Department of Precision and Microsystems Engineering

Equivalence in Experimental Source Characterisation

Tom Spoel

Report no : 2019.044
Coach : Dr. M. V. van der Seijs
Professor : Prof. dr. P. G. Steeneken
Specialisation : Engineering Dynamics
Type of report : MSc. Thesis
Date : 16-12-2019

Equivalence in Experimental Source Characterisation

by

Tom Spoel

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday December 16, 2019 at 01:00 PM.

Student number:	4173945	
Project duration:	October 1, 2018 – December 1, 2019	
Thesis committee:	Prof. dr. P. G. Steeneken,	TU Delft, chairman
	Dr. M. Seyyed Fakhrabadi ,	TU Delft
	Dr. A. M. Aragon,	TU Delft
	Dr. M. V. van der Seijs,	VIBES.technology, supervisor

Abstract

Transfer Path Analysis methods allow to describe the propagation of vibrations from an active source to a passive receiving structure. These methods are based in the experimental modelling domain, where the dynamics of the problem are described based on empirical observations from an experiment, as opposed to a numerical model based on natural and engineering laws. These methods describe the boundary conditions for the experiment, as well as how the obtained data can be used to deduce an equivalent force that describes the vibrations. Within the class of Transfer Path Analysis methods there is a specific class called the Component-based methods that are able to describe the vibrating source in such a way, that the assembled dynamical behaviour of the source attached to any other receiving structure can be predicted. This is by defining the characterisation in terms of an equivalent force on the interface. These methods are said to be a function of the source only. This is called source characterisation.

In this thesis some of the assumptions of these Component-based Transfer Path Analysis are challenged. This will be done based on theoretical research as well as experiments on a numerical model. The reason this research is of interest, is that while these Component-based methods show to be valid in theory, their results in practice are not satisfactory. This is especially apparent when a characterised source is virtually assembled on a different receiving structure. Its predicted behaviour is not in line with what the real world counterpart.

These Component-based Transfer Path Analysis methods involve solving an inverse problem to determine the equivalent force. Due to the experimental nature in which the data are obtained, the data has error-contamination. This happens because of interference effects causing noise on the analog sensors. This noise can have a profound effect on the solution of the inverse problem, since the problem is poorly conditioned.

The main assumption of Component-based Transfer Path Analysis methods is that if an interface force can be defined that excites the interface in the same way as the source does, the interface force can be regarded as equivalent. Due to the linearity of the problem, this equivalent force can be added in negative on top of the source excitation. This would theoretically then cancel out all motion, since both the source excitation and equivalent force have the same effect on the receiving structure.

This thesis proofs that it is not always possible to find an equivalent force that represents the source excitation. This is due to the position of the equivalent force, specifically that it is placed on a single point. Some dynamic modes will show a node of their motion on the interface, meaning the interface force can by definition never excite those modes. This limitation is based from the modelling choices of the interface. It will be uncovered that the modes with such nodal behaviour are the eigenfrequencies of the source for a perfectly fixed interface.

The interface limitation can be expressed as a controllability measure, where the controllability defines to what extend the interface force can excite the possible modes crossing the interface from the source. This controllability is a function of the source, the receiving structure and the sensor placement, where the source determines at which frequencies the controllability problems occur and the receiving structure and sensor placement determine the size of the controllability measure. Both an upper and lower bound of controllability can be defined.

It can be quantified what the effect of the incomplete interface description is on the source characterisation. This is based on comparing the motion from source excitation and the motion from the predicted dynamics from the equivalent force. It is shown that the interface motion cannot be reproduced by the equivalent force at the frequencies where there are controllability problems. This mismatch in interface motion has an effect on the motion of the receiving structure, where it will be clear that these mode

occur at range of frequency bins.

While a portion of the source cannot be characterised, a large part of the behaviour can be. To solve for this equivalent force, an inverse problem has to be solved. This problem is an over-determined problem with a full rank column space. Using a priori information about the levels of noise that are expected on each measurement, the problem can be altered. This makes it possible to reduce the propagation of noise from the measurement data to the solution. A set of guidelines are proposed that can help to dampen or emphasis individual modes of the inverse system. Alternatively, the entire spectrum of modes can be dampened or single modes can be truncated. The regularisation methods will all be judged on how they change the physical relevancy of the problem.

Keywords: Component-based transfer path analysis, in-situ, source characterisation, equivalent force, virtual point transformation, interface limitation, controllability, equivalence, inverse regularisation

Personal contributions

1. Interface Description
   - Inherent limitation of equivalent force on interface uncovered
   - Physical explanation of interface limitations
2. Equivalency
   - Specific controllability defined based on interface modelling
   - Overall controllability defined based on interface modelling
   - Measures to express non-equivalency due to controllability problems
3. Inverse regularisation
   - Guidelines for different regularisation techniques on source characterisation inverse problem based on a priori information

Preface

This thesis is written as a final report for obtaining my Master of Science degree in Mechanical Engineering at the Delft University of Technology. In the summer of 2018 I had my first experiences with the domain of experimental dynamics as part of a course given by Dr. Maarten van der Seijs and Dr. Dennis de Klerk. This course spiked my interest due to the combination of audio, dynamics and signal analysis. Combined with a big focus on automotive applications, I knew that this domain would be something I would enjoy working on.

I am proud to have contributed to a new step in understanding in source characterisation and the potential effects the conclusions of this thesis have on other domains within experimental modelling. It has been very difficult to challenge theories that have been used for many years in academics and industry, but I feel a major step in understanding of some inherent problems in source characterisation has been made.

I would like to thank VIBES.technology for giving me the opportunity to be part of the team for a year and experience all the practical things they encounter in their different projects. It has given me great joy to discuss the topics of experimental modelling with colleagues that are regarded expert in the field. A special thanks to Dr. Maarten van der Seijs who was my daily supervisor and gave me the direction to uncover the possibilities in my research. A big thank you to the chairman of this project, Prof. dr. Peter Steeneken. He has given me guidance in how to translate my knowledge on the topic to the final thesis you see here.

I would also like to thank my parents Corine and René who have given me unconditional support during my entire academic career from primary school to university. Without their loving support, advice and help this journey would not have been possible.

Tom Spoel
Delft, December 2019

Contents

General Introduction

We live in a world full of vibrations. From the rumbling of your phone to the seismic activity in the earth's crust causing earthquakes. Many vibrations have a clear function and are pleasant, like for example the vibrations of air causing audible sound to make music. Others vibrations are an unwanted side effect, such as the vibrations of an unbalanced motor.

To better understand the interplay of forces and motion causing vibrations, it is desired to describe vibrations in a mathematical way and form engineering laws to describe them. This makes it possible to make educated decisions to change design parameters to avoid unwanted side effect. As an extra complexity, some components are deemed too complex to make it worthwhile to model accurately. For such components the dynamics can be described by performing experiments with the real-life product and describing its behaviour based on empirical observations. Dynamic models that are constructed in such a manner are part of a domain dubbed experimental modelling.

An active subdomain of experimental modelling is source characterisation. The goal in source characterisation is to describe the vibrations caused by an active vibrating source in such a way, that its behaviour while attached to any other receiving structure can be predicted. This makes it possible to virtually assemble different components and predict the dynamics of the created system, without physically assembling the full product. In today's industry this can be very useful with multiple design teams and various suppliers working on the same product.

A requirement for the source description is that it should be totally equivalent to the original source. In other words, the source characterisation is successful if the calculated prediction is fully equal to its actual behaviour. The source description should be fully independent of the method in which it is obtained, meaning the description has to be a function of the source only. This makes it possible to characterise a source once and use this characterisation on as many different structures as desired, while still obtaining the correct assembled dynamics.

Problem Statement

Multiple methods have been developed to characterise the vibrations caused by an active source in an experimental manner. These methods are made possible by the work of Jedmundsen's in the late 1980's. His proposed Frequency Based Substructuring method made it possible to describe the assembled dynamics from dynamic models of individual substructures [32]. De Klerk extended the method for coupling and decoupling of structures in different domains [34], as well as finding a suitable description of the interface dynamics [17]. This work was used and extended by Van der Seijs [54] to classify the different methods in a single framework.

A number of the methods developed describe the active source in terms of an equivalent force on the interface. In theory, these methods lead to a fully equivalent description of the source that is independent of the test bench or method it was obtained in. In practice however, the method and test bench do play a role in the characterisation. This can be seen in the results of using the equivalent force to predict responses on the test bench structure, but especially when applying this force on a new receiving structure.

Error Sources

The problem in experimental source characterisation must be due to some error source, either in the experiment or in the logic on which the theory is based. Possible error sources in experimental modelling is an active field of study and many recent publications have been about this topic ([39],[26],[36]). Some of these experimental error sources are inherent due to the experimental nature of obtaining

data, while other experimental error sources are due to practical difficulties of performing the experiment. The error sources for source characterisation can be grouped in 3 categories as seen below:

Systematic Experimental Errors: Due to false experiment methodology. This could be a mismatch in positioning of sensors and impacts in the model and the real-life experiment. Another example is a change in interface conditions leading to different dynamics, as described in [9]. Different boundary conditions can also be regarded as systematic errors, for example the effect of a rigid or totally free interface condition on the measured rigid body modes.

Random Experimental Errors: Due to unwanted error-contamination of measurement data. This could be noise contamination on the measurement due to interference effects of the environment. This can occur due to the analog nature of the sensor, the wiring and the data acquisition device. Another possible random experimental error is inconsistent time-delay of acquisition hardware.

Modelling Errors: Due to false assumptions. These errors are not directly related to the experiment, but rather in how the obtained measurement data is used to deduce other quantities. This is a very large scope of possible errors. Errors range from false assumptions on linearity, to oversimplified interface description and truncation and reduction errors.

While there are many publications on the first two groups of error sources, the scientific community seems hesitant to challenge the supportive theories and verify the assumptions involved. This is questionable, since the source characterisation methods also show unexpected results when the experiment is simulated and the first two types of error sources are fully eliminated. This begs the question if there are other discrepancies between theory and practice that lead to this mismatch.

There are many examples of publication of results of test cases ([21],[35],[5]) where the prediction result is not an exact match with the measurement. These mismatches are explained by experimental errors or by specific anomalies in the test case. Janssens and Verheij [30] recognised the fact that the characterisation is limited by the number of Degrees of Freedom (DoF) of the interface force, but also showed that even a 6 DoF interface force is sometimes not sufficient. No further elaboration was given on this topic. This is not a satisfactory answer.

Goal & Outline

It is clear that especially in the third category of error sources there are opportunities for improving the subdomain of source characterisation, which will be the main scope of the research. The source characterisation method which will be analysed is based on a method that is proposed for ISO standardisation [52]; the In-Situ method. The target is to pinpoint the problems in the method, physically explain the limitations and quantify the effect on the characterisation caused by this possible limitation.

A major assumption in the current source characterisation methods is that if the equivalent force can reproduce the interface motion, all the motion behind the interface is also reproduced. The way this is often defined is by making use of the super-position principle of linear systems. By applying the equivalent force in negative to the system in operation, all the motion must cancel since the equivalent force leads to a fully equal response. For this reason the equivalent force is often called the blocked force. This assumption is based on valid reasoning, but has not been verified in any publications. On top of that, it is unknown what the effects on the source characterisation process are if this assumption does not hold true.

Although the origins of the random experimental errors are known, their effect on the outcome of the source characterisation process is not clearly defined. The source characterisation methods involve solving a poorly conditioned inverse problem from noise-contaminated data. Such inverse problems are typical for many imaging applications and have been a popular research field this century. For source characterisation the rank of the problem is far lower however, which makes methods that work well for imaging not necessarily translatable to this field. Research has been published on regulating the inverse problem for source characterisation in the time domain [58] and for the method of Operational

Transfer Path Analysis [11]. While this research was being performed, a German Master thesis [7] has been published regarding regularising the In-Situ method on an industrial case study. As an additional goal to this thesis, current and old state of the art regularisation methods are used on a noisy inverse problem to relate their functioning to how the problem is physically changed. From this a set of tools is defined to regularise different use-cases.

Research Questions

The research questions that this thesis will try to answer are stated as follows:

1. How blocked is the interface when applying the blocked force?
2. How much motion is there behind the interface when applying the blocked force?
3. What effect does the possible interface motion when applying the blocked force have on the acquired equivalent force?
4. Can the inverse problem be treated to improve conditioning and reduce effect of noise on the characterisation results?

Thesis Outline

The thesis is split up in two main parts. The first part is the theory part and consists of the literature study. It is split up in 3 main chapters. In chapter 1: *Dynamic Modelling* the problem will be specified further to determine what kind of dynamic modelling is possible and relevant. Chapter 2: *Transfer Path Analysis* is an extension of the first chapter, where the dynamic framework in which we express the problem is explained. The Source-Tranmission-Receiver model that will be used extensively is introduced here. As the final chapter of the literature study, 3: *Matrix Inversion* functions as a in-depth literature study on inverse problems to provide a mathematical basis.

The second part of the thesis is can be considered as the main research where the theories of the literature study are used. In chapter 4: *Interface Description* source characterisation will be analysed from the standpoint of the interface. The chapter includes the explanation to answer the first research question. The first main conclusion regarding the inherent limitation of an interface force will follow from this.

The main research continues with chapter 5: *Equivalent Force*. Here the insight of the previous chapter will be used to quantify the effect of the interface limitation on the equivalence of the source characterisation. This will lead to the definition of two controllability measures. The source characterisation results are used to quantify both the motion of the interface and the motion of the rest of the receiving structure. New measures are introduced to quantify these properties. This will lead to answers for the second and third research question.

As a final chapter to the main research, chapter 6: *Robust Inverse Problem Solution* dives into the specifics of solving the error-contaminated inverse problem. A mathematical framework is developed to see the effect of different matrix regularisation techniques on the description of the dynamic problem and how it relates to the physical properties of the real world problem. In this way a set of guidelines to choose the right method for the right use-case will be defined. This will answer the fourth and final research question.

I

Literature Study

1 Dynamic Modelling

1.1. Introduction

The first chapter of this thesis will be all about dynamic modelling. We will start by analysing the real-life problem. This will define the scope and the possibilities that our dynamic model should entail. From this point onwards we will dive into the theory behind describing dynamic properties of the system. We will work our way from a dynamic description for the global problem all the way to a reduced dynamic model that describes the physical phenomena we are interested in.

This knowledge will be used to analyse and compare numerical models and experimental models. For the experimental models the scope is limited to the inherent differences compared to numeric models. The exact methodology of how to obtain such models is outside the scope of this thesis as many publications have been written on the topic ([41],[50],[22]). Rather the goal is to describe them in a similar framework so that we can work in a single model with components of both origins. This will set the baseline for a specific form of Dynamic Modelling called Dynamic Substructuring. We will introduce the concept and discuss its strengths and possibilities. Dynamic Substructuring will set the basis for the rest of this thesis.

1.2. Dynamic Problem Introduction

The dynamic problem of this thesis is to describe the vibration of an active component. This description should be totally equivalent to the vibrations of the actual source. Equivalence means that the source description does not have to be equal to the original excitation, but that the effect on the receiving structure should be equal. On top of that, the description should be only a function of the source, and not of the method in which the description is obtained. We seek a description that relates displacements to forces. In other words this means the relation between motion and excitation.

The component we are trying to describe is an active component which has a complicated cross play of forces and moments inside its structure. It is difficult to find a single force that represents this cross play of forces. The components in question are deemed to be too complex for its dynamics to be modelled accurately. There are many components that hold true to these facts. Examples of such components are gearboxes, where a large number of rotating gears makes contact which creates reaction forces through the entire structure, or compressors, where a complex interplay between gasses and liquids creates forces through the entire structure.

The complexity of these components makes it not feasible to model these components. What is done in practice, is that these components are build and then tested in real life. By the means of experiments and empirical observations, the dynamic relations are defined. The nature of experimental modelling will be further explained in section 1.5. We however still need a theoretical baseline to know how to describe these models.

The most important property of all components discussed in this thesis is that the dynamic behaviour of all components can be described as Linear Time-Invariant Systems. This is a specific class of dynamic models which is not limited to describing physical structures, but for the scope of this thesis we will limit ourselves to structural dynamics. Linear Time-Invariant Systems have a number of useful properties. This means that once you verify that your problem is indeed Linear Time-Invariant, the dynamic analysis becomes a lot more straight forward.

The first property is the linearity. What this implies is that for every input that is scaled by a value, the output of the system is scaled by that same value. In other words, if you would double the force on the structure, the displacement is also doubled. From the linearity follows the superposition principle. This principle says that for two inputs that are added together and then used as an input to the system, the output consists of the sum of the individual inputs. The superposition principle is critical in our analysis and it says that the system behaves identically, no matter what the load case currently is.

The second property of Linear Time-Invariant systems is that they are time-invariant. What this implies is that the system does not act different at a different moment in time. Seen from the other side, this means that a time shift in the input signal leads to an output that has the same shift in time. Related to this property is another property of Linear Time-Invariant systems, which is causality. A system is causal if the output of the system at a time t is completely defined by an input to the system at a time t. This implies that load cases in the past will not have an effect on the system.

The final useful property for Linear Time-Invariant Systems is their stability. For every bounded input, the output is also bounded. This says something about the energy flow in the system and that energy cannot be created from nothing. Additionally, we will see that this implies that the impulse response describing the structure therefore also has a finite norm.

Typical Linear Time-Invariant Systems are lightly-damped structures. Examples are metal structures as for example a car chassis or a metal test bench. Structures which are not Linear Time-Invariant but still often used in experimental modelling, are for example rubbery structures. Rubber has a very clear non-causal behaviour. Although we mentioned that this thesis will focus on structural dynamics, it must be noted that for airborne sound transmission, the dynamics are also Linear Time-Invariant Systems. Therefore airborne sound is also often used in the characterisation of an active source.

1.3. Domains

The interaction between forces and motion can be described in multiple ways. These methods can be used interchangeably and are equivalent. In other words, you can rewrite each method to obtain the other method. These methods can be classified in different domains. In the following section we will analyse 3 of these domains that are of use in dynamic modelling.

Some domains offer the possibility to do a reduction on the system. Once the system is reduced it loses its equivalency to other methods. We will analyse the effects of the reductions for each domain where applicable.

Physical Domain

The physical domain describes the dynamical relation between forces and motion based on the equilibrium between external loads and internal inertia, viscous damping and elasticity loads. The internal inertia, viscous damping and elasticity loads are formulated in terms of system properties and the systems current acceleration, speed and displacement. This explicit description of the system properties make it the physical domain.

The model formulation for the physical domain can be seen in equation 1.1. The formulation is one the most elementary formulas in structural dynamics. The displacements and its two time derivatives are denoted by $\boldsymbol{u}(t)$. The system matrices are represented by $\boldsymbol{M}$, $\boldsymbol{C}$ and $\boldsymbol{K}$ and respectively represent the

linearised mass, damping and stiffness matrix.

$$\boldsymbol{M}\ddot{\boldsymbol{u}}(t) + \boldsymbol{C}\dot{\boldsymbol{u}}(t) + \boldsymbol{K}\boldsymbol{u}(t) = \boldsymbol{f}(t) \tag{1.1}$$

The size of the model is determined by the discretisation of the physical problem. The discretisation determines the physical size of each element that is described by its own mass, damping and stiffness matrix entries respectively, but also by the Degrees of Freedom chosen for each element. Real life 3 dimensional structures can be described with linear motion of each element in the (x, y, z)-direction, but alternatively can also be described by 6 Degree of Freedom motion where the rotations about each axes is also taken into account. The displacement vector $\boldsymbol{u}(t)$ and its derivatives have the according size. The system matrices are obtained using the Finite Element Method.

The physical domain as described in equation 1.1 is a second order differential equation. It describes the system in the most elementary method. This fact makes the method very accurate to describe the problem, but the size of the set of second order differential equations becomes so big, that the equation is deemed impractical for dynamic simulations. Reducing the size of the system matrices by changing the element types or increase the element size is only possible up to a certain extend if the physical description is to be kept.

Frequency Domain

The frequency domain describes the dynamical relation between forces and motion based on the harmonic relation between the two. Whereas the structural domain described force and motion signals as a function of time, in the frequency domain these signals are represented as function that shows their contribution and phase per frequency.

For Linear Time-Invariant systems the physical domain representation can be rewritten to the frequency domain by applying a Fourier transform [13]. This transformation can be seen in equation 1.4, where the relation that $\dot{\boldsymbol{u}}(\omega) = j\omega\boldsymbol{u}(\omega)$ and $\ddot{\boldsymbol{u}}(\omega) = -\omega^2\boldsymbol{u}(\omega)$ are used.

$$\begin{aligned} -\omega^2\boldsymbol{M}\boldsymbol{u}(\omega) + j\omega\boldsymbol{C}\boldsymbol{u}(\omega) + \boldsymbol{K}\boldsymbol{u}(\omega) &= \boldsymbol{f}(\omega) \\ \left[-\omega^2\boldsymbol{M} + j\omega\boldsymbol{C} + \boldsymbol{K}\right]\boldsymbol{u}(\omega) &= \boldsymbol{f}(\omega) \\ \boldsymbol{Z}(\omega)\boldsymbol{u}(\omega) &= \boldsymbol{f}(\omega) \end{aligned} \tag{1.2}$$

The size of the model remains unchanged and is still dependent on the element types and the mesh size of the linearised system matrices. The reason one would like to do this transformation, is sinusoidal fidelity [45]. This is a property of all Linear Time-Invariant dynamical systems, that states that a a sinusoidal input is guaranteed to have a sinusoidal output. The only thing the dynamic system can change, is the amplitude and the phase. This makes it very useful to analyse sinusoidal signals.

In the last line of equation 1.4 it can be seen that the system matrix terms are all collected in a single frequency-dependent matrix $\boldsymbol{Z}(\omega)$. This matrix is called the dynamic stiffness matrix. It describes the frequency-dependent relation between forces and motion in a single operator of a single Degree of Freedom for a single element, while all other Degrees of Freedom are constrained.

The dynamic stiffness matrix representation is not very intuitive. On top of that, we will see later that the experimental world is much more of an inverse relation, where we can describe measured motion from a force input, but not the other way around. We therefore introduce the admittance notation matrix $\boldsymbol{Y}(\omega)$ that describes the motion response for a given harmonic force. Equation 1.3 shows the definition of the receptance matrix.

$$\boldsymbol{u}(\omega) = \boldsymbol{Y}(\omega)\boldsymbol{f}(\omega) \qquad \boldsymbol{Y}(\omega) \equiv (\boldsymbol{Z}(\omega))^{-1} \tag{1.3}$$

The impedance $\boldsymbol{Z}(\omega)$ and admittance $\boldsymbol{Y}(\omega)$ notation will be extensively used in the thesis. Here we defined them both as the relation between displacement and force or vice versa, but they can easily be rewritten in terms of velocity and force or acceleration and force. Since both quantities are defined in the frequency domain, each time derivative is only a factor $(j\omega)^{n_t}$ away, where n_t is the order of the time derivative.

Modal Domain

The modal domain describes the dynamical relation between forces and motion based on the natural vibrations modes. It is an extension of the frequency domain representation. The natural vibration modes, otherwise called eigenmodes, are found at the frequencies where the system oscillates in the absence of any external load or internal damping load. At these frequencies the inertia forces and elastic forces are in equilibrium.

$$\left(\boldsymbol{K} - \omega_r^2 \boldsymbol{M}\right) \boldsymbol{\varphi}_i = \boldsymbol{0} \tag{1.4}$$

The vibration shapes and their corresponding frequency are found by an iterative solver, often dubbed an eigensolver. This leads to a representation as seen in equation 1.4. The size of the problem is still the same and determined by the element Degrees of Freedom and the mesh size. There are as many eigenmodes as there are Degrees of Freedom in the entire system.

These solutions can be used to rewrite the dynamical system in terms of its vibration modes. Equation 1.5 shows the definition of the mode shape matrix $\boldsymbol{\Phi}$ that holds all solutions in order of increasing eigenfrequency. This relation shows the strength of the modal domain, as each mode shape is orthogonal to each other. This creates unique building blocks for the dynamic response.

$$\begin{aligned} &\boldsymbol{\Phi}^T \boldsymbol{M} \boldsymbol{\Phi} = \boldsymbol{I} && \text{with } \boldsymbol{\Phi} = \begin{bmatrix} \boldsymbol{\varphi}_1 & \boldsymbol{\varphi}_2 & \dots & \boldsymbol{\varphi}_n \end{bmatrix} \\ &\boldsymbol{\Phi}^T \boldsymbol{K} \boldsymbol{\Phi} = \text{diag}\left(\omega_1^2, \dots, \omega_n^2\right) \end{aligned} \tag{1.5}$$

Modal Domain Reduction

In the modal domain we can do a reduction of the system size. Since we computed the mode shape matrix $\boldsymbol{\Phi}$ in order of increasing frequency, we can truncate the higher frequent modes. These higher frequent modes will have little to no effect al lower frequencies. If we are only interested in the dynamics up to a certain frequency, we can reduce the size of the problem by a huge amount. Depending on the complexity and size of a structure, their system matrices may easily be in the order size of tens of thousands, whereas the first hundred modes of a system are often sufficient to describe the behaviour up to very high frequencies already.

The way the reduction is implemented is by rewriting the global coordinates $\boldsymbol{u}(t)$ to a smaller subspace that only consists of the first m natural vibration modes. This is done by mapping a reduction matrix $\boldsymbol{R}$ on the so called modal coordinates $\boldsymbol{\eta}(t)$. Equation 1.6 shows the relation between the two, where the reduction matrix $\boldsymbol{R}$ consists of the first m modes of $\boldsymbol{\Phi}$.

$$\boldsymbol{u}(t) \approx \boldsymbol{R}\boldsymbol{\eta}(t) \qquad \text{with } \boldsymbol{\eta} \in \mathbb{R}^m \tag{1.6}$$

With rewriting the global coordinates in terms of the modal coordinates, we can plug this relation in our traditional second order differential equation. This can be seen in equation 1.7. Note that this relation is now an approximation since we truncated the high frequent behaviour. If an excitation has components of these high frequent mode shapes, there will be a residual force, represented in the $\boldsymbol{r}(t)$ residual force vector.

$$\boldsymbol{M}\boldsymbol{R}\ddot{\boldsymbol{\eta}}(t) + \boldsymbol{C}\boldsymbol{R}\dot{\boldsymbol{\eta}}(t) + \boldsymbol{K}\boldsymbol{R}\boldsymbol{\eta}(t) = \boldsymbol{f}(t) + \boldsymbol{r}(t) \approx \boldsymbol{f}(t) \tag{1.7}$$

We can rewrite equation 1.7 to get rid of the residual force vector. This is done by pre-multiplying the entire equation by $\boldsymbol{R}^T$. This will lead to $\boldsymbol{R}^T\boldsymbol{r} = \boldsymbol{0}$. Equation 1.8 shows the reduced dynamical system that now has dimensions m.

$$\boldsymbol{M}_m\ddot{\boldsymbol{\eta}}(t) + \boldsymbol{C}_m\dot{\boldsymbol{\eta}}(t) + \boldsymbol{K}_m\boldsymbol{\eta}(t) = \boldsymbol{f}_m(t) \qquad \text{with } = \begin{cases} \boldsymbol{M}_m &= \boldsymbol{R}^T\boldsymbol{M}\boldsymbol{R} \\ \boldsymbol{C}_m &= \boldsymbol{R}^T\boldsymbol{C}\boldsymbol{R} \\ \boldsymbol{K}_m &= \boldsymbol{R}^T\boldsymbol{K}\boldsymbol{R} \\ \boldsymbol{f}_m &= \boldsymbol{R}^T\boldsymbol{f} \end{cases} \tag{1.8}$$

The final modal domain representation of the dynamic model as seen in equation 1.8 is a very powerful reduction. By projecting the displacement vectors on the vector space consisting of only the free vibration modes we are interested in, we have truncated the problem to the dynamic information we need. There are alternatives to the modal domain that differ in the projection step. Rather than projecting on the free vibration modes, a distinction is made between the internal and external element of the system matrix. Reduction steps like Craig-Bampton or Ruben make use of this, but are left outside of the scope of this thesis.

1.4. Numerical Model

The choice of domain for a numerical model of a dynamic system can be made by identifying the problem type. For Linear Time-Invariant systems there are two clear favourites. The final form of the frequency domain as seen in equation 1.3 is a single line matrix representation of the entire dynamical system. This form can be used to solve the inverse problem too, which will be useful when we start characterising the system. Computational, this form performs very poor because there has been no reduction to the numeric problem. To create a Frequency Response Function in the admittance notation it would mean the full $[n \times n]$ Dynamic Stiffness matrix would have to be inverted for all k frequency bins.

The other favourite it the modal domain, where the reduction step is very powerful and can truncate the higher frequent information of the problem without sacrificing precision at low frequencies. The final form of the modal domain as seen in equation 1.8 is in the form of a second order differential equation however, which also performs poorly computational.

1.4.1. Frequency Response Functions

To combine the best of both worlds, the reduced modal domain can be used to compute Frequency Response Functions in the admittance notation with the method of Mode Synthesis. This leads to a set of equations in the form of equation 1.3 like was desired, while still being able to reduce the size of the system.

Equation 1.9 shows the way to compute the Frequency Response Function from the modal domain. It consists of two summations, ones consisting of the the free body modes and the other consisting of the free vibration modes. The first m_r modes are the rigid body modes and these modes occur at $\omega = 0$ and have no damping ratio assigned to them. The resonance frequencies (eigenfrequencies) of the system are denoted by ω_r here.

$$\boldsymbol{Y}(\omega) \approx -\frac{1}{\omega^2} \sum_{r=1}^{m_r} \boldsymbol{\varphi}_r \boldsymbol{\varphi}_r^T + \sum_{r=m_r+1}^{m} \frac{\boldsymbol{\varphi}_r \boldsymbol{\varphi}_r^T}{-\omega^2 + 2j\omega\zeta_r\omega_r + \omega_r^2} \tag{1.9}$$

The strength of equation 1.9 is that it creates a linear mapping matrix from forces to displacements. Like we have seen earlier, it can easily be rewritten in terms of velocities and accelerations if desired. Another strong point of this method, is that the number of Degrees of Freedom can be truncated without losing any precision. In this manner, only the Degrees of Freedom that are of interest are kept. This leads to a small matrix with exactly the information desired.

1.5. Experimental Model

The experimental model is obtained in a total different fashion than the numerical model. Whereas all the numeric domains originate from a discretised physical description of the problem, the experimental domain originates from a continuous real-life scenario. The goal is to obtain a dynamic model in a similar description form as the numeric model, meaning we are seeking Frequency Response Functions.

The largest difference between the numeric and experimental domain is the practical limitations of the latter. The experimental domain relies on empirical observations. For these observations sensors are needed, either to measure force or some kind of motion (displacement, velocity or acceleration). These sensors often need to be attached to the structure and their data has to be acquired. In this data acquisition process the analog data must be filtered, digitalised at a very high sampling frequency while remaining perfectly synchronised to all other sensors. Specialised hardware is needed is for this.

Besides the difference in the method of obtaining numerical or experimental models, their description is also inherently different. To describe the dynamics between two points on a structure in the numeric world, the entire structure needs to be accounted for since every added mass, stiffness and damping changes the global dynamics. In experimental models all the dynamics are described by what is measured on the sensors. It shows more resemblance to black boxes with known in and outputs behaviour,

which means all internal dynamics are captured in these external measurements.

The fact that all internal dynamics are captured from external measurements is exactly what is desired. One of the reasons the experimental domain is used is that there is no accurate model of the internal dynamics. An added advantage of the experimental domain is that it is more practical to gather admittance descriptions of the problem, whereas for numerical models the description always originates from an impedance description. In other words, for experimental models the motion for a given force input is described. The reason for this is that it is more practical to measure motion for a given force input using an impact hammer and accelerometers than the other way around. This is due to the practical difficulties of measuring forces but especially moments. This can be seen as an advantage, since the admittance notation is more intuitive as was explained earlier.

1.6. Dynamic Substructuring

Dynamic Substructuring is a way to mathematically couple substructures to one another. Equation 1.10 shows the mathematical form for an easy example. The dynamics of each substructure are expressed in an admittance notation. The equation is constructed with the substructure system matrices placed in a block diagonal manner. The Degrees of Freedom of each substructure are explicitly defined as a motion output and a force input. The force is split up in an external forces and reaction forces. Note that all these properties are frequency dependant but for clarity the (ω) term has been omitted.

$$\begin{bmatrix} \boldsymbol{u}_1 \\ \boldsymbol{u}_2^A \\ \boldsymbol{u}_2^B \\ \boldsymbol{u}_3 \end{bmatrix} = \begin{bmatrix} \boldsymbol{Y}_{11}^A & \boldsymbol{Y}_{12}^A & \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{Y}_{21}^A & \boldsymbol{Y}_{22}^A & \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{Y}_{22}^B & \boldsymbol{Y}_{23}^B \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{Y}_{32}^B & \boldsymbol{Y}_{33}^B \end{bmatrix} \left(\begin{bmatrix} \boldsymbol{f}_1 \\ \boldsymbol{0} \\ \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix} + \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{g}_2^A \\ \boldsymbol{g}_2^B \\ \boldsymbol{0} \end{bmatrix} \right) \tag{1.10}$$

The system of equations can be rewritten to cancel out the reaction force and get a system matrix consisting of the coupled admittance $\boldsymbol{Y}_{ij}^{AB}$. This was originally proposed by Jetmundsen [32]. De Klerk suggested a very elegant one-line equation to couple the structures [16] as can be seen in equation 1.11. The exact derivation of the coupling process is outside the scope of the thesis, but the process is closely related to source characterisation as will become clear in the next chapter.

$$\boldsymbol{u} = \boldsymbol{Y}\boldsymbol{f} - \boldsymbol{Y}\boldsymbol{B}^T \left(\boldsymbol{B}\boldsymbol{Y}\boldsymbol{B}^T \right)^{-1} \boldsymbol{B}\boldsymbol{Y}\boldsymbol{f} \tag{1.11}$$

There are two important notions from the coupling process. The first one is that two interface conditions have to be met for the coupling process. These are the coordinate compatibility and the force equilibrium. The coordinate compatibility means that the coupled Degrees of Freedom need to have the same value and sign on both sides of the interface. Similarly, the force equilibrium means that the reaction forces are equal but opposite. This is in line with Newton's third law.

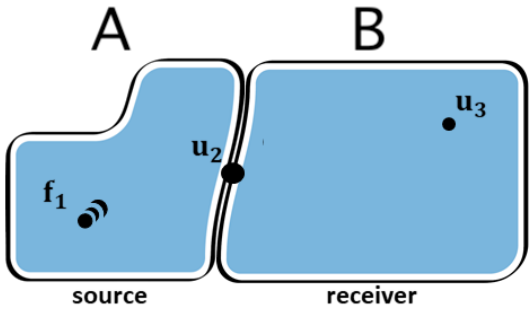

Figure 1.1: Dynamic Substructuring Schematic

The other important conclusion from coupling substructures in this manner is that the coupling cannot be done accurately when coupling only with translational Degrees of Freedom. It has been shown by Drozg that for realistic coupling the rotations and corresponding moments must be accounted for [18]. This is an important fact as this translates directly to source characterisation using an interface force. This force must also be 6 Degrees of Freedom for the result to relate to the real-world situation.

2

Transfer Path Analysis

2.1. Introduction

Chapter 1 gave insight in how to model dynamic behaviour of an active vibrating source. We have seen there are ways to model the dynamics numerically and experimentally and there are ways to use these models in conjunction. In this chapter that knowledge will be extended to the actual problem at hand, namely experimental source characterisation. Whereas for the insights to dynamical modelling was very theoretical, in this chapter some more practical implications of the experimental domain are discussed and analysed. The goal is to lay out the source characterisation problem in the framework of a *Source-Transmission-Receiver* trinity. This framework will be used throughout the thesis, where the source describes the active component, the transmission consists of the interface between the active component and what it is connected to, and the receiver is the receiving structure.

First a more accurate description of the interface will be developed. This is in line with the transmission part. This must be done to find a common interface description for both sides of the interface so coupling as seen in section 1.6 is possible. This description will then be used to dive into Transfer Path Analysis, a method of describing vibrations in terms of the transmission paths. This method could be understood as an extension of the Dynamic Substructuring representation. All relevant, possible methods to describe the active source in terms of another measure will be discussed here within the framework of Transfer Path Analysis.

2.2. Interface Description

There are many types of interface connections. Some interfaces which can be simplified to a discretisation in a single point. These interfaces have a very rigid structure around the interface. Since the structure around the interface moves little, the coupled vibrations will have little effect of the interface compliance. These are connection types like a bolt connection. Other interface connections are interfaces that show more resemblance to a line or a plane. These interfaces cannot be discretised to a single point. The coupled dynamics will be influenced by the compliance in the interface. Examples of these interfaces are adhesively connected panels.

If we limit the scope to interface connections of the first type, theoretically it must be possible to describe the interface in terms of a single discretised points. For this point we must be able to measure 6 Degree of Freedom motion and rotation as well as forces and moments. If such a description is found for both sides of the interface it must be possible to couple experimentally obtained models.

For such an interface description a sensor would be needed that measures translations and rotations. On top of that, it must be possible to apply forces and moments on the sensor position to get the co-located responses. Practically this is not feasible. Rotational sensors have been developed as shown by Jianxin, Su and Mak [33], but these have not been embraced by industry yet, partly due to their lower bandwidth as shown in a case study by Drozg [18]. The biggest problem however is to excite the

moments around the sensor. Multiple methods have been suggested in recent publications ([20],[10]), but these involve very error-prone steps.

2.2.1. Discrete Interface Modelling

We need to develop a method to describe the interface in a discrete manner that represents the real-life connection. We will limit the scope to interfaces of the first kind, which were the interfaces that can be discretised in a single point. To verify this assumption, we will analyse a bolt connection and discuss different options to discretise it.

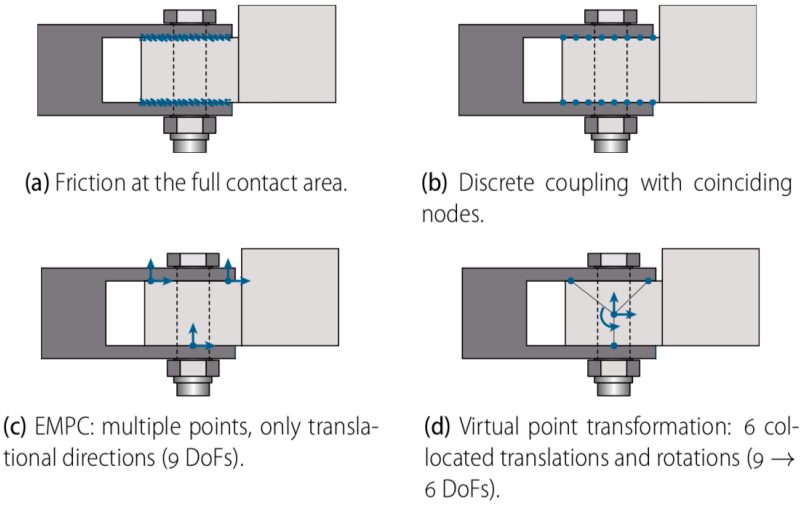

(a) Friction at the full contact area.

(b) Discrete coupling with coinciding nodes.

(c) EMPC: multiple points, only translational directions (9 DoFs).

(d) Virtual point transformation: 6 collocated translations and rotations (9 → 6 DoFs).

Figure 2.1: Interface Modelling

Figure 2.1 shows 4 schematics of a bolt connection. The first picture shows a representation of an actual, real-life bolt connection. Due to the pre-tension of the bolt and nut a large in-plane friction force is created on top of the out-of-plane force due to the pre-tension. The interface stiffness and relative motion is determined by these forces. This goes for in and out-of-plane motion, but also for rotations.

The second picture in figure 2.1 shows a Finite Element Method representation of this same coupling. By choosing the mesh grid equal on both sides of the interface, the motion equilibrium can be obtained. Because the coupling is done with multiple nodes, the motion is constrained in and out-of-plane and rotations are also accounted for. Discrete coupling with coinciding nodes is possible in the numeric domain, but this does not seem to have a dual form in the experimental domain. There are two alternative discrete interface representations that do translate to the experimental domain.

The first experimental discrete interface description is suggested by De Klerk [14] and is called the equivalent multi-point method (EMPC). As the name suggests, it makes use of multiple points to connect both sides of the interface. In practice each side of the interface is measured using a minimum of 3 tri-axial accelerometers on the same spot on the interface. These responses are then set to be equal on both side. Since a minimum of 3 tri-axial accelerometers are used that are preferably not placed in a single line, the rotations are implicitly accounted for. This method has since then been improved and is not used in practice anymore. This is due to the stiffening effect created by forcing 9 Degrees of Freedom to be coupled. This stiffening can create artefacts and spurious peaks in the Frequency Response Function.

The improvement suggested to the EMPC method is called the Virtual Point method and originates from Van der Seijs [53]. The Virtual Point method can be schematically seen in the last picture in figure 2.1. The Virtual Point is the dual version of a Rigid Body Element of the third type (RBE3) for the experimental domain. This can be understood as a post-processing step to the results where all responses are mapped to a single point, in a way taking the average. No extra stiffness or dynamics are

introduced. Similar to the EMPC method it makes use of at least 3 tri-axial accelerometers. It assumes rigid behaviour between the sensors. With this assumptions, the sensor space can be projected onto a 6 Degree of Freedom Virtual Point subspace, consisting of 3 translations and 3 rotations. The notion of a Virtual Point will be further developed in the following section.

2.2.2. Virtual Point

In the following section we will work towards a description of a discretised interface in a single point following the Virtual Point Method. This discretisation is specifically for the experimental domain. Since we are rewriting the measured properties in a different vector space, we can call this method a transformation. In the following section it will be clear that the mapping is to a vector space of a smaller dimension, so the method is also a reduction step.

The method can be seen as a combination two separate transformation. On one side there is a transformation of the dynamic motion behaviour from the sensor measurements to a virtual point motion. This transformation is fully in line with RBE3 transformations in the numerical domain. On the other side there is a transformation of interface forces to the virtual point forces. This reduction is less intuitive at first sight. Both these transformation are independent from each other.

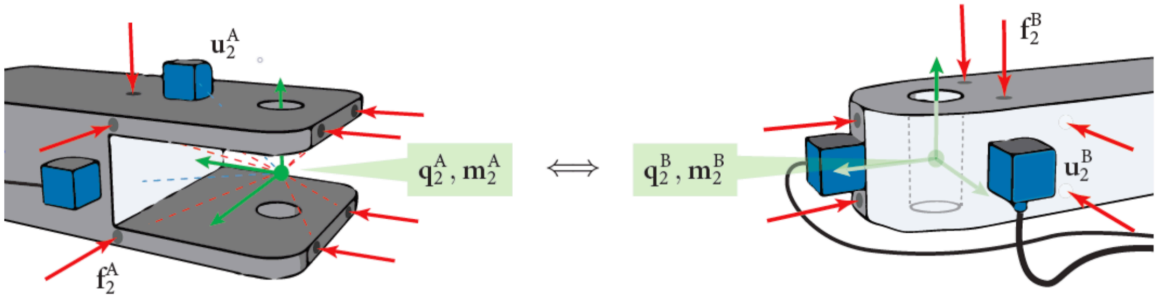

Figure 2.2: Virtual Point Assembly

Interface displacement reduction

The first transformation is to transform the measurements gathered from the sensors to a Virtual Point. For this transformation we have to define the Degrees of Freedom for the Virtual Point. From section 2.2.1 we learned that we wish to take account for three translations and three rotations in the coupling process. For this reason the Virtual Point is chosen to be of 6 Degree of Freedom.

Equation 2.1 shows the relation between the Virtual Point displacements $\boldsymbol{q}$ and the sensors measurements $\boldsymbol{u}$, where $\boldsymbol{q} \in \mathbb{C}^{n_2}$ with $n_2 = 6$ and $\boldsymbol{u}_4 \in \mathbb{C}^{n_4}$. As we have discussed earlier in section 1.5, the number of measurement channels is larger than the number of Virtual Point Degrees of Freedom. In practice the number of measurement channels is chosen to be $n_4 = 9$.

The matrix $\boldsymbol{R}_u$ represents the transformation matrix. It is a frequency-independent mapping of the n_4 sensor channels to the n_2 Virtual Point channels. Its columns consists of the rigid interface displacement modes. This implies that the Virtual Point channels only describe the motion for a rigid interface. All the motion that cannot be mapped to the rigid interface modes is seen in the residual $\boldsymbol{\mu}$. If needed, the transformation matrix can be extended to include flexible interface modes too, as shown by Pasma [42]. While the mapping matrix is frequency independent, the measurement $\boldsymbol{u}$ and Virtual Point motion $\boldsymbol{q}$ are not. For clarity the (ω) term has been omitted.

$$\boldsymbol{u} = \boldsymbol{R}_u \boldsymbol{q} + \boldsymbol{\mu} \tag{2.1}$$

To practically calculate the Virtual Point Channels, we need the inverse relation of the transformation matrix $\boldsymbol{R}_u$ to get a reduction of coordinates. This can be computed by taking the pseudo-inverse of the matrix, denoted by a superscript $+$, as seen in equation 2.2. The exact implications of this is outside the

scope of this thesis, but pseudo-inverses will be discussed in chapter 3. If desired, the inverse solution of the transformation matrix can be extended using a weighting matrix to have certain measurements participate more in the calculation.

$$\boldsymbol{q} = (\boldsymbol{R}_u)^+ \boldsymbol{u} \tag{2.2}$$

Visually this process can be seen in figure 2.2, where the sensors are represented by blue squares attached to the surface of the structure and the Virtual Point motion is represented by green arrows. The figure shows two sides of an interface where the same process is performed. This would make it possible to couple them based on the motion equilibrium.

Interface force reduction

For an interface displacement reduction the exact mapping can be determined. In other words, for a given $\boldsymbol{f}$ there will be an exactly defined $\boldsymbol{q}$. For moments this is not as trivial. This is to be expected, since a moment consists of both a moment arm and a force and these two are not uniquely defined. For this reason the interface force reduction is defined in an inverse manner compared to the displacement reduction. For a given load case $\boldsymbol{f}$, the Virtual Point forces and moments $\boldsymbol{m}$ are defined uniquely.

Another crucial difference between the motion and force reduction, is the fact that the motion reduction was done for sensors that are explicitly placed on the interface to measure the interface. For the force reduction there are no forces on the interface in the problem description at all. These interface forces must be added specifically for the Virtual Point force reduction. In theory a minimum of $n_4 = 6$ interface impacts have to be experimentally tested to reduce them to $n_2 = 6$ Virtual Point forces and moments. In practice, the problem is overdetermined by having at least double the amount of interface impacts. This is done to improve the conditioning pf $\boldsymbol{R}_f$ as will be made clearer in chapter 3.

Equation 2.3 shows the relation between the interface impacts and the Virtual Point forces and moments. Here $\boldsymbol{R}_f$ is the transformation matrix for the force reduction. It is constructed similarly as for the motion reduction and consists of the rigid interface displacement modes. This implies the interface impacts must be located in the circumference of the interface where the structure behaves rigidly. If all interface impacts are chosen on the faces of each sensor, the transformation matrices would be identical. Similar to the the motion reduction, the mapping matrix is frequency independent, the excitation $\boldsymbol{f}$ and Virtual Point forces $\boldsymbol{m}$ are not. For clarity the (ω) term has been omitted.

$$\boldsymbol{m} = \boldsymbol{R}_f^T \boldsymbol{f} \tag{2.3}$$

Figure 2.2 shows the interface force reduction visually. The red arrows represent the position and direction of hammer impacts. These function as a Dirac impulse to excite a broadband of frequencies. These forces are mapped to forces and moments in the interface, depicted by green arrows.

