Imperial College London
Department of Mechanical Engineering

Guided wave monitoring of pipelines

Jacob Dobson

December 2015

Submitted in part fulfilment of the requirements for the degree of
Doctor of Engineering in Mechanical Engineering of Imperial College London

Declaration

This thesis is my own work, produced under the supervision of Professor Peter Cawley. Where I have built on the work of others, I have fully acknowledged this, providing appropriate references.

Jacob Dobson

Abstract

Guided wave sensors are commonly used in the oil and gas industry to inspect pipework for damage. Early sensor were detachable and designed for single inspections, but current best practice is to permanently attach sensors and take repeat readings. Known as a monitoring configuration, this approach offers potential benefits in terms of cost, defect detection capability and safety. There are however implementation challenges that are still to be resolved, particularly the need to compensate collected data for changes in environmental and operational conditions (EOCs). The interaction of guided waves with general corrosion is also poorly understood.

This thesis seeks to address both issues, beginning with a finite element study of the interaction of guided waves with rough surfaces. These rough surfaces are a basic model for general corrosion and this study shows that attenuation seen in commercial inspections can be explained by significant scattering of waves by the rough surface. This scattering is shown to be a function of the rough surface characteristics, pipe diameter and inspection wave frequency. This study was computationally expensive and made possible by recently developed simulation tools, particularly graphics processing unit based solvers. The development and construction of these simulation tools is described in detail.

The second major theme of this thesis is improved signal processing for monitoring data based on Independent Component Analysis [ICA]. ICA was applied to numerous example signals and found to correctly separate structural anomalies from surrounding noise. This technique was then compared to best practice techniques from industry (residual processing and singular value decomposition [SVD]) and it was found that SVD and ICA performed much better than residual processing under a range of conditions. In general ICA performed better than SVD, but was less robust to outlier behaviour. The framework used for this comparison is itself novel and is described.

Dedication

I am very grateful to my supervisor Professor Peter Cawley for his support, guidance and encouragement throughout my research. I would also like to thank Professor Mike Lowe for his valuable discussions and advice, particularly in the area of rough surfaces and finite element modelling. Dr Frederic Cegla and Dr Peter Huthwaite must also be credited for creating a stimulating lab environment to work in.

At BP I would like to thank Dr Tom Knox, Dr Robin Jones and Ian Bradley for their academic and financial support. I would also like to thank those at Guided Ultrasonics Ltd. who supported the project, namely Dr David Alleyne, Dr Tom Vogt and Dr Brian Pavlakovic.

On a more personal note I would like to thank my many colleagues at Imperial College and the wider RCNDE, who made my studies more enjoyable.

Contents

List of Tables

List of Figures

Chapter 1:

Introduction

BP is one of the world's largest oil and gas companies, supplying millions of barrels of oil a day to consumers around the world. A vertically integrated company, it has dealings in all areas of the oil and gas industry including exploration, production, trading, and marketing and distribution. As part of this activity BP operates thousands of miles of pipework which it uses to transport hydrocarbons. This pipework is subject to mechanical and chemical degradation and is regularly inspected to ensure its safety: failure of such pipework can have serious environmental and financial implications [1]. This inspection must be ever more reliable and accurate at a time when budgets in the oil and gas industry are under pressure.

1.1. Guided wave inspection

Inspections are mostly carried out using non-destructive testing, with several localised measurement techniques available based on ultrasound [2] and eddy currents [3]. These techniques inspect a small area of the pipe to a high degree of accuracy and usually require direct access to the section of pipe being inspected. Given that pipe networks in the oil and gas industry stretch for thousands of miles, many such localised measurements would be required if a reasonable proportion of the pipe were to be inspected. In some cases it is also simply not feasible to get direct access to the pipe surface, for example due to thick coatings or insulation. Inspections can also be made using 'pigs', inspection and cleaning platforms which are inserted into the pipe and used for cleaning the pipe of sediment and making wall thickness inspections from the inner surface. Such platforms, typically mounted with ultrasound or eddy current probes, can inspect long lengths of straight pipe very quickly but are not suited to the tight bends typically found in petrochemical facilities. Techniques such as radiography and magnetic methods can also be used for pipelines, but come with their own safety and resolution issues [4]. Guided waves offer a solution to this problem as a long-range, full-volumetric coverage non-destructive testing technique for pipelines. Guided waves are a type of low frequency (<100 kHz) ultrasonic wave that propagate long distances along a pipe using the walls as a waveguide. This means a small number of guided wave sensors widely spaced along a pipe can be used to monitor a large region; a single guided wave sensor can typically inspect 50 m or more from a single location [5]. Although the sensitivity is not as high as point thickness measurements (typically around 5% of pipe cross section) this range means that guided waves can be used as an effective screening tool to identify areas which should be followed-up with a high accuracy, localised technique. This type of sensor can also be used to inspect hard to reach areas such as road crossings or under pipe supports, where direct access is not possible.

The principle of guided wave inspections was demonstrated as early as the seventies [6–8] but its widespread adoption was limited by a number of practical engineering challenges. It was not until the early nineties that advances in mode selection and design of inspection equipment made commercial deployment feasible. Most modern commercial systems consist of a deployable ring and an associated control electronics unit, such as the Guided Ultrasonics Ltd. (GUL) unit shown in figure 1.1. In the GUL system the deployable rings have two rows of piezoelectric transducers which dry-couple to the pipe outer wall, with dry coupling possible because at the low frequencies at which this systems operates there is little loss of transduction [9]. These rows are separated by approximately a quarter of the wavelength of the inspection wave at its center frequency, allowing direction control using constructive or destructive phase interference as necessary [9, 10]. Early commercial systems excited a longitudinal guided wave mode L(0,2), where the (n,m) labelling system of Silk & Bainton is followed [6]. This mode had many of the desirable properties for a guided wave inspection mode: it was virtually non-dispersive at the frequencies typically used in an inspection and had low attenuation in steel. This meant the inspection wave could propagate long distances without changing shape. However, the transduction of this mode was difficult since the L(0,1) mode was often excited at the same time. Instead modern commercial systems use the fundamental torsional mode T(0,1), which is non-dispersive and exists at all frequencies. In the range of frequencies

Figure 1.1.: A commercial guided wave inspection sensor produced by Guided Ultrasonics Ltd. Image courtesy of Guided Ultrasonics Ltd.

typically used in guided wave inspection there are also no other torsional waves and so few unwanted modes are excited alongside the main inspection wave. This is sometimes described as 'clean' transduction.

In a typical inspection the sensor excites a toneburst consisting of several cycles of the T(0,1) mode, with waves excited in both directions. The sensor then waits for echoes from pipe features (i.e. the sensor operates in pulse-echo mode), with echoes generated by both benign pipe features such as welds, flanges, supports and tees as well as pipe anomalies. These pipe features will reflect a distorted, amplitude scaled version of the original T(0,1) toneburst. For those features which are asymmetric about the pipe axis there may also be mode conversion of the T(0,1) mode to other modes, with the extent of this mode conversion proportional to the circumferential extent of the feature [11]. To detect defects in a pipe the operator is typically interested in identifying reflections which occur in unexpected places (both in T(0,1) and other modes) and using the extent of mode conversion and the shape of the reflected wave to infer the shape of the defect. Note that detection of defects close to existing pipe features is difficult in this context, since the reflection from the anomalies will be masked by the reflection from benign pipe features. To associate different reflections in the signal with specific pipe features the operator needs to know how the T(0,1) mode interacts with different common pipe features. There are a number of studies in the literature looking at the reflection from common pipe defects, such as flat bottom holes, notches, complex defects and clusters of pits [12–16]. However most of the pipe features studied so far have been discrete features which act over a small axial extent. Features which extend over a large area have been studied less due to the limitations of experimental and simulation work in guided waves. Until recently it has been prohibitively expensive in time and cost to simulate industrial diameter pipes on the necessary length scales of meters, whilst experiments are difficult to perform due to the difficulty of sourcing appropriate test specimens and the time and cost of machining prepared specimens. In this thesis recent work is presented which significantly extends the modelling capability for guided waves thanks to improvements in Graphics Processing Units (GPUs). This allows the study of large area pipe defects such as general corrosion.

Interpretation of reflections is made more difficult by the presence of coherent noise, a deterministic noise that arises due to the excitation and reception of unwanted guided waves [10]. This coherent noise also arises due to low level reflections from general roughness that is not severe enough to be considered a significant defect. Coherent noise is a complex function of the inspection conditions and cannot be predicted from theory or removed through averaging. The coherent noise typically sets the sensitivity of the inspection; for a pipe feature to be detected it must typically produce a reflection which is larger than the background coherent noise. This coherent noise can be reduced through improved transducer design and better compensation of dispersion or can be reduced through a change in how the sensor is used. Instead of having a detachable sensor which makes only a single measurement, the sensor can be permanently installed and operated in what is known as a monitoring configuration. In a monitoring system repeat measurements are made of the same section of pipe over a time interval of days or weeks. This allows current guided wave signals to be compared to previous historical measurements, which allows partial compensation for coherent noise and benign feature reflections and improves defect detectability [17]. This assumes that the coherent noise and feature reflections are constant over time.

1.2. Guided wave monitoring

The trend away from non-destructive testing towards monitoring is seen in a number of applications and sensor technologies. This trend is partly motivated by potential improvements in defect detection and classification but is also motivated by cost and access issues. In an inspection configuration the sensor must be installed and removed every time a measurement is made. Often it is necessary to inspect pipes which are at elevation or buried and so scaffolding or digging equipment must be employed to access the pipe each time. This can be very costly, as well as having health and safety implications for teams working in the dangerous environment of, for example, a refinery. If the sensor is permanently installed, as in a monitoring configuration, these access costs need only be incurred once. Once installed a cable can be run from the sensor to a more convenient location near the ground, making measurement cheaper and safer.

Commercial guided wave monitoring systems exist, such as the guided-wave Permanently Installed Monitoring System (gPIMS) produced by Guided Ultrasonics Ltd. The design and operation of the sensor is very similar to the inspection system produced by this company. This system is capable of repeat readings with little variability which means that, in principle, coherent noise and feature reflections can be subtracted out of current measurements using the data collected when the pipe was in its undamaged condition. This would allow perfect recovery of any pipe structural changes and hence near perfect detection of defects and anomalies. However, the reality is more challenging than this ideal. Between measurements there will be changes in the environment and operating conditions (EOCs) of the sensor, caused by changes in the ambient environmental conditions surrounding the pipe or changes in the pipe operating conditions. These changes in EOCs will cause changes in the signals from the sensor, due to differences in the time of flight of the inspection signals and differences in contact conditions. This means that any subtraction of previous readings will be imperfect unless the EOCs of collection are exactly the same or perfect compensation of EOCs can be achieved. This is known in the monitoring community as the problem of data normalisation [18].

In practice it is unreasonable to change the EOCs surrounding a pipe to match those of a previous inspection and it is not possible to perfectly compensate for EOCs changes such as temperature, as demonstrated by Galvagni [19]. There are however a number of signal processing techniques that can be employed to partly compensate for this data normalisation problem. Many of these techniques have been demonstrated on example guided wave data but there is as yet no rigorous framework for making a quantitative comparison between competing methods. Work outlined in this thesis proposes a new framework for making quantitative comparison between competing processing methods and proposes a novel technique for processing guided wave monitoring data.

1.3. Thesis objectives

This thesis makes two major contributions which improve the performance of guided wave testing in industry. The first contribution is to significantly extend understanding of the interaction of guided waves with rough surfaces in pipes. This is the first known study to look in detail at the scattering of guided waves by rough surfaces in a pipe and helps explain behaviour seen during practical inspections. The study considers the influence of inspection parameters such as inspection wave frequency, rough surface characteristics and pipe diameter. Such an extensive study is only possible if best practice finite element techniques are used. This thesis makes use of a number of tools such as graphic processing unit based solvers and efficient meshing and post-processing tools. Whilst these techniques are not themselves original to this thesis, this is the first known combined use of such tools. The simulation method is therefore described in detail in this thesis.

The second contribution is to improve the post-processing of guided wave monitoring data. This is achieved through the introduction of a novel post-processing technique based in Independent Component Analysis (ICA) and its demonstration on example guided wave monitoring data. This includes a demonstration of how the ICA technique deals with signals from rough surfaces, enabled by the study performed in the first half of the thesis. The superiority of this new ICA technique is then demonstrated in a quantitative comparison with existing best practice. The framework used for this comparison is novel and has potential for use in the monitoring community generally for the benchmarking of competing signal processing methods. Both the analysis framework and its use to compare ICA to existing best practice were performed in collaboration with a post-doctoral researcher at Imperial College, Chang Liu.

1.4. Thesis structure

The first half of this thesis deals with the simulation of the interaction of guided waves with rough surfaces. Chapter 2 begins with a brief overview of existing finite element modelling techniques for guided waves and recent advances in solution procedures based on graphics processing units (GPUs). These new GPU based solvers enable larger, more complex models to be solved in a shorter space of time, which causes new challenges in terms of construction of models and post-processing of simulation data. This chapter describes some of these challenges and the novel combination of techniques devised by the author.

Chapter 3 makes use of the new modelling capability outlined in chapter 2 to perform a study of the interaction of guided waves with rough surfaces in pipes. This chapter investigates how the scattering of guided waves by these rough surfaces is a function of several variables of the guided wave inspection, such as the inspection frequency, the mean depth and correlation length of the rough surface, and the diameter of the pipe. The results help to explain the apparent attenuation seen during commercial inspection of generally corroded pipework, an attenuation which has been unexplained up to now.

The second half of the thesis begins with chapter 4 which is a literature review of techniques for processing guided wave monitoring data. In particular the chapter focuses on the problem of 'data normalisation', the process of compensating monitoring data for changes in environmental and operating conditions. Some existing compensation procedures are considered, as well as the benefits and limitations of these different techniques. This analysis concludes that there is a need for improved processing techniques, as well as a rigorous framework in which to assess different approaches.

Chapter 5 proposes a new technique for processing guided wave monitoring data which is based on combined application of the novel technique of Independent Component Analysis with the established technique of baseline stretch. This new processing technique is demonstrated on example guided wave monitoring data, including example monitoring data from a rough surface.

Chapter 6 describes a validation of the superposition approach, an approach which can be used to generate large sets of guided wave monitoring data with reduced computation and experimental effort. A comparison is made between superposition data and the equivalent data from a full finite element simulation, with favourable results. The limits of this validation are also discussed.

Chapter 7 then proposes a new framework for comparing competing guided wave signal processing algorithms. This new framework relies on the superposition approach to generate large sets of test data, which in turn allows descriptive statistics. This new framework is then used to compare the new technique described in chapter 5 to the existing best practice described in chapter 4.

Finally, Chapter 8 concludes the thesis and makes suggestions for further work.

Chapter 2:

Efficient large scale modelling of guided waves

2.1. Guided wave modes

For simple problems involving isotropic, elastic solids it is possible to model the propagation of ultrasonic waves using analytical expressions. These analytic expressions come from Newton's second law and Hooke's Law and can be shown [20] to be given by:

$$\mu\nabla^2 u + (\lambda + \mu)\nabla\nabla.u = \rho\left(\frac{\partial^2 u}{\partial t^2}\right) \tag{2.1}$$

where μ and λ are the Lame constants, ρ is the density, u is the displacement and t is the time. For unbounded media this equation can be simplified to give solutions to the familiar wave equations:

$$\begin{aligned} \phi &= a_c e^{i(N_c k_c r - \omega t)} \\ H &= a_s e^{i(N_s k_s r - \omega t)} \end{aligned} \tag{2.2}$$

where ϕ and H are the compressional and shear bulk waves respectively, N_c and N_s are the unit vectors for the direction of propagation, k_c and k_s are the wavenumbers and a_c and a_s are the amplitudes.

When boundary conditions are imposed the solutions become more complicated, although for certain simple geometries analytical solutions are still possible. Gazis for example derived wave equations for an infinitely long hollow cylinder (i.e. a pipe) [21], with these wave equations subsequently shown to be accurate by Finch who validated the equations experimentally [22]. These wave equations can be used to show that for a given pipe geometry there are an infinite number of guided wave modes which can propagate in the structure. These waves can be classified as belonging to one of three types:

- Longitudinal L(0,m) modes are axisymmetric with a displacement field $u_r \neq 0$, $u_\theta = 0$ and $u_z \neq 0$.
- Torsional T(0,m) modes are axisymmetric with a displacement field $u_r = 0$, $u_\theta \neq 0$ and $u_z = 0$.
- Flexural F(n,m) modes are non-axisymmetric with a displacement field $u_r \neq 0$, $u_\theta \neq 0$ and $u_z \neq 0$.

In this description the labelling system (n, m) of Silk and Bainton has been used [6], where the circumferential order n is the wavenumber along the circumferential direction and the counter m reflects the type of harmonic motion through the wall. The frequency dependent behaviour of different modes is usually presented in the form of dispersion curves such as those shown in figure 2.1. This particular set of dispersion curves is for a pipe with a nominal outer diameter of six inches (exact outer diameter is 168 mm) and a wall thickness of 7 mm. The dispersion curves show the phase and group velocity of the modes at different frequencies and the cut-off frequencies (those frequencies at which the mode comes into existence as a propagating mode). For practical guided wave inspection the dispersion curves are a key tool to understanding the feasibility of an inspection and the conditions under which it should be performed. Although isotropic solids are considered in this thesis, recent research now also allows dispersion curves to be calculated for materials with anisotropic behaviour [23].

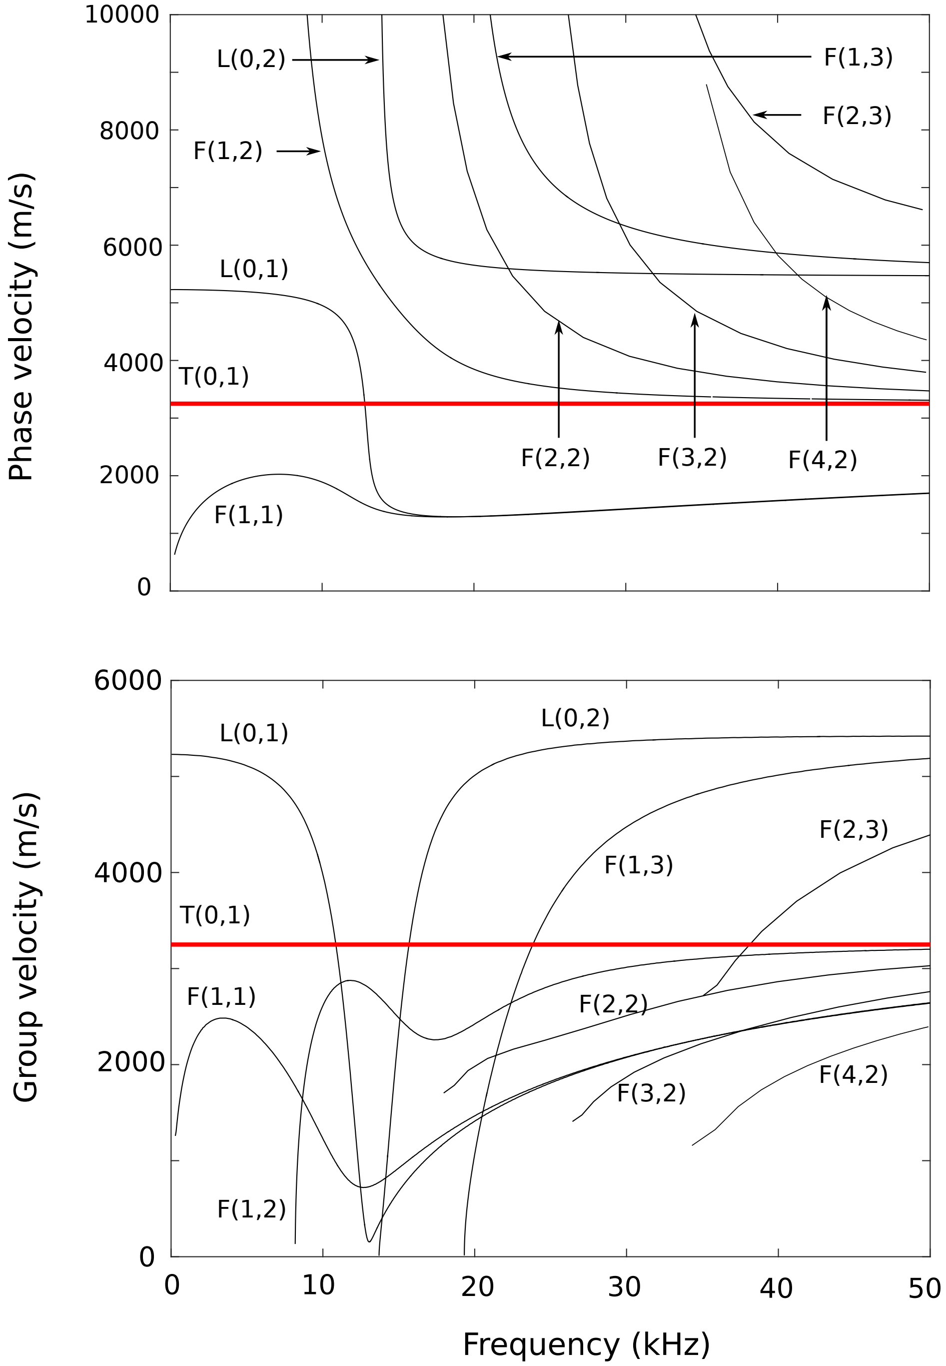

Figure 2.1.: Dispersion curves for a six inch carbon steel pipe with a wall thickness of 7 mm; the fundamental torsional mode T(0,1) is shown in red.

2.2. Modelling techniques for guided waves

As shown in the proceeding section it is possible to calculate wave behaviour in simple solids such as plain pipes. However, once boundary conditions are complicated further, for example to include complex corrosion profiles on the inner surface, analytical equations typically become too complex to solve (unless approximations are used [24]). The analysis of more complex problems is usually done using finite element analysis, a numerical method based on discretization of the solution domain into finite volumes or elements. This is a common technique for engineering analysis and is well described in several sources [25–27]. The accuracy of finite element simulations for modelling of ultrasound will be dependent on several variables of the model, such as the type of element, the characteristic size of the element and the correct specification of boundary conditions. The influence of these parameters has been studied by several authors, in particular Drozdz [28]. In general it holds that accurate simulation of ultrasound requires 10 elements per wavelength in the bulk material, increasing to 30 elements per wavelength close to pipe features. For parameters in the range typical of a commercial inspection this means elements with length scales of the order of 10 mm. However, most of the pipework inspected using guided waves operates on the length scale of meters. If such systems are to be simulated then a very large number of elements would be required, which in turn requires a large amount of computer memory. Simulations would then take a very long time to run, which limits the number of scenarios which can be simulated.

Various techniques have therefore been developed to reduce the expense of finite element simulations. A common approach is to use small elements only where they are needed, for example close to features, and use larger elements in the bulk material. This reduces the number of elements required and hence reduces the time required to solve each time advance of the model. The key challenge with this mesh refinement is that elements with different sizes have different stiffness properties, which creates an impedance mismatch at the boundary between element sizes. This impedance mismatch will introduce a spurious wave reflection which can confuse results and give rise to incorrect conclusions about the feasibility of an inspection. For example, an abrupt change in element size, where the elements have a characteristic length ratio of 1:2, will produce a reflection of around 2 % [28]. This impedance mismatch is difficult to alleviate and is best tackled through optimum choice of element types and thoughtful placement of the transition zone. It is not greatly improved by graduating the element size [28].

Another common simplification to reduce the number of elements is to set up absorbing regions. One classic cause of large models was the need to compensate unwanted boundary reflections from the edges of the model. The brute force solution was to increase the size of the model such that boundary reflections are time separated from the reflections which are being studied. However, the size of models can be greatly reduced if instead absorbing regions are set up at the boundaries of the model. These absorbing regions attenuate any waves passing into their volume, thereby approximating an infinite volume of uniform material with a finite region. This allows the response of a small region to ultrasound to be modelled without interference from other features. Several techniques have been developed over a number of years, with options for both time domain and frequency domain models [29–32]. Time domain absorbing regions based on damping were first mentioned in 1980 by Israeli and Orszag [30] and recently reviewed [31, 33]. In these absorbing regions the stiffness and mass proportional damping are gradually increased through the thickness of absorber. This damping attenuates incident waves but causes an impedance mismatch which generates spurious reflections. This impedance mismatch can be reduced by graduating the damping through the absorber thickness, but this increases the spatial extent of the absorbing region. This spatial extent can be minimised through optimized selection of damping and absorber thickness, with analytical expressions available to guide this parameter choice [28]. An alternative is to simultaneously reduce the stiffness in the absorbing region [34], which attenuates incident waves to the same extent but in a smaller spatial extent. This approach of Pettit et.al. [34] is the approach used in this thesis.

Finite element models can also be reduced in size by substituting regions of the model for analytical expressions which describe the propagation of waves. This substitution is particularly attractive in plain regions of uniform material where the propagation of ultrasound can be well described. An example of this approach is the boundary integration method used, for example, by Shi et al. [35] in the modelling of ultrasound propagation. In this approach small regions surrounding the defect, the source and the monitoring points are modelled using finite element analysis. The solution then proceeds in three stages: source generating signals, defect scattering signals and monitor receiving signals. The complex wave interaction at the source, defect and monitor regions is modelled using finite element modelling. The waves at the boundary of the source box are calculated and then analytical expressions are used to calculate the resultant waves at the edges of the defect box. Finite element analysis is then used to model the waves at the face of the defect (for example a crack) and then boundary integration is used to calculate the resultant waves at the monitoring point. The results of this technique agree well with full finite element simulations and are calculated much faster [35]. However the technique is best suited to small features which are modelled in otherwise large, plain sections of material.

The most significant recent advance in modelling capability for guided waves is in the development of Graphics Processing Unit (GPU) based solvers. Traditionally finite element codes have been solved on the Central Processing Unit (CPU) of a computer. However, the major time constraint in the solution of such models is the expense of loading the data; subsequent operations on the data are relatively inexpensive. This makes the problem suitable for solution on the GPU, where parallel computing and efficient use of bandwidth which under favourable conditions can lead to solution times which are two orders of magnitude faster than the associated CPU solution. Solvers are now available which make use of the GPU for the solution of linear elastic, finite element problems, such as the program POGO [36]. The speed improvement of this program enables much larger finite element models to be solved with reasonable run-times, but at the cost of challenges in terms of meshing and post-processing of data.

2.3. Model meshing and discontinuity smoothing

Although GPU based solvers enable much larger models, this benefit can only be realised if these larger models can be properly discretized with good quality meshes. This requires mesh optimization tools and effective methods of measuring mesh quality, since on the scale of millions of elements visual inspection of the mesh is not always possible. To aid the description of these optimization and measuring tools an example model geometry will be discussed in this section. This model is the same basic geometry as is used in chapter 3 and consists of a six inch, schedule 40 pipe which is 4 m long. This pipe has an outer diameter of 168 mm and a wall thickness of 7 mm with absorbing region at either end, as shown in figure 2.2.

The volume of this model is around $12.7\text{x}10^6\,\text{mm}^2$, which to discretize with elements of 2 mm gives around 6 million degrees of freedom (this gives an indication of the size of the models involved). The first stage in setting up this model is to discretize the volume into tetrahedral elements. Tetrahedral elements were chosen for this model because they are compatible with POGO software and are reasonably well behaved even when distorted

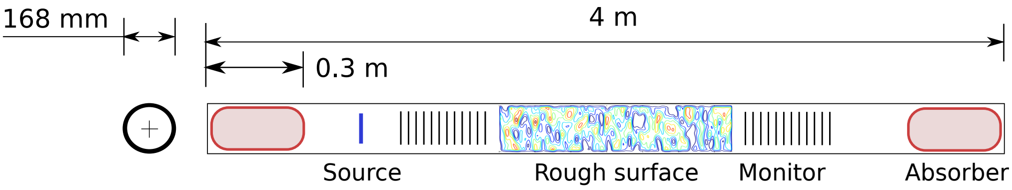

Figure 2.2.: Geometry of a typical large scale guided wave model, solved on the GPU

as shown in [28]. The algorithm used to perform this discretization was Tetgen [37], a standard meshing algorithm for 3D, with meshing performed in Gmsh [38], an open source meshing tool for finite element analysis. To measure the quality of this initial mesh a mesh distortion measure was used which is based on the relative area of different faces of each tetrahedral element. This quality measure γ_k is available in the literature [38] and is given by the equation:

$$\gamma_k = \frac{6\sqrt{6}V_k}{\left(\sum_i^4 a(f_i)\right) \max\limits_{i=1,\ldots,6}(le_i)} \tag{2.3}$$

where V_k is the volume of element k, $a(f_i)$ is the area of the ith face of k and le_i the dimensional length of the ith face of k. A γ_k value of 1 indicates a perfectly distributed tetrahedral element with all angles equal, whilst a value of 0 indicates a fully distorted sliver element. A well formed mesh will have the majority of elements with a high quality measure, with few elements with a γ_k less than 0.5. The distribution of gamma values for the original mesh is shown in figure 2.4.

The next stage in model formation was to distort the inner surface of the pipe to follow a prescribed surface profile. This surface profile is a rough surface with a Gaussian thickness distribution, such as the example surface shown in figure 2.3a. The inner surface of the pipe is made to follow this profile initially simply through adjustment of the node radial position, as shown in figure 2.3b. This adjustment introduces a mesh distortion which decreases the accuracy of the modelling procedure and must therefore be compensated for. The simplest correction procedure is to radially smooth elements, which means that when nodes on the pipe inner surface are radially shifted, nodes in the wall thickness are also radially shifted in proportion to their radial position within the pipe wall, as shown in figure 2.3c. This improves mesh quality but leaves the mesh more distorted than the original mesh; the distribution of mesh quality values for the radially smoothed mesh is shown in green in figure 2.4. This distortion was then further reduced using Netgen

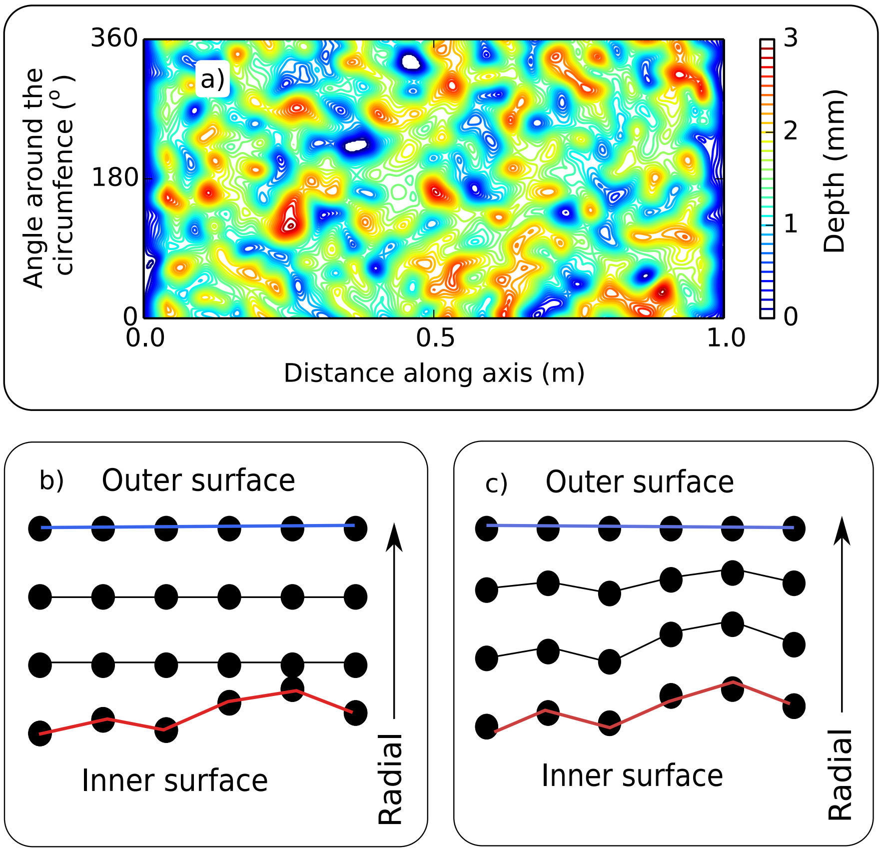

Figure 2.3.: Part a) an example rough surface with mean depth 1.25 mm and correlation length 20 mm, b) an example of node displacement on the inner pipe surface and c) displacement coupled with simple radial smoothing

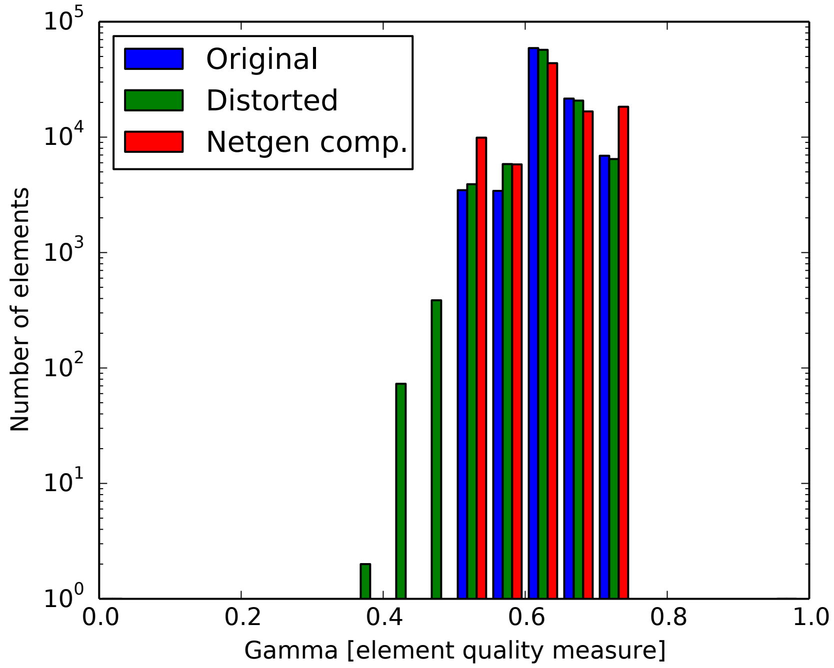

Figure 2.4.: Distribution of γ_k in mesh elements at different stages of mesh processing: the original mesh, after mesh distortion and after compensation

[39], an optimization algorithm for tetrahedral elements. This algorithm works by both adjusting node position within the mesh and swapping links between nodes to give the least distorted element shape. The position adjustment to apply to a node is calculated from the centre of gravity of surrounding nodes, as outlined in [40], and the optimization results in elements which are significantly less distorted than after radial smoothing, as shown in figure 2.4. This figure shows that after compensation using Netgen the mesh quality is similar to the original mesh in the plain pipe.

2.4. Convergence study on element size

Having developed a method for generating the mesh the next stage was to decide on an appropriate element size. It was decided to mesh the entire model with the same size elements to prevent spurious reflections at the boundaries between element densities, since the effects being measured were already relatively small. The element size was chosen through a convergence study which looked at how waves reflected and transmitted from a corrosion patch are a function of the element size in the model. The two parameters used to determine convergence were the maximum signal amplitude and the energy in the signal. Together these two parameters give a good indication of whether or not the model has converged for element size. The maximum amplitude in a signal is a convenient and common measure for convergence, but fails to capture information about the overall shape of signals. The signal energy is used as a more reliable measure of signal amplitude and is ultimately the property of the signal we are interested in.

To perform the convergence study a model was built of a 6 inch, schedule 40 pipe, 2.5 m long with an absorbing region at either end. A rough surface was then generated which was 1 m long and extended for the full inner circumference of the pipe. This rough surface had a mean depth of 2 mm and a correlation length of 7 mm, which is the greatest depth and shortest correlation length considered in the study described in chapter 3. The plain pipe was then meshed with 7 mm elements and the inner surface in the middle of the pipe was perturbed by the rough surface profile. It was decided to begin the convergence study at this extreme large element size as it was obvious no larger element size would be suitable. The scattering of a torsional wave by this rough surface was then measured, both in reflection and transmission, with the first four circumferential orders considered. An example time history for transmitted first order waves is shown in figure 2.5 for different element sizes. This was repeated for models of different characteristic element size, ranging from 7 mm to 1 mm in 1 mm steps. The signals were then assessed against the two criteria of maximum signal amplitude and signal energy. For signal amplitude, the model was said to have converged when the difference in maximum amplitude between successive element sizes was 5% or less, for all circumferential orders in both reflection and transmission. This condition was reached at an element size of 3 mm as shown in figure 2.6.

The energy in each signal was calculated simply as the sum of the squares of the discretised time history. This can be written as:

$$E = \sum_{t_2}^{t_1} u^2 \tag{2.4}$$

where E is the signal energy and u is the discretized time history of the signal which exists between t_1 and t_2. Figure 2.7 shows the energy in signals from different circumferential orders and models with different element sizes. Note that for each circumferential order the maximum amplitude has been set to unity. This was done to keep all results in the

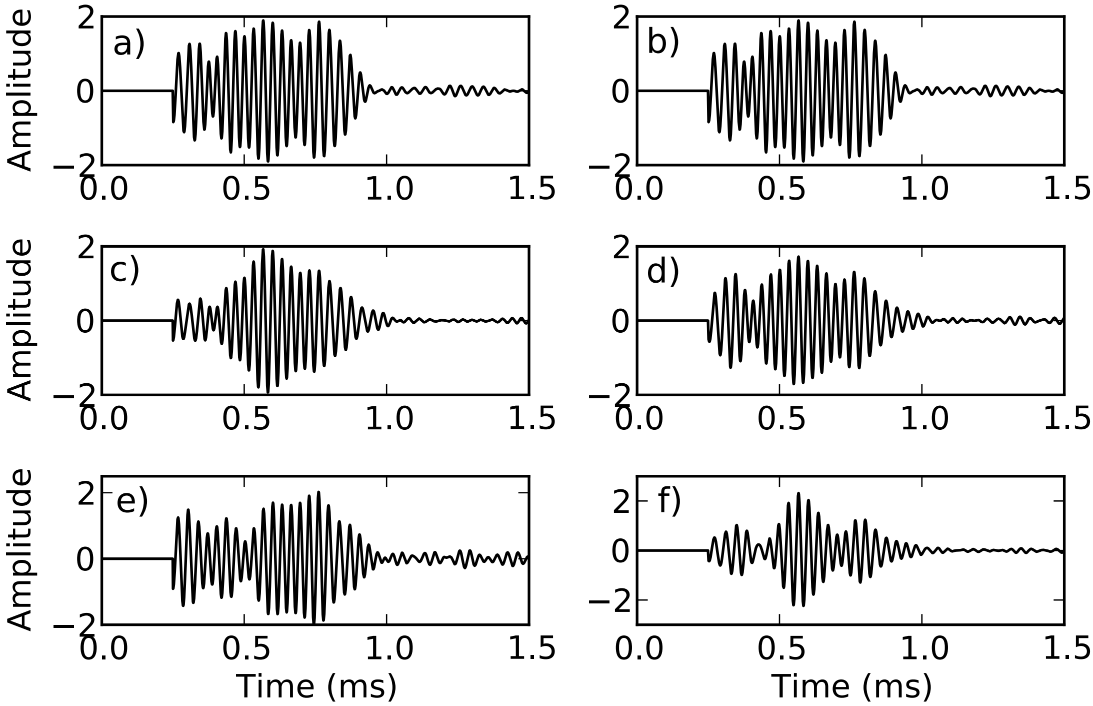

Figure 2.5.: First order reflections for different element sizes, same rough surface profile. Time histories for a) 2 mm, b) 3 mm, c) 4 mm, d) 5 mm, e) 6 mm and f) 7 mm elements. The inspection wave was a six cycle, Hanning window toneburst with a centre frequency of 28 kHz and a wavelength of 114 mm.

