined. Following this process with a desired output noise variance of $6.07 \cdot 10^{-3}$ W, an input noise variance of $61.1 \cdot 10^{-3}$ W is deemed permissible.

One final consideration is the total attenuation of an OFC through the system. The combined frequency responses of the notch and bandpass filters, as well as the effect of the approximation that $\omega_o/\omega_c \approx 1$ made in Equation (7.3.5), result in significant attenuation of low frequency signals, as demonstrated in Figure 7.24. Consequently, an OFC amplitude of 1.4 V is required to cross the 0.5 V lock detection threshold.

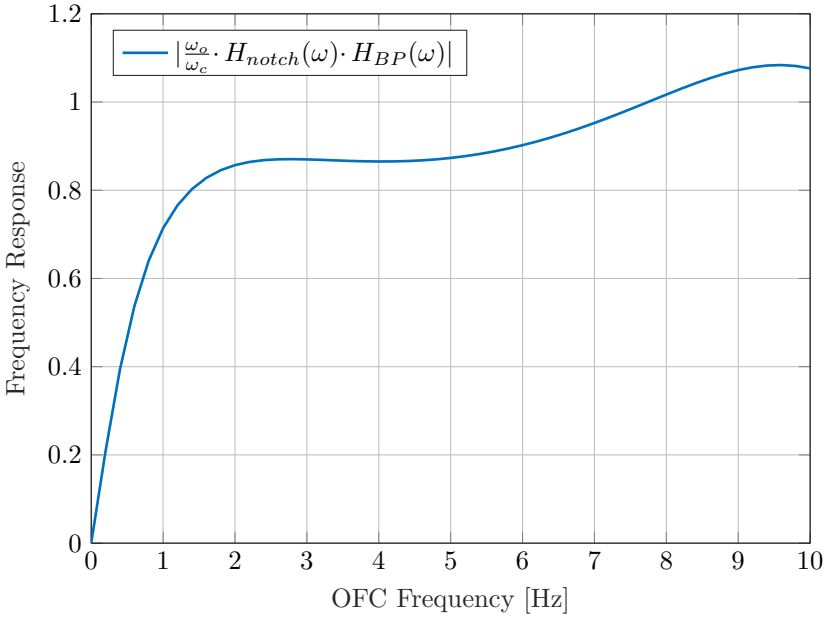

Figure 7.24: Combined Frequency Response of Notch and Bandpass Filters

Finally, given $A_{ofc} = 1.4$ V and ${\sigma_\eta}^2 = 61.1 \cdot 10^{-3}$ W, the theoretical minimum SNR that can be reliably detected is 12 dB.

Of course, this analysis does not explore the effects of modelling errors, the effects of ω_o constantly changing and reacting to the noisy input, or the effects of the double-frequency term in the PLL, but at least provides an indication of the theoretical performance of the PLL detection method. More rigorous threshold training and testing is performed in Chapter 8.

7.5 Threshold Training and Input Limiting

As mentioned before, the gain of the PLL control loop is dependent on A, the amplitude of the PLL input. Additionally, the amplitude of the OFC before the upsampling stage should be 1.4 V to cross a QPD threshold of 0.5 V, according to Section 7.4. However, the OFC is unlikely to exhibit a predictable or guaranteed amplitude. The approach here is therefore to scale the residual signal by a factor K_i so that the smallest expected OFC amplitude is scaled to 1.4 V. A saturation block is then used to limit this amplified residual. This limiting of the residual ensures that large OFCs do not influence the gain of the PLL control loop.

The scaling factor K_i is calculated using the permissible noise power ${\sigma_\eta}^2 = 61.1 \cdot 10^{-3}$ W. Given all the training data sets, the variance of each data set is measured, and the maximum measured variance is assumed to be the expected power of the fault-free residual, ${\sigma_{meas}}^2$. K_i is therefore used to scale the residual signal so that its measured variance matches the desired variance:

$$K_i = \frac{\sigma_\eta}{\sigma_{meas}} \tag{7.5.1}$$

This gain factor will result in an OFC having an amplitude of at least 1.4 V before the upsampling stage. If an OFC has a frequency of 1 Hz, the notch and bandpass filters will reduce its amplitude to 1.11 V. This 1.11 V amplitude ensures that the QPD output reaches at least 0.5 V, based on the scaling factor $\omega_o/\omega_c = 41/45.5$. However, the OFC amplitude would most likely be much greater than 1.4 V, changing the designed characteristics of the PLL. The saturation block therefore ensures that the signal that enters the PLL never exceeds 1.11 V. Figure 7.25 shows the potential change in amplitude of a 1 Hz OFC through the PLL system.

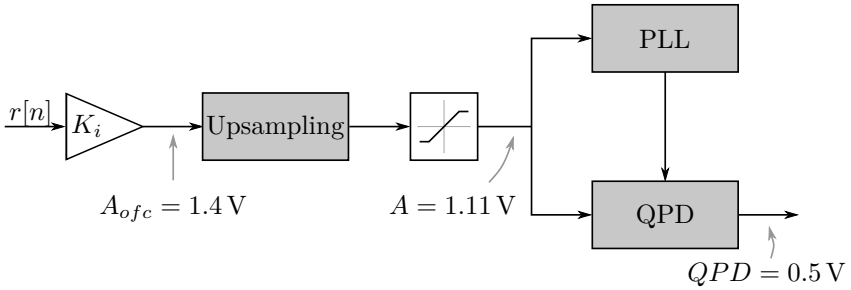

Figure 7.25: PLL Input Gain and Limiter

Its worth noting that if the induced OFC is extremely large, the signal at the output of the saturation would approximate a rectangular waveform. This would result in the PLL input having multiple harmonic frequencies [20]:

$$v_i(t) = 1.11 \times \frac{4}{\pi}\left(\sin(\omega_o t) - \frac{1}{3}\sin(3\omega_o t) + \frac{1}{5}\sin(5\omega_o t) + ...\right) \tag{7.5.2}$$

The PLL tends to lock onto ω_o, ignoring the higher-frequency harmonics. However, a gain factor of $4/\pi$ is introduced, potentially changing the PLL's behaviour. In general however, it was found that this scaling had a negligible effect on the detection performance of the PLL system.

In summary, the threshold of the PLL is trained by adjusting the input gain K_i, and a limiter prevents the PLL control loop gain from varying too much. After calculating K_i, the PLL is tested with the training data again to determine the maximum value that the QPD's output reaches when given fault-free residual. This is designed to be 0.5 V, but it is often found that the threshold can be lowered slightly, allowing for the detection of smaller OFCs.

Lastly, the effective threshold is defined as the smallest amplitude that can be detected within 3 cycles by the PLL. The average effective threshold is approximated as

$$T_{eff} = \frac{2T_{qpd}}{K_i}. \tag{7.5.3}$$

This is the effective threshold that will be assumed in the results of Chapter 8. However, if the effects of the input filters are taken into account, the smallest amplitude that is theoretically guaranteed to be detected is

$$T_{eff} = 1.4\frac{2T_{qpd}}{K_i}. \tag{7.5.4}$$

7.6 Design Considerations

Summarised in this section are additional designs and design considerations investigated throughout the design process. This includes a PLL detection method that avoids the modulation of the residual to higher frequencies, and the use of a higher order loop filter for better noise and high-frequency attenuation. These designs were deemed noteworthy, but ultimately did not provide the same level of performance as the design in Section 7.3.

7.6.1 PLL without Upsampling and Modulation

During the design of the PLL in Section 7.3, the possibility of a solution that did not require the upsampling and modulation stage was investigated. A successful variation was designed, which made use of two PLLs operating in parallel with different frequency bands. This design, however, still required a sample rate of at least 100 Hz. Upsampling by a factor of 3 (without modulation) would be necessary for residuals sampled at 40 Hz.

This implementation made use of two parallel PLLs working in the 1 – 3.2 Hz and 3.2 – 10 Hz bands, and had advantages over the final design. The use of a lower upsampling factor offered a less computationally expensive system overall. The effects of imperfect SSB modulation and DC components were also eliminated.

The phase shifting method used in the QPD had to be revised. Normalising the derivative with the centre frequency was no longer viable, as the assumption that $\omega_o/\omega_c \approx 1$ made in Equation (7.3.5) was no longer valid when working at low frequencies. Instead, the estimate of the frequency deviation provided by the PLL was used for the scaling, so that $\omega_o/(\omega_c + cv_c) \approx 1$. Additionally, a check was inserted to ensure that the scaling value $\omega_c + cv_c$ never dropped below 1 rad/s, as this would introduce large gains that could result in false alarms. This check is represented as a saturation block in the overview of the system in Figure 7.26.

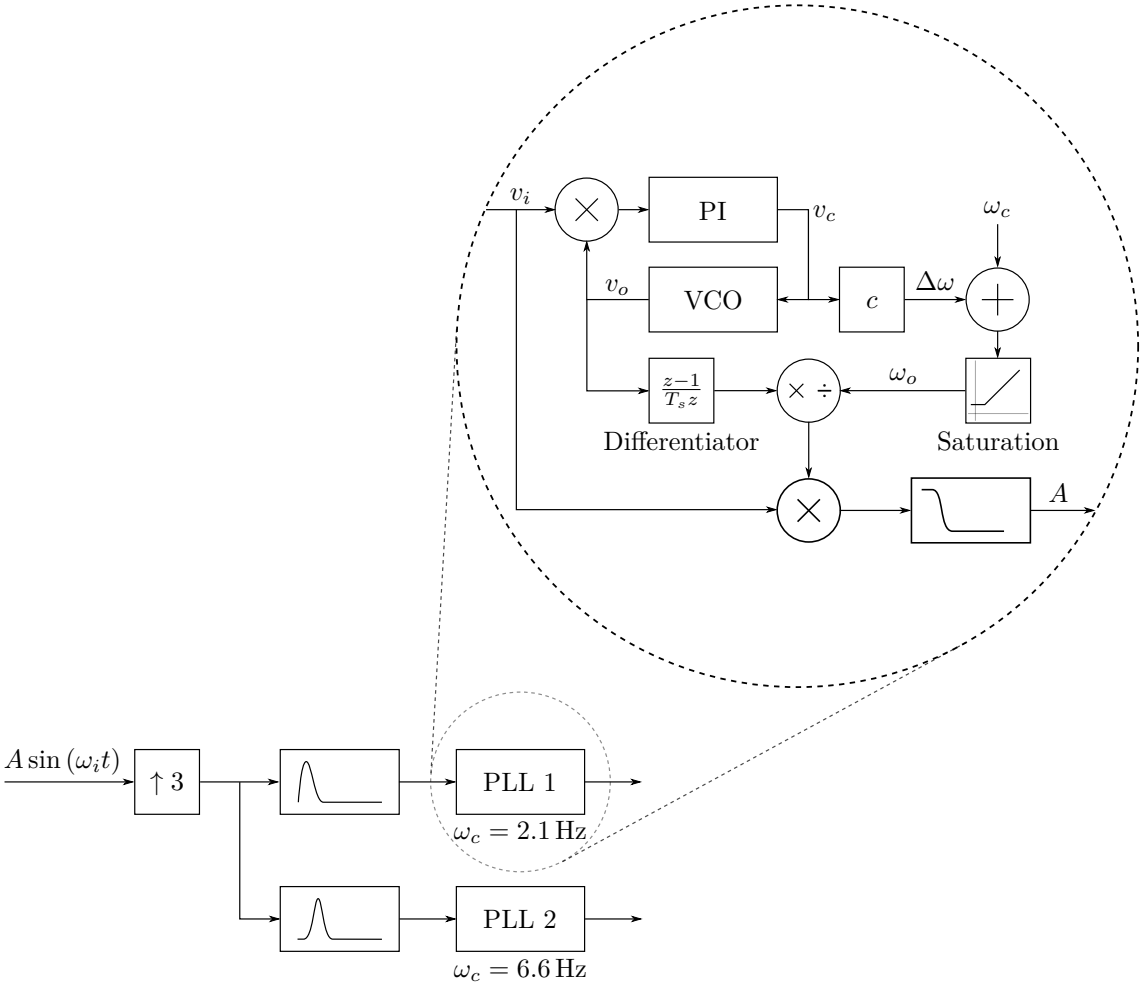

Figure 7.26: Low-frequency PLL Concept

Overall, this design offers a viable implementation of the PLL detection method with slightly lower computational requirements, but with slightly worse performance in terms of detection time

and smallest detectable amplitude. Due to the low-frequency nature of this design, the trade-off between settling time and double frequency attenuation must be considered very carefully.

7.6.2 Second Order Loop Filter

The use of a second order loop filter was introduced to improve the lowpass characteristics of the filter by augmenting the PI controller with a lag filter of the form $K/(\tau s + 1)$. The PI controller ensured accurate frequency tracking, while the lag filter provided better attenuation of the double frequency. This architecture was proposed as a way to avoid any potentially unpredictable or unwanted effects of the double frequency term, and is often recommended in literature [40]. Alternatively, Texas Instruments demonstrates the use of an adaptive notch filter to similar effect [49].