2.2.3. Using Virtual Point in Practice

The Virtual Point method can be used to transform interface motion and forces to a reduced set of Virtual Point Channels and Loads. The method can also be used to do this transformation on the level of the Frequency Response Function Matrices. The usual use-case for this is to couple two substructures on the Frequency Response Function level.

The transformation can be done for both the motion and the forces to get a colocated driving point Frequency Response Function of the interface as seen in equation 2.4. In practice, the reduction matrices are extended to carry along other measurements and loads in the transformation. These extra entries do not participate in the interface reduction. Here they are depicted with the subscript q and m, but this is also often just defined with the subscript 2.

$$\boldsymbol{q} = \boldsymbol{Y}_{qm}\boldsymbol{m} = \left(\boldsymbol{R}_u\right)^+ \boldsymbol{Y}_{u,f} \left(\boldsymbol{R}^T{}_f\right)^+ \boldsymbol{m} \tag{2.4}$$

In the next section we will analyse different methods to obtain a equivalent measure to describe a vibrating source. Many of these measures rely on an interface force as equivalent measure. In some cases it is beneficial to only transform the forces of the Frequency Response Function so that the

measurement channels stays the actual empirical observation without any reductions. Such a one-sided Virtual Point transformation can be seen in equation 2.5.

$$\boldsymbol{u}_4 = \boldsymbol{Y}_{4m}\boldsymbol{m} = \boldsymbol{Y}_{uf}\left(\boldsymbol{R}^T{}_f\right)^+ \boldsymbol{m} \tag{2.5}$$

As a final note on using the Virtual Point Transformation is must be said that there are numerous error sources for this method which have been researched heavily ([15],[3]). Multiple quality indicator measures have been suggested to quantify these error sources [17]. It has also been shown that the result of the transformation is very much dependent on the accuracy of the reduction mapping matrix. Voormeeren showed that for lightly damped structures a slight inaccuracy in this mapping can lead to large inaccuracy of the reduction [56].

2.3. Transfer Path Analysis

Now that we have an accurate description of the interface that can be realised both numerically and experimentally, we can get back to the problem of the thesis; source characterisation. For this we will use the notion of Transfer Path Analysis. Transfer Path Analysis is a set of methods to describe the vibrations from an active source in terms of an equivalent force somewhere else. This is also where the name originates from, as these equivalent forces can be used to analyse how the vibrations split up over different structural paths.

Transfer Path Analysis is the closely related to the experimental nature of source characterisation. It describes the experimental conditions of the active source being tested. We will analyse different methods that have been developed and compare their strengths and weaknesses. Most of these experiments are based on measuring the passive dynamics of the structure to create Frequency Response Functions. After that an operational experiment is done where the source is turned on and the response is measured. The measurement data this provides is used with the earlier obtained functions describing the dynamics of the structure to calculate an equivalent force. All these methods will be described using the framework of source-transmission-receiver. These subscripts will be used throughout the thesis.

2.3.1. Source-Transmission-Receiver Framework

In the Source-Transmission-Receiver framework the structures of the source characterisation problem is split up in 3 distinct parts. This description originates from the acoustic domain. Mondot and Peterson proposed a method in the 1980's to describe the vibration transfer problem in terms of the power coming from the source and some kind of coupling function to describe the transmission [37]. This framework has an explicit mention of the receiving structure in addition.

Figure 2.3 shows a schematic of two coupled substructures. The subscripts used are further elaborated in table 2.1. Although the order of the numbering seems a bit strange at first since 4 comes before 3, it is this way for legacy reasons. Subscript 1 is the location of the source excitation, subscript 2 is the location of the interface and subscript 3 is for the location of the measurement sensors somewhere on the receiving structure away from the interface. Subscript 4 was introduced later and is the locatioin of the sensors used to characterise the interface. For the motion subscript 1 is a reduction from motion subscript 1 as seen in section 2.2.

These definitions make it possible to write the source characterisation problem in terms of a product of a Frequency Response Function and a motion or force. The products are all matrix products since every subscript can have multiple Degrees of Freedom in different directions. In the following section multiple source characterisation methods will be discussed which all originate from the admittance notation $\boldsymbol{u}_i = \boldsymbol{Y}_{ij}\boldsymbol{f}_j$. These vectors are a function of frequency but for clarity the (ω) is omitted.

To make a clear distinction between the source and receiving structure, the structure is denoted in the superscript. The source structure is usually denoted with a superscript A and the receiving structure with B in the case it the receiving structure is part of the product for which A is intended. When the receiving structure is a test bench used purely for the characterisation it is denoted with superscript R

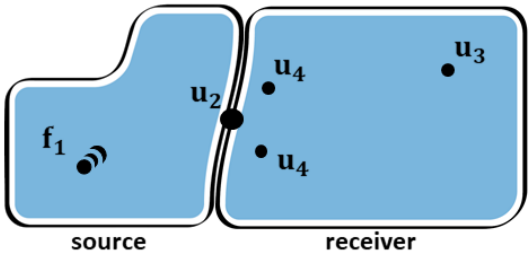

Figure 2.3: Source-Transmission-Receiver model

Degree of Freedom	Subscript	Description	Force / Motion
Internals of source	1	The position of excitation on the source structure. Often unknown.	Force
Interface	2	The position of the interface. Often not measured directly but deduced from other measurements.	Force and Motion
Internals of receiver	3	The position of validation sensors. Placed downstream from interface.	Motion
Indicator around interface	4	The position of indicator sensors. Placed around interface.	Motion

Table 2.1: Degrees of Freedom in Source-Transmission-Receiver Framework

2.3.2. Classical Transfer Path Analysis

There are two main classes of Transfer Path Analysis methods. The first class is the Classical Transfer Path Analysis class. These methods are strong for identifying equivalent forces for existing structures. What this means, is that the equivalent forces found with these methods are equivalent specifically for that structure, but not necessarily for a different receiving structure. The method originates from work of Verheij at the end of last century regarding the transmission of vibrations from ship machinery to the rest of the ship. Verheij managed to experimentally determine the interface forces and moments [55], which paved the way for further academic research.

Direct Force

The first and most obvious method to find an equivalent force is to directly measure the interface force. This is the reaction force due to coupling at the interface. By placing a force transducer between the source and receiving structure, the reaction forces can be measured. The reaction force on the interface would indeed describe the vibration on the passive side, as we have seen in section 1.6 while coupling.

Equation 2.6 shows the mathematical relation of this method. Figure 2.4 shows a schematic representation of the method. The method is valid if there is indeed no interface motion difference between the source and the receiver. In other words, it is valid if $\boldsymbol{u}_2^A = \boldsymbol{u}_2^B$ holds true.

$$\boldsymbol{g}_2^B = -\boldsymbol{\lambda} \tag{2.6}$$

The Direct Force method seems powerful through its simplicity. In theory this is the case, but in practice there are some real-life limitations that make this method not applicable for most problems. First of all you would need a force transducer that measures both 3 Degree of Freedom forces and 3 Degree of Freedom moments. As if that would be not difficult enough, you would also need the transducer to be approximately infinitely stiff in translation and rotation for the boundary displacement condition to hold and to not introduce new dynamics. On top of that, the transducer needs to be of a physical size that makes it practically possible to place between two components. These practical limitations deem it a method which is of little use.

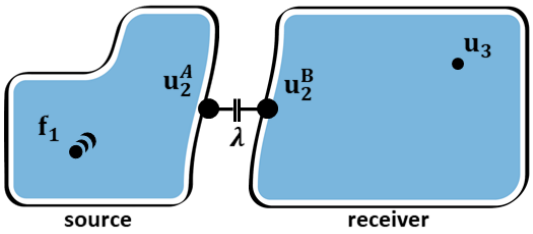

Figure 2.4: Direct Force

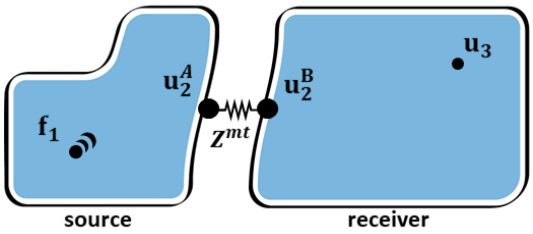

Figure 2.5: Mount Stiffness

Mount Stiffness

The second method of Classical Transfer Path Analysis is a follow-up from the Direct Force method that tries to eliminate one of the practical limitations. Rather than measuring the reaction force due to coupling directly, new and known interface dynamics are introduced. This is where the method gets its name from, as a mount is introduced. These new interface dynamics weaken the interface motion equilibrium. By comparing the difference in interface motion from both components, we can calculate the reaction force by using the new known interface dynamics.

Figure 2.5 shows a schematic representation of the method. Equation 2.7 shows the mathematical relation of this method. For this method to work you need a Dynamic Stiffness description of your interface mount. In contrast to the Direct Force method, the state of the structure is now not measured in terms of the reaction force, but in terms of the motion on both sides of the interface. Determining the motion is a lot easier than determining the force, especially taken into account the Virtual Point Transformation from section 2.2.

$$\boldsymbol{g}_2^B = \boldsymbol{Z}^{mt}\left(\boldsymbol{u}_2^A - \boldsymbol{u}_2^R\right) \tag{2.7}$$

The Mount Stiffness method is a method that takes away some of the impracticalities of the Direct Force method, but introduces some new ones too. First of all, the packaging issue is still present to a certain extend. Although the motion equilibrium condition is weakened, the components still need to practically fit and especially for components that are attached at multiple interface points this becomes troublesome.

It is also not trivial to get a good description of the mount stiffness $\boldsymbol{Z}^{mt}$. In practice these mounts are often rubber mounts that do not hold to the Linear Time-Invariant theory. It is said that for small amplitudes these non-linearities have only little effect and therefore the method is valid [51].

Matrix Inverse

The final method of the Classic Transfer Path Analysis class is the Matrix Inverse method. This method takes advantage of the uncoupled Dynamic Substructuring representation of the Source-Transmission-Receiver model as can be seen in equation 2.8. Using the same logic as during the coupling process, it is reasoned that all motion $\boldsymbol{u}_4$ is due to the reaction force on the receiving structure.

$$\begin{bmatrix}\boldsymbol{u}_1\\ \boldsymbol{u}_2^A\\ \boldsymbol{u}_2^B\\ \boldsymbol{u}_3\\ \boldsymbol{u}_4\end{bmatrix} = \begin{bmatrix}\boldsymbol{Y}_{11}^A & \boldsymbol{Y}_{12}^A & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{0}\\ \boldsymbol{Y}_{21}^A & \boldsymbol{Y}_{22}^A & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{0}\\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{Y}_{22}^B & \boldsymbol{Y}_{23}^B & \boldsymbol{Y}_{24}^B\\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{Y}_{32}^B & \boldsymbol{Y}_{33}^B & \boldsymbol{Y}_{34}^B\\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{Y}_{42}^B & \boldsymbol{Y}_{43}^B & \boldsymbol{Y}_{44}^B\end{bmatrix}\left(\begin{bmatrix}\boldsymbol{f}_1\\ \boldsymbol{0}\\ \boldsymbol{0}\\ \boldsymbol{0}\\ \boldsymbol{0}\end{bmatrix} + \begin{bmatrix}\boldsymbol{0}\\ \boldsymbol{g}_2^A\\ \boldsymbol{g}_2^B\\ \boldsymbol{0}\\ \boldsymbol{0}\end{bmatrix}\right) \tag{2.8}$$

The method works by solving the inverse problem using the Frequency Response Function $\boldsymbol{Y}_{42}^B$ and using the measured response $\boldsymbol{u}_4$. This will lead to the reaction force for this assembly. To determine $\boldsymbol{Y}_{42}^B$, the source and receiving structure must be disassembled. The Frequency Response Functions can then be measured by making use of the Virtual Point Transformation for the forces only. The inverse problem to be solved can be seen in equation 2.9. Note that the mathematical background for this inverse problem will be analysed in chapter 3.

$$\boldsymbol{g}_2^B = \left(\boldsymbol{Y}_{42}^B\right)^+ \boldsymbol{u}_4 \tag{2.9}$$

For this inverse problem to be solvable we need a full rank matrix $\boldsymbol{Y}_{42}^{B}$. The rank of the matrix depends on the sensors placements, a topic which has been heavily researched (for example [59]). In short it depends on observing all possible motion. Similar to the reasoning in section 2.2, it is done by overdetermining the problem to ensure a full rank column space. This makes the Matrix Inverse method the first method of the Classical Transfer Path Analysis methods that is actually feasible to implement in practice. For this reason it is a method that still holds merit amongst the component-based transfer path analysis methods we will discuss next.

2.3.3. Component-based Transfer Path Analysis

The Classical Transfer Path Analysis methods found equivalent forces for the assembled system. This means that when the active source is used in a different assembly, the characterisation would have to be done again. Taking this one step further, this means that this measure is inherently not useful to predict any dynamic behaviour on a different structure. This deems them unusable for source characterisation.

To negate this problem, Component-based Transfer Path Analysis methods are methods to find an equivalent force that is only a function of the source itself. The response for any other assembly with the active source can be predicted. For this prediction you would need the passive dynamics of the new structure with the source turned off. From this we can conclude that it is theoretically possible to characterise a source in ideal testing situations only once, and that this characterisation can be used to predict all possible assemblies.

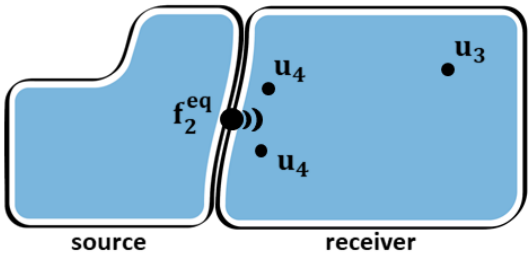

Figure 2.6: Equivalent Source

Equivalant Source

The first type of component-based Transfer Path Analysis methods is the concept of the Equivalent Source. The method uses the reasoning that the response at a sensor is defined by the excitation and the Frequency Response Function between the sensor and the excitation location. An equivalent force somewhere else on the structure should lead to the exact same response for it to be truly equivalent. This concept can be written down as seen in equation 2.10.

$$\boldsymbol{u}_3 = \boldsymbol{Y}_{31}^{AR}\boldsymbol{f}_1 \qquad \Leftrightarrow \qquad \boldsymbol{u}_3 = \boldsymbol{Y}_{32}^{AR}\boldsymbol{f}_2^{eq} \tag{2.10}$$

Equation 2.10 shows the response at $\boldsymbol{u}_3$ due to the operational excitation $\boldsymbol{f}_1$ and the same response due to an equivalent force $\boldsymbol{f}_2^{eq}$. From this definition both excitations should lead to the same response in the absence of the other excitation.

The relation seen in equation 2.10 still seems to be a function of the assembled dynamics. By applying the coupling process of section 1.6 in reverse, the expression can be expanded in terms of its substructure admittance. This can be seen in equation 2.11.

$$\boldsymbol{u}_3 = \left[\boldsymbol{Y}_{32}^{R}\left(\boldsymbol{Y}_{22}^{A}+\boldsymbol{Y}_{22}^{R}\right)^{-1}\boldsymbol{Y}_{21}^{A}\right]\boldsymbol{f}_1 \qquad \Leftrightarrow \qquad \boldsymbol{u}_3 = \left[\boldsymbol{Y}_{32}^{R}\left(\boldsymbol{Y}_{22}^{A}+\boldsymbol{Y}_{22}^{R}\right)^{-1}\boldsymbol{Y}_{22}^{A}\right]\boldsymbol{f}_2^{eq} \tag{2.11}$$

It is seen in equation 2.11 that a large part of both representations is identical. This is the part due to coupling of the structures and indeed, the added interface stiffness due to coupling is identical for

both representations. Eliminating these factors and rewriting the equation in terms of $\boldsymbol{f}_1$ leads to the expression seen in equation 2.12.

$$\boldsymbol{f}_2^{eq} = \left(\boldsymbol{Y}_{22}^{A}\right)^{-1} \boldsymbol{Y}_{21}^{A} \boldsymbol{f}_1 \tag{2.12}$$

Equation 2.12 shows the expression for an equivalent force that is only a function of the source. In the next methods we will encounter ways to simplify this relation with additional measurements.

$$\boldsymbol{u}_2 = \boldsymbol{Y}_{21}^{AR} \boldsymbol{f}_1 - \boldsymbol{Y}_{22}^{AR} \boldsymbol{f}_2^{eq} = \boldsymbol{0} \tag{2.13}$$

The most important notion for the equivalent source method is that the reasoning that when applying the equivalent force in the opposite direction while the source excitation is on, the vibrations are cancelled out due to the superposition principle. This is represented in equation 2.12. Additionally, it is argued that all vibrations downstream of $\boldsymbol{u}_2$ are also cancelled out. Note that the first 3 research questions of this thesis are regarding this assumption.

A final important note on the equivalent source method is that there is a crucial limitation to the equivalent force. This limitation is that the equivalent force describes the motion on the passive side fully equivalent in theory. This is due to the fact that all motion at B is due to a force through the interface. For the motion at A however, the motion is due to the source excitation but also due to a contribution of that force that is reflected through B and the interface back into A.

Blocked Force

The second on the Component-based Transfer Path Analysis methods is the Blocked Force method. For this method you only need the active source and no specific receiving structure. The idea is that the active source is rigidly fixed at its interface and its entire dynamics can be captured due to the reaction forces at the interface.

Figure 2.7 shows a schematic representation of the measurement. By filling in these boundary conditions in the Dynamic Substructuring framework we end up at equation 2.14. By setting $\boldsymbol{u}_2$ to $\boldsymbol{0}$ we can find an expression for the reaction force at the interface as seen in equation 2.15. Note that the derivation in equation 2.15 is just to clarify its origins. In practice the blocked force is measured directly and there is no additional calculations using subsystem Frequency Response Functions.

$$\begin{bmatrix} \boldsymbol{u}_1 \\ \boldsymbol{u}_2^A = \boldsymbol{0} \end{bmatrix} = \begin{bmatrix} \boldsymbol{Y}_{11}^{A} & \boldsymbol{Y}_{12}^{A} \\ \boldsymbol{Y}_{21}^{A} & \boldsymbol{Y}_{22}^{A} \end{bmatrix} \left(\begin{bmatrix} \boldsymbol{f}_1 \\ \boldsymbol{0} \end{bmatrix} + \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{g}_2^A \end{bmatrix} \right) \tag{2.14}$$

$$\boldsymbol{g}_2^A = \left(\boldsymbol{Y}_{22}^{A}\right)^{-1} \boldsymbol{Y}_{21}^{A} \boldsymbol{f}_1 \qquad \rightarrow \qquad \boldsymbol{f}_2^{eq} = \boldsymbol{g}_2^{bl} \tag{2.15}$$

For the blocked force method to work accurately the interface needs to be held still perfectly while also measuring a 6 Degree of Freedom force. This first assumption is very troublesome as even a massive steel block will show some elastic behaviour. Additionally the force transducer will introduce a lot of compliance and dynamics. This makes the blocked force method not very practical and especially for higher frequencies it will be inaccurate due to interface flexibility.

As a final note on the blocked force method is must be said that the name of the method can sometimes be confusing. The force measured is indeed while the interface is blocked. In practice however, the term blocked force is often used to describe the equivalent force. This is due to the equivalent source reasoning we have seen before, where a negative equivalent force on top of the excitation will lead to a blocked interface. For this reason we will always call an equivalent force just that, so that the blocked force can be used for this method.

Free Velocity

Whereas the Blocked Force method assumed the interface to be totally fixed, the Free Velocity method is based on a total free interface. This third Component-based Transfer Path method is based on having the active source operate without any external reaction forces. Instead of measuring reaction forces, we measure the motion at the interface and use that to deduce the equivalent force.

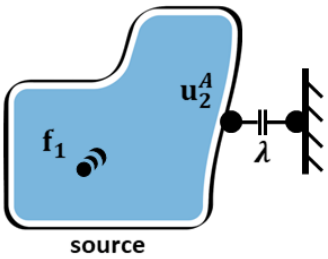

Figure 2.7: Blocked Force

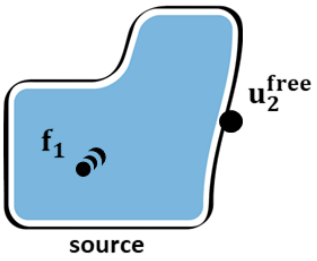

Figure 2.8: Free Velocity

A schematically representation can be seen in figure 2.8. Equation 2.16 shows the method in the Dynamic Substructuring framework, where $\boldsymbol{u}_2^A = \boldsymbol{u}_2^{free}$. Note that the reaction force at the interface is indeed zero. By plugging in the interface measurement $\boldsymbol{u}_2^{free}$ in the equivalent source expression, we find the equivalent force as seen in equation 2.17.

$$\begin{bmatrix} \boldsymbol{u}_1 \\ \boldsymbol{u}_2^A \end{bmatrix} = \begin{bmatrix} \boldsymbol{Y}_{11}^A & \boldsymbol{Y}_{12}^A \\ \boldsymbol{Y}_{21}^A & \boldsymbol{Y}_{22}^A \end{bmatrix} \left(\begin{bmatrix} \boldsymbol{f}_1 \\ \boldsymbol{0} \end{bmatrix} + \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{g}_2^A = \boldsymbol{0} \end{bmatrix} \right) \tag{2.16}$$

$$\boldsymbol{u}_2^{free} = \boldsymbol{Y}_{21}^A \boldsymbol{f}_1 \qquad \rightarrow \qquad \boldsymbol{f}_2^{eq} = \left(\boldsymbol{Y}_{22}^A\right)^{-1} \boldsymbol{u}_2^{free} \tag{2.17}$$

The free velocity measurement is not trivial to perform as an experiment in real life. It is practically not possible to have something suspended without any support. In practice this is solved by suspending the active source with elastic chords. This compliant support ensures the low frequent dynamic behaviour is not blocked.

In-Situ

Both the Blocked Force and Free Velocity method have the problem that the interface conditions are very different to the normal operation for the source being measured. The effects that the interface stiffness has on the source is missing. It could be argued that the source excitation $\boldsymbol{f}_1$ changes due to changing interface stiffness, but to quantify this effect is difficult. This is a topic that has not been researched clearly, but intuitively it makes sense that the boundary conditions change the effective force, since the power of the source is usually determined by its power supply. To understand this effect, a 1D schematic can be found in Appendix REF. This schematic shows a very basic representation of the different degrees of freedom and stiffnesses in the problem.

To negate this problem of changing interface conditions, the In-Situ method is suggested. As the name suggests, the method is based on measuring the active source in its actual situation. As an example, with the In-Situ method the characterisation for an electric steering actuator would be done in the target car it is designed for. The method was suggested by Elliott and Moorhouse [19] and has been used extensively and is now proposed as an ISO standard for source characterisation [52].

As discussed in the General Introduction of this thesis, the goal of source source characterisation from the point of view of Noise-Vibration-Harshness could be to predict the vibrations of the car before it is build. Performing the source characterisation experiment it in the car is a process that comes too late for this to be possible. This means that fundamental design aspects of the car cannot be changed anymore if problems are found in the In-Situ experiment.

Luckily the In-Situ method translates well to performing the experiment on a known test bench. This could be regarded as the next best option besides the actual target assembly. By designing the receiving structure in such a way that the interface stiffness is similar to the target structure, the problems discussed before about changing boundary conditions for the source are managed. Practically this kind of characterisation experiment can also be much easier for the engineers doing the experiment. Everything is accessible which is both beneficial for sensor placement as well as for hammer excitations that are performed to describe the Virtual Point Frequency Response Functions.

Figure 2.6 is a schematic that represents the problem. Equation 2.18 shows a way to deduce the In-Situ formula from a Dynamic Substructuring representation. The motion of the receiving structure is depicted with the subscript 2 here, but since this is a reduction from the $\boldsymbol{u}_4$-space, it can also be written in terms of the indicator sensors. If the coupling is performed the method can be summarised as equation 2.19.

$$\begin{bmatrix} \boldsymbol{u}_1 \\ \boldsymbol{u}_2^A \\ \boldsymbol{u}_2^R \end{bmatrix} = \begin{bmatrix} \boldsymbol{Y}_{11}^A & \boldsymbol{Y}_{12}^A & \boldsymbol{0} \\ \boldsymbol{Y}_{21}^A & \boldsymbol{Y}_{22}^A & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{Y}_{22}^R \end{bmatrix} \left(\begin{bmatrix} \boldsymbol{f}_1 \\ \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix} + \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{g}_2^A \\ \boldsymbol{g}_2^R \end{bmatrix} \right) \qquad \text{with } = \begin{cases} \boldsymbol{u}_2^A & = \boldsymbol{u}_2^R \\ \boldsymbol{g}_2^A & = -\boldsymbol{g}_2^R \end{cases} \tag{2.18}$$

$$\boldsymbol{f}_2^{eq} = \left(\boldsymbol{Y}_{42}^{AR} \right)^+ \boldsymbol{u}_4 \tag{2.19}$$

Figure 2.9 shows a schematic of what a test bench for an In-Situ experiment could look like. This schematic is of a design from the early concept phase and does not represent the final design. The test bench is intended for a compressor, shown as a cylinder in the schematic. This compressor is intended for the air conditioning system of an electric car. It has 3 connection points to its receiving structure. A part of the design cycle for this test bench was performed by me in January 2019, where functional requirements were set and multiple concepts were thought up and compared. The final design was built and used in an experiment performed by colleagues of VIBES.technology and me. This data set is will not be used in this research. Chapter 4 will introduce the numerical case study of an In-Situ experiment that is used.

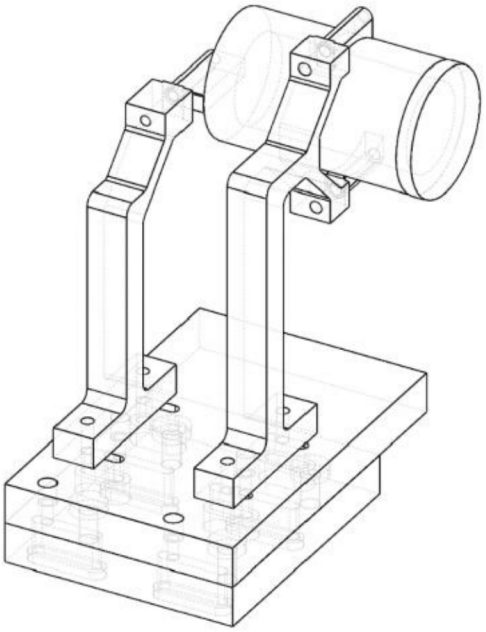

Figure 2.9: Compressor Test Bench

3

Matrix Inversion

3.1. Introduction

The following chapter takes a small detour from the previous chapters. Whereas the first two chapters were all about dynamic descriptions, the following chapter will be more mathematical and based on the inverse problem. To solve the source characterisation problem, an inverse matrix calculation is performed as we have seen in section 2.3. An example of such an inverse relation is seen in equation 3.1. We will see that due to the nature of the problem, which is both overdetermined and error-contaminated, obtaining the solution is not trivial.

In section 3.2 the important mathematical concepts regarding matrix classification are discussed. Some mathematical rules are explained that will be used extensively in the further research. It will become clear that noise on the input data of the inverse problem has the potential to ruin the solution. Based on the Fundamental Theorem of Linear Algebra as defined by Strang [46], the four matrix subspaces are defined and their physical explanation for the source characterisation is explained.

Based on the four subspaces of the matrix we will introduce the Singular Value Decomposition in section 3.3. This can be seen as the most important mathematical concept for this thesis, as it helps us decompose the problem in its fundamental components which have a clear physical representation. It therefore gives insight in how the solution is constructed. It will also lead to a mathematical framework where different inversion techniques can be compared.

Once the mathematical baseline is defined we will dive into multiple different regularisation techniques. These techniques modify the inverse problem by changing the effective singular values of the inverse matrix. Based on literature study a number of promising methods will be highlighted to analyse for the source characterisation problem. The methods introduced here will be tested on a numerical test case in chapter 6. Here the physical relation to these methods will become apparent.

$$\boldsymbol{f}_2^{eq} = (\boldsymbol{Y}_{42})^{+} \boldsymbol{u}_4 \tag{3.1}$$

3.2. Matrix Classification

The following section will act as a basis for the concepts that will be introduced in section 3.4. First some mathematical concepts to evaluate matrices are introduced. After that we will define an exact way of determining whether a problem is well-posed and what the consequences are if this is not the case. Finally, we will take a deep dive into matrix subspaces which will be the basis to define the singular value decomposition in the following section.

3.2.1. Vector and Matrix Norms

Vectors and matrices can be compared in numerous ways. An intuitive way is to define the norm of a vector or a vector space. This defines a size of the vector space which is strictly positive, except

for the zero vector case where it is zero. The most general way of defining the vector norm is seen in equation 3.2. The norm is noted by double vertical bars on either side and a subscript defining what type of norm it is. When no subscript is present, it is assumed it is the 2-norm. The 2-norm is the most intuitive norm as it the Pythagorean length of the vector.

$$\|\boldsymbol{x}\|_p = \left(x_1^p + x_2^p + \cdots + x_n^p\right)^{\frac{1}{p}} \tag{3.2}$$

For matrices the norm definition is similar, as can be seen in equation 3.3. For matrix norms the 2-norm is also the standard norm if no other subscript is provided.

$$\|\boldsymbol{A}\|_p = \left(\sum_{i=1}^{m}\sum_{j=1}^{n}|a_{ij}|^p\right)^{\frac{1}{p}} \tag{3.3}$$

3.2.2. Ill-Conditioned Problems

In mathematics we can differentiate problems by by how well the solution is defined. Jacques Hadamard [25] believed that a mathematical model of any physical phenomena should have the following properties to be well-posed:

1. A solution exists
2. The solution is unique
3. The solution depends continuously on the input.

A well-posed problem is therefore solvable, but more importantly, it is possible to set up the inverse problem. If one of these properties is not met, the problem is said to be ill-posed. Besides this very black and white differentiation, there is also a class of problems that meets all requirements to be well-posed, but still entail issues in solving. These problems are called ill-conditioned problems.

These problems occur when a problem is very sensitive to perturbations. If for a small perturbation on the input, the output changes drastically, solving the inverse problem becomes less trivial. These small perturbations can arise in different situations. They could be caused from numerical instability when solved with finite precision (eg double-precision floating-point values in MATLAB). In experimental modelling the measurement data is always noise contaminated.

This sensitivity for initial conditions can be shown as seen in equation 3.4. The top line shows the perturbation on the input and the corresponding perturbation on the output it causes. If we want to compute the output perturbation, the inverse problem is constructed. To get an upper limit of the output perturbation norm, which is the size of the error, the norm is constructed as seen in the bottom line. Here it is seen that the input perturbation is amplified by the norm of the inverse problem.

$$\begin{aligned} \boldsymbol{A}\left(\boldsymbol{x}+\delta\boldsymbol{x}\right) &= \boldsymbol{b}+\delta\boldsymbol{b} \\ \boldsymbol{A}\delta\boldsymbol{x} &= \delta\boldsymbol{b} \\ \Rightarrow \delta\boldsymbol{b} &= \boldsymbol{A}^{-1}\delta\boldsymbol{x} \\ \|\delta\boldsymbol{b}\| &\leq \|\boldsymbol{A}^{-1}\| \cdot \|\delta\boldsymbol{x}\| \end{aligned} \tag{3.4}$$

The condition number can be computed to quantify how ill-conditioned a problem is. Equation 3.5 shows how the condition number is defined. Roughly speaking one could say the condition number the rate at which the solution will change with respect to a perturbation on the input. Ill-conditions problems have a very large condition number ($\kappa\left(\boldsymbol{A}\right) >> 1$).

$$\kappa\left(\boldsymbol{A}\right) = \lim_{\epsilon\to 0}\sup_{\|\delta\boldsymbol{x}\|<\epsilon}\frac{\|\delta\boldsymbol{b}\|}{\|\delta\boldsymbol{x}\|} \tag{3.5}$$

3.2.3. Matrix Subspaces

As a build-up to section 3.3, the matrix is analysed more closely. From the rank-nullity theorem as defined by Strang [46], it follows that every matrix consists of 4 unique and complementary building blocks that contain all information regarding the matrix. For a matrix $\boldsymbol{A} \in \mathbb{C}^{[m \times n]}$ with rank r the subspaces are defined as seen in table 3.1. These subspaces are complementary vector spaces.

subspace	definition	space	dimension
column space	im($\boldsymbol{A}$)	$\mathbb{C}^m$	r
nullspace	ker($\boldsymbol{A}$)	$\mathbb{C}^n$	$n-r$
row space	im($\boldsymbol{A}^T$)	$\mathbb{C}^n$	r
left nullspace	ker($\boldsymbol{A}^T$)	$\mathbb{C}^m$	$m-r$

Table 3.1: Matrix subspaces

The first matrix subspace of matrix $\boldsymbol{A}$ is the column space, otherwise called the image or range of matrix $\boldsymbol{A}$. The column space of $\boldsymbol{A}$ defines the output that can be reached from any arbitrary type of input. In other words, what vector space does the matrix $\boldsymbol{A}$ span for all possible inputs. The column space is the rank r basis of the columns of $\boldsymbol{A}$.

Translating the notion of the column space to the problem of source characterisation, the column space of $\boldsymbol{Y}_{42}$ forms a rank 6 basis of possible $\boldsymbol{u}_4$ modes for all possible $\boldsymbol{f}_2^{eq}$ inputs. Physically each individual basis vector can be understood as a mode of motion. These modes will be different for each frequency bin. Because the dimension of the column space is defined by the rank of the matrix, the column space is of size 6 for a 6 Degrees of Freedom interface force.

The dual-form of the column space is the nullspace, otherwise called the kernel of matrix $\boldsymbol{A}$. The nullspace of $\boldsymbol{A}$ defines all the inputs that lead to the zero output. The dimension of this vector space is determined by the difference between the rank of the matrix and the number of columns. This is easily explained with the example of a full rank matrix which therefore has independent columns. No linear combinations of these columns will lead to the zero solution, except for the zero input. This means that the nullspace of every matrix $\boldsymbol{A}$ contains at least the zero input. In the scope of this thesis, the nullspace of $\boldsymbol{Y}_{42}$ is the vector space $\boldsymbol{f}_2^{eq}$ that leads to the output $\boldsymbol{u}_4 = \boldsymbol{0}$.

Similarly to the column space and nullspace, the other two matrix subspaces are the row space and the left nullspace. As you would expect, the row space is formed by the rows of matrix $\boldsymbol{A}$. To keep close to the theory of the first two subspaces, these subspaces are computed in the same way as their counterpart, be it using the transposed matrix $\boldsymbol{A}^T$, leading to the informations of the rows, but now in columns again.

The use of the row space and left nullspace is less clear at first sight compared to its counterparts. An intuitive way to see the row space is to see it as the basis to which every solution must comply. The vectors spanning this space can be created by determining the row reduced echelon form of the matrix. Each independent row is a linear equation which must always hold. The row space is therefore an extension of the nullspace, as together they define all possible solutions. The row space and nullspace together span the entire $\mathbb{R}^n$ space and therefore they also must be each others orthogonal complement.

The left nullspace similarly is the orthogonal complement to the column space. In layman's terms it is the vector space that the solution can never reach. It is therefore the extension of the column space and together they span the entire $\mathbb{R}^m$ space. For the problem of the thesis, the left nullspace shows the $\boldsymbol{u}_4$ space that is not possible to reach for any $\boldsymbol{f}_2^{eq}$.

Matrix Classification

Using the definitions of the 4 matrix subspaces as seen in section 3.2.3, every matrix can be classified for what kind of function it is. These classification will help to visualise the mathematical problem back to something intuitive.

The first classification is whether a matrix is *injective* or not. A function is injective if every output has one unique input. In formal terms it is described by equation 3.6. Injection therefore says something about the uniqueness of the inverse function. Injective functions are often called *one-to-one* functions. It is therefore closely linked to the column space and the rank of the matrix. If the column space has the same dimension as the number of columns, in other words if the nullspace has dimension zero (only the trivial nullspace), then the function is said to be injective.

$$\text{for a function } f(\boldsymbol{x}) = \boldsymbol{Ax} \quad \mapsto \quad \text{if } f(\boldsymbol{a}) = f(\boldsymbol{b}), \text{then } \boldsymbol{a} = \boldsymbol{b} \tag{3.6}$$

The second matrix classification is whether a matrix is *surjective* or not. A function is surjective if every possible output can be reached by some input. The formal definition is seen in equation 3.7. Surjective functions are often called *onto* functions, since it maps onto every solution. For a function to be surjective based on the matrix subspaces, we desire the left nullspace to be of dimension zero. Since we are analysing rectangular matrices with matrix dimensions $m > n$, this will never be the case.

$$\text{for a function } f : \mathbb{R}^n \to \mathbb{R}^m \quad \mapsto \quad \text{only if } \boldsymbol{y} \in \mathbb{R}^m, \text{there is } \boldsymbol{x} \in \mathbb{R}^n \text{such that } f(\boldsymbol{x}) = \boldsymbol{y} \tag{3.7}$$

The last matrix classification is whether a matrix is *bijective*. This is the case when a matrix is both injective and surjective. As we already concluded that the system matrices to our interest can not be surjective, it rules out this option.

It is important to grasp the consequences of our function being injective and non-surjective. First of all the non-surjective property is very much logical for a model of a physical phenomena. Certain motions will simply not be possible regardless of the excitation. This fact gives us a tool to analyse experimental obtained data and determine whether we can find an equivalent force to reproduce it. The injective property of the function is a binary requirement for it to be usable for finding an equivalent measure. Both these consequences will be developed and analysed in chapter 5.

3.3. Singular Value Decomposition

The Singular Value Decomposition (SVD) is a factorisation of the system matrix in terms of its subspaces. It can be computed for every positive definite matrix with real or complex entries. Equation 3.8 shows the factorisation in formula form. More intuitive is the graphical representation of figure 3.1, which gives insight in the dimensions of the factorisation.

$$\boldsymbol{Y} = \boldsymbol{U\Sigma V}^T \tag{3.8}$$

The Singular Value Decomposition factorises the matrix into two orthonormal matrices $\boldsymbol{U}$ and $\boldsymbol{V}^T$ which consist of base vectors spanning the column space and row space respectively. These base vectors are called the left-singular vectors and the right-singular vectors respectively. The factorisation also includes a matrix $\boldsymbol{\Sigma}$ with non-negative real numbers on the diagonal which function as the weighting factors. These weightings are known as the singular values.

Table 3.2 shows the matrix subspace overview from section 3.2.3 again, but this time with an additional column that shows how it fits in the Singular Value Decomposition. The strength of the Singular

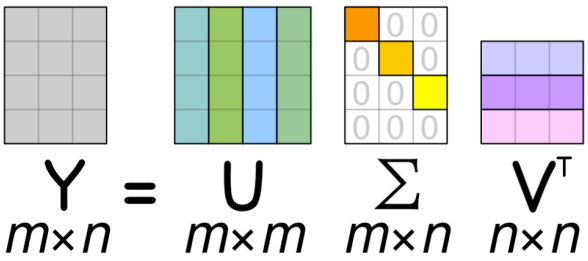

Figure 3.1: Graphical representation of Singular Value Decomposition

Value Decomposition is that all 4 subspaces are present in the decomposition, whereas for the original matrix representation both null spaces are not explicitly present. This means the entire vector space described by the system matrix $\boldsymbol{Y} \in \mathbb{R}^{[m \times n]}$ is accounted for. The singular value matrix $\boldsymbol{\Sigma}$ determine to which extend the base vectors are active in the function. Consequently, the rank of the system matrix $\boldsymbol{Y}$ is equal to the number of non-zero singular values.

subspace	definition	space	dimension	SVD basis
column space	im($\boldsymbol{Y}$)	$\mathbb{C}^m$	r	first r columns of $\boldsymbol{U}$
nullspace	ker($\boldsymbol{Y}$)	$\mathbb{C}^n$	$n-r$	last $(n-r)$ columns of $\boldsymbol{V}$
row space	im($\boldsymbol{Y}^T$)	$\mathbb{C}^n$	r	first r columns of $\boldsymbol{U}$
left nullspace	ker($\boldsymbol{Y}^T$)	$\mathbb{C}^m$	$m-r$	last $(m-r)$ columns of $\boldsymbol{U}$

Table 3.2: Matrix subspaces for SVD

The concept of the Singular Value Decomposition lends itself very well to the problem of source characterisation using $\boldsymbol{Y}_{42}$. It makes it possible to decompose the matrix into the fundamental building blocks of the physical phenomena it describes. This decomposition has to be determined for every frequency bin, but computationally this is no problem since the problem can be regarded as low rank.

The way the Singular Value Decomposition can be understood is as a summation of rank 1 matrices that each describe 1 mode of the column space and its corresponding row space counter part. The norm of this matrix is determined by the corresponding singular value. In this way each of these matrices filter 1 mode of $\boldsymbol{Y}_{42}$. The problem of $\boldsymbol{u} = \boldsymbol{Y}\boldsymbol{f}$ for such a 1 mode matrix can only excite that particular mode in the $\boldsymbol{u}_4$-space. It also means the input to the problem is only sensitive to the corresponding $\boldsymbol{f}_2$-mode. Equation 3.9 shows the summation, where the lower case vectors are the columns of $\boldsymbol{U}$ and $\boldsymbol{V}$. The representation of equation 3.9 will be used in section 3.4 as it will be possible to add regularisation to each rank of the system separately.

$$\boldsymbol{Y} = \sum_{n=1}^{r} \sigma_i \boldsymbol{u_i} \boldsymbol{v_i}^T \tag{3.9}$$

This same logic matrices describing 1 mode is true for the inverse problem. Computing the inverse system matrix for a SVD decomposed matrix is very straight fowards due to the nature of its orthonormal and diagonal matrices. In the inverse the vector spaces are written in inverse relation and the singular values are inverted accordingly. Equation 3.10 shows both SVD representations.

$$\boldsymbol{Y}^{-1} = \boldsymbol{V}\boldsymbol{\Sigma}^{+}\boldsymbol{U}^T \qquad \leftrightarrow \qquad \boldsymbol{Y}^{-1} = \sum_{i=1}^{r} \frac{1}{\sigma_i} \boldsymbol{v_i} \boldsymbol{u_i}^T \tag{3.10}$$

Similarly, the input to the regular or inverse problem can be written down in terms of the left and right singular vectors, as can be seen in equation 3.11. Note that the singular vectors are denoted with a subscript i, whereas the actual response data has no subscript. For all dimension of both spaces, the response or force is projected onto its base vectors.

$$\boldsymbol{f} = \sum_{i=1}^{m} \left(\boldsymbol{v}_i^T \boldsymbol{f}\right) \boldsymbol{v_i} \qquad\qquad \boldsymbol{u} = \sum_{i=1}^{n} \left(\boldsymbol{u}_i^T \boldsymbol{u}\right) \boldsymbol{u_i} \tag{3.11}$$

The concept of writing down the equivalent force or response data in terms of the singular vectors will be further developed in chapter 5. This is due to the fact that this equality will always hold for a summation over the entire vector space, but the equality does not necessarily hold for a summation over the first rank r singular vectors. For now we will assume that the equivalent force and response data can be constructed with the first rank r singular vectors.

Combining equation 3.10 and equation 3.11 we can construct the solution of the inverse problem in terms of the singular vectors and the response data. This leads to a solution in the form of equation

3.12. This representation shows how the solution is constructed in its most basic form. The input data is projected on each mode and based on the norm of the projection and the inverse of the singular value the corresponding equivalent force mode is found.

$$\boldsymbol{f} = \sum_{i=1}^{r} \frac{\left(\boldsymbol{u}_i^T \boldsymbol{u}\right)}{\sigma_i} \boldsymbol{v_i} \tag{3.12}$$

In the next section it will be clear that the assumption $\boldsymbol{u}_4^{meas} \in \begin{bmatrix} \boldsymbol{u}_1 & \dots & \boldsymbol{u}_6 \end{bmatrix}$ does not hold. It means that somes modes of the left nullspace of $\boldsymbol{Y}_{42}$ are present in the measurement data. The reason that this occurs is investigated in chapter 5.

In section 3.4 multiple different regularisation methods will be applied to the inverse problem. These regularisation techniques will be rewritten so that we construct a solution in the form of equation 3.12 where possible. To do this, we introduce the concept of *filter factors*, denoted by ϕ_i [29]. This will lead to a regularised solution as seen in equation 3.13.

$$\boldsymbol{f}_{reg} = \sum_{i=1}^{r} \phi_i \frac{\left(\boldsymbol{u}_i^T \boldsymbol{u}\right)}{\sigma_i} \boldsymbol{v_i} \tag{3.13}$$

3.4. Matrix Inversion Techniques

The following section is an overview of different matrix inversion techniques. With these regularisation techniques we seek a method to robustly find an approximation of the inverse solution while suppressing the effect of noise. The inverse problem is the same as we have seen in equation 3.1. The system matrix is an overdetermined system $\boldsymbol{Y} \in \mathbb{C}^{[m \times n]}$ with $m > n$ with rank r. The challenge in solving the inverse problem is two-fold. First of all, the matrix is non-square, in other words it is overdetermined. As a consequence, the inverse problem will generally also be inconsistent, meaning it has no solution. As a second challenge, the measurement data is error-contaminated.

In equation 3.12 and equation 3.13 we see that the shape of the inverse solution is based on the left and right singular eigenvectors. The amplitude is determined by the singular values and the projection of the measurement data on each left singular mode. All regularisation methods discussed in the following section will be methods that change the singular values of the inverse matrix. What this implies is that the mode shapes of the left and right singular vectors stay unaltered and the physical relevance is kept. Figure B.1 shows an example of the singular values for $(\boldsymbol{Y}_{42})^+$ for a typical source characterisation problem.

The goal of this section is not to derive every regularisation technique from the ground up, but rather to relate their functioning to the physical phenomena they are controlling. In other words, how are the singular values of the inverse matrix modified for different regularisation techniques. As a baseline the non-regularised solution is analysed.

Matrix regularisation methods can be classified in two categories. On the one side there are direct methods that solve the inverse problem in a single calculation. Conversely, there are also iterative methods. The methods discussed in this thesis will all be direct methods. Iterative methods can help in reducing the error amplification as is proven by Biemond [8]. Many publications have been done on the topic, especially with regards to digital image processing. Iterative methods have been tried on the source characterisation problem. Thite showed that while a combination of direct and iterative methods can improve the accuracy of the result, the gain by iterative methods was only marginal while there is considerable extra computational time [49].

Since the system matrices are also obtained from experiments, they have the same potential error-contamination issues as the measurement data. This can lead to a noisy system matrix, or even a system matrix where certain anti-resonances are not captured because they are below the noise floor. Bendat developed a method to statistically determine a corrected system matrix based on multiple

measurements for a user defined normal deviation bound [6]. This method will help for a noisy system matrix, but it will not be able to recover anti-resonance details that are not captured. Thite showed that errors in the system matrix are expected to be smaller than in measurement data and that the noise-contamination of the system matrix can be solved using the singular value truncation [48]. For this reason the scope of the chapter will be limited and we assume the system matrix is correct.