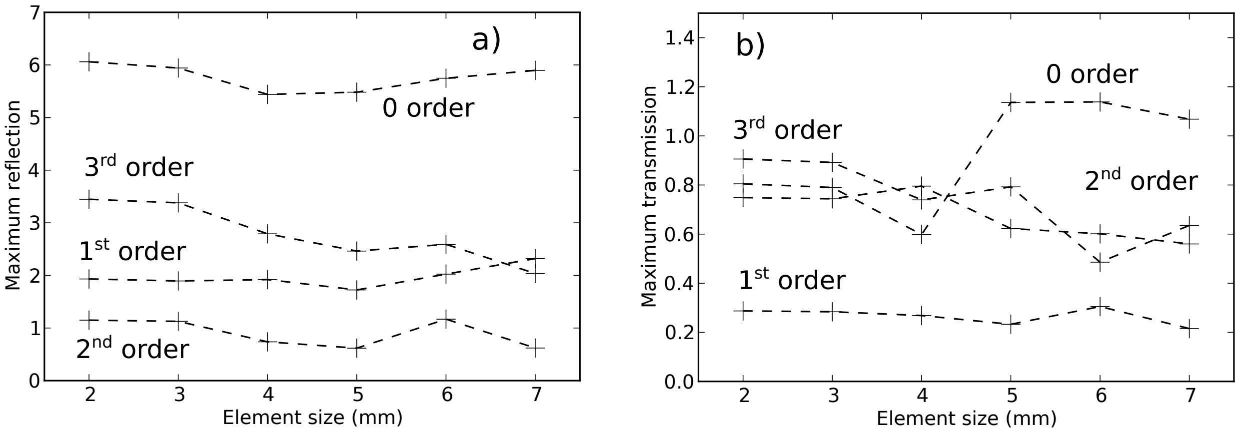

Figure 2.6.: Convergence study for element size, looking at maximum amplitude of a) transmitted and b) reflected components. The inspection signal was a six cycle T(0,1) Hanning window toneburst with 28 kHz centre frequency at unit amplitude.

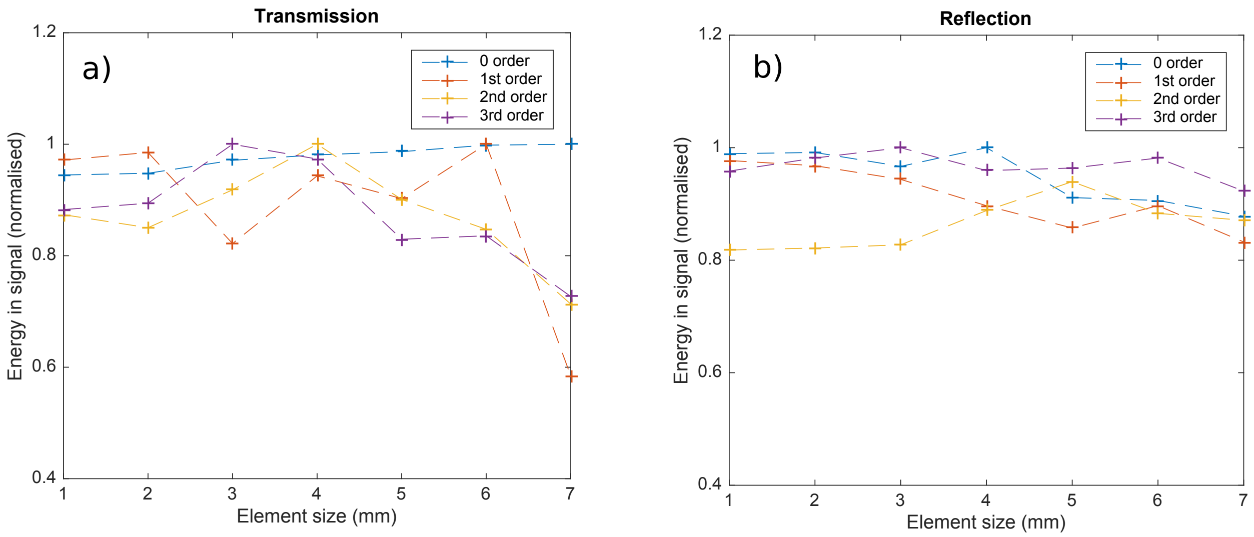

Figure 2.7.: Convergence study for element size, looking at signal energy of a) transmitted and b) reflected components. Results have been normalised so that the maximum energy for each circumferential order is unity. The inspection signal was a six cycle T(0,1) Hanning window toneburst with 28 kHz centre frequency at unit amplitude.

same range and make plotting easier and clearer. As with the maximum amplitude the model was said to have converged when the difference between successive element sizes was less than 5%. According to this criterion the models converge at an element size of 2 mm. Therefore the element size used for all further modelling work was 2 mm.

2.5. Separating guided wave modes and calculating energies

Having defined the model geometry and applied an appropriate discretization, the next stage was to correctly excite and measure guided waves. The most commonly used mode in commercial inspections is T(0,1), the fundamental torsional mode. To excite this mode in our model a ring of nodes was selected, in the case of the rough surfaces study 1.0 m from the rough surface patch, and a uniform tangential force was applied to each node. This force is applied with the time history of a six cycle Hanning window toneburst at a centre frequency of 28 kHz, a typical excitation toneburst used in commercial systems. This force generates a T(0,1) Hanning window toneburst, with no other modes signifi-

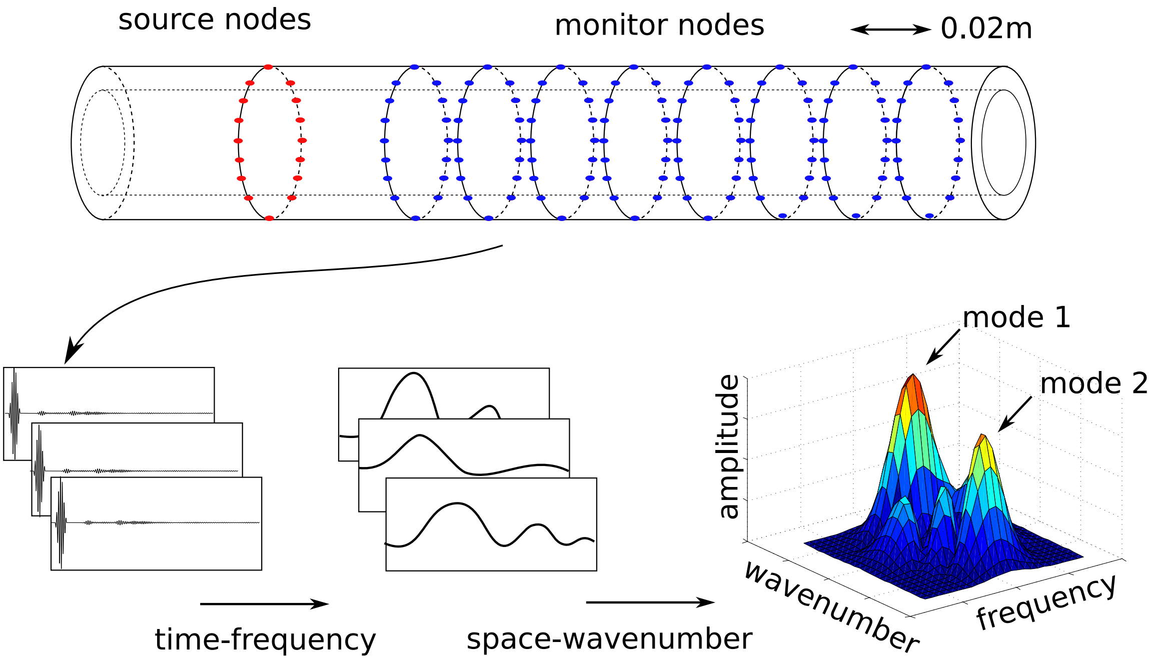

Figure 2.8.: A schematic diagram showing how the three dimensional transform technique can be used to separate guided wave modes by transforming from the time-space domain into the frequency-wavenumber domain. Monitoring nodes distributed in space measure displacement as a function of time. This information distributed in space–time is then converted to a frequency–wavenumber representation using repeat application of the Fourier transform. In this representation different modes are shown as separate peaks and can therefore be isolated.

cantly excited. This selective excitation is possible because no other axisymmetric modes exist with circumferential motion below the T(0,2) cut-off. This model is then solved using an explicit time stepping procedure in POGO. To measure the waves reflected and transmitted from the region of interest in our pipe, in this example the rough surface, a series of monitoring nodes are set up. These monitoring nodes are selected in rings on either side of the region of interest (as shown in figure 2.2), with each ring extending fully around the circumference. There are fifty rings in total, equally spaced along the axis, and at each monitoring node the displacement is measured in each of the three orthogonal directions as a function of time.

The displacements measured at these monitoring nodes will be a complex summation of the contribution from a number of interfering guided wave modes. However, this information is difficult to interpret since it is the total displacement due to a number of overlapping and interfering guided wave modes. A better representation of the data is one in which the contribution from different modes has been separated, since the relative magnitudes of different modes often reveals useful information about the system. In a simple system the modes can be separated using a time of flight analysis, where differences in group velocity between guided wave modes means different modes can be time separated. However, it was not possible to use that method in this study as the different modes are being generated at unknown locations within the corrosion patch and the model is short. These displacements can however be separated using a phased summation around the circumference and a 2D-FFT [41], a technique that makes use of the fact that different modes have different wavenumbers. The monitoring nodes are selected in rings around the full circumference, with fifty of these rings equally spaced along the axis. For each ring, an artificially imposed phase lag followed by a summation around the circumference can be used to separate modes with different circumferential orders, for example separating F(2,3) from F(1,2) but not separating F(1,2) from F(1,3). This is a common technique for separating guided wave modes in simple pipe systems (see for example Ref. [42]). The displacement history associated with each circumferential order is then calculated at fifty axial locations along the length of the monitoring region. For each circumferential order, the fifty time histories are converted from the time-space domain into the frequency–space domain using the Fourier transform. The Fourier transform is then used again to convert to the frequency–wavenumber domain where differences in wavenumber between individual modes means that each mode appears as a separate peak. A schematic of this processing and example peaks in the frequency-wavenumber domain are shown in figure 2.8. Note that an equivalent result is obtained when applying a 3D-FFT to the node data, but the phased summation approach could be used here because of the relatively small number of modes which can exist at the frequencies used in this study. These peaks can then be separated whilst in the frequency–wavenumber domain and the inverse Fourier transform applied twice to recover the time histories for each mode. Note that the width of the peaks in the frequency–wavenumber domain is given by the time step used in the model and the overall axial length covered by the sampling nodes. Since the monitoring region necessarily has a finite extent, the peaks will have a non-zero width and some overlap will occur which introduces small errors in the attribution of displacements to different modes. This is discussed in more detail in chapter 3 when considering figure 3.7.

It was then useful to be able to calculate the energy in the different modes to have a common quantity over which to compare modes. The distribution of energy was calculated by normalising the frequency–displacement curves for each mode by power normalised mode shapes taken from DISPERSE [43]. These normalised spectra are then squared (since energy is related to displacement squared) and divided by the energy spectrum of the input toneburst. This gave the energy in each mode as a fraction of the input energy, both for modes reflected and transmitted from the corrosion patch. Analysis was only performed at those frequencies where the energy input is 10% or greater than the energy input at the centre frequency.

2.6. Energy balance validation

To validate simulations performed with this new modelling procedure an energy balance can be performed. This energy balance is based on the displacement of the waves generated and measured within the model; knowing the amplitude of the T(0,1) wave excited at the end of the model it is possible to calculate the energy put into the system. Since the simulations have been done with no material damping, this energy should be the same as the energy measured in all the waves scattered and transmitted from the patch. A ratio is defined, termed here the energy balance, which gives the ratio of the input and output energies:

$$\text{energy balance} = \frac{\text{energy input}}{\text{transmitted energy} + \text{reflected energy}} \tag{2.5}$$

An energy balance of unity indicates that the input energy and the measured reflected and transmitted energy is the same. A value greater than unity indicates more energy in the transmitted and reflected components, while a value less than unity indicates some energy has been lost in the model. The energy balance is performed in chapter 3 as a validation of the accuracy of the model. An example showing the energy balance under a range of conditions is given in figure 3.13.

2.7. Conclusion

In this chapter the classification of guided waves into torsional, flexural and longitudinal modes has been discussed and a common labelling system for guided wave modes has been introduced. It was recognised that often problems involving these guided wave modes must be solved using numerical techniques and several aspects of finite element modelling techniques for guided waves were introduced. This culminated in the introduction of Graphics Processing Unit based solution procedures, which cause challenges in terms of meshing and distortion compensation due to the number of elements involved. The extent of mesh distortion can be quantified using a simple quality measure applied to each element. This quality measure was then used to show that modelling of rough surfaces causes a distortion of the mesh and that the mesh smoothing algorithms used in this study were appropriate. It was found that for modelling rough surfaces an element size of 2 mm is appropriate, based on a convergence study. It was then shown how to separate different guided wave modes and calculate the energy in different modes. This energy calculation can then be used to perform a simple validation of the model based on an energy balance.

Chapter 3:

The scattering of guided waves by rough surfaces in pipes

3.1. Introduction

When developing signal processing techniques for guided wave monitoring data it is important to test new techniques on example signals. These example signals should represent a wide range of different inspection scenarios. Considerable research effort has gone into understanding the interaction of guided waves with different types of pipe features such as supports, notches, holes and complex defect profiles [13–16]. Example signals for these inspection scenarios are relatively simple to generate through numerical simulation; these pipe features all have a relatively small geometrical extent and so were within the modelling capabilities of previous solution procedures. However, signals resulting from the interaction of guided waves with general corrosion were much more difficult to simulate. Therefore, the new modelling technique developed in the previous chapter was used to model the interaction of guided waves with rough surfaces. This allows rough surface signals to be included in validations of different signal processing techniques.

General corrosion is also important because evidence from the field suggests that this type of defect causes significant scattering of incident waves, leading to increases in the background noise in the reflected signal and increased attenuation of the inspection wave. These two effects are clearly shown in figure 3.1, an example guided wave signal from a laboratory test on a clean section of 3 inch pipe with a 30 year old generally corroded section welded to one end. An increase in the background noise is an issue because it may mask the reflection from more severe localised defects within the generally corroded area, thereby reducing sensitivity. Attenuation is an issue because it reduces the energy reaching features beyond the corrosion and hence reduces the effective inspection range (note for example how in figure 3.1 the reflection from the weld beyond the corroded section is greatly reduced). Another interesting phenomenon from industrial signals is that the attenuation of the T(0,1) inspection wave is greater than would be expected from the energy in the scattered T(0,1) background noise. The weld reflection after the corrosion is approximately a factor of two smaller than would be expected from an extrapolation of the attenuation before the corrosion, with a maximum reflection within the corrosion patch around 10% of the initial weld reflection. Given that energy is related to the square of amplitude, there is clearly energy being lost between the two welds which is not presenting as reflected T(0,1) energy from the corrosion patch. This effect has never been satisfactory explained.

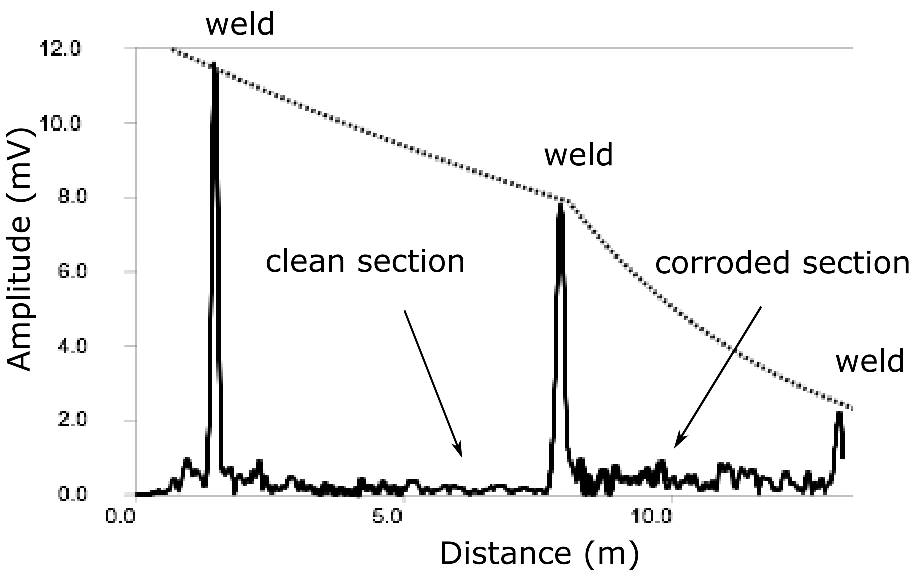

Figure 3.1.: Reflected T(0,1) mode at 28 kHz centre frequency for a 3 inch schedule 40 pipe containing a section of general corrosion. Between the first and second welds there is a clean section of pipe, while between the second and third welds there is a generally corroded section of pipe from industry. Image courtesy of Guided Ultrasonics Ltd.

The ubiquity of this wave–rough surface interaction problem has led to studies in a number of fields, including radar [44], astrophysics [45] and ultrasonics. In ultrasonics, solution procedures are usually split into analytical solutions with several simplifications [46, 47] or full solutions based on numerical methods [48–50]. Studies have looked at a range of different surface types [51] but are largely confined to bulk ultrasonic waves. The problem of guided wave interaction with rough surfaces has received much less attention. Important work in this area was performed by Chimenti and Lobkis [24], who looked at guided wave interaction with rough surfaces in plates. Their analytical approach is based on a phase screen approximation first introduced by Eckart [52] and ignores amplitude effects of the scattering and considers all influence to be felt in the signal phase. Their work achieved generally good agreement between model and experiment but was, due to the approximations used, a partial treatment.

Recent advances in computing power mean that full 3D simulations of guided wave propagation in long lengths of pipe are now feasible [36]. This chapter makes use of these recent advances to carry out a full 3D parametric study of the effect of roughness on guided wave propagation. The approach used in this chapter does not rely on the simplifications of the earlier analytical treatments and allows new physical insights into the mode conversion caused by rough surfaces. The effect of several parameters is investigated, including the correlation length of the surface, the surface roughness, the inspection wave frequency and the pipe diameter. This requires multiple large finite element models to be constructed; a typical model used in this study contains ~6 million degrees of freedom. Section 3.2 introduces the finite element model and the morphology of the rough surfaces, starting with an analysis of a corrosion patch from industry. Section 3.4 shows how the energy in a T(0,1) guided wave scatters when it interacts with a rough surface. This scattering is given as a function of the frequency of the inspection wave, the correlation length of the surface, the mean depth of the surface and the diameter of the pipe. Section 3.5 then describes validation of the simulations and conclusions are drawn in section 3.6. The material in this chapter formed the basis of a journal submission to the Journal of the Acoustic Society of America, currently under review.

3.2. Method

3.2.1. Finite element model

Most of the work in this study was on a 6 inch schedule 40 pipe (a common size in the oil and gas industry), 168 mm outer diameter with a wall thickness of 7 mm and a length of 4.0 m. This model was discretized into linear tetrahedral elements with a

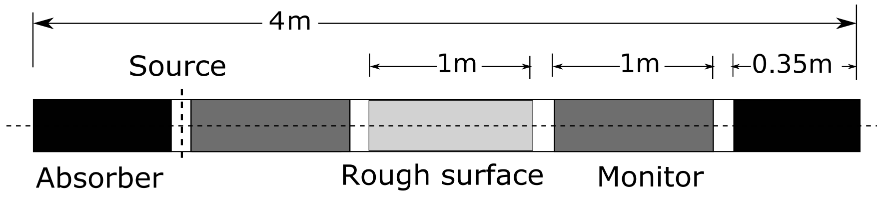

Figure 3.2.: Geometry of the finite element model used in the study. Model separated into five regions: two absorbing regions, two monitoring regions and the rough surface

characteristic length of 2 mm using an implementation of the Netgen algorithm [39] in the Gmsh software. This element characteristic length was chosen from a convergence study which showed that 2 mm elements were sufficient for the frequencies and corrosion geometries used in this study (see chapter 2). Initially, all elements within the mesh were given the properties of a non-attenuative mild steel, Youngs Modulus of 207 GPa, Poissons ratio of 0.27 and density of 7900 kg/m^3. These elements were then grouped into five sections along the length as shown in figure 3.2, with two sections of absorbing region, two monitoring regions and a region of general corrosion. The absorbing regions were 0.3 m long and were set up following Pettit et. al. [34]; the stiffness and damping properties of the model were varied to attenuate any waves passing into these regions, in this way providing a convenient approximation for an infinitely long section of pipe.

3.2.2. Generating synthetic rough surfaces

The corroded region was modelled as a change in geometry of the inner surface so that it followed a given rough surface profile. Although the inner surface was chosen for convenience, the torsional wave used in this study has a similar stress profile through the thickness and hence results will be very similar for corrosion on the outside surface. This rough surface was based on the parameters of corrosion surfaces seen in industry, with an example industrial corrosion patch shown in figure 3.3. This surface comes from a length of steel pipe with general corrosion around the full circumference of the pipe and along a 1 m axial length, caused by intermittent wetting of the surface. A thickness loss map for the corrosion patch is shown in figure 3.3a, where the profile has been measured using a bulk wave ultrasonic thickness probe on a 1 mm grid. To obtain the distribution of this surface the corrosion depths were ordered into the histogram shown in figure 3.3c. Different distributions were fitted to the data and it was found that a normal distribution with a mean of 1.5 mm and a standard deviation of 0.75 mm gave the best fit. The probability density function for such a distribution is shown in figure 3.3c for comparison. The use of a Gaussian distribution to model general corrosion is supported in the literature, where several industrial corrosion patches have been found to follow this type of distribution [53] (although it is noted that experimentally it is possible to misinterpret exponential distributions as Gaussian [54]). The correlation length of the corrosion surface was then calculated, defined as the offset distance at which the autocorrelation drops to $\frac{1}{e}$ of the autocorrelation at zero offset. The autocorrelation for the industrial corrosion surface is shown in figure 3.3b, where the correlation length around the circumference is approximately 10 mm while the correlation length along the pipe axis is approximately 30 mm. From this analysis it was concluded that the rough surface should be modelled as a Gaussian random surface, with a mean depth of the order 1.5 mm and correlation lengths of the order 30 mm.

Simulated rough surfaces were then generated using a process which is described schematically in figure 3.4. The rough surface begins as a grid in two dimensions, with a length and width slightly larger than the axial length and circumference of the desired rough surface and a step size equal to the element size. Each node in the grid is then assigned a number drawn from a standard normal distribution. A separate multivariate normal distribution is then generated which has a probability density function given by:

$$P(x_1, x_2) = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1-\rho^2}} \exp\left[-\frac{z}{2(1-\rho^2)}\right] \tag{3.1}$$

where

$$z \equiv \frac{(x_1-\mu_1)^2}{{\sigma_1}^2} - \frac{2\rho(x_1-\mu_1)(x_2-\mu_2)}{\sigma_1\sigma_2} + \frac{(x_2-\mu_2)^2}{{\sigma_2}^2} \tag{3.2}$$

and

$$\rho \equiv cor(x_1, x_2) = \frac{V_{12}}{\sigma_1\sigma_2} \tag{3.3}$$

In the above equations $P(x_1, x_2)$ is the probability density function, σ_1 and σ_2 are the standard deviations in x_1 and x_2, ρ is the correlation between x_1 and x_2, V_{12} is the covariance and μ_1 and μ_2 are the means in x_1 and x_2. The bivariate normal used in this study had parameters: $\mu_1 = \mu_2 = 0$, $\rho = 0$ and $\sigma_1 = \sigma_2 = CL/2$ where CL is

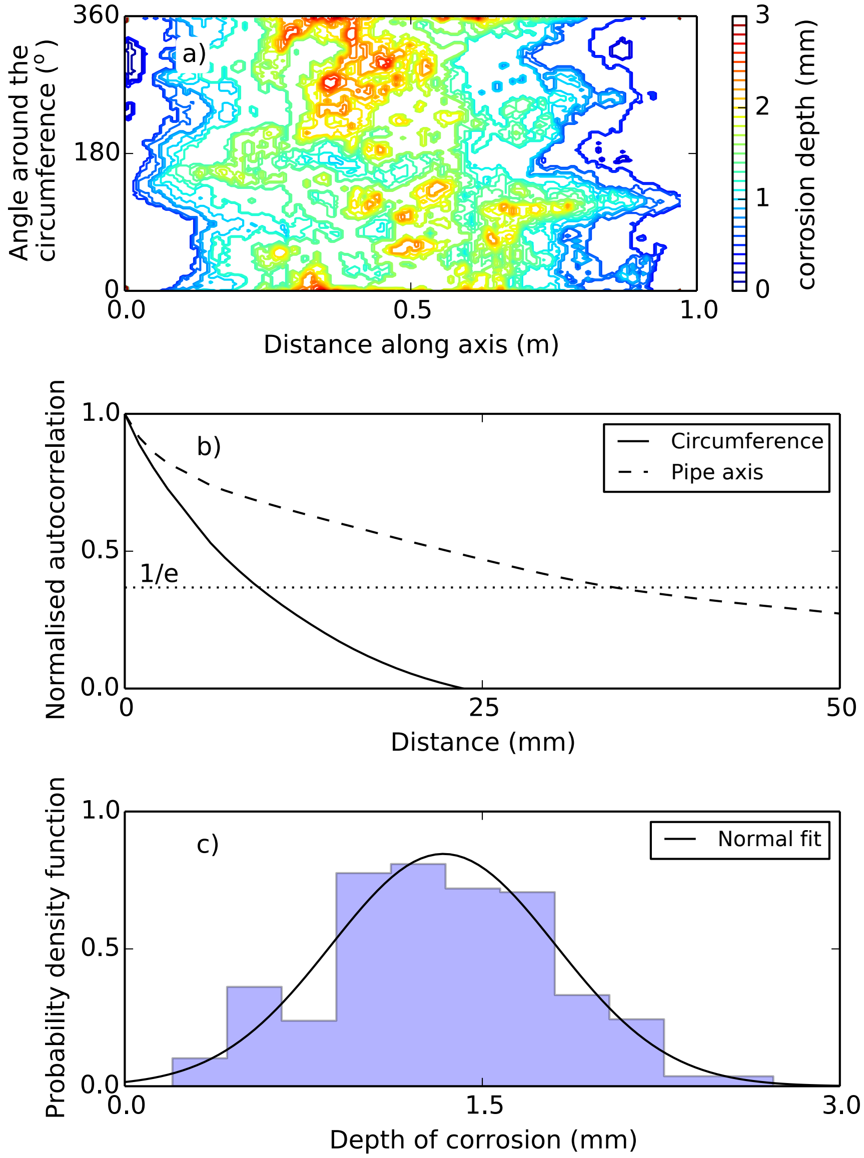

Figure 3.3.: Corrosion surface from industry showing a) the surface morphology, b) the autocorrelation length of the surface in the circumferential and radial directions and c) the distribution of depths. Surface shown has a mean depth of 1.5 mm and correlation lengths of around 10 mm around the circumference and 30 mm along the axis.

the desired correlation length of the generated rough surface. The correlation length is a variable which describes the dominant length scales in the rough surface and in this study the correlation length of the surface is the same around the circumference and along the pipe axis. It is however recognised that many rough surfaces in industry will have difference correlation lengths in either direction. A useful extension of the work in this

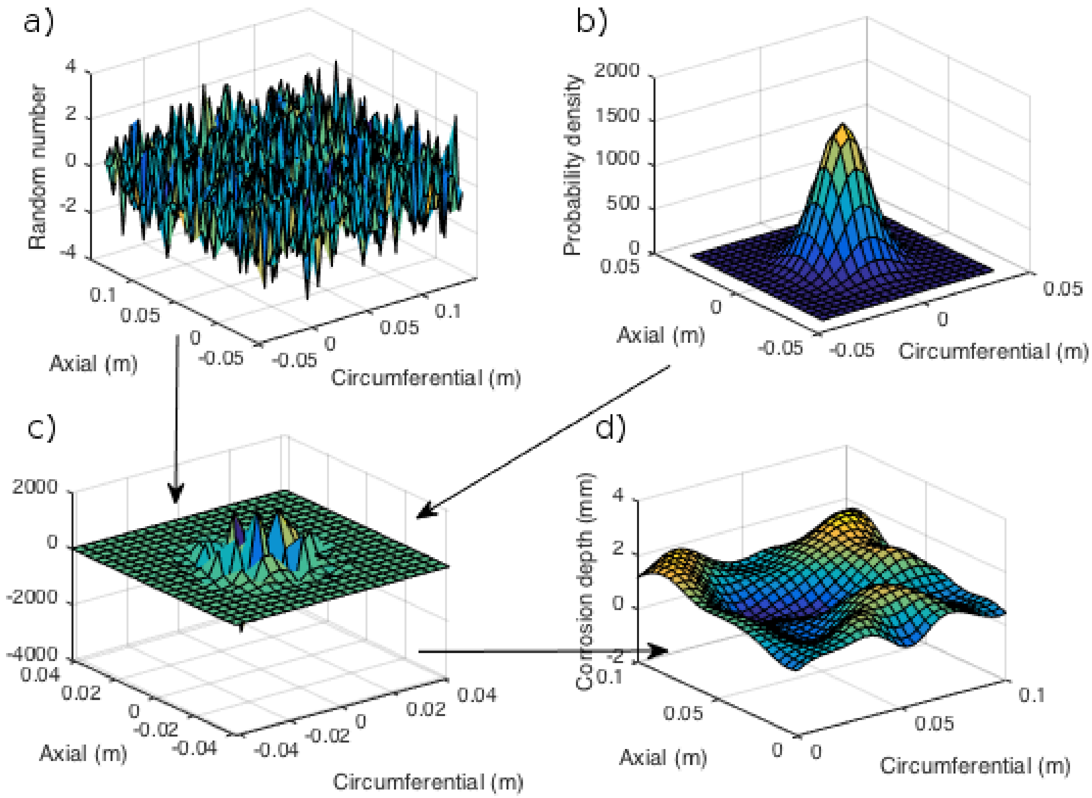

Figure 3.4.: A schematic representation of the processing steps for generating rough surfaces: a) an initial Gaussian random number grid is generated, drawn from a standard normwal distribution, b) a bivariate normal distribution is then generated and c) convolved with the random number grid. The resulting surface d) is then corrected for mean depth and standard deviation. Corrosion surface shown has a mean depth of 1.25 mm and a correlation length of 20 mm. Only a small patch of corrosion is shown to allow greater detail.

chapter would be a sensitivity analysis of how differences in correlation length affect the scattering of guided waves and this is recommended in the future work section of chapter 8.

The random number grid and bivariate normal distribution are then convolved to give a surface with a desired correlation length. This surface is then given an offset equal to the desired mean depth of the corrosion surface and a standard deviation equal to one third of the mean. This generates rough surfaces such as those shown in figure 3.5a, an example surface with a mean depth of 1.25 mm and correlation length of 20 mm. As shown in 3.5c the surface has a Gaussian distribution, although the left tail is distorted due to the fact that all depths less than zero have been set to zero. This was done because corrosion has been modelled here as a thickness loss only. The autocorrelation of the surface is shown in figure 3.5b, with a correlation length of 20 mm around the circumference and along the axis. This surface could then be applied to the pipe model by adjusting mesh node position, with the mesh smoothing algorithm in Gmsh [38] used to minimise mesh distortion.

Models were built with rough surfaces of different depths and correlation lengths, with the study looking at three mean depths: 0.75 mm (11% wall thickness loss), 1.25 mm (18% wall thickness loss) and 2.00 mm (28% wall thickness loss) and fifteen different correlation lengths, from 7 mm in 7 mm steps to 105 mm. These correlation lengths were chosen based on their relationship to the wavelengths of the inspection wave; since the T(0,1) mode is non-dispersive it has a frequency-invariant velocity of 3200 m/s and the centre frequency of the wave used in this model is 28 kHz, giving significant energy in wavelengths ranging from 100 mm to 140 mm. The fifteen correlation lengths therefore correspond to correlation length over wavelength in the range 0.05 to 1.05, depending on the frequency. This then gave a total of 45 model conditions to solve; three corrosion depths, each under fifteen different correlation lengths. However, since the rough surfaces are based on an initial random number grid it is necessary to run multiple models at each condition to determine mean behaviour and so at each condition ten models were generated and then their average behaviour calculated, giving a total of 450 cases to solve. Simulations were also performed for a smaller 3 inch pipe and a larger 12 inch pipe, in both cases with the same wall thickness as for a 6 inch schedule 40 pipe. These models were solved at only one correlation length and depth and their purpose was to investigate diameter affects.

The final stage in setting up the model was to select nodes to excite and measure guided waves. A ring of nodes 0.31 m from one end was selected and a force pattern applied which generated a T(0,1) mode wave with a 6 cycle Hanning window toneburst, centre frequency 28 kHz. Fifty rings of nodes were then selected in each monitoring region for measuring the response of the system to the excitation. These nodes, spaced 0.02 m apart axially, extended around the full circumference and were set up to measure the radial, circumferential and axial displacement of the nodes throughout the propagation of the model. Nodes closest to the excitation point measured the inspection wave and the reflection from the corrosion patch while nodes beyond the corrosion patch measured the transmission. These models were then solved using POGO [36], a finite element solver based on the GPU which is two orders of magnitude faster than similar commercial solvers. This speed improvement enables the current work, making it possible to solve

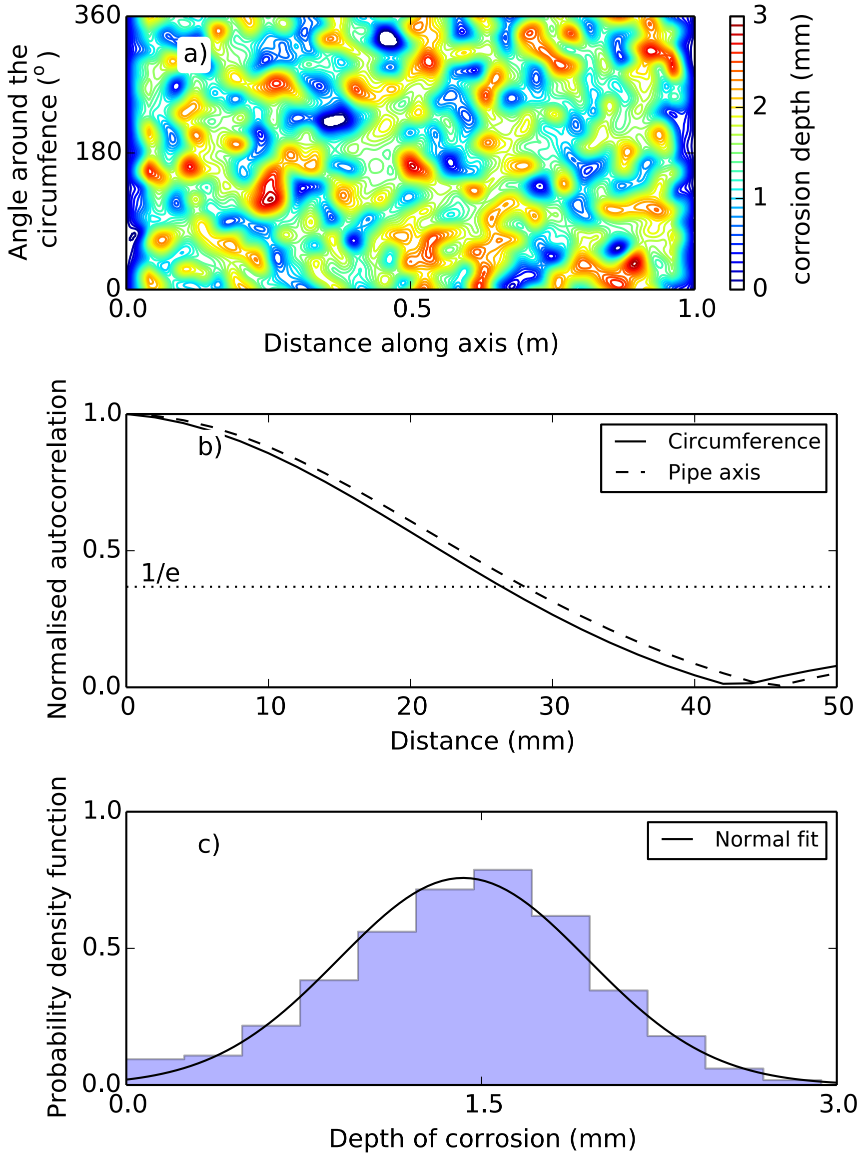

Figure 3.5.: Simulated rough surface showing a) the surface morphology, b) the autocorrelation length of the surface in the circumferential and radial directions and c) the distribution of depths. Surface shown has a mean depth of 1.25 mm and a correlation length of 20 mm.

models with six million degrees of freedom in around 10 minutes. An example time history from a monitoring node 0.2 m from the source is shown in figure 3.6.