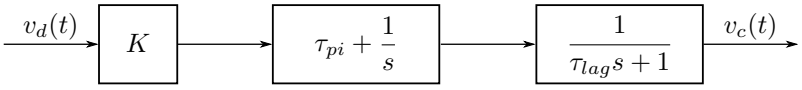

Figure 7.27: Second Order Loop Filter

Provided the PLL's centre frequency is much higher than the PI controller's bandwidth, the pole provided by the lag filter can be placed at a high frequency. This allows the system to approximate the operation of a PI controller, while also attenuating the double-frequency term.

In general, it was discovered that the addition of the lag filter did not improve performance or convergence time, and in general, it increased detection time slightly. The addition of the pole also reduced the lock-in frequency. Finally, it was discovered that the low-frequency ringing effects caused by imperfect SSB modulation had a more significant effect on the PLLs operation than the double-frequency term, and thus the use of this filter could not be justified.

7.7 Summary and Contributions

This chapter investigated the use of telecommunications techniques that are not typically applied to fault detection. A phase-locked loop was designed and its potential application to the oscillatory failure case detection problem was investigated. The PLL attempted to lock onto any sinusoidal signals present in the residual signal, and a quadrature phase detector was used to determine the presence or absence of an OFC.

The following contributions were made in this chapter:

- A detailed literature study on the PLL was performed. This included its operation, design methodologies, and non-linear effects.
- A PLL and QPD system was designed to detect OFCs. It used a PI controller to ensure zero steady state error when tracking frequency changes, and had a settling time of 0.15 seconds. Phase plane analysis was used to confirm that the designed controller would avoid cycle slipping. The bandwidth of the output filter was tuned to ensure 0.3 second detection time for all failure cases.
- An upsampling and modulation stage was developed to modulate the residual to a higher frequency band. This ensured that the initial low sampling rate did not affect the amplitude information of the residual, and the double-frequency term of the phase detector's output would be better attenuated by the loop filter.
- An analysis of the signal-to-noise ratio was performed to determine the theoretical detection performance of the system. This analysis was based on the power density spectrum of the residual, and the effects of the filter applied to the residual signal. The noise analysis was further used to determine an adequate gain value at the PLL's input and to choose an appropriate threshold.

- Two additional designs were implemented. The first was an approach that did not require the modulation stage. It did, however, still require a higher sampling frequency. The second design was the use of a second order loop filter that would provide better attenuation of the double frequency. However, its potential improvements were almost negligible.

Overall, the PLL can detect OFCs with a theoretical signal-to-noise ratio of 12 dB within three cycles. Furthermore, with PLL can be modified to detect much smaller OFCs with longer detection time. However, due to its high sampling rate requirement, this is a computationally expensive approach. The PLL detection method is more thoroughly tested and compared to the other approaches in Chapter 8.

Chapter 8

Results

8.1 Introduction

To recapitulate, the aim of the project is to investigate and devise methods for the detection of oscillatory failures, where each detection method should detect small oscillations within 3 cycles. The preceding chapters have introduced and described five different approaches to OFC detection, and it is now necessary to compare these detection methods based on their ability to detect small amplitudes, and to minimise detection time. This is achieved by running each detection method through a large number of test cases.

The ultimate goal of these tests is to determine the smallest amplitude OFC that each detection method can reliably detect within 3 cycles, and this therefore serves as the primary performance metric. Additional performance metrics include smallest amplitude detectable (within 6 cycles), average detection time, and computational complexity.

8.2 Simulation Setup

The following methods are implemented and tested:

1. Oscillation counting (OC),
2. Integrated absolute error (IAE),
3. Discrete Fourier transform (DFT),
4. Multi-window Fourier transform (MWFT), and
5. Phase-locked loop (PLL).

Each data set is sampled at 40 Hz, with the simulation time set to 30 seconds. The seed of each random number generator source changes for each simulation run. Additionally, the actuator parameters ΔP and K_d are randomised, within some defined bounds, to increase the effects of modelling errors. Both the oscillation counting and the IAE methods upsample the residual signal by a factor of 3, and the DFT and MWFT methods use a zero-padding factor of 5. However, all the detection methods are provided with the same residual signal sampled at 40 Hz.

8.2.1 Training the Detection Methods

1000 training data sets are used to determine the thresholds for each detection method. Half of these sets make use of the chirp load factor command, to stress the aircraft dynamics across all frequencies. The remaining 500 training data sets use the randomised load factor command signal. The detection thresholds for each technique are "trained" on fault-free data to determine the lowest detection thresholds that do not produce false alarms. This provides each method with the best possible opportunity to detect small OFCs.

8.2.2 Testing Approach

Each OFC detection method is tested using a structured approach to find the smallest amplitude OFC that is consistently detected within 3 cycles. The detection methods are tested with a range of OFCs that sweep over an amplitude and a frequency range. Thus, for each discrete frequency over the frequency sweep, a range of amplitudes are tested to find the smallest amplitude that is detectable at that specific frequency. This offers a way to determine the effect of both amplitude and frequency of an OFC for each detection method. The final results are therefore expressed as a function of frequency.

Furthermore, each discrete amplitude-frequency case is tested 10 times, with randomly varied model parameters. Additionally, for each of these 10 tests, the initial phase of the OFC is changed, to test the potential effects of sampling rate, as illustrated earlier in Figure 4.13. Thus, each test data set contains a unique amplitude-frequency-phase combination. An OFC with a specific amplitude and frequency is considered to be reliably detectable if all 10 cases are detected within 3 cycles.

The final simulation generated a total of 60 060 test data sets. These data sets were generated three times for three different cases, which were discussed in Sections 3.4.3 to 3.4.5:

1. Liquid failure at the actuator position sensor,
2. Liquid failure at the command current, and
3. Solid failure at the command current.

The desired metric for OFC detection performance is the smallest amplitude control surface oscillation detectable within 3 cycles. However, the OFCs are induced within the actuator at the rod position sensor, where the amplitude is expressed in mm, or at the servo current in mA. Mathematically translating these OFC amplitudes to control surface deflection in degrees is non-trivial. Therefore, the amplitude of the oscillation in the true control surface deflection is estimated in post-processing using Matlab's FFT function, performed over the entire time span of the OFC. This is shown in Figure 8.1.

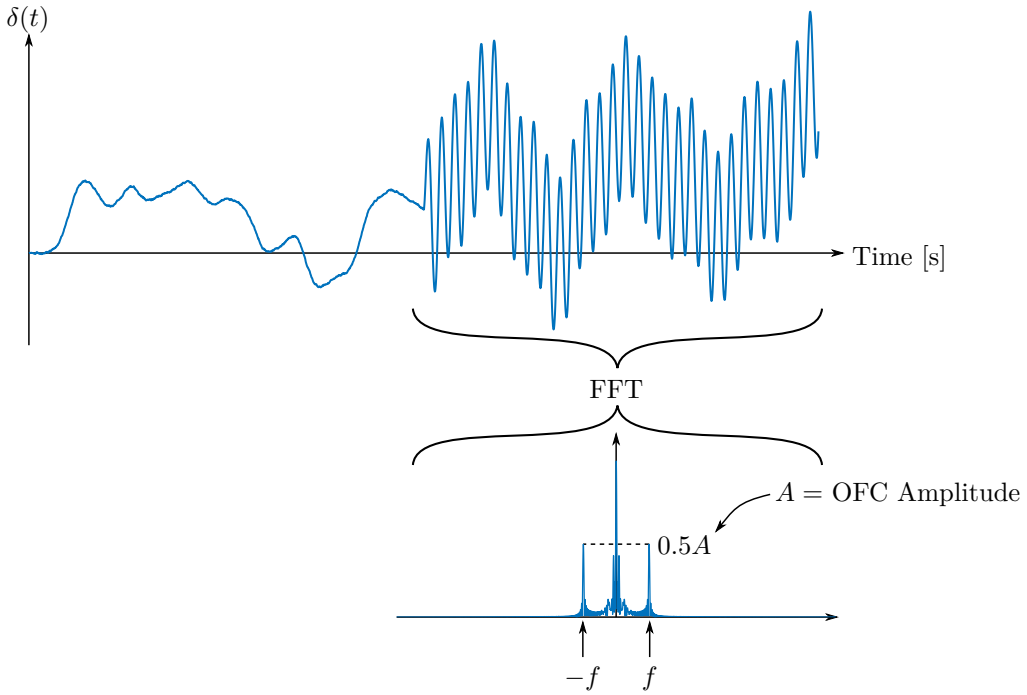

Figure 8.1: Obtaining the True OFC Amplitude at the Control Surface

Figure 8.2 shows the resulting OFC amplitude at the control surface over frequency for different OFC amplitudes A_p at the rod position sensor. At 10 Hz, it can be seen that the resulting control surface oscillation is very small, with amplitude less than 0.8°. At these frequencies, the actuators have reached their physical limits in terms of rod speed, limiting the amplitude of the control surface oscillation.

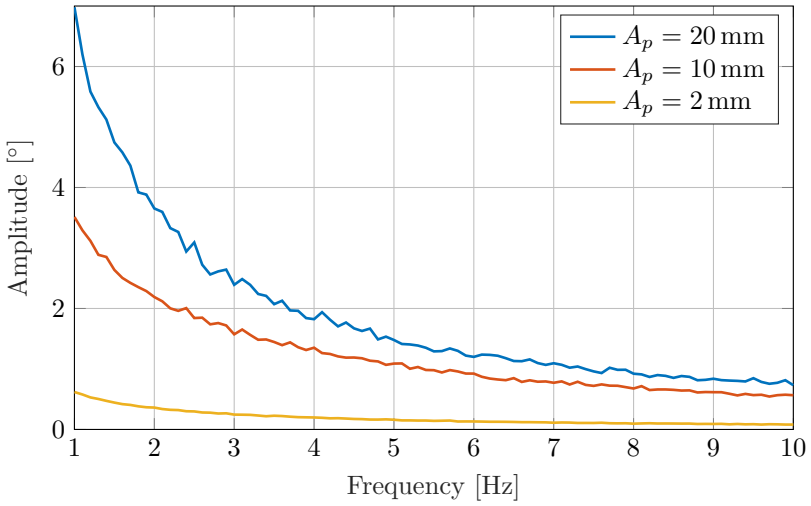

Figure 8.2: Control Surface Oscillation Amplitude for Different Rod Sensor OFC Amplitudes

8.3 Thresholding Results

The first set of notable results is the thresholds that were trained using the training data sets. Figure 8.3 shows the thresholds for each detection method as a function of frequency, and represents the smallest amplitude OFCs that can theoretically be detected, but not necessarily within 3 cycles. Note that here the thresholds of the FFT and MWFT are doubled to show the true amplitudes, since the frequency component of an OFC with an amplitude of A manifests in the frequency domain with a magnitude of $A/2$.

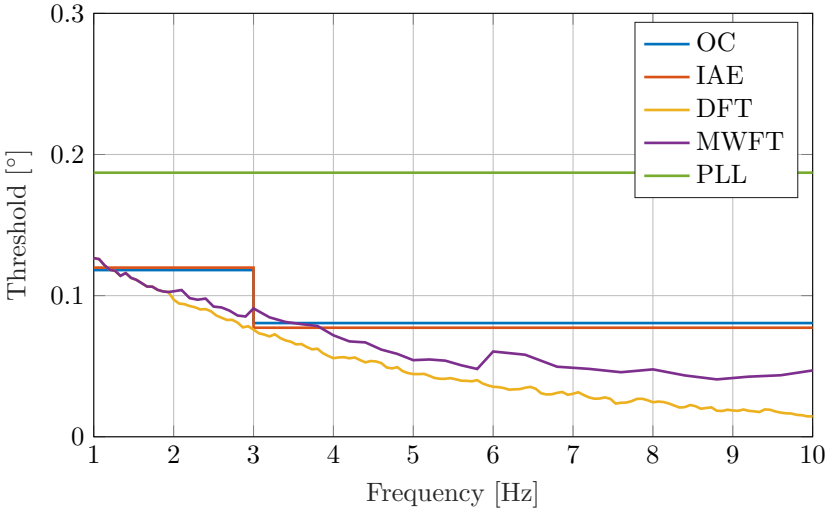

Figure 8.3: Smallest Amplitudes Theoretically Detectable over Frequency

These thresholds are essentially the frequency-dependent thresholds for each method. The oscillation counting and IAE methods divide the residual into two frequency sub-bands, where each

sub-band is evaluated with a different threshold. This is illustrated in Figure 8.3 with the step at 3 Hz. The thresholds for oscillation counting and the IAE are extremely similar, suggesting that both methods should offer similar performance. These two thresholds are generally the largest of the different thresholds, with the exception of the PLL. The PLL has by far the largest threshold, and uses only one threshold value for all frequencies. The DFT method consistently has the smallest threshold across all frequencies, and should therefore be able to detect the smallest OFCs. The threshold of the MWFT is slightly higher than the DFT, due to the spectral leakage and smearing that occurs due to the shorter windows. The DFT and MWFT have the advantage of having a unique threshold for each frequency component. Essentially, these frequency-domain techniques have high-resolution frequency-domain thresholds, while the others have low-resolution thresholds.