3.4.1. Least Square Method

The first type of matrix inversion techniques is the Moore-Penrose inverse. This method is the most general way of performing a pseudo-inverse of an overdetermined system. It could be described as the non-regularised solution. The solution will be the least squares approximation to the system of equations. It is often called the naive solution, since it does not take into account any system properties, but rather bluntly finds the best fit approximation of equation 3.14. For clarity, the solution has been given a subscript, whereas the subscript for the system matrix is omitted.

$$\min_{\boldsymbol{f}_{lsq} \in \mathbb{C}^n} \|\boldsymbol{u} - \boldsymbol{Y}\boldsymbol{f}_{lsq}\| \tag{3.14}$$

The fact that the overdetermined system has no solution, means that the solution of $\boldsymbol{Y}\boldsymbol{f} = \boldsymbol{u}$ is not inside the column space of $\boldsymbol{Y}$. The closest approximation in the sense of closest distance (smallest 2-norm), will be the projection of $\boldsymbol{u}$ onto the column space. Equation 3.15 shows this relation, where ker($\boldsymbol{Y}$) is the kernel of the matrix, which is equal to the column space.

$$\boldsymbol{Y}\boldsymbol{f}_{lsq} = \text{proj}_{\text{ker}(\boldsymbol{Y})}(\boldsymbol{u}) \tag{3.15}$$

To find the closest approximation possible, we plug in equation 3.15 into equation 3.14, as can be seen in equation 3.16. Looking at the first line, the right hand side can be simplified, since the projection of $\boldsymbol{u}$ onto the column space minus the vector $\boldsymbol{u}$, is by definition equal to the the orthogonal complement of the column space. This is equal to the left nullspace, as can be seen in the second line of equation 3.16. This insight is crucial in understanding the residual of the inverse problem. If the left null space is zero, the residual will be zero too.

$$\begin{aligned} \boldsymbol{u} - \boldsymbol{Y}\boldsymbol{f}_{lsq} &= \boldsymbol{u} - \text{proj}_{\text{ker}(\boldsymbol{Y})}(\boldsymbol{u}) = \ker(\boldsymbol{Y})^{\perp} \\ \boldsymbol{u} - \boldsymbol{Y}\boldsymbol{f}_{lsq} &\in \text{null}\left(\boldsymbol{Y}^T\right) \end{aligned} \tag{3.16}$$

If we pre-multiply the result of equation 3.16 with $\boldsymbol{Y}^T$, we get a result equal to zero, since the result is in the null space of $\boldsymbol{Y}^T$. Line 2 of equation 3.17 shows the outcome we obtain. This equation is identical

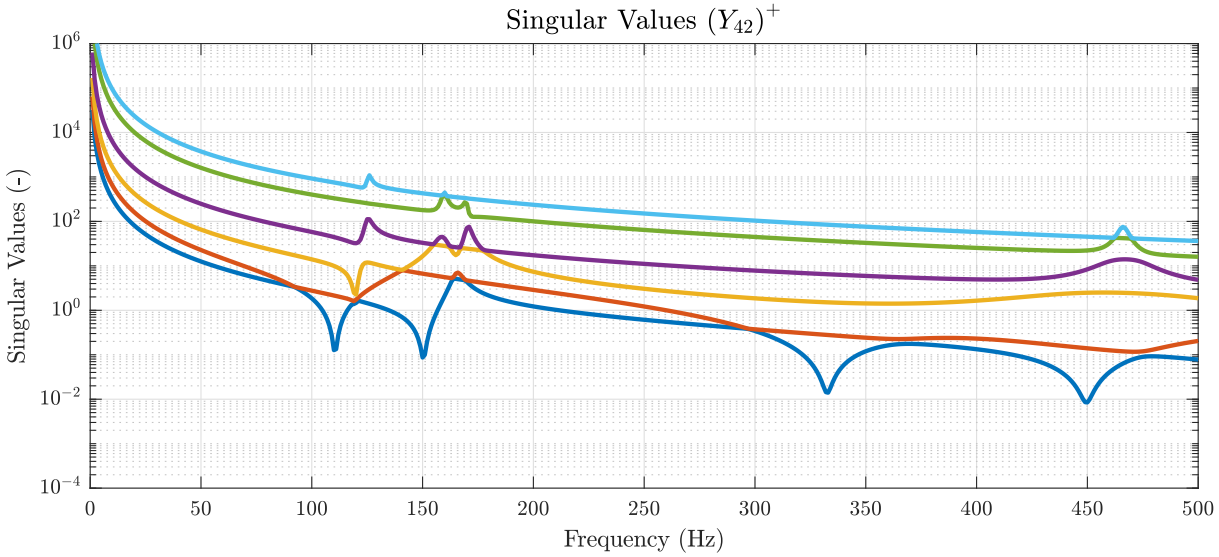

Figure 3.2: Singular Values for inverse problem

to the standard form of the problem, but now pre-multiplied with $\boldsymbol{Y}^T$. The advantage of this is that the problem is changed from a pseudo-inverse problem, to an exact inverse problem. This is due to the fact that for every $[m \times n]$ matrix with rank $r = n$, the relation $\left(\boldsymbol{Y}^T\boldsymbol{Y}\right)$ will have an exact inverse. Note that the conjugate transpose is used here. This leads to an analytic relation we find for the least square solution.

$$\begin{aligned} \boldsymbol{Y}^T\left(\boldsymbol{u} - \boldsymbol{Y}\boldsymbol{f}_{lsq}\right) &= \boldsymbol{0} \\ \boldsymbol{Y}^T\boldsymbol{u} - \boldsymbol{Y}^T\boldsymbol{Y}\boldsymbol{f}_{lsq} &= \boldsymbol{0} \\ \rightarrow \boldsymbol{f}_{lsq} &= \left(\boldsymbol{Y}^T\boldsymbol{Y}\right)^{-1}\boldsymbol{Y}^T\boldsymbol{u} \end{aligned} \tag{3.17}$$

From the analytic solution we can define an inverse. This inverse is identical to the inverse seen in equation 3.10. The advantage of this form is that no Singular Value Decomposition has to be performed. Besides the practical advantage of this form, it is also shows the elegance of the Least Squares Solution, as it is now an analytically determined solution to a minimisation problem.

$$\boldsymbol{Y}^+ = \left(\boldsymbol{Y}^T\boldsymbol{Y}\right)^{-1}\boldsymbol{Y}^T \tag{3.18}$$

3.4.2. Weighted Least Square Method

The Weighted Least Square Method is an extension of the regular Least Squares Method. Whereas the Least Square Solution computes the solution based on the overall best fit of the data, the Weighted Least Square Method is an adaptation to give each measurement channel its own weight.

The mathematical definition of the method can be seen in equation 6.7, which originates from a derivation from Strang [47]. It can be seen to be very similar to the Least Square solution of equation 3.18, but extended with a weighting matrix. The weighting matrix $\boldsymbol{W} \in \mathbb{R}^{[m \times m]}$ is a positive definite matrix.

$$\boldsymbol{Y}^T\boldsymbol{W}\boldsymbol{Y}\boldsymbol{f} = \boldsymbol{Y}^T\boldsymbol{W}\boldsymbol{u} \tag{3.19}$$

The weighting matrix can be chosen in multiple ways. Powell suggests using the signal to noise ratio for each measurement channel as a diagonal weighting matrix [43]. Since the measurements with more noise will have a lower signal to noise ratio, these measurements are taken to be less important. This method can be useful if some measurement channels have very poor signal to noise, which possibly leads to a solution that is fitted on noise data instead of real physical phenomena. This method was originally suggested for the case where multiple different channels are used with different noise levels and even with different units (for example accelerometers and optical motion sensors). The different channels will be non-dimensionalised in this manner so that they can be used together.

The weighting matrix can also be chosen from a more statistical point of view, as suggested by Aitken [1]. This method is based on the fact that we can do multiple independent measurements. It is assumed that the physicality of the problem stays the same and that all difference in measurement is due to the uncorrelated white noise. Aitken proved that for such cases a Best Linear Unbiased Estimator can be obtained when using the reciprocal of the variance of the measurements as the weighting matrix entries. These measurements are assumed to be uncorrelated due to the uncorrelated white noise.

3.4.3. Truncated Singular Value Decomposition

The inversion techniques so far were based on different methods of data fitting. An alternative strategy is to adapt the system matrix before the inverse is performed. By setting certain singular values to zero, we discard their effect on the problem. This is called the Truncated Singular Value Decomposition. There are multiple ways of deciding which singular values should be neglected.

To relate this method to the physical phenomena occurring, we remember that the system matrix $\boldsymbol{Y}$ is a receptance matrix that describes the displacements for a unit of force. This implies that the modes related to the smallest singular value is the mode that is the most stiff and therefore occurs for the relative most amount of force. For the inverse problem, this mode becomes the mode related to the largest singular value of the inverse problem, and in the dynamic stiffness description is indeed the

stiffest mode. On this mode the noise will have the largest influence, since the signal to noise ratio will be the worst here. The assumption that the stiffest mode will have the worst signal to noise holds only true for broadband excitations. It might very well be the case that all motion occurs in this mode for some source excitation.

A simple way to truncate the system matrix is to only keep the first n_{sing} number of singular values. Equation 3.20 shows the filter factors for this kind of truncation. This method can be very strong if the system matrix has a rank that is higher than the actual physical problem it describes. In these cases, the singular values for these high ranks do not describe physical phenomena. They originate from a bad measurement and truncating these mode can lead to a better result. For low rank problems, this method is often too blunt, as it changes the physical description without taking the specific excitation into account.

$$\boldsymbol{f}_{tsvd} = \sum_{i=1}^{r} \phi_i \frac{\left(\boldsymbol{u}_i^T \boldsymbol{u}\right)}{\sigma_i} \boldsymbol{v_i} \qquad \phi_i = \begin{cases} 1 & \text{if } i \leq n_{sing} \\ 0 & \text{otherwise} \end{cases} \tag{3.20}$$

A refinement of this method is to decide which singular values to neglect on the basis of how much the corresponding modes are operational. By doing this, we abandon the logic that stiff modes are inherently bad news for the characterisation, but relate the problem to what is occurring. To do this, the measurement data is projected onto the left singular values. For modes where this projection is below a user defined threshold, the singular values is set to zero using the filter factors. This can be seen in equation 3.21.

$$\boldsymbol{f}_{tsvd} = \sum_{i=1}^{r} \phi_i \frac{\left(\boldsymbol{u}_i^T \boldsymbol{u}\right)}{\sigma_i} \boldsymbol{v_i} \qquad \phi_i = \begin{cases} 1 & \text{if } \left(\boldsymbol{u}_i^T \boldsymbol{u}\right) \geq \tau \\ 0 & \text{otherwise} \end{cases} \tag{3.21}$$

It is clear that the second type of Truncated Singular Value Decomposition method based on a threshold value is more robust for the source characterisation problem. Choosing the threshold value is not trivial however. The projection of the measurement data on each left singular mode is not normalised in the current form, meaning the truncation will be a function of how large the excitation is, rather than truncating based on physical modes.

Significant Rank Criterion

Janssen and Verheij suggested a method to choose the threshold based on the noise level and the system properties, which they dubbed the Significant Rank Criterion [31]. Their suggested method does not depend on the projection of the data as we have seen in equation 3.21, but rather at the size of the singular value. The method takes advantage of the notion that the measurement can be written as a summation of contributions from the individual modes. For modes that contribute less than the estimated error the singular values are rejected.

Equation 3.22 shows a way to write the measurement data in terms of the noiseless, 'true' response (denoted with a hat accent) and a separate noise contribution. This concept will be further developed and explained in section 6.2.

$$\boldsymbol{u} = \hat{\boldsymbol{u}} + \boldsymbol{e} \tag{3.22}$$

The derivation of the criterion makes use of the Singular Value Decomposition, specifically that the norm of every mode is determined by its singular value. Starting from the relation $\boldsymbol{u} = \boldsymbol{Y}\boldsymbol{f}$, Janssen and Verheij take a bit of leeway in the derivation and assume that the norm of the product is approximately the same as the product of the norm, as can be seen in equation 3.23. Note that this equation assumes that the force can control the motion space fully, something we will challenge in section 5.2.

$$\|\hat{\boldsymbol{u}}\| + \|\boldsymbol{e}\| \approx \sum_{i=1}^{m} \sigma_i \|\boldsymbol{f}\| \tag{3.23}$$

Looking at equation 3.23, the singular values can be split in two groups based on their order of magnitude. A group resulting in responses with a order of magnitude larger than the error magnitude $\|\boldsymbol{e}\|$, and responses with a smaller order of magnitude. This last group can be regarded as insignificant and

these will be the singular values to reject. In the mathematical sense, equation 3.24 shows the relation for the rejected singular values group.

$$\sigma_i \|\boldsymbol{f}\| \leq \|\boldsymbol{e}\| \tag{3.24}$$

As a final step the error is normalised using the measurement data as can be seen on the right hand side of equation 3.25. The left hand side is normalised in a similar fashion, assuming that the norm of the measurement can be approximated by the contribution of the first singular value.

$$\frac{\sigma_i \|\boldsymbol{f}\|}{\sigma_1 \|\boldsymbol{f}\|} \leq \frac{\|\boldsymbol{e}\|}{\|\boldsymbol{u}\|} \tag{3.25}$$

Rewriting equation 3.25 to cancel out some terms, leads to equation 3.26. The singular values below this upper bound are deemed to be not significant and can be rejected. This leads to an overall threshold for the singular values.

$$\sigma_i \leq \frac{\|\boldsymbol{e}\|}{\|\boldsymbol{u}\|} \sigma_1 \tag{3.26}$$

3.4.4. Tikhonov Regularisation

A method to have more control over the inverse problem is to add extra terms to the minimiser that was introduced in equation 3.15. A widely used method (cite) is the Tikhonov regularisation method, where a penalty term is added to minimiser in terms of the norm of the solution. This solution norm is premultiplied with a regularisation parameter α that gives a weighting to the penalty term. The most general case of the Tikhonov regularised problem can be seen in equation 3.27.

$$\min_{\boldsymbol{f} \in \mathbb{C}^n} \left\{ \|\boldsymbol{u} - \boldsymbol{Y}\boldsymbol{f}\|^2 + \alpha \|\boldsymbol{f}\|^2 \right\} \tag{3.27}$$

Equation 3.27 should be interpreted as a minimiser of two terms. The first term is the residual term and this defines the mismatch of the solution mapped with the system matrix compared to the measurement. The second term is the regularisation term. The regularisation parameter balances the importance between these two terms in the minimiser. For large regularisation parameter choices the minimiser will be dominated by the regularisation term, meaning a small norm is more important than a small residual error. It smoothens the problem. Conversely, small amounts of regularisation will mean a smaller residual, but therefore also a solution that is more sensitive to noise.

$$\boldsymbol{Y}_\alpha^{-1} = \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha^2 \boldsymbol{I}\right)^{-1} \boldsymbol{Y}^T \tag{3.28}$$

It can be shown that equation 3.27 can be rewritten in the form of equation 3.28 [28]. This form shows a close relation to the least square solution. It is often called the 'stacked' Tikhonov regularisation form.

$$\boldsymbol{f}_\alpha = \sum_{i=1}^{r} \phi_i \frac{\left(\boldsymbol{u}_i^T \boldsymbol{u}\right)}{\sigma_i} \boldsymbol{v_i} \qquad \phi_i = \frac{\sigma_i^2}{\sigma_i^2 + \alpha^2} \tag{3.29}$$

The Tikhonov solution can be written in terms of filter factors as seen in equation 3.29. This representation is strong as it shows the regularisation only has an effect on the weighting of different modes, not on the rank or vectorial subspaces of the system. For no regularisation the filter factor will be unity as expected. For larger regularisation parameters the weighting of that mode goes down. This effect is stronger for smaller singular values. This means the Tikhonov regularisation dampens the influence of the singular vectors corresponding to small singular values the most.

$$\phi_i = \frac{\sigma_i^2}{\sigma_i^2 + \alpha^2} \approx \begin{cases} 1 & \sigma_i \gg \alpha \\ \sigma_i^2/\alpha^2 & \sigma_i \ll \alpha \end{cases} \tag{3.30}$$

3.5. Tikhonov Regularisation Parameter Choice

Tikhonov Regularisation from section 3.4.4 gives us a parameter which we can tune to change the problem. For each different parameter choice, a different answer is found. Since only the weight of the

singular values are altered, the function remains an injective function, as defined in section 3.2. This means there is a unique answer for each regularisation parameter choice. Since the actual solution is unknown however, it is not possible to tune this parameter by hand.

In Appendix E a quantitative analysis of the error due to Tikhonov Regularisation is made. The error is split up in a measurement error due to noise contamination and a regularisation error due to a non-zero regularisation parameter α. Upper bounds for both errors are defined, giving insight in how both errors are a function of the regularisation parameter.

Choosing the regularisation parameter is not a trivial task. The goal is to find middle ground between describing the true solution and being too sensitive to noise contamination. Research has been done to compare the popular different techniques. Choi compared using the L-Curve, Ordinary Cross-Validation and General Cross-Validation for an inverse force determination problem for multiple noise levels. His conclusion was that while all three methods show potential for specific cases, the L-Curve method performs the best, especially for high noise levels [12].

3.5.1. L-Curve Method

An insightful method to compare the measurement and regularisation error, is the L-Curve Method. The L-Curve Method is a visual representation of the two terms in Tikhonov regulrisation as we have seen in equation 3.27, where the first term is dubbed the residual norm and the second term the solution norm. For a range of regularisation values, the residual norm and corresponding solution norm are plotted against each other on a logarithmic scale. A typical shape of an L-Curve is seen in figure 3.3.

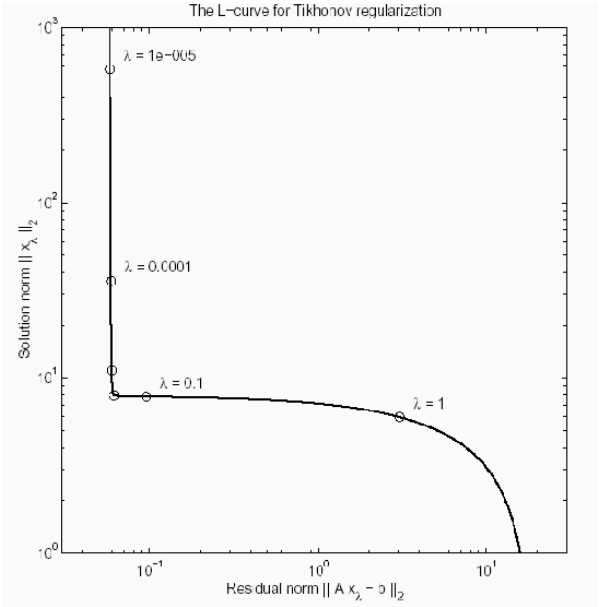

Figure 3.3: Typical L-Curve

Figure 3.3 shows where the method derives its name from. The L-Curve visualises the two regions of regularisation. Under-regularisation occurs for small values of α as can be seen in the vertical part of the curve, and conversely over-regularisation occurs for larger values of α as can be seen in the horizontal part of the curve.

The under-regularisation is characterised by a very sensitive behaviour in terms of the solution norm. For a small decrease in α, the solution norm increases rapidly while the residual norm stays relatively constant. Relating this to the source characterisation problem, it means that in the under-regularised region, a large increase in the force is needed to get a small decrease in the residual motion. Intuitively the smallest solution norm that leads to the smallest residual norm seems the most likely true answer.

Similarly, the over-regularised region is characterised by an inverse relation, where the residual norm is sensitive to changes in α, while the solution norm stays relatively constant. For the source characterisation problem this would mean that for a slight increase of the equivalent force norm $\|\boldsymbol{f}\|$, the residual motion norm $\|\boldsymbol{u} - \boldsymbol{Yf}\|$ decreases rapidly. Intuitively a smaller residual norm for an almost equal solution norm seems most likely as the true solution.

The over- and under-regularised region transition in the corner of the L-Curve. The regularisation parameter for which this occurs is the optimal regularisation according to Hanssen [27]. It is not trivial to numerically find the transition point of the two regions since the shape of the L-Curve is not always an ideal L-shape. A lot of research is done to implement an algorithm that robustly finds the parameter.

3.5.2. Morozov's Discrepancy Principle

Morozov's Discrepancy Principle is a method to determine the regularisation parameter α bases on prior knowledge of the noise level. The reasoning behind this principle is that since the measurement data $\boldsymbol{u}$ is error-contaminated with an error norm of size $\|\boldsymbol{e}\|$, it does not make sense to solve for a residual norm that is smaller than that error norm. This is because the maximum accuracy of the measurement data is the noise level. Finding a solution that has a higher accuracy is wrong in terms of significant figures.

$$\min_{\boldsymbol{f}_{mor} \in \mathbb{C}^n} \|\boldsymbol{u} - \boldsymbol{Yf}_{mor}\| \geq \|\boldsymbol{e}\| \tag{3.31}$$

The implementation of Morozov's Discrepancy Principle is a constraint to the minimiser as seen in equation 6.3 in terms of an upper bound, as can be seen in equation 3.31. To solve this constrained problem, it is rewritten in the form of Tikhonov Regularisation. The regularisation parameter is controlled to have the residual be larger than the noise floor.

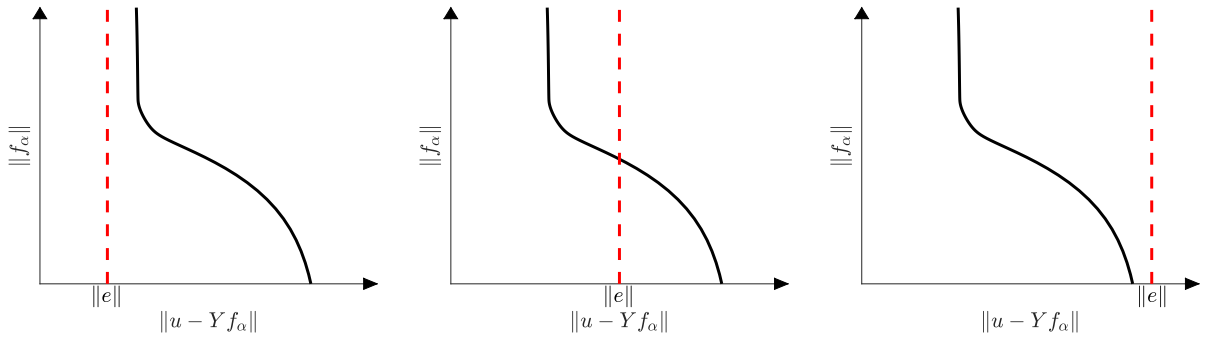

(a) Case 1: residual above noise level

(b) Case 2: residual on noise level

(c) Case 3: residual below noise level

Figure 3.4: Morozozov possible cases for residual noise floor shown on L-Curve

Morozov's Discrepancy Principle can be visualised in the L-Curve. Figure 3.4 shows a schematic view of the 3 possible cases of the relation between the residual and solution norm and the noise floor depicted in red. From these cases it can be determined whether the Discrepancy Principle should be used to determine the regularisation parameter.

In the first case (figure 3.4a) the residual motion is above the noise level for all possible regularisation parameters. These are cases where the equivalent force is not able to span the entire vector space and a residual motion is left, no matter how large the equivalent force is chosen. For these cases, the Discrepancy Principle is not a useful method. These cases will be analysed in more detail in chapter 5.

The third case (figure 4.4) shows the case where all residual motion is below the noise floor. Even for extremely large regularisation values for which the solution norm is close to zero, the residual is still below the noise floor. This case occurs when the noise level is very high which shifts the noise

boundary to the right. It can also occur if there is almost no motion occuring. Either way, these cases should not occur in a proper experiment.

Finally the second case (figure 4.3) shows the situation where the Discrepancy Principle can be of use. The inverse problem can be solved well below the noise floor. The Discrepancy Principle shows its strength at these cases, since the norm of the solution can be controlled to ensure the solution is not overfitted to the noise. The optimal regularisation parameter for solving up to the noise level can be found in an iterative scheme for problems as case 2. The inverse problem is solved for a slowly increasing α up to the point where the residual norm meets the noise floor.

Morozov's Discrepancy Principle provides a robust result in the case where the residual error is due to noise on the measurement. A good estimate of the noise norm is needed. This is not always possible in practice. For the experimental modelling problems there are ways of determining the noise level by either a separate experiment without any excitation to measure only the noise on the sensors, or by the data sheet of the hardware used.

3.5.3. Wiener Filter

The Wiener Filter is a method of determining the Tikhonov regularisation parameter based on the norm of the excitation and the norm of the noise level. Unlike Morozov's Discrepancy Principle it is a value that can be obtained via a direct calculation. The methods finds it origins in inverse imaging problems and its use was suggested by Murli [40].

$$\alpha_{wien}(\omega) = \frac{\|\boldsymbol{e}\|}{\|\boldsymbol{u}\|} \tag{3.32}$$

Equation 3.32 shows the definition of the regularisation parameter from Wiener. It is the reciprocal of the signal to noise ratio. In the case where the signal to noise ratio is large and hence a lot of excitation is occurring, the Wiener regularisation parameter will be small and hence will not regularise the problem by much. In the cases where the signal to noise is low however, a lot of regularisation happens and the solution norm will be small in comparison.

Choosing the regularisation parameter using the Wiener Filter is especially powerful for harmonic excitations [59]. The excitations cause responses on certain frequencies bins only. For the frequency bins where there is little to no excitation, the measurement consists of only the noise measurement. The inverse solution would then also only be comprised of noise data which has no relation to the actual physical problem.

II

Source Characterisation

4

Interface Description

4.1. Introduction

The following chapter elaborates on the Transmission of the *Source-Transmission-Receiver*-model. The transmission from source to receiver occurs at the interface. We will work towards a description of the interface that will make it possible to evaluate the assumptions made while modelling the interface. In this chapter we will see the first signs that a traditional 6 Degree of Freedom interface equivalent force measure is perhaps not a valid description of the real-life problem.

Firstly we will analyse the real-life interface from an observatory point of view. This will be done by overloading the interface and observing what is measured on the receiving side. Using our developed knowledge about matrix subspaces from chapter 3 we can describe this motion in terms of a summation of rank 1 matrices. With this we will seek a quantified measures to describe the effective rank of the interface. The effective rank describes how much of the vibrations from the source passes to the receiving structure, without discretising the interface. This is therefore a proper representation of the real-life problem.

Once we know how much dynamics passes the interface in the real-life problem, we can start with the analysis of the modelled interface. The interface as we model it will be analysed in a few ways. First of all the rank will be compared to the effective rank of the real-life problem. The will be done by making use of the so-called Round-Way Trip to find an expression of the dynamics passing over the discretised interface.

With a quantified measure of the rank of the interface, we will start to develop a method to verify whether the rank of our equivalent measure is sufficient. This will be taken a step further by making use of the so-called Interface Completeness Criterion as was introduced by Meggit [36]. A new version of this Criterion will be developed followed up with an explanation of the discrepancies between the real-life interface and our modelled interface found here.

In section 4.4 the insights from the Interface Completeness Criterion are translated to what is happening physically at the interface. This will answer the first question of how blocked the interface is when applying the blocked force. It will turn out that for certain frequencies the equivalent force on the interface cannot be a valid representation of the source. By using the analogy of a guitar string, the problem is explained in layman terms.

As a validation step to the new found insight, we seek a mathematical way to show that it is indeed true. This chapter will close with an analysis based on the co-located interface Frequency Response Function $\boldsymbol{Y}_{22}$. This response, otherwise called the Driving Point Frequency Response Function, will proof that these troublesome frequency bins are inherently caused by the choice of equivalent force location.

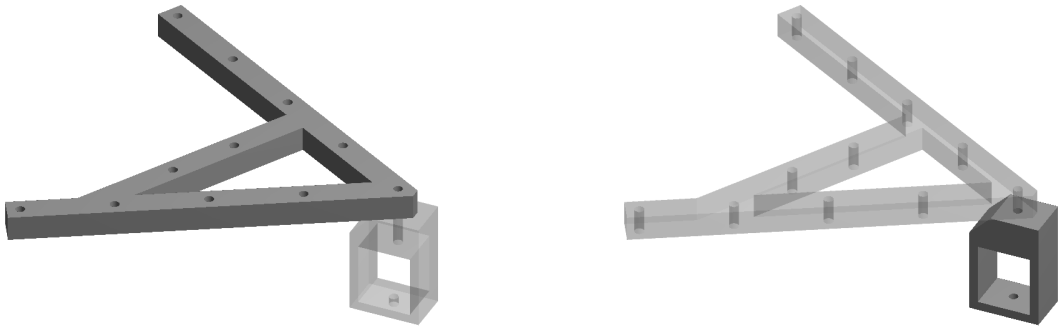

Figure 4.1: Source and receiving structure

The analysis in this section is supported with experiments done on a numerical model. This model is constructed using the VIBES.technology Matlab Toolbox and the specifics can be found in appendix A. It consists of a source structure that is connected to a receiving structure as can be seen in figure 4.1. The numerical model simulates experiments with which Frequency Response Functions can be constructed and responses can be measured from chosen excitations.

4.2. Effective Interface Rank

In the following section we will find a measure to describe the effective rank of the interface. The effective rank differs from the actual rank in the sense that the actual rank can be very large with many modes that have very little participation in the system, and for the effective rank we only look at the most dominant modes. This analysis can be done using the Singular Value Decomposition knowledge that was developed in section 3.3.

Section 3.3 showed us that we can represent the Frequency Response Function as a summation of rank 1 matrices as can be seen in equation 4.1. As a quick reminder, each rank 1 matrix has a column and row space consisting of the mode shape of that rank for the sensors and impacts respectively. The weight of this matrix is determined by the singular value. The decomposition is done in such a way that the singular values are decreasing for every mode.

$$\boldsymbol{Y} = \sum_{i=1}^{r} \sigma_i \boldsymbol{u_i} \boldsymbol{v_i}^T \tag{4.1}$$

From the theory discussed in chapter 2 we expect the interface to be of rank 6. We have seen this in Dynamic Substructuring that there are countless publications regarding the coupling of two substructures using a 6 Degree of Freedom interface. We have also seen in this in section 2.3 concerning Transfer Path Analysis specifically, where the characterisation was also done using 3 forces and 3 moments.

To verify the rank 6 nature of a structural interface we define a new measure that helps us determine the effective rank of the interface. This is done by comparing the size of the first n_{sing} singular values to the entire range of singular values r. This will give a value bounded between 0 and 1 that indicates the effective rank. We will call this measure the participation factor and it can be seen in equation 4.2.

$$\text{participation}\left(n_{\text{sing}}\right) = \frac{\sum_{i=1}^{n_{\text{sing}}} \sigma_i}{\sum_{i=1}^{r} \sigma_i} \tag{4.2}$$

4.2.1. Real-Life Interface

We will start by trying to define the effective rank of the real-life interface. What is meant with real-life here is not that the experiment will be performed in real-life, but that the interface is modelled without forcing it to be of any specific form. The specifics of how the interface is modelled for it to represent a real-life connection can be seen in section A.1 and is based on the knowledge from 2.2.

This experiment makes heavy use of the Singular Value Decomposition method. We will overload the interface with a very large number of impacts on the active side and a very large number of sensors on the passive side. The goal is to create a very large Frequency Response Function matrix that describes motion over the interface. This matrix could be called $\boldsymbol{Y}_{31}^{AR}$, where the overload of impacts represent excitations in position 1 (somewhere on the active source) and the measurements happen at position 3 (somewhere on the passive side, either close or far away from the interface). This matrix shows more similarity to a matrix with pseudo forces as suggested by Janssen [31], so the correct subscript should be $3ps$. To minimise confusion and make the derivations in following sections more intuitive, the matrix will be dubbed with subscript for this chapter 31.

For every impact added to $\boldsymbol{Y}_{31}^{AR}$, the matrix will gain an extra row. This extra information will not necessarily increase the rank of the matrix or the size of the left nullspace as the impact might excite exactly the same modes as a previous impact. In practice however, every impact will lead to an increase in the rank of the matrix and increase the row space. This is due to flexibility caused by every impact that is very much dependant on the positioning of that impact. So maybe a large portion of the motion added by the impact can be described by modes already present in the matrix, but the flexible modes

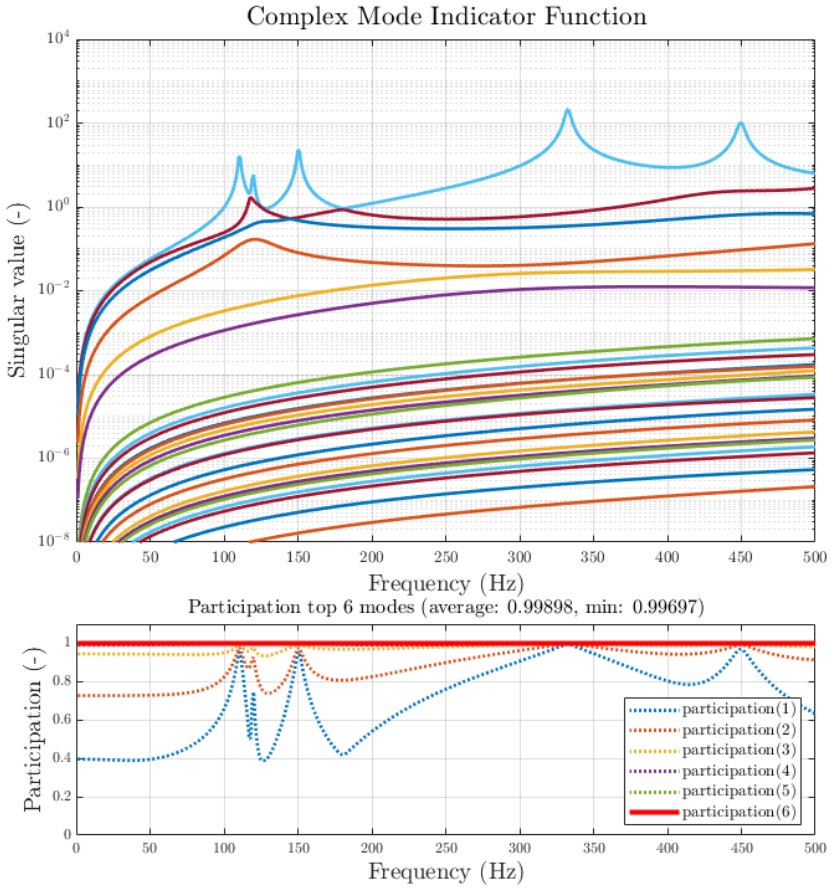

Figure 4.2: Participation real-life interface

will most likely not be. To ensure saturation of the interface, the number of impacts is increased until some converging behaviour is observed.

For adding extra sensors the explanation is identical to the story of the impacts, only this time an extra sensor will cause 3 extra columns in the system matrix due to the tri-axial nature of each sensor. Additionally, the information from these sensor measurements is contained in the left singular vectors, whereas the information for the impacts is contained in the right singular vectors.

Figure 4.2 shows the singular values of $\boldsymbol{Y}_{31}^{AR}$ and its participation for the first 6 modes for the experiment where 40 impacts were located on the source and 9 sensors (meaning $9 \times 3 = 27$ sensor channels) were placed on the receiving side. As you can see if you can carefully count the number of singular values, there is a total of 27 singular values. This is because the rank of the matrix is 27, irregardless of the amount of impacts. This is as expected, since the rank of the matrix will always be limited by either the Degrees of Freedom of the impacts or by the Degrees of Freedom of the sensor channels. The behaviour of the matrix for such a load case and sensor set-up converges and adding extra information, either by impacts or more sensor channels, does not change the system significantly.

The bottom plot of figure 4.2 shows the participation as we have defined it earlier. At first glance it seems the first 6 modes of $\boldsymbol{Y}_{31}^{AR}$ are describing the full behaviour. From the participation of the first few modes it can be seen that the first 3 modes actually already describe the behaviour quite well. The participation for the first 6 modes has a maximum mismatch to the full problem of 0.3%. As described earlier, this mismatch is due to flexibility in the structure. The question arises if this mismatch is significant enough for it to be not a 6 Degree of Freedom problem.

To answer this question we must remember that the behaviour of $\boldsymbol{Y}_{31}^{AR}$ describes the behaviour of dynamics passing over the interface of the structure. It does not take into account any specific excitation that is occurring, it only describes the dynamics that the system passively is able to describe. The participation is fully independent of the load case of the actual source. With this in mind, you could think of specific source excitations that will only excite the structure in such a manner that it lies outside of the rank 6 representation of $\boldsymbol{Y}_{31}^{AR}$. Realistically speaking, this will likely not be the case. For this reason we can indeed conclude that the effective rank of the interface is 6.

4.2.2. Modelled Interface

We can do a similar analysis as seen for the real-life interface for the modelled interface. The modelled interface, or discretised interface as you will, is based on the Virtual Point Transformation. We have modelled the interface to be 6 Degrees of Freedom. Doing exactly the same analysis as we did for the real-life interface does not make sense, since we enforce a rank of the interface in the way we model it. The participation of the first 6 modes will by definition describe the problem fully. Nevertheless the behaviour of the individual modes as can be seen in the figures gives a lot of insight.

There are a few possibilities to analyse the dynamics going across the discretised interface. To choose which method has our preference, we remember the reason that this experiment is performed lies in source characterisation. Therefore we are interested in an expression of how much of the measured response we can describe with a 6 Degree of Freedom force on the interface.

We will rewrite $\boldsymbol{Y}_{31}^{AR}$ in terms of the modelled interface. The derivation is based on the Round-Way Trip as originally proposed by Moorhouse and Elliot [38]. It was developed as an alternative way to obtain interface Frequency Response Functions, but we will see that we can use their way of reasoning to describe the modelled interface.

The derivation makes use of the equivalent source principle as seen in section 2.3. We want to rewrite the original source excitation $\boldsymbol{f}_1$ in terms of the equivalent force $\boldsymbol{f}_2^{eq}$. This mapping is a transmissibility transformation. This method is are normally used to identify the dominant paths in the assembled structure. Equation 4.3 shows the equivalent source principle and the transmissibility transformation

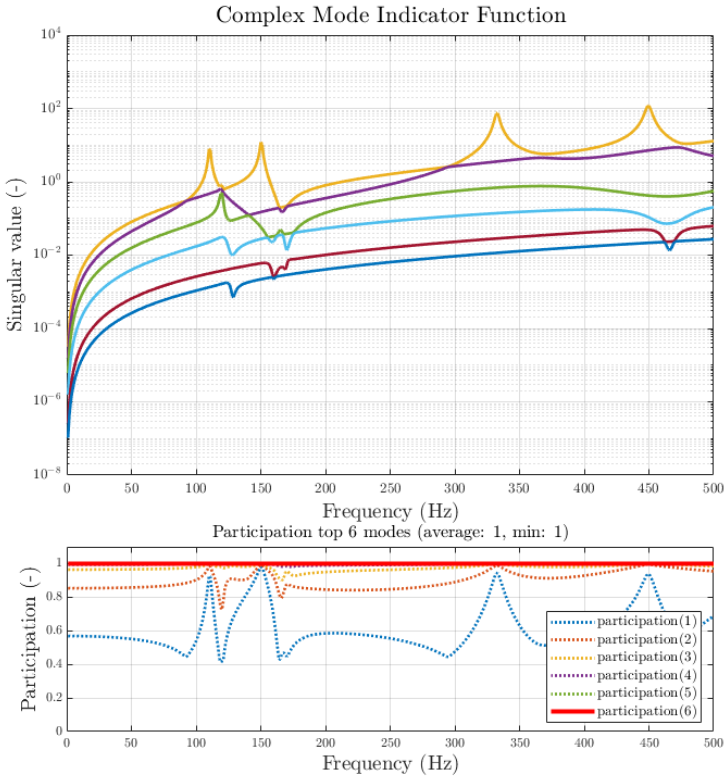

Figure 4.3: $\boldsymbol{Y}_{42}$

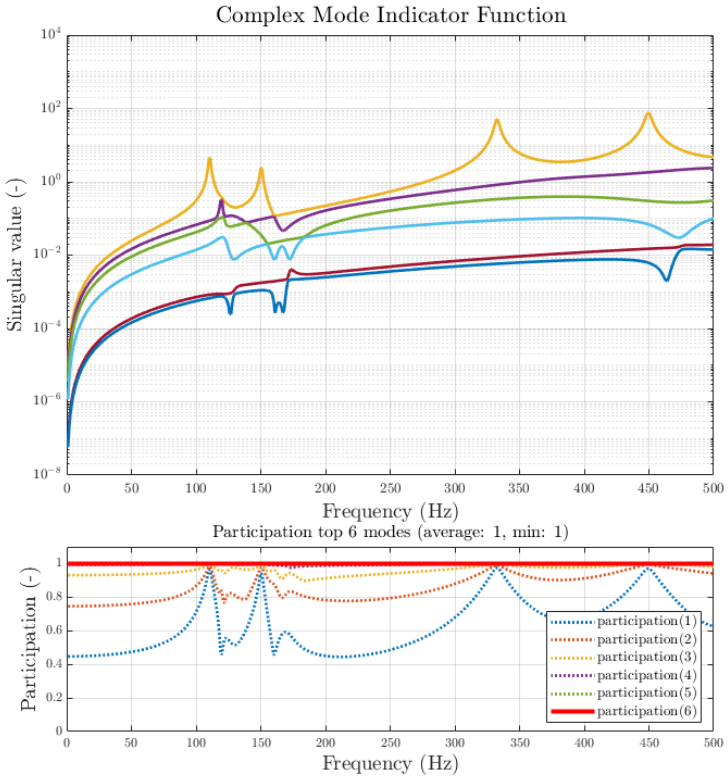

Figure 4.4: $\boldsymbol{Y}_{32}$

Figure 4.5: Singular Values for Discretised Interface (ph)

that follows from it.

$$\begin{aligned} \boldsymbol{u}_3 &= \boldsymbol{Y}_{32}\boldsymbol{f}_2^{eq} \\ \boldsymbol{u}_3 &= \boldsymbol{Y}_{31}\boldsymbol{f}_1 \end{aligned} \qquad \rightarrow \qquad \boldsymbol{f}_1 = \underbrace{(\boldsymbol{Y}_{31})^+ \boldsymbol{Y}_{32}}_{\boldsymbol{T}_{\boldsymbol{f}_1\boldsymbol{f}_2}} \boldsymbol{f}_2^{eq} \tag{4.3}$$

The result from equation 4.3 can be used to determine response $\boldsymbol{u}_4$ as seen in equation 4.4. Here we make use of the equivalent source principle again, but now we substitute the equivalent force found in equation 4.3. This will lead to an alternative way to determine Frequency Response Function of the information passing the discretised interface $\boldsymbol{Y}_{42}$.

$$\begin{aligned} \boldsymbol{u}_4 &= \boldsymbol{Y}_{42}\boldsymbol{f}_2^{eq} \\ \boldsymbol{u}_4 &= \boldsymbol{Y}_{41}\boldsymbol{f}_1 \end{aligned} \qquad \rightarrow \qquad \boldsymbol{u}_4 = \underbrace{\boldsymbol{Y}_{41}(\boldsymbol{Y}_{31})^+ \boldsymbol{Y}_{32}}_{\boldsymbol{Y}_{42}} \boldsymbol{f}_2^{eq} \tag{4.4}$$

The result from equation 4.4 is used to analyse the participation. Note that this exact derivation is also possible in terms of $\boldsymbol{u}_3$. For now we have not quantified the differences between the $\boldsymbol{u}_3$ and $\boldsymbol{u}_4$ space, so we will analyse both. Figure 4.5 shows both the singular values and participation plot for $\boldsymbol{Y}_{42}^{AR}$ and $\boldsymbol{Y}_{32}^{AR}$. Note that the subscript 2 indicated the interface force.

When comparing the dynamics of the real-life interface in figure 4.2 and the modelled interface in figure 4.5, it seems they show a lot of similar behaviour. The resonance frequencies of the largest mode is identical and the peak value is in the same order of magnitude. On top of that, the participation of the first 6 modes shows similar behaviour for both cases. Since we are only looking at the absolute size of each mode and not the shape of their left and right singular vector, we cannot exactly compare them, but this rough analysis seems to enforce the opinion that this discretisation of the interface is valid.

A big difference in the real-life interface and the modelled interface can be seen at certain frequencies however. If you look around $f \approx 125$Hz, $f \approx 160$Hz and $f \approx 460$Hz, the modelled interface seems to show some anti-resonance behaviour while the real-life interface shows nothing of that. These frequencies are actually not that random, as they are the eigenfrequencies of the source for a perfectly clamped interface. In other words, the eigenfrequencies of the source when all 6 Degrees of Freedom on the interface are fixed.

4.3. Interface Completeness

Now that we have a description for both the real-life and discretised interface, we will think of a method to compare them. This will be done by defining a blockedness of the interface. The analogy is that for an equivalent force $\boldsymbol{f}_2^{eq}$ that is truly equivalent to $\boldsymbol{f}_1$, the responses must be the same. Similarly, if the equivalent force is excited in the opposite direction, all motion must cancel out. For this reason, an actual equivalent measure must be able to block the interface and everything behind it.

Mathematically this means that we are interested in the motion $\boldsymbol{u}_2$. Since this motion is a reduced Hilbert Space from $\boldsymbol{u}_4$, we can work directly with the actual measurement data. In theory these two measures are equivalent, but in practice there is always some minor deviation. For that reason we define the blockedness of the interface as the residual motion of $\boldsymbol{u}_4$ when the equivalent force is applied in negative. This can be seen in equation 4.5.

$$\|\boldsymbol{u}_4^{res}\| = \|\boldsymbol{u}_4^{meas} - \boldsymbol{u}_4^{pred}\| \qquad \text{with } \boldsymbol{u}_4^{pred} = \boldsymbol{Y}_{42}^{AR}\boldsymbol{f}_2^{eq} \tag{4.5}$$

To analyse the blockedness of the interface we will first introduce an analogy to explain the issue at hand. With this analogy we will proof that non-equivalency indeed causes issues for the prediction, but also that this issue is a function of the test bench. This has consequences for transferability.

This analogy was originally suggested by Meggitt in [36] in a slightly different form to verify the choice of how many Degrees of Freedom the interface is modelled as. The similarity measure that follows from his publication was dubbed the Interface Completeness Criterion. We will change the mathematical definition slightly, but stick to the name that explains the problem quite well.