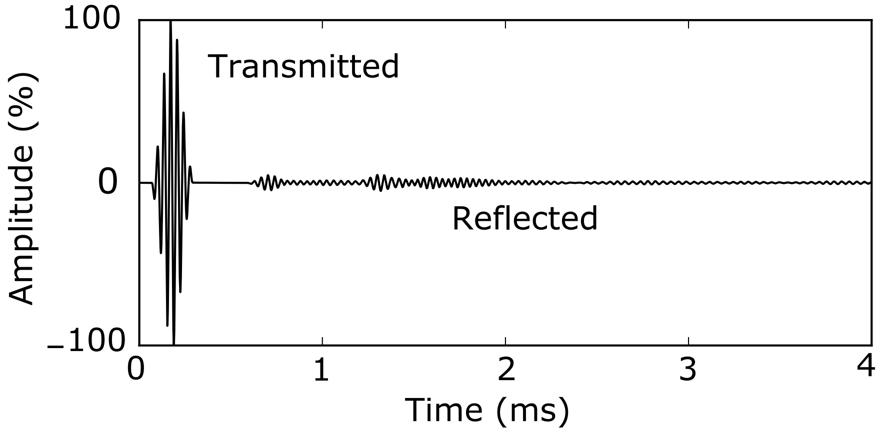

Figure 3.6.: Raw circumferential displacement history for a monitoring node 0.2 m from the source. The initial large wave is the passing inspection signal while smaller subsequent signals are due to backscatter from the rough region.

3.3. Separating modes and calculating energies

The modes transmitted and reflected from the rough surfaces were then separated using the phased summation and 2D-FFT procedure outlined in chapter 2. The separated time histories from an example model are shown in figure 3.7. Note that the small reflected T(0,1) wave at 0.1 ms is spurious and arises because at this point in time there is a very large amplitude (100 %) transmitted T(0,1) wave. Due to the way in which the total displacement is attributed to different modes travelling in different directions, there can be small numerical errors. This is due to wave spreading in the frequency–wavenumber domain and the difficulty in distinguishing between the transmitted wave (with positive wavenumber) and reflected waves (with negative wavenumber). The error in this case is small and this is the only mode for which this spreading is a significant problem. A fuller description of this numerical error is given in chapter 2. It should also be noted that the time trace in part f) is correct; in this example no reflected F(2,2) waves are generated. The energies in the different modes were then calculated as per the procedure in chapter 2.

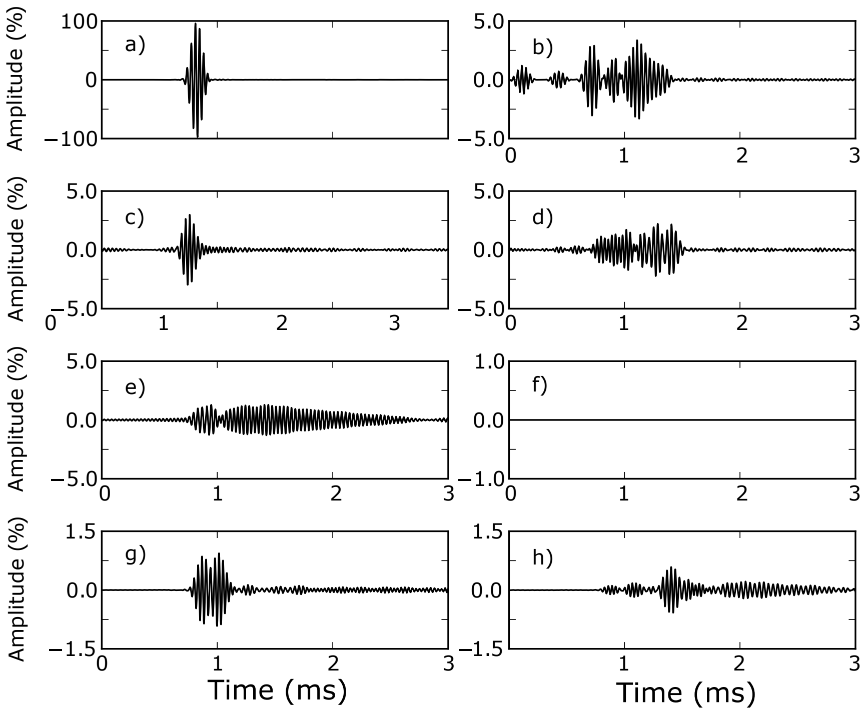

Figure 3.7.: Example transmitted and reflected components from a surface with a correlation length of 28 mm and mean depth of 0.75 mm showing a) transmitted T(0,1), b) reflected T(0,1), c) transmitted F(1,2), d) reflected F(1,2), e) transmitted F(2,2), f) reflected F(2,2), g) transmitted F(3,2) and h) reflected F(3,2). Modes extracted using a phased summation around the circumference followed by a 2D-FFT. The inspection signal was a 6 cycle Hanning window toneburst, centre frequency 28 kHz

3.4. Results and discussion

3.4.1. Effect of inspection wave frequency

We first look at the interaction of guided waves with rough surfaces as a function of the frequency of the inspection wave. Figure 3.8 shows the distribution of energies for a rough surface with a correlation length of 28 mm and mean depth of 1.25 mm, where the average across ten random surfaces is plotted. This figure shows two behaviours which are observed across a range of correlation lengths and mean depths: significant decreases in transmitted T(0,1) energy at frequencies of 24 kHz, 30 kHz and 32 kHz and a gradual decrease in transmitted T(0,1) energy with increasing frequency. At the frequencies where there is a significant dip in transmitted T(0,1) there is a corresponding increase in energy in the F(3,2), F(2,3) and F(4,2) modes. Analysis of the dispersion curves shows that these increases occur at the cut-off frequencies of the higher order modes. A large mode conversion to a mode at its cut-off frequency has been reported in the literature for other types of guided wave interaction and is usually attributed to wall-thickness resonance effects [55]. The study of ref. [55] was at higher frequencies than are considered here so at the cut-off frequencies of higher order modes the resonance condition is through the thickness of the waveguide. In the low frequency propagation along a pipe considered here, the cut-offs correspond to standing waves around the circumference of the pipe; however the physics is analogous. This resonance condition encourages mode conversion and successive cycles pump more energy into this cut-off mode, giving the peaks in the scattered modes close to their cut off frequencies as seen in figure 3.8. The other main feature of this graph is that the energy transmitted in the T(0,1) mode decreases as the frequency increases. This is because more modes can exist at higher frequencies and so it is more likely that energy will find a suitable mode to scatter into at higher frequencies.

3.4.2. Effect of corrosion correlation length

The energy in different modes as a function of correlation length is given in figure 3.9. Note that in this figure the energy has been calculated at a frequency of 34 kHz with this frequency chosen to be well away from the cut-off frequencies of any modes. The behaviour is shown at one particular frequency rather than on a normalised frequency scale because different modes can exist at different frequencies and so the energy scattering will be different in each case. Results at other frequencies are very similar, with the difference that the energy is scattered into different modes. Also note that the correlation length has been converted to a non-dimensional parameter through division by the wavelength of the T(0,1) mode at the respective frequency. The results have been plotted here as the average over ten random surfaces for each correlation length, with the mean indicated as a solid symbol and the error bars showing one standard deviation.

The main feature of these graphs is that the minimum transmission of T(0,1) occurs when the correlation length over wavelength is in the region 0.2-0.25. At this correlation length to wavelength ratio the greatest energy is being scattered to other modes. This is consistent with scattering from many other pipe features which have their maximum interaction with guided waves at length scales of around one quarter wavelength, for example axi-symmetric notches [14]. In the case of notches this maximum scattering is due to constructive interference effects and it is thought a similar mechanism is responsible

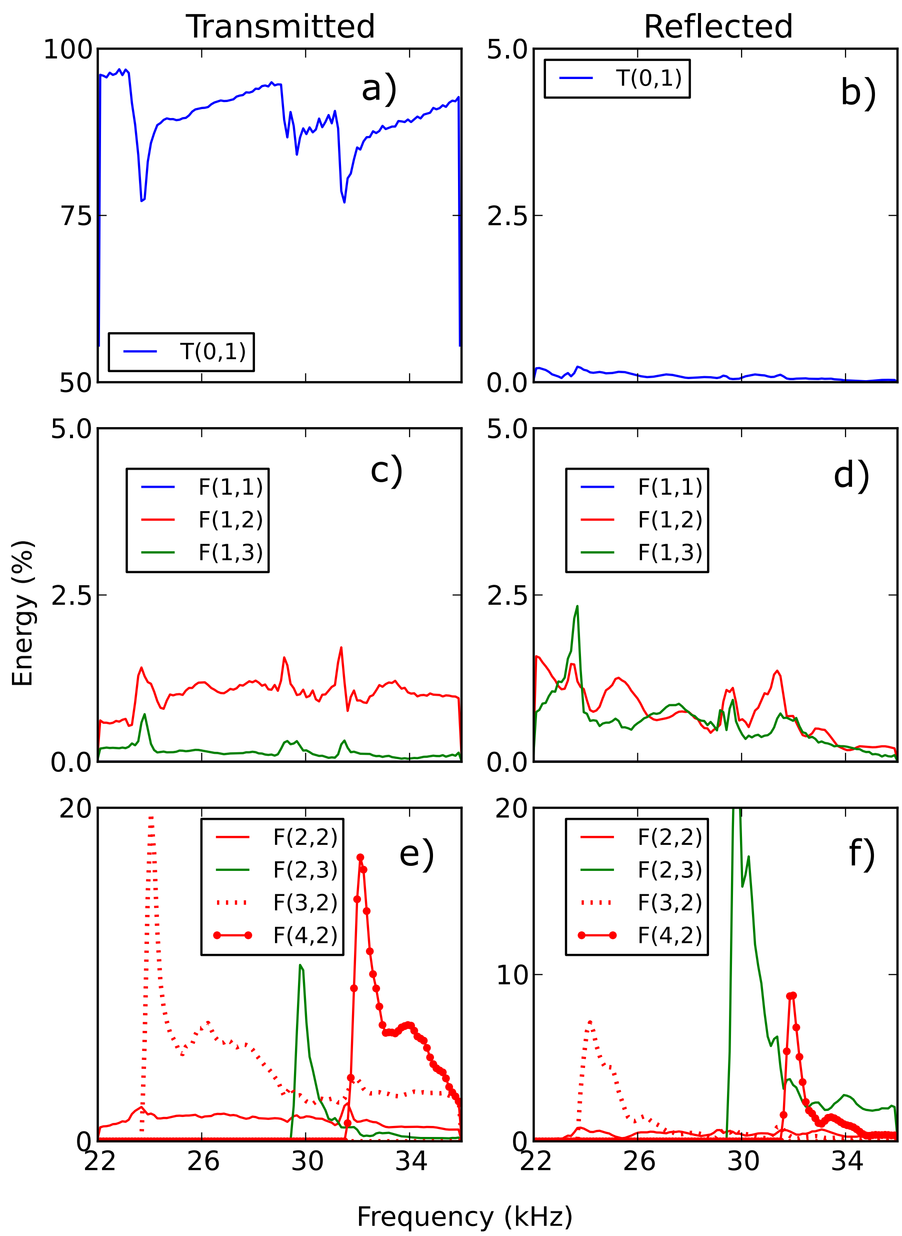

Figure 3.8.: Example energy distribution in transmission and reflection for a rough surface as a function of frequency. Showing a) transmitted T(0,1), b) reflected T(0,1), c) transmitted first order modes, d) reflected first order modes, e) transmitted higher order modes and f) reflected higher order modes. Values shown are for a 1 m length of surface with root mean square depth of 1.25 mm and a correlation length of 28 mm

for the behaviour in rough surfaces. Away from frequencies close to cut-off the scattered energy is transferred mainly to the higher order (F(*,2)) modes in transmission. This preference for higher order flexural modes is due to the similarity in displacement profile between the fundamental torsional mode and these flexural modes, as shown in figure 3.10. Since these modes have the greatest circumferential displacement for a given power flow they will couple most easily to the torsional mode. It is interesting that the F(*,3) modes have relatively high circumferential motion close to cut-off which together with the resonance phenomenon discussed above gives peaks in the conversion to these modes. The proportions in which the energy scatters to these higher order flexural modes is also a function of frequency since these modes cannot exist below a threshold frequency, a threshold which is different for each mode. It is also interesting to note that most of the scattered energy is transmitted through the corrosion patch rather than being reflected towards the source.

3.4.3. Effect of corrosion depth

The energy in the transmitted T(0,1) mode as a function of correlation length and mean corrosion depth is shown in figure 3.11. Average results have been plotted here, with the typical standard deviation 10% of the mean. This shows that the amount of scattering from the corrosion surface increases as the depth of corrosion increases. This is in agreement with previous studies, which in general find that the interaction of guided waves with pipe features is roughly proportional to the cross sectional loss represented by that feature.

3.4.4. Effect of pipe diameter

A comparison of the transmitted T(0,1) energy for 12 inch, 6 inch and 3 inch diameter pipes is shown in figure 3.12, where all pipes have been modelled with a wall thickness of 7 mm and all rough surfaces have a correlation length of 50 mm and a mean depth of 18 % of the wall thickness. The results have been plotted as a function of frequency-diameter to highlight the influence of pipe diameter on the transmitted T(0,1) energy. It is clear that the attenuation of T(0,1) waves is a strong function of pipe size. Significant, localised reductions in transmitted T(0,1) which are due to cut–off frequencies of higher order modes occur at different frequencies since these cut–off frequencies will be a function of pipe diameter. Away from these frequencies the transmission reduces with pipe diameter. This diameter dependence is due to the existence of more higher order flexural modes at a given frequency for the energy to scatter into. In commercial pipework, the wall

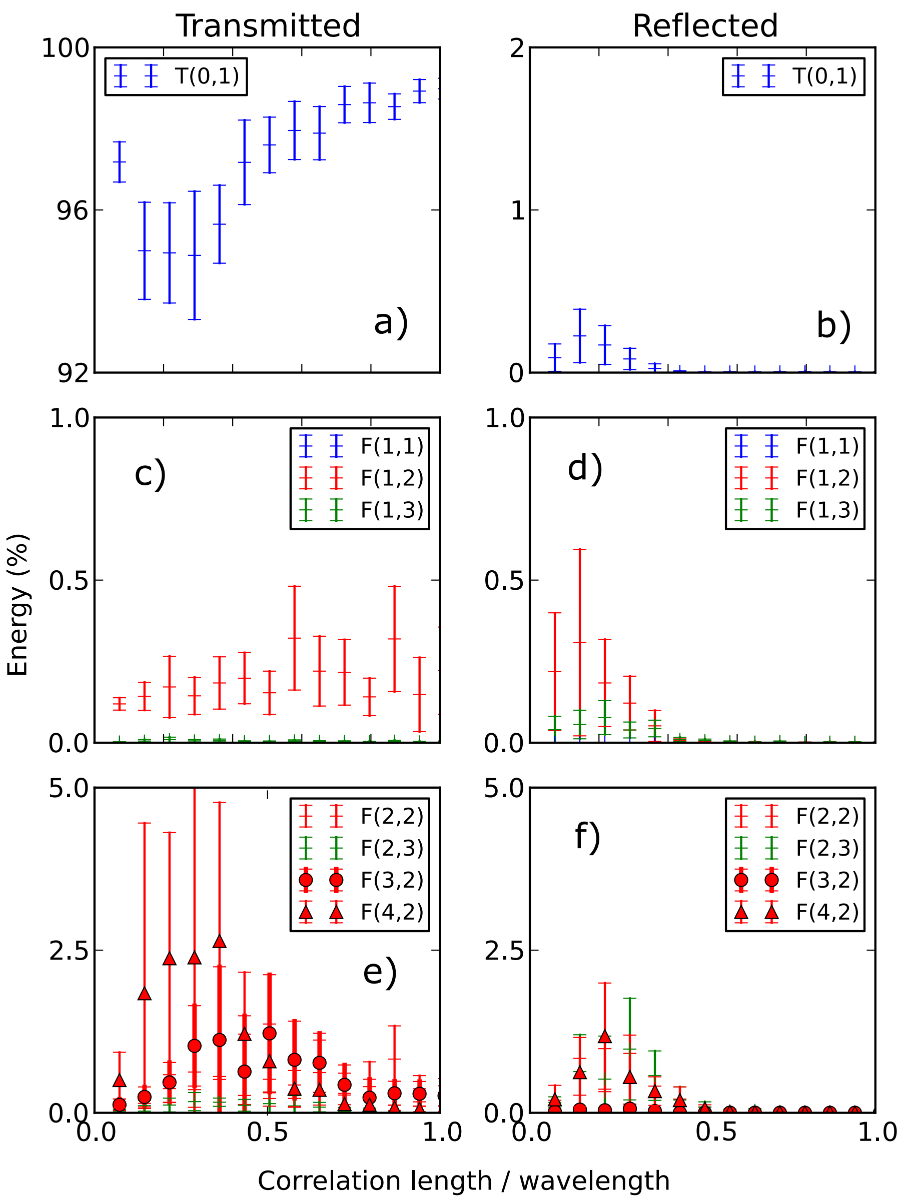

Figure 3.9.: Example energy distribution in transmission and reflection for a corrosion surface as a function of correlation length. Showing a) transmitted T(0,1), b) reflected T(0,1), c) transmitted first order modes, d) reflected first order modes, e) transmitted higher order modes and f) reflected higher order modes. Values calculated at 34 kHz. Values are plotted as the average across ten surfaces showing the mean as a solid symbol and errorbars as one standard deviation.

thickness for a given pipe schedule increases with pipe diameter but not in proportion to it. The analysis was therefore repeated for schedule 40 pipes of the three diameters and the results were very similar to those of figure 11 (maximum difference 5%), showing that

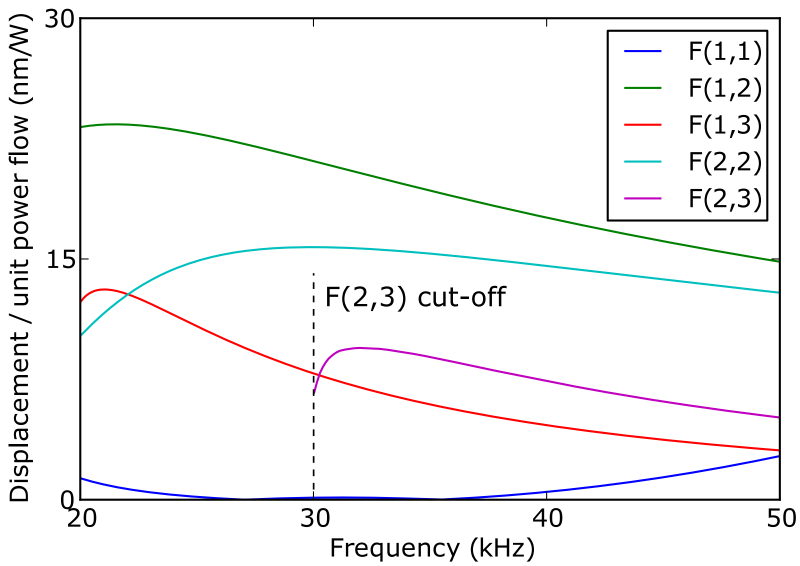

Figure 3.10.: Torsional displacement per unit power flow for different modes. The F(*,2) modes have the greatest displacement and hence couple most easily to T(0,1). The F(1,3) mode has a cut-off at 19 kHz, just outside the plotted range

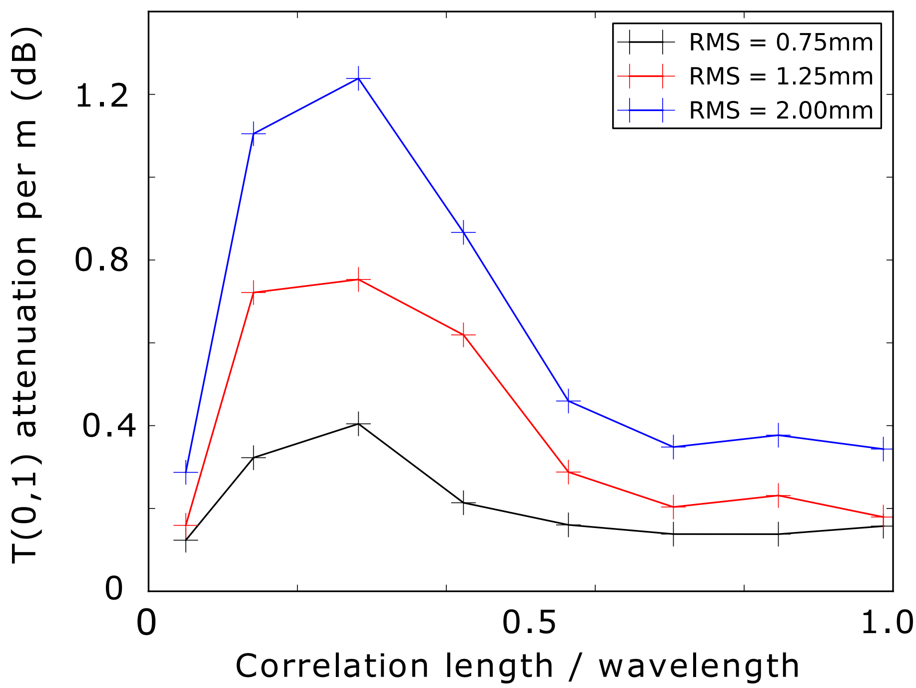

Figure 3.11.: Attenuation of T(0,1) as a function of the correlation length of the surface and the mean depth.

the effect of frequency-thickness product is much smaller than that of frequency-diameter.

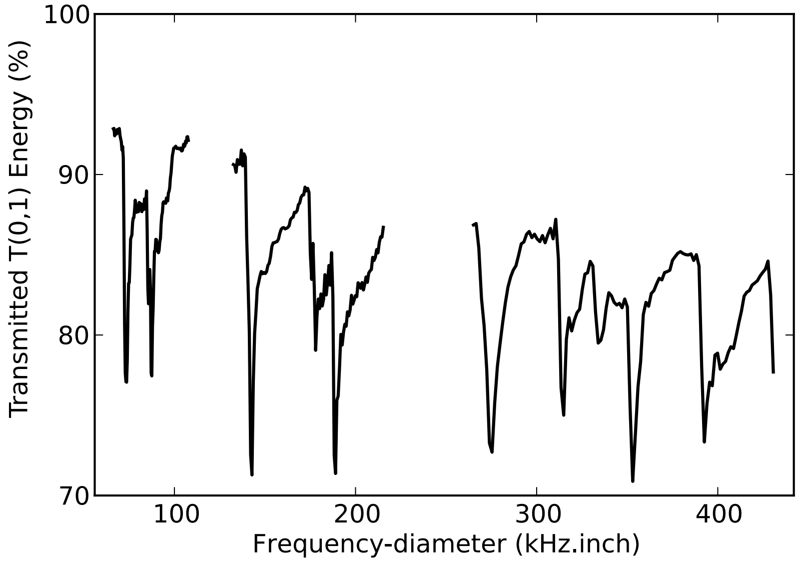

Figure 3.12.: Variation in transmitted T(0,1) energy with pipe diameter, showing results for a 3 inch, 6 inch and 12 inch pipes all at a wall thickness of 7 mm. All corrosion surfaces have a mean depth of 18 % of wall thickness and a correlation length of 50 mm.

3.5. Validation

3.5.1. Validation through energy balance

As described in chapter 2 an energy balance can be used as a simple validation of the simulation results. Figure 3.13 shows the energy balance sampled across a range of correlation lengths and a range of frequencies. Overall the energy balance is very good, with a maximum discrepancy of 3 % from the ideal. The energy balance is worst at lower correlation lengths, the conditions under which the greatest scattering occurs.

3.5.2. Validation through experiment

A 3.5 m length of 3 inch diameter, schedule 40 pipe was machined on the outer surface with a rough profile which had an overall length of 0.75 m, a correlation length of 50 mm and a

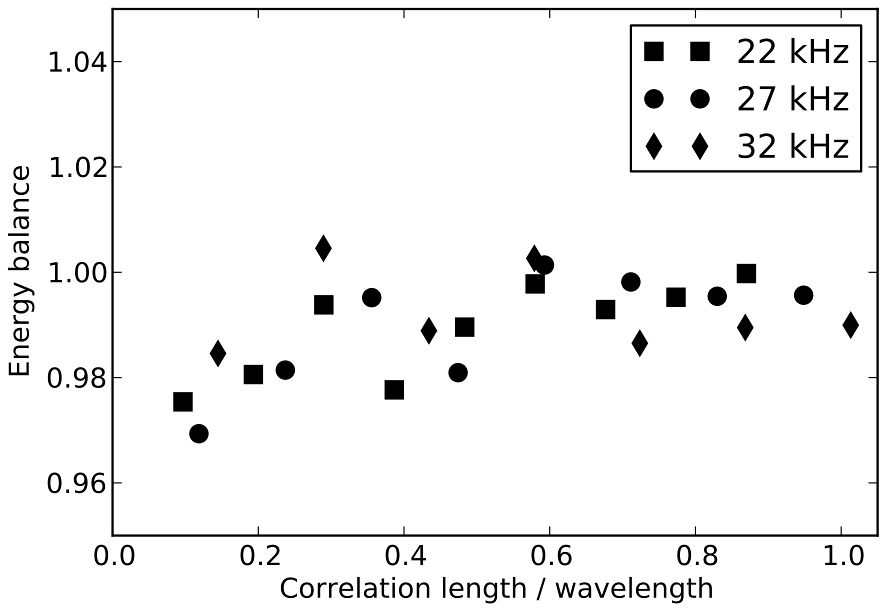

Figure 3.13.: Energy balance of the models across a range of correlation lengths and frequencies. A value of unity indicates perfect energy recovery.

mean depth of 2.0 mm (36 % wall thickness), generated using the same procedure as for the simulations. This patch started 1.1 m from the end of the pipe end and extended for over half of the circumference, with half circumference machining chosen in preference to full circumference so that the pipe did not have to be moved during the machining operation which was done on a numerically controlled flat bed milling machine. This patch also had a linear ramp at its edges to provide a smooth transition to the surrounding plain pipe. A thickness map for the patch is shown in figure 3.14 which shows the distribution of peaks and troughs. A commercial transducer ring from Guided Ultrasonics Ltd was then clamped to the pipe at one end. The ring comprises 32 dry coupled transducers that apply a circumferential traction to the pipe wall, which when excited simultaneously generate a pure T(0,1) mode and when linked in reception receive this mode preferentially. Measurements were made with this sensor at five different centre frequencies: 23 kHz, 28 kHz, 34 kHz, 41 kHz and 50 kHz. This inspection set-up, shown in figure 3.15, was then modelled using finite element simulations based on the same procedure as used for the rest of the results presented in this chapter. For accurate comparison with experiments the finite element results were corrected for the slight attenuation which occurs as the waves pass under the ring itself. This was done by applying an attenuation to the time histories from the model; the correct attenuation was calculated by performing a guided wave test on an undamaged length of 3 inch schedule 40 pipe and measuring the attenuation of the inspection wave with successive passes under the sensor.

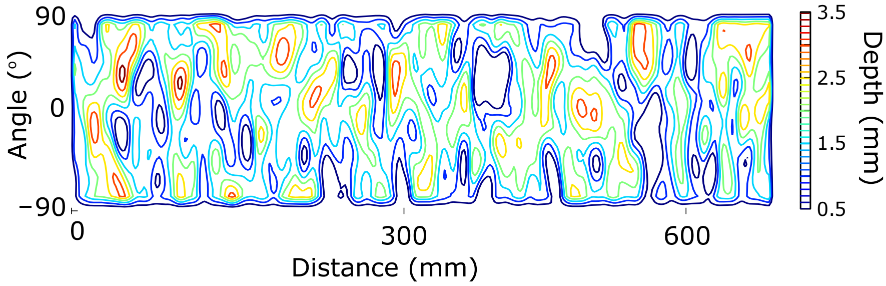

Figure 3.14.: Corrosion patch machined into a 3 inch pipe. Patch has a correlation length of 50 mm and a mean depth of 2 mm.

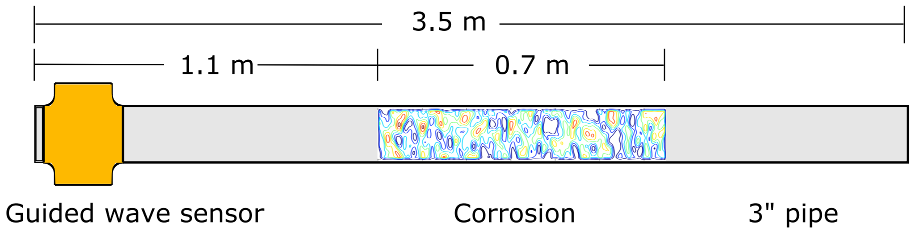

Figure 3.15.: Validation experiment equipment set-up

Example time histories at a centre frequency of 28 kHz are shown in figure 3.16 for both the experiment and the simulations. These time histories show a series of large reflections which are due to successive end reverberations from the pipe. The amplitude of these reverberations is reducing due to mode conversion within the corrosion patch, while the signals between end reverberations are due to scattering of T(0,1) waves from the patch. Note that in figure 3.16 the measured time history has been set to zero for the first 5 m, which explains the discrepancy between predicted and measured time histories in this region. This zeroing was done because the guided wave sensors do not give accurate measurements immediately after transmission due to ring-down of the transducers.

By comparing the amplitudes of successive reverberations for both the experiment and simulations it is possible to compare the mode conversion in both. The amplitudes of the end reflections for both the simulations and experiments are shown in figure 3.17 at a range of centre frequencies. This figure shows that there is excellent agreement between experiments and simulations at all frequencies. It is worth noting that the attenuations seen in the experiment and corresponding simulations are larger than the attenuations predicted in the earlier simulations. The reason for this is that multiple patches with similar statistics were modelled and one with a large attenuation was chosen for the experiment in order to make the measurements easier. The patch also has a mean depth

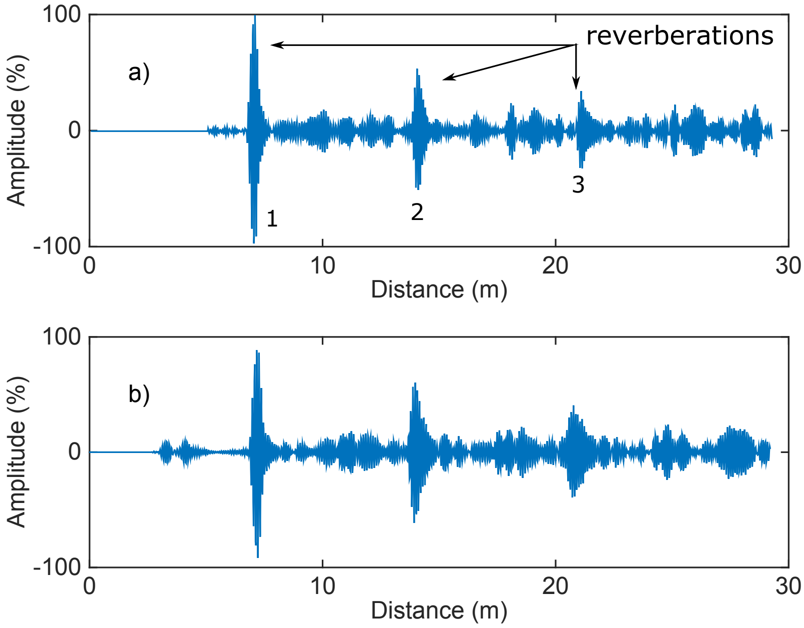

Figure 3.16.: Comparison of the a) measured and b) predicted time history of the T(0,1) mode in the setup of figure 3.15. Large peaks in the signal come from reflections from the pipe end. The first 5 m of the measured signal has been set to zero since the sensors are unreliable immediately after excitation due to ring-down.

greater than was used in the earlier simulations and since the sensor is configured in pulse–echo mode, the transit distance of the wave through the corrosion patch is 1.5 m compared with the 1.0 m used in the 6 inch pipe simulations. It is also possible that machining the random surface over part of the pipe circumference rather than the whole circumference increases asymmetry and so promotes mode conversion. Further investigation of this possibility is beyond the scope of this chapter.

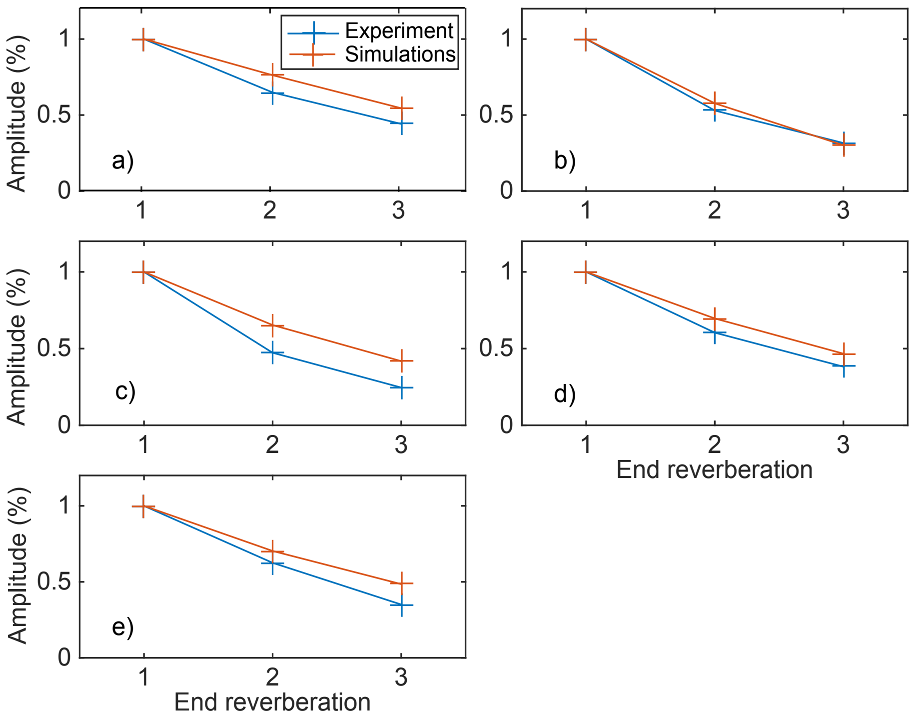

Figure 3.17.: Comparison of reverberation amplitudes (shown in figure 3.16) of the T(0,1) wave between simulations and experiments at centre frequencies of a) 23 kHz, b) 28 kHz, c) 34 kHz, d) 41 kHz and e) 50 kHz

3.6. Conclusions

This chapter has used finite element simulations to investigate the mode conversion of torsional guided waves passing through rough surfaces in pipes. These rough surfaces, modelled as a Gaussian, random wall thickness loss were designed to investigate the reason for the greatly increased attenuation seen in the inspection of generally corroded pipe. It was found that these rough surfaces cause mode conversion from the torsional wave to non-axisymmetric higher order flexural modes, which explains the significant attenuation seen during commercial inspections.

This mode conversion was found to be a function of several variables: the frequency of the inspection wave, the characteristics of the rough surface and the diameter of the pipe. In general it was found that the mode conversion increases as the inspection frequency increases, since at higher frequencies more modes can exist and hence there are more modes for the energy to scatter into. However, this trend breaks down close to the cut-off frequencies of the higher order flexural modes where the mode conversion is particularly strong; at their cut-off frequency flexural modes form standing waves around the pipe circumference and this resonance condition encourages mode conversion. The mean depth and correlation length of the surface also has an influence on mode conversion, with attenuation greatest for surfaces where the correlation length over wavelength is in the region 0.2-0.25; the attenuation approaches zero at low and high values of correlation length relative to the wavelength. The deeper the corrosion the more severe the mode conversion, since deeper corrosion will lead to a reduced remaining wall thickness and hence a greater impedance mismatch with the plain pipe. It was also found that mode conversion is a strong function of pipe diameter; for pipes with the same wall thickness and rough surface statistics the attenuation increases significantly with frequency-diameter product. This is because larger diameter pipes can support more modes at a given frequency, so there are more possible modes for the energy to scatter into. Away from cut-off frequencies the mode conversion is mostly to the F(*,2) modes since these modes have the greatest circumferential displacement for unit power flow; the T(0,1) mode therefore couples most easily into these modes.

The finite element simulations on which these results are based have been validated using an energy balance and experimental validation. The energy balance was performed across the inputs and outputs from the finite element model and gave a worst–case discrepancy of 3 %. Experimental validation of the finite element modelling was performed on a 3 inch diameter pipe which had a corrosion surface machined in the middle of the pipe over half the circumference. Comparisons of attenuation in the time histories from the experiment and a simulation of this inspection set-up show very good agreement across a range of frequencies.

This mode conversion can cause severe attenuation of the inspection signal; for example a mean wall thickness loss of 28 % can cause 1.35 dB per m attenuation of the torsional wave, corresponding to 2.7 dB loss per metre in a pulse-echo configuration. This is for a frequency range typical of commercial deployment. This does not include the effect of any adhered corrosion product which could further increase attenuation, an effect which has not been modelled in the simulations. To minimise this attenuation the inspection should be performed at the lowest possible frequency while avoiding cut-off frequencies and the region where the correlation length is around a quarter wavelength.

The ability to simulate rough surfaces also means that we now have a method for evaluating how different processing techniques for monitoring data deal with rough surface signals.

Chapter 4:

Compensating monitoring data for environmental effects

4.1. Introduction

As noted in chapter 1 it is becoming more common for guided wave sensors to be operated in a monitoring configuration rather than an inspection configuration. This trend is partly motivated by potential improvements in defect detection and classification due to the higher volumes of repeatable data. However, extracting useful information from this repeat data is difficult. Variation in the data will be due to both structural changes and environmental changes and these two factors must be separated. Various signal processing techniques have been developed to perform this separation, with no clear superior method having been identified. In section 4.2 some of the issues relating to changing environmental and operational conditions are introduced. Section 4.3 then goes on to describe one class of techniques devised to overcome these issues, whilst section 4.4 then goes on to discuss another class. Conclusions are given in section 4.5.

4.2. Environmental effects

A monitoring procedure is based on the assumption that the growth of damage over time will cause changes in the output from monitoring sensors. This is typically true, but is complicated by the fact that benign changes in the environment will also cause changes in the sensor output. Therefore, a way must be found of identifying those changes which occur due to damage and those changes occurring due to changing environmental conditions. This is known in Structural Health Monitoring (SHM) as the problem of data normalisation [18]. Looking specifically at the monitoring of pipelines using guided waves, it is possible to identify three contributing factors to environmental noise following Sohn [56]: temperature, changes in boundary conditions and changes in loading.