However, these thresholds do not guarantee 3 cycle detection time, but can be used to infer the amplitudes that are (theoretically) guaranteed to be detected within 3 cycles. Figure 8.4 shows these "effective thresholds", which are the smallest amplitude OFCs that can theoretically be detected within 3 cycles.

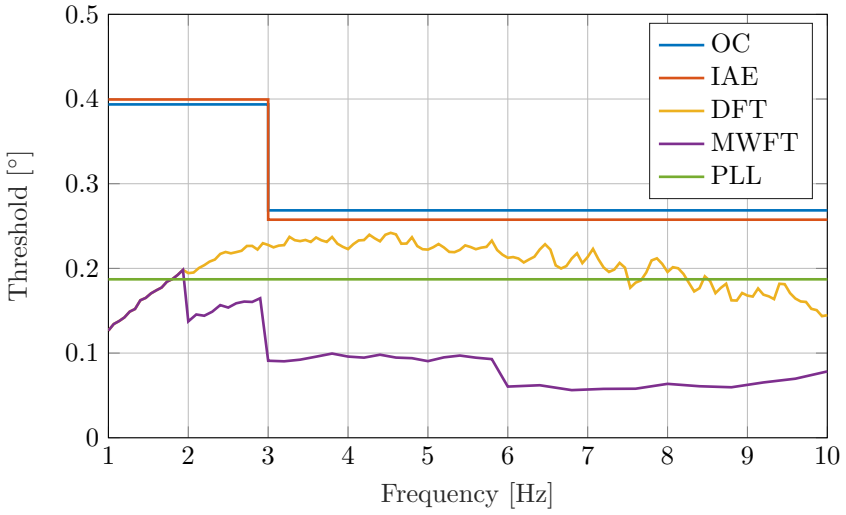

Figure 8.4: Smallest Amplitudes Theoretically Detectable within 3 Cycles

Here, the oscillation counting and IAE methods have the largest thresholds. These thresholds are adjusted based on the 3.3 factor rule-of-thumb explained in Section 4.3.2. The PLL shows no threshold change, due to the fact that it is designed to ensure 0.3 s detection time. Its threshold is comparable to the DFT threshold, which shows a significant increase over the original threshold. This increase is due to the fact that the DFT's detection time, from Equation (6.4.15), is a function of the OFC's amplitude and frequency, as well as window size. This window size dependency provides the MWFT with a significant advantage over all the other methods, where it uses shorter windows for higher frequencies. The MWFT thus features the smallest effective threshold.

In summary, and in theory, the DFT should be able to detect the smallest OFCs in general, but the MWFT should be able to detect the smallest OFCs within 3 cycles. The thresholds and effective thresholds represent the theoretical performance of each OFC detection methods. In the next sections, the theoretical performance will be verified by testing all five OFC detection techniques on the testing data provided by the simulation framework.

8.4 Test Case: Liquid failure at the Rod Position Sensor

8.4.1 Smallest OFCs Detected within 3 Cycles

The first case that is investigated is a liquid failure at the position sensor. Figure 8.5 shows the smallest amplitudes versus frequency that each detection method detected within 3 cycles. Note

that at if a plot is broken at some frequency, no OFCs at that frequency could be reliably detected within 3 cycles for that detection method.

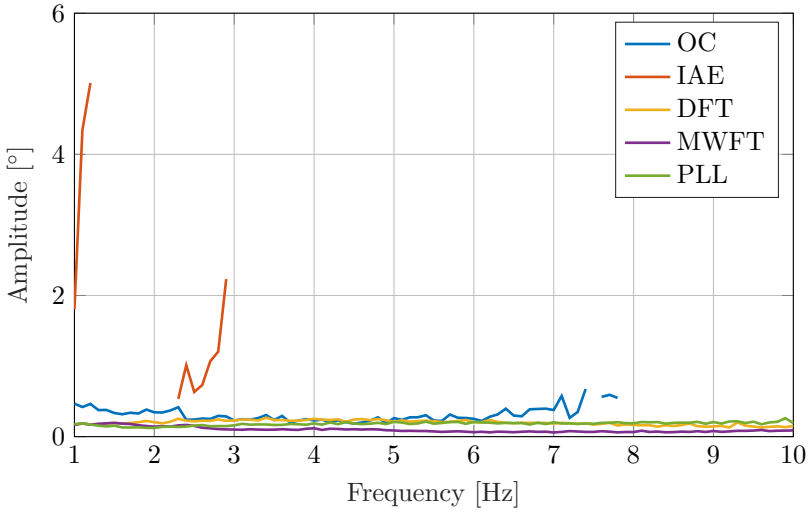

Figure 8.5: Smallest Amplitudes Detected within 3 Cycles (Liquid Failure at Rod Sensor)

It is immediately seen that the IAE shows the worst performance, with timely detections occurring in a very narrow range of frequencies, and requiring enormous amplitudes at these frequencies. This is due to the effects of filtering, combined with the fact that thresholds are counted at the end of a cycle. The IAE does manage to detect OFCs between 1 and 1.2 Hz, and again between 2.3 and 3 Hz. Over the first interval, the filters induce a phase lead effect, allowing the OFC to be detected earlier, while over the second interval, the smaller window used by the higher-frequency sub-band begins to detect the OFCs earlier than the larger window. In conclusion, the IAE is the worst performing method.

To get a closer look at the other results, the IAE result is removed from Figure 8.5, and the remaining results are shown in Figure 8.6.

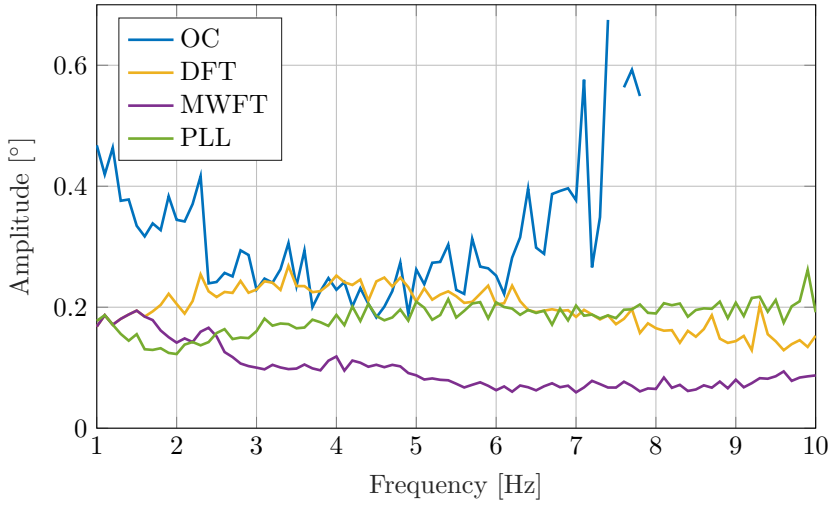

Figure 8.6: Smallest Amplitudes Detected within 3 Cycles (Liquid Failure at Rod Sensor, IAE Excluded)

Here, it can be seen that, in general, the oscillation counting approach requires the largest amplitudes to guarantee detection time. Moreover, at frequencies higher than 7.5 Hz, oscillation counting struggles to detect any OFCs in time, and can therefore not guarantee detection time for these higher frequency failures. Once again, this is likely due to the effects of filtering. This could also be caused by the limited control surface oscillation amplitude at higher frequencies as illustrated in Figure 8.2, but these amplitudes are still greater than the effective thresholds, making this theory unlikely.

The DFT, MWFT, and PLL all show reliable detection across all frequencies. The DFT shows in general higher amplitudes than the MWFT and PLL, as suggested by the effective thresholds in Figure 8.4. The PLL shows a fairly consistent amplitude of 0.2° across all frequencies. Finally, the MWFT shows considerably better detection amplitudes than the other methods. This is due to the variable window size that decreases for higher frequencies.

These results can now be compared to the theoretical effective thresholds to confirm that the experimental results match the theoretical predictions. Figure 8.7 plots the experimental results for each method together with their theoretically predicted effective thresholds, which are represented with the red dashed lines. Due to the IAE's poor performance, it is not shown.

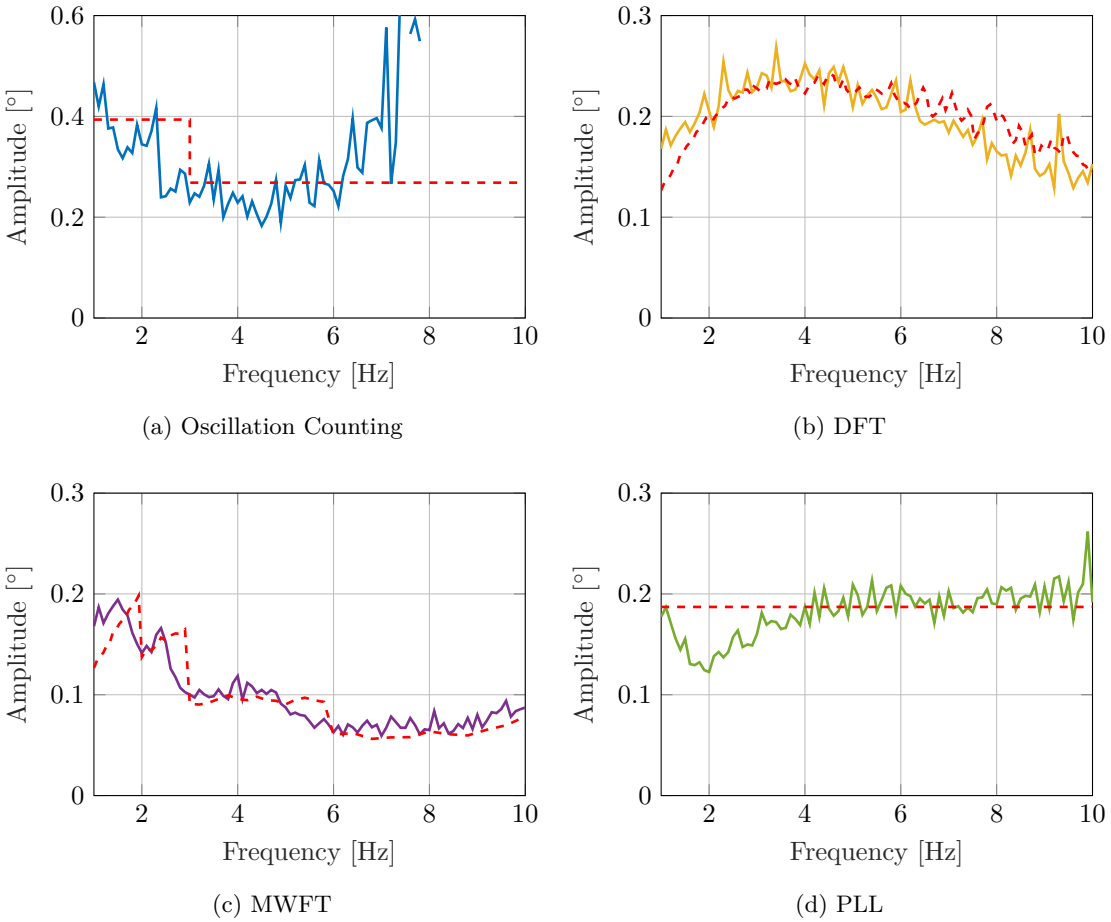

(a) Oscillation Counting

(b) DFT

(c) MWFT

(d) PLL

Figure 8.7: Minimum Amplitudes Detected within 3 Cycles and Corresponding Effective Thresholds (Liquid Failure at Rod Sensor)

In general, the experimental results agree with the effective thresholds predicted by the theory. The oscillation counting result shows that the rule-of-thumb established for determining the effective threshold does provide a reasonably good indication of the expected performance. Between 1 and 2.2 Hz, the experimental data seems to agree with the theoretical data. At frequencies higher than 2.2 Hz, the data follows the threshold for the 3 to 10 Hz sub-band, suggesting that OFCs between

2.2 and 3 Hz are being detected by the high frequency window sooner than with the low frequency window. This is likely due to the phase lead effect that low frequencies experience at the 3 to 10 Hz sub-band filter, and these frequencies are not being completely attenuated. Unfortunately, the detection method begins to diverge from the theory at frequencies above 6 Hz.

The DFT and MWFT experimental results tend to fit with the theoretical effective thresholds rather well, confirming the theory proposed in Chapter 6.

Finally, the PLL is expected to detect all OFC frequencies equally well, since a single threshold value is applied to all frequencies, and a detection time of 0.3 s is guaranteed. Indeed, the experimental data seems to match the theoretical amplitude for frequencies between 4 and 10 Hz. At lower frequencies, the PLL shows improved performance. This occurs despite the attenuation and ω_o/ω_c scaling explored in Sections 7.4 and 7.5. A possible reason for this discrepancy is that at these low frequencies, the upsampling and modulating stage at the PLL input causes low-frequency oscillation that allows lower amplitudes to cross the PLL's threshold.