Analogy Introduction

By definition the dynamics of the problem are fully captured by equation 4.6. It consists of the actual measured data from the sensors. In practice the system matrix is not known, but with the reasoning as seen in the real-life interface experiment in section 4.2.1, we can find a basis of the matrix that describes the mapping from $\boldsymbol{f}_1$ to $\boldsymbol{u}_4$.

$$\boldsymbol{u}_4 = \boldsymbol{Y}_{41}^{AR}\,\boldsymbol{f}_1 \tag{4.6}$$

In our analogy we will split up the interface from the Source-Transmission-Receiver model into two parts. The first part is the interface that is known and described. This is the interface as we have been working with so far and is described for a response at $\boldsymbol{u}_4$ by $\boldsymbol{Y}_{42}\boldsymbol{f}_2$. The other interface describes the behaviour that we are not capturing with the way we modelled the interface. These measures are noted with a tilde accent. Equation 4.7 shows the new description for $\boldsymbol{u}_4$. Note that the superscript from the equivalent force is omitted for both the known and unknown interface force, since it is assumed that neither properties are equivalent.

$$\boldsymbol{u}_4 \quad = \quad \underbrace{\boldsymbol{Y}_{42}^{AR}\,\boldsymbol{f}_2}_{\text{known}} \quad + \quad \underbrace{\widetilde{\boldsymbol{Y}_{42}^{AR}}\,\widetilde{\boldsymbol{f}_2}}_{\text{unknown}} \tag{4.7}$$

With the expansion of the interface in a known and unknown part, we can solve for the equivalent force in the In-Situ way. Note that the inversion is done for the Frequency Response Function that is known. Equation 4.8 shows the result of solving the inverse problem by pre-multiplying equation 4.7 with $\left(\boldsymbol{Y}_{42}^{AR}\right)^{+}$.

$$\left(\boldsymbol{Y}_{42}^{AR}\right)^{+}\boldsymbol{u}_4 \quad = \quad \boldsymbol{f}_2^{eq} \quad = \quad \boldsymbol{f}_2 + \left(\boldsymbol{Y}_{42}^{AR}\right)^{+}\widetilde{\boldsymbol{Y}_{42}^{AR}}\,\widetilde{\boldsymbol{f}_2} \tag{4.8}$$

It already seems clear from equation 4.8 that the equivalent force consists of a term that we can describe, and a term that we cannot describe. If this second term has any significant size, the characterisation will be both wrong for this structure, but especially for any other receiving structure. To proof this last statement, we apply the acquired equivalent force from structure AR to a new receiving structure AB as seen below in equation 4.9.

$$\begin{aligned}\boldsymbol{u}_4^{AB} &= \boldsymbol{Y}_{42}^{AB}\,\boldsymbol{f}_2^{eq} \\ &= \boldsymbol{Y}_{42}^{AB}\,\boldsymbol{f}_2 \quad + \quad \boldsymbol{Y}_{42}^{AB}\left(\boldsymbol{Y}_{42}^{AR}\right)^{+}\widetilde{\boldsymbol{Y}_{42}^{AR}}\,\widetilde{\boldsymbol{f}_2}\end{aligned} \tag{4.9}$$

Analogy Experiment

From equation 4.9 it is clear that the source characterisation will not be successful if the interface does not describe the full problem. To take this logic a step further, we will find a description of the motion behind the interface when the equivalent force is applied in negative. In this scenario the interface for which the description is known is fully blocked, meaning all known Degrees of Freedom are constrained. This scenario can be seen in figure 4.6 and its mathematical description is seen in equation 4.10.

$$\begin{bmatrix}\boldsymbol{u}_1\\ \boldsymbol{0}\\ \widetilde{\boldsymbol{u}_2}\\ \boldsymbol{u}_3\end{bmatrix} = \begin{bmatrix}\boldsymbol{Y}_{11} & \boldsymbol{Y}_{12} & \widetilde{\boldsymbol{Y}_{12}} & \boldsymbol{Y}_{13}\\ \boldsymbol{Y}_{21} & \boldsymbol{Y}_{22} & \widetilde{\boldsymbol{Y}_{22}} & \boldsymbol{Y}_{23}\\ \widetilde{\boldsymbol{Y}_{21}} & \widetilde{\boldsymbol{Y}_{22}} & \widetilde{\boldsymbol{Y}_{22}} & \widetilde{\boldsymbol{Y}_{23}}\\ \boldsymbol{Y}_{31} & \boldsymbol{Y}_{32} & \widetilde{\boldsymbol{Y}_{32}} & \boldsymbol{Y}_{33}\end{bmatrix}\begin{bmatrix}\boldsymbol{f}_1\\ \boldsymbol{f}_2\\ \boldsymbol{0}\\ \boldsymbol{0}\end{bmatrix} \tag{4.10}$$

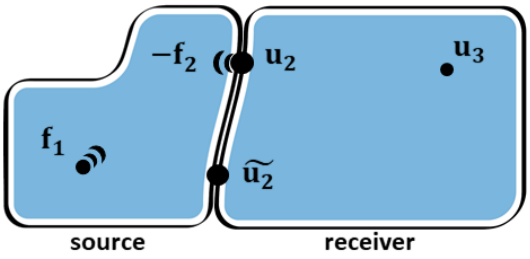

Figure 4.6: Known and Unknown Transmission

From equation 4.10 we can derive an expression of the motion behind the interface for when the discretised interface is fully blocked. In this case the interface $\boldsymbol{u}_2$ is indeed $\mathbf{0}$. The only forces acting on the system are the excitation force $\boldsymbol{f}_1$ and the known part of the equivalent force $\boldsymbol{f}_2$. We assume the interface is partly constrained and that the force on the interface that we have not described is zero ($\widetilde{\boldsymbol{f}_2} = \boldsymbol{f}_3 = \mathbf{0}$).

To verify this scheme is correct, we take line of equation 4.10 and rewrite it to obtain an expression of the equivalent force for the known interface. Equation 4.11 shows this expression and it shows that the force is indeed in the form as derived in section 2.3, but with a minus sign signifying we are blocking the interface.

$$\begin{aligned} \mathbf{0} &= \boldsymbol{Y}_{\mathbf{21}}\, \boldsymbol{f}_1 \quad + \quad \boldsymbol{Y}_{\mathbf{22}}\, \boldsymbol{f}_2 \\ \boldsymbol{f}_2 &= -\left(\boldsymbol{Y}_{\mathbf{22}}\right)^{-1} \boldsymbol{Y}_{\mathbf{21}}\, \boldsymbol{f}_1 \end{aligned} \tag{4.11}$$

If we now look at the last line of equation 4.10 we can find an expression of the motion behind the interface while the known interface is constrained. This means that all dynamics passes through the interface that we have not described. The result is seen in equation 4.12.

$$\widetilde{\boldsymbol{u}_{\mathbf{3}}} = \boldsymbol{Y}_{\mathbf{31}}\, \boldsymbol{f}_1 \quad + \quad \boldsymbol{Y}_{\mathbf{32}}\, \boldsymbol{f}_2 \tag{4.12}$$

Plugging the result of equation 4.11 into equation 4.12 leads to:

$$\begin{aligned} \widetilde{\boldsymbol{u}_{\mathbf{3}}} &= \boldsymbol{Y}_{\mathbf{31}}\, \boldsymbol{f}_1 \quad - \quad \boldsymbol{Y}_{\mathbf{32}}\, \left(\boldsymbol{Y}_{\mathbf{22}}\right)^{-1} \boldsymbol{Y}_{\mathbf{21}}\, \boldsymbol{f}_1 \\ &= \left[\boldsymbol{Y}_{\mathbf{31}} \quad - \quad \boldsymbol{Y}_{\mathbf{32}}\, \left(\boldsymbol{Y}_{\mathbf{22}}\right)^{-1} \boldsymbol{Y}_{\mathbf{21}}\right] \boldsymbol{f}_1 \end{aligned} \tag{4.13}$$

The bracketed mobility term of equation 4.13 is therefore admittance across the interface of the assembly where the known Degrees of Freedom are constrained. We can interpret this equation as seen in equation 4.14, where we recognise the individual admittance terms. The LHS is the admittance across the interface for the case where the unknown DoFs are unconstrained. The RHS consists of the full unconstrained case, and the case here the known DoFs are unconstrained respectively. Also note that this last term is based on the round-way-trip.

$$\underbrace{\widetilde{\boldsymbol{Y}_{31}}}_{\text{unknown DoFs unconstrained}} = \underbrace{\boldsymbol{Y}_{\mathbf{31}}}_{\text{all DoFs unconstrained}} - \underbrace{\boldsymbol{Y}_{32}\, \left(\boldsymbol{Y}_{22}\right)^{-1} \boldsymbol{Y}_{21}}_{\text{known DoFs unconstrained}} \tag{4.14}$$

To make it easier to calculate with these quantities, new symbols are defined for all three receptances as seen in equation 4.15. Here it becomes clear that the full dynamics over the interface are described by the Frequency Response Function created by the actual measurement, which is equal to the sum of our known modelled interface and a the unknown modelled interface.

$$\begin{aligned} \text{unknown DoFs unconstrained: } \boldsymbol{Y}_{31}^{c_j} &= \widetilde{\boldsymbol{Y}_{31}} \\ \text{all DoFs unconstrained: } \boldsymbol{Y}_{31}^{c} &= \boldsymbol{Y}_{31} \\ \text{known DoFs unconstrained: } \boldsymbol{Y}_{31}^{c_i} &= \boldsymbol{Y}_{32}\, \left(\boldsymbol{Y}_{22}\right)^{-1} \boldsymbol{Y}_{21} \end{aligned} \tag{4.15}$$

Interface Completeness Criterion

This small example makes it clear that we can express the difference between the real-life interface and the modelled interface in terms of a combination of Frequency Response Functions. Meggitt proposed to define a similarity measure based on the Modal Assurance Criterion (MAC) to compare these two measures [36], as can be seen in equation 4.11. This expressions yields a frequency dependent scalar value bound between 0 and 1, where 1 is a full complete interface description and 0 a total incomplete interface description.

$$\text{ICC} = \frac{\| \left(\boldsymbol{Y}_{31}^{c}\right)^{H} \boldsymbol{Y}_{31}^{c_i} \|^2}{\left(\boldsymbol{Y}_{31}^{c}\right)^{H} \boldsymbol{Y}_{31}^{c} \left(\boldsymbol{Y}_{31}^{c_i}\right)^{H} \boldsymbol{Y}_{31}^{c_i}} \tag{4.16}$$

The Interface Completeness Criterion can be adapted in two places to make more sense for the Transfer Path Analysis problem. Firstly, in the current Modal Assurance Criterion (MAC) form, it purely

evaluates the similarity in the direction of each matrix since the MAC is invariant of amplitude differences. Although this a valid way of using the MAC according to Allemang [2], it can be adapted to a more general coherence similarity function so that the amplitude is also weighted in the comparison. The similarity function we will use is as suggested by Van der Seijs from VIBES.technology which is an in-house used function to define coherence where the phase and magnitude can be weighted in their importance. This method is dubbed the log coherence function by VIBES.technology. Coherence is used to define the similarity between two values, taking into account their magnitude and phase separately [4].

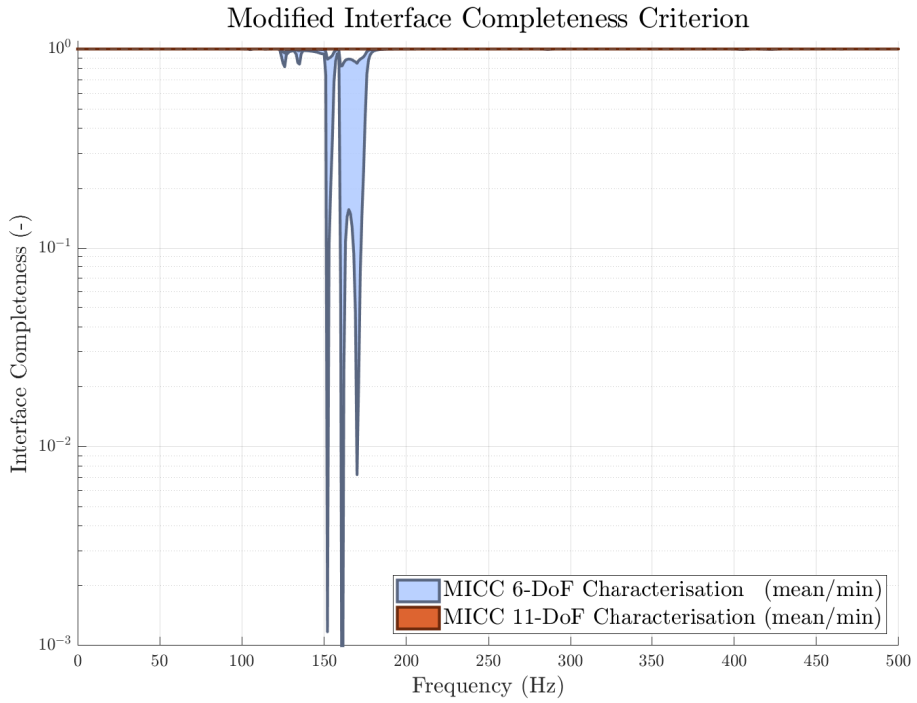

Figure 4.7: Interface Completeness Criterion

As a second proposed change to the ICC, we rewrite the known interface model $\boldsymbol{Y}_{31}^{c_i}$ in terms of the round-way trip similar as seen in equation 4.4, making it a function of the $\boldsymbol{u}_4$-vector space. This makes more sense since the equivalent force is deduced from this space. To make a clear distinction between the old and new formulation, the criterion is called the Modified Interface Completeness Criterion (MICC). The definition of MICC is seen in equation 4.17.

$$\text{MICC} = \log \text{coherence}\left(\boldsymbol{Y}_{41}, \boldsymbol{Y}_{42}\left(\boldsymbol{Y}_{42}\right)^{+}\boldsymbol{Y}_{41}\right) \tag{4.17}$$

If we apply the MICC to the system matrices of section 4.2, we can quantify the completeness of the interface. The excitation case that we chose is a single impact on the source angled with respect to all three axes to get a complicated load case where many modes are excited. In figure 4.7 in blue we see the completeness for a 6 Degree of Freedom force on the interface. We plot both the mean and minimum value of the completeness per frequency bin to get some insight in the severity of the incompleteness.

It is clear that at certain frequency bins the interface is not described fully with a 6 Degree of Freedom interface force. The frequencies were we see the mismatch are the same frequencies that we saw the differences between the singular values of $\boldsymbol{Y}_{41}$ and $\boldsymbol{Y}_{42}$ in section 4.2, which were the eigenfrequencies for a fixed interface.

In the next section we will start working towards an explanation of why these frequencies are the troublesome frequencies and what kind of motion occurs here. To make most use of figure 4.7 however, we also plot the MICC results for an 11-DoF interface, where the traditional interface force is extended with 5 pseudo-impacts placed on the source away from the interface. As seen in red, this addition makes the interface description complete.

4.4. Interface Force Limitation

We have seen that the modelled interface with 6 Degrees of Freedom does not describe the $\boldsymbol{u}_4$-space to the same extend as the real-life interface. Although we have not defined the effect of this on the characterisation, we can say the mismatch originates from a lack of controllability. In other words, the equivalent force cannot excite the entire $\boldsymbol{u}_4$-space. This problem only arises at the frequencies for which the source would have its fixed interface eigenfrequencies. These frequencies came to the surface during the effective rank research, but with the Modified Interface Completeness Criterion we can conclude with certainty that these are the problem frequencies.

The fixed interface eigenfrequencies of a structure is not dependent on the structure behind the interface. This means that the interface limitation will be present regardless of the receiving structure. To validate this some alternative receiving structures were modelled and added to the numerical model. There are two alternative structures which can both be seen in appendix D. One structure is designed to show very similar dynamics while the other is a very compliant design. Both structures indeed show the same controllability issues as the original case.

This conclusion is both good and bad news for the method of Source Characterisation with a 6 Degree of Freedom interface force. The good aspect of this lack of controllability is that it is a function of the active source, specifically the passive dynamics. These dynamics will be the same in any receiving structure, meaning the non-equivalency is purely a function of the active source. The bad news is that the issue seems to arise from an inherent problem with an interface force. This issue is not just a problem for source characterisation, but for every method making use of Virtual Points transformation, as for example the coupling of substructures in Dynamic Substructuring.

In the following section we will dive deeper into these frequencies and will try to come up with a physical explanation that an interface force cannot excite the entire $\boldsymbol{u}_4$-space. It does not seem intuitive that frequencies which describe the fixed interface dynamics, a totally different dynamic problem then what we are working with, has such an effect on our problem.

The physical explanation will start by explaining where this limitation originates from. By comparing the problem to a very simple other dynamic problem, we can reason a physical explanation. As a closing remark to this chapter, we will take a closer look at the Frequency Response Function of the interface itself to mathematically verify the reasoning of the physical explanation. We are looking at the driving point function $\boldsymbol{Y}_{22}$ specifically and will analyse how changing the receiving structure indeed does not change the frequencies for which the inherent non-equivalency occurs. If this holds true for the general case, it provides a valid proof that the fixed interface eigenfrequencies are the troublesome frequency bins.

4.4.1. Mode Excitation

We seek an intuitive explanation for the lack of controllability with an interface force. There are two aspects to explain. How come energy is flowing through the interface when the interface motion is zero and how come this occurs at the frequencies at which is does. To explain both phenomena, we take a look at a guitar string.

Figure 4.8 shows a schematic of a guitar string. This guitar string is a Linear Time-Invariant system, and here depicted as a planar problem where the string only has vertical motion. The string is excited by the player by plugging the string on the right, depicted with a finer to imply a force excitation. This

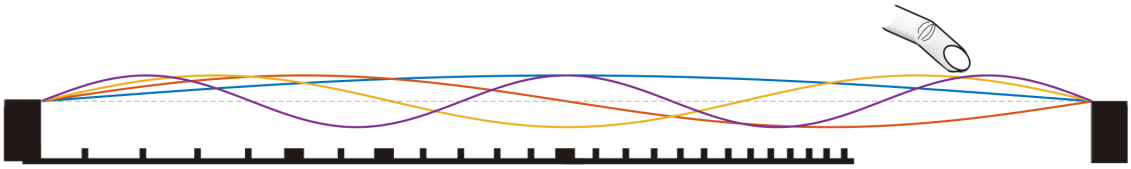

Figure 4.8: Guitar String Excitation

excitation can be seen as a broadband impulse that excites a very wide frequency range. The string has its own Frequency Response Function due to the length and tension of the string. This causes only the frequencies which operate well in this system to resonate, while the other frequencies are damped out almost immediately. The first 4 resonance frequencies can be seen.

Now imagine that the guitar string is an assembly of two substructures where the interface lies exactly at $1/3^{\text{rd}}$ of the length of the string, above the 7^{th} fret. The right side is where we excite and the left side is now where we sense the motion. In the fashion of source characterisation we want to describe the source excitation with a new force at the interface. This force has to lead to the same motion behind the interface.

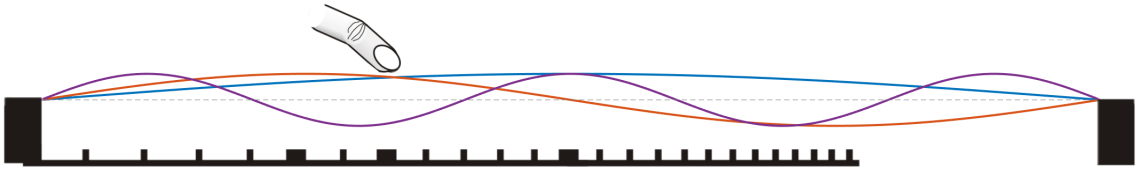

Figure 4.9: Guitar String Interface Force Excitation

Figure 4.9 shows the equivalent force found and its response on the string. The force excitation now occurs at the interface, as can be seen by the finger placement. The responses are drawn in the same colour as in figure 4.8. The equivalent force can excite all modes, except for the yellow mode. This yellow mode has a node exactly at the interface. A force from the nodal point of a mode can never excite that mode. This is a vital conclusion in understanding the limitations of using an interface force. The controllability issue stems from the choice of location of the equivalent force. Since they are all located on the same Virtual Point, it becomes a possibility that certain modes have no amplitude at that position and hence cannot be excited.

This explains why there is a lack of controllability, but does not yet explain why this occurs at the fixed interface eigenfrequency. To explain this, we excite the string with both the original excitation as well as the equivalent force on the interface, but with the equivalent force in negative to block the interface. Figure 4.10 shows that for both excitations simultaneous, the interface is standing still. However, although most modes indeed cancel out, the mode which we did not capture in the equivalent force is not cancelled out. There is motion behind the interface while the interface is standing still. And as expected, figure 4.10 shows that the mode shape in the residual is the first resonant frequency of the interface if its vertical Degree of Freedom would be fixed.

Although the guitar string analogy is a one-dimensional problem, it translates well to the more complex real-life situation in 3D-space. At the fixed interface eigenfrequencies the preferred mode shape of the source structure is one where the interface is standing still, yet there is still energy flowing through the interface to have a continuous energy flow. This causes motion behind the interface while the interface is still. During the source characterisation this mode can simply not be represented with only a force on the interface.

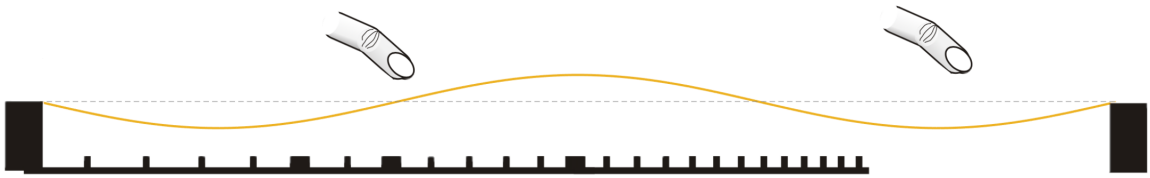

Figure 4.10: Guitar String Blocked Interface

4.4.2. Driving Point Anti-Resonances

The conclusions drawn in the previous section are based on reasoning. The controllability issue is made intuitive by expressing it in terms of a common 1D problem. To verify this logic, a mathematical approach is chosen. In the following section we will analyse the driving point Frequency Response Function of the interface $\boldsymbol{Y}_{22}$. The driving point is the colocated Response Function where the force and motion are at the same location. We will see how the anti-resonances of this local function say something about the global dynamics of the source characterisation.

Anti-resonances of a driving point Frequency Response Function can be interpreted as the resonance frequencies of the system fixed at the excitation points in the excitation direction [23]. As a consequence, the anti-resonances will stay unchanged for added mass or stiffness at the driving point in the excitation direction [57]. This is as expected, since the fixed interface eigenfrequencies of a structure do not change if mass or stiffness is added to the interface. The structure does not feel these changes as the interface is fixed. We can verify this by adding mass and stiffness to the structure and observing the change in zeros of the function, as these will describe the anti-resonances.

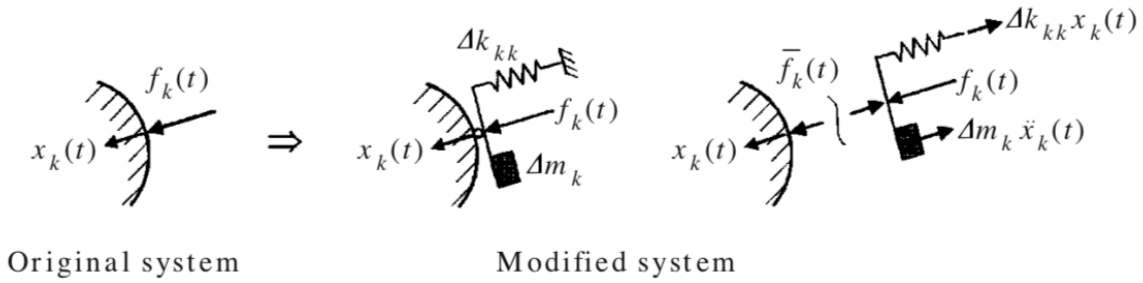

Figure 4.11: Mass and stiffness difference (from [57])

There are multiple ways to go about showing this effect. We want to derive a new Frequency Response Function matrix with extra mass and stiffness at the interface and compare this to the original system. To make the derivation easier, we choose to compare the cases as seen in equation 4.18. Here we denote the altered system with a hat-accent and the corresponding force without any accent.

$$\begin{aligned} \text{Modified system}: \quad \boldsymbol{u}_2 &= \bar{\boldsymbol{Y}}_{22}\boldsymbol{f}_2 \\ \text{Original system}: \quad \boldsymbol{u}_2 &= \boldsymbol{Y}_{22}\bar{\boldsymbol{f}}_2 \end{aligned} \tag{4.18}$$

The force felt by the original system is different compared to the altered system. The effective excitation force $\bar{\boldsymbol{f}}_2$ is smaller due to the added mass and stiffness. We can make use of the linearised Mass and Stiffness matrix to account for these changes. Equation 4.19 shows the change in interface force. The derivation originates from the physical domain, but is depicted here as its frequency domain counterpart.

$$\bar{\boldsymbol{f}}_2 = \boldsymbol{f}_2 - \left[-\omega^2 \Delta\boldsymbol{M} + \Delta\boldsymbol{K}\right]\boldsymbol{u}_2 \tag{4.19}$$

We now seek an expression for the modified system matrix in terms of the change in mass and stiffness at the interface. This derivation is performed by Wahl [57] and is based on using the inherent reciprocity of a co-located Frequency Response Function. Wahl shows that the modified system matrix can be

rewritten in terms of the original system matrix with an additional factor for the change in force. This leads to the final modified system matrix representation as seen in 4.20.

$$\bar{\boldsymbol{Y}}_{ij} = \frac{\boldsymbol{Y}_{ij}}{1 + \boldsymbol{Y}_{jj}\Delta\boldsymbol{b}_{kk}} \quad \text{with } \Delta\boldsymbol{b}_{kk} = \left(-\omega^2\Delta\boldsymbol{m}_k + \Delta\boldsymbol{k}_{kk}\right) \tag{4.20}$$

Equation 4.20 shows that additional mass or stiffness on the excitation direction has an effect on the poles of the system, but not on the zeros. For undamped or lightly damped structures the anti-resonances are determined by the zeros. This effect of the zero shifting happens for any Frequency Response Function where additional mass and stiffness is added at the excitation point, and not just for the co-located function.

When looking specific at the co-located Frequency Response Function as seen in equation 4.21, this same conclusion is made. The mass and stiffness loading indeed has no effect on the anti-resonances. From this we can conclude that the resonance frequencies of the fixed interface source stay the same, no matter what the source is connected to.

$$\bar{\boldsymbol{Y}}_{22} = \frac{\boldsymbol{Y}_{22}}{1 + \boldsymbol{Y}_{22}\Delta\boldsymbol{b}_{22}} \quad \text{with } \Delta\boldsymbol{b}_{22} = \left(-\omega^2\Delta\boldsymbol{m}_2 + \Delta\boldsymbol{k}_{22}\right) \tag{4.21}$$

4.5. Conclusion

In this chapter the source characterisation problem has been analysed from the standpoint of the interface. The analysis started with verifying whether the 6 Degrees of Freedom interface force for source characterisation is a valid option to characterise the vibrating source. This was done by defining an effective rank of the interface and see how much dynamics can possibly pass. We concluded that the effective rank of the interface is indeed 6.

This analysis was extended to verify whether the Frequency Response Functions that describe the real-life interface and the modelled interface are indeed describing the same physical behaviour. To do this a clear system matrix representation of the modelled interface is created and this is tested using the Modified Interface Completeness Criterion. This Criterion is a statistical measure that compares the matrices in terms of mode shape, amplitude and phase.

From the results of the Modified Interface Completeness Criterion we can distill the troublesome frequencies of the modelled interface. It turns out the frequency bins where the interface can be regarded as incomplete are the frequency bins where the fixed interface eigenfrequencies of the source occur. These frequencies are a function of the source only, so to proof the incompleteness is a function of the source only, the experiment was verified by changing the receiving structure. For these the same results were obtained leading to the conclusion that the fixed interface eigenfrequencies of a source are troublesome to model using the existing interface methods.

Because this limitation did not feel intuitive, we set out to find a physical explanation of the problem occurring. This way done by analogy of a guitar string. From this it became clear that the position of the equivalent force has a big role in what possible modes it can excite. We concluded that due to the fact that all equivalent forces are positioned on the same virtual point, the mode shapes that have a node at this same location can by definition not be excited.

As a final confirmation to the analogy the interface limitation was also expressed in terms of the Frequency Response Function of the interface. The receiving structure was changed in mass and stiffness to observe the effect of the fixed interface eigenfrequencies. Here it was verified that the receiving structure has no influence on the frequency bins for which this problem occurs.

5

Equivalent Force

5.1. Introduction

The following chapter will translate the knowledge obtained from the previous chapter to its consequences on the equivalent force for source characterisation. It was seen in figure 4.7 that a 6 Degree of Freedom interface force is not complete, it has some limitations. Whereas the Modified Interface Completeness Criterion is in terms of Frequency Response Functions which are often not obtainable, we seek a way to describe the same interface limitations and their effect on the final characterisation in a quantified manner.

As a first step we will introduce quantified measures for controllability. This controllability is defined as to what extend the equivalent force $\boldsymbol{f}_2^{eq}$ can excite the $\boldsymbol{u}_4$-vector space. In other words, these controllability measures quantify the percentage of possible mode shapes that can be controlled. Two different measures will be defined here. The first one is a controllability measure that defines the specific controllability for the obtained measurement. The other one is a global one that defines the overall controllability limitations of the equivalent force for the entire $\boldsymbol{u}_4$-vector space.

Once the controllability measures are defined, we will work towards a way of expressing the lack of controllability in terms of the effect on the source characterisation. We will analyse measures to quantify the mismatch of the characterisation based on the difference of the measurement and the prediction. The reason this is done is to quantify whether the controllability issues we have uncovered in chapter 4 and described in section 5.2 have an effect on the final outcome of the characterisation.

The mismatch measures seen so far quantify the absolute mismatch of the characterisation. We will introduce two new measures that can relate the mismatch to the original excitation. In this way a quality measure is created that expresses the mismatch as a percentage. This gives an actual quantification on how successful the characterisation is, or in other words, how equivalent the description is.

5.2. Controllability

In the following section we will define two new measures for controllability of the equivalent force $\boldsymbol{f}_2^{eq}$. The term controllability should be understood as a measure to which extent the equivalent force can excite the $\boldsymbol{u}_4$-vector space for which the characterisation is done. In other words, if there is no possible equivalent force that will lead to a certain motion in the $\boldsymbol{u}_4$-vector space, we say there is no full controllability.

We have seen that the controllability issue stems from the location of the equivalent force, specifically that they are all positioned in the same point. In the physical sense it means that certain modes cannot be excited from the interface. To which extend this lack of controllability has an effect on the characterisation depends all on the source excitation. If the source excites many of the unobtainable modes,

the characterisation is troublesome, but if none of the unobtainable modes are excited, there is no controllability issue whatsoever, even though the interface is not complete in the MICC sense.

This reasoning leads to two descriptions of controllability. We can define a Specific Controllability which describes the controllability for the specific source excitation at hand. This describes the controllability of the equivalent force for the test case at hand. On the other hand, we can also define a global controllability. Here we compare the equivalent force to all possible excitations that cross the interface in a similar manner to how it was done in section 4.2.

The power of having two kinds of controllability measures is that one can define the limitations of the entire test set-up using the Global Controllability. When the excitation is unknown or varies greatly for different use cases, it gives the lower bound of the possible controllability. When the source excitation is known and measurement data is available, an upper bound of the controllability can be defined using the Specific Controllability.

5.2.1. Specific Controllability

The Specific Controllability will be constructed to be a measure that describes to what extend the $\boldsymbol{u}_4$-vector space of the measurement can be excited in the $\boldsymbol{u}_4$-vector space by the equivalent force. Previous work by Wernsen [59] has tried to define a controllability measure in a similar fashion. This measure was based on describing the ratio of the kernel of $\boldsymbol{Y}_{42}$ compared to the full $\boldsymbol{u}_4$-vector space spanned by $\boldsymbol{Y}_{42}$. This controllability measure therefore gives an indication on what the effect of left nullspace is, but this paints the wrong picture. We have shown that the lack of controllability is not due to a non-zero left nullspace, but due to the location of the equivalent force.

To define the controllability we must look at the system matrix $\boldsymbol{Y}_{42}$. This is the right side Virtual Point transformed admittance matrix, meaning the sensor channels are unaltered, while the excitation is reduced to the 6 Degree of Freedom Virtual Point force. The matrix can be decomposed using the Singular Value Decomposition, leading to 6 orthogonal left singular eigenvectors. These vectors describe the entire $\boldsymbol{u}_4$-space that the equivalent force can excite. The notation for this vector space is in line with chapter 3 and is depicted with a capital $\boldsymbol{U}^{42}$ where each column consists of a mode shape in $\boldsymbol{u}_4$-space.

Equation 5.2 shows the definition for Specific Controllability. For every frequency bin ((ω) omitted for clarity here) we do a projection of the measurement data $\boldsymbol{u}_4$ on the i$^{\text{th}}$ left singular eigenvalue. There are multiple equivalent ways to do such a summation, especially for the normalisation step. The reason we normalise each summation with the squared norm of the projection vectors is that this leads to a

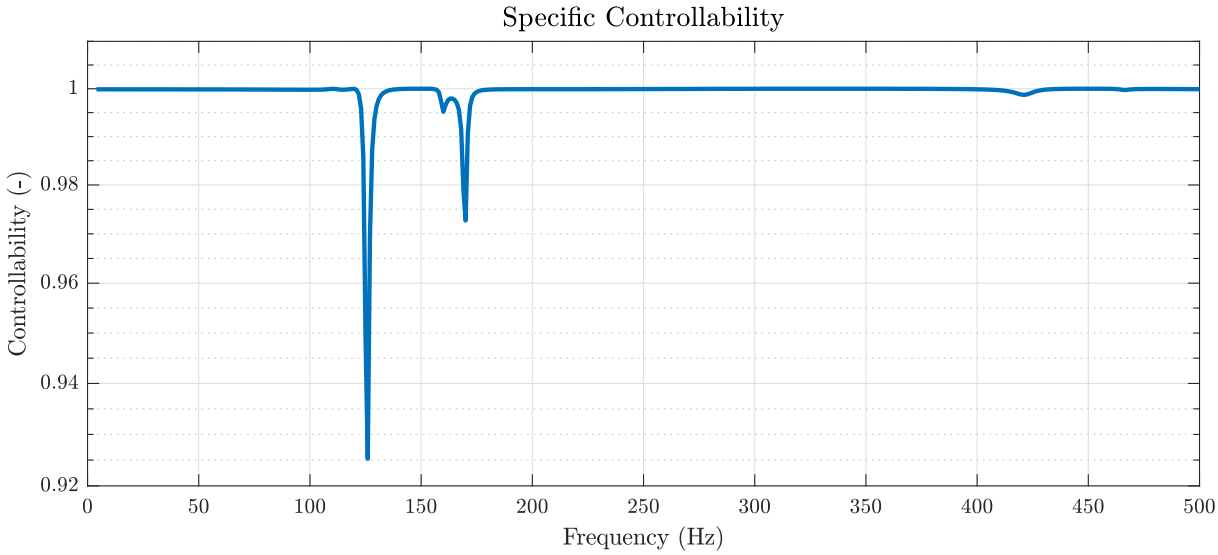

Figure 5.1: Specific Controllability

summation over all modes of maximum value 1. Additionally, in this manner it is more in line with the Global Controllability measure. As an added benefit, you could omit the summation and analyse the participation of each mode on the measurement data.

$$\text{Specific Controllability} = \rho_s = \sum_{i=1}^{r} \frac{\left(\boldsymbol{u}_4 \cdot \boldsymbol{U}_i^{42}\right)^2}{\|\boldsymbol{u}_4\|^2} \tag{5.1}$$

In the way the controllability is defined it quantifies what percentage of motion in the $\boldsymbol{u}_4$-space is not reachable for any equivalent force $\boldsymbol{f}_2^{eq}$. This means that it does not show how much of the measurement is controlled, but rather to which extend the mode shapes can be controlled. This is done by normalising with the norm of the measurement.

In practice a typical experiment consists of testing multiple different load cases. This leads to a range measurement sets that will all have their own specific controllability.
The results shown in figure 5.1 are from same load case as chosen for the MICC analysis. The Specific Controllability is unity for almost all frequency bins, but has clear drops at some frequencies. These are again the same frequencies that were the fixed interface eigenfrequencies of the source. For this load case we see the lowest controllability at around $f \approx 128$Hz. This means that the equivalent force $\boldsymbol{f}_2^{eq}$ can excite approximately 92.5% of the modes found in the measurement at this frequency bin.

5.2.2. Global Controllability

The Global Controllability is defined in a similar fashion to the Specific Controllability. It will be constructed to be a measure that describes how well the equivalent force can represent all possible mode shapes across the interface. The full dynamics are described as we have seen in section 4.2, where we overload the interface with a large number of pseudo-forces or impacts and truncate this to the first 6 modes.

The implementation can be seen in equation 5.2. Compared to the Specific Controllability, the difference lies in the fact that the left singular eigenvectors of $\boldsymbol{Y}_{42}$ are compared to the left singular eigenvalues of $\boldsymbol{Y}_{4ps}$, which are defined as the columns of $\boldsymbol{U}^{4ps}$.

$$\text{Global Controllability} = \rho_o^i = \sum_{k=1}^{r} \left(\boldsymbol{U}_i^{42} \cdot \boldsymbol{U}_k^{4ps}\right)^2 \tag{5.2}$$

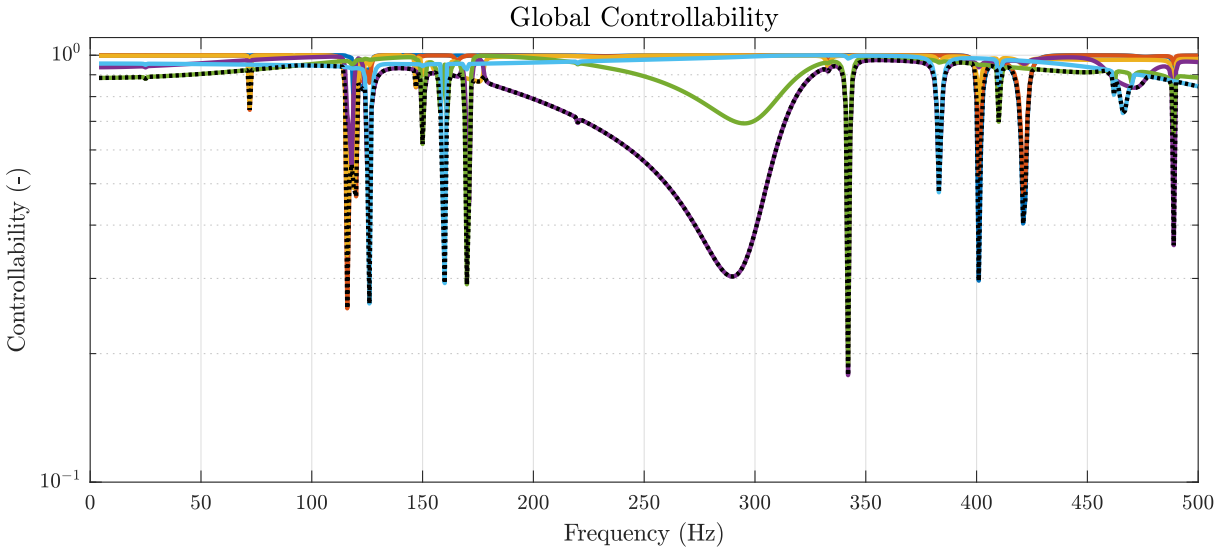

Figure 5.2: Global Controllability (ph, replace with Global)

A big assumption for this relation to be true, is that the $\boldsymbol{u}_4$-space that $\boldsymbol{Y}_{42}$ can excite is not larger than the space $\boldsymbol{Y}_{4ps}$ can excite. In other words, the lack of controllability we see here is due to motion the equivalent force cannot replicate. This opposed to an equivalent force that excites modes that cannot originate from across the interface.

This assumption should hold true by the way in which the Virtual Point Transformation is done. The description of the interface originates from a transformation of impacts around the interface on the source side. This same motion is described in $\boldsymbol{Y}_{4ps}$. Nevertheless figure 5.2 shows some unexpected peaks in lack of controllability. Some of the peaks are indeed caused by the controllability issue we have defined, but the others are the passive eigenfrequencies of the assembled structure. This begs the question if the assumption holds true.

Solution to Controllability

This thesis set out in identifying and quantifying the equivalence problems and therefore a solution to the controllability problem is outside of the scope. The origins of the problem are so clear however that two methods to fix it seem obvious. Both methods depend on extending the equivalent force with pseudo forces on the source structure. These forces can contribute to the modes that are not captured by the interface force. The method differ in the fact that the inverse problem can be solved in 1 go with a system matrix describing the response from the interface force and pseudo forces, or that the pseudo forces are only used to solve for the residual motion of the interface force.

The results of these methods can be found in appendix C. The experiment shows the methods do not improve the characterisation. Many alternative things can be tried, but since this is outside of the scope, the answer to the solution to controllability is left for now.

5.3. Non-Equivalency Quantification

In the previous section we determined the controllability for an interface force. This limitation was based on how well the equivalent force can possible represent the original motion. We now seek to translate the controllability limitation to how successful the source characterisation process is. In other words, how equivalent is the equivalent force compared to the original excitation.

The notion of equivalence is crucial here and should be defined clearer. The goal of source characterisation is to describe an active vibration source in terms of a force on the interface. In this way, one could numerically predict the motion on any receiving structure by coupling the passive dynamics of the new structure and applying the interface excitation. The equivalence can be understood as how equal the response is from the original source excitation or from our calculated equivalent force.

If the measured and predicted $\boldsymbol{u}_4$ response is exactly the same, it means the equivalent measure can predict the interface motion exactly. The possible mismatch is not related to the actual motion of the entire receiving structure however, but rather only to the interface. It is therefore difficult to qualify a characterisation as good or bad based just on this measure.

An equivalent measure is defined as a measure that leads to identical results on the receiving structure. For this reason the match in the $\boldsymbol{u}_3$-space seems leading in defining equivalency. We have seen in the decomposition of the discretised interface in section 4.2.2 that the dynamics describing the $\boldsymbol{u}_3$-space are not necessarily the same as for the $\boldsymbol{u}_4$-space, although intuitively it seems they should show similar dynamics. To verify this, they are compared in the following subsection.

Measured vs. Predicted Responses

Each In-Situ measurement campaign will consist of n_4 indicator $\boldsymbol{u}_4$ sensor channelss and n_3 validation $\boldsymbol{u}_3$ sensor channels. In the source characterisation calculations only the measurement data from the indicator sensors is used. The measurement data from the validation sensors is then typically used to validate the equivalent measure. In this way the measurement and prediction are compared.

Equation 5.3 and equation 5.4 show the definition of the residual motion. Each measured sensor channel is compared to its predicted counter part. For the $\boldsymbol{u}_4$-space this measure shows the residual of what the inverse solution method could not capture. In the case of a Least Squares approximation this residual is even the definition. For the $\boldsymbol{u}_3$-space residual the relation to the equivalent force is less intuitive.

$$\boldsymbol{u}_4^{res} = \boldsymbol{u}_4^{meas} - \boldsymbol{u}_4^{pred} \qquad \text{with } \boldsymbol{u}_4^{pred} = \boldsymbol{Y}_{42}\boldsymbol{f}_2^{eq} \tag{5.3}$$

$$\boldsymbol{u}_3^{res} = \boldsymbol{u}_3^{meas} - \boldsymbol{u}_3^{pred} \qquad \text{with } \boldsymbol{u}_3^{pred} = \boldsymbol{Y}_{32}\boldsymbol{f}_2^{eq} \tag{5.4}$$

The result of these residual motions can be seen in figure 5.4. For a clearer representation, the norm of both is plotted instead of each sensor channel. The excitation is a broadband excitation and the measurement data is not noise-contaminated. We clearly see the a residual peak in both vector norms at the same frequencies were we observed the controllability problems. From these measures we cannot qualify the characterisation as being good or bad however, since this only shows the absolute residual, but not related to the original motion.

An important conclusion we can make from these plots however, is that the residual motion for both vector spaces shows the same resonant behaviour. This same relation is seen when the $\boldsymbol{u}_4$-space is reduced to the interface $\boldsymbol{u}_2$-space. In other words, when the interface is not fully blocked, the residual motion further downstream the receiving structure is also the largest. Extending this notion further, we can conclude that the interface motion does describe the receiving structure fully. This conclusion can be explained by the fact that the $\boldsymbol{u}_4$-space sensors are placed behind the interface, but not on the interface. The controllability issues of the equivalent force were caused by its single position, but for the sensors this is not the case.

5.3.1. Effect of Mismatch in Prediction

The residual motion as described in equation 5.3 can be interpreted as a measure of how successful the characterisation is. If the prediction and the original motion are equal, then the difference between the two should be zero. This can also be explained by the superposition principle for linear systems, where in this case the equivalent force is applied in negative to the system. We wish to relate this mismatch on the interface to the overall equivalence on the entire receiving structure. In other words, how successful is the characterisation, how equal is the equivalent force. Whereas for equation 5.3 and equation 5.4 the mismatch was defined as an absolute measure, it would be better to have a measure that defines the equivalence relative to the original excitation.

The effect of the mismatch in the prediction will be quantified in two steps, both involving $\boldsymbol{u}_4^{res}$. This measure can be reduced to the $\boldsymbol{u}_2$-space as seen in section 2.2 to describe the actual interface motion. Similarly, it turns out $\boldsymbol{u}_4^{res}$ can also be transformed in the other direction to the $\boldsymbol{u}_3$-space. This will make it possible to describe the non-equivalence of the source characterisation in terms of how blocked the interface is. To show the prediction data can be used in both direction, the figures are plotted on the last page of the chapter.

Interface Space

The $\boldsymbol{u}_4$-space can be reduced to a $\boldsymbol{u}_2$-space using the Virtual Point Transformation theory from section 2.2, as seen in equation 5.5. For this reduction to be a valid reduction, two assumptions must hold true. The number of measurement responses n_4 must be larger than the number of Virtual Point resonses n_2 and the structure around the $\boldsymbol{u}_4$ sensors must behave rigidly. For the case study these assumptions are tested in section A.1 and it is shown they hold true.

$$\boldsymbol{u}_2 = (\boldsymbol{R}_u)^+ \boldsymbol{u}_4 \tag{5.5}$$

We can relate the residual interface motion in the $\boldsymbol{u}_2$-space to the original motion in that space to get a relative measure. There are many ways to relate the two and it is difficult to pick a best solution. The implementation I suggest is to define two measures to describe the interface, one for the maximum mismatch and one for the mean mismatch.

The first measure is the maximum $\boldsymbol{u}_2$ blockedness and is defined using the the largest value of the $\boldsymbol{u}_2$ measurement and its relation to that same measurement channel in the residual motion $\boldsymbol{u}_2^{res}$. This gives an upper bound for the definition of the interface stillness. Equation 5.6 shows the mathematical definition. The mean mismatch can be defined using equation 5.7 and compares the overall interface motion to the residual.

$$\boldsymbol{u}_{2,\text{max}}^{\text{blocked}} = 1 - \frac{|\boldsymbol{u}_2^{res}\,(i_{\text{max}})\,|}{|\boldsymbol{u}_2\,(i_{\text{max}})\,|} \qquad \text{with} \quad \max\left(|\boldsymbol{u}_2|\right) = |\boldsymbol{u}_2\,(i_{\text{max}})\,| \tag{5.6}$$

$$\boldsymbol{u}_{2,\text{mean}}^{\text{blocked}} = 1 - \frac{\text{mean}\left(|\boldsymbol{u}_2^{res}|\right)}{\text{mean}\left(|\boldsymbol{u}_2|\right)} \tag{5.7}$$

This measure is chosen to be very simplistic on purpose, as it will be solely used to give insight for which frequency bins the interface is not blocked. There is certainly room for improvement in defining this measure, for example by also extending the maximum mismatch to more than one channel. That being said, in the way the measure is defined currently it gives a good indication of the blockedness of the interface. It clearly shows resemblance to the specific controllability that was defined earlier.

Receiving Structure Space

The $\boldsymbol{u}_4$-space can be transformed to the $\boldsymbol{u}_3$-space in a similar manner. The relation between the two spaces can be described by a Transmissibility matrix as seen in equation 5.8. This method was first suggested by Ribeiro [44] for a different use case, namely to identify different transmission paths and to rank their importance. The method was extended by Gajdatsy to a MIMO case [24].

$$\boldsymbol{u}_3 = \boldsymbol{T}_{34}\boldsymbol{u}_4 \tag{5.8}$$

As suggested by Van der Seijs [54], the concept of Transmissibility can be used for Transfer Path Analysis use cases too. The danger in using the method is that the relation of equation 5.8 does not hold if there is a transmission path from the source to the $\boldsymbol{u}_3$-space without it crossing the $\boldsymbol{u}_4$-space or when there are observability problemsx. For substructures where there is only a single interface, this issue can be prevented. The transmissibility from $\boldsymbol{u}_4$-space to $\boldsymbol{u}_3$-space can be exact if the Transmissibility matrix $\boldsymbol{T}_{34}$ is constructed using the Frequency Response Function that we created for the real-life interface in section 4.2. This makes it possible to describe the full $\boldsymbol{u}_3$-space from the $\boldsymbol{u}_4$-space.