Of these three, temperature effects make the greatest contribution to environmental noise [57]. Temperature may adversely affect the performance of the sensor and will certainly change the material properties of the pipe which govern guided wave propagation. These changes in material properties are unavoidable and at present the issue of sensor stability can be significant, although it is suggested that in time sensor effects could be negated through better design [57]. Changes in material properties will cause changes in geometry as well as changes in the propagation of ultrasonic waves [57]. Of these, changes in the velocity of the ultrasonic waves will be the most significant effect with changes occurring in both the phase velocity and group velocity. Most compensation schemes try to address this group velocity effect, a non-trivial task given that the correction to apply will be a function of the frequency content of inspection and the degree of interference between feature reflections [58]. The importance of temperature in guided wave monitoring is reinforced by the large number of papers on the subject, see [59–61].

A perhaps less well studied environmental effect is the effect of changing boundary conditions on the structure. Some of these changes are direct, such as a change in the contents of the pipe. Such a change, either in the fluid being transported or the fluid level will cause changes in the dispersion and attenuation of a guided wave signal [62]. Indirect effects arise as a secondary effect of another change. Temperature changes for example can cause thermal expansion which tightens clamped supports or brings other supports into contact. Supports can also be brought into contact through changes in loading. Research has shown that these supports can then act as high pass filters to the guided wave signal [16], affecting the ability of inspectors to correctly assess reflections from beyond a support. The final effect, mass loading, has not been studied in great detail for guided waves.

4.3. Compensation procedures based on signal manipulation

There have been many proposals in the literature of ways by which to compensate for these environmental effects. A useful framework for considering these different proposals is to follow the structure of Sohn [56] and separate these solutions into: regression analysis, novelty detection, neural networks (particularly auto-associative) and component analysis methods. The first two techniques on this list compensate for environmental effects by collecting data about environmental variations when the system was first installed. They then use an understanding of these early variations to compensate future readings. The second pair of techniques projects the original sensor data onto a new set of axes. If this projection is done correctly some of these new representations of the data will be sensitive to damage and insensitive to environmental variations. This allows the two sources of change, environmental effects and structural change, to be separated.

Regression analysis is the most straightforward technique for mitigating environmental effects. The technique starts with the collection of data from the structure in its undamaged state. Several measurements are then made under different environmental conditions which are representative of those conditions expected in service. This data will be termed baselines henceforth for convenience. Some kind of regression analysis is then used to relate different sensor outputs to certain environmental conditions. When the system comes on-line, an inspection is made with the environmental conditions measured at the same time. Using the baseline data some expected sensor output is determined and is then compared to the actual sensor output and if there is significant discrepancy it is deemed that some structural change has occurred. This technique was used for example in the monitoring of the Z24 bridge in Switzerland [63].

A variant of this technique is the optimal baseline subtraction (OBS) technique. This technique has been used in a number of studies looking at removal of environmental variations in guided wave signals [60, 64, 65]. A series of baseline guided wave signals are collected under different environmental conditions. When an on-line measurement is made, the baseline signals are examined to find the baseline that gives the smallest residual after subtraction of the on-line measurement. This then effectively finds the baseline which matches the on-line measurement best in terms of conditions of collection. Note that explicit measurement of environmental conditions is not required at any stage as it is inferred from the minimisation of the residual. If a significant discrepancy is then found between the on-line measurement and the selected baseline then it is determined that a structural change has occurred. Note that some authors use other measures of discrepancy, such as the scale-invariant correlation coefficient [66].

Optimal baseline subtraction has been applied successfully in a number of studies but requires a large number of baselines to be effective. Fewer baselines are required if a technique known as baseline signal stretch (BSS) is used. In this technique baseline signals are manipulated to reverse some of the effects which result from temperature changes. This manipulation typically involves some combination of stretching, amplitude scaling and time delay applied to the signal to bring it closer to a reference on-line measurement. These manipulations can be performed in isolation, for example the approach of [67] in which only amplitude scaling is applied to compensate for water loading based attenuation. Likewise stretch can be applied in isolation, such as used in [65]. Manipulating the signal in only one aspect has the advantage that simple metrics can be used to determine the optimum stretch/delay/scale factor to apply; a common approach is to minimise the residual between the stretched signal and the on-line measurement. Improved compensation can be achieved if all three manipulations are performed in tandem, but this approach requires some form of simultaneous optimisation. The temperature change being compensated is not generally known beforehand and so it is not possible to calculate the required delay/ scale/ stretch analytically. There are various forms of optimisation procedure which can be employed such as: genetic algorithms, simulated annealing or patterns search [68–70]. The choice of optimisation procedure is guided by the need to achieve a globally optimal solution and authors such as Galvagni have proposed guidelines as to the conditions under which each technique is most appropriate [19]. Simulated annealing is used in chapter 5 in this thesis because it is capable of dealing with local minima (which is a problem in optimisation procedures for baseline stretch[65]). It is also easy to implement in software packages such as Matlab. In this approach two signals are selected; a baseline and a reading, where the reading is the signal which will be manipulated to match the baseline signal. Initially both signals are transformed from the time domain to the frequency domain using a Fourier transform. This is because the stretch procedure amounts to a re-sampling of the time trace at a different sampling frequency, which is most conveniently achieved in the frequency domain [65]. A delay, scale and stretch are then iteratively applied to the reading and the residual between the frequency spectra of the baseline and reading is calculated. The delay, scale and stretch are iterated together until the residual is minimised (the minimum is found using the optimisation procedure), at which point both signals are reverted to the time domain.

A slightly different stretch procedure is used in chapter 7 to be consistent with the method of a collaborator. In this procedure the stretch, scale and delay are applied independently, which makes the compensation faster to perform compared with simultaneous optimisation procedures. However, independently optimal stretch, scale and delay factors do not necessarily lead to a globally optimal combination of these manipulations. However, no direct comparison of the two techniques has been published and hence it is not possible to quantify the discrepancy in performance, if any. The stretch factor in this procedure is calculated very efficiently using a variant of the Mellin transform as described in [71]. This uses the correlation between signals to calculate the optimal stretch factor, which is more efficient than calculating the residual after taking the difference between two signals.

For all the techniques used in this thesis the signal manipulation is applied globally, which means that the same stretch, amplitude scale and delay is applied to the whole signal. This is appropriate for the type of signals considered in this thesis where the temperature variation is constant throughout the pipe. However, it is also possible to apply different signal manipulations to different sections of the signal to compensate for localised environmental variations. This approach is outlined in [19] and could easily be implemented as an addition to the work outlined in this thesis, where appropriate.

Although originally developed separately it is also possible to combine optimal baseline selection and baseline stretch to make the most efficient use of the available data. These combined techniques can be used to create a simple monitoring system which allows detection of cross sectional changes of less than one percent [64]. These approaches do however have a major drawback in that they are dependent on the efficient collection of baseline data. To be truly effective this baseline data set should be collected under the full range of environmental conditions, but with no structural changes occurring. In practice these requirements are difficult to meet, especially on already operational systems.

A technique closely related to regression analysis is novelty detection. Novelty detection also relies on a set of baseline data but makes different use of it. Rather than simply implying a regression relationship between the variables, novelty detection uses the baselines to build a model of the system. The technique used to form this model will vary depending on the implementation but can include simple distribution fitting [72] or a more sophisticated method such as neural networks [73] or support vector machines [74].

When the monitoring system comes on-line, each new reading is compared to the model. If a reading deviates from the model by a certain threshold amount (as determined using some measure such as the Mahalanobis distance [75]) it contains some novelty. If the baselines have been collected correctly this novelty will be due to damage. Indeed, it has been shown in some experiments that changes in a structure can be detected using this approach. In one case simulated damage to an aircraft panel was detected [76] whilst in another damage in an aero-engine was detected using vibration responses [73]. As with regression analysis, the main problem with this approach is that it only works if the baseline data completely captures all of the environmental variations whilst at the same time not including any structural changes. This is very difficult to achieve in a practical system.

4.4. Compensation procedures based on data deconstruction

Recognising that the collection of baseline signals is problematic, other techniques have been developed which attempt to remove the need to define baselines. One of these approaches is known as auto-associative neural networks. A neural network consists of a series of decision points or nodes which are linked together in a large network. The degree to which various nodes are linked together is determined by a weighting factor assigned to each link. These links and nodes are then structured in layers, typically input layer, output layer and multiple hidden layers with each layer of nodes acting as an input to the next layer. Feeding example input and output data to the network allows the weighting values of the links to be learned. The nodes and their associated weighted links then represent some underlying structure which relates the inputs to the outputs. Examining this structure therefore allows hidden relationships to be uncovered. For more information, see for example [77, 78].

In an auto-associative neural network the input layer and output layer are the same. These layers are connected by multiple hidden layers, one of which is constrained to have fewer nodes than the input layer [79]. In this way, the constrained hidden layer is forced to switch to a more basic representation of the input data. When used on the output from a monitoring system, one of the lower dimensional representations may be a new feature which is sensitive to damage and insensitive to environmental changes. In tracking this new feature, damage can be detected without false calls from environmental changes. The principle of this technique has been demonstrated experimentally and produced reasonable results [80] but is yet to be trialled in an industrial environment.

The main limitation of this technique is that neural networks require large sets of data in order to correctly train the weights of the network. Large sets of data are typically not available in most monitoring situations. This technique also has the limitation that it is not based on any traceable physical or statistical reasoning and it is therefore difficult to apply with confidence in new applications.

Another baseline free approach is Principal Component Analysis (PCA), which is effectively the well known singular value decomposition (SVD) with a post-processing stage [81]. When used in monitoring, this algorithm takes as its input a series of signals from the monitoring system. The algorithm then takes this data and transforms it to reveal a new representation of the data which consists of a group of uncorrelated signals. Some of these new signal representations will change as a result of structural changes but not due to environmental changes. Tracking these representations allows damage identification without false calls from environmental changes. The details of this transformation process are available in the literature but it is worth noting that this decomposition exists for all sets of input signals [82].

Principal component analysis and singular value decomposition have been proposed for a number of monitoring systems [66, 83–86] and in these applications they have proved to be effective. Compared to neural network equivalents these techniques have the advantage of being applicable to small sets of data as well as having a firm statistical basis. They are however based on the assumption that the underlying source signals which the data is being decomposed into have a Gaussian distribution [87]. This condition may or may not be met for guided wave monitoring data, depending on the collection conditions. A related technique is independent component analysis (ICA), a technique which similarly performs a decomposition of the data into constituent components. The difference however is that ICA assumes the underlying source signals have a non-Gaussian distribution [88]. It was felt that this may be a more general assumption to apply, although without a direct comparison it is difficult to know which assumption is better (such a comparison is performed in chapter 7). This independent component analysis technique was therefore developed further.

One common issue throughout studies looking at these techniques is the difficulty in making a quantitative comparison between competing methods. Without a framework to make such a comparison it is impossible to converge to best practice for the monitoring community.

4.5. Conclusion

In this chapter the problem of data normalisation has been introduced. This is the need to compensate monitoring data for environmental and operational conditions, whilst retaining useful information regarding the structural changes which are taking place. It was seen that for guided wave monitoring system the major issue is temperature, which causes a change in the propagation velocity of the waves and changes in boundary conditions. Some common compensation techniques have been introduced and reviewed, with the techniques broadly separated into compensation procedures based on signal manipulation and compensation procedures based on data deconstruction. Many of these techniques were demonstrated as useful in given situations, but there has been little consideration in the literature as to which technique is superior in a given scenario. This highlights the need for a quantitative framework within which to compare monitoring techniques. It was also found that although several techniques have proven useful there is scope for further improvement. A new technique, independent component analysis, was highlighted as having potential for application to guided wave monitoring data.

Chapter 5:

Processing guided wave monitoring data using independent component analysis

5.1. Introduction

Independent component analysis is a statistical technique for decomposing mixed data into its constituent elements. It has potential for use with guided wave monitoring data, where the defect signal is to be separated from background feature reflections and coherent noise during the decomposition. This chapter introduces the independent component analysis method and gives examples of its application on representative guided wave industrial signals. The purpose of this chapter is to show whether independent component analysis can be successfully applied to a range of guided wave monitoring signals. Section 5.2 introduces Independent Component Analysis and the implementation used in this chapter, as well as the procedure for selecting the damage sensitive feature. Section 5.3 then introduces the guided wave data that will be passed to the ICA algorithm and section 5.4 shows the results of this analysis. The signals analysed in section 5.4 all relate to discrete echoes from features with small axial extent so echoes are well separated. In section 5.5 we consider the more challenging case of general corrosion growth within which there is a region of more severe growth. Section 5.6 presents the conclusions of the work and a discussion of the potential of the technique. The work presented in this chapter has been published in a journal article [89].

5.2. ICA methodology

Independent Component Analysis (ICA) is a statistical technique for revealing the hidden trends and groupings which underlie a set of data. The technique takes a set of multidimensional data and transforms it into components which are as statistically independent as possible [88, 90]. The intention of using ICA on guided wave data is that after applying the transform one of the independent components will contain data relating to the defect, whilst most coherent noise will be rejected to other independent components; so giving a clearer representation of the defect. Numerous implementations of ICA are available [91, 92] but the FastICA algorithm [88] is used in this study because it achieves similar performance to other algorithms but at greater speed [93]. A brief overview of the FastICA method is given here but the interested reader can find full details in [88].

The FastICA method assumes that the input data $\mathbf{X}$ can be described by the model:

$$\mathbf{X} = \mathbf{AS} \tag{5.1}$$

where $\mathbf{S}$ is a statistically independent representation of the data and $\mathbf{A}$ is a matrix of scalar values which relates $\mathbf{S}$ to $\mathbf{X}$. In the case of guided wave analysis $\mathbf{X}$ is an $[m \times n]$ matrix of guided wave signals collected from a sensor. Each of the m rows is a guided wave response signal as a function of time (which can be converted to distance knowing the wave velocity), with each of the n columns containing the amplitude of that signal at a certain time. The signals which are the rows of $\mathbf{S}$ will be the new representation we want in which the defect signal and the coherent noise have been separated into different components. Independent component analysis allows us to find $\mathbf{S}$ by performing the operation $\mathbf{S} = \mathbf{WX}$, where $\mathbf{W}$ is the inverse of the mixing matrix $\mathbf{A}$. To calculate $\mathbf{W}$ the algorithm uses the property that the signals in $\mathbf{S}$ must, by definition, be maximally statistically independent. The algorithm iterates through possible values of $\mathbf{W}$ to find the matrix which minimises mutual information between the rows of $\mathbf{S}$.

When FastICA is applied to guided wave data the rows of $\mathbf{S}$ will be guided wave signals containing the reflections from different features, whilst the columns of $\mathbf{A}$ show the amplitude of these signals across the input data. The columns of $\mathbf{A}$ therefore track the amplitude progression of the rows of $\mathbf{S}$, the components. In this chapter all columns of

A are scaled such that the maximum value in each column is one, with the amplitude of the rows of **S** adjusted accordingly. This allows the weighting functions to be compared easily across different components. Note that the product of the row of **S** and the weighting function **A** will give the true amplitude of that component of the signal. How the amplitude is apportioned between **A** and **S** is a matter of presentational clarity. It is then necessary to know which row(s) of **S** contain information about defect growth. In this initial investigation, since the defect location is known we merely select the row that is most similar to the defect reflection response. However, if the method is to be used in practice with unknown defects it will be necessary to produce an automatic method to distinguish components relating to true signal growth from those relating to random noise or environmental effects. This issue is address in chapter 7 where the Mann-Kendal test is used to reliably identify the defect component.

5.3. Datasets

5.3.1. Laboratory experiment

The first set of data comes from an experiment with a 6 m length of 8 inch schedule 40 piping that was temperature cycled in the laboratory. This data was collected by Guided Ultrasonics Ltd and subsequently given to Andrea Galvagni for research use. The purpose of this experiment was to collect commercial quality guided wave signals from a length of plain pipe experiencing known changes in environmental condition. Since temperature is the main contributor to signal changes in guided wave inspection, this was the condition which was varied. With a sufficient pool of data collected in the undamaged condition, the second stage was to introduce a small point defect and see how this changes the received signal.

The pipe was heated via a resistive heating element which was suspended in the centre using sheet metal inserts. In this way the element is prevented from contacting the pipe wall whilst the sheet metal inserts produce negligible reflection during the guided wave inspection. A commercial guided wave sensor (gPIMS unit manufactured by Guided Ultrasonics Ltd) was then installed on the outside of the pipe 2 m from one end as well as seven thermocouples distributed across the outer surface. This assembly was then wrapped in an insulating material. A diagram of this set-up is given in figure 5.1.

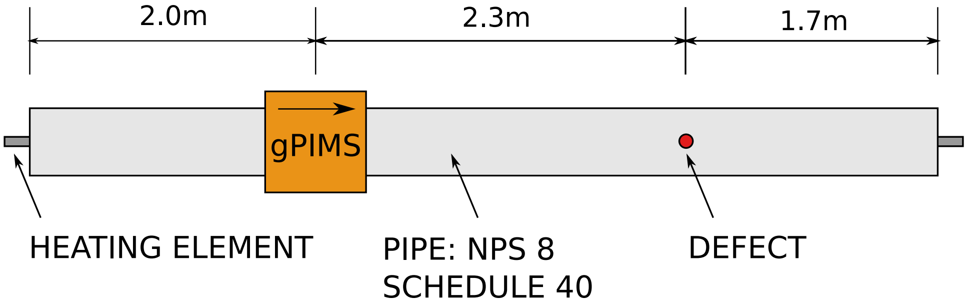

Figure 5.1.: Equipment for the laboratory experiment: NPS8 pipe with internal heating element and guided wave sensor (gPIMS) permanently attached. Note the non-contacting heating element through the centre of the pipe. Not to scale.

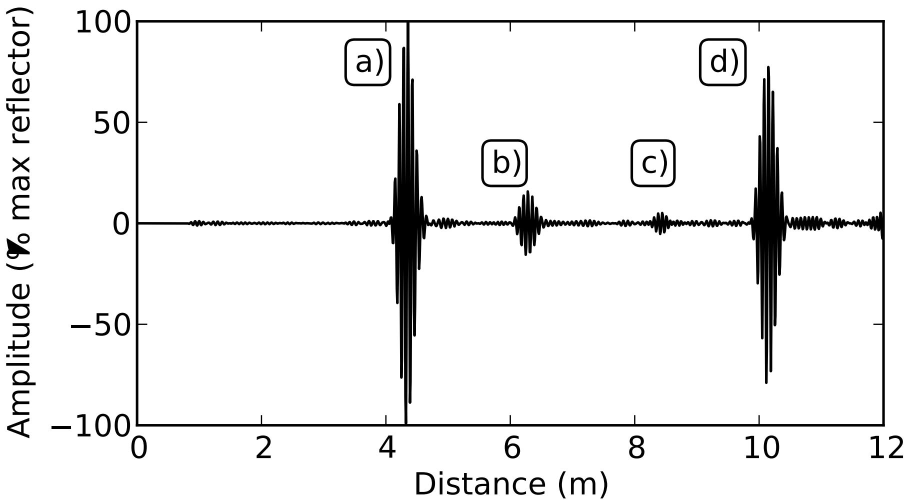

Figure 5.2.: Example pulse-echo measurement from the laboratory experiment, showing reflected T(0,1) mode in the forward direction (forward direction indicated by arrow in figure 5.1). Sensor located at the origin, with reflections caused by a) end reflection, b) imperfect direction control, c) ring reflection and d) first reverberation.

Using the thermocouples, the heating element and an appropriate feedback control loop the temperature of the pipe was varied. The pipe was heated to 90^{o}C, held to allow the temperature to equalise throughout the pipe and then cooled to 30^{o}C. During this cooling cycle 42 measurements were made using the gPIMS sensor in approximately equal steps of time. The measurements were made using an eight cycle Hanning window toneburst, centre frequency 31 kHz, T(0,1) mode. This process was repeated for six temperature cycles to give a total of 252 measurements across a range of temperatures. An example measurement from the sensor is given in figure 5.2 where only the forward direction is shown (forward direction indicated by arrow in figure 5.1). Note that we consider even those parts of the signal which indicate a distance further than 4 m. These are not caused by physical features beyond the end of the pipe, but arise due to reverberations between the pipe ends. These signal components serve as a useful model for a longer pipe with multiple features, which is how they are used in this study.

After this initial period of temperature cycling the insulation was removed and a flat bottom hole was drilled 2.3 m from the gPIMS sensor. This hole had a depth of 4 mm and a diameter of 7 mm, representing a cross sectional change of 0.5% at its maximum extent. In the experiment the hole was drilled in stages but in this chapter only the signals collected before drilling and at the end of drilling have been used. The signals collected whilst the hole was being enlarged are used in chapter 7. The insulation was then reapplied and the same process used to collect readings from the damaged pipe between 90^{o}C and 30^{o}C.

5.3.2. Industrial pipe loop experiment

The second set of data comes from a pipe loop installed at an industrial site. This data was collected by Andrea Galvagni with the assistance of Guided Ultrasonics Ltd. [19]. The purpose of this experiment was to collect commercial quality guided wave signals from a section of pipe with features such as welds, bends and flanges. The other purpose of the experiment was to understand how these signals change when defects are introduced in both plain sections of pipe and at welds, before and after bends.

Experiments were performed on a section of pipework made from NPS8 schedule 40 piping. At one end the pipework begins with a blanked flange, connected to a straight section of pipe 5.79 m long. A weld then joins this to another section of straight pipe which is 6.82 m long. A Guided Ultrasonics Ltd gPIMS unit was permanently attached to this length of pipe, 11.17 m from the flange. This in turn is connected by a weld to a 90^o bend, beyond which there is a further section of straight pipe which extends beyond the region of interest. The pipe is a closed system filled with a liquid whose temperature can be controlled; in this way the temperature of the pipe could be controlled during the experiment. This set-up is shown in figure 5.3.

Initially the temperature of the pipe loop was set to 38^{o}C and four measurements were made under the no damage condition. The measurement was made using an eight cycle Hanning window toneburst, centre frequency 31 kHz, T(0,1) mode. The temperature

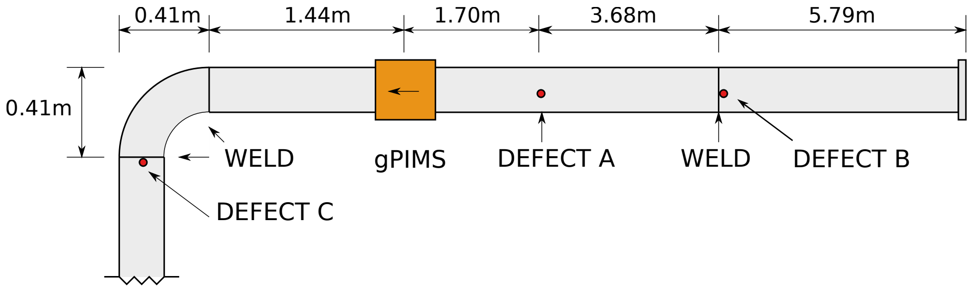

Figure 5.3.: Equipment for the industrial pipe loop experiment: feature-rich pipe with welds and bends, filled with temperature controlled fluid and with a guided wave sensor (gPIMS) permanently attached. Not to scale.

of the loop was then ramped to 90^{o}C and allowed to cool again to 38^{o}C. The purpose of this ramping was to cause thermal stress in the sensor and study the effect of sensor changes due to environmental cycling. Three 1 inch diameter flat bottom holes were then drilled to give defect A, B and C indicated in figure 5.3. Defect A was drilled in a clean section of pipe and should be the easiest to detect. Defect B was drilled adjacent to a weld and will be more difficult to detect due to masking by the weld reflection. Defect C is adjacent to a weld and beyond a bend, making it the most difficult to detect due to masking and the fact that bends tend to exacerbate environmental effects. Initially the holes in these three locations were drilled to a depth which gives a 0.25% cross sectional change. Four measurements were made using the sensor and the loop is again ramped to 90^{o}C and cooled to 38^{o}C. The holes were enlarged to 0.5% cross sectional change and another four measurements made. This process was repeated a final time for a hole with 0.75% cross sectional loss. An example output from the sensor is given in figure 5.4, with positive direction indicated by the arrow in figure 5.3.

5.3.3. Data post-processing

It is possible to apply Independent Component Analysis (ICA) directly to the signals from the laboratory and pipe-loop experiments. However, it was found that the ICA performed better if the signals were first partly compensated for environmental changes. For the results presented in this chapter each signal has been compensated using a stretch algorithm. This stretch involves dilation or compression of the signal in the time domain, an amplitude scaling and a delay, the purpose being to bring the signal as close as possible to a reference signal. This manipulation was performed in the frequency domain and in this study was achieved using a simulated annealing optimisation to optimise across the

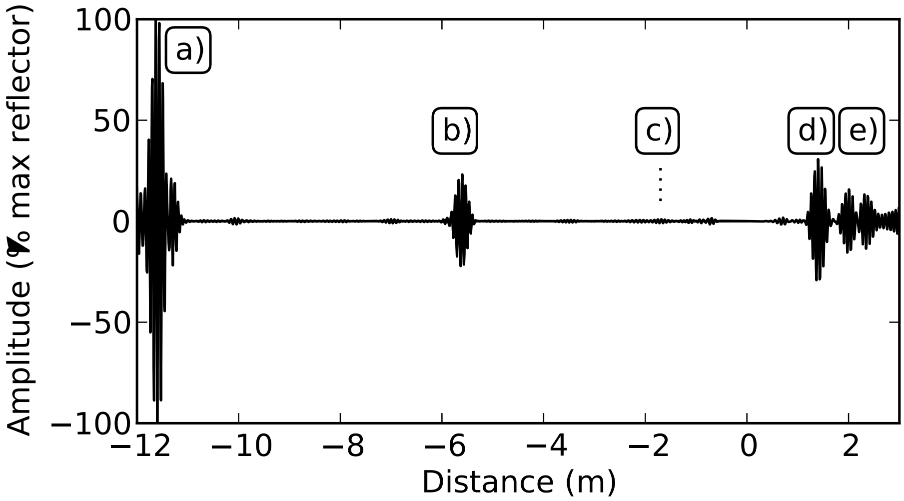

Figure 5.4.: Example pulse-echo measurement from the industrial pipe loop experiment showing reflected T(0,1) mode in the forward and backward directions (forward indicated by arrow in figure 5.3). Sensor located at the origin (0 m), with reflections caused by a) flange, b) weld and defect B, c) defect A, d) weld before bend and e) weld after bend and defect C.

three variables simultaneously. For details on this stretch procedure see chapter 4. For both the laboratory experiment and industrial pipe loop experiment a suitable reference signal had to be chosen to stretch to. This reference was arbitrarily chosen as the signal with a temperature closest to the mean.

5.4. Results of applying ICA

5.4.1. Laboratory experiment with simulated simple step defect

The first test of the algorithm was to see if a simple step change in a pipe undergoing temperature changes could be detected. This scenario was modelled by taking data from the undamaged pipe in the laboratory experiment and adding a signal which models the reflection from a point defect. This first test looked at three scenarios: case A which uses twenty experiment signals collected in the temperature range 60-63^{o}C, case B which uses twenty signals collected in the temperature range 60-70^{o}C and case C which uses ten signals in the temperature range 60-70^{o}C and ten signals in the temperature range 71-80^{o}C. The specific temperatures for each case are given in table 5.1. Note that

Table 5.1.: Temperature distribution of signals across three example cases: A, B and C

	Temperature ^oC (+/- 0.5^oC)		
Signal Number	Case A	Case B	Case C
1	60	70	70
2	60	68	68
3	60	60	60
4	61	68	68
5	61	69	69
6	61	61	61
7	62	64	64
8	62	63	63
9	63	68	68
10	63	64	62
11	60	69	72
12	60	62	74
13	60	63	71
14	61	61	71
15	61	61	79
16	61	69	80
17	62	66	76
18	62	65	71
19	63	61	72
20	63	60	74

these temperatures do not come directly from thermocouple measurements but have been inferred from the dilation or compression between signals within a known temperature range. In all three cases the first ten signals were left in the undamaged condition whilst the second ten signals had another signal added. This signal is a model of the reflection from a point defect and is simply an amplitude scaled version of the excitation signal used by the sensor: an eight cycle Hanning window toneburst, centre frequency 31 kHz. This signal was added at a constant amplitude of 0.75%, thus simulating a simple step change in the structural condition of the pipe. The amplitude is given here as a percentage of the amplitude of the largest reflector in the signal (the first end reflection). This signal was added at a distance of 7.1 m from the transduction, remote from the major feature reflections. Note that this is not a physically possible damage location in the 6 m pipe, but instead is a portion of the signal that arises due to reverberation between the pipe ends.

Some of the outputs from the ICA algorithm for case A are shown in figure 5.5. Parts a) and b) are the signal (row of matrix s) and weighting function (column of matrix A)

for the defect signal, showing a defect signal which is clearly separated from coherent noise and a weighting function which correctly tracks the true amplitude of the defect [dotted line in part b)]. Parts c) to f) show some of the other output signals from the algorithm. Parts c) and d) show the signal and weighting for end reverberations and ring reflections, these parts having constant weighting because these parts of the signal do not change across the input signals. Parts e) and f) show the signal and weighting for one of the components of noise introduced by imperfectly compensated environmental effects. The noise component with the largest amplitude has been shown, but the ICA algorithm outputs a total of 6 noise signals for this set of inputs. These noise signals all have similar oscillating weighting functions but their maximum amplitude varies from 2% in the case shown in figure 5.5 e) to a minimum of 0.9%. Other noise components are shown in appendix A. Although in some cases the noise components have a larger amplitude than the defect signal, their weighting function is different to that of the defect. This difference in weighting function is one way of distinguishing between damage and noise components as shown in chapter 7.

The ICA output for temperature case B is shown in figure 5.6, parts a) and b), where only the defect signal has been shown. The other outputs of the algorithm are similar to those shown in figure 5.5 although there are now 10 noise components due to imperfectly compensated environmental variations. These noise components have a maximum amplitude which varies from 4.0% to 0.2% across the 10 noise outputs. This increase in the number of components is due to the greater temperature difference amongst the input signals, which causes an overall increase in the background noise. Although the maximum number of outputs the algorithm can calculate is equal to the number of inputs, in some cases where there is little variability in the input data a smaller number of components will be outputted. These noise components all have non-monotonic weighting functions that vary in a pattern similar to that shown in figure 5.5 f). Similarly, the defect signal for temperature case B has more coherent noise than temperature case A. This is again thought to be because there is greater temperature difference amongst the input signals and hence the algorithm has greater difficulty separating the defect signal. The ICA output for temperature case C is shown in figure 5.6 parts c) and d), where only the defect signal has been shown. As before, the other outputs of the algorithm are similar to those shown in figure 5.5 although there are now 14 noise components in the ICA output. These noise components have a maximum amplitude which varies from 6.0% to 0.1%, with the maximum noise occurring at the end reflections. Again, the weighting functions for these noise components are non-monotonic, varying in patterns similar to those shown in figure 5.5 f). The noise on the defect signal is also higher than temperature case A or B, again because there is greater temperature difference on the input signals.

5.4.2. Laboratory experiment with simulated growth of a flat bottom hole

The second test of the algorithm was to look at a more realistic model for defect growth. Instead of modelling a simple step change, this test modelled the reflection from a flat bottom hole with increasing cross sectional loss over time, but a constant depth to diameter ratio. This is a model of, for example, growing localised corrosion. This scenario was modelled by taking twenty signals from the undamaged pipe in the laboratory experiment and adding a defect signal to the final ten. For the sake of brevity only temperature case B is presented in this chapter, where all signals have a temperature within the range 60-70^{o}C (see table 5.1 for details). The added defect signals were generated using the analytical model of Cegla et. al. [94]. Using this model the reflection coefficient vs. frequency behaviour of ten different flat bottom holes was calculated. Each hole had a depth to diameter ratio of 1:2 and represented a cross sectional loss of between 0.1% and 1.0% in the 8 inch pipe. These reflection coefficient curves are shown in figure 5.7. To generate a signal to add to the laboratory data the frequency spectrum of the inspection signal (eight cycle Hanning window toneburst, 31 kHz centre frequency) was multiplied by these curves. Transforming this data back into the time domain gave a signal with the correct amplitude and frequency content for a reflection from each hole. Note that as the hole size increases the frequency content of reflection changes, so the defect signal is no longer simply an increasing amplitude of the same signal. It is of interest to see whether ICA can still identify this hole as a single component at all hole sizes. Although this model is for flat bottom holes in plates, the ratio of pipe radius to hole radius was large enough to justify this approximation. This signal was again added at a distance of 7.1 m, away from major feature reflections.

The results from the ICA algorithm are given in figure 5.8 which shows the defect signal and associated weighting function. The figure shows that the reflection from the defect can be easily identified and the weighting function correctly tracks the true weighting function [dotted line in figure 5.7 b)], albeit with some noise. The other outputs from the algorithm are similar to those shown in figure 5.5, with a total of 10 noise outputs in this case. The maximum amplitude of these noise components varies from 3.7% to 0.4% and their weighting functions vary non-monotonically in a pattern similar to figure 5.5 f).

5.4.3. Laboratory experiment with drilled flat bottom hole

The previous tests with laboratory data have used signals from the undamaged pipe with added artificial defect signals. This is a reasonable approach given that the interaction of guided waves with defects in pipes is well understood [94]. However, there may be some subtlety that arises with a physically introduced defect that is not captured in this previous approach. For that reason the next test of the ICA algorithm involved data from both the undamaged and damaged pipe in the laboratory experiments. Thirty signals were selected, fifteen from a pipe in the undamaged condition and fifteen from a pipe containing a 7 mm diameter, 4 mm deep flat bottom hole; such a hole represents a cross sectional loss of 0.75% This hole was drilled 2.3 m from the sensor. These thirty signals were collected in the temperature range 60-70^{o}C, with the temperatures covering this full range. These signals were then passed to the ICA algorithm, although note that only the first 4 m of the signals were processed to keep the analysis physical.

The results of the ICA algorithm are given in figure 5.9 which shows a) the reflection from the flat bottom hole (at a distance of 2.3 m) and b) the associated weighting function. The algorithm has extracted the defect signal with a fairly low level of coherent noise given the size of the defect and has accurately tracked the amplitude growth of the signal. Note that in this case the algorithm is trying to extract a defect reflection of around 0.75%, whereas in one off guided wave non-destructive testing the operator would typically look to detect changes of 5% or greater. The other outputs of the ICA algorithm are similar to those shown in figure 5.5, with a total of 12 noise components. These noise components have a maximum amplitude which varies from 1.4% to 0.2% and again the corresponding weighting functions vary non-monotonically in a pattern similar to figure 5.5 f).

5.4.4. Industrial pipe loop experiment with physical flat bottom holes

The last analysis on signals with discrete echoes looks at the reflection from holes physically introduced in the industrial pipe loop experiment. As with the previous test the defects in this case have been physically drilled in the pipe. The difference compared with the previous case is that these holes have been grown in stages and have been drilled not only in clean sections of pipe but also close to welds and beyond bends. Fourteen signals were selected, with this group representing four different damage conditions: four signals in the undamaged condition, four signals with holes representing a 0.25% cross sectional change, four signals with holes representing a 0.50% cross sectional change and three signals representing a 0.75% cross sectional change. All of these signals were collected at a temperature of 38^{o}C but with temperature cycling in between each damage condition to simulate ageing of the sensor and hence possible changes in sensitivity or frequency response

The analysis of this data is split into two: the data from the 'forward' direction looking towards the bend and the data in the 'backwards' direction which looks towards the flange (forward direction indicated by arrow in figure 5.9). Looking in the backwards direction first [figure 5.10 parts a) and b)], the ICA algorithm extracts defects A and B, which are in a clean section of pipe and close to the weld respectively. Both defects occur on the same output because the holes were drilled at the same time. All signals therefore either have neither or both defects present. Although the defects have the same cross sectional loss, they have different amplitudes in the output from the ICA. It is thought that the amplitude of defect B is underestimated because of the difficulty in separating it from the reflection from the neighbouring weld. The weighting function [part b)] correctly tracks the known weighting function of the defect. The other outputs from the ICA algorithm are the unchanging weld and flange reflections and 12 noise components. These noise components have a maximum amplitude between 4% and 0.2% and weighting functions which vary non-monotonically in patterns similar to figure 5.5 f). In the forwards direction [figure 5.10 parts c) and d)] the algorithm extracts the reflection from the drilled hole with little coherent noise. The algorithm also correctly tracks the weighting of the defect, albeit underestimating the amplitude when the defect is small. The fact that the weighting function is poorer in the forward direction than the backwards direction is to be expected: a defect at a weld beyond a bend is a challenging inspection scenario, especially for this size of defect. The other outputs from the ICA algorithm are the unchanging reflection from the welds before and after the bend and 12 noise components. These noise components have a maximum amplitude between 4% and 0.2% and weighting functions which vary non-monotonically in a pattern similar to figure 5.5 f).