In summary, the experimental performance of all the methods seem to confirm the theoretical performance provided by the effective thresholds, with the exception of the IAE. The MWFT provides the best amplitude and detection time performance, followed by the PLL and DFT. Oscillation counting cannot guarantee detection time for frequencies above 7.5 Hz, due to filtering effects, and possibly because the dynamics of the system cannot allow for a high frequency oscillation large enough for timely detection. The IAE performs extremely poorly in terms of detection time.

The box-and-whisker graph in Figure 8.8 visualises the performance of the detection methods in terms of smallest amplitude OFCs detectable within 3 cycles, based on the experimental results. The figure shows the distribution of the smallest amplitudes detectable within 3 cycles across the frequency range. In general, the maximum of these plots would provide an indication of the largest OFC that an aircraft's structural support should accommodate. The MWFT has the lowest maximum, and could therefore allow for the greatest weight-savings in the aircraft structural design. Note that the red mark represents the median of the data, and the bottom and top edges of the box indicate the 25^{th} and 75^{th} percentiles, respectively.

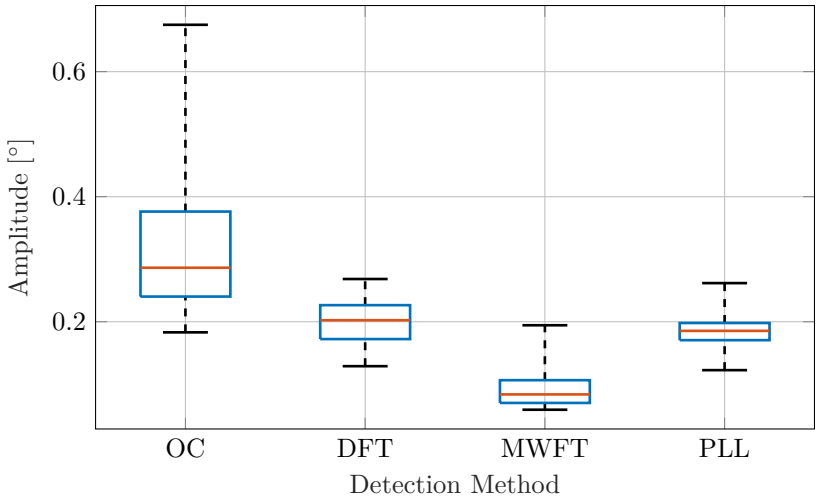

Figure 8.8: Range of Amplitudes Detectable within 3 Cycles (Liquid Failure at Rod Sensor)

8.4.2 Smallest OFCs Detected within 6 Cycles

Figure 8.9 shows the smallest amplitude OFCs that were detected within 6 cycles. This showcases the potential performance of each method if the detection time requirement was relaxed. Here, a detection time of longer than 6 cycles is considered to be a missed detection.

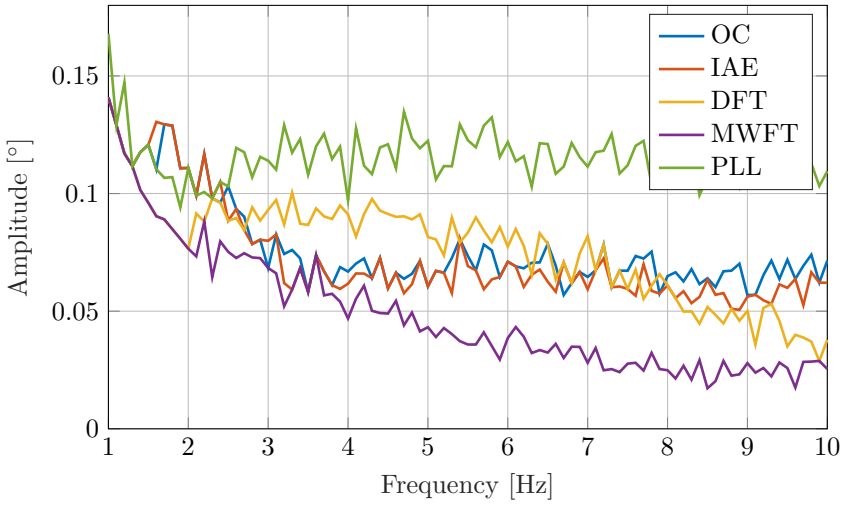

Figure 8.9: Smallest Amplitudes Detected within 6 Cycles (Liquid Failure at Rod Sensor)

Briefly, the MWFT appears to detect the smallest OFCs consistently across all frequencies, while the PLL shows the worst performance overall. Figure 8.10 compares these results to the thresholds of each detection method.

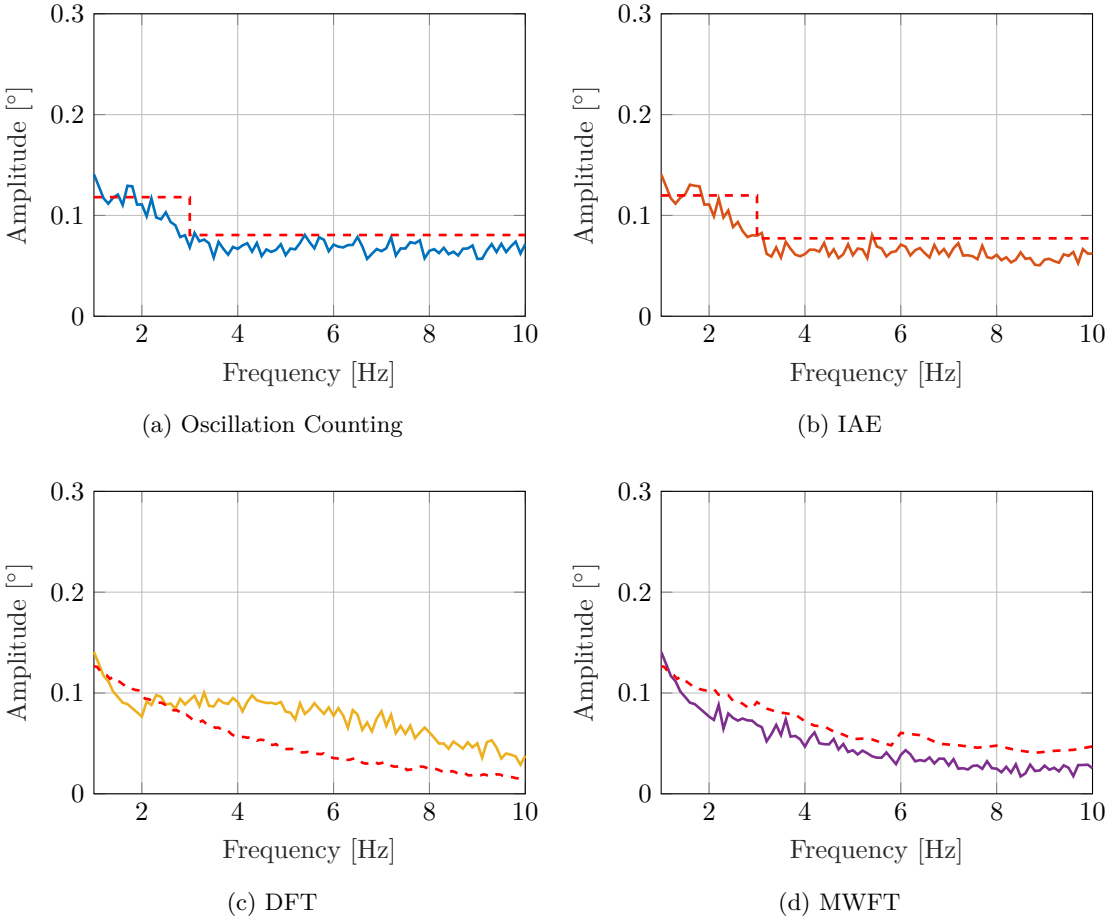

(a) Oscillation Counting

(b) IAE

(c) DFT

(d) MWFT

Figure 8.10: Minimum Amplitudes Detected and Corresponding Thresholds (Liquid Failure at Rod Sensor)

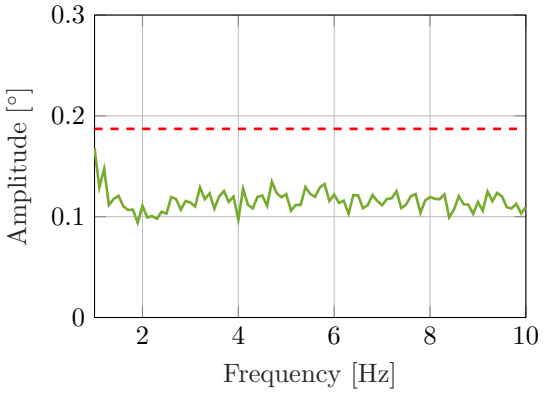

(e) PLL

Figure 8.10: Minimum Amplitudes Detected and Corresponding Thresholds (Liquid Failure at Rod Sensor) (Continued)

The oscillation counting and the IAE methods show very similar performance. This proves that the IAE is a competent approach to oscillation detection, and its only drawback is that it struggles to meet the detection time requirements of the OFC detection problem. Here, the MWFT continues to outperform the other detection methods.

For the most part, the various methods are capable of detecting OFCs that are slightly smaller than their thresholds suggest. Here, the noise and modelling errors "assist" the detection methods by adding energy to the failure. The exception to this is the DFT approach. The DFT uses a large window, and has to wait for multiple cycles, up to 30, to enter the observation window before the magnitude of the frequency component reaches its maximum value. The 6 cycle limit used here is therefore detrimental to the detection performance of the DFT.

8.4.3 Detection Time

Figure 8.11 shows the spread of the detection times for each detection method across all the tests in this test case. This takes into account all detections that occur within 6 cycles.

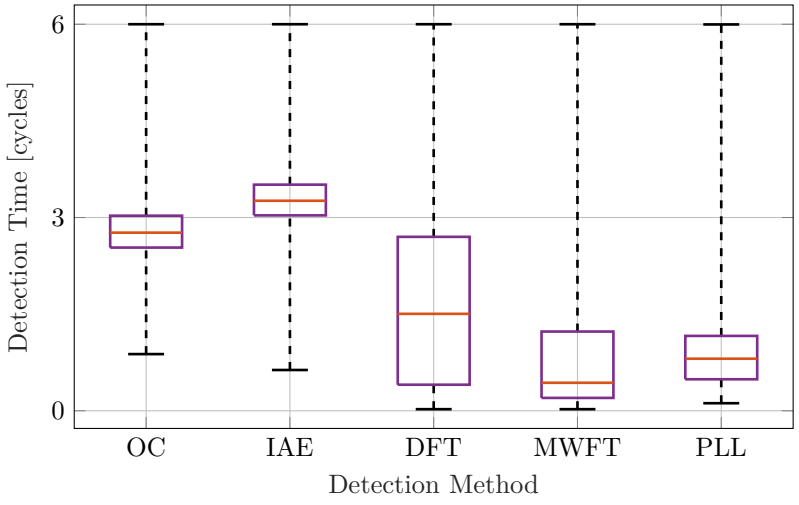

Figure 8.11: Detection Time of OFC Detection Methods in Cycles (Liquid Failure at Rod Sensor)

The MWFT outperforms all other techniques with the smallest median detection time of 0.44 cycles. In other words, 50 % of the failures that were detected achieved a detection time of 0.44

cycles or less. It manages to achieve this performance by using shorter windows to detect higher frequencies. The PLL attains a close second place with a median detection time of 0.81 cycles. This fast detection is the result of the PLL design that attempts to detect all OFCs within 0.3 seconds. The DFT also manages to achieve early detections due to its detection time being dependent on OFC amplitude, where larger amplitudes result in faster detections. However, it shows the greatest spread of all the methods. The oscillation counting approach manages a median detection time of just under 3 cycles, as expected given the fact that the 6^{th} threshold crossing occurs before the end of the 3^{rd} cycle. Finally, the IAE shows a median detection time of 3.26 cycles, proving that it offers a competent OFC detection technique, but suffers from the fact that the 6^{th} threshold crossing occurs at the very end of the 3^{rd} cycle. It is also noted that the oscillation counting and the IAE methods show the least spread, indicating that they provide the most consistent detection times.

Each detection method has a maximum detection time of 6 cycles. As mentioned, these longer detections likely occur in cases where the amplitude of the failure is just below the threshold level, and the failure is eventually detected due to the added energy of noise and modelling errors. The 6 cycle detection time result for the PLL, which is specifically designed for 0.3 second detection time, may also be indicative of cycle slipping.

Finally, the DFT and MWFT appear to show cases where the failure is detected almost instantaneously. Since the detection time of the DFT and MWFT is inversely proportional to the amplitude of an oscillatory failure, these extremely short detection times occur in the cases where the amplitude of the failure is exceptionally large.

8.5 Test Case: Liquid Failure at the Command Current

The second test case is the liquid failure at the command current. The results here are very similar to those of the position sensor liquid failure, and the results and discussions in the previous section are all applicable here. For completeness, some results are shown here. Figure 8.12 shows the smallest amplitudes detected within 3 cycles for each method, excluding the IAE. Figure 8.13 shows the distribution of the detection times.

