

Department of Precision and Microsystems Engineering

Adjoint Based Guided Wave Tomography for Corrosion Imaging in Pipelines

Marn Klein Holkenborg

Report no : 2025.009
Supervisors : Jieun Yang (TUD)
Emiel Hassefras (TNO)
Arno Volker (TNO)
Specialisation : Mechatronic System Design
Type of report : MSc. Thesis
Date : 7 March 2025

Adjoint Based Guided Wave Tomography for Corrosion Imaging in Pipelines

Marn Klein Holkenborg

Delft University of Technology & TNO


TU Delft

TNO innovation
for life

Adjoint Based Guided Wave Tomography for Corrosion Imaging in Pipelines

by

Marn Klein Holkenborg

4660269

TNO supervisor:	Emiel Hassefras
TNO supervisor:	Arno Volker
Department:	Acoustic & Underwater Warfare
TU Delft supervisor:	Jieun Yang
Faculty:	Mechanical Engineering
Department:	Precision and Microsystems Engineering

Cover: *Liu, H. (2024, March 29). pipeline-technology. Encyclopedia Britannica. (Modified)*

Abstract

In industries such as oil and gas, power generation, and chemical processing, the structural integrity of piping systems is crucial for operational safety and reliability. Various factors such as corrosion, erosion, and manufacturing defects can result in a loss of wall thickness, potentially leading to catastrophic failures. Detecting and assessing the integrity of piping systems without causing disruption to operations or compromising their structural integrity is a challenge in asset management and maintenance.

Guided Wave Tomography (GWT) is a promising non-destructive testing (NDT) technique for assessing the integrity of a pipeline by mapping its wall thickness. Using tomographic algorithms that use forward wavefield extrapolators, based on the gradient of the misfit a defect reconstruction can be made. Although the forward model is computationally inexpensive, the currently implemented gradient calculation is costly due to the large number of simulations required for each inversion parameter. This study explores the feasibility of using the adjoint state method for gradient computation in GWT, requiring only a single forward simulation independent of the number of model parameters.

A nonlinear conjugate gradient optimization method is implemented to iteratively update the wall thickness based on the misfit between simulated and observed data. The inversion process makes use of the dispersion curves to establish the relationship between phase velocity, frequency, and wall thickness. The study specifically focuses on the S_0 mode, analysing its suitability for complex defect reconstruction. Numerical experiments demonstrate that the adjoint method effectively localizes defects in both noise-free and noise environments, provided that there is a sufficient contrast in the sensitivity kernel between the defect and its surroundings. The technique achieves a resolution of approximately 1 to 1.5λ , with frequency continuation and multimodal inversions improving reconstruction accuracy. However, challenges arise in handling noise and wall thickness interchanges, which can be mitigated using adaptive regularization strategies.

The results suggest that the adjoint method is a viable alternative to traditional finite-difference sensitivity calculations. Future research should focus on experimental validation with real-world or as noise-free alternative 3D elastodynamic data evaluating the method on complex defect geometries and varying pipe configurations. Finally, the algorithm can be enhanced by integrating Hessian-based optimization techniques, thereby implementing a second-order adjoint method.

Preface

This thesis concludes my master's degree in High-Tech Engineering, a track of the Mechanical Engineering master at the Delft University of Technology, and therefore also my time as a student.

When I started my studies I did not image myself graduating on the topic of corrosion detection in pipeline systems. However, although being a very niche topic, in the past year I have seen that it is a whole fascinating world on itself, and one can be easily lost in all aspects that influence a high resolution reconstruction.

Therefore I would like to thank Emiel and Arno, my supervisors at TNO, who have introduced me to the topic in a very nice and structured way. The expertise and passion that you show to the project is very impressive and has been very motivating throughout the year.

Also my appreciation goes to Jieun, who has been very helpful throughout the year by having recurring meetings in which I try to make sense of my own findings, while you asking the correct questions to guide me to a way better understanding of the matter.

Finally, I want to thank my family for their limitless support throughout my whole studies. Without them the unforgettable experience at Forze Hydrogen Racing would not have been possible, aswell as the countless other fun moments throughout my time as student in Delft.

Enjoy reading!