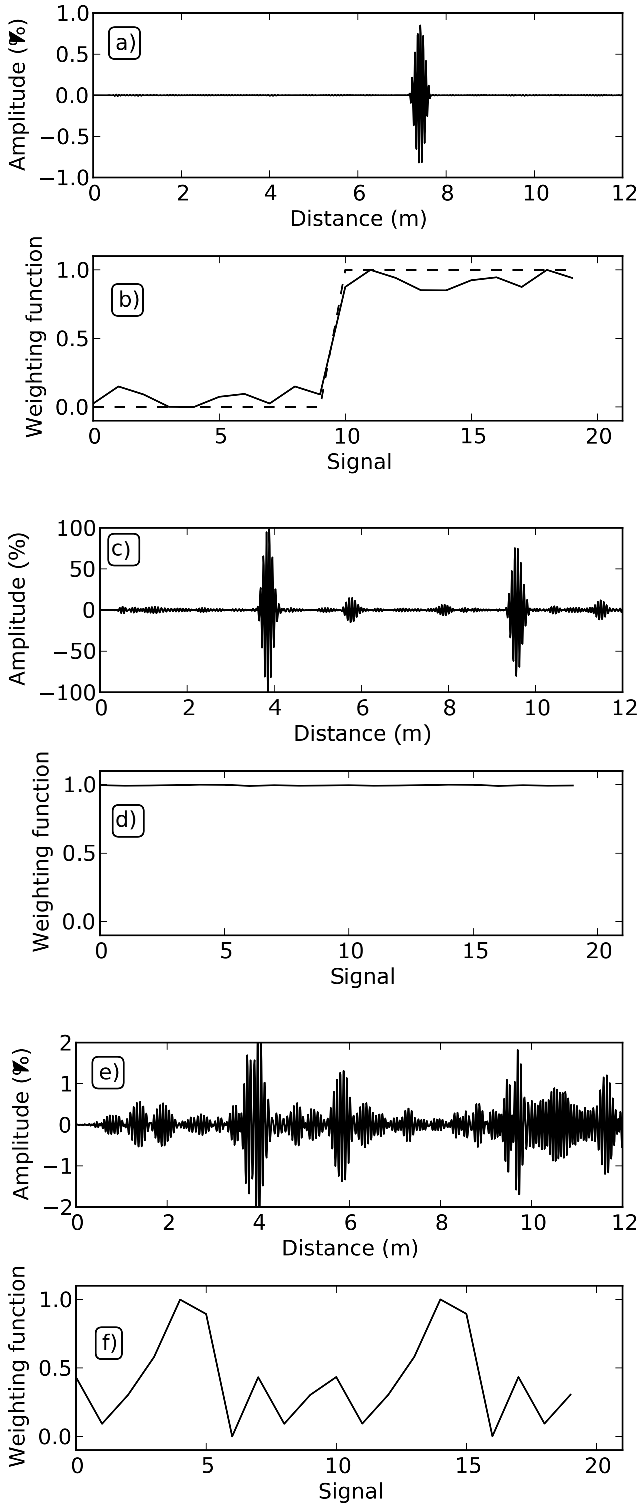

Figure 5.5.: Output from the ICA algorithm for a simulated step defect in the laboratory experiment (temperature case A), showing: a) the recovered defect signal and b) the associated weighting (solid line is the weighting calculated by ICA, whilst the dotted line is the true), c) the major feature reflections and d) the associated weighting function , e) one component of the noise introduced from imperfectly compensated environmental effects and f) the associated weighting function

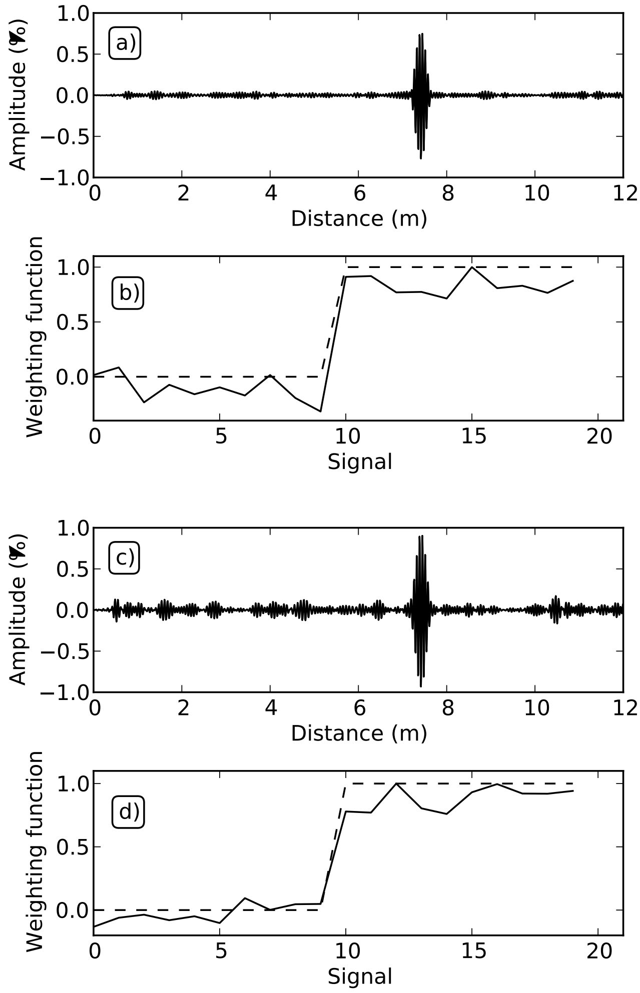

Figure 5.6.: Example output of the ICA algorithm for simulated step defects in the laboratory experiment, showing: a) defect signal for temperature case B and b) the associated weighting function, c) the defect signal for temperature case C and d) the associated weighting function

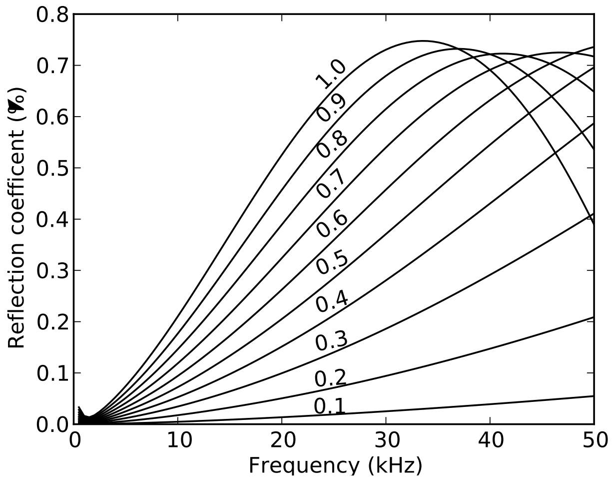

Figure 5.7.: Reflection coefficient vs. frequency behaviour for a flat bottom hole. Values along each curve are the percentage cross sectional change represented by each curve. All holes have a constant depth to diameter ratio of 1:2

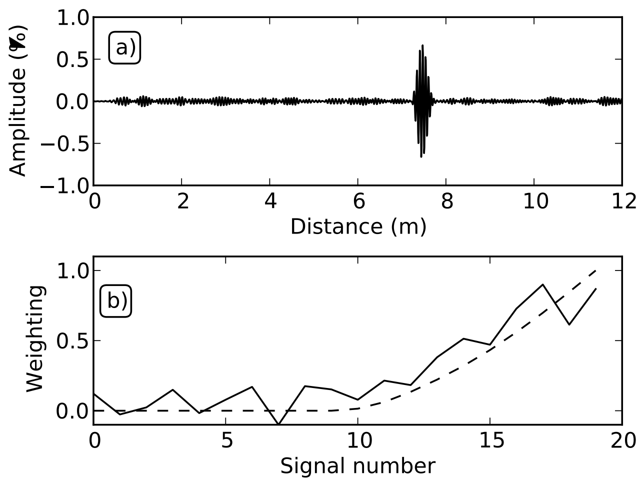

Figure 5.8.: Output from the ICA algorithm for a simulated flat bottom hole grown in stages in the laboratory experiment data, showing: a) defect signal and b) the associated weighting function. Solid line is recovered weighting function, whilst dotted line is the true. Signals have the temperature variation described by case B.

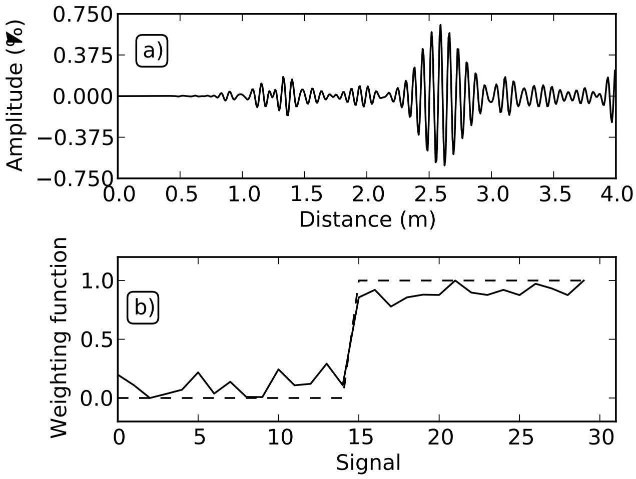

Figure 5.9.: Result of applying the ICA algorithm for a drilled flat bottom hole in the laboratory experiments, showing: a) the reflection from the flat bottom hole, beginning at 2.4 m and b) the associated weighting function (dotted line is true). All signals at 60-70^{o}C.

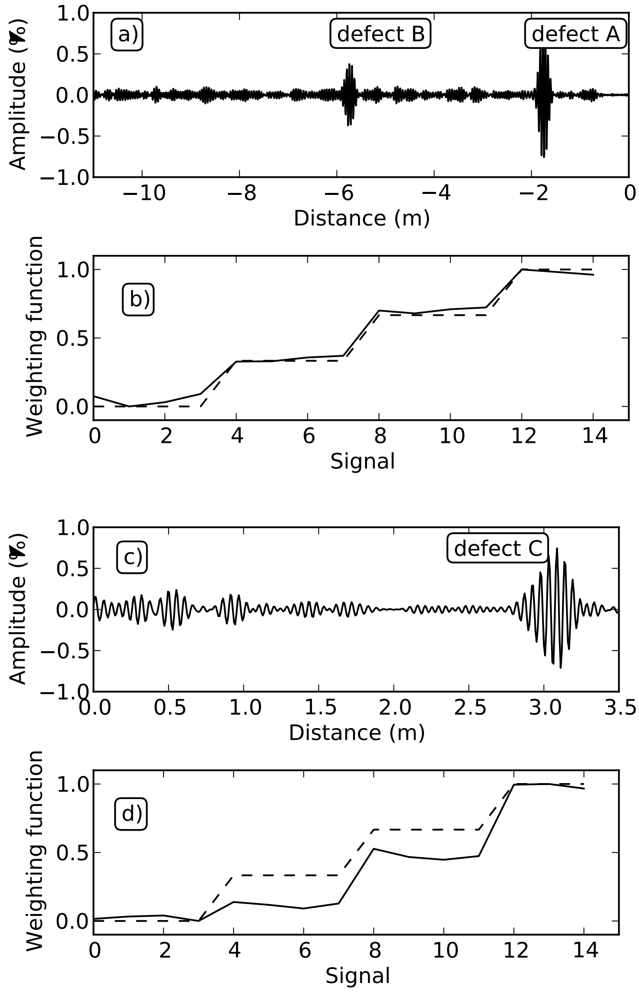

Figure 5.10.: Results of applying ICA to pipe loop data showing: a) defect A and B, b) the associated weighting function, c) defect C and d) the associated weighting function. All signals collected at 38^oC.

5.5. Simulated general corrosion

5.5.1. Method

The analysis so far has focused on discrete echoes from features with small axial extent. The purpose of this section is to consider the case of interacting echoes from distributed general corrosion. Specifically we are interested in whether the ICA algorithm is able to identify the reflection from a region of deep corrosion growing within a patch of shallower corrosion i.e is the algorithm capable of isolating severe damage from benign damage. Finite element simulations were used to study this issue given the difficulty of preparing long sections of pipe with a prescribed surface profile.

The finite element model consisted of a 3.0 m long section of pipe, 168 mm outer diameter and a wall thickness of 7 mm (NPS 6 schedule 40). This model was discretized into tetrahedral elements with a characteristic length of 2 mm using the Netgen algorithm [39]. The elements were then given the properties of steel except for the first 0.3 m which was set up as an absorbing region following Pettit et. al. [34]. The inner surface of the pipe from 0.7 m to 2.7 m was then adjusted to follow a given corrosion profile. The corrosion profile was generated separately as a grid 2.0 m by 0.484 m (inner circumference of the pipe) broken into 2 mm elements. Each point in the grid was given a value taken from a Gaussian distribution with specified mean and given standard deviation. An analysis of industrial corrosion patches indicates that the Gaussian distribution is a suitable model for distributed general corrosion [53]. This surface was then multiplied by a 3D-Gaussian to give the surface a specific correlation length in both directions and all heights greater than zero are set to zero (no increase in pipe thickness). The nodes of the pipe model were then adjusted to match this corrosion profile, with smoothing applied in the radial direction to reduce mesh distortion. For full details of the simulation process for rough surfaces see chapter 3.

A uniform tangential force was then applied to a ring of nodes on the outside surface of the pipe, 0.4 m from one end. This changing force generated a T(0,1) mode, eight cycle Hanning window toneburst, centre frequency 31 kHz. This model was then stepped in time using POGO, a finite element solver based on the GPU [36]. The T(0,1) mode reflection from the corrosion patch was then measured at the location where the T(0,1) mode was generated. This model is shown in figure 5.11.

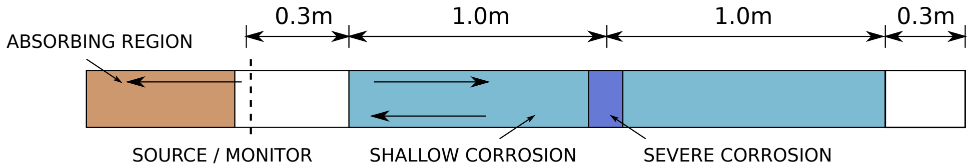

Figure 5.11.: Geometry of the model used in general corrosion simulations. Not to scale

5.5.2. Dataset

Initially ten random corrosion surfaces were generated having a correlation length of 5 mm and a root mean square depth of 0.1 mm. These shallow surfaces were used to generate a low level of background noise to model the low levels of coherent noise observed when inspecting a nominally clean section of pipe. A further ten surfaces were then generated to model a patch of general corrosion gradually increasing in depth and severity.

The first surface was a Gaussian random surface with a correlation length of 5 mm and a root mean square depth of 0.15 mm. However, a small region of this surface was given a root mean square depth of 0.3 mm. This region extended from 0.95 m to 1.05 m along the axis of the pipe and around half of the circumference, with Gaussian smoothing applied at the interface with the rest of the corrosion surface to prevent discontinuities.

To generate the second surface a perturbing surface was then added to this original surface. The perturbing surface had an increased correlation length and root mean square depth and again had a small region from 0.95 m to 1.05 m where the corrosion was deeper. This perturbing surface used a random number grid which was a 50:50 weighting of the previous random number grid and a new random number grid. This process of taking the previous surface and perturbing it was continued nine times to give a total of ten random surfaces with increasing mean depth and correlation length. The correlation length of the perturbing surface was increased from 5 mm to 50 mm and the mean depth was increased from 0.1 mm to 0.5 mm across the nine surfaces. This gave the final corrosion surface shown in figure 5.12, where the shallow general corrosion had a maximum depth of 1.5 mm whilst the deep corrosion pits had a maximum depth of 3 mm. Note that these are absolute limits on the depth; outlier regions deeper than this have been set to the maximum depth. Although convoluted, this process of perturbing previous surfaces was preferable to generating completely new random surfaces for each increment. This is because a perturbation more closely resembles the physical corrosion process. In the real corrosion of pipework, each new corrosion surface must necessarily be an increment of the corrosion surface that preceded it.

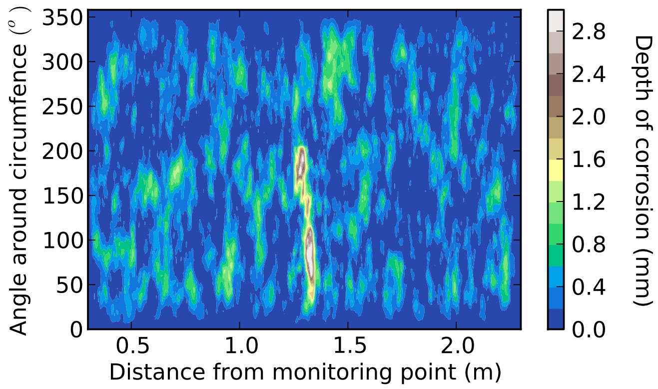

Figure 5.12.: Thickness map of the final state of the simulated corrosion patch

The amplitude of T(0,1) reflection from the general corrosion surface in figure 5.12 is shown as the solid black line in figure 5.13. This reflection amplitude is shown alongside the cross sectional loss represented by the corrosion. As expected there is broad agreement between the two variables, although the agreement will not be perfect due to the complex constructive and destructive interference between overlapping waves. This figure shows clearly the increased reflection from the more severely corroded section of pipe. Note that this figure does not show the reflection from the end of the pipe (2.6 m from the monitoring point) which dominates the received signal.

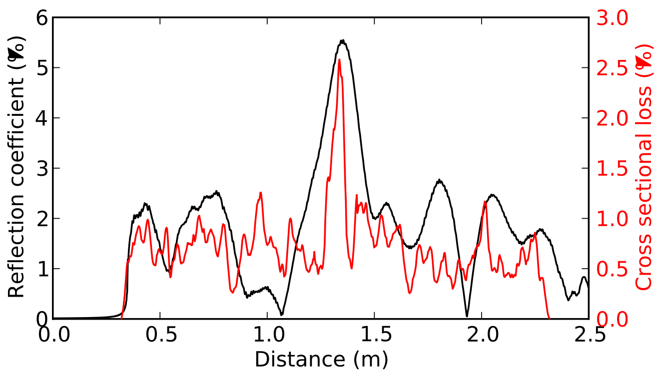

Figure 5.13.: Overlay of cross sectional loss and reflection coefficient for the corrosion patch. Note that the corroded surface starts at a distance of 0.3 m from the monitoring point.

5.5.3. Results of applying ICA algorithm

Some example outputs from the ICA algorithm are given in figure 5.14, where the input to the ICA algorithm is the first 3.0 m of the signals from the model. Parts a) and b) show the signal and weighting function for the severely corroded section of pipe. This signal contains a major reflection from the severely corroded region and is the only signal of this type in the ICA output. This is a very encouraging result, showing that ICA tracks the severe corrosion, keeping it in single component even though the corrosion is growing at more than one point and its shape is changing. Parts c) and d) show the signal and weighting function for the pipe end reflection. This weighting function correctly shows a decrease in the amplitude of the reflection across the final ten signals as the corrosion grows. This happens because the increase in corrosion causes a reduction in the T(0,1) mode reaching the pipe end. Parts e) and f) show the signal and weighting function for one of the signal components due to the shallow general corrosion. Note that the algorithm separates the reflection from the shallow corrosion into 14 components, with the largest amplitude component shown in figure 5.14. The other components have a maximum amplitude from 2.4% to 0.6% and these components constructively and destructively interfere to give the total reflection seen in figure 5.13.

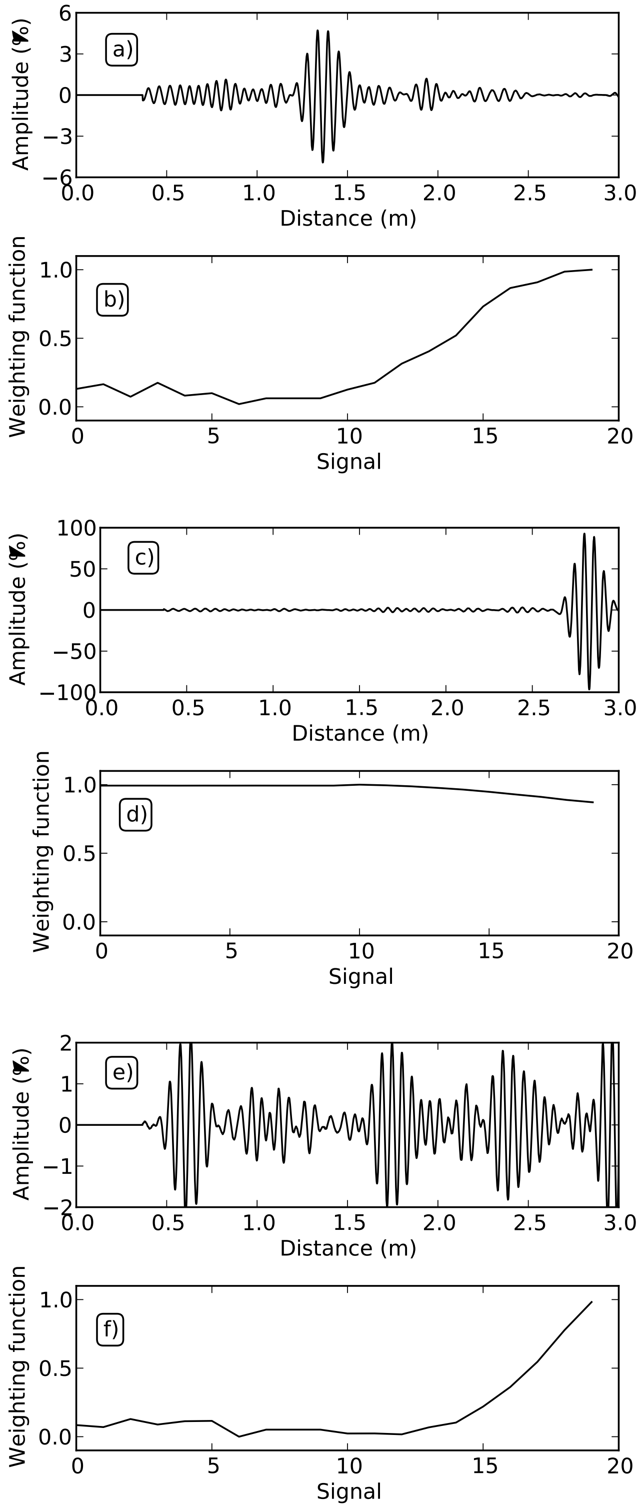

Figure 5.14.: Results of applying ICA to the general corrosion signals, showing: a) the reflection from the severely corroded region and b) the associated weighting function, c) the end reflection and d) the associated weighting function and e) one of the components of the general shallow corrosion and f) the associated weighting function

5.6. Conclusions

In this chapter we have discussed the use of Independent Component Analysis for improved defect detection in guided wave monitoring. Through a combination of the traditional baseline stretch approach and the novel use of ICA we are able to extract a representation of the defect which has lower coherent noise than the original signal. The algorithm can also be used to accurately track the amplitude history of the defect and hence understand its growth pattern. This is in the context of changing environmental conditions.

This chapter has looked at experimental data from both a plain section of pipe and an industrial pipe loop containing welds, bends and flanges. Data was collected from these pipes whilst they were in their undamaged and damaged condition and undergoing temperature variations. Initially model defect signals were added to the data from the undamaged plain pipe to simulate step changes and gradual defect growth. The ICA algorithm was able to separate these defect signals from coherent noise, even in the context of significant changes in temperature and changing frequency content of the defect signal. The ICA algorithm was then applied to data from the undamaged and damaged pipes and again the algorithm was able to separate the defect signal from coherent noise. This was for physically introduced defects in clean sections of pipe, at welds and at welds beyond bends. This chapter also used finite element simulations to investigate generally corroded sections of pipe with a patch of severe corrosion. It was found that the ICA algorithm correctly separated the reflection due to the severe corrosion and recognised it as a single evolving defect. This is in the context of the corrosion undergoing a complex growth pattern where its shape and amplitude is changing.

It therefore appears that ICA is a very promising method for the automated processing of guided wave data. The next stage is to look at ways of automating identification of the component containing information about the defect. The ICA algorithm outputs several components, some of which are noise and some of which contain useful information. A method must be developed for automatically identifying the components containing useful information and work is in progress on this issue. One promising approach is to use the weighting functions of the components; in all the cases reported here the noise components have strongly non-monotonic weighting functions whilst the defect signals have near-monotonic weighting functions. Such a technique is presented in chapter 7, known as the Mann-Kendal test. The ICA technique is also only one of several possible techniques for processing guided wave monitoring data. A quantitative comparison of the different techniques would indicate when each method is most appropriate. Such a comparison is presented in chapter 7.

Chapter 6:

A validation of defect superposition

6.1. Introduction

To evaluate signal processing techniques for guided wave data it is common to apply the techniques to a range of example signals. This was the approach taken in chapter 5, where independent component analysis was applied to a range of example signals. These example signals can be created by making measurements from an undamaged system, physically adding damage, and taking further measurements in the damaged state. These measurements must also be made under a range of environmental and operational conditions (EOCs) for the signals to be representative of field monitoring conditions. However, in structural health monitoring the performance of signal processing techniques may change with different histories of EOC variations and different damage progression. Therefore, the two effects are coupled and need to be sampled simultaneously in order for us to accurately assess the performance of these signal processing techniques. To achieve this, we need to create damage on many structures, which is often not practical since such tests can be expensive and time consuming to perform, especially for the large structures typically of interest in guided wave systems. Hence, we need an alternative way of generating datasets for testing signal processing algorithms.

When long range guided wave inspection was in its infancy it was only possible to do 2D simulations [13], whilst in the last few years full 3D simulations of realistic corrosion patches have become feasible (as detailed in chapter 3). These 3D simulations are now very efficient with the advent of graphics card processing schemes [36]. However, reliable prediction of signal changes due to environmental and other variability is much more difficult because of the complexity of these effects [56, 59, 95]. On the other hand, obtaining experimental data with environmental variation from an undamaged structure is easy. Therefore, we propose a method of generating example guided wave data by superposing artificial defect signals on experimental signals from undamaged pipes undergoing multiple environmental cycles. This hybrid approach enables us to add damage at different locations with different growth patterns, and to easily investigate the effects of other practical parameters such as the degree of EOCs, damage severity and location, frequency of readings, etc. Also, by repetitively randomly selecting from the records that are collected in specific EOC ranges, we can gather sufficient data to statistically evaluate the performance of different damage detection schemes in these practical settings. However, this superposition approach must be validated to ensure that it results in signals which are representative of the equivalent damaged structure.

This chapter presents such a validation. In section 6.2 the modelling approach behind this validation is outlined, whilst in section 6.3 the results of this validation are presented. Section 6.4 then goes on to discuss the limitations of this validation whilst conclusions are given in section 6.5.

6.2. Modelling approach

6.2.1. Generic scheme

The basic approach of this validation is to create two sets of data which represent the same inspection scenario, created using different approaches. The inspection scenario being modelled is a length of pipe containing pipe features, a simple flat bottom hole defect and a source of coherent noise. The first set of data comes from a full finite element model of this inspection scenario. The second set of data comes from a finite element model without a defect, with a synthetic defect signal superposed at the correct time in the signal. This is the superposition approach. A comparison is then made between the two sets of data, with equivalence or near equivalence demonstrating that the superposition approach is valid.

In this chapter the superposition approach is validated for three different monitoring scenarios. The first scenario is for a plain section of pipe with a simple flat bottom hole defect, at constant temperature with background coherent noise. The second scenario is for a plain section of pipe with a simple flat bottom hole defect, at varying temperatures and coherent noise. The signals from these models are processed with a temperature compensation scheme prior to comparison. The third scenario is for a pipe with a weld, with a flat bottom hole defect growing in the weld, varying temperatures and coherent noise. Again, the signals from these models are processed with a temperature compensation scheme prior to comparison. Together these three scenarios cover a wide range of potential monitoring conditions including defects in clean pipe and at pipe features, undergoing changes in EOC and undergoing representative signal post-processing.

These three scenarios were simulated using a finite element model of a 6 inch schedule 40 pipe; outer diameter 168 mm and wall thickness 7 mm. This cylinder was 2.4 m long and was discretized into tetrahedral elements with a characteristic length of 2 mm. The resulting elements were then given the properties of mild steel, with these properties modelled as temperature dependent according to [96]. In this approach, the properties of steel at a certain temperature are interpolated from tables of standard data. It was then necessary to model the coherent noise seen in commercial guided wave inspections. It was important to include this coherent noise, since superposition would not be required on a plain pipe giving no reflected signal in the absence of a defect since it would only be necessary to predict the defect reflection. However, in practical guided wave inspections there is a coherent noise which overlays the reflection from defects and other pipe features. Although coherent noise has numerous sources in practice, it can be simulated in these models by introducing a roughness on the inside surface of the pipe. This roughness has a Gaussian thickness distribution, a mean depth of 2 mm and a correlation length of 50 mm. As seen in chapter 3 of the thesis, such rough surfaces cause scattering of guided waves which in turn generates a low level background noise. This background noise approximates coherent noise. The maximum amplitude of this background noise is typically 2.0% of the inspection wave amplitude, a level consistent with coherent noise levels seen in commercial inspections. The pipe was then set up with absorbing regions following Pettit [34] and a ring of source and monitoring nodes were set up 0.4 m from the pipe end.

In those cases where a defect was included it was modelled as a flat bottom hole in the pipe outer wall, 1.0 m from the guided wave source. The cross section of this hole varied from 0.2 - 2.0% of the pipe cross section but had a constant diameter to depth ratio of 2:1. In some cases a weld was also included and in such cases was modelled as a thickness increase with a ramp-flat-ramp profile, a total length of 10 mm, a rise angle

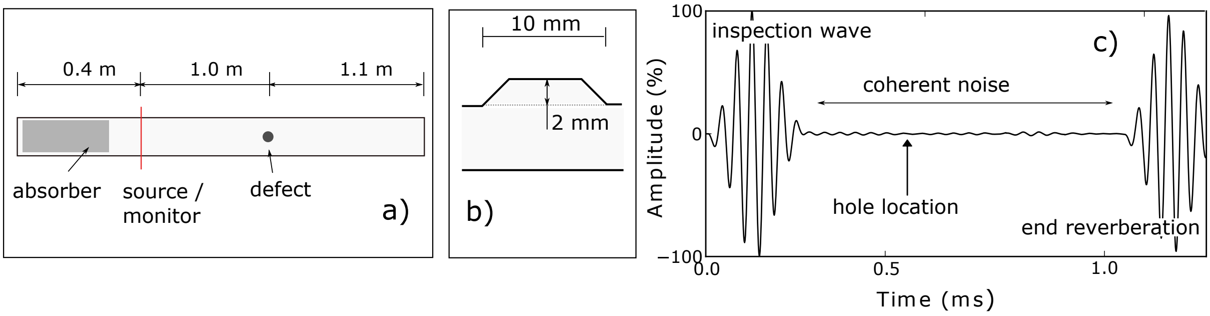

Figure 6.1.: The model used in the validation showing: a) the model set-up, b) detail of the weld structure and c) a typical time trace from the model, with main signal features identified.

of 45^o and a height of 2 mm starting at a distance of 0.995 m from the source of guided waves. In models with both a weld and a hole this means the hole was removing material from the weld cap. An example model of a plain pipe with a flat bottom hole is shown in Fig 6.1a and the weld geometry is shown in Fig 6.1b. These models were then solved using POGO [36], a GPU based finite element solver and where temperature compensation was employed the baseline stretch procedure was used (as detailed in chapter 4).

6.2.2. Datasets

For the first inspection scenario a model of a plain pipe was built. A rough surface profile was then generated based on an initial random seed and this profile applied to the inside surface of the pipe. A second copy of the model was then generated, into which a flat bottom hole was introduced. Initially the cross section of this hole was 0.2%, with a depth to diameter ratio of 1:2 and the hole on the outer surface of the pipe. It is possible to place the hole on the outer surface since the fundamental torsional wave used in commercial inspections has a near-constant stress profile through the thickness. This type of wave is therefore equally sensitive to defects on the inner and outer surface of the pipe.

Both models are then solved to obtain the time history at the monitoring nodes, with an example time history shown in figure 6.1c. An artificial defect signal is then added to the time history for the undamaged pipe, with the artificial defect signal calculated for a flat bottom hole of 0.2% cross section. We generate the artificial defect signal by first generating reflection coefficients of flat bottom holes of growing size on a homogeneous plate using the analytical model developed in [94], modified to compensate for the curvature of the pipe geometry [14]. We then multiply the reflection coefficients from the hole by the Hanning-windowed inspection toneburst in the frequency domain; this frequency domain signal is then inverse Fourier transformed to obtain the reflection signals from the hole at that particular time. To simulate damage at a location of interest, we compute the arrival time of the damage reflection knowing the damage location and then superpose the damage reflection onto the undamaged signals. The maximum amplitude in the time window 0.55-0.6 ms is then compared across the two signals, with this time window corresponding to the time at which the defect reflection appears. This process was then repeated for ten different random rough surfaces to generate ten data points at this hole size. The hole was then increased in size in ten steps from a cross sectional loss of 0.2% to 2.0% in equal steps of cross section. Note that although the cross section increases linearly with time, the amplitude of the damage reflection increases non-linearly and is frequency dependent [94].

For the second inspection scenario the modelling procedure was the same except the models were adjusted for temperature variations. For each pair of models a temperature was chosen at random from within the range 10 - 40^{o}C. The material properties of the model were then updated accordingly before the model was solved. When the artificial defect signal was superposed care was taken to insert the defect signal at the correct time; since temperature variations change the time of flight of ultrasonic waves the defect needed to be inserted at such a time as to ensure that the defect was in the same location. Before the comparison of amplitudes was done, the signals were stretched to 25^{o}C using the temperature compensation procedure outlined in chapter 4. The third inspection scenario was the same as the second except the model included a weld, which produced a reflection amplitude of around 20% which is typical of the weld reflection seen in commercial inspections.

6.3. Results of comparison

The results for the first scenario are plotted in figure 6.2 where the data has been plotted in two different ways. In figure 6.2a the data has been plotted as the maximum amplitude of the defect signal and coherent noise superposed. The data points do not necessarily overlap because the coherent noise has a different amplitude for each rough surface profile. In figure 6.2b the data has been plotted as the amplitude of the defect signal with coherent noise subtracted. This means that each superposed signal has the

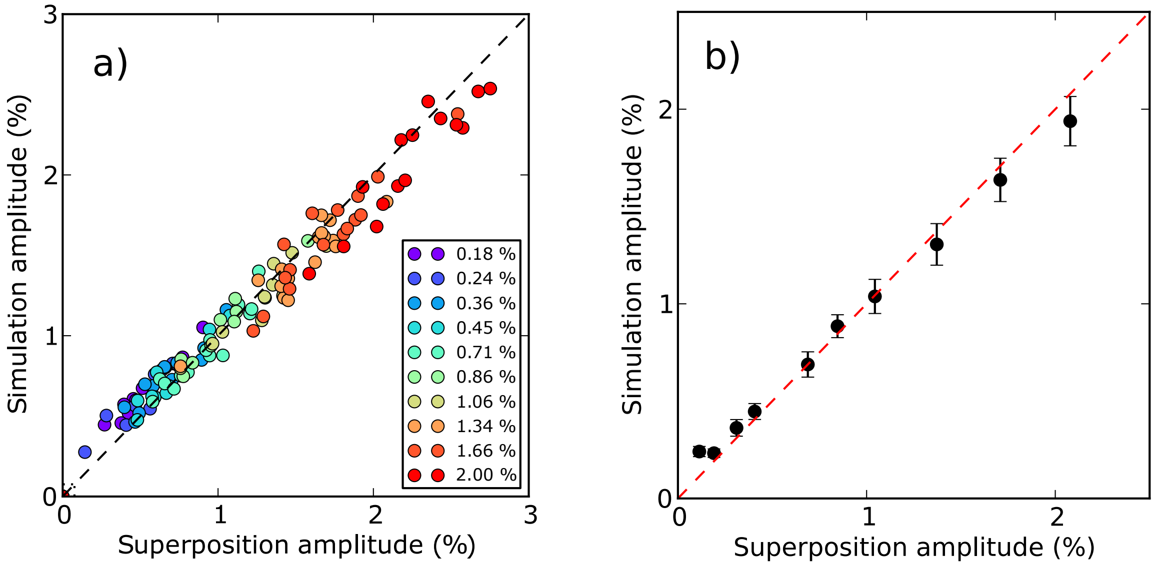

Figure 6.2.: Comparison of signals from the full finite element model with signals from superposition for the first scenario showing: a) results for defect + coherent noise for different defect amplitudes and corrosion profiles and b) results for defect only (after subtraction of the coherent noise). Error bars in b) are the standard deviation across the ten iterations at each defect size.

same amplitude for a given hole size. Plotting the data in this way allows us to see more clearly the variability caused by different rough surface profiles, but with a fixed defect size. In general there is very good agreement between the superposition and full modelling approach. Although there is some difference in the final amplitudes, this difference is small compared with the amplitude of coherent noise in the signals (typically around 2%). Perfect agreement would not be expected as the reflection from a hole of a given size is not strictly independent of the cross section profile onto which it is placed. At smaller defect sizes the superposition approach seems to underestimate the amplitude of the defect reflection, whilst at large defect sizes the superposition approach overestimates the defect amplitude. The variability of results is also a function of defect size as shown in figure 6.2b, where larger defects have greater variability.

The results for the second scenario are shown in figure 6.3, with very similar behaviour as was seen in figure 6.2. This suggests that the validity of the superposition approach is not affected by baseline stretch signal post-processing. It also shows that the validity of the superposition approach is not a function of temperature. The results for the third scenario are shown in figure 6.4. This figure shows that the superposition approach is equally valid at existing pipe features as it is in clean sections of pipe with only coherent

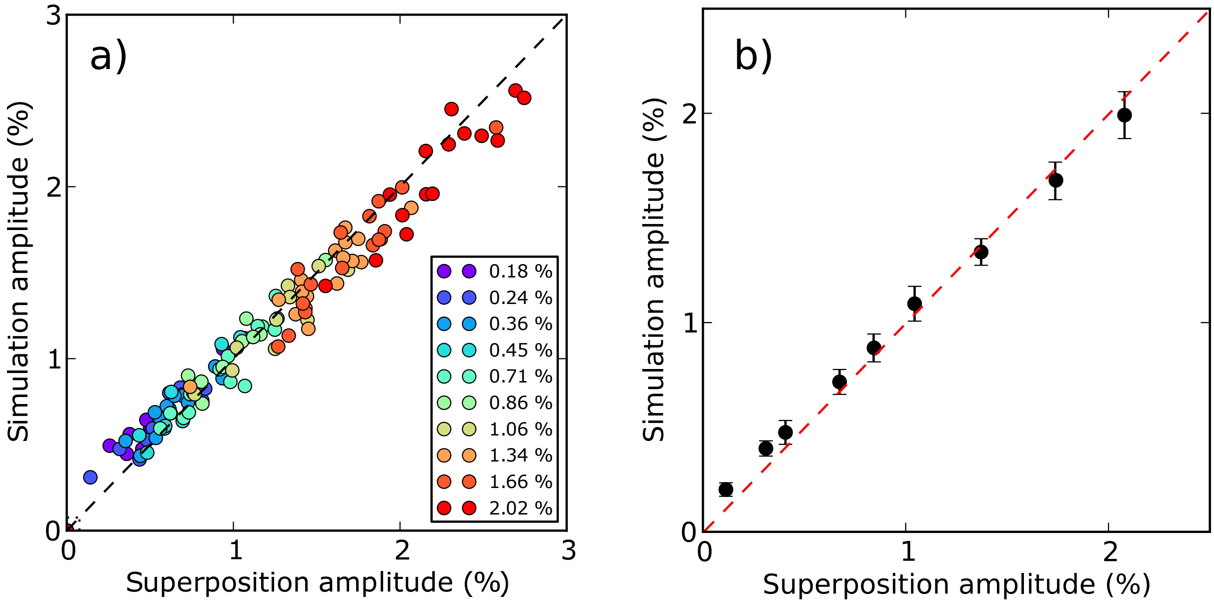

Figure 6.3.: Comparison of signals from the full finite element model with signals from superposition for the second scenario showing: a) results for defect and coherent noise for different defect amplitudes and corrosion profiles and b) results for defect only (after subtraction of the coherent noise). Error bars in b) are the standard deviation across the ten iterations at each defect size.

noise. Note that in this figure the larger defect sizes result in smaller signal amplitudes. This is because the hole is removing material from the weld (which is an increase in the cross section), with the net effect that the overall cross sectional change is reduced.