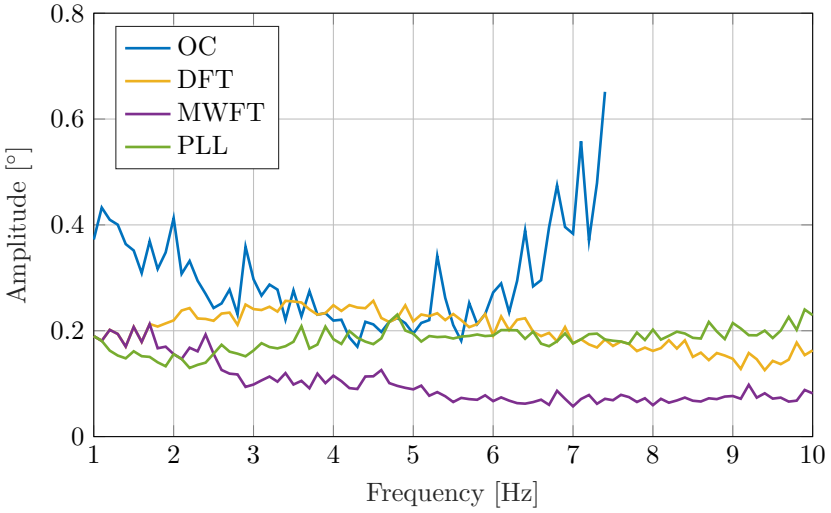

Figure 8.12: Smallest Amplitudes Detected within 3 Cycles (Liquid Failure at Command Current)

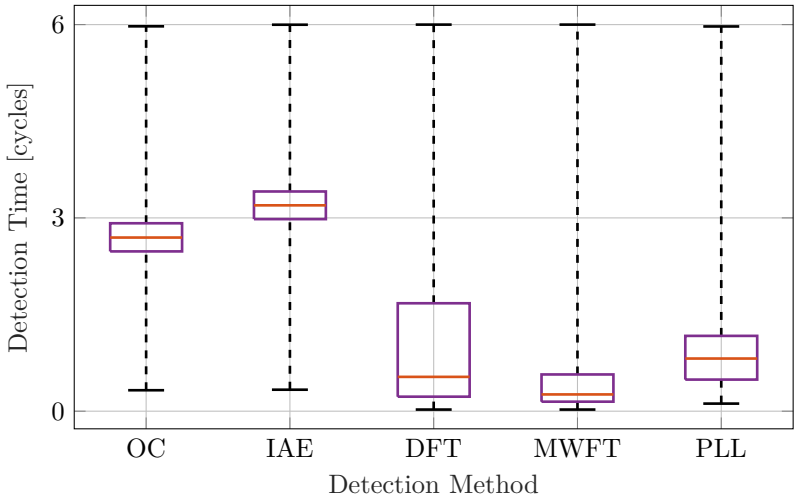

Figure 8.13: Detection Time of OFC Detection Methods in Cycles (Liquid Failure at Command Current)

It is noted here that the overall detection times for oscillation counting, the IAE, and the PLL are similar to the previous set of results, but the median detection times of the DFT and MWFT are considerably lower in this failure case. This is simply because the resultant OFC amplitude at the control surface tends to be larger when the OFC is injected at the command current than at the rod position sensor. This effect can be seen when comparing Sections 3.4.3 and 3.4.4. The DFT and MWFT therefore detect the OFCs more quickly because their detection times decrease with larger amplitudes.

8.6 Test Case: Solid Failure at the Command Current

8.6.1 Smallest OFCs Detected within 3 Cycles

This section presents the detection results for a solid failure at the command current. Figure 8.14 shows the smallest amplitudes detected in 3 cycles.

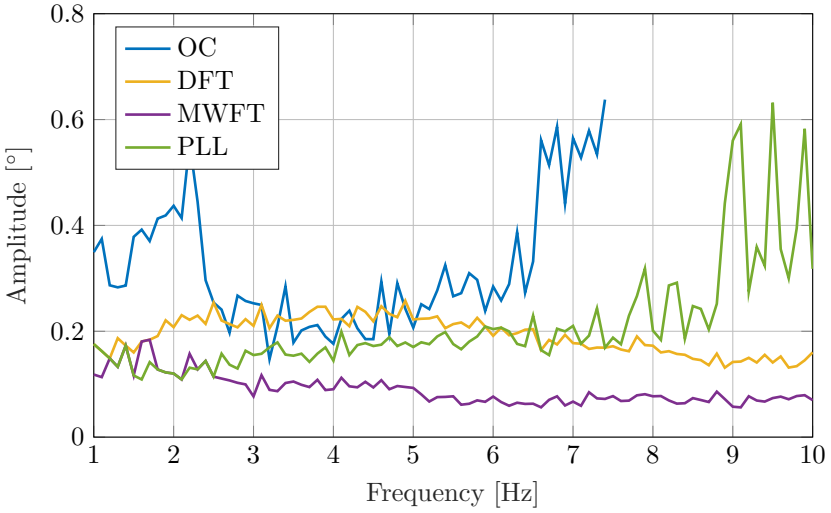

Figure 8.14: Smallest Amplitudes Detected within 3 Cycles (Solid Failure at Command Current)

For the most part, the results are similar to those in the previous cases. One exception that is immediately evident is the large spikes between 9 and 10 Hz for the PLL. This unfortunately occurs due to cycle slipping. The large frequency offset combined with the transient effects of the solid failure and the effects of the upsampling stage ultimately results in the PLL slipping one cycle. This trend can also be seen in Figure 8.6 for liquid failures, albeit to a much smaller extent. However, these late detections typically don't exceed 4 cycles, and the average detection time for the PLL in Figure 8.17 doesn't show any significant deterioration in comparison to the previous cases.

In terms of the other methods, the IAE (not shown) continues to show poor detection time performance, the oscillation counting cannot guarantee timely detection at frequencies higher than 7.5 Hz, and the MWFT continues to offer the best overall performance. The results of Figure 8.14 are summarised in the box-and-whisker diagram in Figure 8.15, which shows the distribution of the smallest amplitudes detectable within 3 cycles across the frequency range.

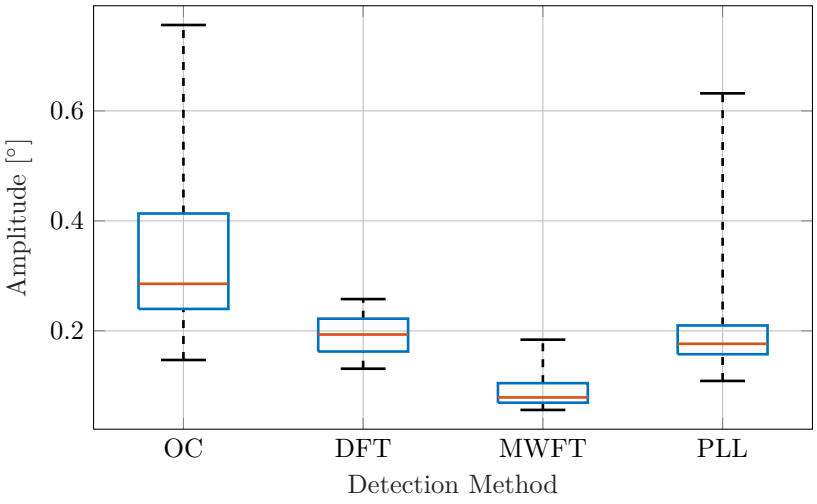

Figure 8.15: Range of Amplitudes Detected within 3 Cycles (Solid Failure at Command Current)

8.6.2 Smallest OFCs Detected within 6 Cycles

The smallest amplitudes for solid failures detected within 6 cycles are shown in Figure 8.16. This figure shows a curious result at low frequencies, where most of the methods appear to detect low-frequency OFCs with 0° amplitude. This is a case where the amplitude of the induced OFC is zero, but due to the nature of the solid failure, the actuator stops responding to reference commands. This is essentially a jammed actuator. Many of the oscillation detection methods manage to detect the failure due to filter transients, or in the case of the DFT methods, the frequency content of the position estimate crossing the thresholds. These jamming detections seem to occur only at low frequencies simply because of the 6 cycle detection time specified, and 6 cycles of the low frequencies stretches over more time than at high frequencies, providing the detection methods with longer opportunities to detect these jamming failures.

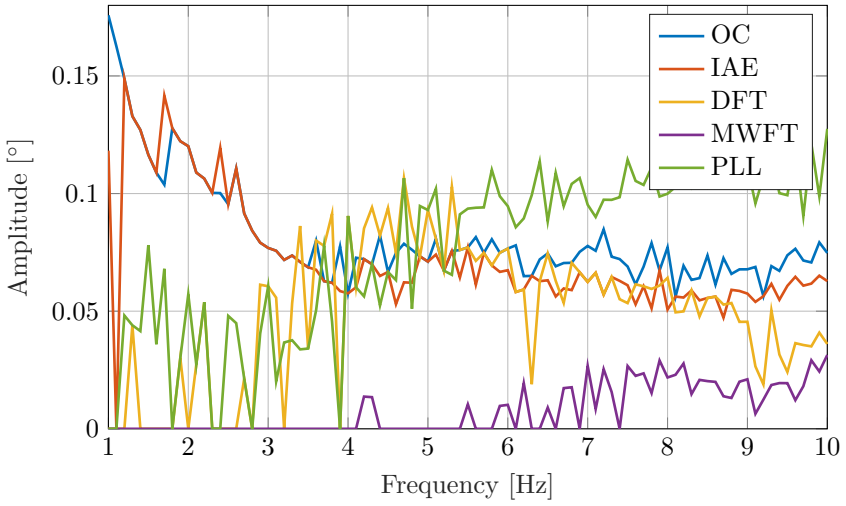

Figure 8.16: Smallest Amplitudes Detected within 6 Cycles (Solid Failure at Command Current)

8.6.3 Detection Time

Finally, the distribution of the detection times for each method is summarised in Figure 8.17 below. These results are very similar to those shown for the liquid failure in the command current.

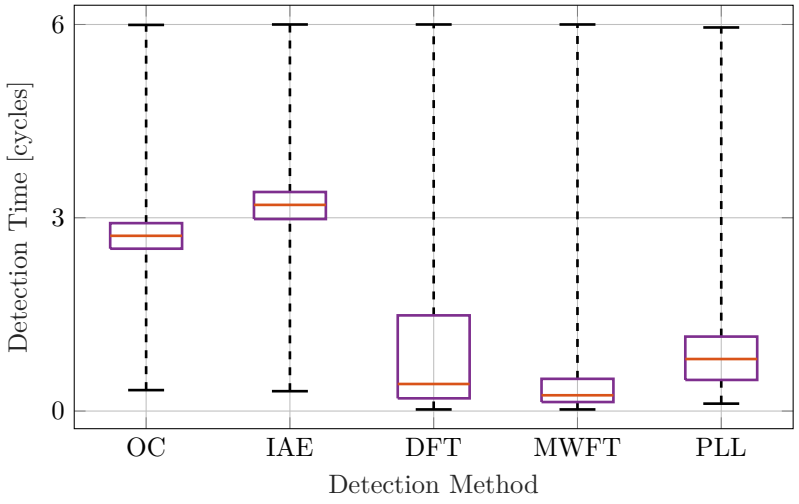

Figure 8.17: Detection Time of OFC Detection Methods in Cycles (Solid Failure at Command Current)

8.7 False Detections

With all the detection methods exhaustively and robustly trained with 1000 training data sets, it is expected that no false detections should occur for any detection methods. A false alarm occurs if a detection method flags a failure before the OFC is injected into the system. Table 8.1 lists the number of false alarms that occurred for each detection method over all the 180 180 test data sets. The DFT shows a large number of false alarms, with a false alarm rate of almost 0.04 %, while the MWFT suffered from just one false alarm. All the other methods experienced no false alarms.

Table 8.1: Number of False Detections during Testing

Detection Method	False Detections	False Detection Rate
Oscillation Counting	0	0 %
Integrated Absolute Error	0	0 %
Discrete Fourier Transform	68	$37.7 \cdot 10^{-3}$ %
Multi-window Fourier Transform	1	$0.55 \cdot 10^{-3}$ %
Phase-locked Loop	0	0 %

The DFT's false alarms occurred due to high-frequency components that originate from sensor noise. This is primarily due to the high-resolution threshold of the DFT that assigns a unique threshold for each frequency component, and the threshold values at high frequencies are not large enough to avoid false alarms, since most of the energy of the residual lies at low frequencies. The training data therefore failed to train the DFT's high-frequency threshold values sufficiently. Methods such as oscillation counting, on the other hand, only have two unique threshold values, and the threshold for high frequencies has to be high enough to avoid lower-frequency threshold crossings as well. The MWFT has a considerably higher resolution threshold than other methods, but lower than the DFT, which is likely the reason for its improved robustness.

To prevent false detections from occurring, the thresholds for the DFT and the MWFT could be scaled by some safety margin, at the cost of the detection of smaller OFCs.