Similar to how we described the interface blockedness in equation 5.6, we can also define a blockedness or stillness for the $\boldsymbol{u}_3$-space behind the interface. We define the same two measures, only this time for the $\boldsymbol{u}_3$-space. The maximum stillness will give proper insight if the largest motion on the receiving structure is predicted properly by the equivalent force. The mean stillness gives good insight in how close the overall prediction is in the source characterisation. The relations can be seen in equation 5.9 and equation 5.10

$$\boldsymbol{u}_{3\text{max}}^{\text{stillness}} = 1 - \frac{|\boldsymbol{u}_3^{res}\,(i_{\text{max}})\,|}{|\boldsymbol{u}_3\,(i_{\text{max}})\,|} \qquad \text{with} \quad \max\left(|\boldsymbol{u}_3|\right) = |\boldsymbol{u}_3\,(i_{\text{max}})\,| \tag{5.9}$$

$$\boldsymbol{u}_{3,\text{mean}}^{\text{stillness}} = 1 - \frac{\text{mean}\left(|\boldsymbol{u}_3^{res}|\right)}{\text{mean}\left(|\boldsymbol{u}_3|\right)} \tag{5.10}$$

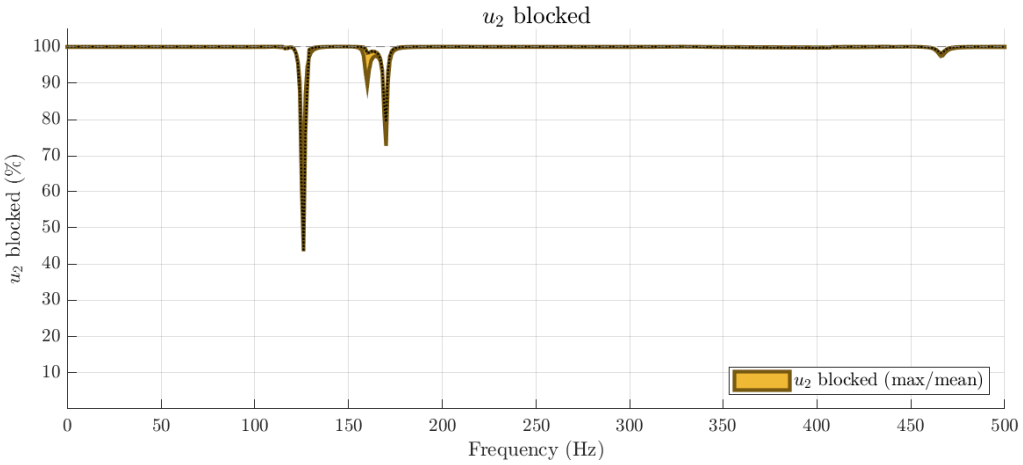

Figure 5.3: u_2 blocked

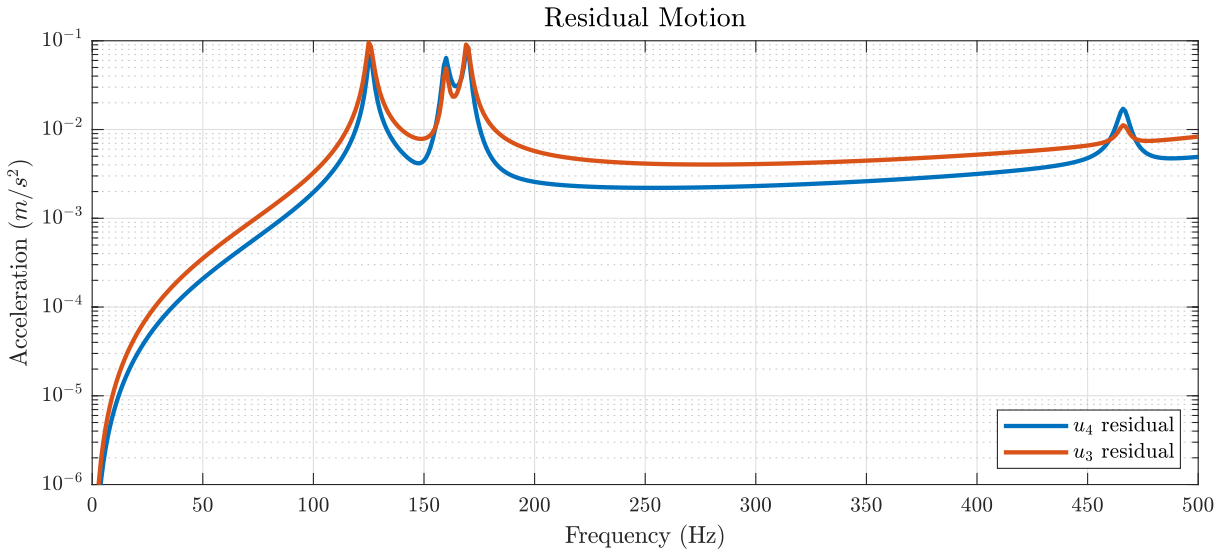

Figure 5.4: Residual motion in $\boldsymbol{u}_4$-space and $\boldsymbol{u}_3$-space

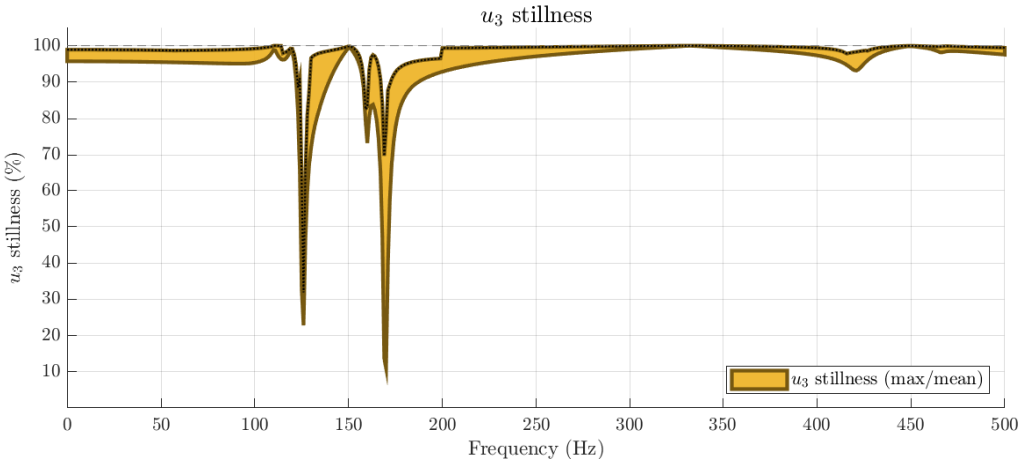

Figure 5.5: u_3 stillness

5.4. Conclusion

This chapter set out to translate the interface problems uncovered in chapter 4 to its effect on the equivalent force in source characterisation. The interface incompleteness was described in terms of Frequency Response Functions that are often unobtainable and did not show the actual difference in motion on the receiving structure.

Two new controllability measures were introduced. These measures define the controllability as a measure to which extend the equivalent force can excite the $\boldsymbol{u}_4$-space. This can be done for a specific measurement, leading the Specific Controllability. This analyses the modes that are present in the measurement signal to give a percentage of how many of these are controllable from the interface. The results of the Specific Controllability give an upper bound of how well the different mode shapes can be predicted.

A similar approach is chosen for the Global Controllability, but as the name suggests, this measure defines the controllability of the system regardless of the excitation load case. In this manner the lower bound of which different mode shapes can be predicted from the interface is defined. The results of the Global Controllability is in terms of each of the 6 modes of the equivalent force per frequency. The results showed some unexpected peaks in lack of controllability that cannot be explained from the way in which the interface is modelled. This begs the question of there are other modelling assumptions that have to be revisited.

As a final step the controllability issues are translated to their effect on the characterisation. To do this, two new measures are introduced that judge the quality of the characterisation process. The equivalence is clearly defined as being the motion behind the interface in the $\boldsymbol{u}_3$-space. We see that this space is linearly dependant on the $\boldsymbol{u}_4$-space, which means the residual of the inverse problem can be used directly to define the equivalence. Similarly, the residual motion of the inverse problem an also be reduced to the $\boldsymbol{u}_2$-space to define how blocked the interface is. The conclusion made here is that the non-equivalence is due to the interface not being blocked in the blocked force case. The modes occurring at these frequencies are not reproduced by the equivalent force. While the troublesome modes resonate at these frequencies, it becomes clear that they are also present in a much wider frequency range around the resonance frequency. This means the equivalence can be worse at the frequencies surrounding the fixed interface eigenfrequencies.

6

Robust Inverse Problem Solution

6.1. Introduction

The following chapter is the final chapter in the source characterisation part of this thesis and will focus on the final research question. We will zero in on the solving of the inverse problem from error-contaminated data. The goal of this chapter is to set guidelines that help in choosing the right regularisation tool for the right purpose.

First we will expand the mathematical description of the inverse problem to understand the problem at hand. Here we will see that the final solution to the noisy inverse problem will be an approximation of the actual solution. Depending on the noise level, this approximation can be only a slight deviation from the true solution for low relative noise levels, up to an approximation that has little connection to the physical problem for large relative noise levels.

We will try multiple methods that we have seen in chapter 3 to the inverse problem. The goal will be to translate these theoretical concepts to a physical explanation for the problem at hand. We recognise that we can categorise the regularisation techniques based on two pieces of a priori knowledge. First of all we have a way of finding the noise level of our measurement. On top of that we also have a way of describing the measured response data in terms of the individual modes of the matrix. From this we can find the exact building blocks of motion that passes the interface and its respective equivalent force.

The inverse problem we investigate in this section will be the In-Situ source characterisation problem from the Component-based Transfer Path Analysis methods. The reason we choose this method is that it is often used and in theory it is the most transferable method. On top of that the problem becomes an overdetermined problem, which has its own difficulties associated with it. Nevertheless, the methods tried here on the In-Situ method will also translate to other methods. The pseudo-force method that could help the inherent interface force controllability problem of chapter 4 is a similar overdetermined problem with far worse conditioning in general.

The specifics regarding the experimental model can be found in section **??**. The structure we use will be the same as for the other chapters. For the excitation we will test 2 load cases. The first one is a broadband signal that has unity gain across the frequency range. This type of signal is typical for wind noise for example. The other load case is a tonal or harmonic excitation which would typically be seen in a rotating component. The signal is extracted from a typical tonal excitation from a case study done by me and VIBES.technology on a E-Compressor for an automotive application. Some white noise is added to the simulated response data to simulate a real-life experiment for both cases.

The regularisation techniques are split up in different classes that act in a different way on the problem. Their effect can all be expressed in what the effective singular values are of the regularised inverse matrix. This gives direct physical relation to how the problem is changed due to regularisation. The inverse matrix describes a dynamic stiffness matrix where the singular values determine the stiffness of each

mode. The different methods will be compared by analysing the noise propagation of the measurement data to the solution. As we have described in chapter 5, the equivalence of the characterisation can be described in terms of the $\boldsymbol{u}_3$-motion. This is also the measure of choice to classify the successfulness of each method.

6.2. Mathematical Description

The mathematical model of the real-life Dynamic Substructuring system is similar to what we have seen in chapter 2, but differs due to some real-life limitations. There is still a linear relation between the forces and accelerations of the system, but due to the experimental nature in which both the system matrix and response data is gathered, they are both discritized and error-contaminated. The scope of this thesis is limited to the noise contamination of the response data, and therefore we will assume for now that the system matrix is not affected by noise, as was explained in section 3.4.

$$\boldsymbol{u} = \boldsymbol{Y}\boldsymbol{f} \tag{6.1}$$

We consider a linear system in the form of equation 6.1 with a linear operator $\boldsymbol{Y} \in \mathbb{C}^{[m \times n]}$ with $m > n$, meaning it is an over-determined system. The vector $\boldsymbol{u}$ represents the measurement data where $\boldsymbol{u} \in \mathbb{C}^{[n]}$. This measurement data is error-contaminated and we define its parts as seen in equation 6.2, where $\boldsymbol{u}$ without label stands for the measurement data, $\hat{\boldsymbol{u}}$ stands for the true, noiseless response and $\boldsymbol{e}$ stands for the noise and has the same dimension as $\boldsymbol{u}$.

$$\boldsymbol{u} = \hat{\boldsymbol{u}} + \boldsymbol{e} \tag{6.2}$$

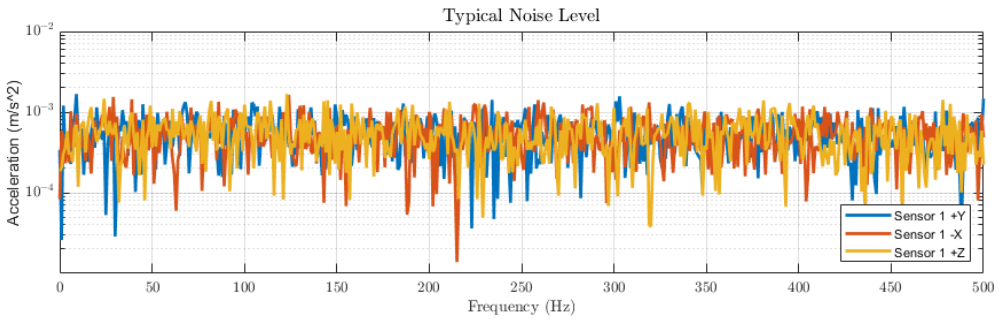

Figure 6.1: Typical Noise Level

Figure 6.1 shows the effect of just the noise term $\boldsymbol{e}$ on the measurement $\boldsymbol{u}$. The noise is random white noise, meaning it has a flat response over frequency. For all frequencies the noise has the same amplitude. For each sensor channel a noise level of $\|\boldsymbol{e}\| = 0.01(\mathrm{m/s^2})$ rms is added. This noise level is chosen such that the signal to noise ratio is of a typical real-life experiment where most dynamics are measured above the noise level but some areas fall just below it, meaning they are difficult to capture in the inverse problem.

6.2.1. Solution

Equation 6.3 shows the inverse problem written down in the most basic form as a least-square problem. We seek a method for finding an approximate solution $\boldsymbol{f}$ stably from $\boldsymbol{u}$ using $\boldsymbol{Y}$. Since the measurement data $\boldsymbol{u}$ is build up from the true response and a noise term, solving equation 6.3 naively will lead to a solution that will also have a true solution term and some noise related term.

$$\min_{\boldsymbol{f} \in \mathbb{C}^n} \|\boldsymbol{u} - \boldsymbol{Y}\boldsymbol{f}\| \tag{6.3}$$

To show the effect of the noise on the inverse problem, equation 6.4 shows a naive solution of the problem. It is computed by pre-multiplying both sides of equation 6.1 with the inverse of the system

matrix $\boldsymbol{Y}$ and using the split up measurement term from equation 6.2. Since the system matrix is non-square, the Moore-Penrose pseudo-inverse which is denoted by $\boldsymbol{Y}^+$ is used, as we have seen in section 3.4.1.

$$\begin{aligned} \boldsymbol{Y}\boldsymbol{f} &= \hat{\boldsymbol{u}} + \boldsymbol{e} \\ \boldsymbol{Y}^+\boldsymbol{Y}\boldsymbol{f} &= \boldsymbol{Y}^+\hat{\boldsymbol{u}} + \boldsymbol{Y}^+\boldsymbol{e} \\ \rightarrow \boldsymbol{f} &= \hat{\boldsymbol{f}} + \boldsymbol{Y}^+\boldsymbol{e} \end{aligned} \tag{6.4}$$

Equation 6.4 shows that the solution obtained from naively solving the problem includes the true solution denoted by $\hat{\boldsymbol{f}} \in \mathbb{C}^{[n]}$, but there is also a term related to the noise. To express the effect of this term, we calculate the norm of the solution as seen in equation 6.5. We use the triangle inequality principle for the summation of norms to find an expression using both terms. On top of that we assume that the noise $\boldsymbol{e}$ measurement has a constant norm ϵ that we can pull out of the norm expression.

$$\begin{aligned} \|\boldsymbol{f}\| &= \|\hat{\boldsymbol{f}} + \boldsymbol{Y}^+\boldsymbol{e}\| \\ &\leq \|\hat{\boldsymbol{f}}\| + \|\boldsymbol{Y}^+\boldsymbol{e}\| = \|\hat{\boldsymbol{f}}\| + \|\boldsymbol{Y}^+\|\sqrt{n_c}\epsilon \end{aligned} \tag{6.5}$$

We have seen that the norm of the inverse matrix is equal to one over the smallest singular value of the original system matrix. This singular value can be very small, meaning its inverse is very large, meaning the error term is amplified. This causes the upper bound for the solution norm to be very large and this in turn leads to the solution deviating from the true solution and often this is not a meaningful approximation.

6.3. No Regularisation

As a baseline we will solve the inverse problem with the Moore-Penrose pseudo-inverse. This inverse, which is also dubbed the Least Squares Solution, was derived in section 3.4.1 and is the most basic form of solving an overdetermined set of equation. It is often called the naive solution, as it bluntly finds the solution that has minimum error without taking any specifics of the problem into consideration. The solution is found in the form of equation 6.6.

$$\boldsymbol{f} = \sum_{i=1}^{r} \frac{\left(\boldsymbol{u}_i^T\boldsymbol{u}\right)}{\sigma_i}\boldsymbol{v_i} \tag{6.6}$$

Before looking at any of the results, we already know what kind of solution to expect from the naive method. The matrix inverse is fully determined by the left and right singular vectors and their weighting is simply the inverse of the singular values of the addmitance matrix $\boldsymbol{Y}_{42}$ and the projection of the measurement data on each mode. It uses no a priori information about the noise level and simply maps the noisy measurement to a noisy equivalent force. We expect a high level of noise propagation in the final solution as well as a solution that is potentially overfitted to the noise, leading to a non-smooth behaviour when plotted over frequency.

Since this is the first case analysis, some insight in how to read the overview plots is provided. The top 6 figures show the forces and moments of the the solution, the equivalent force $\boldsymbol{f}_2^{eq}$. The true noiseless solution is plotted with a dashed line. The noise propagation of just the noise part of the measurement can be seen in an orange area plot. It becomes clear that all solutions that are found that are below this noise level are not computed correctly and become lost in the noise. The bottom three plots show the solution in a very compressed form, where we first see the norm of the solution which gives insight in how big the solution is. The singular values plot is for the inverse problem and can be regarded as the effective singular values of the inverse problem. The colours used for the singular values correspond to all figures depicting information about the modes. Finally the last plot shows the equivalence how we have defined it earlier. It can be argued that this is the most important plot since it indicates how successful the source characterisation process is.

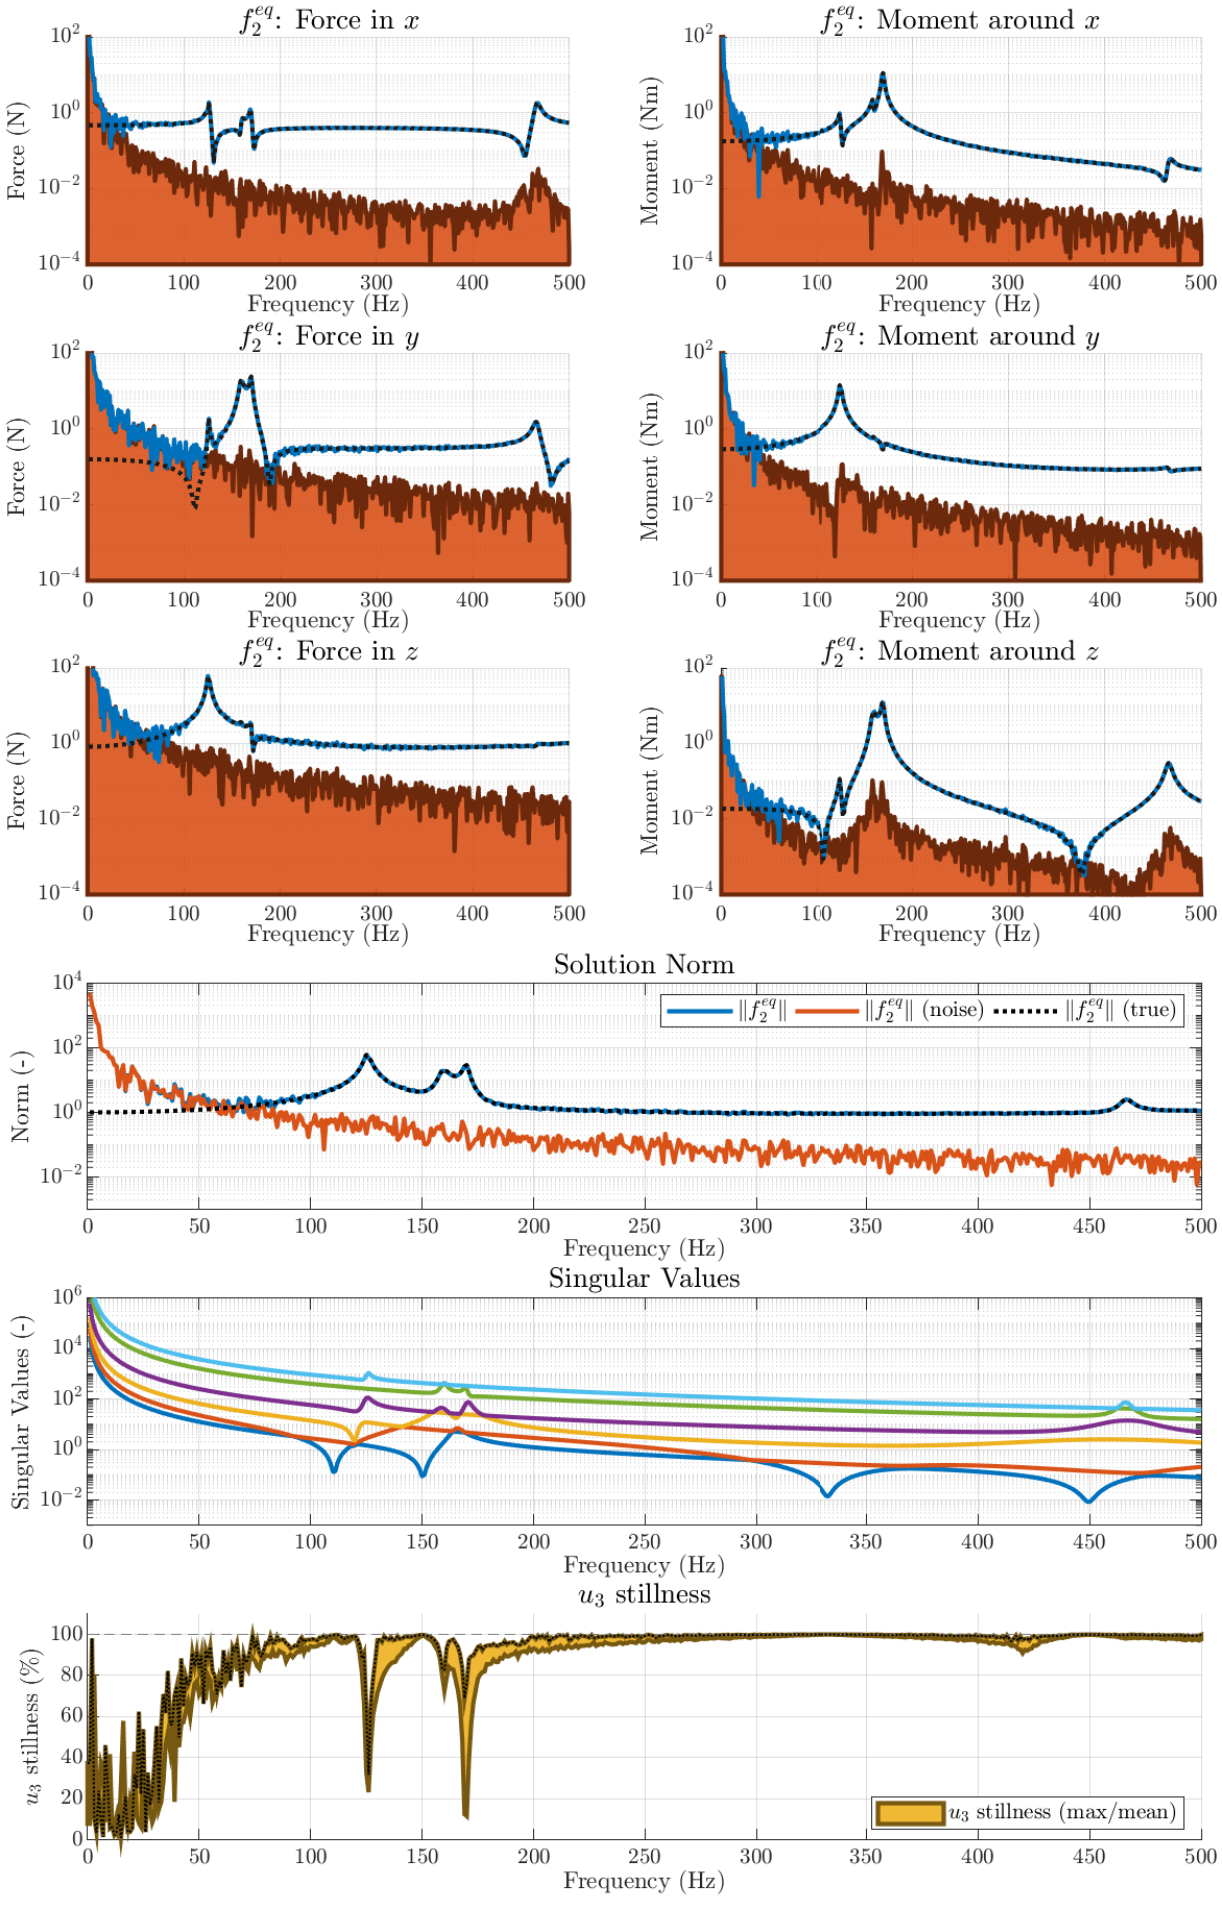

Figure 6.2: Overview Naive inverse (broadband excitation)

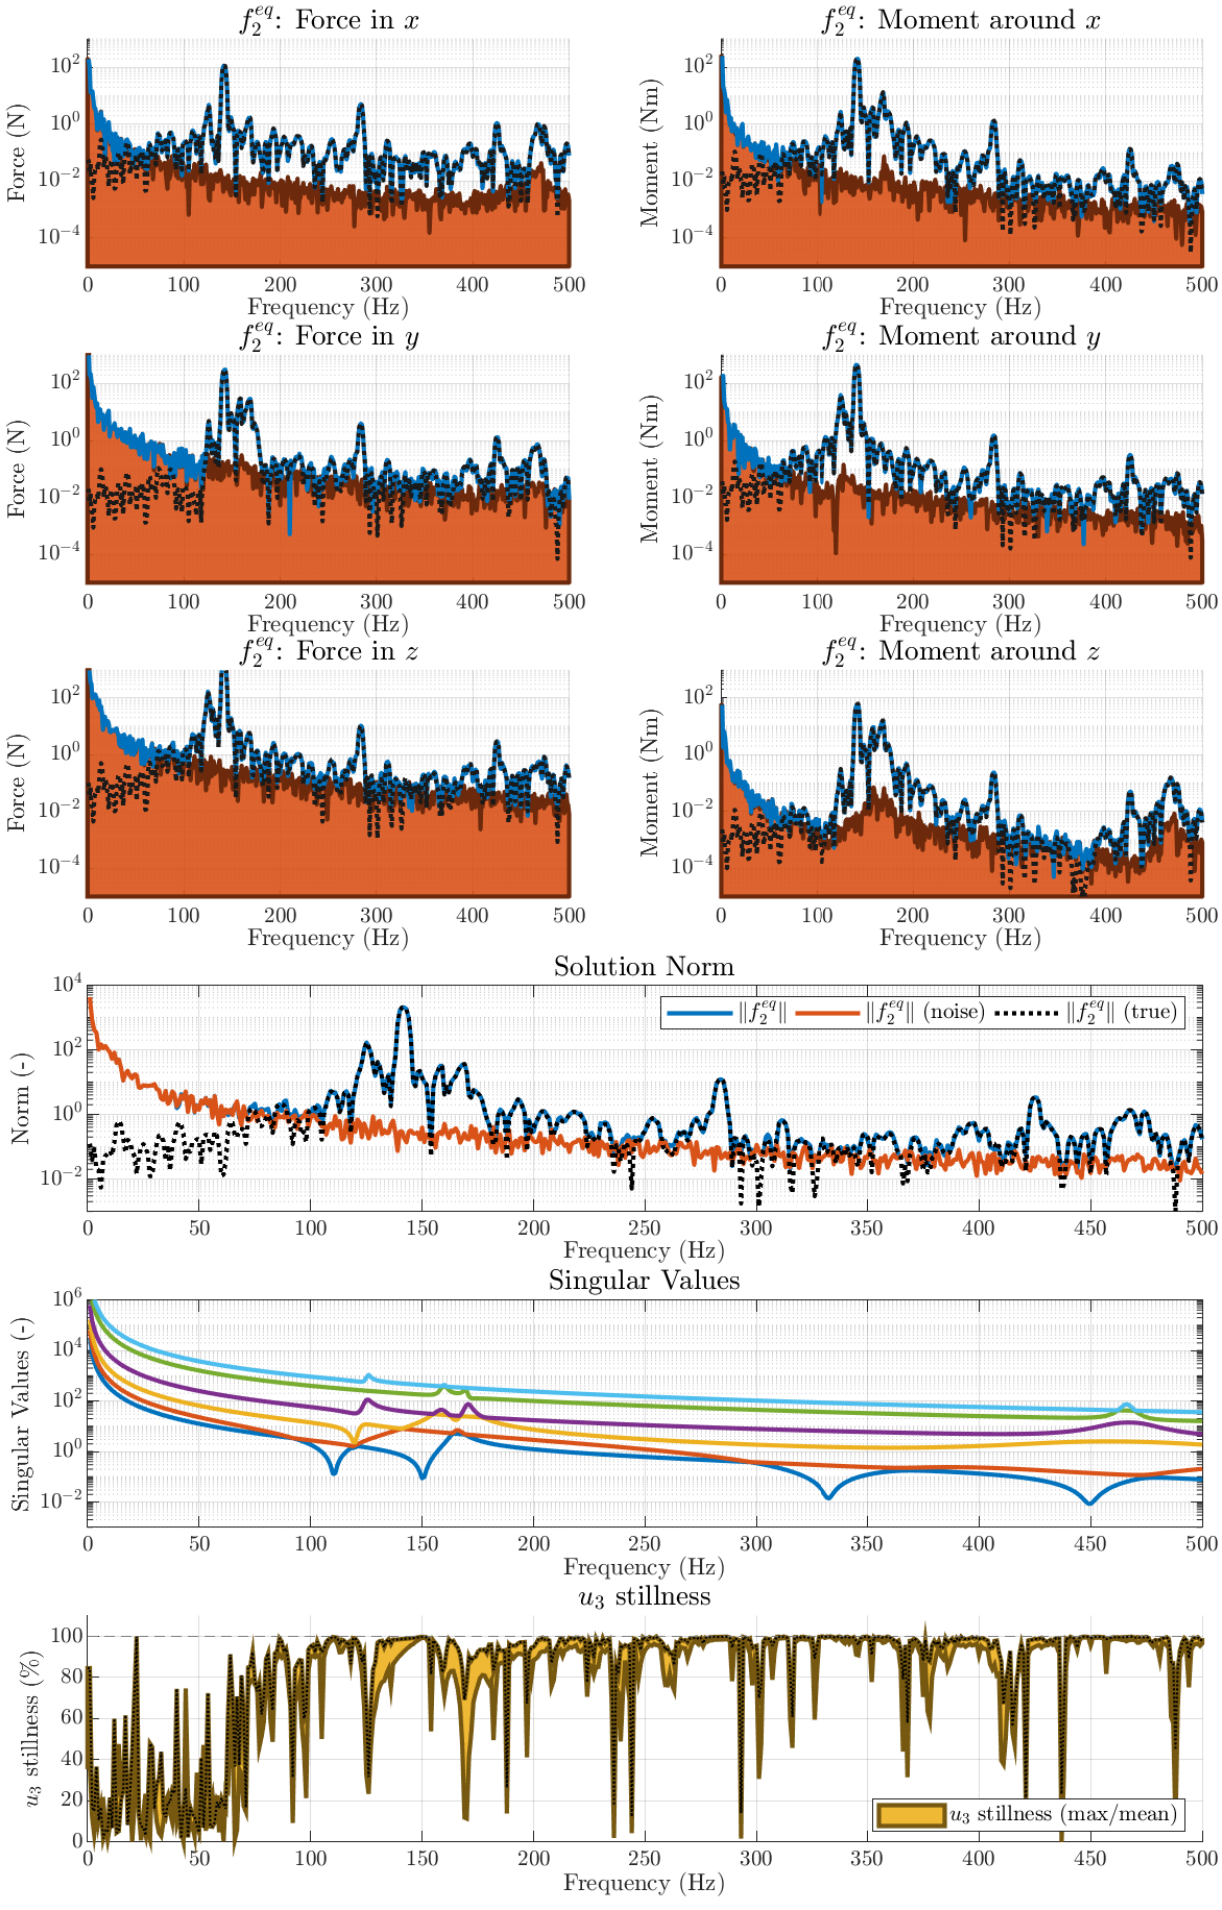

Figure 6.3: Overview Naive inverse (tonal excitation)

Broadband Excitation

For the first test case we see the results in figure 6.2. The equivalent force has a high noise floor. This especially seen in the low frequent region, where the solution is fully dominated by the noise. The solution can simply not describe the true solution in these places due to the amplification of the noise. This can also be seen in the u_3-stillness plot, where the low frequent region shows almost no equivalency. This behaviour is due to the large stiffness at low frequencies. This inherently occurs for structures that are attached to the fixed world, as the test bench for these experiments is.

Once the equivalent force is above its own noise level, the naive method performs a lot better. The solution norm is above the noise level from around $f \approx 100$Hz, but there is still some non-equivalency. The first peaks of non-equivalency around $f \approx 120$Hz and $f \approx 160$Hz are due to the controllability issue of an interface force. No matter the method, we will always see some non-equivalence here due to the choice of an interface force.

From about $f \approx 200$Hz and higher, we expect a good result since we are above the noise level and there are no controllability issues. The true solution is followed quite closely, although the solution is very noisy itself. This high frequent jitter around the true solution is due to the overfitting on the noisy data. We see this back in the u_3-stillness plot, where over the whole spectrum the result is quite shaky. This non-smooth behaviour over the frequency range originates from overfitting of every frequency bins on the noisy data.

Harmonic Excitation

For the second test case we see the results in figure 6.3. The results show a similar behaviour to the first test case. This is as expected, since the singular values of the inverse problem are the same. Unlike for the broadband excitation case, the tonal case has an equivalent force that is very close to the noise floor.

For low frequent the solution is dominated by noise which is seen in the poor equivalency. For higher frequencies the solution surpasses the noise floor, but drops below it at a few instances. At these frequency bins we see the poorer equivalency result as well. The controllability issues are also clearly visible for the tonal excitation. The excitation has been shaped in such a way, that the third harmonic of the main resonant frequency is around the same frequency as a controllability problem frequency ($f \approx 425$Hz). Although the noise solution follows the true solution very well here, the equivalence plot still shows a significant mismatch. This shows once more that the problem of source characterisation is not purely originated from noisy measurements, but also from the limitations of the method.

As a final note on the tonal case, it must be mentioned that the $\boldsymbol{u}_3$-stillness plot can be slightly misleading in judging whether the characterisation is successful. For a very tonal excitation it could be argued that the peak amplitude of the harmonics are very important to capture and the other frequency bins are of minor importance. This is not shown in the figure clearly, as it shows the full frequency spectrum.

Evaluation: Least Squares

We introduced the Least Squares method as the naive solution, which implicitly passes judgement on the method as it being no good. When we look at the results of the broadband excitation however, the equivalency is actually not that poor. The frequency ranges where it struggles are the regions due to controllability issues and regions where the signal to noise is very low. No regularisation method will solve the controllability issues. As for the low signal to noise ratio, there are some alternative methods which will be analysed later.

The quality of the result of using the Least Squares method for solving the overdetermined inverse problem depends very much on the specific system and sensor set-up. For experiments where all measurement channels have data which operate in the same order of magnitude, the method will find a good approximation of the solution. In the case where certain dynamics are only captured by a sensor channel which has relative low signal, this is not the case. These dynamics get lost in the least squares method since they have such low weighting compared to the other measurement channels.

Problems due to noise propagation will become very apparent with the Least Squares Method. The Least Squares inverse matrix has very stiff modes. These modes might not be participating in the true noiseless measurement $\boldsymbol{u}$, but since we are working with the noise contaminated data it means that there will be a projection on these modes and therefore only noise is propagated to the solution while no extra physicality of the problem is described.

6.4. Regularisation: Based on Noise Level

The first class of regularisation methods will be regularisation methods that are based on the noise level. For these methods we assume we know the noise norm ϵ as we have defined it in section 6.2. This is a realistic assumption, as in practice it is almost always possible to do a measurement with the source excitation turned off to get some pure noise measurements.

From the theory of chapter 3 we recognise a few methods which could be specifically used for noise suppression. The first method is the method which is closest to the naive solution, namely the Weighted Least Squares method. In this method we can determine the weight of every measurement channel in the mapping of the inverse solution. The second method we recognise is a specific class of Tikhonov Regularisation methods. Using the L-Curve, we can visualise the solution space and the corresponding residual. Summarising, the following methods are tested for both excitations:

1. Weighted Least Squares
2. L-Curve
    - Morozov Discrepancy Principle
    - Regularisation Choice by Wiener

6.4.1. Weighted Least Squares Method

The Weighted Least Squares method is a method that shows much resemblance to the Least Squares Method. We have seen in its derivation in section 3.4.2 that it originates from the same derivation as the Least Squares Method, with the only change being the weighting of the measurement data used.

$$\boldsymbol{f} = \left(\boldsymbol{Y}^T\boldsymbol{W}\boldsymbol{Y}\right)^{-1}\boldsymbol{Y}^T\boldsymbol{W}\boldsymbol{u} \tag{6.7}$$

Equation 6.7 shows the mathematical relation to obtain the solution, where $\boldsymbol{W}$ is the weighting matrix ($\boldsymbol{W} \in \mathbb{R}^{[m\times m]}$). This weighting matrix is a diagonal matrix which determines which channels are the most important in the approximation. This feature seems exactly to be what the usual Least Squares method was lacking. Unfortunately there is no elegant way to write down this method in a Singular Value Decomposition form with filter facts. We can however determine what the effective singular values of the new inverse matrix are and get an overview in that way.

There are multiple options to choose what weighting factor is given in $\boldsymbol{W}$. Two options are of the most interest to us, being the Signal to Noise ratio as weighting factors or using the reciprocal of the variance of a set of measurements. Both these methods lead to a diagonal Weighting Matrix that affect an individual measurement channel each. The way the method can be understood physically is that the weighting matrix only changes the relative size of each data point in the vector $\boldsymbol{u}$ to the others. Intuitively this makes sense when thinking of how a least squares solution finds the best middle ground between the data points, but now some data points are changed in importance.

In chapter 3 we have seen two methods of choosing the weighting matrix based on the noise level, being the Signal to Noise ratio as a Weighting Matrix and the Reciprocal of the Variance as a Weighting Matrix. Both these methods originate from the same reasoning in that the signal and noise level are related to each other and the weighting is determined from that. The effect that they have on the inverse problem is exactly opposite however. Figure 6.8 shows the weighting for each channel where the same colours for the same channels are used in both plots. As expected the relation is similar, but inverse from each other. The physical explanation from the singular values is not as intuitive for the

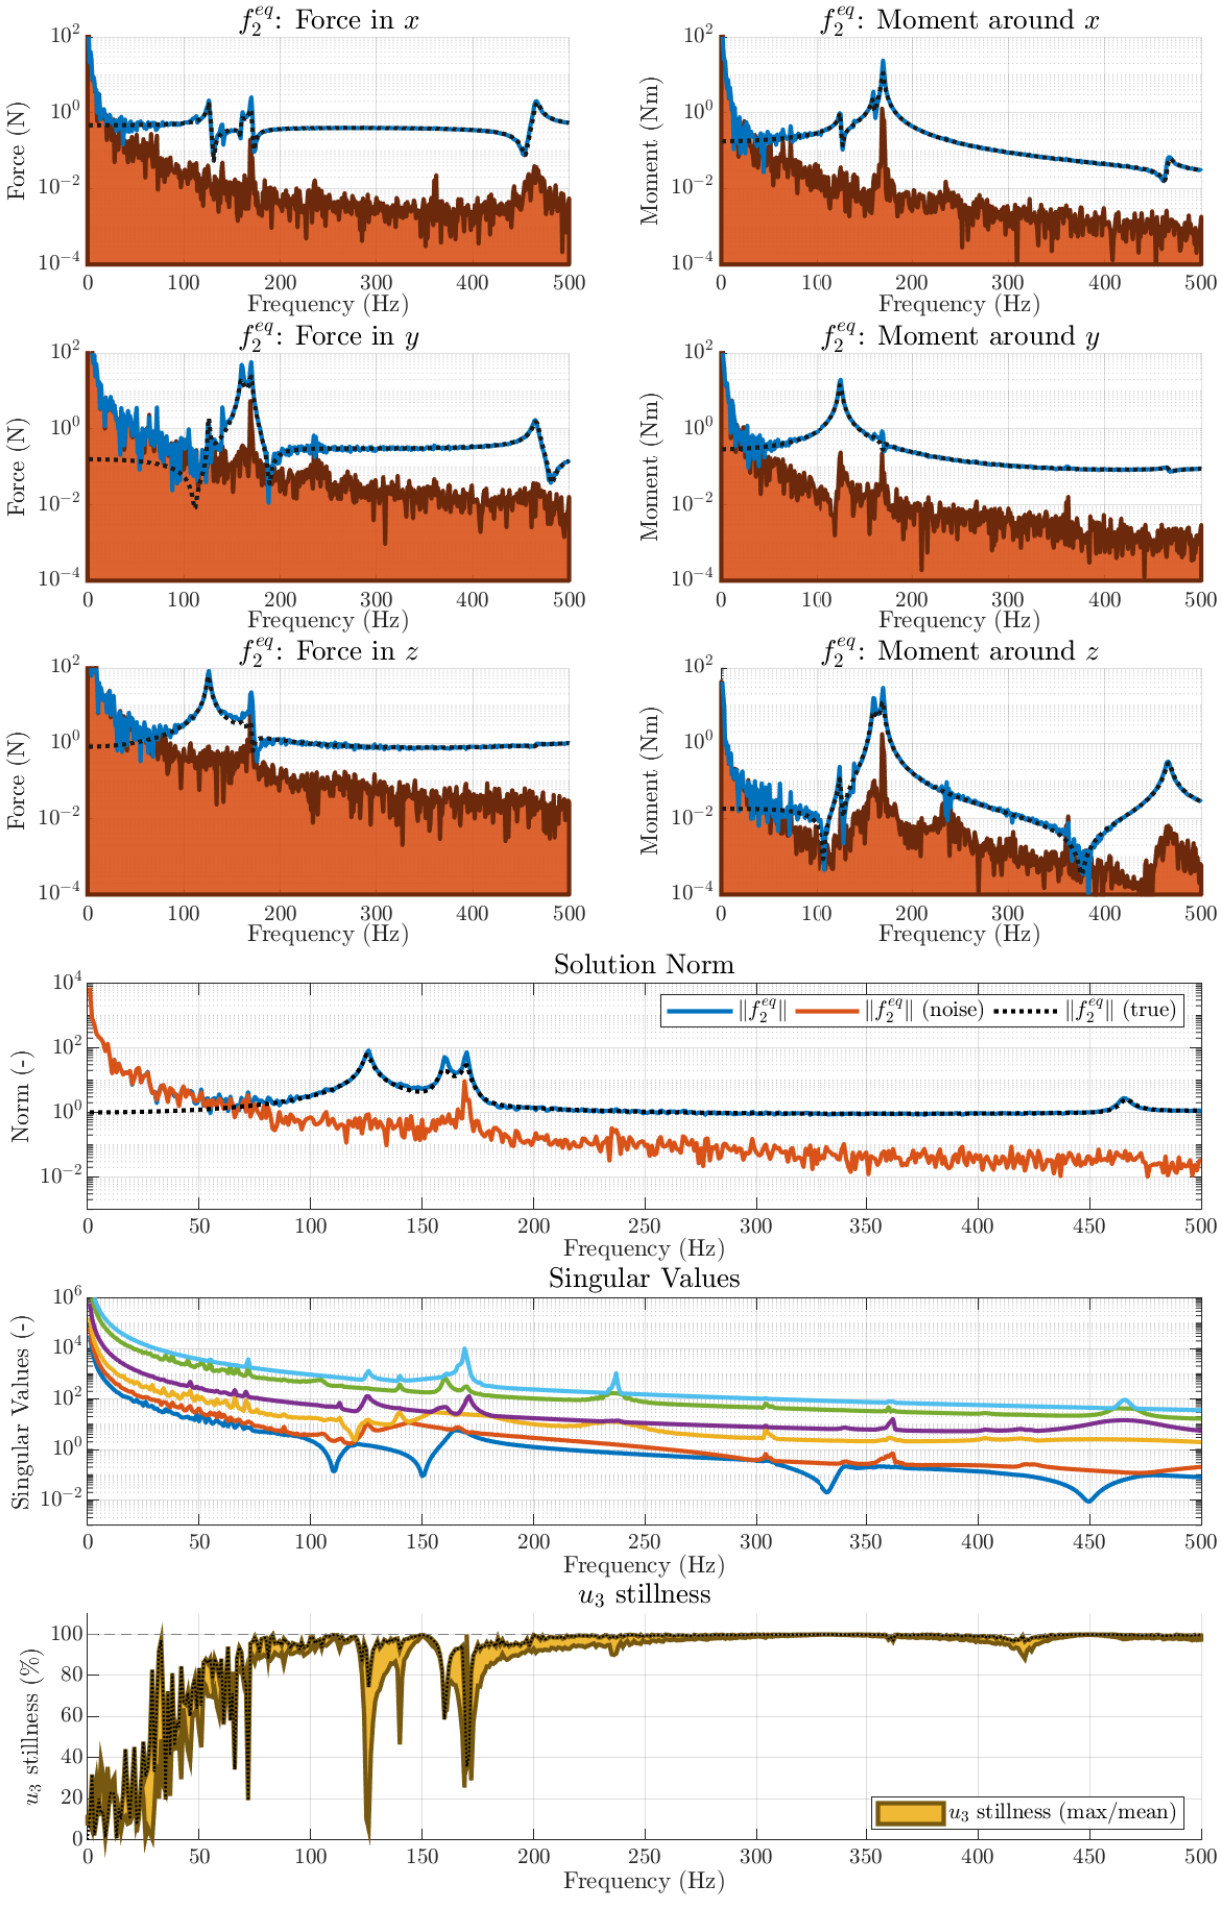

Figure 6.4: Overview Weighted Least Squares inverse (Variance) (broadband excitation)