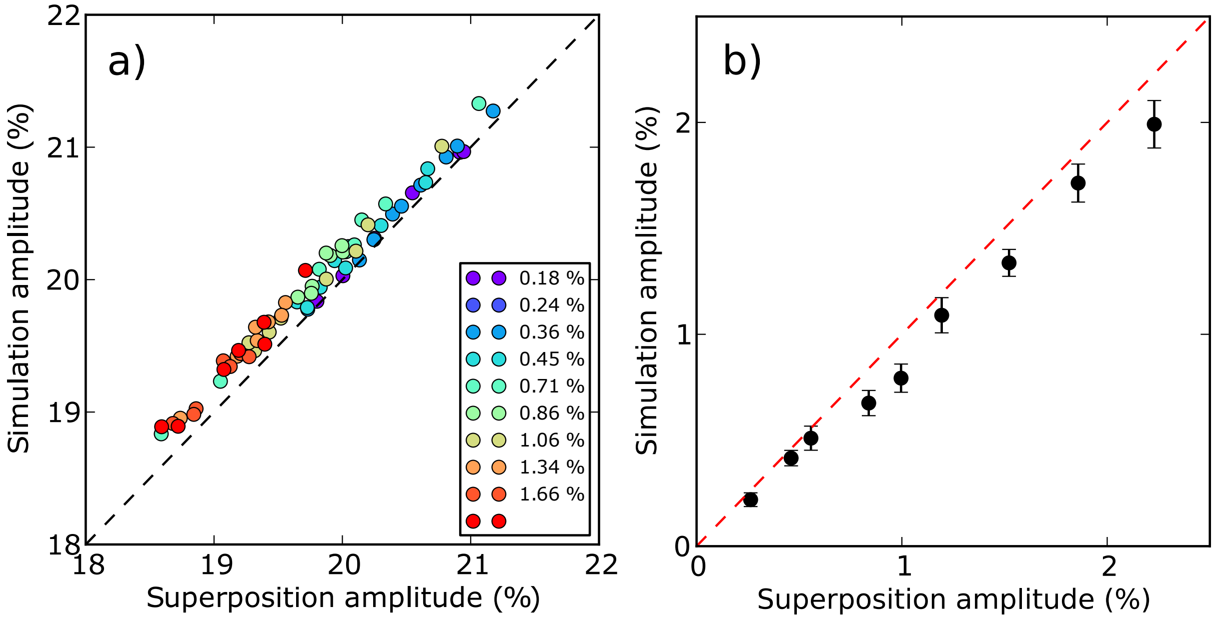

Figure 6.4.: Comparison of signals from the full finite element model with signals from superposition for the third scenario showing: a) results for defect + coherent noise for different defect amplitudes and corrosion profiles and b) results for defect only (after subtraction of the coherent noise). Error bars in b) are the standard deviation across the ten iterations at each defect size.

6.4. Discussion of limitations

In the analysis of the superposition approach certain assumptions have been made about the conditions of the inspection. It is worth considering these assumptions and the subsequent limits of the analysis of this chapter. In the superposition approach an assumption is made that the defect is essentially a small perturbation of the signal. This assumption arises from the fact that superposition only considers reflection from a defect. It does not consider any effect the signal may have on the inspection signal transmitted through the defect. For example, a large flat bottom hole will lead to an appreciable attenuation of the main inspection wave. This will lead to a reduced amplitude of reflection for those pipe features beyond the defect. This effect will also be seen where the defect is, for example, a rough surface corrosion patch. As shown in chapter 3 such rough surfaces lead to a strong attenuation of the transmitted inspection wave, even when the reflection amplitude is small. This shadowing of features downstream of the defect is not considered in the present superposition approach.

The analysis in this chapter also assumes relatively featureless pipelines. It does not, for example, consider the effect of any reverberations of the defect signal. If there are several strong reflectors in the pipe, such as welds near the defect, the reflection from the defect may be bounced between these other reflectors before reaching the sensor. This can lead to several signals appearing in the time trace for only one defect, an effect that is not considered in the current analysis. The analysis also assumes that the inspection signal that reaches the defect is undistorted with respect to the original inspection wave. This may not be the case, if for example the pipe contains several simple supports before the defect location. Pipe supports have been shown to act as high pass filters for guided waves [16] and so can distort inspection waves passing through them.

However, these limitation generally apply to extreme inspection cases for large defects or feature rich pipelines. As shown in this chapter, for reasonable defect sizes the superposition analysis works well.

6.5. Conclusions

This chapter has formalised the superposition approach and sought to compare results obtained with superposition to results from full finite element simulations. This comparison has been performed for three representative inspection scenarios including: simple flat-bottom hole defects in plain pipe and at welds, with and without temperature variations and with representative subsequent signal processing. This is in the context of background coherent noise that produces a low level noise throughout the signal. It was found that superposition worked well and gave very similar results compared with the full modelling approach. As expected, there was some difference between the two approaches, which can be explained by the fact that the reflection from a hole is not totally independent of the cross section into which it is being added. In general it was found that superposition slightly underestimates the defect amplitude for small defects and overestimates defect amplitude for large defects. It was also discussed that the results of this validation are applicable only to small defects in relatively feature-free pipes. For large defects which affect the transmission of the inspection wave, or feature rich pipes where reverberation of the inspection signal is possible, this superposition approach may not work as well.

Chapter 7:

A quantitative comparison of guided wave monitoring algorithms

7.1. Introduction

Having proposed Independent Component Analysis (ICA) for processing monitoring data the next stage is to evaluate the performance of ICA relative to other techniques. Some evaluation methods exists for making such a comparison, such as a probability of detection (POD) analysis [97–99] which requires data from laboratory experiments under varying EOCs and physically growing defects on realistic structures. However, these laboratory experiments are difficult and costly to perform, especially on the large systems typically of interest in guided wave monitoring. Instead in this chapter a new framework is proposed to assess damage detection performance based on the receiver operating characteristic (ROC) curve, which is used in statistics to illustrate the performance of a binary classifier [100], and has been also widely adopted in NDE/SHM [101, 102]. This new framework makes use of the superposition approach validated in the previous chapter.

The methodology can be applied to any guided wave inspection system but in order to make the processing involved clearer, it is presented in the context of a controlled trial of a permanently installed monitoring system. This monitoring system is installed on a length of 8 inch diameter pipe that is subject to temperature cycling, as described in section 7.2. In section 7.3 the new ROC generation scheme is presented and then in section 7.4 three damage detection schemes are reviewed and a new technique is introduced which can be

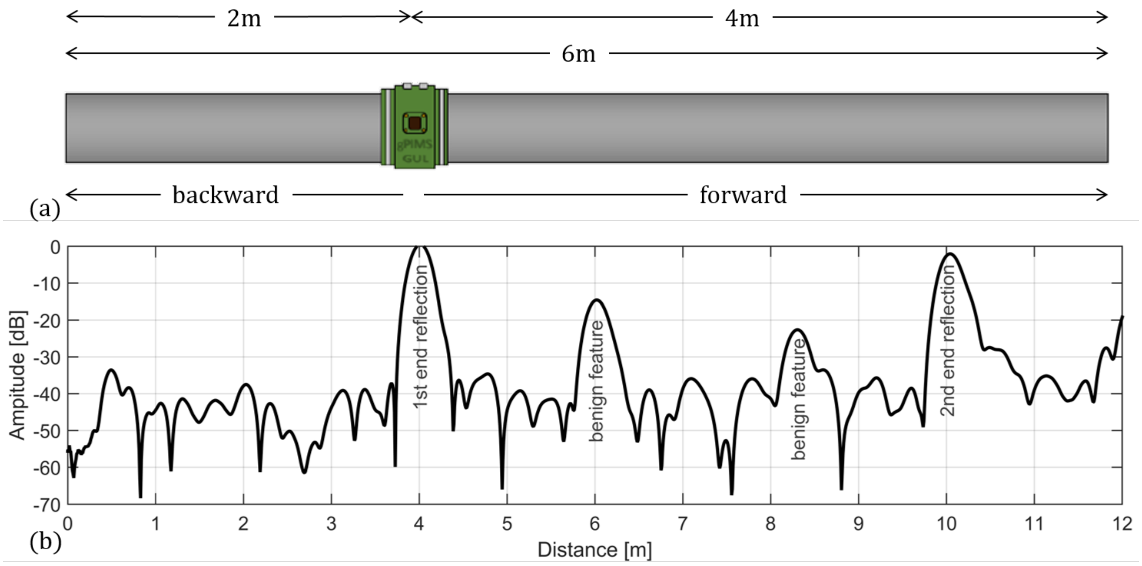

Figure 7.1.: (a) Experimental setup on a 6m 8 inch schedule 40 steel pipe instrumented with a Guided Ultrasonics, Ltd. permanently installed monitoring system (PIMS). (b) A typical T(0, 1) mode, forward-direction A-scan of the pipe including reverberations.

used to identify the damage component in component analysis methods. The performance of the three methods on one synthetic scenario is then compared in section 7.5.1, and the area under curve (AUC) is introduced to compare the ROC curves across different test scenarios in sections 7.5.2. In section 7.5.3 the three methods are compared over the test scenarios listed in table 7.1. The more difficult case where damage is added near a structural feature is discussed in section 7.5.4. An experimental validation of the procedure is then outlined in section 7.6.

The research presented in this chapter was performed in collaboration with Dr Chang Liu, a postdoctoral researcher in the non-destructive testing lab at Imperial College. The concept of trend analysis for detecting the defect component was originated by the author, as was the processing procedure for the independent component analysis technique. Other work in this chapter is the result of analysis by Dr Chang Liu, of which discussion with the author formed a part. The work presented in this chapter has been prepared for a journal submission, expected to be submitted shortly.

7.2. Tests on a Pipe Monitoring System

In this section, an experimental setup is demonstrated which aims to collect commercial quality guided wave signals from a length of plain pipe experiencing realistic variations in environmental condition. Fig 7.1(a) shows a schematic drawing of the experimental setup on a 6 m long, 8 inch schedule 40 carbon steel pipe (wall thickness 8.2 mm). A resistive heating element was inserted and suspended in the centre of the pipe to uniformly cycle the temperature of the pipe. A commercial guided wave sensor (gPIMS unit manufactured by Guided Ultrasonics Ltd.) was installed on the outside of the pipe, 2 m from one end; this sends and receives the torsional (T(0, 1)) mode in the pipe. The whole setup was then wrapped in an insulating material. The experimental setup is used to investigate the effect of temperature variations on commercial quality guided wave monitoring signals. Temperature is one of the main contributors to signal changes in guided wave monitoring,

In each temperature cycle, the pipe was heated to 90^oC, held to allow the temperature to equalize throughout the pipe and then allowed to naturally cool to 30^oC in about ten hours. Guided wave measurements were taken every 15 minutes, resulting in about 40 measurements per cycle, covering a range of temperatures. The excitation used in this analysis was an eight cycle Hanning-windowed toneburst with a centre frequency of 23.5 kHz. The data used in this chapter was the same as that used in chapter 5.

Fig 7.1(b) shows the envelope of a typical forward-direction (defined in Fig 7.1(a)) T(0, 1) mode A-scan; distance is measured from the sensor, which is located at 0 m. The signal is plotted on a dB scale, and is normalized such that the amplitude of the first cut end reflection, seen at 4 m, is 0dB. Note that even those parts of the signal which indicate a distance further than 4 m are considered. These are not caused by physical features beyond the end of the pipe, but arise due to reverberations between the pipe ends and interactions with the transducer ring. These signal components serve as a useful model for a longer pipe with multiple features, such as welds, supports, or bends.

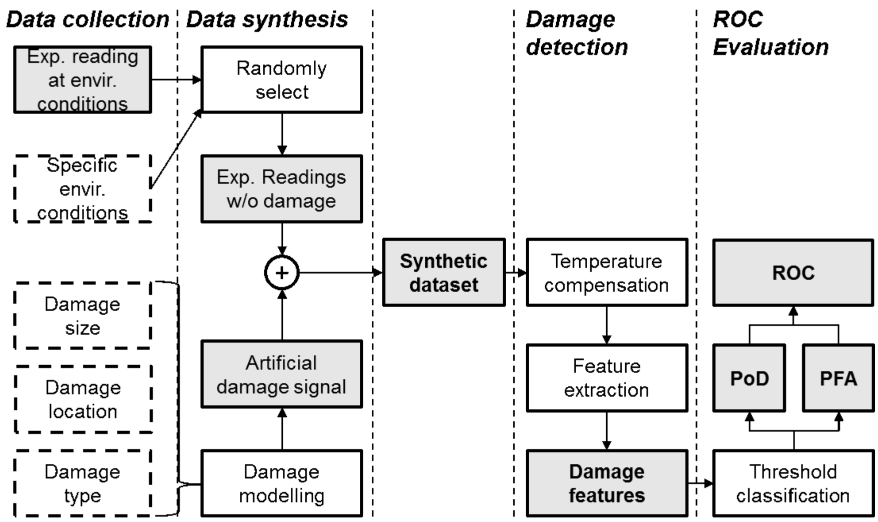

Figure 7.2.: Schematic of the proposed evaluation framework. Synthetic datasets are generated using undamaged experimental signals and predicted damage signals. The ROC curve is then generated after damage detection is performed on the synthetic datasets.

7.3. Evaluation framework for SHM system

7.3.1. Generic scheme

An evaluation framework is proposed which uses experimental data over multiple environmental cycles on an undamaged structure and synthetic damage signatures. Fig 7.2 shows the flow diagram of the proposed methodology. Experimental data is first collected on an undamaged structure whilst environmental and operational conditions (EOCs) are varied to follow those that are likely to occur in its operation. Such data collection can be completed in a relatively short time by physically applying EOC variations that may occur more slowly in industrial operation, and setting the monitoring interval accordingly. The experimental data then serves as the undamaged monitoring signals that would be seen in the absence of damage growth.

The EOCs of interest are then specified and randomly selected from experimental records collected under such conditions. For each test scenario, the randomization process is repeated multiple times to generate multiple synthetic datasets in order to sample from the possible histories of EOCs. Artificial damage signals are then superposed on each of the undamaged datasets to create the synthetic datasets, with the damage following a specified growth pattern. These datasets then enable the mean and variability of the receiver operating characteristic to be determined.

The generated synthetic datasets are first compensated for any temperature variations, using a stretch-based temperature compensation algorithm [103], then processed to extract any damage features. Three damage feature extraction methods, namely, baseline subtraction residual, singular value decomposition (SVD), and independent component analysis (ICA), are implemented and compared in this study, as detailed in section 7.4. Synthetic datasets are created at different EOC and damage conditions, and the receiver operating characteristic (ROC) curves are computed for the different methods under different degrees of EOC variation and at different levels of damage growth.

Section 7.3.2 illustrates the process to predict damage signals and superpose them at the locations of interest. This superposition approach was validated in chapter 6.

7.3.2. Prediction of artificial damage signal and superposition onto undamaged signals

Unlike conventional one-off inspection, in SHM a large set of monitoring records is processed which where collected over time, and the progressive trend of damage growth is identified, even if the absolute damage amplitude is small. In our evaluation framework, synthetic datasets are created to mimic this process. In this chapter, growing corrosion is represented by the constant growth rate of the cross sectional area loss (CSAL) due to flat bottom holes of increasing cross-sectional area loss with a constant diameter-to-depth ratio of two. The CSAL is increased linearly from zero to simulate damage growth, and the diameters and depths of the flat-bottom holes are calculated accordingly during the damage progression.

The damage signature is predicted by first generating reflection coefficients of flat bottom holes of growing size on a homogeneous plate using the analytical model developed in [94], modified to compensate for the curvature of the pipe geometry [14]. Then at each time step during the damage progression, the reflection coefficient from the hole is multiplied by the Hanning-windowed inspection toneburst in the frequency domain; this frequency domain signal is then inverse Fourier transformed to obtain the reflection signals from the hole at that particular time. Note that although the CSAL increases linearly with time, the amplitude of the damage reflection increases non-linearly and is

Table 7.1.: Test parameters used to generate synthetic data

Test Parameter	Value	Unit
Temp. Range over which undamaged signals are selected	2, 10, 30, 60	^oC
Damage (CSAL) at the end of ramp	$0.3, 0.5, 0.7, 1.0, 3.0$	%
Number of readings over which damage grows	$10, 30, 50, 70, 100$	-
Damage location (away from features)	2.3, 7.5, 11.3	m
Damage location (near to a feature)	6.2, 10.1	

frequency dependent [94].

To simulate damage at a location of interest, the arrival time of the damage reflection is computed knowing the damage location and then the damage reflection is superposed onto the undamaged signals collected in the experimental setup in section 7.2. The undamaged signals are randomly selected from experimental signals collected over a range of EOCs, to mimic the monitoring records affected by temperature variations that could occur in the real world. The arrival time of damage is computed as a proportion of that of a significant structural feature, e.g. the second end reflection; this ensures that the damage reflection is inserted at the same spatial location in each record obtained at different temperatures. The damage reflections are then delayed and superposed on the undamaged experimental signals to create a synthetic dataset.

Table 7.1 summarizes the test parameters used to generate the synthetic datasets. Each synthetic dataset is generated with one value for each parameter. For example, a combination of 10^oC temperature variation, 1% damage amplitude, 50 readings, and 10.1 m means a synthetic dataset is generated that consists of 50 signals, where the undamaged signals are randomly selected from experimental records collected over a 10^oC range; the 50 damage signals linearly increasing from 0 to 1% CSAL are predicted and added to the undamaged experimental signals at 10.1 m, slightly after the second cut end reflection.

For each combination of test parameters, this synthetic data generation process is repeated 500 times. The resulting 500 synthetic datasets mimic the same damage progression but under different temperature histories, which enables us to statistically assess the damage detection performance in different practical settings. In total, about four million synthetic records are generated for 136 test scenarios with combinations of the test parameters shown in Table 7.1.

To ensure the independence of these synthetic datasets, the distance correlation coefficients are calculated [104] and the two-dimensional mutual information [105] between each pair of synthetic datasets generated for one test scenario. The distance correlation coefficient is similar to the Pearson correlation coefficient but is zero if and only if statistical independence is satisfied. The mutual information of two random variables X and Y measures the information they share, and can be expressed as $I(X;Y) = H(X,Y) - H(X|Y) - H(Y|X)$, where $H(X|Y)$ and $H(Y|X)$ are the conditional entropies, and H(X,Y) is the joint entropy of X and Y. Both the distance correlation coefficients and the ratio of mutual information over joint entropy ($I(X;Y)/H(X,Y)$) are scalars in (0, 1), where zero means they are independent, and unity means one completely determines the other.

Both measures show weak dependence between datasets when the number of measurements in a dataset is greater than 10. For example, for 500 datasets each containing 50 measurements randomly drawn from a 30^oC range, the 95% upper bound of the two independence measures were, respectively, 0.30 and 0.12. For all datasets having at least 30 measurements, the two independence measures are all below 0.4 and 0.14, respectively. When the number of measurements decreases to ten, the dependence between different datasets becomes much stronger, and the independence assumption no longer holds. Therefore, in the later sections, only the results using synthetic datasets with more than 10 measurements are analysed, which allows us to evaluate not only the mean but also the variance of the damage detection performance.

7.3.3. Receiver operating characteristic (ROC)

In statistics, ROC curves show the true positive rate (TPR) against the false positive rate (FPR), at various threshold values. In the field of NDE/SHM, ROC curves are also often used to illustrate the performance of damage detection, in which the two axes are termed the probability of detection (POD) and the probability of false alarm (PFA), both in the range (0, 1).

A representative baseline subtracted residual signal obtained as discussed in Section 7.4.1 below is used here to illustrate the generation of ROC curves. A random selection of 100 undamaged signals is made from the experimental data in section 7.2, covering 10^oC of temperature variation. Damage signatures of a linearly growing flat-bottom hole with maximum size of 1% are then predicted and added to the undamaged signals at 7.5 m, following the procedure of section 7.3.2. The temperature differences are first compensated

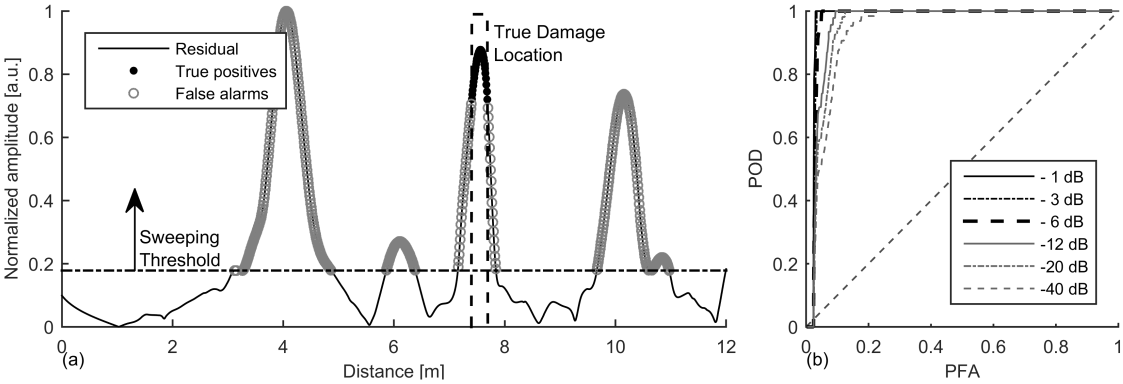

Figure 7.3.: (a) Example residual signal obtained from a set of 100 measurements under temperature variation of 10^oC, over which damage grows from zero to 1% of the cross-sectional area. (b) Receiver operating characteristic curves computed from the residual in (a) with different tolerance level.

using stretch-based scale-transform temperature compensation [103]; all 100 records are stretched or compressed to the baseline signal collected at the median temperature. The residual signals are then obtained using the procedure described in section 7.4.1.

Fig 7.3(a) plots the residual signal as a solid line, with its maximum amplitude normalized to unity; the dashed line shows the true damage location. Since the excitation signal used in guided wave monitoring is typically a Hanning-windowed toneburst, the reflection from features in the pipe are expected to extend over at least this length in time; therefore when the time axis of the received signal is converted to a length scale, the signals from features occur over a finite length of the pipe. Assuming a reflector has small axial extent compared to the wavelength, the reflected signal length will be approximately that of the input toneburst so it would be expected that a signal would be over this distance; the damage location on Fig. 7.3(a) is therefore shown as a finite length of pipe that is much larger than the true damage extent. To calculate the ROC, the threshold (dot-dashed line) is swept from 0 to 1, and all the values above the threshold are classified as positives. The positives that fall within the region expected for the true damage are labelled as true positives or detection (solid circles), and the positives elsewhere are labelled as false alarms (hollow circles). At every threshold, the probability of detection and false alarm rate are calculated, which becomes one point on the ROC curve. The collection of the points becomes the ROC curve for this test scenario.

In guided wave monitoring, damage can be masked by large residuals caused by EOC variations, even after compensation; this problem is most severe where there is a large signal in the undamaged structure, i.e. at structural features. In practical guided wave inspection, the inspector assesses the whole test length for defects and the probability of detection and false call rates are defined over the whole length. We are interested in monitoring the growth of damage over multiple readings, but the concept of assessing the whole pipe with one criterion is retained. This is a harsh measure for feature-free areas since many more false calls are generated at the locations of features; it would be possible to segment the pipe and compare performance in similar regions. However, since the purpose of this chapter is to set out the general methodology and to compare different feature extraction methods, this extra sophistication was not considered.

Since the excitation signal is a Hanning-windowed toneburst, the amplitude approaches zero at its beginning and end. If detection is required over the whole theoretical length shown as the damage location in Fig. 7.3 a high false call rate will be incurred at the tails. A defect is therefore expected to give a reflection above the 'call level' over the spatial region where the signal is above a specified fraction (tolerance level) of the peak signal. This is expressed as a fraction in dB so a -6dB tolerance level corresponds to requiring the defect signal to be above the threshold at all points either side of the peak where the envelope of the input toneburst is above half its peak amplitude. If this criterion is met then the POD is unity; a lower tolerance level means more points above the threshold are required for unity POD.

The ROC curve is an indicator of detection performance. The perfect detector yields a curve that goes through point (0, 1), indicating no false alarms and perfect detection (100% probability of detection). A random guess yields a ROC curve following the 45^o diagonal line, where the POD and PFA are equal. Fig 7.3(b) shows the ROC curves corresponding to different tolerance levels from -1 to -40dB. As the tolerance level goes down, the less tolerant the measure is to noise, and more false alarms will occur at any given POD level. Here a -6dB tolerance level is used (corresponding to the ROC curve with a thicker dashed line) to be consistent with common industrial practice.

7.4. Damage detection methodology for SHM

The conventional baseline subtracted residual method is first reviewed in section 7.4.1, and two advanced data-driven methods, singular value decomposition (SVD) and independent component analysis (ICA), in section 7.4.2. Section 7.4.3 discusses the use of the Mann-Kendall test to identify a monotonic trend in the data.

7.4.1. Temperature compensation and baseline subtraction

One common method to enhance signal-to-noise ratio (SNR) for damage detection is baseline subtraction. By subtracting a baseline signal that is collected in a known undamaged state, the residual signal should in principle resemble the signature created by the defect, if one exists. Signals are often pre-processed before the subtraction to eliminate any difference between the current measurement and the baseline due to temperature change. In this chapter, a stretch-based temperature compensation method [103] is implemented, as these techniques allow comparison of all the measured signals with a single reference signal.

The stretch-based temperature compensation methods approximate the temperature change as a stretch or compression in time, and reverse the effect by first finding the optimal stretching factor that makes the measurement and the baseline most alike, and then stretching or compressing one to the other. Although the stretch approximation is usually valid when the temperature difference is small, the stretching can introduce frequency distortion when the temperature difference is large [65]. To minimize the distortion, measurement are stretched to a baseline collected at the median temperature to minimize the maximum temperature difference to be compensated. This can be achieved even without temperature measurements, because the estimated stretching factors are effectively a measure of the temperature.

In our processing of a synthetic dataset, all the measurements are first compressed/stretched to the measurement at the median temperature in the particular dataset. Residual signals are then obtained by subtracting the first signal in the dataset from all the temperature compensated signals. Since the reflection from small defects roughly replicates the inspection toneburst, the obtained residual signals are cross-correlated with the eight cycle, Hanning-windowed toneburst to enhance the signal-to-noise ratio.

In conventional one-off inspection, the obtained residual signals are processed and assessed separately until a damage feature larger than a threshold is detected. However, in long term monitoring many records are obtained over which the damage grows progressively, which can be used to improve sensitivity. Specifically, the residual at the damage location would change monotonically, while values elsewhere would not due to random changes. If such a monotonic trend can be detected, we are more likely to identify the presence of damage at an earlier stage than in one-off inspection. In this chapter the case where damage grows linearly from the start of monitoring is considered. This is used to illustrate the method; it is not limited to this case.

The residual at location x can be written as a linear regression model of variable time t,

$$r_{xi} = \beta_{x0} + \beta_{x1} t_i + \epsilon_i \tag{7.1}$$

where β_{x1} is the rate of change at location x over time t, β_{x0} is the intercept corresponding to the residual at time zero, and ϵ is the error term which is assumed to follow the same distribution at all locations. The error consists of not only random system noise but also residual noise caused by imperfect compensation of the temperature differences, which follows a generalized Gaussian distribution if the monitoring period is sufficiently long and the temperatures at which readings are taken are randomly distributed.

The parameters β_{x0} and β_{x1} are estimated using least squares linear regression. At locations where damage grows progressively, the slope β_{x1} represents the growth rate of damage over time; elsewhere, no systematic trend should exist so β_{x1} is expected to be close to zero. On the other hand, β_{x0} might not be zero because of imperfect baseline removal, but this is not connected to the progression of damage. Therefore, the damage progression can be characterised solely using the slope β_{x1}, i.e., the values of β_{x1} are used at all locations along the pipe to represent the presence of damage and its progression if it is present.

The variance of the estimated damage growth rate β_{x1} is given by,

$$Var(\hat{\beta_{x1}}) = \frac{\sigma^2}{\sum_{i=1}^{N}(t_i - \bar{t})^2} \tag{7.2}$$

where the numerator is the variance of the noise term ϵ, and the denominator is the variance of the time at which all the measurements are taken. Therefore, there is greater confidence about the estimated slope if more observations are taken, or measurements are taken over longer period of time, making this suitable for processing long-term monitoring data.

7.4.2. Damage detection based on matrix decomposition

A recently developed group of methods is based on component analysis [66, 89], which decomposes the data matrix into different components to represent variations produced by different sources. The damage reflections and the signals caused by other EOC variations vary with position and have different varying patterns over time, and will be separated into different components. By identifying the components related to damage, variations from other sources are effectively eliminated and the damage detection performance is significantly improved.

The general form of decomposition of long-term monitoring records is

$$X = WA^T \tag{7.3}$$

where X, W, and A are matrices that contain the measurements over time, the weights, and the components, respectively. X is the $N \times D$ data matrix that samples the structural status both spatially and temporally, where N is the number of signals measured over the monitoring period, and D is the number of sample points in each signal. The two dimensions correspond to two different sampling frequencies – the frequency at which measurements are acquired from the monitoring system (temporal sampling interval Δt), and the sampling frequency of the A/D converter. The time axis of each measurement is converted to a distance scale by dividing the group velocity of the guided wave used by the sampling frequency to give a spatial sampling interval Δd. Then, each column of X represents one record collected at time t from $[0, \Delta t, \ldots, (N-1)\Delta t]$, and each row represents the wave reflected from a certain location d from $[0, \Delta d, \ldots, (D-1)\Delta d]$. The decomposition of X yields two matrices $W^{N \times R}$ and $A^{D \times R}$, each decomposing the variations along one dimension into R components, with $R \ll N, D$ representing the reduced dimensionality. Weight matrix W separates the variations over time, where each column $w^{N \times 1}$ represents the trend over time corresponding to one source. Component matrix A separates the variations in space, where each column $a^{R \times 1}$ represents the signal components along the distance axis.

The decomposition can be achieved in different ways. Two different decomposition methods, singular value decomposition (SVD), and independent component analysis (ICA)[90] were implemented. The two methods are widely used in fields such as face recognition [106] and speech separation [107], and their application in structural health monitoring has been illustrated in [66] and [89]. The major difference between the two

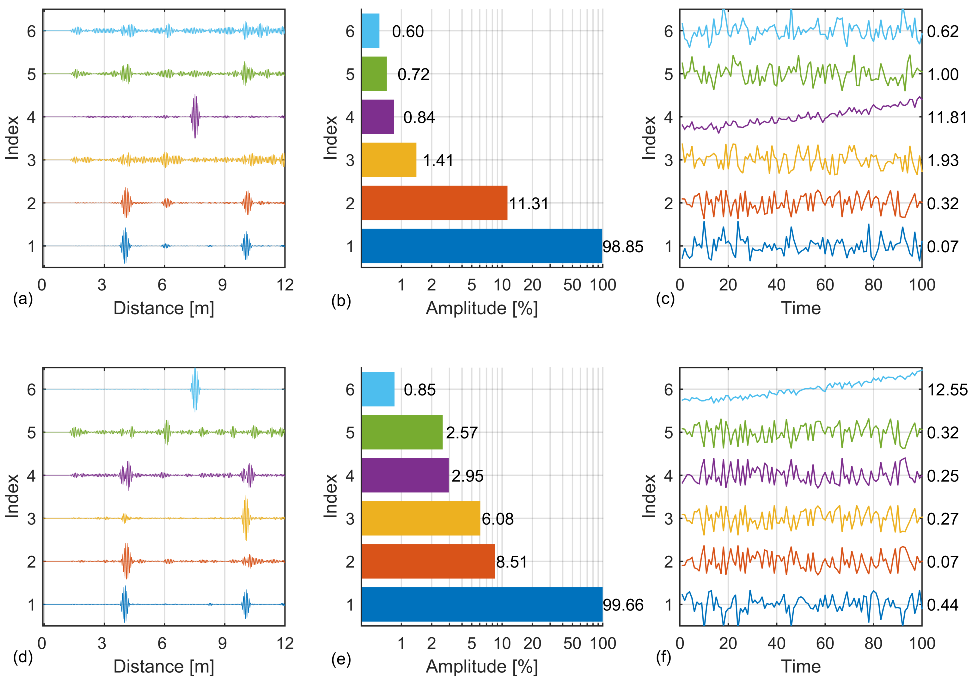

Figure 7.4.: Decomposition of 100 synthetic records collected in a 10^oC temperature range, over which damage grows linearly from 0 to 1% CSAL at 7.5 m. Six components from SVD (a) and ICA (d) are ordered by their absolute amplitudes shown in (b) and (e). The corresponding normalized weights over time (c) and (f) show the progressive trend over time. The MK test scores Z_{MK} (detailed in section 7.4.3) are labeled to the right of each set of weights.

methods is the criteria by which the components are separated: SVD enforces the orthonormality of both the component and the weight matrices, while ICA maximizes the statistical independence only between the components.

Fig 7.4 shows example results of applying SVD and ICA on the same synthetic dataset processed with the baseline subtraction residual method in section 7.3.3. The dataset contains 100 measurements with undamaged experimental signals over a temperature range of 10^oC, superposed with artificial damage progressing to 1% CSAL at its maximum extent, and was pre-processed with the same temperature compensation procedure as in section 7.3.3 and section 7.4.1. Both SVD and ICA decomposed the records into multiple components. Fig 7.4 shows the first six components and corresponding weights, as well as their absolute amplitudes at the end of the monitoring period.

Fig 7.4(a-c) shows the decomposition results using SVD, where (a) shows the first six normalized components plotted against distance from the sensor, ordered by their absolute amplitudes, which is shown in (b) on a logarithmic scale, and (c) shows their zero-centered and normalized weights over the 100 measurements. Fig 7.4(d-f) shows the corresponding results obtained using ICA.

The first components shown in Fig 7.4 (a) and (d) both resemble the undamaged A-scan of the pipe, with absolute amplitudes close to 100%. The corresponding weights in Fig 7.4 (c) and (f) vary randomly around the average amplitude. It should be stressed that the amplitudes of the weights shown in Fig 7.4(c) and (f) are normalized to their individual maxima so that their trends can be seen; their relative magnitudes can be seen in Fig 7.4(b) and (e). The fourth SVD component in (a) and sixth ICA component in (d) are clearly related to damage, each showing a dominant peak at around 7.5 m where damage is added. The corresponding final amplitude of the two components, as shown in Fig 7.4(b) and (e), are 0.85% and 0.84%, which are consistent with the predicted reflection coefficient from the 1% CSAL defect at the end of the damage growth. Moreover, their corresponding weights over the 100 measurements in Fig 7.4(c) and (f) clearly follow a monotonically increasing trend, with an increasing growth rate as CSAL progresses. This is consistent with our analytical model of a growing flat-bottom hole which predicts a non-linear relation between the reflection coefficient and the CSAL [14, 94].

Note that even with an increasing damage amplitude, the decomposed weights can decrease with time, in which case the phase of the corresponding component will shift by 180^o. This is because $w_i a_i^T$ is equivalent to $(-w_i)(-a_i^T)$, and both are valid results of the decomposition. Therefore, the damage related component can be identified by its monotonicity, not by its directionality. In section 7.4.3, the Mann-Kendall test is demonstrated, which is used to determine whether a monotonic trend (either increasing or decreasing) exists in the data.

The other components in Fig 7.4(a) and (d) capture features at different locations, and have various amplitudes. They represent other sources of variation that exist in the data that are either orthogonal to (in the case of SVD) or statistically independent of (in the cases of ICA) the undamaged experimental signal and the damage signature. Note that the amplitudes of some other components are larger than that of the damage related components. Nevertheless, the monotonic damage signatures are clearly distinguishable from the other components whose variation is either due to noise or related to the random temperature variation.

7.4.3. Identifying monotonic trend using Mann-Kendall trend test

The Mann-Kendall (MK) trend test [108, 109] is a hypothesis test to determine whether a monotonic trend exists in the data. A monotonic trend means that the variable consistently increases or decreases through time, but the trend may or may not be linear. The MK test is a simple, non-parametric test that only requires the data to be distributed similarly over time, and does not depend on the magnitude of data or its exact distribution. The test and its variant forms are extensively used in the fields of geological and environmental studies [110]. Some variants of the test also take seasonal effects into consideration [111], which could potentially benefit SHM applications since the EOC affecting applications are often seasonal or cyclic.

The MK test assesses whether to reject the null hypothesis (H_0: no monotonic trend) and accept the alternative hypothesis (H_a: monotonic trend is present), where

$$\begin{cases} H_0 : Prob[x_j > x_k] = 0.5 \\ H_a : Prob[x_j > x_k] \neq 0.5 \end{cases} , \forall j > k \tag{7.4}$$

To evaluate the test, $sign(x_j - x_k)$ is determined for all $n(n-1)/2$ pairs of j, k where $j > k$, and compute S by,

$$S = \sum_{k=1}^{n-1} \sum_{j=k+1}^{n} sign(x_j - x_k) \tag{7.5}$$

S essentially calculates the difference between the number of increases and the number of decreases. A positive S means the observations obtained later in time tend to be larger than observations made earlier. To determine whether the difference is statistically significant, the MK test statistic Z_{MK} is computed,

$$Z_{MK} = \frac{S - sign(S)}{\sqrt{VAR(S)}} \tag{7.6}$$

where the variation of S [112] is computed as,

$$Var(S) = \frac{n(n-1)(2n+5)}{18} \tag{7.7}$$

The p-value of the test statistic is then compared to the significance level α, which is the probability of wrongly rejecting the null hypothesis given that it is true (type I error). With $\alpha = 0.05$, the test is significant if $|Z_{MK}| \geq Z_{1-\alpha/2} = 1.96$, in which case the null hypothesis H_0 is rejected and it is concluded that there is a monotonic trend over time. Otherwise, the null hypothesis is accepted and it is concluded that no monotonic trend exists in the data. As an example, this methodology is used to test the weights obtained using SVD and ICA in section 7.4.2, each representing the trend of one component over the 100 measurements. The MK test statistics Z_{MK} are computed and labelled to the right of the weights in Fig 7.4(c) and (f). The Z_{MK} of the sixth component of ICA and the fourth component of SVD are, respectively, 11.81 and 12.55, corresponding to p-values close to zero. This suggests that the alternative hypothesis should be accepted and it is concluded that a monotonic trend exists in the data. Note that after decomposition, the monotonic trend is represented by a single component, while the other components, including but not limited to those shown in Fig 7.4(c) and (f), all have Z_{MK} that are smaller than $Z_{1-\alpha/2} = 1.96$ with a significance level of 0.05, marking the absence of a monotonic trend. This example shows that SVD and ICA clearly separate the damage related component from those corresponding to other variations, and that the damage related component can be identified using the MK test.