8.8 Computational Complexity

The final performance metric that must be quantified is the computational complexity of each OFC detection method. Here, in Table 8.2, the computational requirements of each detection method is given in terms of the number of mathematical operations per 40 Hz time step that must be performed, specifically the number of real multiplications and real additions. The derivations of these results can be found in Appendices B to E.

Table 8.2: Mathematical Operations Required for Each Detection Method

Detection Method	Real Multiplications	Real Additions	Other
Oscillation Counting	120	120	
Integrated Absolute Error	120	132	
Discrete Fourier Transform	1088	2040	
Multi-window Fourier Transform	408	765	
Phase-locked Loop	285	254	Linear/Binary Search

Oscillation counting manages to be the most computationally efficient approach for OFC detection, despite the upsampling factor of 3. The IAE is a close second with only 12 more additions over the oscillation counting.

The PLL requires little more than double the number of operations, assuming a lookup table is used to calculate the cosine of the VCO. However, the use of a lookup table requires a search algorithm to obtain the correct lookup value pair. Two popular algorithms that could be applied here are linear or binary search algorithms. The computational order of a linear search is $\mathcal{O}(N)$, and a binary search has an order of $\mathcal{O}(\log N)$. The computational cost of both of these approaches is dependent on the chosen length of the lookup table.

The DFT and MWFT are considerably more expensive than the rest of the methods. The DFT requires 9 times the number of multiplications versus oscillation counting, and 17 times the number of additions. The MWFT, however, only requires an increase of 3.4 times the number of multiplications and 6.4 times the number of additions over oscillation counting, making it more viable

for embedded applications than the DFT. It should also be noted that the structure of the sliding DFT allows it to be parallelised, leading to faster calculations provided the hardware allows for it.

Note that the computational complexity of each method given here may be more useful when converted to the number of clock cycles in an embedded computer. Typically, it can be assumed that a real addition may take 1 clock cycle, while multiplication may take 3 or 4 clock cycles. However, this is dependent on the hardware on which these techniques are implemented.

8.9 Summary

This chapter detailed the results of a rigorous test campaign that tested and compared five different OFC detection methods. The results were compared based on smallest amplitude reliably detected within 3 cycles, smallest amplitude detected within 6 cycles, average detection time, and false alarm rate. Finally, the computational complexity of each method was compared.

Oscillation counting and the integrated absolute error (IAE), the two oscillation detection techniques extracted from literature, provided the worst performance in terms of smallest amplitude detected within 3 cycles. Neither methods could guarantee 3 cycle detection time for the entire 1 to 10 Hz spectrum of OFC frequencies. They also had the longest average detection times. Oscillation counting was, in general, better than the IAE by having faster detection times and using slightly fewer mathematical operations. Additionally, neither of these methods triggered any false alarms, and both are extremely computationally efficient.

The phase-locked loop (PLL) provided good overall performance, and is approximately twice as computationally expensive as oscillation counting. It was designed to ensure 0.3 second detection time for all frequencies, and thus had extremely short detection times. It also offered more consistent performance than the first two methods in terms of smallest amplitude detectable within 3 cycles. However, despite being designed to avoid cycle slipping, the PLL did suffer from cycle slipping at higher frequencies, especially during solid failures.

The discrete Fourier transform (DFT) and multi-window Fourier transform (MWFT) are by far the most computationally expensive approaches. The DFT performed worse than the PLL in terms of detection time and amplitude, but still managed to improve on the oscillation counting and IAE methods in these regards. Unfortunately, it suffered from many false detections caused by high-frequency sensor noise. The MWFT in contrast is considerably more efficient than the DFT. This is achieved by using smaller windows for higher frequencies. These smaller windows also allowed the MWFT to detect the smallest OFCs within 3 cycles, and within 6 cycles, and it maintained the smallest average detection time of all the tested methods. The MWFT also suffered from fewer false alarms than the DFT. Due to its superior performance, the MWFT is the recommended approach to OFC detection.

8.10 Contributions

The following contributions were made in this chapter:

- A Monte Carlo simulation was set up to test all the different approaches. This was based on the longitudinal model developed in Chapter 3. 1000 training data sets were generated to train the thresholds of each detection method. A total of 180 180 unique testing data sets were generated and used to test each detection method rigorously.
- Each method was evaluated and compared based on their ability to detect small amplitude OFCs within 3 cycles and within 6 cycles. The detection time statistics of each method were also compared and discussed.
- An analysis of the computational complexity of each method was performed and detailed in Appendices B to E.

Chapter 9

Conclusions and Recommendations

This thesis investigated detection strategies for the oscillatory failure case (OFC), a specific type of failure that causes aircraft control surfaces to oscillate uncontrollably. These oscillatory failures cause additional loads to be placed on the structure of the aircraft, and to mitigate this, structural engineers have to reinforce the body of the aircraft to handle these large loads. The structural reinforcements increase the overall weight of the aircraft, reducing fuel efficiency, flight time, and handling quality. Therefore, if a fault detection system could detect and pacify these oscillations quickly, the structural reinforcement could be reduced.

Fault detection and reconfiguration systems are already integrated into the flight actuator control systems. These fault detection systems make use of analytical redundancy to detect faults by simulating an actuator, and subtracting the position of the simulated actuator from the measurements of the real actuator. This results in a residual signal that represents the difference between the real actuator and the simulated actuator. Residual evaluation techniques are then used to check for the presence of failures.

The goal of this project was to investigate residual evaluation techniques that can detect OFCs. These techniques should be able to detect OFCs with as small amplitudes as possible, and the OFCs should be detected within 3 cycles to prevent sustained loads. Furthermore, the online and embedded nature of the fault detection system limits the computational resources available, and thus residual evaluation techniques need to be computationally efficient.

9.1 Summary of Work Done

Two fault detection methods for oscillation detection found in literature were investigated and implemented. These methods are oscillation counting [1] and integrated absolute error (IAE) [17]. Three more methods were designed and implemented, based on the discrete Fourier transform (DFT) and phase-locked loop (PLL) techniques. These five approaches are summarised as follows:

1. Oscillation counting detects oscillatory failures by setting a time-domain threshold, and counting the number of times the residual signal crosses the threshold. An OFC is detected if six threshold crossings occur over a certain time period. Two bandpass filters are used to divide the residual into different frequency sub-bands, and to remove noise. However, these filters add transient effects and phase delays to the residual, creating the potential for late detections.

2. The IAE technique integrates the residual signal between consecutive zero crossings. If the IAE is greater than some threshold value, a disturbance is detected. If six disturbances are detected, an OFC is inferred. Unfortunately, due to the integration, the magnitude of the IAE is dependent on the frequency of the OFC, with higher frequencies showing smaller magnitudes. Therefore, a new dynamic threshold selector is implemented that scales the threshold based on the time between the two zero crossings. Sub-band filters are used here as well to ensure fair comparison with oscillation counting.

3. The DFT approach calculates the magnitude of the frequency content in the residual signal. If the magnitude of any frequency components exceed some threshold, then an OFC is inferred. A new approach to threshold selection is proposed here, where a frequency-dependent threshold is used. This threshold allows for higher thresholds at frequencies where the fault-free residual typically contains more energy, and lower at frequencies that typically contain less energy. The detection time of the DFT approach is mathematically proven to be dependent on OFC amplitude, frequency, window size, and the selected threshold. The use of a computationally efficient approach to calculating the DFT, known as the sliding-DFT, is investigated and recommended in preference to the fast Fourier transform.

4. The nature of the sliding-DFT is exploited to allow for a multi-window Fourier transform (MWFT), where shorter window sizes are used to calculate higher frequency components to improve detection time.

5. The PLL technique detects OFCs by attempting to lock onto any dominant frequencies that exist in the residual signal. A quadrature phase detector then multiplies the PLL's output with the residual signal, and OFCs are detected by checking for a DC offset at the quadrature phase detector's output. The PLL is designed using classical control theory and non-linear phase plane analysis to guarantee 0.3 second detection time. However, the PLL struggled to guarantee detection time at low frequencies, and thus relied on an upsampling and modulation stage to increase the perceived frequency of the OFC.

A simulation framework was created to serve as a testbed for the training and testing of the different OFC detection techniques. The simulation framework contains models for the physical actuator, the analytically redundant actuator, the oscillatory failures (both liquid and solid failures), the flight control system, and the aircraft longitudinal dynamics. It is important to model the aircraft response and the flight control system response to the oscillatory actuator failure, since it affects the performance of the OFC detection. The simulation framework is designed so that the actuator parameters and the oscillatory failure parameters (amplitude, frequency, type, and injection point) can be varied randomly to produce training and testing data that is sufficiently rich.

9.2 Summary of Results

Each detection technique was subjected to a rigorous test campaign to determine the lowest amplitude OFC that could be detected within 3 cycles. The results of the tests showed that the oscillation counting technique could not guarantee detection time for frequencies higher than 7.5 Hz. The second approach, the IAE, essentially failed to guarantee detection time for any OFCs, and was therefore deemed unsuitable for the application of OFC detection. However, if the detection time requirements were relaxed, both the oscillation counting and IAE methods would be able to detect small OFCs, with detection performance comparable to the DFT and the PLL. They are also the most computationally efficient approaches.

Also tested were the PLL and two implementations of the DFT, namely the single-window and multi-window DFTs. These three methods managed to detect OFCs of all frequencies within the required detection time, unlike oscillation counting or the IAE. They also managed to detect smaller OFCs within 3 cycles. Finally, the performance of all three methods in the experimental tests matched or improved on the performance predicted in theory. The exception was the PLL that occasionally showed cycle slipping at high frequencies.

Ultimately, the MWFT shows the most promising results of all the methods, with superior detection time and smallest detectable amplitudes. This, combined with its computational optimisations over the single-window DFT, makes it the recommended approach to OFC detection.

9.3 Conclusions

The oscillation counting approach and the IAE approach have both been used successfully in industry, and offer simple and computationally efficient algorithms for oscillation detection. They work especially well when detection time is not a critical requirement. Unfortunately, their use of

bandpass filters for improved robustness causes time delays. This combined with their reliance on multiple threshold crossings places them at a disadvantage when applied to systems where strict detection time is a necessity. These methods require that 3 cycles pass before the system is alerted to the failure. This is proven with the average detection times of 2.8 and 3.3 cycles for oscillation counting and IAE respectively.

The DFT and PLL methods each rely on a single threshold crossing for OFC detection, and they therefore inherently allow for shorter detection times than the methods that require six threshold crossings. This does potentially come with an increased susceptibility to false alarms, as seen with the DFT that suffered from a false alarm rate of 0.04 % in the tests.

The PLL is computationally efficient, and offers an improvement over oscillation counting in terms of detection time. However, of all the methods, it is by far the most complex to implement as a system, with multiple interconnected components such as the upsampling and modulation stage, the control loop, and the quadrature phase detector. Designing the system is also more challenging, with many aspects that have to be taken into account, such as settling time, cycle slipping, and gain variation based on the input amplitude.

The MWFT offers a computationally efficient approach to OFC detection, and offers the best performance in terms of amplitude and detection time. It relies on the sliding-DFT, which is a very simple algorithm to implement, and is more efficient than the fast Fourier transform. Additionally, it does not require a high sampling rate to function correctly. The MWFT is therefore the recommended approach to OFC detection.

9.4 Recommendations and Future Work

- This study has shown that improvements to the OFC detection problem can be achieved by reviewing and redesigning the residual evaluation stage. The techniques in this study has been limited to time-domain thresholding, the DFT, and the PLL, but many other residual evaluation techniques may prove to be applicable. Further areas of investigation could include:
  - Signal estimation: The literature study briefly mentioned signal estimation techniques such as MUSIC and adaptive notch filters for frequency and amplitude estimation of sinusoids.
  - Telecommunications techniques: As mentioned, the telecoms field has a strong background in the detection of signals in noise, and while many of these techniques may prove to be too computationally expensive, valuable insights and techniques may be obtained.
  - Neural networks: Machine learning and deep learning techniques are steadily becoming more and more common in many different applications, and may offer unique techniques for residual generation and/or evaluation.
- Ultimately, residual evaluation techniques will be limited by the quality of the residual generation stage. More accurate actuator models will reduce the effect of modelling errors in the residual, allowing for lower thresholds, and will consequently enable the detection of smaller amplitude failures. Another approach to improve the performance and robustness of the system is the use of adaptive thresholds. This would involve dynamically changing the threshold based on the characteristics of the command signal, aircraft states, or actuator states. This is motivated by the observations in Section 3.4, where the energy in the residual is often dependent on the command signal.
- The methods investigated in this thesis could also be improved, specifically the DFT-based methods. The DFT technique suffers from an increase in the OFC detection threshold when the window size is decreased. This effect is evident in the detection thresholds of the DFT and the MWFT. On the other hand, smaller windows result in a decrease in computational cost, and ensures shorter detection times. Future investigations could attempt to optimise the size of the window for low thresholds, fast detection times, and minimum number of bins

to be computed. This optimisation would be applicable for both the DFT and the MWFT techniques. The MWFT could also be optimised by increasing the number of different window sizes, rather than using just four window sizes as in the proof of concept presented in this thesis.