Marn Klein Holkenborg
Delft, March 2025

Contents

Abstract	i
Preface	ii
1 Literature review	1
1.1 Introduction	1
1.2 State of the art	2
1.3 Research gap	3
1.4 Conclusion	4
2 Methodology	5
2.1 Guided waves	5
2.2 Wave mode selection	8
2.3 Forward modeling	8
2.3.1 One-way wavefield extrapolation	9
2.4 Full waveform inversion	10
2.5 Gradient optimization	12
2.5.1 Perturbation method	12
2.5.2 Adjoint state method	12
2.5.3 Limited view	14
2.6 Constrained non linear optimization methods	15
2.6.1 Conjugate Gradient	17
3 Numerical simulations	19
3.1 Model configuration	19
3.1.1 Plate and defect configuration	19
3.1.2 Frequency selection	21
3.1.3 Transducer configuration	23
3.1.4 Reducing non-uniqueness	26
3.1.5 Performance indication	27
3.2 Inversion results for the Gaussian defect	28
3.2.1 Single frequency inversion	28
3.2.2 Multi frequency inversion	31
3.2.3 Multi mode inversion	33
3.2.4 Parametric study	35
3.3 Inversion results for the Echo defect	37
3.3.1 Model configuration	37
3.3.2 Single frequency inversion	38
3.3.3 Multi frequency inversion	41
3.3.4 Parametric study	42
3.4 Inversion performance under noise	44

4	Discussion	51
4.1	Inverse crime	51
4.2	Scaling and Regularization	52
4.2.1	Kernel scaling	52
4.2.2	Regularization	52
5	Conclusion	55
5.1	Recommendations and further research	56
5.1.1	Hessian-Based Optimization and trust region	56
5.1.2	Validation on Real-World and 3D Elastodynamic Data	56
5.1.3	Validation on different configurations	56
A	Derivation of the Gradient of the Objective Function	58
B	Smoothing defect shapes by wavenumber domain filtering	59
C	Varying the $\alpha_{0,k}$ maximal growth rate	60
D	Initial kernels of non circular defects	62
	References	64

Literature review

1.1. Introduction

In various industries, the integrity of piping systems is important to ensure operational safety, reliability, and durability. Piping networks, serving critical functions in sectors such as oil and gas, chemical processing, power generation and many more, are vulnerable to degradation over time due to factors such as corrosion, erosion, fatigue, dents and manufacturing defects, as can be seen in figure 1.1. The structural failure of a piping system can potentially lead to disastrous events. For example, in March 2006, a leaking pipe spilled approximately 5054 barrels of crude oil on the North Slope in Alaska. A second spill occurred in August 2006, spilling approximately 24 barrels of crude oil [1]. In order to prevent future disasters, accurate thickness mapping of large piping structures containing corrosion damages is critical to assess the integrity and residual life of the components. However, detecting and assessing the integrity of piping systems without causing disruption to operations or compromising their structural integrity is a fundamental challenge in asset management and maintenance.

Non destructive testing (NDT) techniques play a vital role in the assessment and monitoring of piping integrity. One of the NDT methods available, besides methods such as eddy current testing [3], magnetic flux leakage testing [4, 5] or radiography testing [6], is ultrasonic wave testing [7, 8]. The goal of these methods is to create wall thickness maps of pipes and plates, thereby ensuring the structural integrity of the system. Conventional ultrasonic thickness mappings are generated through the systematic traversal of a probe across the entirety of the pipe surface, creating the thickness profiles based on the differences in wave reflections. Guided wave tomography (GWT) has emerged as a powerful and versatile approach for inspecting long lengths of piping due to the fact that it can monitor areas




Figure 1.1: Example of typical corrosion defects (a) and dents (b) on a pipe [2].

with restricted access efficiently and effectively [9, 10]. GWT removes the requirement to have direct access to all points on the surface of the pipe. Furthermore, it also has the benefit of being faster since there is no need for a point-by-point measurement. The inherent dispersive characteristics exhibited by ultrasonic guided waves in plates have been widely utilised to create thickness maps. Specifically, Lamb waves propagating at a defined frequency exhibit a change in phase and group velocities when they interact with a region containing a reduced wall thickness caused by corrosion [11]. The accuracy of GWT is highly dependent on the underlying assumptions used in the reconstruction of the wave propagation.