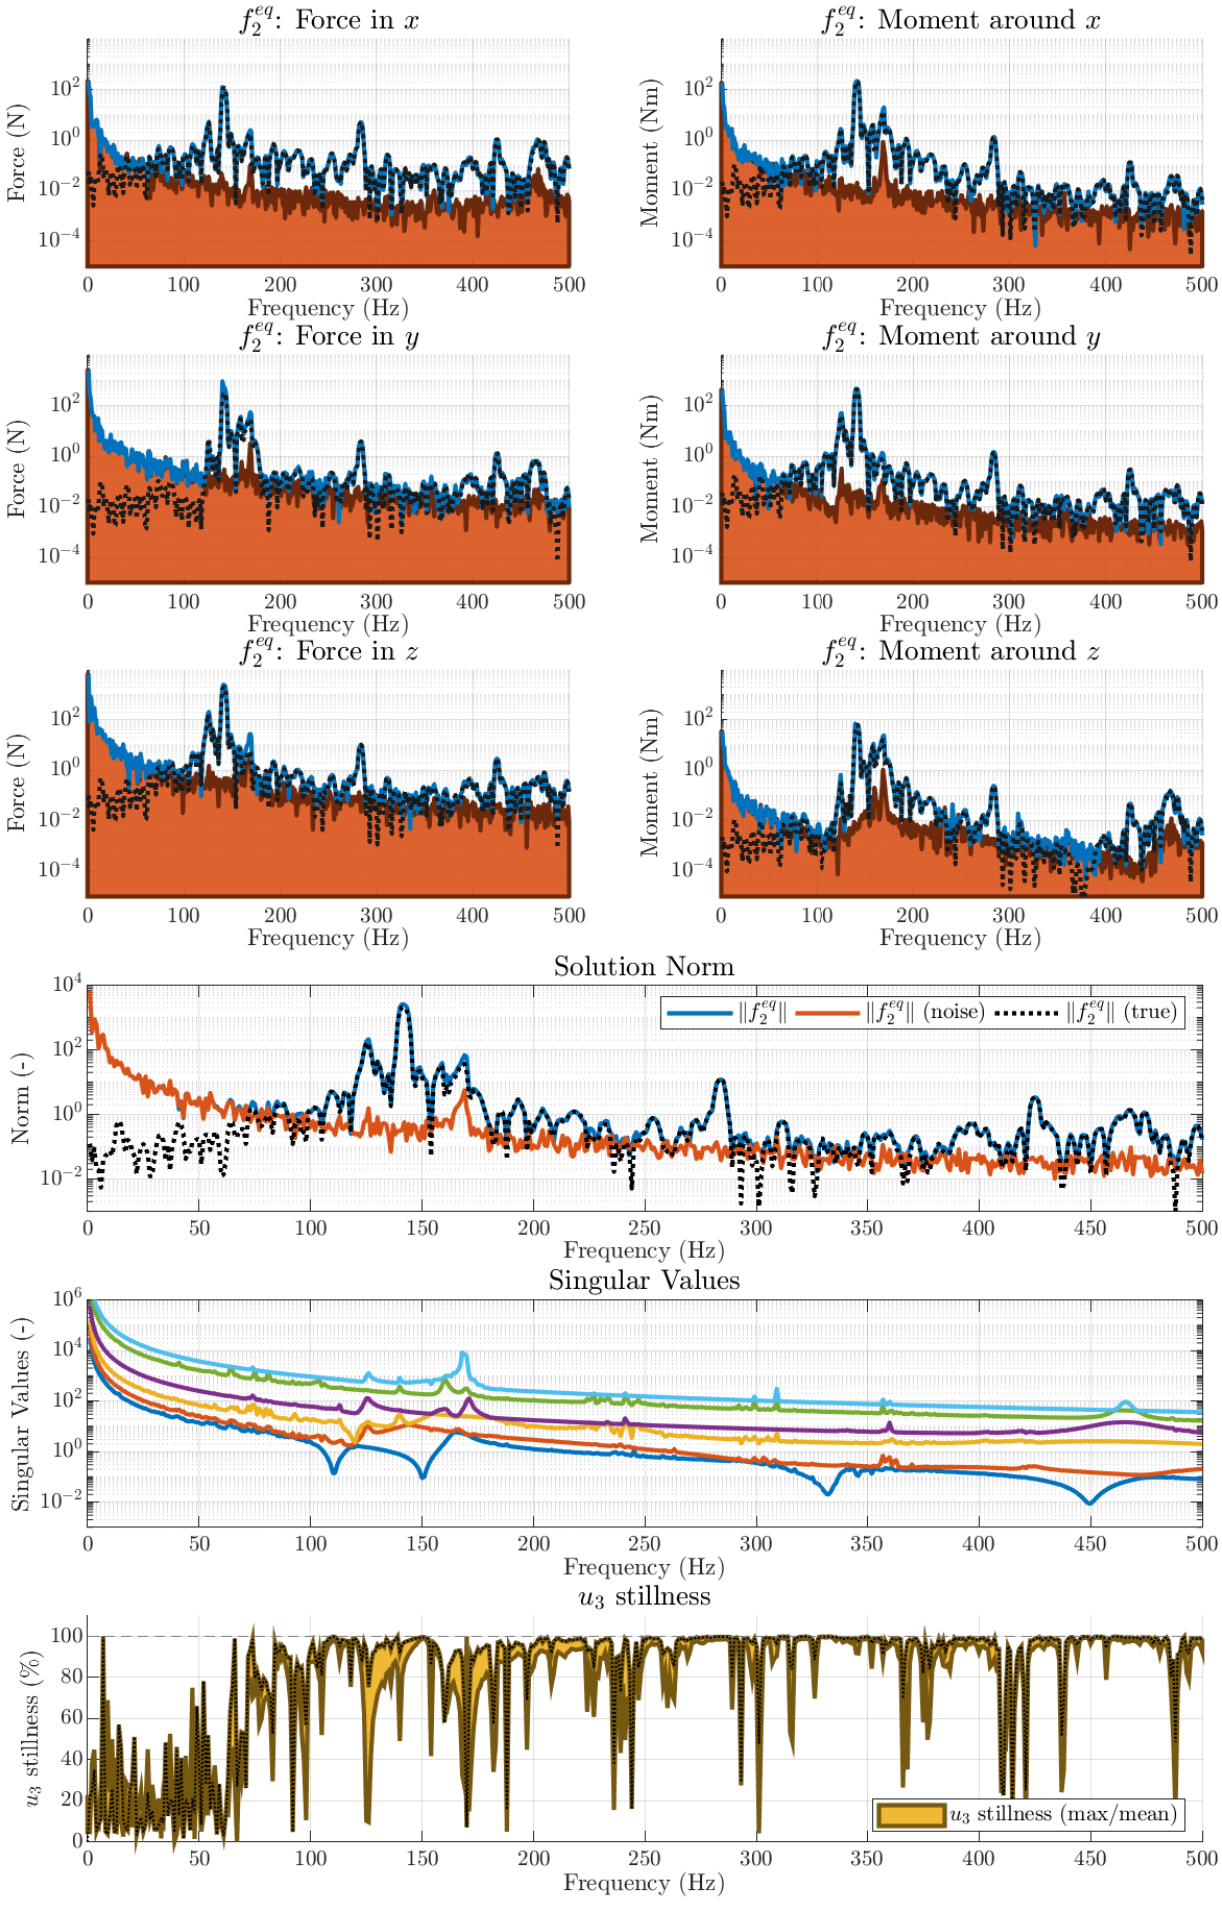

Figure 6.5: Overview Weighted Least Squares inverse (Variance) (tonal excitation)

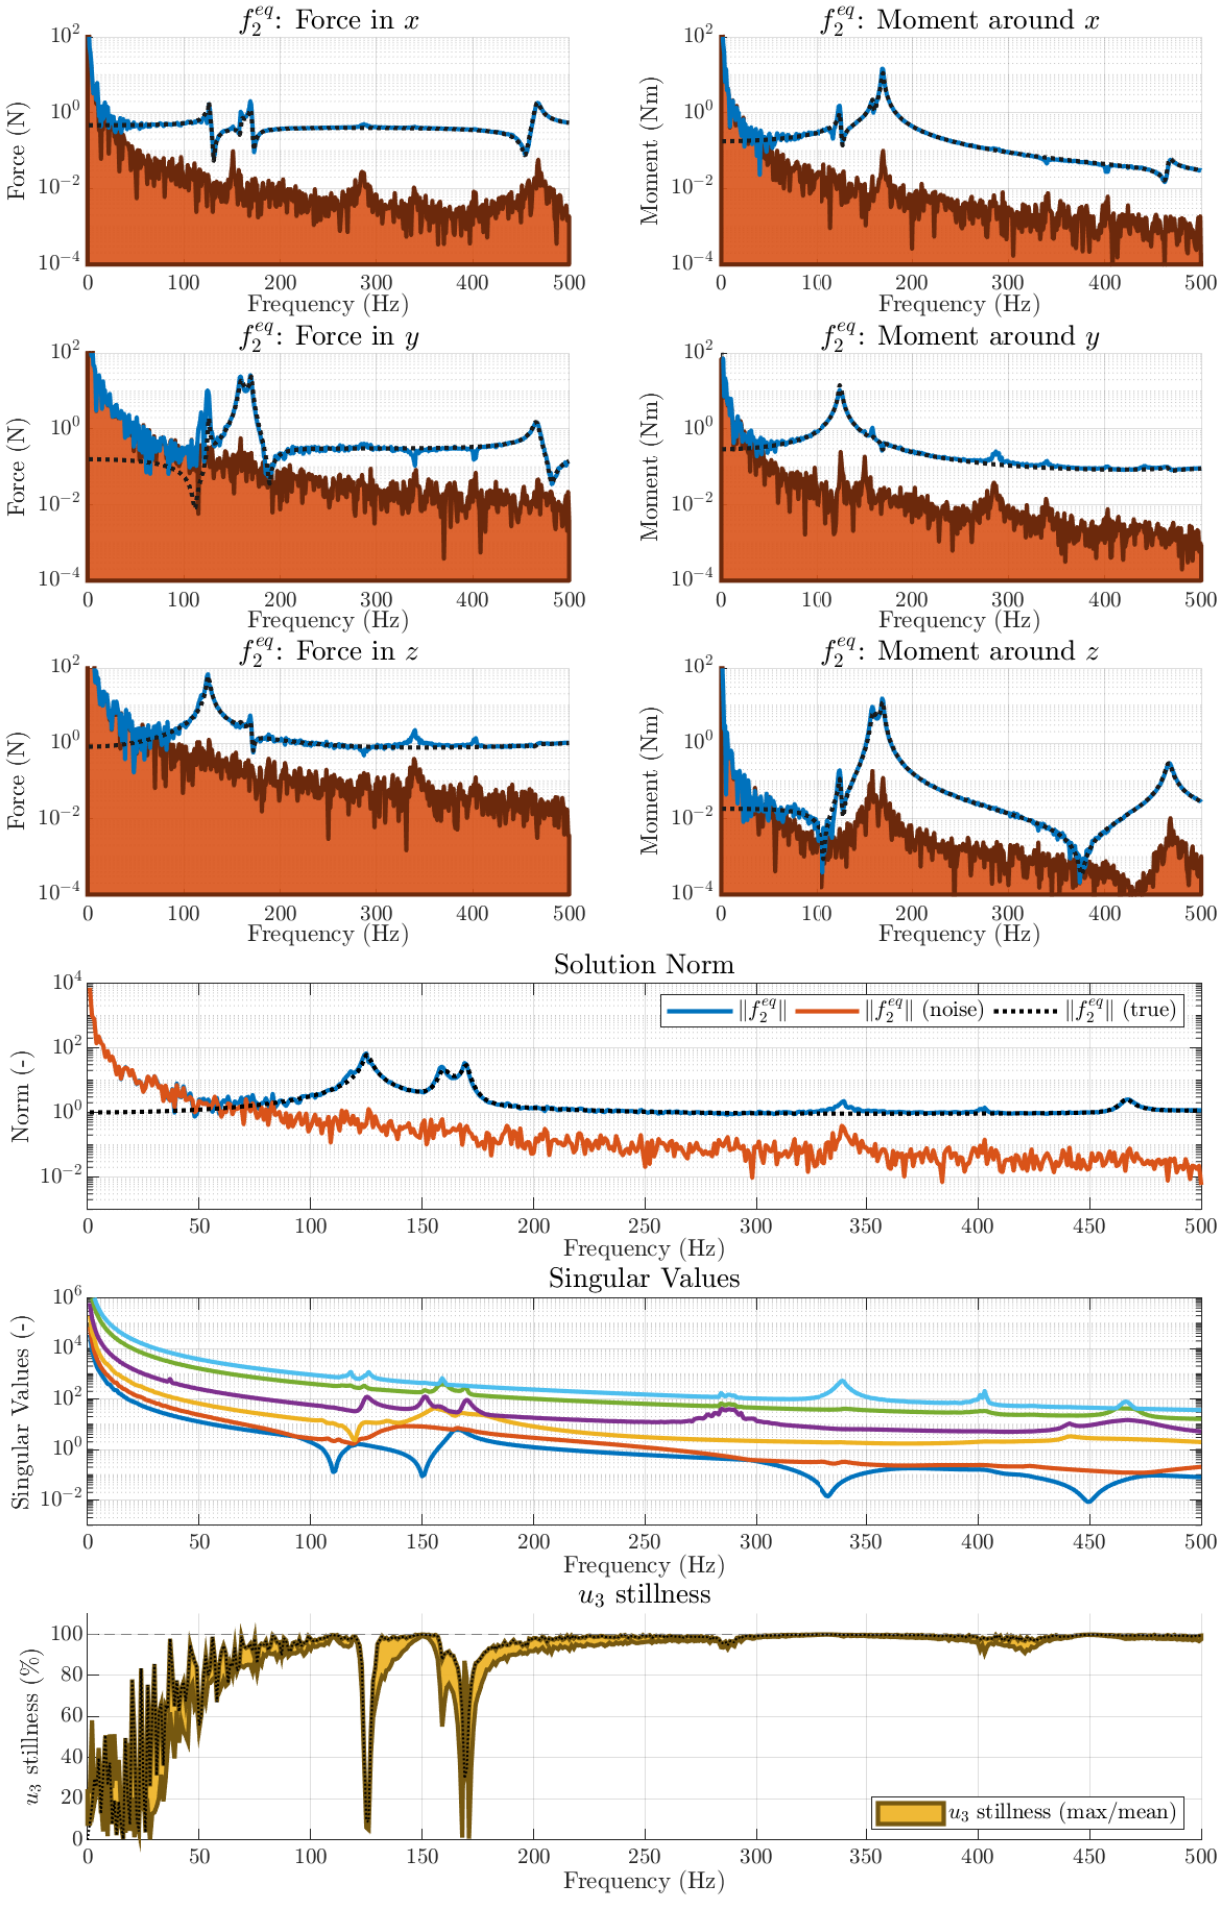

Figure 6.6: Overview Weighted Least Squares inverse (SNR) (broadband excitation)

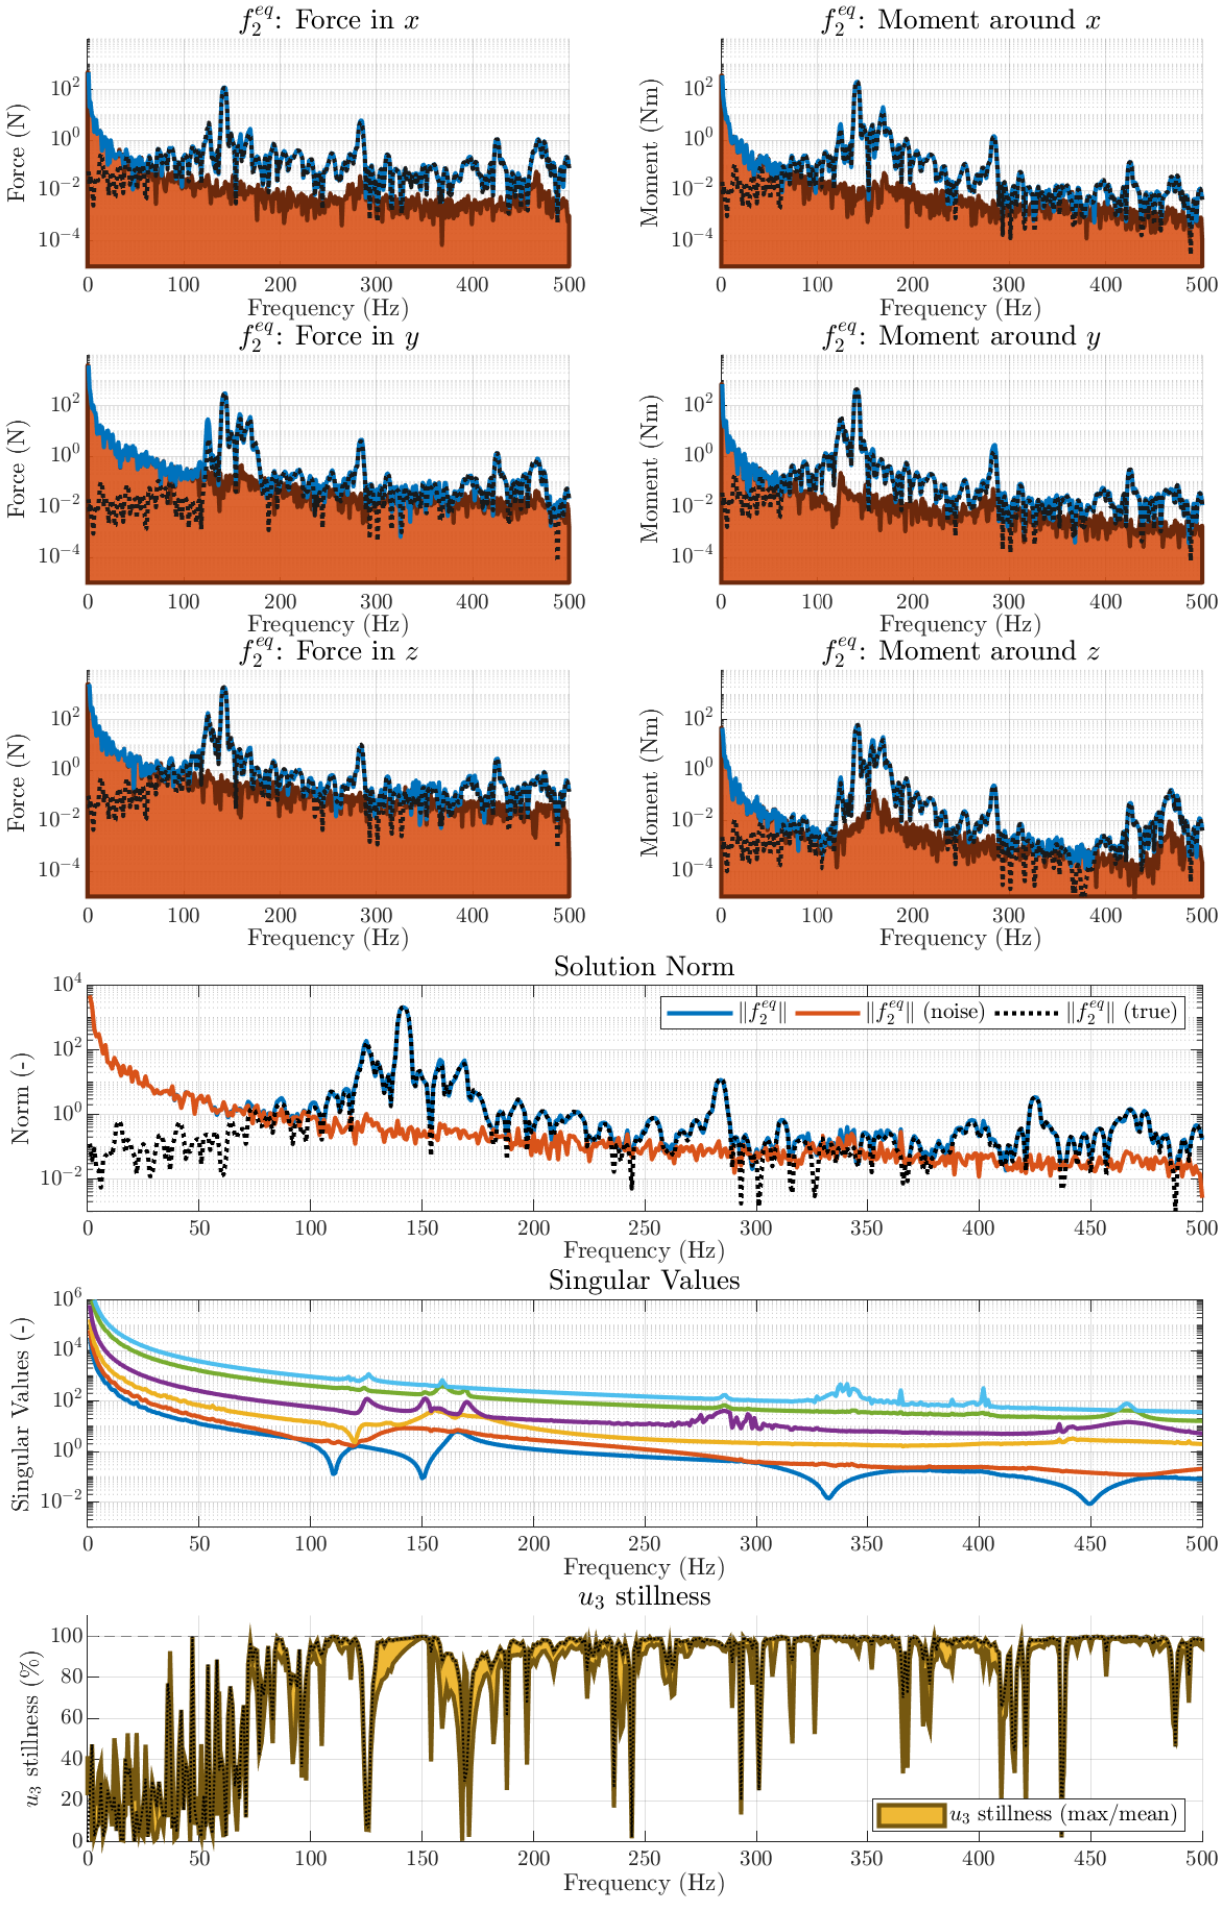

Figure 6.7: Overview Weighted Least Squares inverse (SNR) (tonal excitation)

Weighted Least Squares Method as for the other methods.

When using the *Signal to Noise ratio as a Weighting Matrix*, the measurement channels that have the lowest relative noise get a heavier weight in the solution. What this does is that certain dynamics are ignored when their measurement had a relative poor Signal to Noise. Therefore the method dampens out the the participation of noisy data channels. You can reason that therefore the effect of noise is indeed minimised. This motion in turn gets solved by extra participation of larger modes. If you look at the effective singular values (not depicted here), you see that the the largest modes are smoothened. This leads to the over representation of certain dynamics. Although this is beneficial for noise propagation, some physical representation to the real problem is lost.

An alternative is to use the *Reciprocal of the Variance as a Weighting Matrix* as suggested by Aitken. The variance is in regards to multiple measurements. This method is very applicable for experimental dynamics. By using different measurement sets and determining the variance between these measurements for each channel, we will find which channels at which frequencies are coherent (and therefore have little effect of noise) and which channels show the opposite behaviour. By using the reciprocal of the variance, we give the low signal to noise measurements extra weight. This is bad for noise propagation, but good for capturing effects that are hidden in the noise.

Although the initial target of regularising the inverse problem was to dampen the effects of noise, we have found a method that works well to emphasise poorly captured features in the data. This makes sense when looking at the situation for which both were developed by Powell and Aitken respectively. Powell suggested the method of using the Signal to Noise from a source characterisation experiment similar to the one in this thesis. Here the effect of noise is negative and hence it needs to be dampened. For Aitken however, the goal was to get the best estimation of the true solution from noisy data. This was more from a statistical point of view.

Broadband Excitation

For the test cases both Weighting Matrices to investigate their differences on the result. For the first test case and using the Reciprocal of the Variance as a Weighting Matrix we see the results in figure 6.4. For the Weighting Matrix from Signal to Noise the results are seen in figure 6.6.

Starting for the case using the reciprocal of the variance, the variance is determined for 20 independent measurements. When looking at the effective singular values it is clear that this method affects

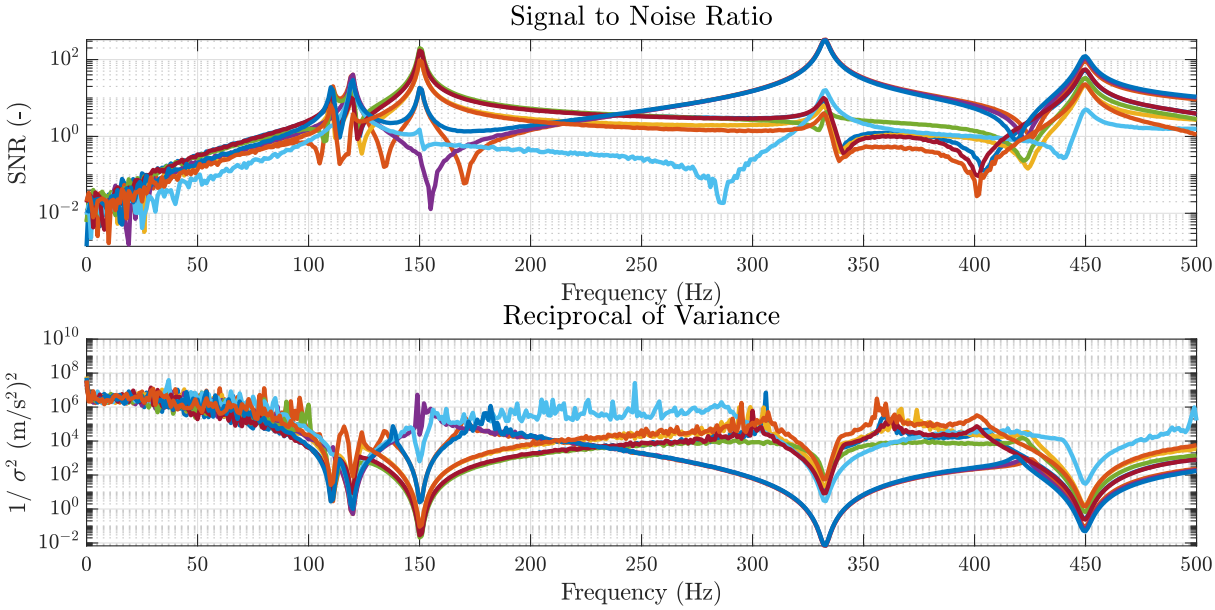

Figure 6.8: Weighting Matrix diagonal entires

the weighting of all dynamic modes. This is also seen in the equivalent force. Due to the increase of the singular values, the noise propagation is also larger. This is well observable at for example $f \approx 160$Hz. We see that the noise level of the solution norm is larger, but also that the equivalent force is approximated to be larger. In this case, we know that this frequency is a problematic frequency due to the lack of controllability. With this method the equivalence result in the $\boldsymbol{u}_3$ stillness plot are actually better than for the unregularised case.

Comparing it to the results from the weighting through the Signal to Noise in figure 6.6, the difference becomes apparent. The solution has less noise contamination from the input data. When comparing the $\boldsymbol{u}_3$ equivalence, the SNR weighting method performs has a less successful equivalent force. Due to the dampening of certain motion in the input data they are also not represented in the equivalent force found. This effect can be clearly seen at the first frequency bins with controllability issues around $f \approx 128$Hz where the equivalence is very poor. Therefore applying the Signal to Noise for the inverse problem can be strong in areas where the solution is close to the noise floor of the solution, but for frequencies where the noise propagation is of less importance it is better to not use it.

Harmonic Excitation

For the second test case we test both methods like we did for the broadband case. The result for using the reiprocal of the variance is seen in figure 6.5 and for the signal to noise ratio in figure 6.7. Similar to the broadband case, the variance is determined from 20 independent measurements. Note that these variances are recalculated for this load case.

The conclusion of the broadband case applies to the tonal case as well. For frequency bins where the solution is around the noise floor the Signal to Noise Weighting can improve the solution. For frequencies where this is not the case, the equivalence is less.

Weighted Least Squares Evaluation

We have shown that using the Weighted Least Squares method can lead to two different results. On one side it is possible to dampen out the specific measurement channels that have poor signal to noise. This can be helpful when the noise propagation becomes a problem and the solution is hidden in the noise level.

On the other side it is possible to do the opposite and focus more on the channels that vary more due to noise. This can help to emphasis dynamics by stiffening the problem for the modes corresponding to that motion. This stiffening shows a larger noise transmission, which is perhaps an unwanted effect, so it should be used only when needed.

Although both methods are similar, they originate from different properties, namely the Signal to Noise (SNR) ratio and the Variance. This similarity is also seen in figure 6.8 from the shape of each channel. Although these both describe the same phenomena, they are quite different. The SNR is dimensionless which makes the physicality of the problem stay more in tact. When using the Reciprocal of the Variance the units do not add up to an equivalent force anymore.

Nevertheless there are effective ways to use the Reciprocal of the Variance in the problem of source characterisation. Since the system matrix is measured from operational data too, it becomes very troublesome to observe the anti-resonant behaviour since these motions are masked by the noise floor of the sensor. These anti-resonances show up as anti-resonances in the singular values. If these are not captured, it means the inverse will also not show these as peaks. This means the problem is approximated to be more compliant than it actually is. With this weighting method it is possible to emphasis these frequencies again, causing a stiffening leading to a larger equivalent force.

6.4.2. L-Curve Method

Another potentially strong method to negate the effects of noise on the inverse problem are the methods related to the L-Curve Method. The L-Curve method is a way to represent the result of Tikhonov Regularisation for many regularisation parameters. Its strength lies in the fact that it visually shows in what region of the Tikhonov solution space you are operating for a chosen regularisation parameter.

$$\boldsymbol{f}_{tikh} = \sum_{i=1}^{r} \phi_i \frac{\left(\boldsymbol{u}_i^T \boldsymbol{u}\right)}{\sigma_i} \boldsymbol{v_i} \qquad \phi_i = \frac{\sigma_i^2}{\sigma_i^2 + \alpha^2} \approx \begin{cases} 1 & \sigma_i \gg \alpha \\ \sigma_i^2/\alpha^2 & \sigma_i \ll \alpha \end{cases} \tag{6.8}$$

Tikhonov Regularisation was discussed in section 3.4.4 and is a method that gives a penalty term to the solution norm. Equation 6.8 shows the form of the solution in its Singular Value Decomposition form, including the filter factors. From the filter factors it can be seen that the method has a big effect on the smaller singular values. To analyse the use of the L-Curve for noise supression we choose a regularisation parameters using both Morozov's Discrepancy Principle and a Wiener filter.

Morozov's Discrepancy Principle

Morozov's Discrepancy Principle states that we desire to solve the traditional least squares problem not in a minimal sense, but up to a certain residual norm. This can be seen in equation 6.9. Note that we can define the exact norm of the noise floor from our noise measurements and the number of channels.

$$\min_{\boldsymbol{f}_{mor} \in \mathbb{C}^n} \|\boldsymbol{u} - \boldsymbol{Y}\boldsymbol{f}_{mor}\| \geq \|\boldsymbol{e}\| \qquad \text{with } \|\boldsymbol{e}\| = \sqrt{n_c}\epsilon \tag{6.9}$$

In section 3.5.2 we have seen that the solution can be of three different kinds, depending on where the solution space is in regards to the minimum noise residual. These three cases are purely determined by the residual and the residual noise floor. We concluded that there are 2 cases where the L-Curve is interesting, and those were the situation where the solution space lies on the bounds (case 2, as seen in figure 6.9) or where the solution space os above the bounds (case 1, as seen in figure 6.10). The third case where the residual norm is below the noise floor for the entire solution space are of no interest, since that would indicate either no motion, or such a poor signal to noise that practically no viable information is measured.

Morozov's Discrepancy Principle originates from significance arithmetic, meaning we cannot solve for more accuracy than the accuracy of the input data. The notion can be extended by also defining a lower bound for the equivalent force based on the system properties and noise levels. In equation 6.10 the system matrix norm is described in terms of the maximum norm of the motion and force as suggested by Strang [47], and this can be rewritten to find an lower bound for the equivalent force norm.

$$\|\boldsymbol{Y}\|^2 = \max \frac{\|\boldsymbol{u}\|^2}{\|\boldsymbol{f}\|^2} \qquad \rightarrow \qquad \|\boldsymbol{f}\|^2 \geq \frac{\|\boldsymbol{u}\|^2}{\sigma_{max}^2} \tag{6.10}$$

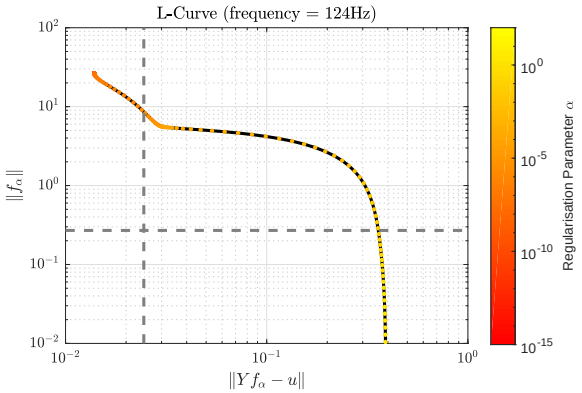

Figure 6.9: L-Curve at $f = 124$Hz

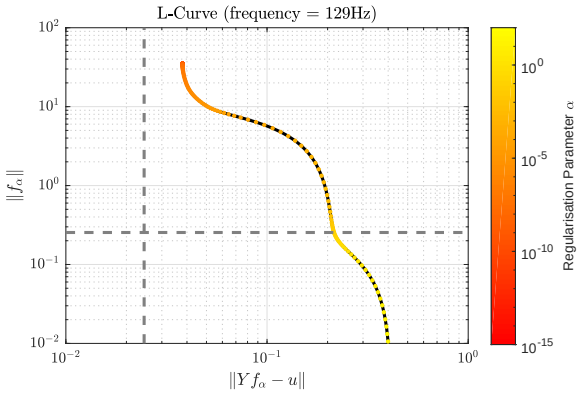

Figure 6.10: L-Curve at $f = 129$Hz

Figure 6.9 and figure 6.10 show two L-Curves for two different frequency bins. The chosen regularisation parameter α is depicted in a range from yellow to red, where red stands for less regularisation. The residual motion noise floor and solution norm noise floor are shown as a vertical and horizontal dotted line respectively. The frequencies depicted are around the frequency area where we have a controllability problem.

Following Morozov's Discrepancy Principle, the best solution is found for the smallest regularisation parameter where the residual norm noise floor is just met. In the case of figure 6.9 this would lead to a solution norm that is around 1 order of magnitude smaller than the unregularised solution. This deviation seems very large. At this frequency the overall signal to noise is around a factor 50. Putting it into perspective, that means for this frequency the solution is dampened by a factor 10 to suppress $1/50^{\text{th}}$ of the motion.

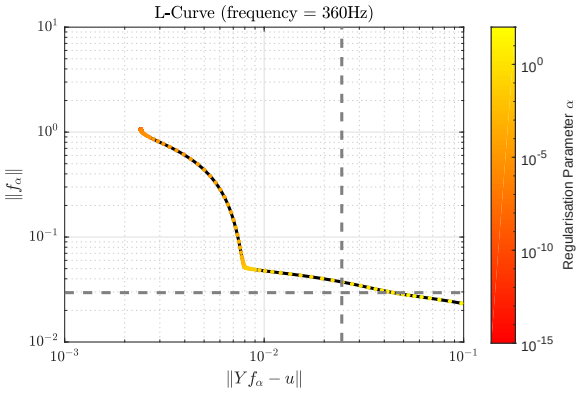

Figure 6.11: L-Curve at $f = 360$Hz

This effect is only worsened if we look at more typical cases of the L-Curve. The L-Curves around the frequencies with the no full controllability the L-Curve are a case 1 situation, and for frequencies further away it becomes more and more a case 2 situation. Figure 6.11 shows a typical case of an L-Curve around a frequency where there are no controllability issues. The solution norm in line with Morozov would be more than an order of magnitude smaller than the true solution.

For cases like in figure 6.10 Morozov's Discrepancy Principle is of little use. It will lead to an unregularised solution to minimise the residual norm. The L-Curve visualisation is of use here however, since we see that the residual norm is staying pretty much constant while the solution norm is still increasing rapidly. We expect the unregularised solution to be estimated too large here, since we are tying to solve an uncontrollable frequency range.

Regularisation Parameter Choice through Wiener

Morozov's Discrepancy Principle seems to be a too aggressive method to suppress the effects of noise. The difference in the solution norm are very large, which begs the question if it is an accurate representation of what is physically happening. A popular alternative to determine the regularisation parameter is the Wiener filter. This method determines the regularisation parameter based on the Signal to Noise ratio, as can be seen in equation 6.11. The Wiener filter method of choosing the regularisation parameter will be used to analyse the two load cases.

$$\alpha_{wien} = \frac{\|\boldsymbol{e}\|}{\|\boldsymbol{u}\|} = \frac{\sqrt{n_c}\epsilon}{\|\boldsymbol{u}\|} \tag{6.11}$$

The reason this method is tried out is two-fold. First of all it seems intuitively a method that makes physical sense, in the way that the solution is regularised stronger for frequencies where there is less signal to noise. As already discussed in the theory section, we expect this to work very well for harmonic source excitations where little is happening on the frequency bins between the harmonics. This method will ensure the equivalent force is estimated to be small in those areas.

The other reason we try the Wiener filter is related to the theory regards the L-Curve. We have seen in section 3.5.1 that the transition between large and little regularisation occurs at the corner of the L-Curve and this approximates the best theoretical solution [29]. The Wiener filter approximates the regularisation parameters to be in the corner quite closely for frequency bins similar to the case we have seen in figure 6.11. This means that the regularisation choice through Wiener has less regularisation than through Morozov. To show this is indeed the case, a number of L-Curves have been plotted in a 3D fashion where the frequency bin increases over depth, as seen in figure 6.14.

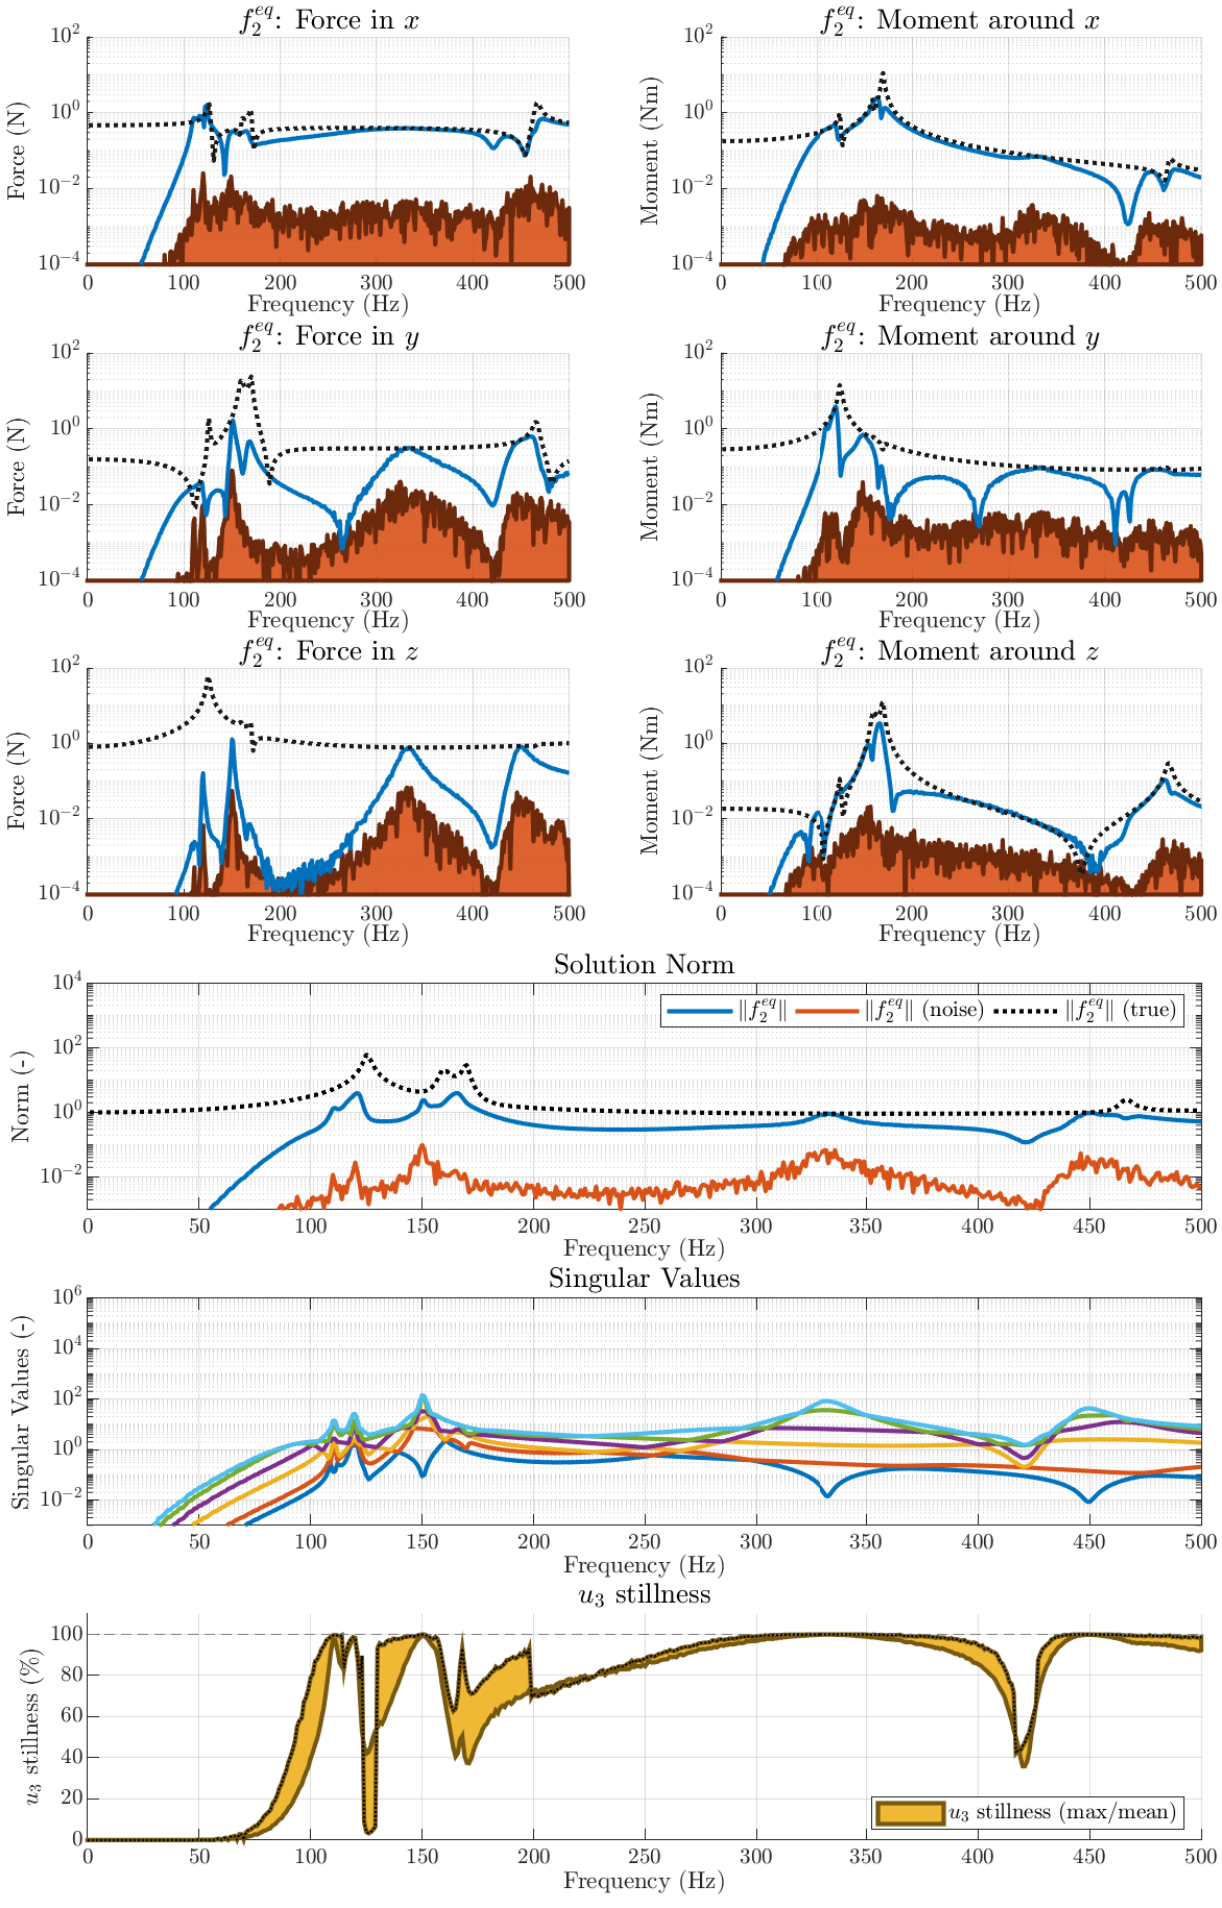

Figure 6.12: Overview Wiener Filters inverse (broadband excitation)