7.5. Comparison

7.5.1. Evaluation of different methods in a single scenario

Damage detection results are compared for residual processing, SVD, and ICA side-by-side on the same test scenario previously discussed in sections 7.3.3 and 7.4.2. Here, instead of processing a single dataset generated from that test scenario, the 500 synthetic datasets are processed with the same test parameters, where each dataset contains 100 undamaged measurements randomly selected over a temperature range of 10^oC, superposed with damage growing linearly from 0 to 1% CSAL. Fig 7.5 shows the extracted damage signatures, their amplitudes over time, and the ROC curves.

Fig 7.5(a) shows the baseline subtracted residuals, obtained using the procedure described in section 7.4.1 on the 500 datasets. The same datasets are also processed with SVD and ICA with the decomposition in section 7.4.2 and then the MK trend test in

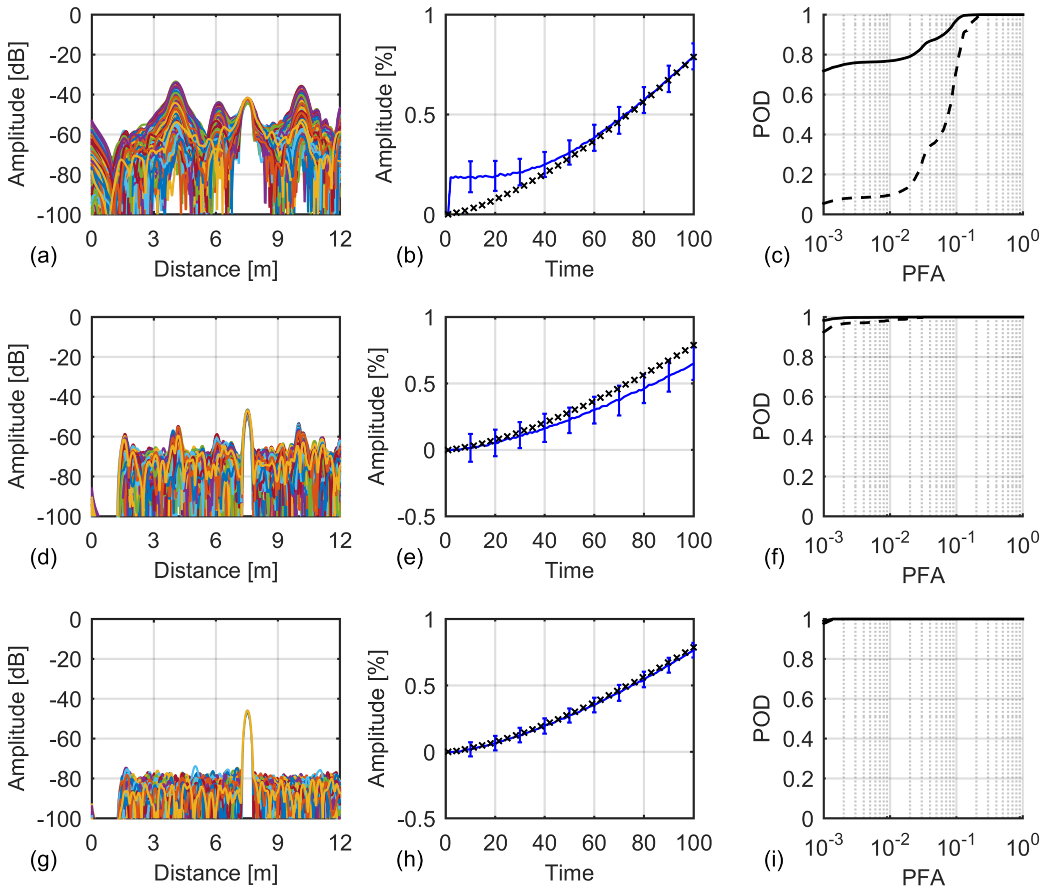

Figure 7.5.: Damage detection results using residual (a-c), SVD (d-f), and ICA (g-i) on 500 synthetic datasets containing 100 measurements collected over a 10^oC temperature range, while damage grew linearly from 0 to 1% CSAL at 7.5 m. (a) Baseline subtracted residuals. (d, g) Damage related component from SVD and ICA. (b) Amplitudes of baseline-subtracted residual at damage location over time (blue) and true damage amplitude (black). (e, h) Weights from SVD and ICA (blue), along with true damage amplitude (black). (c, f, and i) ROC curves (solid) and the 95% lower one-sided prediction bound (dashed).

section 7.4.3, and the resulting damage related components are shown in Fig 7.5(d) and (g).

Fig 7.5(a, d, and g) all show a clear peak at 7.5 m, where the damage is introduced. However, the damage signature in the residual signal in (a) is sometimes masked by the large residual noise, especially at structural features at 4 m and 10 m. The components obtained from SVD Fig 7.5(d) and ICA Fig 7.5(g) both clearly show damage features that are much larger than the ambient noise; the signal-to-noise ratio in the ICA damage component (Fig 7.5(g)) is on average 30dB higher than that of the residual (Fig 7.5(a)), and is roughly 10dB higher than that of SVD (Fig 7.5(d)), indicating that the performance is best with the ICA method.

Fig 7.5(b) shows the residual amplitude over time at the damage location. For the baseline subtraction method, the damage amplitude is usually represented by the maximum or the root mean square (RMS) of the residual signal; provided the baseline is perfectly removed, those measures will accurately track the damage growth. However, in the presence of EOC variations, these approaches are prone to errors resulting from imperfect baseline removal. Therefore, the residual amplitudes at each location are processed separately using the MK trend test, then the residual amplitudes at the location associated with the highest $Z_{MK}$are plotted as the representation of damage growth. The error bars show the standard deviations of the amplitudes. The true damage amplitude is also shown in (b), e) and h) as black crosses for comparison. The residual amplitude is initially masked by the noise floor but approximates the true damage amplitude (black crosses) as damage grows and produces a reflection larger than the noise. Note the error bars in Fig 7.5(b, e, and h) are the results of processing the whole 100 measurements, not of processing the observations up to that particular time step. Therefore, it shows how the history of damage growth would be seen after the 100 measurements over the growth to 1% CSAL. The detectability of different levels of damage and the effect of changing the number of readings are considered in section 7.5.3. Fig 7.5(e) and (h) show the weights obtained from SVD and ICA, respectively, both of which show clear monotonic increasing trends with time. The weights of SVD in Fig 7.5(e) follow a similar growth pattern as the true damage progression, but the growth rate is underestimated by about 15%. The weights of ICA in Fig 7.5(h) closely follow the true damage amplitudes, and show the smallest variation of the three, providing better tracking of the damage growth.

We then obtain the ROC curves by processing the extracted damage signatures with the procedure described in section 7.3.3. Fig 7.5(c), (f), and (i) show the ROC curves for, respectively, residual processing, SVD, and ICA. The ROC curves of the 500 synthetic datasets are 'vertically' averaged (giving the averaged POD at a given PFA)[100] and plotted as a solid line, while the dashed line shows the lower one-sided 95% prediction bound of POD, indicating that with a 95% probability future observations will fall above this value. The ROC curves demonstrate the performance in the case of linear damage growth (giving non linear growth in reflection amplitude) to a maximum 1% CSAL with 100 readings over a 10^oC temperature range. The baseline-subtracted residual successfully indicates the presence of damage most of the time, with some false alarms caused by the residual noise that often masks the damage signature. The wide gap between the average performance and the lower 95% prediction bound suggests that it is also less robust to the temperature variations. As an example, along the dashed curve in Fig 7.5(c) that represents the 95% one-sided lower prediction bound, if a POD of 90% is required, then the corresponding PFA is 13%. In contrast, with SVD and ICA a lower 95% prediction bound POD of 99% is obtained at PFA of 1.8% and 0.12%, respectively.

7.5.2. Comparison of ROC curves across different scenarios

Section 7.5.1 demonstrates the operation of our proposed evaluation framework when comparing the three different damage detection methods in a single damage growth - EOC variation case (1% linear CSAL growth and 10^oC temperature range). However, ROC curves are a two-dimensional depiction of classifier performance, and are difficult to compare across multiple scenarios. Therefore, the area under the ROC curve (AUC)[113] is used as a scalar metric to facilitate a comparison of the performance of the different methods across multiple scenarios. AUC is a single scalar metric whose value lies between 0 and 1, representing the proportion of the area of the unit square. An AUC of unity corresponds to a perfect detector and 0.5 corresponds to a random guess.

Here, the AUC is demonstrated on the results of residual processing on four synthetic datasets generated with different damage severities, including the one processed in section 7.5.1. For each of the four scenarios, 500 synthetic datasets were generated, each containing 100 randomly selected experimental signals collected over a 10^oC temperature range. Fig 7.6 shows the four cases with different colours where damage grows linearly from zero to, 0.3%, 0.5%, 0.7%, and 1.0%.

Fig 7.6(a) shows the mean ROC curves on a linear PFA scale as solid lines and their 95% lower one-sided prediction bounds as dashed lines. Unsurprisingly, the larger the damage growth rate is, the more the ROC curve approaches the top left hand corner of the plot, and the smaller is the variance. Fig 7.6(b) shows four box plots of AUCs corresponding to the ROC curves in Fig 7.6(a), where the box, the horizontal line, and the vertical whisker indicate, respectively, the interquartile, the median, and the 90% two-sided prediction interval of the 500 AUCs. Two-sided prediction interval is used here such that the lower bound, which is of most interest in the evaluation of the performance, is consistent with the one-sided 95% bound shown in Fig 7.6(a). The AUC box plots clearly show the performance improves as damage severity increases, with the boxes moving up the AUC axis, corresponding to higher POD and lower PFA, and their smaller extent

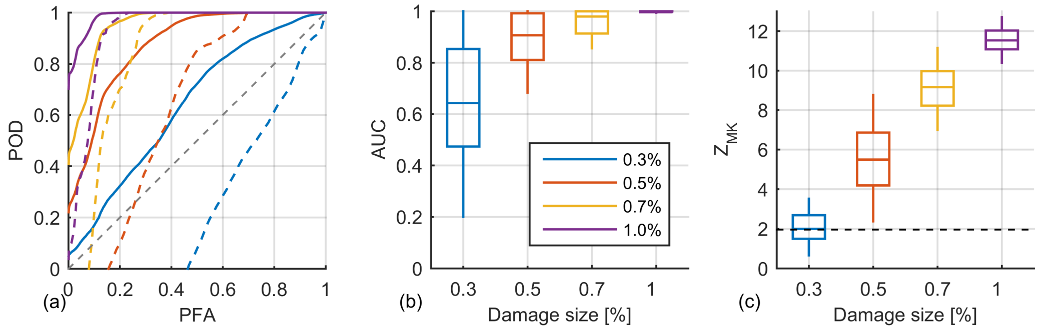

Figure 7.6.: Comparison of damage detection using baseline subtracted residual across different scenarios; different colours represent synthetic datasets generated with different damage growth rates over 100 records over a 10^oC temperature range. The purple curves are results of the same case shown in Fig 7.5(c). (a) ROC curves (solid) and lower one-sided 95% prediction bound (dashed), (b) Box plots of area under curve with boxes indicating interquartile and whiskers indicating 90% two-sided prediction interval, and (c) Box plots of Z_{MK} of the weights.

along the AUC axis corresponding to smaller variance. Note that in the case of 0.3% damage growth (blue), some of the AUCs drop below 0.5, corresponding to ROC curves below the diagonal line, indicting that the performance is worse than a random guess. This happens when the extracted damage feature is similar in amplitude to the ambient noise and is much smaller than the residual noise at structural features.

Fig 7.6(c) shows the MK test statistic Z_{MK} indicating the presence of a monotonic trend. The indications from Z_{MK} are consistent with those from the ROC curves and AUCs. As the damage growth over the 100 readings increases, the Z_{MK} test statistic increases, and at damage of 0.5% or higher CSAL after 100 readings, in over 95% of the 500 cases Z_{MK} exceeds the 1.96 threshold associated with a significance level of 0.05, as marked as a dashed line in Fig 7.6(c). Since the trend of Z_{MK} is consistent with that of AUC, in the next section, the AUC is used to summarize the results and compare performance across different scenarios.

7.5.3. Evaluation of different methods across different scenarios

The analysis is now extended to compare the three methods across scenarios with different combinations of test parameters from Table 7.1. Fig 7.7 shows AUC box plots for each test scenario, where the three colors correspond to the results of residual processing

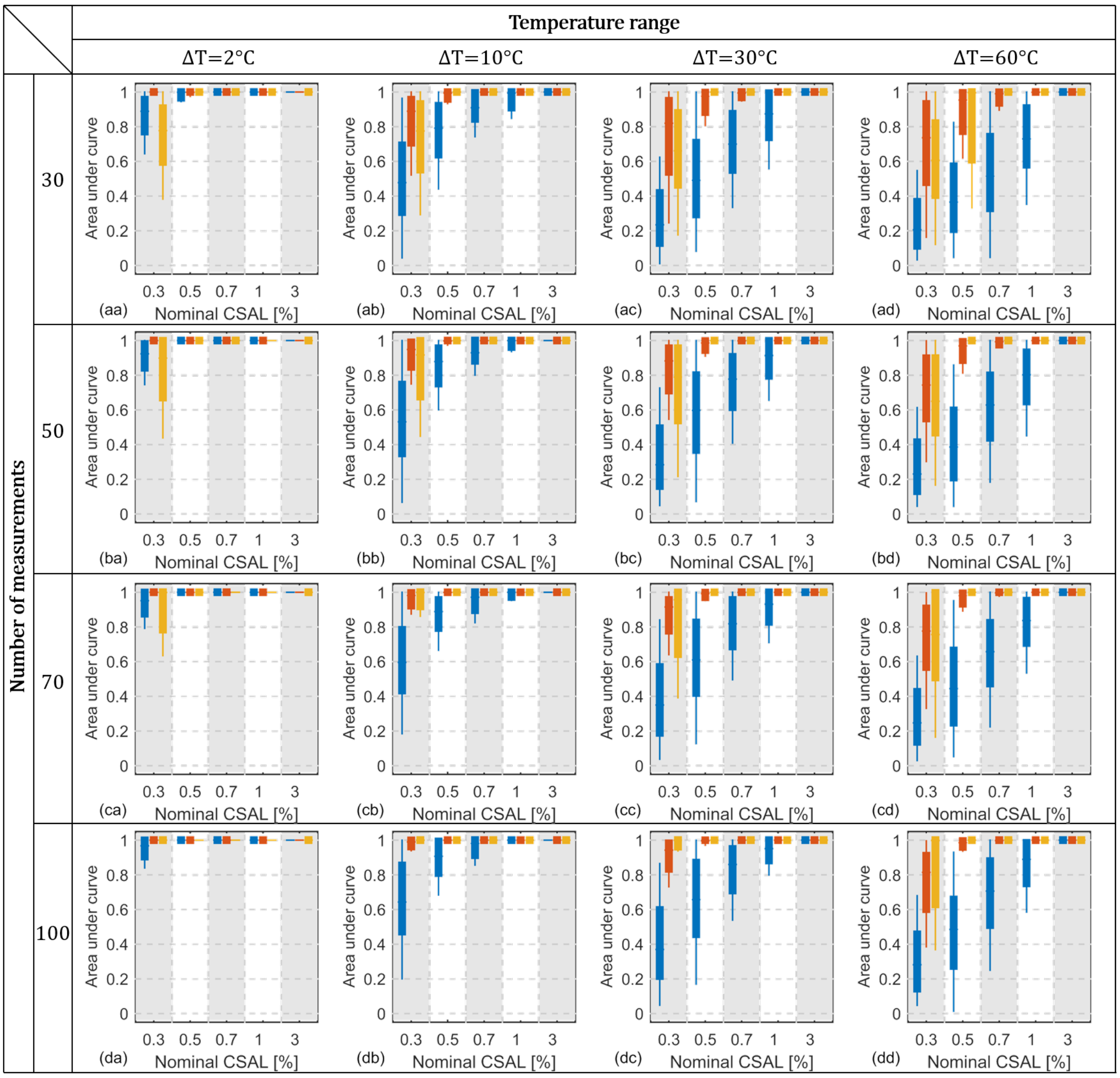

Figure 7.7.: Box plots of AUCs on synthetic datasets generated with different test parameters. Rows: different numbers of readings. Columns: different temperature ranges. Horizontal axis: nominal CSAL at the end of the linear damage growth. The blue, red, and yellow boxes in each stripe represent results of, respectively, residual processing, SVD, and ICA.

(blue), SVD (red), and ICA (yellow). The four rows correspond to different numbers of readings taken during the damage growth, and the four columns correspond to different temperature ranges. Within each plot, the horizontal axis indicates the nominal CSAL at the end of the linear damage growth, where the stripes mark the results of the three methods on the same test scenario.

Unsurprisingly, with all other parameters equal, the methods achieve better performance (higher AUCs) when damage is larger, or when temperature variation is smaller.

The performance also improves with increasing number of measurements taken during the damage growth. Comparing the three methods, SVD and ICA produce much better results than the conventional residual processing in all test scenarios, especially when the temperature variation is large. ICA generally outperforms SVD except in cases where damage is small (0.3% CSAL).

Fig 7.7 can then be used to assess the damage detection performance as a function of any of the three varying parameters, while keeping the other two unchanged. For example, Fig 7.7(cc) shows the performance of detecting different levels of damage growth across 70 measurements under $30^o C$ temperature variations. Assuming a detector is satisfactory if its AUC exceeds 0.95, which roughly corresponds to 90% POD with 2% PFA (this level of PFA may not be practically acceptable, but is used here only to compare the methods), then at a 95% prediction level SVD and ICA can be used to detect the damage before it grows to 0.5% CSAL. On the other hand, although the residual performance improves as damage severity increases, it does not perform sufficiently well until damage grows to 3%. Similarly, comparing Fig 7.7 (ad, bd, cd, and dd), it is seen that detecting 0.5% CSAL growth with a 95% prediction bound AUC of 0.9 requires at least 70 measurements for SVD, or at least 50 for ICA.

To successfully implement these methods, it is critically important and of most interest to estimate the lower prediction limit of the damage detection performance in the practical settings of the SHM system, which bounds the worst-case performance of a future test. The results of Fig 7.7 show that the SHM system performance improves with damage size and degrades with degree of temperature variation. It is therefore interesting to plot performance versus the signal-to-noise ratio (SNR), which is defined as the reflection coefficient of the damage at the end of the measurement period divided by the maximum amplitude of the coherent noise resulting from the temperature range. The SNR indicates how difficult it is to detect damage in this scenario – a small SNR corresponds to the growth of small damage over a large temperature range, and a large ratio corresponds to the growth of large damage over a small temperature range. Fig 7.8 plots the 95% one-sided lower prediction bound of AUC against the SNR. The three colors correspond to the three damage detection methods, and solid and dashed lines represent the results from synthetic datasets of, respectively, 50 and 100 measurements. Note that different combinations of damage amplitude and temperature range may result in similar SNR. For example, three test cases with SNR of around -20dB correspond to, respectively, 0.3% CSAL over $10^o C$ ($SNR = -18.8$dB), 0.7% CSAL over $30^o C$ ($SNR = -19.5$dB), and 1.0% CSAL over $60^o C$ ($SNR = -20.1$dB).

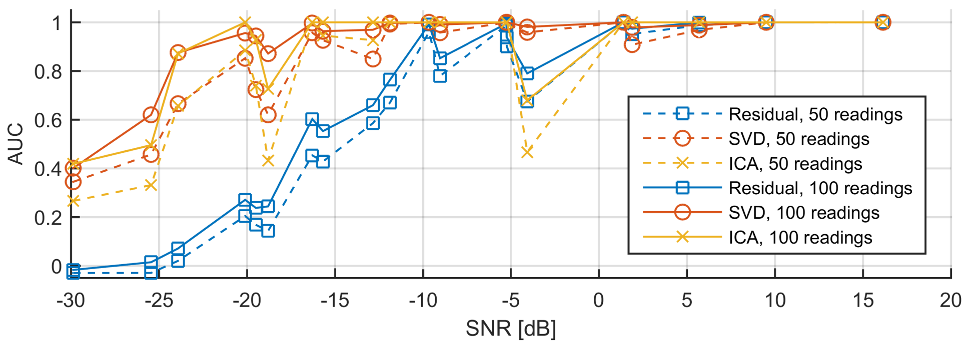

Figure 7.8.: The lower 95% prediction bound of AUC of three methods against the ratio of damage amplitude over temperature noise. Dashed lines and solid lines are generated with synthetic datasets containing, respectively, 50 readings and 100 readings.

The damage detection applications can be separated into three regimes by the SNR: with SNR larger than 0dB, generally good detection results are obtained with any of the methods; with SNR in between 0 and -18 dB, sophisticated processing such as SVD and ICA gives detection gains; with SNR smaller than -18dB, satisfactory results cannot be obtained with any of the methods.

The anomalies in Fig 7.8 at -4.1dB and -18.8dB correspond to a damage level of 0.3% over different temperature ranges, which suggests that there is a small source of coherent noise that is not suppressed by temperature compensation and this affects the detectability of small defects. This requires further investigation, but it should be stressed that 0.3% cross section loss is more than an order of magnitude smaller than the wall loss typically detected in guided wave inspection [17]. Comparing SVD (red circles) and ICA (yellow crosses), SVD seems to be more robust to this effect.

The dash lines and the solid lines represent results with, respectively, 50 and 100 readings over the damage growth. For a year of monitoring, 50 readings roughly corresponds to observations made weekly, while 100 readings corresponds to 2 measurements per week. Comparing the two sets of results it can be concluded that more robust damage detection can be achieved with more measurements.

7.5.4. Evaluation of different methods with the damage near to structural features

In previous sections, damage detection performance of the three methods was evaluated and compared on different scenarios where damage was added at 7.5m, in a plain section of the pipe. Here further synthetic datasets are processed with damage added at other locations including ones close to structural features.

Fig 7.9 shows four pairs of AUC (a-d) and Z_{MK} (e-h) of the three methods detecting 1% damage growth over 100 measurements. The four stripes in each plot mark temperature ranges of 2, 10, 30, and 60^oC, respectively. Fig 7.9(a, e) and (b, f) show results when damage is added at, respectively, 2.3 m and 11.3 m, both in plain sections of the pipe. As with the results shown in Fig 7.7(da-dd) where damage was added at 7.5 m, here the SVD and the ICA both yield AUCs close to unity, indicating successful detection of damage with few false alarms. The residual processing also achieves AUCs close to unity when the temperature range is small (2 or 10^oC), but its performance quickly degrades with increasing temperature range. The corresponding Z_{MK} displays the same trend as the AUC, as in Fig 7.6.

Fig 7.9(c, g) and (d, h) show the cases where damage is added at locations close to structural features. In these cases, the AUC in itself cannot represent the detection performance, and needs to be interpreted together with the Z_{MK}. The two cases display similar behaviour, so only Fig 7.9(d, h) is discussed for conciseness. Fig 7.9(d, h) corresponds to the cases where damage is added at 10.1 m, slightly after the second end reflection. Recall that reverberating signals are used to represent a longer pipe in which the end reflections would be large features. In all four cases over different temperature ranges in Fig 7.9(d), the AUCs of SVD and ICA are close to unity, and the AUCs of residual processing are also over 0.9. The AUC with the residual method appears high because features tend to generate relatively large residuals; at a feature the residual is the difference between two large reflections in the baseline and current, temperature compensated signal, whereas away from features the signals are smaller so it is easier to obtain a small residual. In contrast, the Z_{MK} in Fig 7.9(h) suggests that when the temperature range exceeds 10^oC, the residual processing fails the MK trend test. In this case, no damage is detected along the pipe at any location, and the ROC should not be calculated. This interpretation should also be extended to SVD and ICA, where the MK test is used to automatically determine the damage-related component: when the Z_{MK} of all

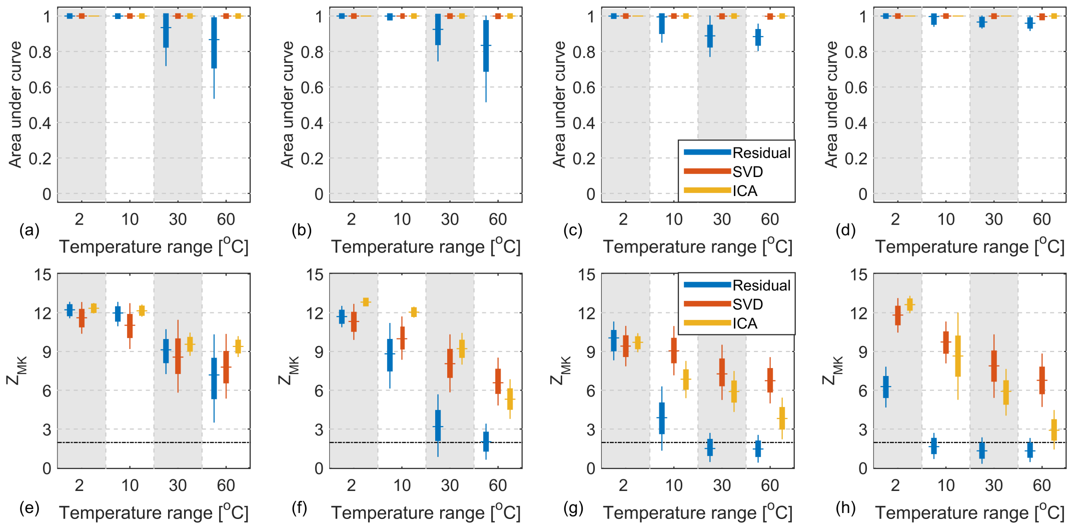

Figure 7.9.: Box plots of AUC and Z_{MK} of detecting 1% damage growth over 100 measurements over different temperature ranges, where damage is added at four different locations in Table 7.1. Plots (a, e) and (b, f) show results when damage is added at, respectively, 2.3 m and 11.3 m, away from structural features. Plots (c, g) and (d, h) show results when damage is added at, respectively, 6.2 m or 10.1 m, close to structural features. The dot-dash line in (e - h) shows the threshold of 1.96.

the components are below the threshold, it should be concluded that no single component represents the damage, and that damage is not present in the dataset being processed.

As an example, Fig 7.10 plots results of the three methods on synthetic datasets corresponding to the third stripes in Fig 7.9(d and h). Each synthetic dataset contains 100 measurements collected over a 30^oC range, while damage grew linearly from 0 to 1% CSAL at 10.1 m (slightly after second end reflection). The residuals in Fig 7.10(a) are dominated by the large residual errors at structural features that mask the damage signature. The weights shown in Fig 7.10(b) display a 15% standard deviation, and no monotonic trend can be identified by the MK test. Fig 7.10(c) shows the 95% lower prediction bound of ROC curve with large PFAs even for small PODs.

Fig 7.10(d, e, f) show the results of SVD, which successfully extracts the damage signature and suppresses the benign features. Although the components shown in Fig 7.10(d) are still noisy, the damage signature at 10.1 m can still be identified, which is at least 6dB larger than the noise elsewhere. Correspondingly, the ROC curve in Fig 7.10(f) shows that a 95% lower prediction bound POD of 90% corresponds to a PFA of 3.3%. The weights in Fig 7.10(e) significantly underestimate the damage growth rate, but its trend

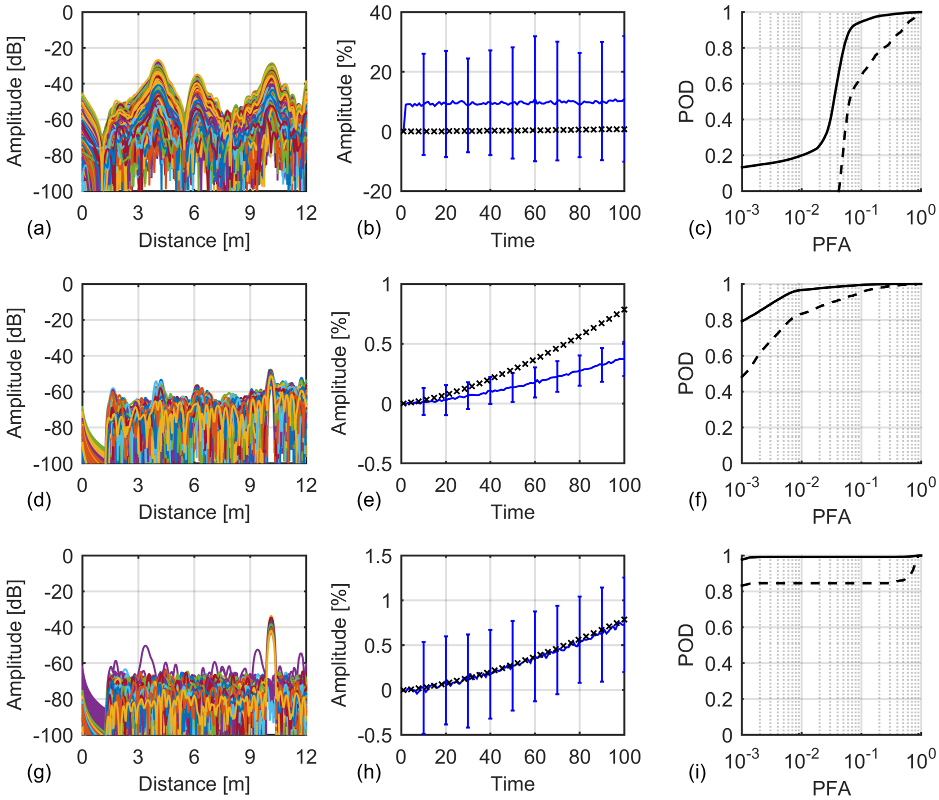

Figure 7.10.: Damage detection results using residual (a-c), SVD (d-f), and ICA (g-i) on synthetic datasets each containing 100 measurements collected over a 30^oC temperature range, while damage grew linearly from 0 to 1% CSAL at 10.1m (slightly after second end reflection). (a) Baseline subtracted residuals. (d, g) Damage related component from SVD and ICA. (b) Amplitudes of baseline-subtracted residual at damage location over time (blue) and true damage amplitude (black). (e, h) Weights from SVD and ICA (blue), along with true damage amplitude (black). (c, f, and i) ROC curves calculated from the damage signatures in (a, d, and g) and the 95% lower one-sided prediction bound.

is qualitatively consistent with the true damage growth.

The ICA components in Fig 7.10(g) show very clear damage signatures in most cases. However, on four of the 500 datasets (eg. the purple line in Fig 7.10(g) with peaks at 3.5 m and 9.5 m), ICA fails to extract the damage component and yields components corresponding to noise from other random variations. These outlier cases produce the large variations in the weights as shown in Fig 7.10(h), and the large drop in the 95% prediction bound of POD in Fig 7.10(i) from its mean. When the temperature range increases from 30^oC to 60^oC, the proportion of the outlier cases increases from 1% to 20%, producing a significant degradation of Z_{MK} as can be seen in Fig 7.9(h). As a consequence, better results were obtained using SVD than using ICA when damage was added close to a structural feature under large temperature variations. A probable cause of the outliers is that ICA is essentially an optimization process to find a set of components that minimize the mutual information, and is therefore prone to local minima. As the temperature range increases, the damage signature becomes a smaller portion of all the variations in the data, so the optimization is more likely to get stuck in local minima.

7.6. Experimental validation

The evaluation framework is now validated on experimental data monitoring actual damage growth. A flat-bottom hole was drilled on the same pipe specimen shown in Fig 7.1(a) at 2.3 m to the right of the PIMS transducer ring. The hole was 4 mm deep and was extended from 1 mm to 7 mm diameter in 1 mm steps, simulating progressively growing damage from 0 to roughly 0.5% CSAL. At each damage step, the pipe was heated to 90^oC and naturally cooled to 30^oC, and data was collected over the temperature cycle as described in section 7.2. Data from the last temperature cycle prior to introducing the hole, together with that from the cycles after each damage set was used, giving a total of 338 measurements over the damage growth. This gave a reasonable approximation to linear damage growth from 0 to 0.5% CSAL.

As with the synthetic datasets, the performance of damage detection is assessed over different monitoring records under different levels of temperature variation. For temperature ranges from 10 to 60^oC in 10^oC steps, n signals are randomly selected from each of the 8 cycles, forming datasets of $8n$ monitoring records over the 0.5% damage growth, where n =1, 4, 6, and 9 correspond to 8, 32, 48, and 72 records, respectively. Note that here the chronological order of the selected measurements from different damage steps is retained, but the measurements from the same damage step are randomly ordered, such that the temperature associated with the measurements varies randomly. This random selection process was repeated 50 times for each combination of test parameters (temperature range and number of records) to generate statistics.

Fig 7.11 shows the results of residual processing (blue), SVD (red) and ICA (yellow) on different numbers of experimental records over different temperature ranges. Fig 7.11(a) to (d) correspond to the results using, respectively, 8, 32, 48, and 72 measurements, while the horizontal axis in each plot indicates the temperature range over which those measurements are selected.

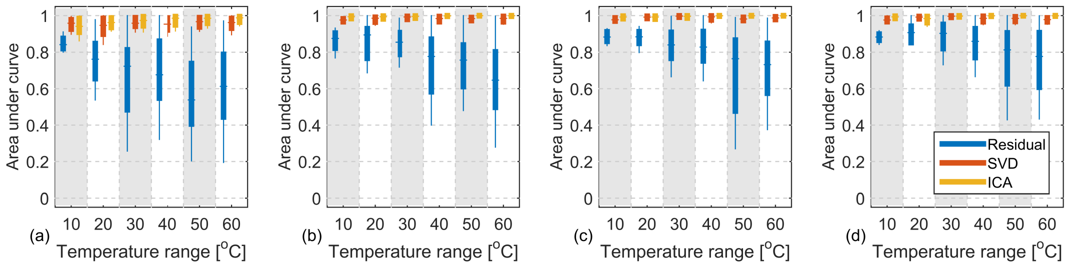

Figure 7.11.: Results of baseline-subtracted residual (blue), SVD (red), and ICA (yellow) on experimental records. Plots (a) to (d) correspond to the results of using, respectively, 8, 32, 48, and 72 measurements. The horizontal axis indicates the temperature range over which measurements are selected.

The baseline-subtracted residual performs poorly when detecting the 0.5% damage, with AUCs mostly lower than 0.9, and shows only small improvement when more experimental records are processed. In comparison, both SVD and ICA perform reasonably well using 32 records, and better with more records; the performance generally degrades as the temperature range increases. Overall, ICA outperforms SVD, especially when many monitoring records are available.

Comparing Fig 7.11 with the results using synthetic data in Fig 7.7, the performance is qualitatively consistent. As with the synthetic data, the component methods perform significantly better than the residual method, with ICA generally the better of the two, indicating that the proposed framework of using experimental data on an undamaged system coupled with predicted damage signals is effective. The performance on experimental data is slightly better than previous results with synthetic datasets as shown in section 7.5.3, especially at large temperature ranges. This is probably partly because the data was only processed up to and including the first end reflection, as would be done in practice; this reduces the number of large, benign reflectors in the signal, and hence the likelihood of false calls.

7.7. Conclusion

In this chapter an evaluation framework has been proposed to assess the damage detection performance of guided wave structural health monitoring systems under varying environmental and operational conditions (EOC). By synthetically superposing damage reflection signals from a growing defect onto experimental signals from an undamaged structure measured under practical environmental variations, the performance of an SHM system can be assessed under various scenarios of interest. This enables the effect of temperature range over which measurements are taken, damage location, damage amplitude, the number of observations over the monitoring period, etc. to be studied. Moreover, by generating multiple datasets with randomly selected experimental records, the performance can be statistically assessed in these practical settings. The framework has been demonstrated using undamaged data taken on an 8 inch Schedule 40 pipe in the laboratory under up to 60^oC temperature variations.

The performance of three damage detection methods has then been assessed using receiver operating characteristic (ROC) curves, which plot the probability of detection (POD) against the probability of false alarm (PFA). The three damage detection detection methods, conventional baseline-subtracted residual processing, singular value decomposition (SVD), and independent component analysis (ICA), were first reviewed. For the SVD and ICA techniques a new method was proposed for identifying the damage component, based on trend analysis. The three techniques were then compared across multiple defect detection scenarios. Under all investigated circumstances, the component methods performed significantly better than the residual method, with ICA generally better than SVD except for a small number of outliers. For all three methods, the performance improves with larger damage size and more observations, and degrades with larger range of temperature variations. The results were then validated using experimental data monitoring a flat-bottom hole with growing diameter, the results being consistent with the synthetic data.

The results show that we can use undamaged data from any installed monitoring system to predict its damage detection performance. This makes it possible, for example, to assess whether a defect of a given size at a particular location would be detected reliably at a given data collection frequency. This enables the system to be tuned to meet particular requirements and the performance demonstration will be valuable in building safety cases.

Chapter 8:

Conclusions and future work

8.1. Thesis review

This thesis began in chapter 1 with a review of guided wave inspection and its development into a commercially available inspection technology. Some of the challenges associated with this development were also discussed. This chapter then went on to review the trend in industry towards operating sensors in a monitoring configuration rather than an inspection configuration. Some of the benefits of this shift were discussed, as well as some of the current challenges with extracting the maximum use from this monitoring approach. The chapter then ended with an outline of the thesis and the major contributions made.

The remainder of the thesis could conceptually be broken into two parts. The first half of this thesis dealt with finite element modelling of guided waves and contributions made by the author in this area. Chapter 2 began with a brief overview of existing finite element modelling techniques for guided waves and recent advances in solution procedures based on graphics processing units (GPUs). These new GPU based solvers enable larger, more complex models to be solved in a shorter space of time, which causes new challenges in terms of construction of models and post-processing of simulation data. Although the tools described in this chapter were not themselves novel, this was the first known instance of their combined use.

Chapter 3 made use of the new modelling capability outlined in chapter 2 to perform a study of the interaction of guided waves with rough surfaces in pipes. This chapter investigated how the scattering of guided waves by these rough surfaces was a function of several variables of the guided wave inspection, such as the inspection frequency, the mean depth and correlation length of the rough surface, and the diameter of the pipe. The results help to explain the apparent attenuation seen during commercial inspection of generally corroded pipework, an attenuation which has been unexplained up to now.