- The MWFT relies on the sliding-DFT, which is a marginally stable filter, and its stability on aircraft embedded computers must be evaluated. The use of a stable implementation, as mention in Section 6.6, and its effect on accuracy should also be further investigated.
- One final recommendation is that a sampling frequency of at least 100 Hz should be used for any time-domain detection techniques, as explained in Section 4.3.3. This will ensure that sufficient amplitude information is available for the detection methods.

Appendix A

Derivation of the Sliding Discrete Fourier Transform

The sliding discrete Fourier transform is derived from the discrete Fourier transform formula. Starting at some time step n, the k^{th} frequency bin is described as

$$R[k,n] = \sum_{n=n_o}^{n_o+N-1} r[n]e^{-j2\pi k(n-n_o)/N}.$$

At the next time step,

$$\begin{aligned}
R[k,n+1] &= \sum_{n=n_o}^{n_o+N-1} r[n+1]e^{-j2\pi k(n-n_o)/N} \\
&= \sum_{n=n_o+1}^{n_o+N} r[n]e^{-j2\pi k(n-n_o-1)/N} \\
&= \sum_{n=n_o}^{n_o+N-1} r[n]e^{-j2\pi k(n-n_o-1)/N} - r[n_o]e^{-j2\pi k(-1)/N} + r[n_o+N]e^{-j2\pi k(N-1)/N} \\
&= e^{j2\pi k/N}\left(\sum_{n=n_o}^{n_o+N-1} r[n]e^{-j2\pi k(n-n_o)/N} - r[n_o] + r[n_o+N]\right) \\
&= e^{j2\pi k/N}\left(R[k,n] - r[n_o] + r[n_o+N]\right)
\end{aligned}$$

Appendix B

Oscillation Counting Computational Evaluation

This appendix evaluates the number of mathematical operations required per time step to implement oscillation counting, based on the system overview in Figure B.1.

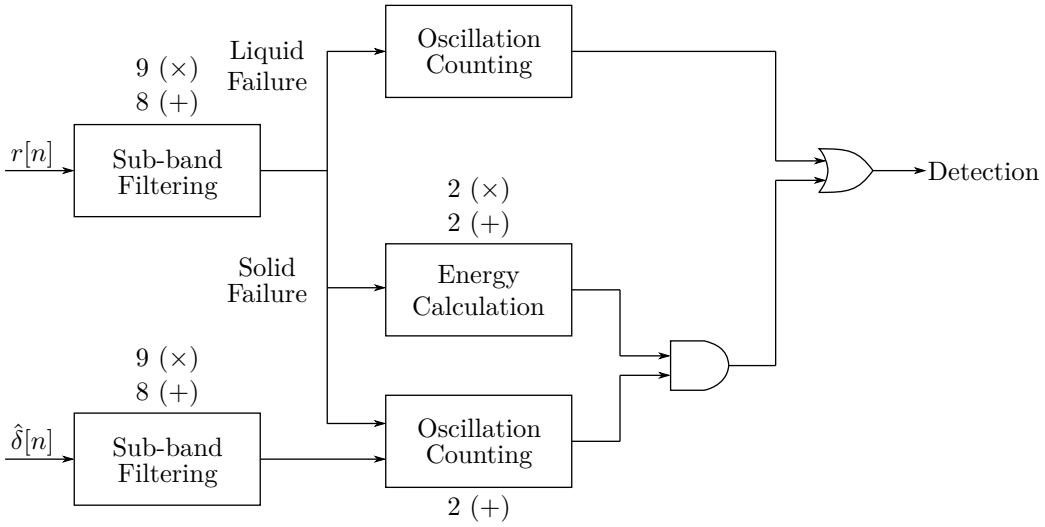

Figure B.1: Real Additions and Multiplications in Oscillation Counting

- The residual and the estimated control surface deflection are filtered using 4^{th} order infinite impulse filters. The difference equations of these filters can be expressed as:

$$y[n] = \sum_{i=0}^{4} b_i r[n-i] - \sum_{j=1}^{4} a_j y[n-j] \tag{B.0.1}$$

  At each sample time, Equation (B.0.1) executes 9 multiplications and 8 additions. Figure 4.3 includes two such filters, and thus the total number of filtering-related calculations are 18 multiplications and 16 additions.

- Two oscillation counting blocks are present. The liquid failure counting requires no mathematical operations, but the solid counting offsets the positive and negative thresholds by the deflection estimate. This executes 2 additions.

- The energy calculation can be implemented with the following iterative filter, consisting of two multiplications and two additions:

$$E[n] = r^2[n] + E[n-1] - r^2[n-1] \tag{B.0.2}$$

Thus the total number of operations are 20 multiplications and 20 additions. However, the entire process is performed twice for two frequency sub-bands, and an upsampling factor of 3 further increases the number of computations per 40 Hz sample time. This results in a final result of 120 real multiplications and 120 real additions.

Appendix C

Integrated Absolute Error Computational Evaluation

The IAE has a very similar computational requirement to the oscillation counting, as the structure of the system is nearly identical, if separate checks are in place for liquid and solid failures. The computational requirements for oscillation counting are detailed in Appendix B

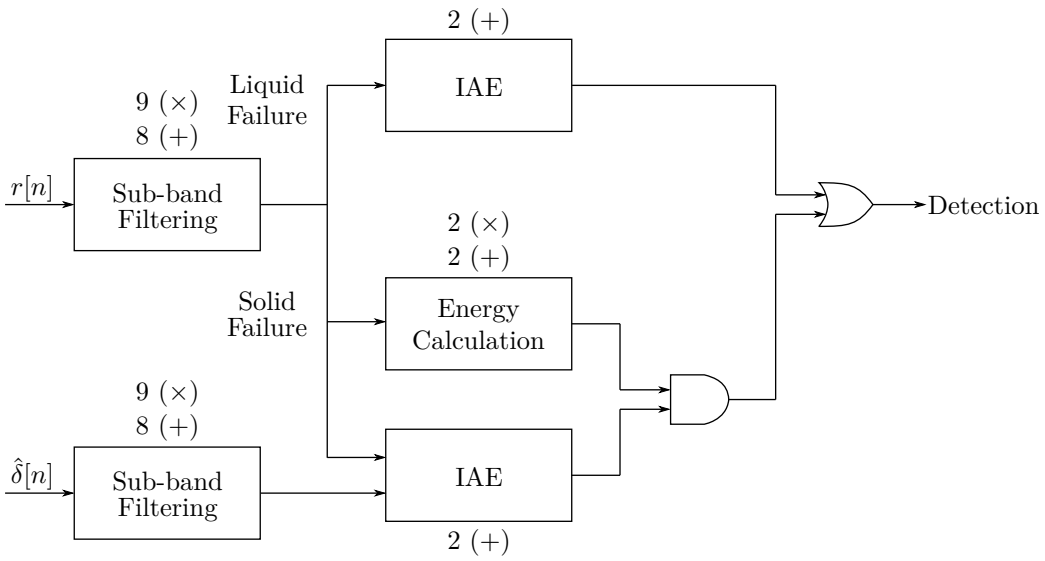

Figure C.1: Real Additions and Multiplications in the IAE Method

The IAE trades the two oscillation counting blocks for two integration stages. Each of these contain two additions: one to integrate the residual and one to increment the integration time. Additionally, the solid failure check does not have to offset the threshold, since the zero-crossing checks will occur around the estimated deflection rather than around zero.

Note that when a zero-crossing is detected, the threshold value must be calculated. Because this does not happen at every sample instance, it is ignored for this analysis.

The total number of operations for the IAE is therefore 20 multiplications and 22 additions. Taking upsampling and the two frequency sub-bands into account, the totals rise to 120 and 132 multiplications and additions respectively.

Appendix D

Sliding DFT Computational Evaluation

To reiterate, the sliding DFT, with zero-padding, can be mathematically expressed as:

$$R[k,n] = e^{j2\pi k/m_z N}\left(R[k,n-1] - r[n-N] + r[n]e^{-j2\pi k/m_z}\right). \tag{D.0.1}$$

This shows two complex multiplications and two complex additions per frequency bin at each time step. Here, the two exponentials are assumed to be pre-calculated and stored as lookup tables.

For fair comparison to the other implementations, the number of real operations is required. Each complex multiplication consists of the following operations [50]:

$$\Re[(a+jb)(c+jd)] = ac - bd \tag{D.0.2}$$
$$\Im[(a+jb)(c+jd)] = (a+b)(c+d) - ac - bd \tag{D.0.3}$$

This operation contains three multiplies and five additions. Complex addition uses two real additions. Therefore, the total number of real operations per frequency bin is 6 multiplications and 14 additions.

The result of Equation (D.0.1), $R[k,n]$, is complex, and the absolute value is required for threshold checking. This typically involves the Euclidean distance equation, which includes the use of a square-root operator:

$$|R[k,n]| = \sqrt{\Re(R[k,n])^2 + \Im(R[k,n])^2} \tag{D.0.4}$$

However, the detection algorithm could make use of the periodogram, which is defined as the square of the DFT's magnitude. This would cancel out the square root, simplifying the calculation to just two multiplications and one addition. Thus, the number of operations increases to 8 multiplications and 15 additions.

Section 6.6 states that with a zero-padding factor of 5, a total of 136 frequency bins must be computed. Thus, the total number of computations are 1088 multiplications and 2040 additions.

The multi-window implementation of the DFT requires only 51 frequency bins to be calculated, reducing the total number of operations to 408 multiplications and 765 additions.

Appendix E

PLL Computational Evaluation

The phase-locked loop, Figure E.1 consists of the following mathematical operations:

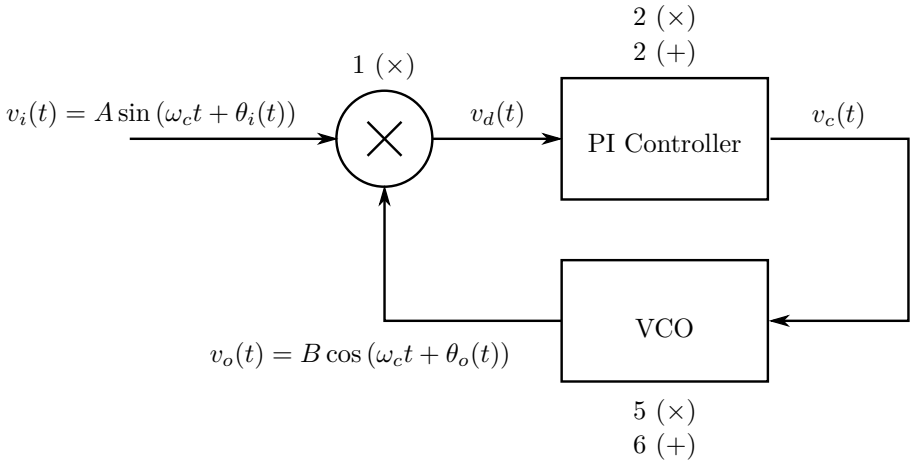

Figure E.1: Phase-locked Loop

- The phase detector contributes a single real multiplier.
- The PLL's PI controller is described in the Z-domain as:

$$D(z) = P + I\frac{1+z^{-1}}{1-z^{-1}} \tag{E.0.1}$$

This translates to the following difference equation:

$$y[n] = (I+P)x[n] + (I-P)x[n-1] + y[n-1] \tag{E.0.2}$$

The values $(I+P)$ and $(I-P)$ can be hard-coded, thereby requiring two additions and two multiplications.

- The VCO has to generate a sinusoid based on:

$$v_o(t) = \cos\left(\omega_c t + \int v_c(t)dt\right) \tag{E.0.3}$$

The phase of the cosine above can be calculated with two multiplications and three additions:

$$\theta[n] = \omega_c F_s n + \theta_o[n], \quad \theta_o[n] = \frac{T_s}{2}(v_c[n] + v_c[n+1]) + \theta_o[n-1] \tag{E.0.4}$$

The trigonometric calculation is more difficult to quantify. An efficient implementation could be to save a precomputed cosine as a lookup table, at the cost of accuracy. The accuracy can

be improved to an extent by interpolating between the saved samples. Linear interpolation between two points (x_1, y_1) and (x_2, y_2) can be expressed as

$$y = \frac{y_1(x_2 - x) - y_2(x_1 - x)}{x_2 - x_1}, \tag{E.0.5}$$

with three multiplications and three additions, if the step size $(x_2 - x_1)$ is known beforehand. Linear or binary search is also required here to obtain the cosine value, but this is dependent on the total number of points stored, and is ignored for this analysis.

OFC detection is performed using a quadrature phase detector, shown in Figure E.2.