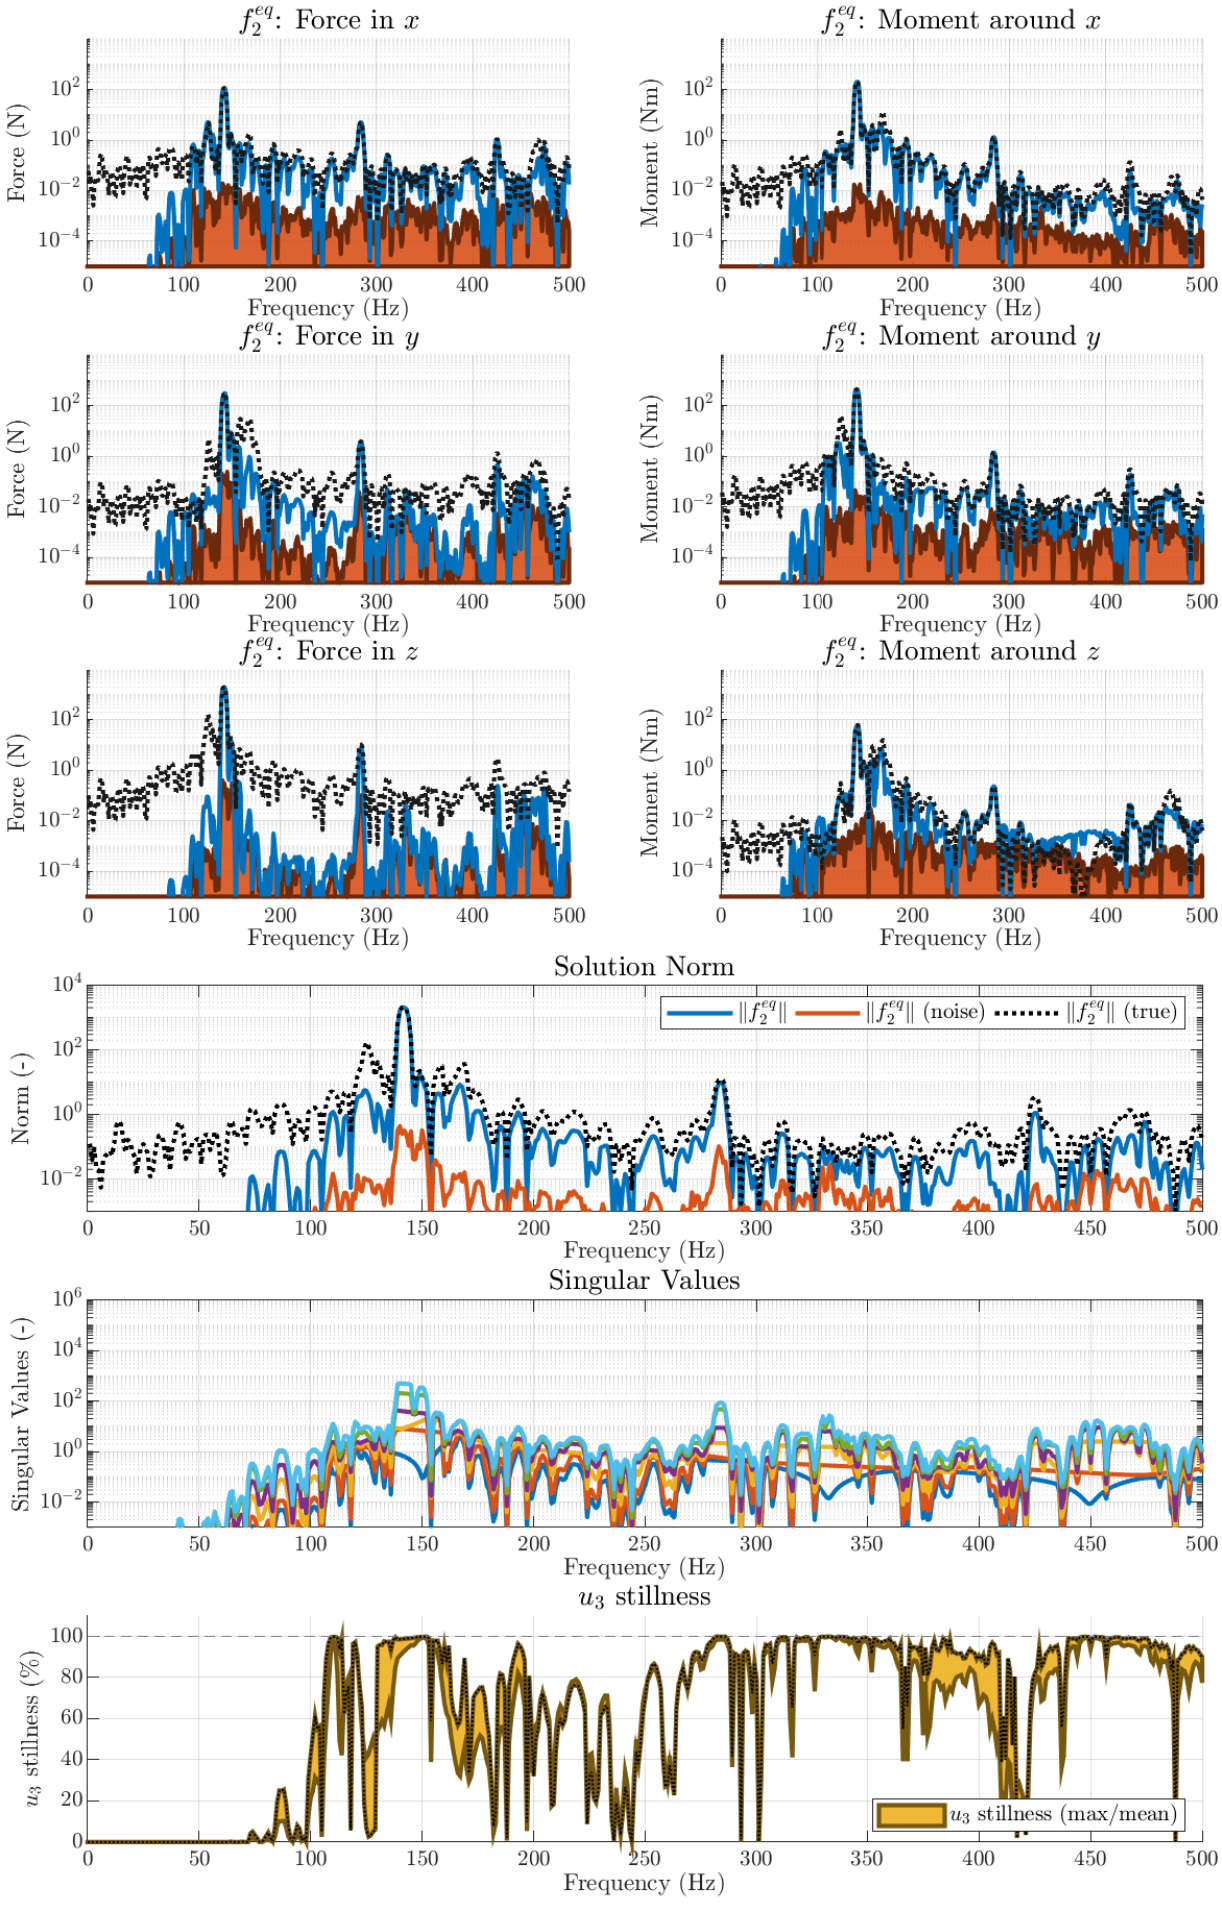

Figure 6.13: Overview Wiener Filter inverse (tonal excitation)

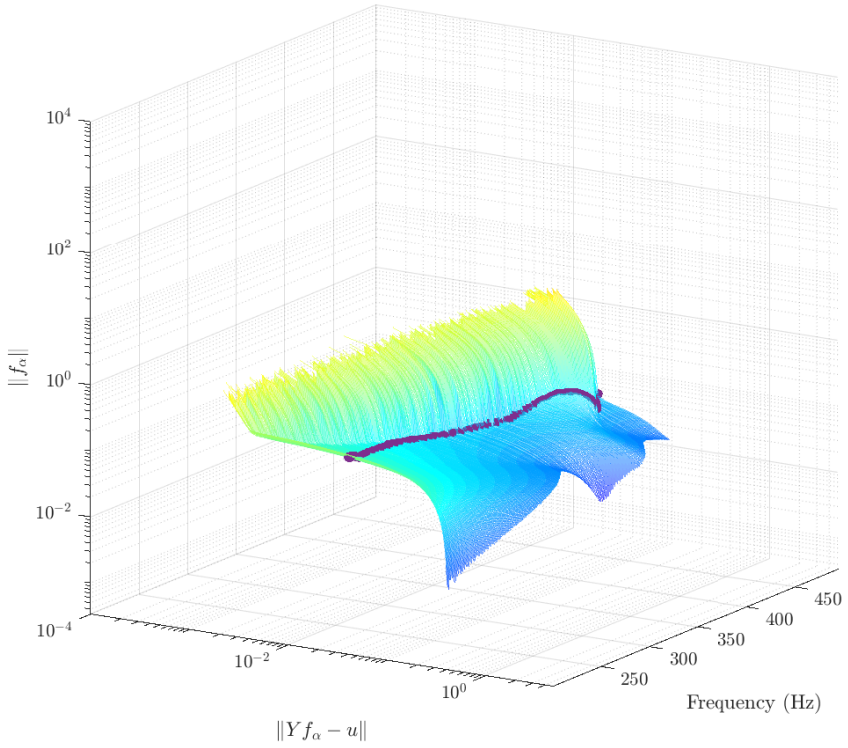

Figure 6.14: Surf plot L-Curve over frequency with Wiener Alpha

Broadband Excitation

For the first test case we see the results in figure 6.12. Right away it is clear that through the Wiener filter method the problem is regularised greatly. This is especially clear at lower frequencies. In previous methods we have seen equivalent forces that were totally determined by the noise on the measurement, whereas for the Wiener method the solution is regularised away. This has the unwelcomed side effect that the result at these frequencies is not equivalent by definition, but one could argue that no solution is better than a wrong solution.

This effect is also seen in the effective singular values. For the modes with the smallest singular values the regularisation has the most effect on the inverse solution. Unlike the Weighted Least Squares method, where we saw the weighting of individual modes being changed, for Tikhonov Regularisation all modes are regularised. For the inverse matrix this effect is larger for larger singular values. It shows that it only dampens the response and at no instances a mode has a larger participation. This is as expected from the way Tikhonov regularisation is defined.

At some specific frequency bins the Wiener approximation is equal to the true solution, whereas at other frequency bins the mismatch seems to show a maximum. These dynamics in the solution are due to the Signal to Noise ratio. As an example, at around $f \approx 350$Hz there is a peak in the Signal to Noise, whereas at around $f \approx 420$Hz there is a valley in the Signal to Noise on top of the already present controllability issue. This begs the questions for what noise levels the Wiener filter is able to approximate the solution accurately.

Harmonic Excitation

For the second test case we see the results in figure 6.13. The effect of the Wiener can most clearly be seen at the harmonics of the excitation at $f \approx 140$Hz, $f \approx 280$Hz and $f \approx 420$Hz. At these frequencies the Wiener approximation is equal to the true solution. At the frequency bins between the harmonics the Wiener approximation is far away from the true solution which can be seen in both the equivalent

force plots and the equivalence plot.

Although the first two resonance frequencies are matched very well by the Wiener filter, the same cannot be said about the third resonance at $f \approx 420$Hz. This is due to a the controllability issues at this frequency. This last feat is a very interesting fact. It shows that while the Wiener filter is very strong for harmonic excitations due to the Signal to Noise ratio these excitations inherently cause, it is not failure proof.

Evaluation L-Curve Method

The L-Curve is a strong method to get insight in the parameter choice for Tikhonov Regularisation. Multiple methods to choose the Tikhonov Parameter can be easily compared. It gives a tool for engineers to use and judge the regularisation parameter choice.

Using *Morozov's Discrepancy Principle* to determine the regularisation parameter to solve exactly to the noise level seems not very promising. It over regulates the problem, leading to poor equivalence. Using the noise floor it can be a tool to identify frequency bins for which there is a controllability problem, as we have seen by defining the 3 different possible cases.

A reason for the poor performance of Morozov's Discrepancy Principle could be pinpointed to how the noise floor is defined. We assume that the norm of the residual motion cannot be smaller than the noise floor, while the intent of the method is actually based on significance arithmetic of each individual sensor channel. Because of this the noise floor is estimated too large. It would be beneficial to define a threshold of the noise floor based on the largest specific residual channel to take into account the cases where there will be little signal in some directions, leading to the residual norm becoming smaller due to the nature of the signal, not because of the force.

Using *Wiener Filter* to determine the regularisation parameter regularises the problem to a lesser extend. According to the L-Curve theory, this method gives the ideal regularisation parameter to be on the transition between under- and over-regularisation. Although this might be the case, it is seen that for the broadband case it over regulates the solution significantly, leading to a poor characterisation although the noise is not a real issue.

For harmonic excitations the Wiener Filter seems a strong method, although there are still some open questions. It is clear that the harmonics are captured better than the noise dominated frequency bins in between them. However, the regularisation is purely based on the noise level. The filter has the tendency to over regulate, so when a harmonic will not have a very high amplitude it might very well be dampened out.

6.4.3. Conclusion on Regularisation: Bases on Noise Level

We have seen two different techniques to regularise the inverse problem based on a priori knowledge of the noise level. For the first method the weighting of the input data is changed. This can be done to emphasis or dampen certain motions. The second method used a different approach where the the full system matrix is dampened out for the frequency bins where little motion is measured and the noise is a relative large portion of the signal. Therefore this method focusses on areas with high signal to noise.

With the weighted least square method it's possible to relate each measurement channel directly to its noise level. It relates the inverse problem regularisation directly to the specific noise on the data which keep the physical relevance of the method. For this reason its effect on the inverse matrix can also be explained in terms of dynamics. The measurement channels that were either emphasised or dampened belong to modes that had relative less motion, meaning they were stiffer modes. By emphasising these channels, we basically say the measurement was larger. In turn the equivalent force becomes larger to reproduce this emphasised motion. Similarly the opposite occurs when damping these modes.

Extending the evaluation to the potential uses of the method in a real life experiment, the weighted least squares method can be of multiple use. The first use could be for when certain dynamics are not

captured in the initial system matrix measurement during an In-Situ experiment. This can occur due to poor signal to noise for certain modes, especially around anti resonances. When manually exciting the structure using an impact hammer, these stiff modes are barely excited. In the operational measurement the excitation level is generally higher so these stiff modes might very well be present. To solve the underestimation of these modes in the system matrix, they can be emphasised using this method. This can be useful when the equivalent force is believed to be underestimated. For situations where the noise propagation in the solution becomes a problem, it is desired to dampen the input data.

For the Tikhonov Regularisation based methods we can use the L-Curve which gives a very nice engineering tool to evaluate the characterisation and regularisation choice. The big disadvantage is that all singular values are changed simultaneously to a large extend. The physical representation of the real problem is lost in this manner. For this reason the transferability of the equivalent force created with these methods will perform worse.

6.5. Regularisation: Based on System Properties

The second class of regularisation methods will be regularisation methods that are based on the system properties. The system properties are determined by the physical properties of the source and receiver structures, the specific source excitation, the modelling of the interface and the sensor set-up.

From the theory of chapter 3 we recognise a few methods which could be specifically used to regularise the system based on its properties. The first method that will be discussed is the method of Truncated Singular Values. In this method the rank r system is truncated to a $r - n_{tsvd}$ system where n_{tsvd} is the number of truncated modes. This method has different implementations which will be analysed.

1. - Truncated Singular Value Decomposition
   - Number of singular values
   - Threshold

Mode Participation

As a first step of the analysis we take a look at what dynamics are captured in the measurement data. The reason this is done is to get a quantification of the effect of regularising modes. If we look closely at equation 6.6, we see that this mapping is at the basis of how the inverse solution is found. The only extra step needed is to map it in the corresponding right singular vector mode shape and weight it using the inverse of the singular value.

$$\boldsymbol{u}_{4i}^{\text{participation}} = \boldsymbol{u}_4 \cdot \boldsymbol{U}_i \tag{6.12}$$

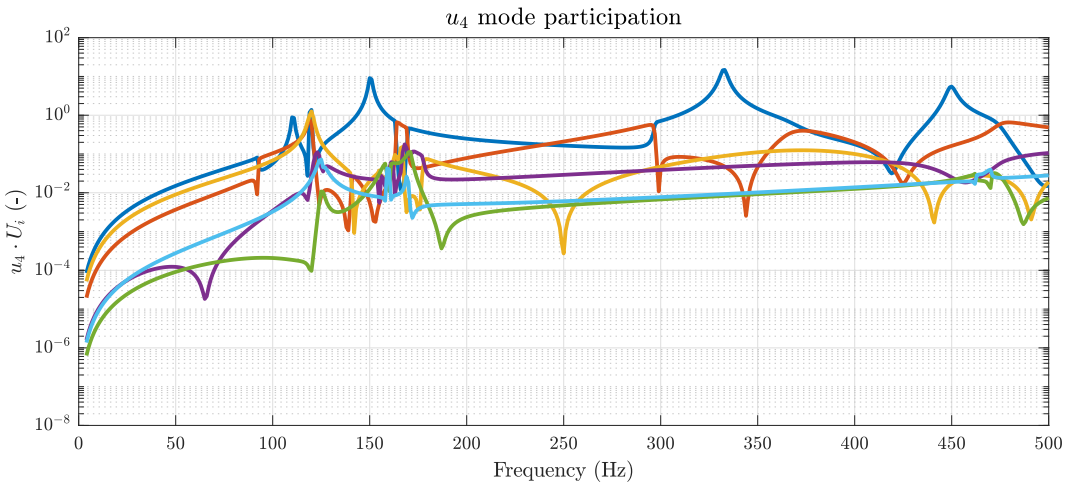

Figure 6.15: u_4 mode participation

Equation 6.12 shows how the participation of each mode is defined. This can be seen in figure 6.15 for a broadband excitation. The colours of each specific mode are chosen equal to the singular value plots as seen in the overview plots. Analysing the plot, it shows that at many frequency bins a very large part of the measurement is described by only 1 or 2 modes. The three regions with controllability issues show participation of multiple modes with very abrupt transitions.

6.5.1. Truncated Singular Value Decomposition Method

The first method of Regularisation Methods based on the System Properties is the Truncated Singular Value Decomposition Method. We have seen this method in section 3.4.3. This method truncates entire modes from the solution space. This method seems promising since it crops dynamics from the problem without changing any of the other dynamics. This keeps the physical relevance high. We have discussed the two possible implementations of the method.

The first method to choose which singular values to reject is by simply truncating to a user specified rank. In the In-Situ experiment we are describing an effective rank 6 Degree of Freedom interface with a 6 Degree of Freedom force. All the modes are therefore significant and cannot simply be cropped. A manual step one could take is to manually crop modes based on analysis of the $\boldsymbol{u}_4$-participation of the measurement data.

Equation 6.13 shows the implementation where the truncation is determined by Significant Rank Criterion. This method seems very promising as it is directly related to the signal to noise of the measurement. In practice it turns out the threshold is very aggressive and only few singular values are kept, leading to poor equivalence.

$$\boldsymbol{f}_{tsvd} = \sum_{i=1}^{r} \phi_i \frac{\left(\boldsymbol{u}_i^T \boldsymbol{u}\right)}{\sigma_i} \boldsymbol{v_i} \qquad \phi_i = \begin{cases} 1 & \text{if } \sigma_i \leq \frac{\|\boldsymbol{e}\|}{\|\boldsymbol{u}\|}\sigma_1 \\ 0 & \text{otherwise} \end{cases} \tag{6.13}$$

As a solution it is possible to add a user defined weighting to the threshold value to increase its size. The threshold criterion is not a precise factor since in its derivation the norm of the measurement is approximated. To add a user specified weighting is therefore not a big departure from the physical relevance of the real life problem. It turns out this method is very effective as an engineers tool to manually determine the amount of regularisation. Both experiments are therefore done by manually tuning the threshold addition of the significant rank criterion.

Using the Significant Rank Criterion leads to an over regulated inverse problem. Almost all singular values are truncated and at some frequency bins all modes are truncated. This shows the method is too aggressive and is of little use for source characterisation. An adaptation that is suggested is to change the threshold definition from a limit on the singular value size to a limit on the mode participation that determines the noise on each mode. This can be done by making use of the $\boldsymbol{u}_4$-participation. This leads to an expression as seen in equation 6.14.

$$\boldsymbol{f}_{tsvd} = \sum_{i=1}^{r} \phi_i \frac{\left(\boldsymbol{u}_i^T \boldsymbol{u}\right)}{\sigma_i} \boldsymbol{v_i} \qquad \phi_i = \begin{cases} 1 & \text{if } \boldsymbol{u}_4 \cdot \boldsymbol{U}_i \leq \|\boldsymbol{e}\| \\ 0 & \text{otherwise} \end{cases} \tag{6.14}$$

This method is very close related to the physical problem we are describing. It filters out exactly the modes that cause noise propagation while not contributing to solution in a physical sense. This method has not been described in literature regarding source characterisation, so for now it will be dubbed the Singular Value Truncation based on $\boldsymbol{u}_4$-participation. This method will be tested for both excitations.

Broadband Excitation

For the first test case we see the results in figure 6.16. Looking at the equivalent forces it is very clear that truncating singular values abruptly causes discontinuous forces and moments over frequency. This can also be seen in the noise floor that jumps up and down. This is not realistic so the question can be

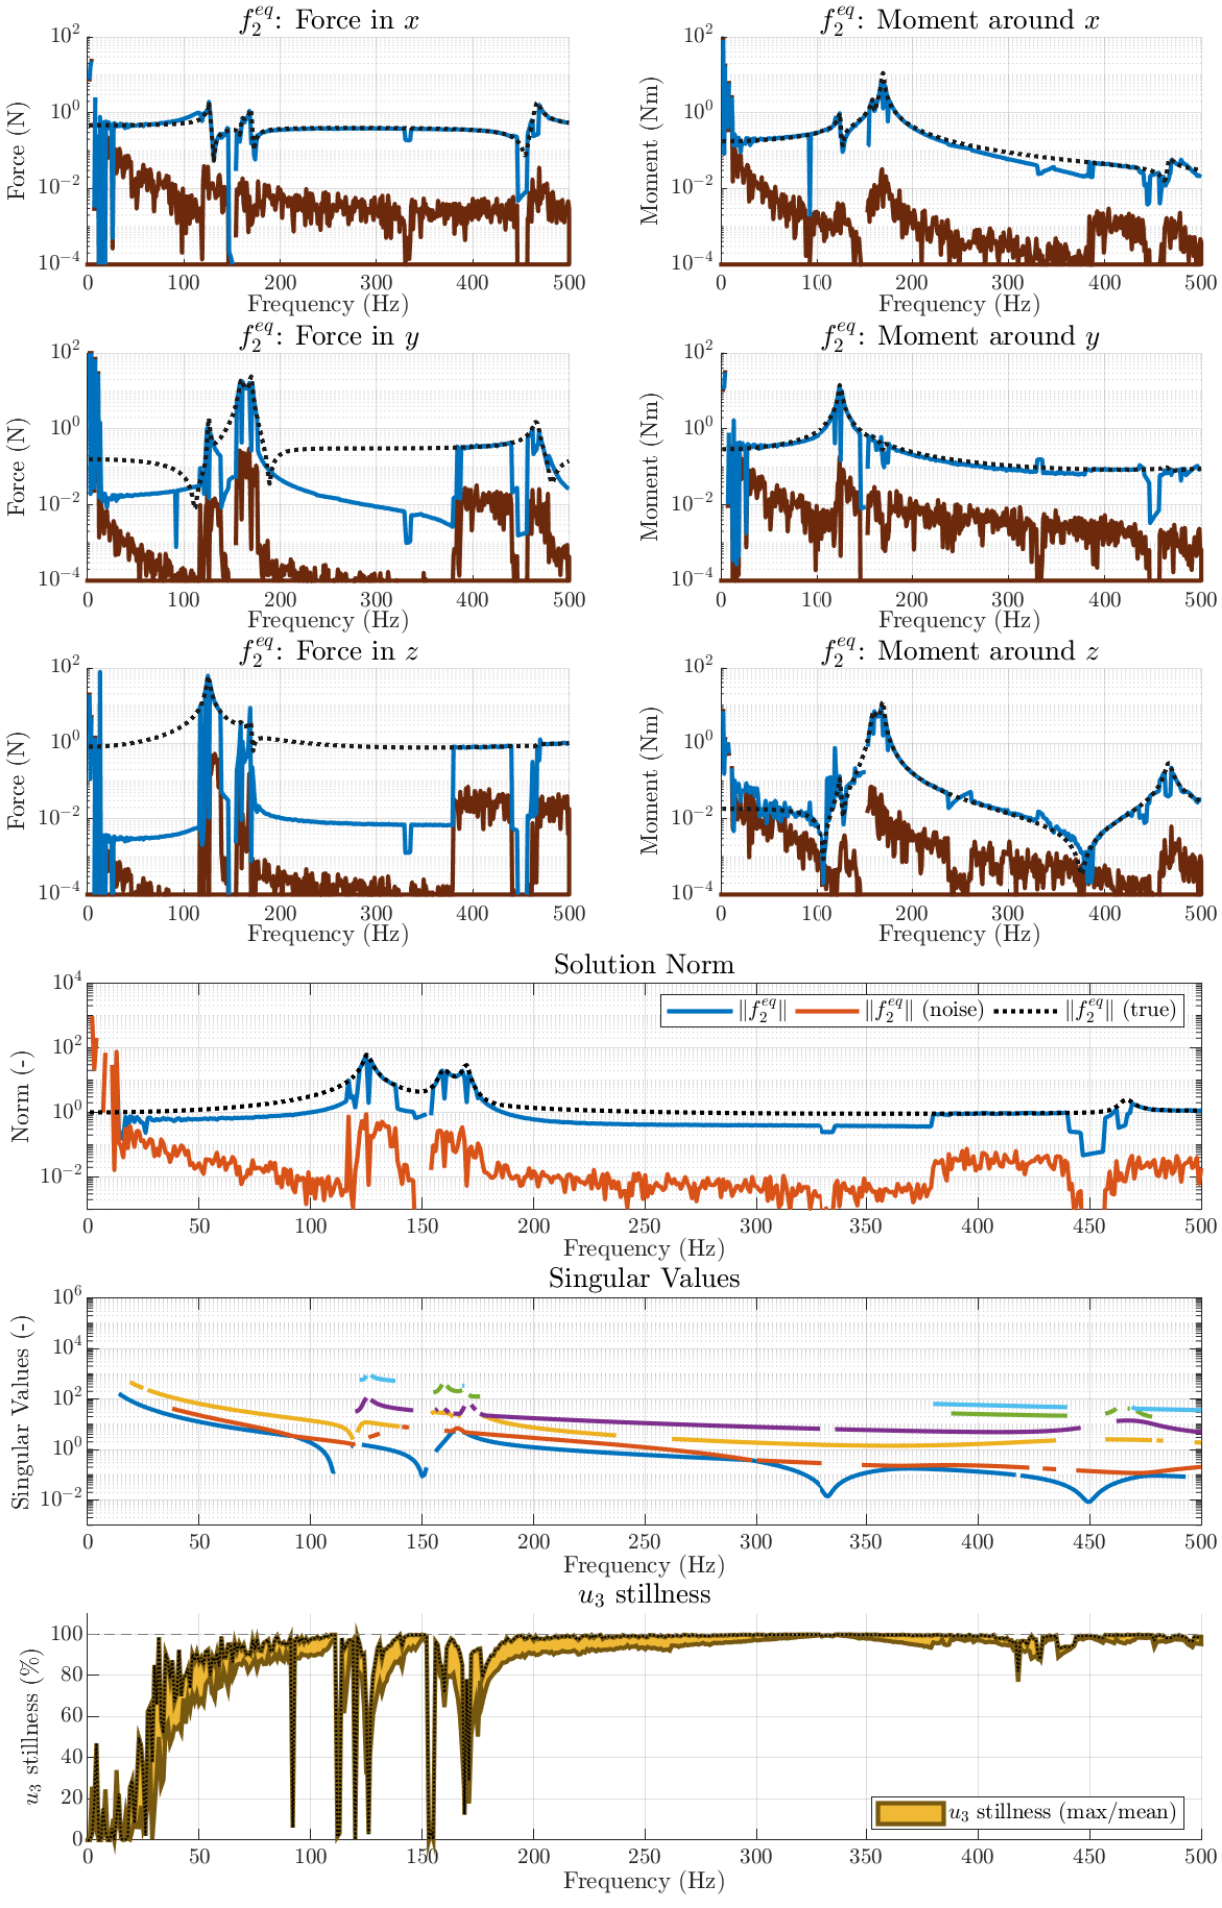

Figure 6.16: Overview Truncated Singular Value Decomposition (broadband excitation)

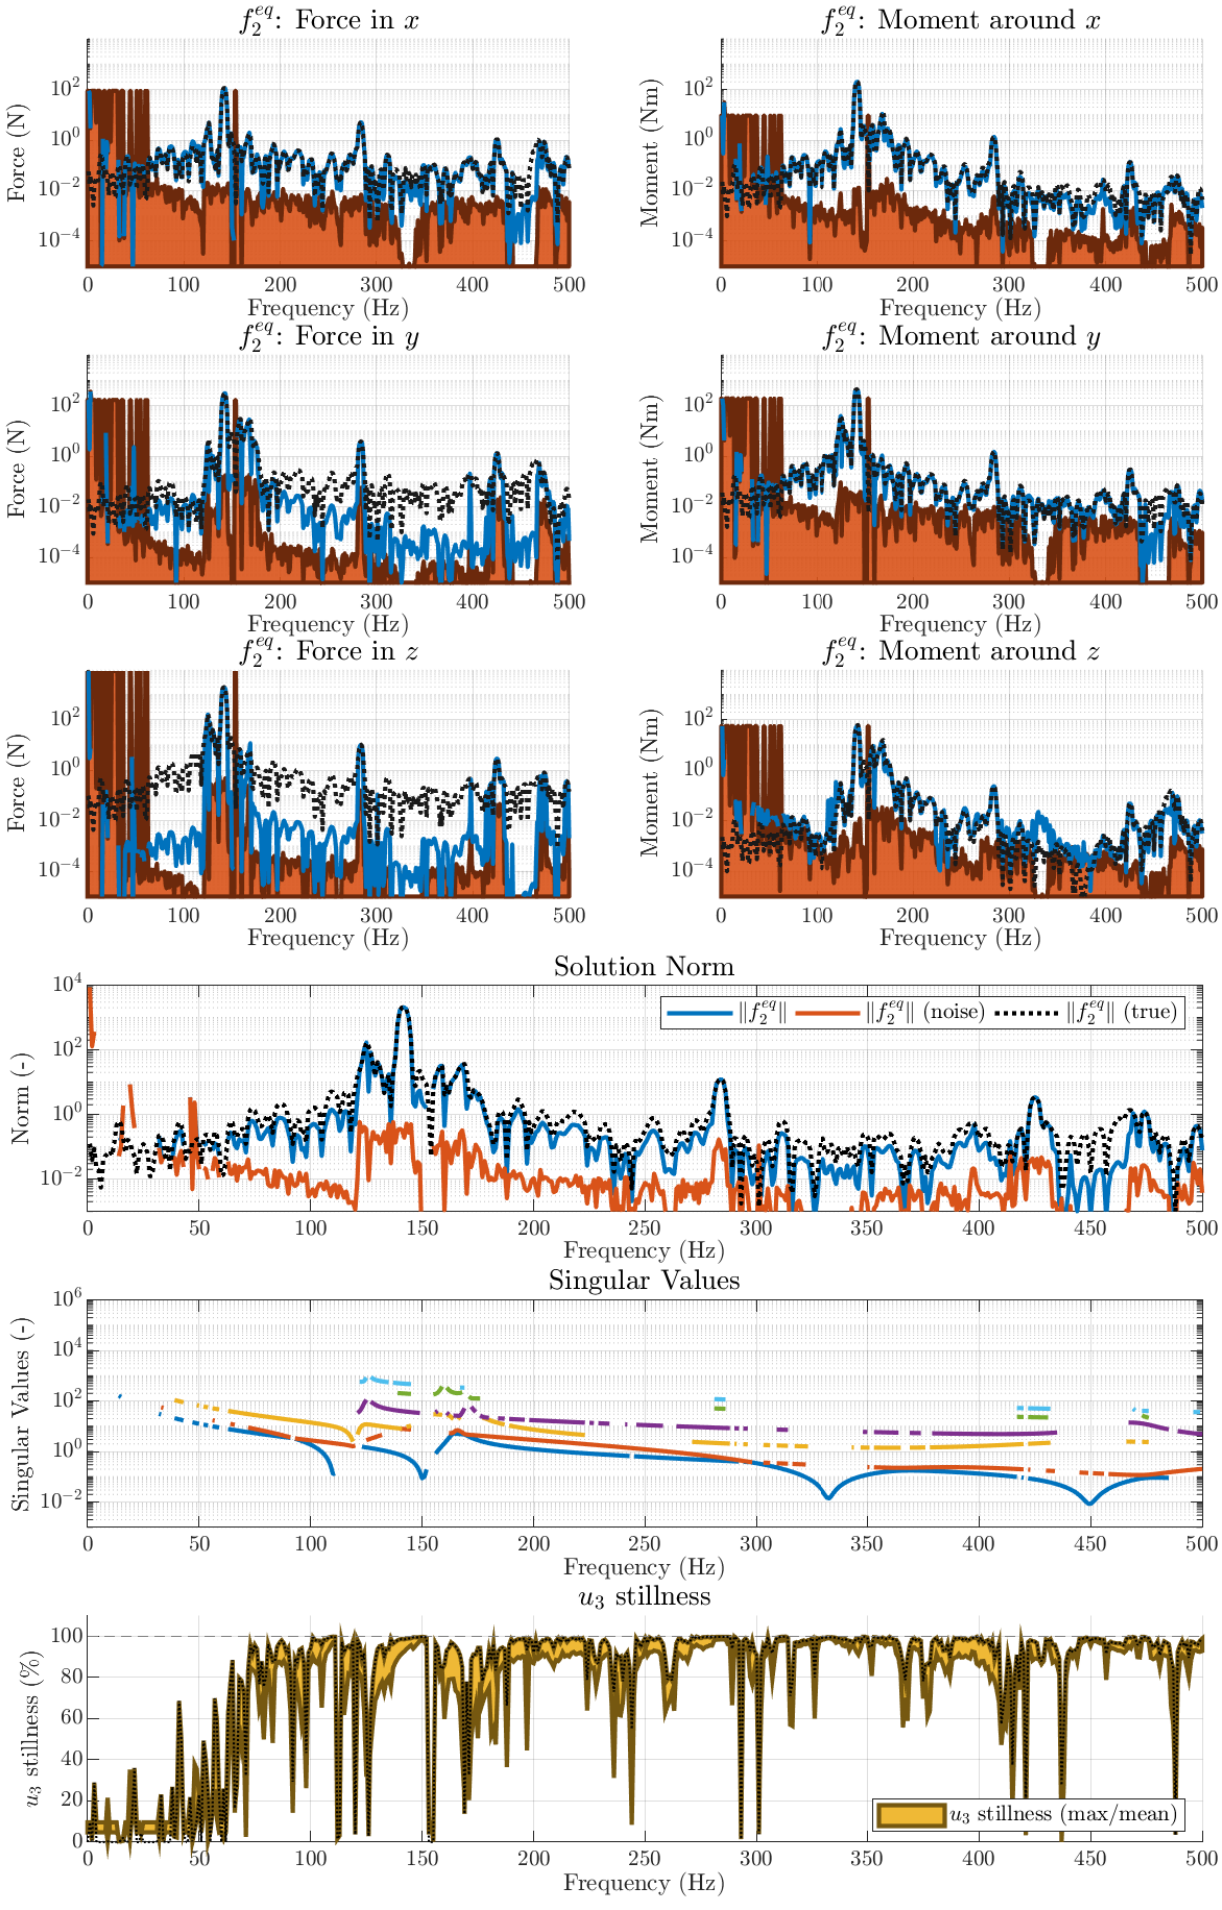

Figure 6.17: Overview Truncated Singular Value Decomposition (tonal excitation)

asked whether truncating singular values is a desired method.

Conversely, when looking at the $\boldsymbol{u}_3$ equivalence plot the characterisation seems to be very successful since the prediction follows the true motion very closely. It can therefore be argues that although the method in which this result is obtained loses some physical representation to the real problem, the result certainly does not and seems better than for the other methods analysed in this chapter.

When looking at the singular values it shows the strength of the method. In all Truncated Singular Value methods seen before, the truncation is based on the size of the singular value. Therefore by definition the truncation happens from above. Here it is clear that some stiffer modes are kept while less stiff modes are truncated, as for example around $f \approx 250$Hz. This means that the yellow mode at this frequency is not participating enough in the measurement for it to be regarded important in the inverse problem. The equivalence plot shows that indeed this truncation has no negative effect on the quality of the characterisation.

Harmonic Excitation

For the second test case we see the results in figure 6.17. Similar to the broadband case the solution is discontinuous due to the abrupt truncation of singular values. Much clearer than for the broadband case, it is seen that all the low frequent modes are truncated and deemed below the noise level.

The same conclusion can be made here that the physicality of the solution found departs from the real world situation in the sense that this solution does not describe what is actually happening. When looking at how equivalent the result is however, the solution seems very accurate. For the harmonics of the source excitation the amplitudes are matched. In between the harmonics the effects of noise are very much dampened. This causes the solution to be estimated lower than the noiseless solution at these frequencies, but the shape of the solution stays correct.

6.5.2. Conclusion on Regularisation: Bases on System Properties

We have seen multiple possible implementations of the Truncated Singular Value Decomposition. The implementations differed in the way the threshold value that controls the truncation is determined. The method that seemed most promising from the literature study was the Significant Rank Criterion, which relates the singular value size to the signal to noise rate and the largest singular value of the system. It turns out in practice this method does not work well, regardless if the problem has a high or low noise level.

A new method is suggested to not truncate modes from the inverse problem based on the size of the singular value, but to look at the participation of each mode in the measurement data. A threshold can be defined from this projection using the measured noise level. Although this is close related to the first class of regularisation techniques, the key difference here is that the system is truncated and therefore altered to a different extend.

The suggested method has the benefit that the truncation is not just dependant on the stiffness of the particular mode. To which extend modes are excited from a load case is not just dependant on how stiff those modes are, but also on the specific direction, position and Frequency Response Function of the excitation. This means that some stiffer modes might participate while other more compliant modes are excited to a lesser extend. Judging on which modes to truncate based on the participation makes it possible to be very precise in the truncation.

The effect this truncation has on the result can be split in two parts. On the one side the solution is found to be very discontinuous due to the discontinuous nature of the singular values. This is a disadvantage of this problem, since some physical relevance is lost. The forces found might describe the problem, but they deviate from the actual force. On the other side, the results this force obtains is very successful in terms of source characterisation. This can be seen from the equivalence plots. Although many modes are truncated, the solution found can represent the source excitation up to the controllability upper bound.

6.6. Conclusion

This chapter set out with the goal to modify the inverse problem with known methods from literature while relating the effects of these methods to the physical problem of source characterisation. By doing this, the regularisation techniques are expressed in the physics that they change of the problem. The increase or decrease of singular values changes the stiffness of the structure and therefore the equivalent force found.

Not all the methods proposed in chapter 3 turned out to be useful methods for source characterisation. For some of these methods this is because they originate from a different domain and are used for different types of problems, whereas for others these are specifically defined for source characterisation with an equivalent force, but still perform poorly. Summarising all methods, they can have the following effects:

Dampen Individual Modes Based on Noise: The inverse problem can be adapted so that measurement data that is deemed less reliable due to error-contamination can be given less weight in the least squares solution. This can be achieved by making use of the Weighted Least Squares Method where the Weighting Matrix is determined by the Signal to Noise ratio of each sensor channel on the diagonal.

This method is very strong in reducing the noise propagation from the input data to the solution space. Due to the overdetermined nature of the inverse problem it makes sense to focus on the input data that has the least influence from noise, in other words the input data that describes the actual occurring physical phenomena the best.

The disadvantage of the method is that the effect it has on the inverse problem is not just that of damping certain modes, but also emphasising other modes that have a better Signal to Noise Ratio. Although desired for noise propagation, it means the solution will capture certain dynamics to a bigger extend than what should physically be occurring. Nevertheless the method shows promise for use in practice. It can help reducing the noise level of the solution which is especially useful when the solution is close to the noise floor. It can also be used a tool to recognise sensor placements that lead to poor Signal to Noise.

Emphasis Individual Modes Based on Noise: Using the opposite logic from the previous type of regularisation, the inverse problem can also be adapted so that dynamics that are captured poorly in the measurement data is emphasised. By making use of the Weighted Least Squares Method where the diagonal Weighting Matrix is determined from the variance of each measurement channel for multiple measurements.

This method finds it origins in statistical estimators for random data. It gives the tool to emphasis the dynamics occurring at low signal levels so that these dynamics are also captured. This is very useful in practical experiments, since the measurement matrix describes the problem the least accurate for the anti resonances of the admittance matrix. In this way these anti resonances are amplified in the inverse problem, leading to a stiffer response. This could be used as a tool when the solution is deemed too small.

Using this method has the potential danger of over fitting the solution to noise. By stiffening the problem, the noise propagation is also larger. For frequencies where the Signal to Noise is really poor, the entire measurement is defined by the noise. For these frequency bins it makes no sense to emphasis these modes, because it will not emphasis any real dynamics occurring. Therefore this method should be used in areas of poor Signal to Noise, but not for areas where there is effectively no Signal to Noise.

Dampen All Modes Based on Noise: The inverse problem can be changed to a Tikhonov Regularisation problem. This adaptation makes it possible to have an extra term in the minimiser which functions as a penalty term on the solution norm. This makes it possible to balance size of the equivalent force with the equivalence of the source characterisation.

To implement this method an extensive analysis of the L-Curve representation is done. From Morozov's Discrepancy Principle a lower limit of the residual motion can be defined in this L-Curve. This is based on the expected norm of the error on the residual motion which is based on the size of the error on the measurement. As an extension, we proposed to also define a lower limit for the solution norm in a similar manner to how the residual norm limit is determined.

From this analysis the method of choosing the regularisation parameter from Morozov's Principle is deemed to be too aggressive for the inverse source characterisation problem. A less aggressive regularisation parameter method is chosen with the Wiener Filter. This method approximates the theoretical sweet spot of the L-Curve really closely, which in theory means the ideal balance between under- and over-regularisation is found.

In practise the method of Tikhonov is not ideal for source characterisation. The method has an effect on all singular values at the same time, meaning the physical representation of the problem is lost. This is seen in the results for both load cases, where the singular values are changed significantly and the equivalence is poor. There are cases where this method can be useful, and that is when it is desired to dampen all modes for frequency bins where the measurement is dominated by noise. Here it can be argued that no solution or a heavily damped solution is better than a solution that tracks the noise from the measurement.

Truncate Individual Modes Based on Noise: The inverse problem can be adapted in more rigorous manner by reducing the rank of the problem. This can be done with the method of Truncated Singular Value Decomposition. There are multiple methods to determine the threshold for truncating modes.

A method has been suggested to determine the threshold based on the projection of the measurement data on each mode and relating this to the noise level. In this way only the modes that are participating below the threshold and therefore describe purely the noise of the measurement, can be neglected. This is a more robust way than the suggested Significant Rank Criterion which based the threshold on the system properties and not specifically the measurement.

The effects this method has on the characterisation are very strong. The noise propagation is a lot lower than for the unregularised inverse, while all significant motion is still accounted for. This can be seen in the equivalence figures, where this method performs very well while the amount of regularisation is very high. On the other hand, the solution found is to be of a very discontinuous nature due to the abrupt truncation of singular values. The solution therefore has little representation to the physical phenomena occurring. Since the goal of source characterisation is to predict the motion on any receiving structure, it can be argued that the equivalence results outweigh the fact that the solution lost some physical meaning.

Overall Conclusion on Regularising for Source Characterisation

We have shown that there are multiple useful tools in adapting the inverse problem. These tools each have their specific use case and can help in finding a better equivalent force to describe the active source.

Within the grand scheme of source characterisation these methods all have the same problem, and that is that the transferability of the characterisation is lost to a certain extent. In other words, when an equivalent force is found from a heavily regularised inverse problem, it takes a step away from being a property of the source only. This is due to the nature of the experiment where the receiving structure and sensor set-up play such a big role in how much noise contamination is on the measurement. Due to this the regularisation will be different on each different test set-up.

For In-Situ experiments this does not have to be a dealbreaker. The equivalent force found for the target assembly could be from a regularised problem and still describe that problem case very well. For this reason it is useful to have a clear understanding of the possibilities in regularisation. The first 3 classes of tools discussed in the Conclusion should be regarded as tools that can be used on specific

areas. This should be judged by the engineer doing the characterisation. The final method can be regarded as a general way to improve conditioning by removing the modes that are not participating in the dynamics of the problem.

III

Conclusions & Recommendations

7

Conclusions & Recommendations

The success of experimental source characterisation depends on many aspects. On one side there are many practical error source due to the experimental nature of how the process is done. On the other hand there are also modelling choices to be made. These choices determine how experiment is performed, but also how the equivalent force is deduced from the measured quantities. In this thesis some of the assumptions made in the source characterisation methods have been challenged; and with success. An inherent limitation to using an interface force for source characterisation is uncovered. The effects this has on the characterisation can be quantified with the new measures introduced in this thesis. Additionally, the more practical side of solving the inverse problem is analysed. From this a set of guidelines are created which help to determine how to adapt the inverse problem to negate the effects of noise while still obtaining an equivalent measure of the source.

7.1. Conclusions

The conclusions of this work can be summarised in two categories. The first category is regarding the limitations of the interface and the effect this has on the characterisation. This is what followed from chapter 4 and chapter 5. The second category is regarding solving the inverse problem and consists of the work of chapter 6.

Inherent Interface Force Limitation: For a single connection point problem, an equivalent force on the interface has an inherent limitation caused by its position. This limitation is a controllability problem, which means certain modes of motion cannot be excited by a force from the interface.

The effective rank of a real-life interface is 6. Although the rank of an interface force is also 6, it does not describe the interface in the same manner. This can be seen in the difference of the Frequency Response Function that describes the motion over the interface and the Frequency Response Function that for describes motion from the interface. This difference can be expressed in the Modified Interface Completeness Criterion form, from which both the maximum and average mismatch in shape, amplitude and phase can be determined.

The interface limitation can be translated from the mathematical framework to a physical explanation that describes what is occurring. The limitation originates from the fact that the equivalent force is positioned in one location. Motion modes that have a node at this position cannot be excited from the interface force. These frequencies are the eigenfrequencies for the dynamic problem where the source structure interface is perfectly fixed.

The interface limitation can also be expressed as a controllability measure which is defined as to which extend the equivalent force can excite the vector space of the measurement. There are two possibilities to express the controllability, a specific controllability related to the measurement and a global controlla-

bility related to the test set-up. In this manner an upper and lower bound of controllability can be defined.

The effect that the controllability limitation has on source characterisation can be quantified. The residual of the inverse problem can be reduced to the interface space to define the blockedness of the interface. It can also be transformed in the other direction to quantify the motion on the receiving structure. Both measures are defined in such a way that the residual motion is normalised and that the degrees of freedom with the biggest effect are described. It is shown that the non-equivalence of the equivalent force is due to non-equivalent motion of the interface.

The In-Situ source characterisation problem with a single connection point is therefore limited in its accuracy. Depending on the source excitation, an interface force that describes the source fully equivalent is not possible. The troublesome frequencies bins can be determined and the effects it has on the characterisation can be quantified.

Inverse Regularisation Guidelines: The inverse problem can be adapted for multiple use cases based on the noise level and system properties. Guidelines are determined to dampen or emphasis specific modes, to dampen all modes and to truncate entire modes from the system matrix.

The inverse problem can be adapted so that measurement data that is deemed less reliable due to error-contamination can be given less weight in the least squares solution. This can be achieved by making use of the Weighted Least Squares Method where the Weighting Matrix is determined by the Signal to Noise ratio of each sensor channel on the diagonal. This method is strong in reducing the noise propagation from the input data to the solution space. This can be useful when certain measurement channels have poor Signal to Noise.

Conversely, the inverse problem can also be adapted so that dynamics that are captured poorly in the measurement data are emphasised. This can be achieved by making use of the Weighted Least Squares Method where the diagonal Weighting Matrix is determined from the variance of each measurement channel for multiple measurements. This method is strong in emphasising dynamics that are poorly captured. It can help when anti-resonances of the admittance matrix obtained from experiment are not captured well, but their behaviour is deemed important for the inverse problem.

The inverse problem can be changed to a Tikhonov Regularisation problem. This adaptation makes it possible to have an extra term in the minimiser which functions as a penalty term on the solution norm. This makes it possible to balance size of the equivalent force with the equivalence of the source characterisation. Although it is shown that choosing the regularisation parameter through Wiener estimates the theoretical ideal amount of regularisation based on the L-Curve, the method has the tendency to overregulate. It can be of use for very harmonic sources where the motion between the resonances is judged of less importance. Using Tikhonov makes it possible to dampen all modes for these frequencies.

It is also possible to truncate individual modes from the inverse problem which reduces the rank of the problem. This can be done with the method of Truncated Singular Value Decomposition. A new method to choose the threshold is suggested based on the projection of the measurement data on each mode and relating this to the noise level. In this way only the modes that are participating below the threshold and therefore describe purely the noise of the measurement, can be neglected. The noise propagation is declined while the equivalence of the equivalent force is very high. The solution found is to be of a very discontinuous nature due to the abrupt truncation of singular values. It can be argued that the equivalence results outweigh the fact that the solution lost some physical meaning.

7.2. Recommendations

During the research process some questions arose that can be of interest for future work. Additionally, the conclusions of the this work can be extended to multiple other fields which also gives the chance for future research.

Interface Limitation Extended to Multiple Connection Points: This research was based on a single connection point problem. The interface force limitation we uncovered is inherently caused by the way the interface is modelled and its effect can be quantified. The question is how this translates to a problem of multiple connection points. The fixed interface eigenfrequencies will likely be much higher and potentially it causes no problems therefore. The mathematical representation of the problem changes very much for multiple connection points and some of the assumptions that we used, for example that the system matrix is by injective function, does not necessarily hold true for multi connection point cases.