The second half of the thesis began with chapter 4 which considered the problems associated with the processing of guided wave monitoring data. In particular the chapter focused on the problem of 'data normalisation', the process of compensating monitoring data for changes in environmental and operating conditions. Some existing compensation procedures were considered, as well as the benefits and limitations of these different techniques. This analysis concluded that there is a need for improved processing techniques, as well as a rigorous framework in which to assess different approaches.

Chapter 5 proposed a new technique for processing guided wave monitoring data which was based on combined application of the novel technique of Independent Component Analysis with the established technique of baseline stretch. This new processing technique was demonstrated on example guided wave monitoring data, including example monitoring data from a rough surface.

Chapter 6 described a validation of the superposition approach, an approach which can be used to generate large sets of guided wave monitoring data with reduced computation and experimental effort. A comparison was made between superposition data and the equivalent data from a full finite element simulation, with favourable results. The limits of this validation were also discussed.

Chapter 7 then proposed a new framework for comparing competing guided wave signal processing algorithms. This new framework relied on the superposition approach to generate large sets of test data, which in turn allowed descriptive statistics. This new framework was then used to compare the new technique described in chapter 5 to the existing best practice described in chapter 4.

8.2. Main contributions

8.2.1. The interaction of guided waves with rough surfaces

The first main contribution of this thesis is in improved modelling capability for guided waves. This thesis makes use of graphics processing unit (GPU) based solvers to model guided waves on large length scales. This makes it possible to model large area pipe features such as rough surfaces, models which would not have been possible with previous techniques. Large models do however cause challenges in terms of the discretization of the model volume, as models often involve millions of elements which cannot be verified through visual inspection. This thesis outlines tools for efficient meshing, mesh smoothing and quantification of element quality which enable these large models.

This modelling capability made it possible to study the interaction of guided waves with rough surfaces, an interaction which to date had received little consideration in the literature. Such a study is only made possible through the use of best practice modelling techniques and gives new insight into the scattering mechanisms for guided waves and rough surfaces. It was shown that rough surfaces scatter guided waves, which explains the attenuation seen during commercial inspection of corroded pipes. This scattering was found to be dependent on several variables of the inspection set-up; inspection wave frequency, pipe diameter and rough surface characteristics. In general it was found that the scattering increases as the inspection frequency increases, since at higher frequencies more modes can exist and hence there are more modes for the energy to scatter into. However, this trend breaks down close to the cut-off frequencies of the higher order flexural modes where the mode conversion is particularly strong; at their cut-off frequency, flexural modes form standing waves around the pipe circumference and this resonance condition encourages mode conversion. The mean depth and correlation length of the surface also has an influence on mode conversion, with attenuation greatest for surfaces where the correlation length over wavelength is in the region 0.2-0.25; the attenuation approaches zero at low and high values of correlation length relative to the wavelength. The deeper the corrosion the more severe the mode conversion, since deeper corrosion will lead to a reduced remaining wall thickness and hence a greater impedance mismatch with the plain pipe. It was also found that mode conversion is a strong function of pipe diameter; for pipes with the same wall thickness and rough surface statistics the attenuation increases significantly with frequency-diameter product. This is because larger diameter pipes can support more modes at a given frequency, so there are more possible modes for the energy to scatter into. Away from cut-off frequencies the mode conversion is mostly to the F(*,2) modes since these modes have the greatest circumferential displacement for unit power flow; the T(0,1) mode therefore couples most easily into these modes.

8.2.2. Independent component analysis applied to monitoring data

This thesis has introduced a novel technique for processing guided wave monitoring data based on independent component analysis. In chapter 5 this new technique was applied to sets of data which approximate the data collected from guided wave monitoring systems. This data included experimental data from both a plain section of pipe and an industrial pipe loop containing welds, bends and flanges. Data was collected from these pipes whilst they were in their undamaged and damaged condition and undergoing temperature variations. Initially model defect signals were added to the data from the undamaged plain pipe to simulate step changes and gradual defect growth. The ICA algorithm was able to separate these defect signals from coherent noise, even in the context of significant changes in temperature and changing frequency content of the defect signal. The ICA algorithm was then applied to data from the undamaged and damaged pipes and again the algorithm was able to separate the defect signal from coherent noise. This was for physically introduced defects in clean sections of pipe, at welds and at welds beyond bends. This chapter also used finite element simulations to investigate generally corroded sections of pipe with a patch of severe corrosion. It was found that the ICA algorithm correctly separated the reflection due to the severe corrosion and recognised it as a single evolving defect. This is in the context of the corrosion undergoing a complex growth pattern where its shape and amplitude is changing.

8.2.3. Comparing guided wave monitoring algorithms

In chapter 6 a new approach for generating large sets of guided wave monitoring data was presented. This chapter formalised the superposition approach and sought to compare results obtained with superposition to results from full finite element simulations. This comparison has been performed for three representative inspection scenarios including: simple flat-bottom hole defects in plain pipe and at welds, with and without temperature variations and with representative subsequent signal processing. This is in the context of background coherent noise that produces a low level noise throughout the signal. It was found that superposition worked well and gave very similar results compared with the full modelling approach. As expected, there was some difference between the two approaches, which can be explained by the fact that the reflection from a hole is not totally independent of the cross section into which it is being added. In general is was found that superposition slightly underestimates the defect amplitude for small defects and overestimates defect amplitude for large defects. It was also discussed that the results of this validation are applicable only to small defects in relatively feature-free pipes.

The validated superposition approach from chapter 6 was then used in a new framework for comparing monitoring algorithms, outlined in chapter 7. This framework is based on the generation of large sets of data which can then be processed using competing signal processing algorithms. The performance of these different algorithms is then assessed using receiver operator characteristics, a standard technique for assessing binary classifiers. This framework was then used to compare three common signal processing techniques for guided wave monitoring data: residual processing, singular value decomposition (SVD) and independent component analysis (ICA). A new, automated technique for identifying the damage component in component analysis methods was also described. This technique, based on the Mann-Kendal test, is effectively a trend analysis.

In the comparison it was found that under all investigated circumstances the component methods performed significantly better than the residual method, with ICA generally better than SVD except for a small number of outliers. For all three methods, the performance improved with larger damage size and more observations, and degraded with larger range of temperature variations. The results were then validated using experimental data monitoring a flat-bottom hole with growing diameter, the results being consistent with the synthetic data. These results showed that undamaged data from any installed monitoring system can be used to predict its damage detection performance. This makes it possible, for example, to assess whether a defect of a given size at a particular location would be detected reliably at a given data collection frequency. This enables the system to be tuned to meet particular requirements and the performance demonstration will be valuable in building safety cases.

8.3. Future work

It was not possible, within the limits of this project, to consider all the aspects which would ideally have been covered. The purpose if this section therefore is to highlight some areas where further work would be beneficial. This future work is presented in terms of the two main contributions of this thesis, in the area of the modelling of large area defect such as corrosion and signal processing techniques for guided wave monitoring data.

8.3.1. Rough surfaces

In the work presented in this thesis only rough surfaces with a Gaussian distribution were considered. This is the most general type of distribution to consider, but there is some evidence from the field that rough surfaces in pipe often have an exponential distribution. It therefore makes sense to expand the analysis of rough surfaces to look at rough surfaces with other types of distribution and other ranges of descriptive statistics. In particular it would be useful to look at rough surfaces where the correlation length is different around the circumference and along the pipe axis. Evidence in example surfaces from industry suggests that unequal correlation lengths do occur. There was also only a short consideration in this thesis of how detectability of other defects is affected by the presence of rough surfaces. This will be a key issue for future work, particularly how shallow generally rough surfaces can obscure more localised, deeper corrosion pits. There has also been little analysis of have automatic algorithms for processing guided wave monitoring data might deal with general corrosion. In particular it would be of interest to investigate how algorithms such as independent component analysis decompose rough surface signals and how this decomposition is a function of the variables of the surface.

8.3.2. Comparing guided wave monitoring algorithms

So far the framework outlined in this thesis has only been used to compare three different approaches for the processing of guided wave monitoring data. However, this framework could be used much more widely to compare many different processing methods. This is particularly appropriate given the new data sets which it is possible to construct using the recently validated superposition approach. There is also still work to be done in validating the independent component analysis algorithm. In this thesis it was shown that ICA was capable of identifying defect signals for simple types of defect such as flat bottom holes. It was also demonstrated with temperature as the model for environmental noise. However, there are many other types of defect which can exist in pipe systems and these defects will also need to be validated with ICA. Such defects might include pit clusters, shallow scallop types defects, distributed rough surfaces or large cross sectional loss pitting corrosion. So far these defects have only be modelled with step type defect growth or linear ramps in their cross sectional loss. It would also be worthwhile modelling other types of defect growth. Eventually the ICA algorithm will need to be validated with field data from operators of monitoring equipment.

Appendices

Appendix A:

Other outputs of Independent Component Analysis

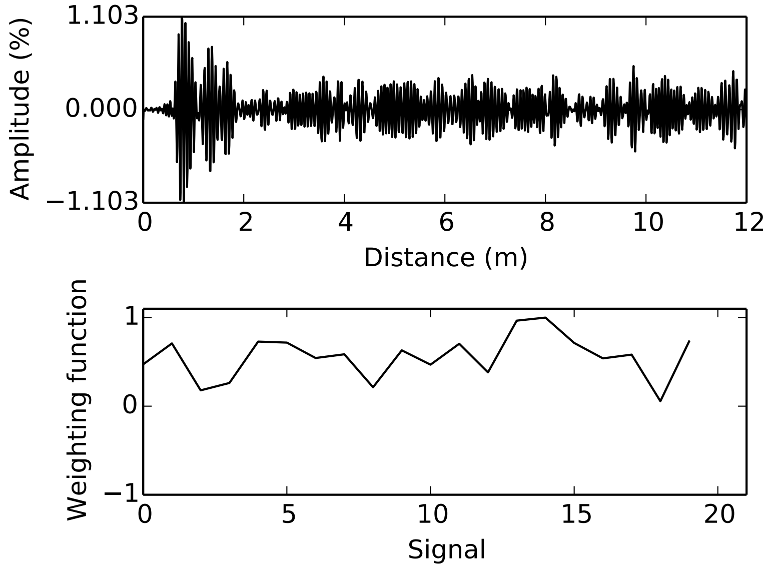

Figure A.1.: Other noise outputs from processing of the laboratory experiment with a simple step defect

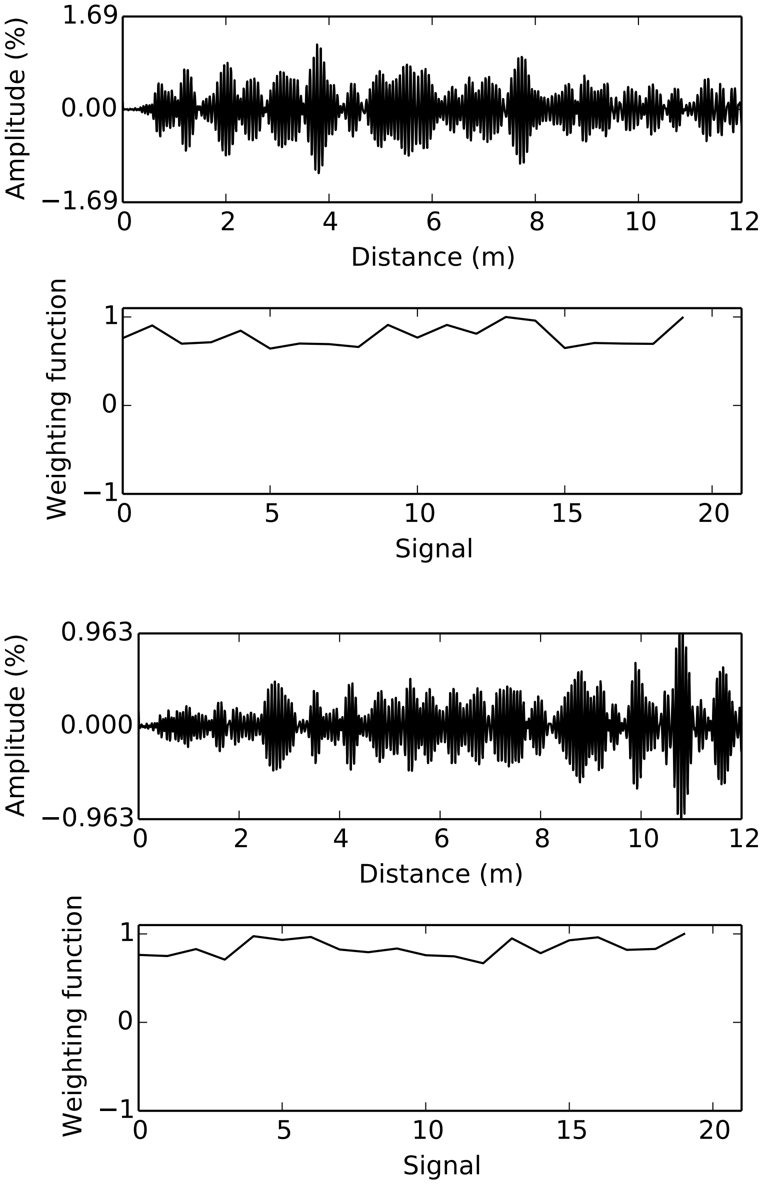

Figure A.2.: Other noise outputs from processing of the laboratory experiment with a simple step defect

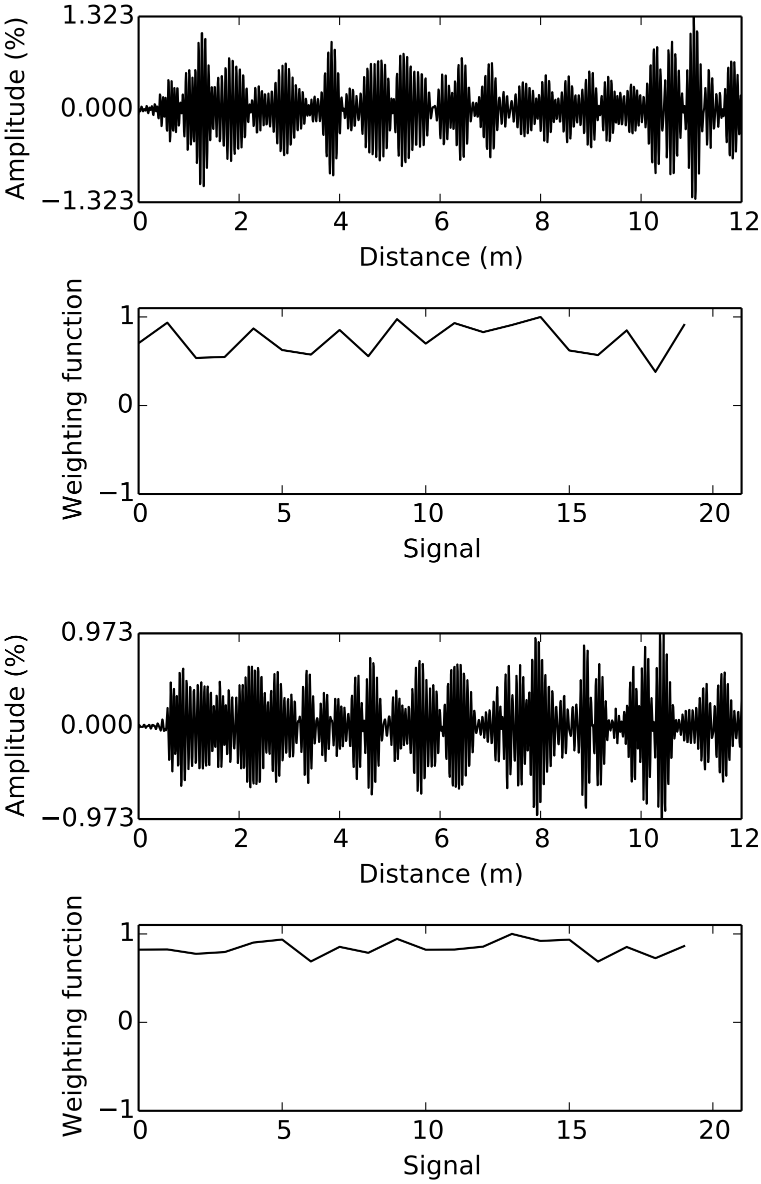

Figure A.3.: Other noise outputs from processing of the laboratory experiment with a simple step defect

Bibliography

[1] C. Restrepo, J. Simonoff, and R. Zimmerman. Causes, cost consequences and risk implications of accidents in US hazardous liquid pipeline infrastructure. *Critical Infrastructure Protection*, 2:38–50, 2009.

[2] F. Cegla, P. Cawley, J. Allin, and J. Davies. High-temperature (> 500 ^{o}C) wall thickness monitoring using dry-coupled ultrasonic waveguide transducers. *IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control*, 58:156–167, 2011.

[3] R. Hughes, Y. Fan, and S. Dixon. Near Electrical Resonance Signal Enhancement (NERSE) in eddy-current crack detection. *NDT & E International*, 66:82–89, 2014.

[4] Health and Safety Executive UK. Industrial radiography: managing radiation risks. Safety recommendations no. 1, Ionising Radiation Protection Series, 2015.

[5] M. Lowe and P. Cawley. Long range guided wave inspection usage – current commercial capabilities and research directions. Technical report, Imperial College London, 2006. URL `http://www3.imperial.ac.uk/pls/portallive/docs/1/55745699.PDF`.

[6] M. Silk and K. Bainton. The propagation in metal tubing of ultrasonic wave modes equivalent to Lamb waves. *Ultrasonics*, 17:11–19, 1979.

[7] W. Mohr and P. Holler. On inspection of thin-walled tubes for transverse and longitudinal flaws by guided ultrasonic waves. *IEEE Transactions on Sonics and Ultrasonics*, 23(5):369–374, 1976.

[8] R. B. Thompson, G. A. Alers, and M. A. Tennison. Application of direct electromagnetic Lamb wave generation to gas pipeline inspection. *Ultrasonics Symposium*, pages 91–94, 1972.

[9] D. Alleyne and P. Cawley. The excitation of Lamb waves in pipes using dry-coupled piezoelectric transducers. *Journal of Nondestructive Evaluation*, 15:11–20, 1996.

[10] P. Cawley. Practical long range guided wave inspection - managing complexity. In D. Thompson and D. Chimenti, editors, *Review of Progress in Quantitative Nondestructive Evaluation*, volume 657, pages 22–37. American Institute of Physics, 2003.

[11] A. Demma, P. Cawley, and M. Lowe. The reflection of the fundamental torsional mode from cracks and notches in pipes. *The Journal of the Acoustical Society of America*, 114:611–625, 2003.

[12] A. Lovstad and P. Cawley. The reflection of the fundamental torsional guided wave from multiple circular holes in pipes. *NDT & E International*, 44:553–562, 2011.

[13] D. Alleyne, M. Lowe, and P. Cawley. The reflection of guided waves from circumferential notches in pipes. *Journal of Applied Mechanics*, pages 635–641, 1998.

[14] R. Carandente and P. Cawley. The scattering of the fundamental torsional mode from axi-symmetric defects with varying depth profiles in pipes. *The Journal of the Acoustical Society of America*, 127:3440–3448, 2010.

[15] A. Demma, B. Pavlakovic, M. Lowe, P. Cawley, and A. G. Roosenbrand. The reflection of guided waves from notches in pipes: a guide for interpreting corrosion measurements. *NDT & E International*, 37:167–180, 2003.

[16] A. Galvagni and P. Cawley. The reflection of guided waves from simple supports in pipes. *The Journal of the Acoustical Society of America*, 129:1869–1880, 2011.

[17] P. Cawley, F. Cegla, and A. Galvagni. Guided waves for NDT and permanently-installed monitoring. *INSIGHT*, 54:594–601, 2012.

[18] C. Farrar, H. Sohn, and K. Worden. Data normalization: a key for structural health monitoring. Technical report, Los Alamos National Laboritory, 2004. URL `http://permalink.lanl.gov/object/tr?what=info:lanl-repo/lareport/LA-UR-01-4212`.

[19] A. Galvagni. *Pipeline health monitoring*. PhD thesis, Imperial College London, 2013.

[20] H. F. Pollard. *Sound waves in solids*. Pion, London, 1977.

[21] D. C. Gazis. Three-dimensional investigation of the propogation of wave in hollow circular cylinders. *The Journal of the Acoustical Society of America*, 31(5):568–574, 1958.

[22] A. H. Fitch. Observation of elastic-pulse propagation in axially symmetric and nonaxially symmetric longitudinal modes of hollow cylinders. *The Journal of the Acoustical Society of America*, 35:706, 1963.

[23] F. H. Quintanilla, Z. Fan, M. Lowe, and R. Craster. Guided waves' dispersion curves in anisotropic viscoelastic single and multi-layered media. *Proceedings of the Royal Society A: Mathematical Pyscial and Engineering Sciences*, 471(2183), 2015.

[24] D. E. Chimenti and O. Lobkis. Effect of rough surfaces on guided waves in plates. In D. Thompson and D. Chimenti, editors, *Review of Progress in Quantitative Nondestructive Evaluation*, volume 16, pages 169–176, 1997.

[25] G. Allaire. *Numerical Analysis and Optimization: An Introduction to Mathematical Modelling and Numerical Simulation*. Numerical mathematics and scientific computation. Oxford University Press, 2007.

[26] K.-J. Bathe and E. Wilson. *Numerical methods in finite element analysis,*. Prentice Hall, 1976.

[27] J. Chaskalovic. *Finite element methods for engineering sciences*. Springer, 2008.

[28] M. Drozdz. *Finite element modelling of ultrasound waves in elastic media*. PhD thesis, Imperial College London, 2008.

[29] M. Castaings, C. Bacon, B. Hosten, and M. Predoi. Finite element predictions for the dynamic response of thermo-viscoelastic material structures. *The Journal of the Acoustical Society of America*, 115:1125, 2004.

[30] M. Israeli and S. Orszag. Approximation of radiation boundary conditions. *Journal of Computational Physics*, 41:115–135, 1981.

[31] G. Liu and S. S. Quek. A non-reflecting boundary for analysing wave propagation using the finite element method. *Finite Elements in Analysis and Design*, 39:403–417, 2003.

[32] J. Sochacki, R. Kubichek, J. George, W. R. Fletcher, and S. Smithson. Absorbing boundary condition and surface waves. *Geophysics*, 52:60–71, 1987.

[33] M. Drozdz, M. Lowe, E. Skelton, and R. Craster. Modeling bulk and guided wave propagation in unbounded elastic media using absorbing layers in commercial finite element packages. *Review of Progress in Quantitative Nondestructive Evaluation*, 26:87–94, 2007.

[34] J. Pettit, A. Walker, P. Cawley, and M. Lowe. A stiffness reduction method for efficient absorption of waves at boundaries for use in commercial finite element codes. *Ultrasonics*, pages 1868–1879, 2014.

[35] F. Shi, W. Choi, E. Skelton, M. Lowe, and R. Craster. A time-domain finite element boundary integration method for ultrasonic nondestructive evaluation. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 61(12):2054–2066, 2014.

[36] P. Huthwaite. Accelerated finite element elastodynamic simulations using the GPU. *Journal of Computational Physics*, 257, Part A:687 – 707, 2014.

[37] H. Si. TetGen, a delaunay-based quality tetrahedral mesh generator. *ACM Transactions on Mathematical Software*, 41(11), 2015.

[38] C. Geuzaine and J. F. Remacle. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79:1309–1331, 2009.

[39] J. Schoberl. Netgen, an advancing front 2D/3D-mesh generator based on abstract rules. *Computing and Visualization in Science*, 1:41–52, 1997.

[40] N. Arnold. *Netgen user manual*. University of Zurich, 2013.

[41] P. Cawley and D. Alleyne. A two-dimensional Fourier transform method for the measurement of propagating multimode signals. *The Journal of the Acoustical Society of America*, 89:1159–1168, 1991.

[42] M. J. S. Lowe, D. N. Alleyne, and P. Cawley. The mode conversion of a guided wave by a part-circumferential notch in a pipe. *Journal of Applied Mechanics*, 65(3):649–656, 1998. URL `http://cat.inist.fr/?aModele=afficheN&cpsidt=1583686`.

[43] B. Pavlakovic, M. Lowe, D. Alleyne, and P. Cawley. Disperse: A general purpose program for creating dispersion curves. In D. O. Thompson and D. E. Chimenti, editors, *Review of Progress in Quantitative Nondestructive Evaluation*, volume 16 A of *16*, pages 185–192. Springer US, 1997.

[44] A. G. Voronovich and V. U. Zavorotny. Theoretical model for scattering of radar signals in ku- and c-bands from a rough sea surface with breaking waves. *Waves in Random Media*, 11:247–269, 2001.

[45] F. Bass and I. Fuks. *Wave scattering from statistically rough surfaces*, volume 93 of *International series in natural philosophy*. Pergamon, 1979.

[46] T. M. Elfouhaily and C.-A. Guerin. A critial survey of approximate scattering wave theories from random rough surfaces. *Waves in Random Media*, 14(R1-R40), 2004.

[47] D. McCammon. Application of a new theoretical treatment to an old problem; sinusoidal pressure release boundary scattering. *The Journal of the Acoustical Society of America*, 78(149), 1985.

[48] J. A. Ogilvy. Computer simulation of acoustic wave scattering from rough surfaces. *Journal of Physics D: Applied physics*, 21:260–277, 1988.

[49] A. Jarvis and F. Cegla. Application of the distributed point source method to rough surface scattering and ultrasonic wall thickness measurement. *The Journal of the Acoustical Society of America*, 132:1325–1335, 2012.

[50] W. Choi, E. Skelton, M. Lowe, and R. Craster. Unit cell finite element modelling for ultrasonic scattering from periodic surfaces. In D. Thompson and D. Chimenti, editors, *Review of Progress in Quantitative Nondestructive Evaluation*, volume 1511, pages 83–90, 2013.

[51] J. A. Ogilvy. Wave scattering from rough surfaces. *Reports of Progress in Physics*, 50:1553–1608, 1987.

[52] C. Eckart. The scattering of sound from the sea surface. *The Journal of the Acoustical Society of America*, 25:566–570, 1953.

[53] M. Stone. Wall thickness distributions for steels in corrosive environments and determination of suitable statistical analysis methods. In *4th European-American Workshop on Reliability of NDE*, 2009.

[54] J. Ogilvy and J. Foster. Rough surfaces: Gaussian or exponential statistics? *Journal of Physics D: Applied physics*, 22:1243–1251, 1989.

[55] B. Masserey and P. Fromme. On the reflection of coupled Rayleigh-like waves at surface defects in plates. *The Journal of the Acoustical Society of America*, 123: 88–98, 2008.

[56] H. Sohn. Effects of environmental and operational variability on structural health monitoring. *Philosophical Transactions of the Royal Society A*, 365:539–560, 2007.

[57] A. Croxford, P. Wilcox, and B. Drinkwater. Strategies for guided-wave structural health monitoring. *Proceedings of the Royal Society A*, 463:2961–2981, 2007.

[58] T. Clarke. *Guided wave health monitoring of complex structures.* PhD thesis, Imperial College, 2009.

[59] T. Clarke, F. Simonetti, and P. Cawley. Guided wave health monitoring of complex structures by sparse array systems: influence of temperature changes on performance. *Journal of Sound and Vibration*, 329:2306–2322, 2010.

[60] G. Konstantinidis, B. Drinkwater, and P. D. Wilcox. The temperature stability of guided wave structural health monitoring. *Smart Material and Structures*, 15: 967–976, 2006.

[61] Y. Lu and J. Michaels. A methodology for structural health monitoring with diffuse ultrasonic waves. *Ultrasonics*, 43:717–731, 2005.

[62] J. Ma. *On-line measurement of contents inside pipes using guided ultrasonic waves.* PhD thesis, Imperial College London, 2007.

[63] B. Peeters and G. D. Roeck. One year monitoring of the Z24-bridge: environmental influences versus damage events. In *18th International modal analysis conference*, volume 4062, pages 1570–1576, San Antonio, 2000. Society for experimental mechanics.

[64] P. Cawley, F. Cegla, and M. Stone. Corrosion monitoring strategies—choice between area and point measurements. *Journal of Nondestructive Evaluation*, 32: 156–163, 2013.

[65] A. Croxford, J. Moll, P. Wilcox, and J. Michaels. Efficient temperature compensation strategies for guided wave structural health monitoring. *Ultrasonics*, 50(45): 517–528, 2010.

[66] C. Liu, J. Harley, M. Berges, D. Greve, and I. Oppenheim. Robust ultrasonic damage detection under complex environmental conditions using singular value decomposition. *Ultrasonics*, 58:75–86, 2015.

[67] T. Cicero. *Signal processing for guided wave structural health monitoring.* PhD thesis, Imperial College London, 2009.

[68] S. Boyd and L. Vandenberghe. *Convex optimization.* Cambridge university press, 2004.

[69] D. E. Goldberg. *Genetic algorithms in search, optimization and machine learning.* Addison-Wesley Longman publishing, 1989.

[70] J. Nocedal and S. Wright. *Numercial Optimization.* Springer, 2000.

[71] J. Harley and J. Moura. Guided wave temperature compensation with the scale-invariant correlation coefficient. In K.-Y. Hashimoto, editor, *Ultrasonics Symposium 2011, IEEE International*, pages 1068–1071. IEEE, 2011.

[72] K. Worden, G. Manson, and N. Fieller. Damge detection using outlier analysis. *Journal of Sound and Vibration*, 229:647–667, 2000.

[73] A. Nairac, T. Corbett-Clark, R. Ripley, N. W. Townsend, and L. Tarassenko. Choosing an appropriate model for novelty detection. In *Proceedings of the 5th IEE International conference on artificial neural networks*, pages 117–122. Institue of Electrical Engineers, 1997.

[74] L. Hu, N. Hu, X. Zhang, F. Gu, and M. Gao. Novelty detection methods for online health monitoring and post data analysis of turbopumps. *Journal of Mechanical Science and Technology*, 27:1933–1942, 2013.

[75] P. Mahalanobis. On the generalised distance in statistics. In *Proceedings National Institute of Science India*, volume 2, pages 49–55, 1936.

[76] K. Worden, C. R. Farrar, G. Manson, and G. Park. The fundamental axioms in structural health monitoring. *Proceedings of the Royal Society A*, 463:1639–1664, 2007.

[77] A. Goh. Back-propagation neural networks for modeling complex systems. *Artificial Intelligence in Engineering*, 9:143–151, 1995.

[78] D. Plaut and G. Hinton. Learning sets of filters using back-propagation. *Computer Speech and Language*, 2:35–61, 1987.

[79] M. Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, 37:233–243, 1991.

[80] N. Dervilis, R. Barthorpe, I. Antoniadou, and K. Worden. Impact damage detection for composite material typical of wind turbine blades using novelty detection. In *6th European workshop on structural health monitoring*, Dresden, Germany, 2012. EWSHM.

[81] H. Abdi and L. Williams. Principal component analysis. *Wiley Interdisciplinary Reviews: computational statistics*, 2:433–459, 2010.

[82] G. Golub and C. V. Loan. *Matrix computations*. John Hopkins university press, 1996.

[83] D. Tibaduiza, L. Mujica, and J. Rodellar. Damage classification in structural health monitoring using principal component analysis and self-organizing maps. *Structural Control and Health Monitoring*, 20:1303–1316, 2013.

[84] L. Mujica, J. Vehi, M. Ruiz, M. Verleysen, W. Staszewski, and K. Worden. Multivariate statistics process control for dimensionality reduction in structural assessment. *Mechanical Systems and Signal Processing*, 22:155–171, 2008.

[85] I. Buethe, M. Torres-Arredondo, L. Mujica, J. Rodellar, and C. Fritzen. Damage detection in piping systems using pattern recognition techniques. In *6th European workshop on structural health monitoring*, pages 1–8, Dresden, Germany, 2012. EWSHM.

[86] C. liu, J. Harley, Y. Ying, I. Oppenheim, M. Berges, D. Greve, and J. Garrett. Singular value decomposition for novelty detection in ultrasonic pipe monitoring. In J. P. Lynch, C.-B. Yun, and K.-W. Wang, editors, *Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems*, volume 8692, San Diego, California, USA, 2013.

[87] L. Trefethen and D. Bau. *Numerical linear algebra*. Society of Industrial and Applied Mathematics, 1997.

[88] A. Hyvarinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Trans. Neural Netw.*, 10:626–634, 1999.

[89] J. Dobson and P. Cawley. Independent component analysis for improved defect detection in guided wave monitoring. *Proceedings of the IEEE*, pages 1–12, 2015.

[90] A. Hyvarinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. Wiley Interscience, 2001.

[91] A. Delrome and S. Makeig. EEGLAB: an open source toolbox for analysis of single analysis EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134:9–21, 2004.

[92] P. Hojen-Sorenson, O. Winther, and L. K. Hansen. Mean field approaches to independent component analysis. *Neural Computation*, 134:889–918, 2002.

[93] D. Langlois, S. Chartier, and D. Gosselein. An introduction to independent component analysis: InfoMax and FastICA algorithms. *Tutorials in Quantitative Methods for Psychology*, 6:31–38, 2010.

[94] F. Cegla, A. Rohde, and M. Veidt. Analytical prediction and experimental measurement for mode conversion and scattering of plate waves at non-symmetric circular blind holes in isotropic plates. *Wave Motion*, 45:162–177, 2008.

[95] A. Raghavan and C. E. Cesnik. Effects of elevated temperature on guided-wave structural health monitoring. *Journal of Intelligent Material Systems and Structures*, 19(12):1383–1398, 2008. URL `http://jim.sagepub.com/content/19/12/1383.short`.

[96] L. G. Zhuravlev. Temperature dependence of the mechanical properties of high-carbon austenitic steel. *Metal Science and Heat Treatment*, 6(2):92–93, 1964. ISSN 0026-0673, 1573-8973. doi: 10.1007/BF00655383. URL `http://link.springer.com/10.1007/BF00655383`.

[97] J. A. Ogilvy. Model for predicting ultrasonic pulse-echo probability of detection. *NDT & E International*, 26(1):19–29, 1993. ISSN 0963-8695. doi: 10.1016/0963-8695(93)90160-V. URL `http://www.sciencedirect.com/science/article/pii/096386959390160V`.

[98] J. R. Davis. *ASM Handbook: Nondestructive evaluation and quality control (3rd ed.)*, volume 17. ASM International, 1994.

[99] F. Schoefs, A. Clément, and A. Nouy. Assessment of ROC curves for inspection of random fields. *Structural Safety*, 31(5):409–419, 2009. ISSN 0167-4730. doi: 10.1016/j.strusafe.2009.01.004. URL `http://www.sciencedirect.com/science/article/pii/S0167473009000034`.

[100] T. Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8): 861–874, 2006. ISSN 0167-8655. doi: 10.1016/j.patrec.2005.10.010. URL `http://www.sciencedirect.com/science/article/pii/S016786550500303X`.

[101] Y. Lu and J. Michaels. Feature extraction and sensor fusion for ultrasonic structural health monitoring under changing environmental conditions. *Sensors Journal, IEEE*, pages 142–1471, 2009.

[102] C. Haynes, M. Todd, E. Flynn, and A. Croxford. Statistically-based damage detection in geometrically-complex structures using ultrasonic interrogation. *Structural Health Monitoring*, 12:141–152, 2013.

[103] J. B. Harley and J. M. F. Moura. Scale transform signal processing for optimal ultrasonic temperature compensation. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 59(10):2226–2236, 2012. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6327494.

[104] G. J. Székely, M. L. Rizzo, and N. K. Bakirov. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35(6):2769–2794, 2007. URL http://projecteuclid.org/euclid.aos/1201012979.

[105] A. M. Fraser and H. L. Swinney. Independent coordinates for strange attractors from mutual information. *Physical Review A*, 33(2):1134, 1986. URL http://journals.aps.org/pra/abstract/10.1103/PhysRevA.33.1134.

[106] M. Turk and A. Pentland. Face recognition using eigenfaces. In *, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 1991. Proceedings CVPR '91*, pages 586–591, 1991. doi: 10.1109/CVPR.1991.139758.

[107] A. Hyvarinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13:411–430, 2000.

[108] H. B. Mann. Nonparametric tests against trend. *Econometrica: Journal of the Econometric Society*, pages 245–259, 1945. URL http://www.jstor.org/stable/1907187.

[109] M. Kendall. *Multivariate analysis*. Charles Griffin, 1975. URL http://infoscience.epfl.ch/record/2795.

[110] R. O. Gilbert. *Statistical methods for environmental pollution monitoring*. John Wiley & Sons, 1987. URL https://books.google.co.uk/books?hl=en&lr=&id=lEo1rvDGUEkC&oi=fnd&pg=PR9&dq=book+gilbert+1987+mann-kendall&ots=F91eOT5Bs4&sig=tLBy2s3PZed70vVBYUrE8pNeIPc.

[111] R. M. Hirsch and J. R. Slack. A nonparametric trend test for seasonal data with serial dependence. *Water Resources Research*, 20:727–732, 1984. ISSN 0043-1397. doi: 10.1029/WR020i006p00727. URL http://adsabs.harvard.edu/abs/1984WRR....20..727H.

[112] P. D. Valz and A. I. McLeod. A simplified derivation of the variance of Kendall's rank correlation coefficient. *The American Statistician*, 44(1):39–40, 1990. URL http://www.tandfonline.com/doi/pdf/10.1080/00031305.1990.10475691.

[113] A. P. Bradley. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159, 1997. URL http://www.sciencedirect.com/science/article/pii/S0031320396001422.

List of my Publications

[1] J. Dobson and P. Cawley. Independent Component Analysis for improved defect detection in guided wave monitoring. *Proceedings of the IEEE*, PP:1–12, 2015.

[2] J. Dobson and P. Cawley. Independent Component Analysis for improved defect detection in guided wave monitoring. pages 1878–1885. DESTECH PUBLICATIONS, INC, 2015.

[3] C. Liu, J. Dobson, and P. Cawley. Practical evaluation of SHM damage detection under complex environmental conditions using receiver operating characteristics. pages 1949–1956. DESTECH PUBLICATIONS, INC, 2015.

[4] J. Dobson and P. Cawley. The scattering of torsional guided waves from Gaussian distributed rough surfaces in pipework. *Journal of the Acoustical Society of America*, To appear, 2016.