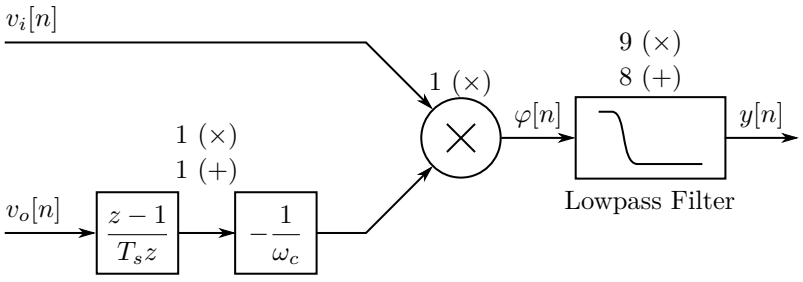

Figure E.2: Quadrature Phase Detector

- The QPD also has a phase detector, which is a single multiplication.
- The differentiator can be implemented as

$$y[n] = \frac{-1}{T_s \omega_c}(v_o[n] - v_o[n-1]) \tag{E.0.6}$$

With the fractional constant pre-calculated as a single value, this is one multiplication and one addition.

- Finally, the output is filtered with a 4^{th} order IIR filter, with nine multiplications and eight additions, as described in Appendix B.

Finally, the upsampling stage is analysed in Figure E.3.

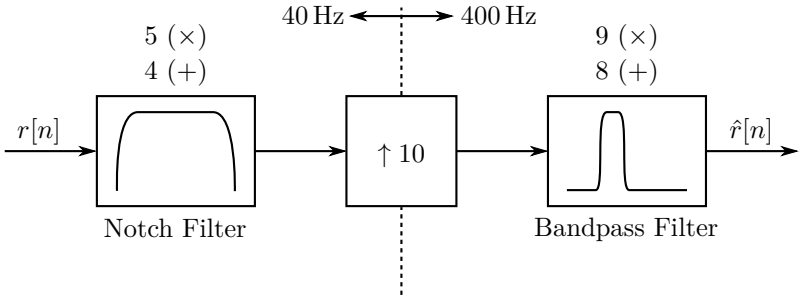

Figure E.3: Upsampling Stage Block Diagram

The total number of computations of the preceding elements, and the bandpass filter in Figure E.3 total to 28 multiplications and 25 additions. However, these elements run at an upsampled frequency, resulting in a total of 280 multiplications and 250 additions per 40 Hz time step.

Before the upsampling, there is a 2^{nd} order notch filter that runs at 40 Hz. The notch filter has five multiplications and four additions. In conclusion, the PLL with an upscaling factor of 10 uses a total of 285 multiplications and 254 additions.

Bibliography

[1] P. Goupil, "Oscillatory failure case detection in the A380 electrical flight control system by analytical redundancy", *Control Engineering Practice*, vol. 18, no. 9, pp. 1110–1119, 2010 (Cited on pages iii, iv, 1–3, 7, 8, 16, 27, 29, 30, 39, 99).

[2] D. Van Den Bossche, "The a380 flight control electrohydrostatic actuators, achievements and lessons learnt", *25th International Congress of the Aeronautical Sciences*, Jan. 2006 (Cited on page 1).

[3] AIRBUS, "A380 Aircraft Characteristics – Airport and Maintenance Planning", 2005 (Cited on pages 1, 12).

[4] H. Sachs, U. Carl and F. Thielecke, "Monitoring of oscillatory failure cases for the reduction of structural loads in an aeroelastic aircraft", *26th International Congress of the Aeronautical Sciences*, 2008 (Cited on page 2).

[5] E. Chow and A. Willsky, "Analytical redundancy and the design of robust failure detection systems", *IEEE Transactions on Automatic Control*, vol. 29, no. 7, pp. 603–614, 1984 (Cited on page 3).

[6] H. Alwi and C. Edwards, "Oscillatory failure case detection for aircraft using an adaptive sliding mode differentiator scheme", in *Proceedings of the 2011 American Control Conference*, 2011, pp. 1384–1389 (Cited on page 7).

[7] A. Levant, "Robust exact differentiation via sliding mode technique", *Automatica*, vol. 34, no. 3, pp. 379–384, 1998 (Cited on page 7).

[8] E. Alcorta-Garcia, A. Zolghadri and P. Goupil, "A Nonlinear Observer-Based Strategy for Aircraft Oscillatory Failure Detection: A380 Case Study", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 47, no. 4, pp. 2792–2806, 2011 (Cited on page 7).

[9] L. Lavigne, A. Zolghadri, P. Goupil and P. Simon, "Oscillatory Failure Case detection for new generation airbus aircraft: A model-based challenge", in *47th IEEE Conference on Decision and Control*, 2008, pp. 1249–1254 (Cited on pages 7, 8).

[10] A. Varga and D. Ossmann, "Lpv model-based robust diagnosis of flight actuator faults", *Control Engineering Practice*, vol. 31, pp. 135–147, 2014 (Cited on page 8).

[11] K. S. Narendra and J. Balakrishnan, "Adaptive control using multiple models", *IEEE Transactions on Automatic Control*, vol. 42, no. 2, pp. 171–187, 1997 (Cited on page 8).

[12] E. A. Morelli, "Real-time parameter estimation in the frequency domain", *Journal of Guidance, Control, and Dynamics*, vol. 23, no. 5, pp. 812–818, 2000 (Cited on page 8).

[13] L. Lavigne, F. Cazaurang, L. Fadiga and P. Goupil, "New sequential probability ratio test: Validation on a380 flight data", *Control Engineering Practice*, vol. 22, pp. 1–9, 2014 (Cited on page 8).

[14] A. Wald, "Sequential Tests of Statistical Hypotheses", *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117–186, 1945 (Cited on page 8).

[15] H. Khan, S. C. Abou and N. Sepehri, "Nonlinear observer-based fault detection technique for electro-hydraulic servo-positioning systems", *Mechatronics*, vol. 15, no. 9, pp. 1037–1059, 2005 (Cited on page 8).

[16] K. Uosaki and T. Kagawa, "Backward sprt failure detection in multivariate observation case", in *Proceedings IECON '91: 1991 International Conference on Industrial Electronics, Control and Instrumentation*, 1991, 2169–2174 vol.3 (Cited on page 8).

[17] T. Hägglund, "A control-loop performance monitor", *Control Engineering Practice*, vol. 3, no. 11, pp. 1543–1551, 1995 (Cited on pages 8, 38, 39, 99).

[18] K. Forsman and A. Stattin, "A new criterion for detecting oscillations in control loops", in *1999 European Control Conference (ECC)*, 1999, pp. 2313–2316 (Cited on page 8).

[19] N. Thornhill, B. Huang and H. Zhang, "Detection of multiple oscillations in control loops", *Journal of Process Control*, vol. 13, no. 1, pp. 91–100, 2003 (Cited on page 9).

[20] B. P. Lathi, *Modern Digital and Analog Communication Systems*, International 4th ed. New York: Oxford University Press, 2010 (Cited on pages 9, 61, 80).

[21] R. Srinivasan, R. Rengaswamy and R. Miller, "A modified empirical mode decomposition (emd) process for oscillation characterization in control loops", *Control Engineering Practice*, vol. 15, no. 9, pp. 1135–1148, 2007 (Cited on pages 9, 10).

[22] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N.-C. Yen, C. C. Tung and H. H. Liu, "The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time series analysis", *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 454, no. 1971, pp. 903–995, 1998 (Cited on page 9).

[23] P. F. Odgaard and K. Trangbaek, "Comparison of Methods for Oscillation Detection – Case Study on a Coal-fired Power Plant", *IFAC Proceedings Volumes*, vol. 39, no. 7, pp. 297–302, 2006, 5th IFAC Symposium on Power Plants and Power Systems Control (Cited on page 9).

[24] C. Kar and A. Mohanty, "Vibration and current transient monitoring for gearbox fault detection using multiresolution fourier transform", *Journal of Sound and Vibration*, vol. 311, no. 1, pp. 109–132, 2008 (Cited on page 10).

[25] T. Chen, C. Wang and D. J. Hill, "Small oscillation fault detection for a class of nonlinear systems with output measurements using deterministic learning", *Systems & Control Letters*, vol. 79, pp. 39–46, 2015 (Cited on page 10).

[26] X. Guo and M. Bodson, "Adaptive rejection of multiple sinusoids of unknown frequency", in *2007 European Control Conference (ECC)*, 2007, pp. 121–128 (Cited on page 10).

[27] R. Schmidt, "Multiple emitter location and signal parameter estimation", *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, 1986 (Cited on page 10).

[28] R. Roy and T. Kailath, "Esprit-estimation of signal parameters via rotational invariance techniques", *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 7, pp. 984–995, 1989 (Cited on page 10).

[29] P. Stoica, H. Li and J. Li, "Amplitude estimation of sinusoidal signals: Survey, new results, and an application", *IEEE Transactions on Signal Processing*, vol. 48, no. 2, pp. 338–352, 2000 (Cited on page 11).

[30] L. Hsu, R. Ortega and G. Damm, "A globally convergent frequency estimator", *IEEE Transactions on Automatic Control*, vol. 44, no. 4, pp. 698–713, 1999 (Cited on page 11).

[31] M. Hou, "Estimation of sinusoidal frequencies and amplitudes using adaptive identifier and observer", *IEEE Transactions on Automatic Control*, vol. 52, no. 3, pp. 493–499, 2007 (Cited on page 11).

[32] H. C. So, Y. T. Chan, Q. Ma and P. C. Ching, "Comparison of various periodograms for sinusoid detection and frequency estimation", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 35, no. 3, pp. 945–952, 1999 (Cited on page 11).

[33] B. Messner, R. Hill and J. D. Taylor. (1st May 2018). Aircraft pitch: System modeling, [Online]. Available: `http://ctms.engin.umich.edu/CTMS/index.php?example=AircraftPitch&section=SystemModeling` (Cited on page 12).

[34] I. K. Peddle and J. A. A. Engelbrecht, "Advance Automation 833: Introductory Course to Aircraft Dynamics", 2017 (Cited on page 12).

[35] J. Cieslak, D. Efimov, A. Zolghadri, D. Henry and P. Goupil, "Design of a non-homogeneous differentiator for actuator oscillatory failure case reconstruction in noisy environment", *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, vol. 229, no. 3, pp. 266–275, 2015 (Cited on pages 17, 28).

[36] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing*, 4th ed. Pearson, 2014 (Cited on pages 43, 51, 54, 68).

[37] R. G. Lyons, *Understanding Digital Signal Processing*, 3rd ed. 2010 (Cited on page 53).

[38] K. Duda, "Accurate, Guaranteed Stable, Sliding Discrete Fourier Transform", *IEEE Signal Processing Magazine*, vol. 27, no. 6, pp. 124–127, 2010 (Cited on page 55).

[39] S. C. Douglas and J. K. Soh, "A numerically-stable sliding-window estimator and its application to adaptive filters", vol. 1, 111–115 vol.1, 1997 (Cited on page 55).

[40] D. Abramovitch, "Phase-locked loops: A control centric tutorial", in *Proceedings of the 2002 American Control Conference*, vol. 1, 2002, pp. 1–15 (Cited on pages 60, 64, 67, 82).

[41] F. M. Gardner, *Phaselock Techniques*, 2nd ed. New York: Wiley, 1979 (Cited on pages 61, 63, 65).

[42] P. Lutus. (2011). Understanding phase-locked loops, An in-depth look at a fascinating technique, [Online]. Available: `https://arachnoid.com/phase_locked_loop/` (visited on 15/08/2018) (Cited on page 63).

[43] G. A. Leonov, N. V. Kuznetsov, M. V. Yuldashev and R. V. Yuldashev, "Hold-In, Pull-In, and Lock-In Ranges of PLL Circuits: Rigorous Mathematical Definitions and Limitations of Classical Theory", *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 62, no. 10, pp. 2454–2464, 2015 (Cited on pages 63, 65, 74).

[44] D. Y. Abramovitch, "Lyapunov redesign of analog phase-lock loops", *IEEE Transactions on Communications*, vol. 38, no. 12, pp. 2197–2202, 1990 (Cited on page 64).

[45] N. E. Wu, "Analog phaselock loop design using popov criterion", in *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*, vol. 1, 2002, 16–18 vol.1 (Cited on page 64).

[46] G.-C. Hsieh and J. C. Hung, "Phase-locked loop techniques. A survey", *IEEE Transactions on Industrial Electronics*, vol. 43, no. 6, pp. 609–615, 1996 (Cited on page 65).

[47] J. G. van de Groenendaal and R. M. Braun, "Phase-Plane Analysis of Phase-Locked Loops used for Clock Recovery", in *Proceedings of COMSIG '94 - 1994 South African Symposium on Communications and Signal Processing*, 1994, pp. 99–104 (Cited on page 73).

[48] P. Z. Peebles and B. Shi, *Probability, Random Variables, and Random Signal Principles*. McGraw-Hill Education, 2015 (Cited on pages 77, 78).

[49] M. Bhardwaj, "Software Phase Locked Loop Design Using C2000 Microcontrollers for Single Phase Grid Connected Inverter", Texas Instruments, 2013 (Cited on page 82).

[50] E. Weisstein. (2018). Complex multiplication, MathWorld–A Wolfram Web Resource, [Online]. Available: `http://mathworld.wolfram.com/ComplexMultiplication.html` (Cited on page 107).