Interface Limitation Effect on Dynamic Substructuring: The interface limitation was found from the point of source characterisation. It can be pinpointed to the way the interface is modelled with the Virtual Point Transformation. In Dynamic Substructuring this same method is used to couple structures into an assembly. This begs the question if controllability problems of the interface force has an effect on the outcome. It is also interesting to investigate how this non-equivalency propagates through multiple connected structures.

Effect of Interface Stiffness on Source Excitation: During the literature research on the different existing Component-Based Transfer Path Analysis methods it became clear that while these methods make it possible to find a source description as a function of the source only, a major assumption must hold true. This assumption is that the source excitation does not change for changing interface conditions, or in other words changing boundary conditions. For some sources this might be the case, but in general it can be assumed that the power of an active source is determined by its power source and that the force felt on the interface is very much a function of the interface stiffness.

Pre-condition Measurement Data for Inverse Problem: The different regularisation techniques show potential for usage in the inverse source characterisation problem. It is worth to investigate whether there are possibilities to not alter the inverse problem, but change the measurement data. This measurement data originates from time domain measurements that are transformed to the frequency domain. Cleaning up the data in one of these intermediate steps can make it possible to keep the system matrix unaltered so that it keeps its maximum physical relevance, while the noise propagation is taken care of before the mapping on the inverse solution.

A

Numerical Case Study: Single Connection Point

The following appendix will be regarding the numerical case study for a single connection point source characterisation problem that is being described in this thesis. For this problem we will construct a model of an active, vibrating source and a passive receiving structure. These structures will be coupled as they would in real life. The goal of the model is to perform virtual experiments and gather data as if we would perform the real life experiment.

In section A.1 the entire model and its set-up will be introduced. We will start by introducing the structures and how we get to a set of Frequency Response Functions from their physical properties. The model will start from a very detailed dynamical description of the problem that will be reduced in a few steps to create a workable model. All assumptions that are made in the model are discussed and analysed. To negate some of the inherent error sources of the Virtual Point transformation, we make use of the numeric nature of the model and make use of Rigid Body Elements of the second type. The model is be used to verify the findings of chapter 6, chapter 4 and chapter 5.

A.1. Set-up

The numerical model is created in MATLAB 2018a making use of the VIBES.technology toolbox version 3.0.0. This toolbox is specifically made for performing virtual experiments. Its strengths lies in its ability to simulate experiments and and combine different domains in one framework. Together with a transparent mathematical basis and the visualisation possibilities it offers it provides a very powerful tool for experimental modelling.

A.1.1. Structures

Figure A.3 shows the substructures used for the experiment. The left side of figure A.3 shows the structure which will function as the source, whereas the right side of figure A.3 shows the receiving structure. The structures are aluminium structures which are connected at a single point with a bolt connection.

These structures were developed by Wernsen [59] and studied extensively in multiple publications by VIBES.technology. The reason these structures are chosen for this research are twofold. First of all the structures have many eigenmodes at low frequencies, meaning a lot of dynamics are occuring at these frequencies which makes the source characterisation problem interesting. The second reason to choose these structures is the practical reason that these structures are available for real-life experimental testing. Although these real-life structures have quite some differences to their numeric counterpart, mostly due to poor tolerances and poor welding in the construction, they still resemble the numeric problem.

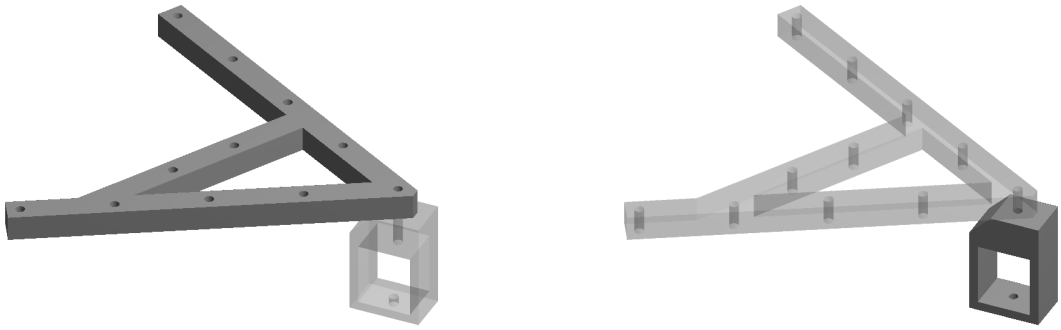

Figure A.1: Source and receiving structure

A.1.2. Structural Dynamics

We seek a dynamic description of the structures in terms of frequency response functions. Structures A and R are modelled as monolithic structures. First the shape is used to discretise the problem in Finite Element Method software. The mesh to discretise is chosen very fine to ensure a converged description of the problem. As discussed in chapter 1, we can linearise the problem to a linearised Mass Matrix $\boldsymbol{M}$ and a linearised Stiffness Matrix $\boldsymbol{K}$.

The first reduction step is to go to the modal domain from the linearised mass and stiffness matrix description. This is done via equation 1.4 as seen in 1.3. The power of this reduction step is that this method builds up an accurate frequency description depending on how many modes you decide to reduce to. What this implies, is that this method loses accurate physical description at high frequencies, while the lower frequency range stays accurate. After this reduction we have a model with the same amount of Degrees of Freedom as in the linearised mass and stiffness matrix approach, but we truncated the high frequent behaviour.

In the modal reduction step we also take account for all the boundary conditions. The receiving structure in this case study is fixed to the ground. This means that the linearised mass and stiffness matrix entries for these nodes are fixed in the modal reduction.

From the modal model we can compute the frequency response functions as seen in equation 1.9. A second reduction step is done here, this time in the number of Degrees of Freedom that are described. The Frequency Response Function is only determined for a select few locations on the structures. These locations represent the excitation locations and the measurement locations. These will be explained shortly.

Coupling

An essential part of dynamic substructuring is the coupling of the two structures. We have seen in section 2.3 that structures can be coupled by using the Virtual Point Transformation on both structures. This makes the coupling of two experimentally obtained models feasible, but also introduces a very big assumption that this coupling is equivalent to the realistic coupling.

To negate this assumption, we couple the structures in a more realistic fashion. In real life the two structures are attached by the use of a bolt connection. This connection will be under significant amount of pre-tension, meaning the surfaces are also linked in-surface to each other due to the friction forces.

This coupling can be done at the substructure Frequency Response Function level, meaning two separate uncoupled structural descriptions are made from the substructures and the coupling happens at

the end. Equivalently, the coupling can be done on the linearised mass and stiffness matrix level. It has been shown that these methods are equivalent and the choice is therefore rather arbitrary.

Figure A.2 shows a close-up of the mesh of both substructures at the interface. To resemble the pre-tensioned bolt connection, a ring of nodes around the bolt hole are chosen to be in force and displacement equilibrium on both sides. By meshing both substructures in the same grid at the interface, this step becomes possible on the linearised mass and stiffness matrix level.

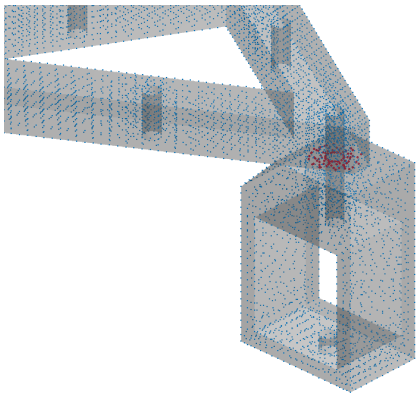

Figure A.2: Coupling nodes of substructures AR

It can be noted that this coupling may seem very blunt and lacks a proper link to the physical representation. This is true to a certain extend, but that is done on purpose. When doing a source characterisation experiment, the interface stiffness will have a big effect on the characterisation. As long as the interface conditions in the characterisation experiment and the real-life application are identical, this problem is negated.

Damping

Up to this point there has been no mention of any damping. In the linearised mass and stiffness matrix representation there was no damping matrix. Since we are not interested in the transient response of the system, damping is of less importance. For a lightly damped system, it will suffice to implement modal damping - damping equal for all modes - or Rayleigh damping - damping dependent on the mass and stiffness matrix.

Since this model is used to characterise and validate the source characterisation, it is trivial what kind of damping is implemented. We choose a constant modal damping for both substructures. There are underdamped with a damping ratio of $\zeta = 0.05$.

A.1.3. Virtual Experiment

With the structural dynamics in place we can perform the virtual experiment. The virtual experiment consists of two main choices, being the excitation and the measurement. The excitations are frequency dependent force vectors somewhere on the source side of the assembly. The measurements are done via tri-axial accelerometers.

Measurements

The measurements are done in accordance to the Transfer Path Analysis theory of section **??**. This means that there will be a set of indicator sensors u_4 and a set of validation sensors u_3. In the virtual experiment these sensors and impacts are modelled as a Virtual Point to close neighbouring nodes of the linearised mass and stiffness matrix.
The indicator sensors are placed closely to the interface and a total of 3 tri-axial sensors are used. This is in line with Wernsen [59], who's conclusion for the sensor placement for Transfer Path Analysis is that the sensors must be located close to the equivalent force and that the amount of measurement

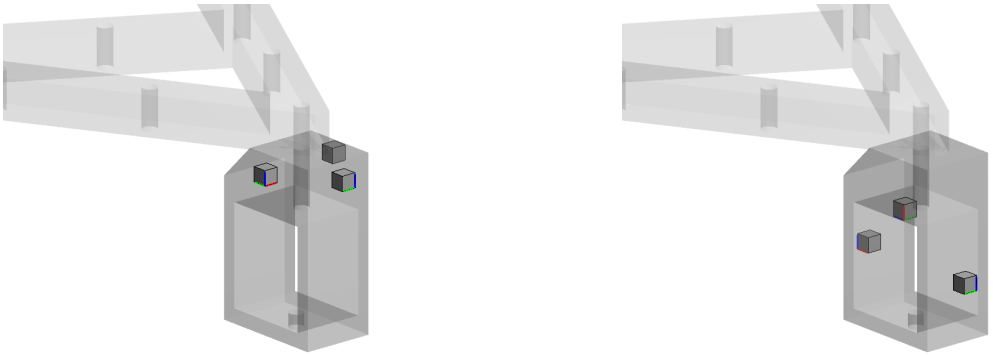

Figure A.3: Measurement Sensors u_4 and u_3

channels must be larger than the equivalent force DoF number. By placing the indicator sensors close to the interface we improve the conditioning of the inverse problem. By using more measurement channels than the DoF number of the inverse problem solution, we create an overdetermined problem.

The validation sensors are placed quite arbitrary somewhere further on the structure. There are no real guidelines for the placement of these sensors as they are only used to validate the equivalent force found. In practise only one tri-axial sensor is used. For this virtual experiment, we can extend our validation to more sensors.

B

1D Schematic Source-Transmission-Receiver

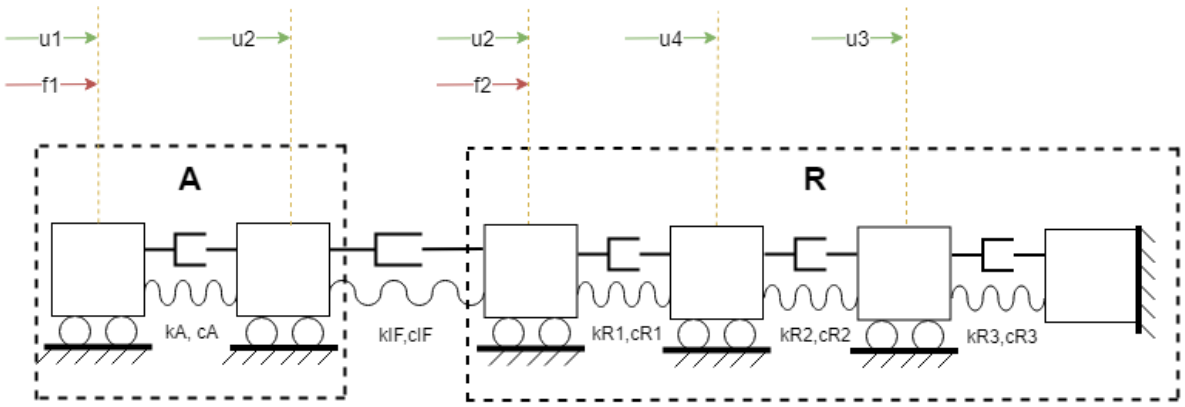

Figure B.1: 1D Schematic Source-Transmission-Receiver

C

Pseudo Force Extension

The controllability issues are caused by the position of the equivalent force. Since the structure is only getting force and moment input from its interface, certain motion modes cannot be excited. This problem seems to be fixable by describing the source with an additional equivalent force that is located somewhere else, a pseudo force.

The position of the forces on the receiving structure are chosen as such that the modes shapes that cannot be excited from the interface are controllable. Since we know that these mode shapes are the mode shapes of the fixed interface resonance frequencies, it becomes easy to choose locations. In this appendix two different methods are tried to characterise the source using additional pseudo forces. The first method is an extension of the traditional method, whereas the second method is fully different.

Method 1: Traditional 6 DoF Interface Force Extension

The idea of the first method is to solve the inverse problem as usual and try to solve the residual motion found with just the pseudo forces. The number of pseudo forces is not limited, but the problem becomes underdetermined if too many are added. For this reason we limit ourselves to 3 pseudo forces.

The method is summarised as follows: Equation C.1 shows the formulation of the known 6 Degrees of Freedom interface force. The residual motion that is left is seen in equation C.2 and follows from the known definition. This is used as the input data for the system seen in equation C.3 that describes the pseudo forces.

Step 1

$$\boldsymbol{f}_{eq}^{trad} = \left(\boldsymbol{Y}_{42}^{AR}\right)^{+} \boldsymbol{u}_4 \qquad \boldsymbol{f}_{eq}^{trad} \in \mathbb{C}^m \text{ with } m = 6 \tag{C.1}$$

$$\boldsymbol{u}_4^{res} = \boldsymbol{u}_4^{meas} - \boldsymbol{Y}_{42}^{AR} \boldsymbol{f}_{eq}^{trad} \tag{C.2}$$

Step 2

$$\boldsymbol{f}_{eq}^{ps} = \left(\boldsymbol{Y}_{4ps}^{AR}\right)^{+} \boldsymbol{u}_4^{res} \qquad \boldsymbol{f}_{eq}^{ps} \in \mathbb{C}_{ps}^{n} \text{ with } n = 3 \tag{C.3}$$

Method 2: 9 DoF Interface Force

The idea of the second method is to solve the inverse problem using the full range of interface force and pseudo forces in one go. Equation C.4 shows the definition of the inverse problem.

$$\boldsymbol{f}_{eq}^{m2} = \left(\boldsymbol{Y}_{4t}^{AR}\right)^{+} \boldsymbol{u}_4 \qquad \boldsymbol{f}_{eq}^{ps} \in \mathbb{C}^m \text{ with } m = 9 \tag{C.4}$$

Analysis

The top plot of figure C.1 shows the results of the first method. It shows that the residual motion in the $\boldsymbol{u}_4$-space is only marginally improved with the extra step. The effects this has on the $\boldsymbol{u}_3$-space is very minor. This extra step can not really be regarded as an improvement. Figure C.2 shows the equivalent forces found. The pseudo forces show resonances on the troublesome frequencies, but their amplitude is very small. Therefore the effect is also minor.

For the second case the $\boldsymbol{u}_4$-residual is not plotted since there is no residual. This is as expected from the Modified Interface Completeness Criterion, where we've seen that the interface can be described perfectly when extending the Degrees of Freedom of the equivalent force. Looking at the bottom plot of figure C.1 however, we see the residual motion in the $\boldsymbol{u}_3$-space. This residual is huge which means the second method found a force that is totally not equivalent.

Figure C.3 shows an overview of the equivalent forces for the second method. The shape of the solution seems like a realistic description of a vibrating source. The interface forces are around 1 to 2 orders of magnitude bigger than the moments which seems logic when comparing the units of both measures. The pseudo forces clearly participate the most at the frequencies where there are interface controllability problems. This analysis has also been performed on the alternative receiving structures that are seen in Appendix REF and it showed that a total different force spectrum was found. The results when transferring the characterisation from one structure to another were totally false. This kind of behaviour was already expected due to the large $\boldsymbol{u}_3$-space residual.

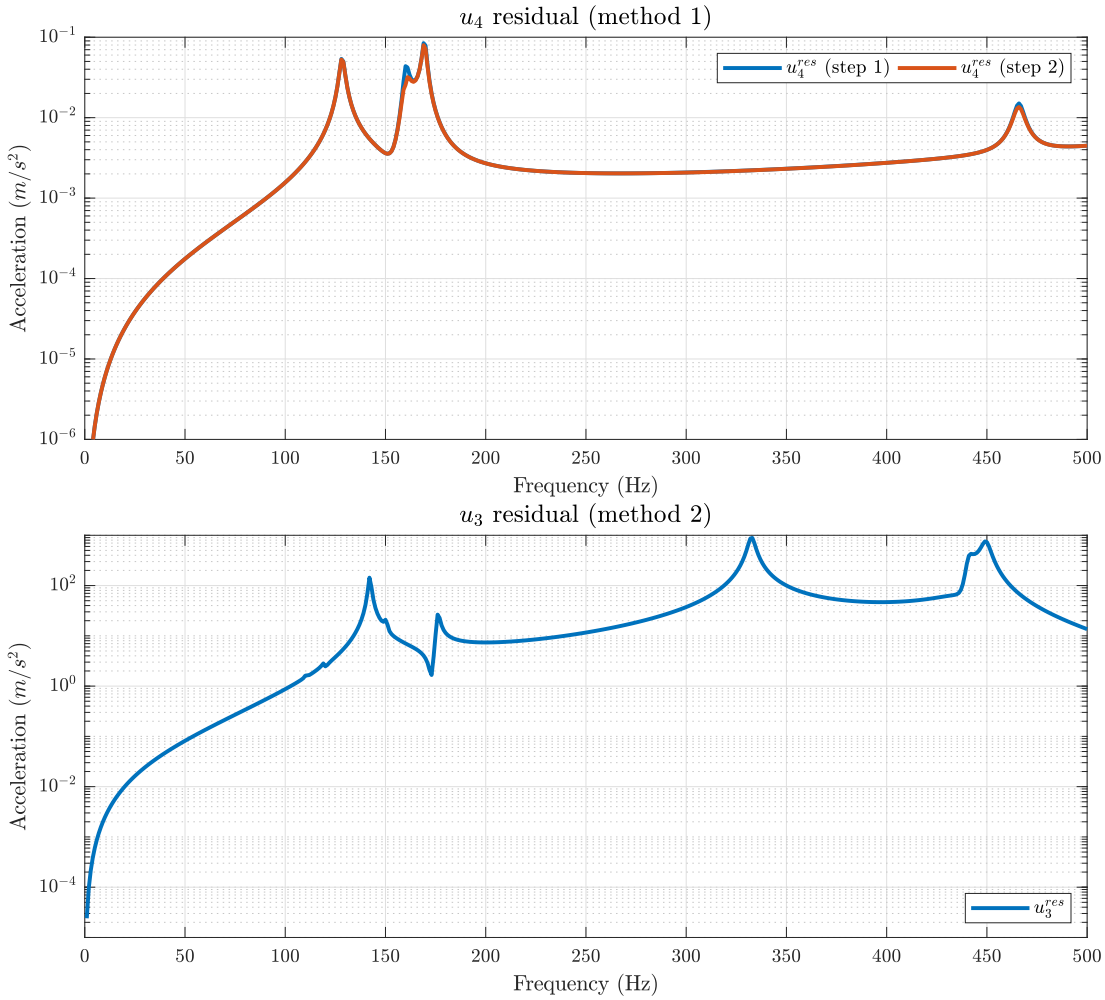

Figure C.1: Residual motion in for method 1 and method 2 (ph)

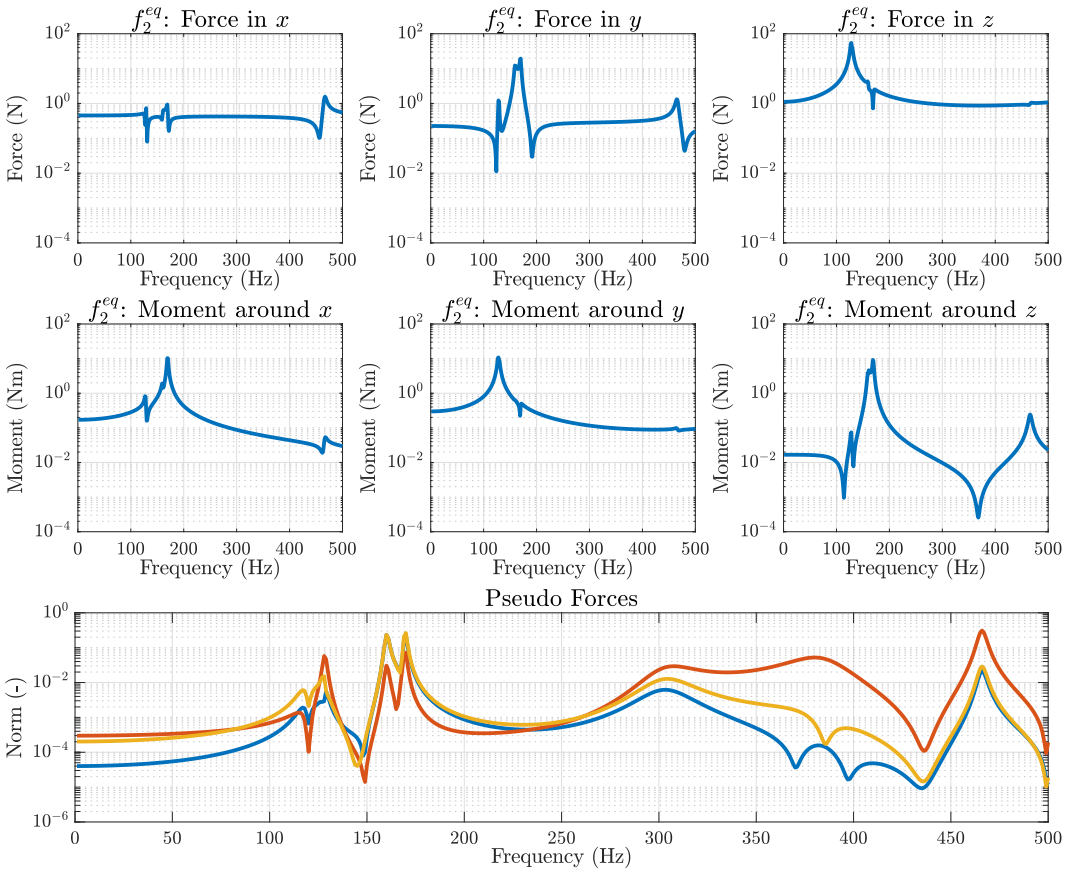

Figure C.2: Overview equivalent force (method 1)

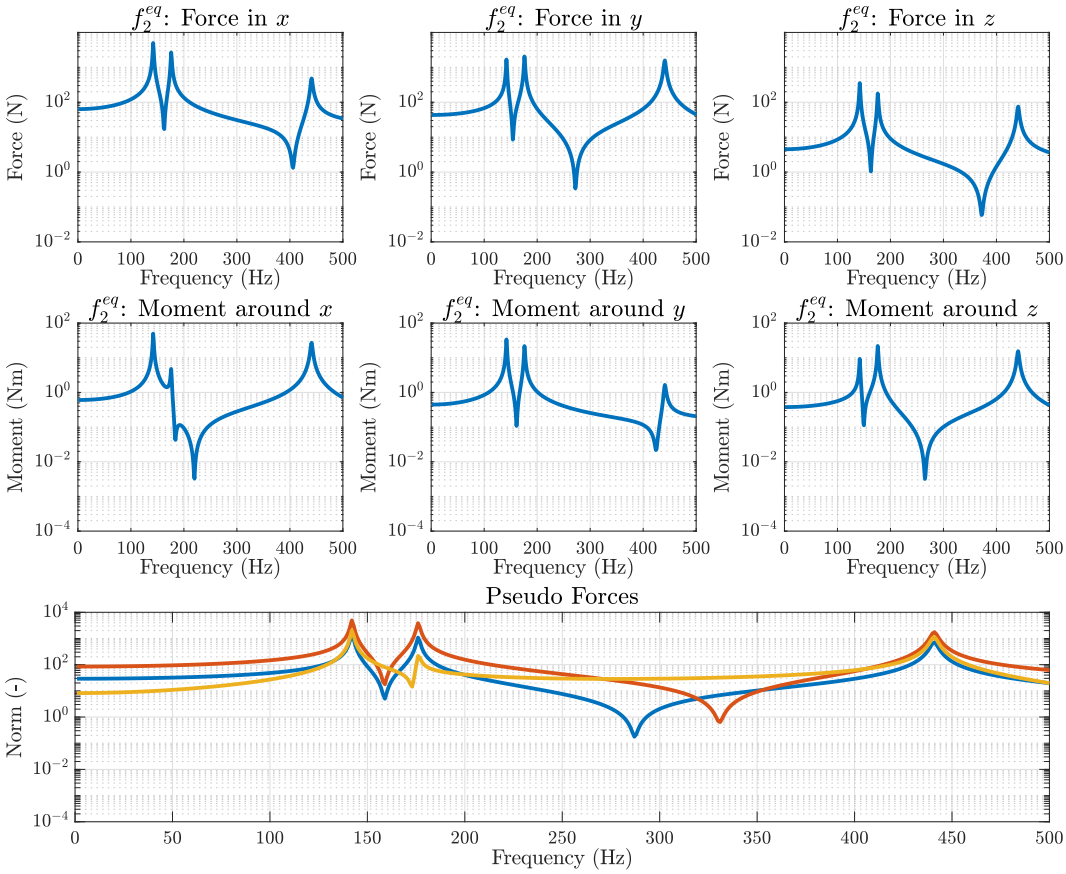

Figure C.3: Overview equivalent force (method 2)

D

Alternative Receiving Structures

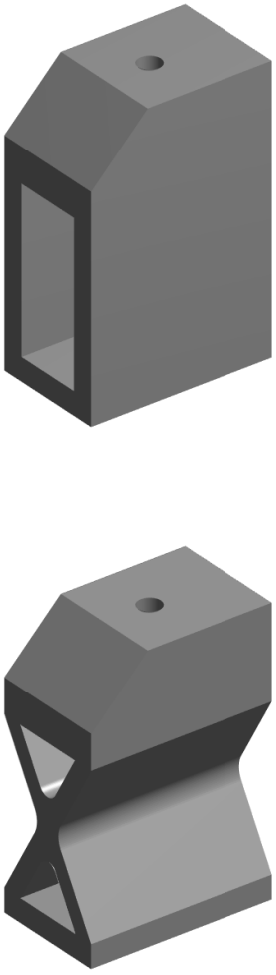

Figure D.1: Alternative Receiving Structures

E

Tikhonov Regularisation Error Estimation

When solving the inverse problem without any regularisation, the error in the solution originates purely from the measurement error. When regularising the inverse problem however, the error is split-up in two parts. One part of the error is due to the measurement error, while another part is due to the regularisation. Because regularisation alters the system matrix, certain information on the system is lost and the solution therefore deviates from the true solution.

To quantify each error term, we seek a description for both terms as a function of the regularisation parameter α and noise level $\boldsymbol{u}^0$, as well as the system properties. Equation E.1 shows the error, both in its most basic form and split up in its two parts. We define the regularisation error as being the difference between the true solution ($\hat{\boldsymbol{f}}$) and the noiseless regularised solution ($\boldsymbol{f}_\alpha^0$). This will give an indication how much the solution deviates from the true solution. The measurement error is defined as the difference between the noiseless regularised solution ($\boldsymbol{f}_\alpha^0$) and the regularised solution including noise ($\boldsymbol{f}_\alpha^e$). In this way the measurement error only takes into account the error due to noise that propagates into the solution.

$$\begin{aligned} \|\boldsymbol{e_f}\| &= \|\hat{\boldsymbol{f}} - \boldsymbol{f_\alpha^e}\| \\ &= \| \underbrace{\hat{\boldsymbol{f}} - \boldsymbol{f_\alpha^0}}_{\text{Regularisation error}} + \underbrace{\boldsymbol{f_\alpha^0} - \boldsymbol{f_\alpha^e}}_{\text{Measurement error}} \| \end{aligned} \tag{E.1}$$

In practise we do not know the true solution or the noiseless regularised solution. However, we can get an estimate of the norm of each error as a function of the regularisation parameter α as will be demonstrated in the next subsections.

Regularisation Error

The regularisation error originates from the difference between the 'normal' inverse and the regularised inverse. This error term is bounded from both sides for all regularisation parameters. For $\alpha \to 0$ both solutions will converge and hence there is no error. For $\alpha \to \infty$ the regularised solution converges to 0. The error will hence be the norm of the true solution. This means the regularisation error is bounded between 0 and $\|\hat{\boldsymbol{f}}\|$.

The unregularised inverse solution norm is calculated using the least square solution as seen in section 3.4.1 and in equation E.2.

$$\begin{aligned} \boldsymbol{Y}^T\boldsymbol{Y}\hat{\boldsymbol{f}} &= \boldsymbol{Y}^T\boldsymbol{u}^0 \\ \hat{\boldsymbol{f}} &= \left(\boldsymbol{Y}^T\boldsymbol{Y}\right)^{-1}\boldsymbol{Y}^T\boldsymbol{u}^0 \end{aligned} \tag{E.2}$$

Using the norm rules from section 3.2.1, the upper bound of the norm can be determined as seen in equation E.3. As expected, the norm of the true solution does not depend on any regularisation

parameters, but solely on system properties (singular values) and the norm of the measurement without noise.

$$\begin{aligned} \|\hat{\boldsymbol{f}}\| &= \| \left(\boldsymbol{Y}^T\boldsymbol{Y}\right)^{-1}\boldsymbol{Y}^T\boldsymbol{u}^0\| \\ &\leq \| \left(\boldsymbol{Y}^T\boldsymbol{Y}\right)^{-1}\| \cdot \|\boldsymbol{Y}^T\| \cdot \|\boldsymbol{u}^0\| \\ &\leq \frac{1}{\sigma_{min}^2} \cdot \sigma_{max} \cdot \|\boldsymbol{u}^0\| \\ &\leq \frac{\kappa}{\sigma_{min}} \cdot \|\boldsymbol{u}^0\| \end{aligned} \tag{E.3}$$

Using the stacked version of the regularised inverse, we can determine the norm of the regularised solution. Equation E.4 shows the regularised solution.

$$\begin{aligned} \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)\boldsymbol{f}_\alpha^0 &= \boldsymbol{Y}^T\boldsymbol{u}^0 \\ \boldsymbol{f}_\alpha^0 &= \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)^{-1}\boldsymbol{Y}^T\boldsymbol{u}^0 \end{aligned} \tag{E.4}$$

It is not as trivial to compute the norm for equation E.4 as it was for equation E.2. This is due to the inverse of the sum of two matrices. The exact solution could be determined using the Woodbury matrix identity. Alternatively, the troublesome term can be replaced by the filter factors (denoted by ϕ_i) as we have seen in section 3.4.4 and seen in equation E.5.

$$\begin{aligned} \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)^{-1} &= \phi \cdot \left(\boldsymbol{\Sigma}^{-1}\right)^2 \\ \| \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)^{-1}\| &= \|\phi \cdot \left(\boldsymbol{\Sigma}^{-1}\right)^2\| \\ \| \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)^{-1}\| &= \frac{\sigma_{min}^2}{\alpha + \sigma_{min}^2} \cdot \frac{1}{\sigma_{min}^2} = \frac{1}{\alpha + \sigma_{min}^2} \end{aligned} \tag{E.5}$$

Using equation E.5 the norm of the regularised solution can be computed as seen in equation E.6. The relation tells us that for a bigger amount of regularisation, the solution norm becomes smaller. This is exactly what we expected.

$$\begin{aligned} \|\boldsymbol{f}_\alpha^0\| &= \| \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)^{-1}\boldsymbol{Y}^T\boldsymbol{u}^0\| \\ &\leq \| \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)^{-1}\| \cdot \|\boldsymbol{Y}^T\| \cdot \|\boldsymbol{u}^0\| \\ &\leq \frac{1}{\alpha + \sigma_{min}^2} \cdot \sigma_{max} \cdot \|\boldsymbol{u}^0\| \end{aligned} \tag{E.6}$$

If we now compute the regularisation error, we will get an indication of the behaviour with respect to the regularisation parameter. Equation E.7 shows the upper bound of the regularisation error. As expected, the error converges to the true solution for $\alpha \to \infty$.

$$\begin{aligned} \|\hat{\boldsymbol{f}} - \boldsymbol{f}_\alpha^0\| &\leq \|\hat{\boldsymbol{f}}\| + \|\boldsymbol{f}_\alpha^0\| \\ &\leq \left(\frac{\sigma_{max}}{\sigma_{min}^2} + \frac{\sigma_{max}}{\alpha + \sigma_{min}^2}\right)\|\boldsymbol{u}^0\| \\ &\leq \left(\frac{1}{\sigma_{min}^2} + \frac{1}{\alpha + \sigma_{min}^2}\right)\sigma_{max} \cdot \|\boldsymbol{u}^0\| \end{aligned} \tag{E.7}$$

If we want to say something about the lower bound, we write out the norm as seen in equation E.8. For regularisation parameters $\alpha \to 0$ the second term also goes to 0 and hence the error is 0.

$$\begin{aligned} \|\hat{\boldsymbol{f}} - \boldsymbol{f}_\alpha^0\| &\geq \|\hat{\boldsymbol{f}}\| - \|\boldsymbol{f}_\alpha^0\| \\ &\geq \left(\frac{\sigma_{max}}{\sigma_{min}^2} - \frac{\sigma_{max}}{\alpha + \sigma_{min}^2}\right)\|\boldsymbol{u}^0\| \\ &\geq \left(\frac{1}{\sigma_{min}^2} - \frac{1}{\alpha + \sigma_{min}^2}\right)\sigma_{max} \cdot \|\boldsymbol{u}^0\| \end{aligned} \tag{E.8}$$

Concluding, the regularisation error is bounded from both sides for all values of regularisation parameter α. It increases for more regularisation and decreases for less regularisation.

Measurement Error

The measurement error originates from the difference between the noiseless regularised solution and the regularised solution including noise. We expect the opposite behaviour than for the regularisation error. For an increase in regularisation, the measurement error is expected to decrease and vice versa. The measurement error is expected to be bound by 0 for $\alpha \to \infty$ and by the noise level $\|\boldsymbol{e}\|$ for $\alpha \to 0$.

To quantify the error, we use the same method as for the regularisation error. The first term in the measurement error will not change from previous derivations in equation E.6. Similarly, the second term is very similar to what we have seen in equation E.6, with the only difference being the measurement term. This term now consists of a true measurement term and a pure noise term.

$$\begin{aligned} \|\boldsymbol{f}_\alpha^e\| &= \| \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)^{-1} \boldsymbol{Y}^T\boldsymbol{u}^e \| \\ &\leq \| \left(\boldsymbol{Y}^T\boldsymbol{Y} + \alpha\boldsymbol{I}\right)^{-1} \| \cdot \|\boldsymbol{Y}^T\| \cdot \|\boldsymbol{u}^e\| \\ &\leq \frac{1}{\alpha + \sigma_{min}^2} \cdot \sigma_{max} \cdot \|\boldsymbol{u}^e\| = \frac{\sigma_{max}}{\alpha + \sigma_{min}^2} \cdot \| \left(\boldsymbol{u}^0 + \boldsymbol{e}\right) \| \end{aligned} \tag{E.9}$$

Using equation E.9 the measurement error can be analysed in a similar way as the regularisation error. Equation E.10 shows the upper bound for the measurement error. For $\alpha \to \infty$ the upper bound converges to zero. This is as expected, since for very large amounts of regularisation the solution norm goes to zero, regardless of whether any noise is present.

The total measurement error can now be computed by taking the difference of equation E.6 and equation E.9. Equation E.10 shows the upper limit of the measurement error. Note that equality sign has switched in the derivation to get a positive valued bound. For $\alpha \to 0$ the upper bound is linear proportional to the size of the noise norm. This is in line with the solution norm we saw in equation E.3. This conclusion is as expected, since the measurement error in the case of no regularisation is indeed a function of the noise norm.

$$\begin{aligned} \|\boldsymbol{f}_\alpha^0 - \boldsymbol{f}_\alpha^e\| &\geq \|\boldsymbol{f}_\alpha^0\| - \|\boldsymbol{f}_\alpha^e\| \\ &\geq \frac{\sigma_{max}}{\alpha + \sigma_{min}^2} \cdot \|\boldsymbol{u}^0\| - \frac{\sigma_{max}}{\alpha + \sigma_{min}^2} \cdot \| \left(\boldsymbol{u}^0 + \boldsymbol{e}\right) \| \\ &\leq \left(\frac{\sigma_{max}}{\alpha + \sigma_{min}^2} \right) (\|\boldsymbol{e}\|) \end{aligned} \tag{E.10}$$

Bibliography

[1] AC Aitken and H Silverstone. Xv.—on the estimation of statistical parameters. *Proceedings of the Royal Society of Edinburgh Section A: Mathematics*, 61(2):186–194, 1942.

[2] Randall J Allemang. The modal assurance criterion–twenty years of use and abuse. *Sound and vibration*, 37(8):14–23, 2003.

[3] Matthew S Allen, Daniel C Kammer, and Randall L Mayes. Uncertainty in experimental/analytical substructuring predictions: a review with illustrative examples. In *International Conference on Noise and Vibration Engineering (ISMA), Leuven, Belgium, Sept*, pages 20–22, 2010.

[4] Christine M Anderson-Cook. The cambridge dictionary of statistics. *Journal of the American Statistical Association*, 98(463):777–779, 2003.

[5] G Banwell and R Faventi. Assessment of experimental techniques to characterise the vibration source strength of a motor radially mounted with resilient elements. *ISMA, Leuven*, 2016.

[6] Julius S Bendat and Allan G Piersol. Engineering applications of correlation and spectral analysis. *New York, Wiley-Interscience, 1980. 315 p.*, 1980.

[7] Felix Benner. Vergleich von regularisierungsmethoden zur bestimmung von blocked forces. 2019.

[8] Jan Biemond, Reginald L Lagendijk, and Russell M Mersereau. Iterative methods for image deblurring. *Proceedings of the IEEE*, 78(5):856–883, 1990.

[9] S Bograd, P Reuss, A Schmidt, L Gaul, and M Mayer. Modeling the dynamics of mechanical joints. *Mechanical Systems and Signal Processing*, 25(8):2801–2826, 2011.

[10] Y Champoux, V Cotoni, B Paillard, and O Beslin. Moment excitation of structures using two synchronized impact hammers. *Journal of sound and vibration*, 263(3):515–533, 2003.

[11] Wei Cheng, Yingying Lu, and Zhousuo Zhang. Tikhonov regularization-based operational transfer path analysis. *Mechanical Systems and Signal Processing*, 75:494–514, 2016.

[12] Hyoung Gil Choi, Anand N Thite, and David J Thompson. Comparison of methods for parameter selection in tikhonov regularization with application to inverse force determination. *Journal of Sound and Vibration*, 304(3-5):894–917, 2007.

[13] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90):297–301, 1965.

[14] Dennis De Klerk. Dynamic response characterization of complex systems through operational identification and dynamic substructuring-an application to gear noise propagation in the automotive industry. 2009.

[15] Dennis De Klerk. How bias errors affect experimental dynamic substructuring. In *Structural Dynamics, Volume 3*, pages 1101–1112. Springer, 2011.

[16] Dennis de Klerk, Daniel J Rixen, and Jasper de Jong. The frequency based substructuring (fbs) method reformulated according to the dual domain decomposition method. In *24th International Modal Analysis Conference, St. Louis, MO*, 2006.

[17] Dennis de Klerk, D Rixen, SN Voormeeren, and Fs Pasteuning. Solving the rdof problem in experimental dynamic substructuring. In *26th International Modal Analysis Conference (IMAC XXVI), Orlando, FL*, 2008.

[18] Armin Drozg, Gregor Čepon, and Miha Boltežar. Full-degrees-of-freedom frequency based substructuring. *Mechanical Systems and Signal Processing*, 98:570–579, 2018.

[19] Andrew Elliott and Andy T Moorhouse. Characterisation of structure borne sound sources from measurement in-situ. *Journal of the Acoustical Society of America*, 123(5):3176, 2008.

[20] AS Elliott, AT Moorhouse, and G Pavić. Moment excitation and the measurement of moment mobilities. *Journal of sound and vibration*, 331(11):2499–2519, 2012.

[21] AS Elliott, AT Moorhouse, T Huntley, and S Tate. In-situ source path contribution analysis of structure borne road noise. *Journal of Sound and Vibration*, 332(24):6276–6295, 2013.

[22] David J Ewins. *Modal testing: theory and practice*, volume 15. Research studies press Letchworth, 1984.

[23] Wang Feng-Quan, Han Xiao-Ling, and Guo Ying-Zheng. Analysis of the characteristics of pseudo-resonance and anti-resonance. 1996.

[24] Peter Gajdatsy, Karl Janssens, Wim Desmet, and Herman Van Der Auweraer. Application of the transmissibility concept in transfer path analysis. *Mechanical Systems and Signal Processing*, 24 (7):1963–1976, 2010.

[25] Jacques Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton university bulletin*, pages 49–52, 1902.

[26] M Haeussler, S Sendlbeck, and D Rixen. Automated correction of sensor orientation in experimental dynamic substructuring. In *Dynamics of Coupled Structures, Volume 4*, pages 65–70. Springer, 2018.

[27] Per Christian Hansen. Analysis of discrete ill-posed problems by means of the l-curve. *SIAM review*, 34(4):561–580, 1992.

[28] Per Christian Hansen. Regularization tools: a matlab package for analysis and solution of discrete ill-posed problems. *Numerical algorithms*, 6(1):1–35, 1994.

[29] Per Christian Hansen. *Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion*, volume 4. Siam, 2005.

[30] Marcel HA Janssens, Jan Willem Verheij, and DJ Thompson. The use of an equivalent forces method for the experimental quantification of structural sound transmission in ships. *Journal of sound and vibration*, 226(2):305–328, 1999.

[31] MHA Janssens and JW Verheij. A pseudo-forces methodology to be used in characterization of structure-borne sound sources. *Applied Acoustics*, 61(3):285–308, 2000.

[32] Bjorn Jetmundsen, Richard L Bielawa, and William G Flannelly. Generalized frequency domain substructure synthesis. *Journal of the American Helicopter Society*, 33(1):55–64, 1988.

[33] Su Jianxin and CM Mak. Direct measurement of moment mobility and a moment excitation system. *Applied Acoustics*, 63(2):139–151, 2002.

[34] D de Klerk, Daniel J Rixen, and SN Voormeeren. General framework for dynamic substructuring: history, review and classification of techniques. *AIAA journal*, 46(5):1169–1181, 2008.

[35] David Lennström, Magnus Olsson, Frédéric Wullens, and Arne Nykänen. Validation of the blocked force method for various boundary conditions for automotive source characterization. *Applied Acoustics*, 102:108–119, 2016.

[36] JWR Meggitt, AT Moorhouse, and AS Elliott. On the problem of describing the coupling interface between sub-structures: An experimental test for 'completeness'. In *Dynamics of Coupled Structures, Volume 4*, pages 171–182. Springer, 2018.

[37] JM Mondot and B Petersson. Characterization of structure-borne sound sources: the source descriptor and the coupling function. *Journal of sound and vibration*, 114(3):507–518, 1987.

[38] Andy Moorhouse and Andy Elliott. The “round trip” theory for reconstruction of green’s functions at passive locations. *The Journal of the Acoustical Society of America*, 134(5):3605–3612, 2013.

[39] Andy Moorhouse, Joshua Meggitt, and Andrew Elliott. Evaluation of uncertainties in classical and component (blocked force) transfer path analysis (tpa). Technical report, SAE Technical Paper, 2019.

[40] Almerico Murli, Luisa D’Amore, and Valentina De Simone. The wiener filter and regularization methods for image restoration problems. In *Proceedings 10th International Conference on Image Analysis and Processing*, pages 394–399. Ieee, 1999.

[41] Michael Papadopoulos and Ephrahim Garcia. Sensor placement methodologies for dynamic testing. *AIAA journal*, 36(2):256–263, 1998.

[42] EA Pasma, MV van der Seijs, SWB Klaassen, and MW van der Kooij. Frequency based substructuring with the virtual point transformation, flexible interface modes and a transmission simulator. In *Dynamics of Coupled Structures, Volume 4*, pages 205–213. Springer, 2018.

[43] Robert E Powell. *Multichannel inverse filtering of machinery vibration signals*. PhD thesis, Massachusetts Institute of Technology, 1982.

[44] AMR Ribeiro, JMM Silva, and NMM Maia. On the generalisation of the transmissibility concept. *Mechanical Systems and Signal Processing*, 14(1):29–35, 2000.

[45] Steven W Smith et al. The scientist and engineer’s guide to digital signal processing. 1997.

[46] Gilbert Strang. The fundamental theorem of linear algebra. *The American Mathematical Monthly*, 100(9):848–855, 1993.

[47] Gilbert Strang. *Computational science and engineering*, volume 791. Wellesley-Cambridge Press Wellesley, 2007.

[48] AN Thite and DJ Thompson. The quantification of structure-borne transmission paths by inverse methods. part 1: Improved singular value rejection methods. *Journal of Sound and Vibration*, 264 (2):411–431, 2003.

[49] AN Thite and DJ Thompson. The quantification of structure-borne transmission paths by inverse methods. part 2: Use of regularization techniques. *Journal of Sound and Vibration*, 264(2):433–451, 2003.

[50] AN Thite and DJ Thompson. Selection of response measurement locations to improve inverse force determination. *Applied Acoustics*, 67(8):797–818, 2006.

[51] DJ Thompson, WJ Van Vliet, and JW Verheij. Developments of the indirect method for measuring the high frequency dynamic stiffness of resilient elements. *Journal of Sound and Vibration*, 213 (1):169–188, 1998.

[52] Daniël van den Bosch, Maarten van der Seijs, and Dennis de Klerk. A comparison of two source characterisation techniques proposed for standardisation. Technical report, SAE Technical Paper, 2019.

[53] Maarten V van der Seijs, Daniël D van den Bosch, Daniel J Rixen, and Dennis de Klerk. An improved methodology for the virtual point transformation of measured frequency response functions in dynamic substructuring. *COMPDYN*, 2013.

[54] MV Van der Seijs. Experimental dynamic substructuring: analysis and design strategies for vehicle development. 2016.

[55] JW Verheij. Measuring sound transfer through resilient mountings for separate excitation with orthogonal translations and rotations. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, volume 1980, pages 723–726. Institute of Noise Control Engineering, 1980.

[56] SN Voormeeren, D De Klerk, and DJ Rixen. Uncertainty quantification in experimental frequency based substructuring. *Mechanical Systems and Signal Processing*, 24(1):106–118, 2010.

[57] F Wahl, G Schmidt, and L Forrai. On the significance of antiresonance frequencies in experimental structural analysis. *Journal of Sound and Vibration*, 219(3):379–394, 1999.

[58] Lei Wang, Xiaojun Wang, and Xiao Li. Inverse system method for dynamic loads identification via noisy measured dynamic responses. *Engineering Computations*, 33(4):1070–1094, 2016.

[59] MWF Wernsen. Observability and transferability of in-situ blocked force characterisation. 2017.