In the context of pipe measurements, the defect is not encompassed by transducers; instead, it is flanked by two ring arrays, as depicted in Figure 1.2a. In many industrial applications, the diameter of a pipe is much larger than its wall thickness. In this case, a pipe can be considered as an unwrapped plate [12]. The problem can then be conceptually transformed into a plate with the defect positioned between two linear arrays. Reconstruction from limited view data presents an ill posed problem, where multiple solutions (defect geometries) can result in the same measured data. The cyclic nature of the pipe allows waves to travel numerous times around its circumference, these helical paths enable the same mode to reach a receiver at different times, which is visualised in Figure 1.2b. The information contained in higher order helical modes can be exploited in GWT, increasing depth estimation accuracy [13].




Figure 1.2: The unwrapping of a pipe section. as per [14] (a) Three distinct wave paths traverse a pipe section, originating from a singular source and converging at a common receiver. (b) The unwrapped structure exhibits the direct path alongside two helical modes, characterized by replication values of $n = [1 \ 0 \ 1]$. Note that the red squares all represent the same receiver, where the signal reaches it at different times.

1.2. State of the art

In the process of creating quantitative wall thickness maps, the resolution is strongly affected by the model used for wave propagation. Therefore, it is of great importance to identify a suitable model to describe the behavior of the wave propagation. One of these possible models is based on the ray theory of geometric optics, which is a widely used assumption and can be used as a forward model in an acoustic wavefield inversion algorithm. The waves are assumed to propagate as rays, with the arrival time being a line integral of the slowness (which is the reciprocal of the velocity profile) along the ray path. Wong et al. [15] introduced a method using straight rays parallel to a borehole, but also methods using bent rays exist to account for refraction effects [16]. Although these ray methods are widely used in GWT [17], they do not account for any diffraction effects. Furthermore, the resolution of this method is limited by the width of the first Fresnel zone $\sqrt{\lambda L}$, where L is the distance between the source and the receiver, and λ is the wavelength of the illuminating wavefield [18].

In order to improve performance, several attempts have been made using Born approximations as forward modeling assumptions. The Born approximation simplifies the analysis by considering only the first order scattering of waves by the medium. This means that the interaction between waves and the medium remains sufficiently weak, allowing the treatment of scattered waves as if they solely come from the incident waves, avoiding the calculations of the interactions among the scattered waves themselves. Therefore, this method accounts for diffraction effects. However, experimental results demonstrate that, ultimately, the scattering from the array of transducers is a major source of error in the tomographic reconstruction. When there is no scattering from the array of transducers, the reconstructions are very similar to the inversions with modeled data [19]. However, the method is only applicable when the perturbation is relatively small and of low contrast, thereby being weakly scattering. This given makes the method impractical for most field applications.

Huthwaite et al. proposed the hybrid algorithm for robust breast ultrasound tomography (HARBUT), which utilises the strengths of both ray and diffraction tomography. The algorithm is shown to produce accurate reconstructions with realistic data from a complex three dimensional simulation. An in plane resolution of 1 mm was achieved with a slice thickness of 9 mm. However, the known resolution is limited to the diffraction limit [20] for idealized cases, and to around $1.5-2\lambda$ for realistic test data.

In order to deal with the challenges stated above, a comprehensive full waveform inversion (FWI) technique is employed within guided wave tomography to map thickness profiles. Initially established within geophysics for seismic wave imaging [21, 22, 23], this method utilises a computational forward model and the dispersive nature of guided waves through defect regions, along with an iterative inverse model to reconstruct defect profiles. Throughout each iteration, numerical simulations are conducted with the objective of minimising the least squared misfit between modelled and observed data. The achieved resolution by FWI strongly depends on the essential ingredients: an efficient forward modelling scheme, as well as an efficient gradient estimator.

There exists a wide range of methodologies employed for forward modeling in FWI. The most popular method to discretise the wave equation in the time or frequency domain is the finite difference method [24, 25]. However, techniques such as finite element methods [26, 27], finite volume methods [28], discrete wave number methods [29] and more [21] are all used in literature. Recently, the utilisation of recursive extrapolation for forward modelling in the frequency space (FX) domain has proven to enable fast forward modelling for extensive pipe sections [14]. Using one way wave equations and assuming zero reflection, the method approaches the speed of ray tracing while accurately accounting for diffraction phenomena. It employs a table driven approach, with recomputed extrapolation operators stored across a spectrum of wave numbers, giving computational benefits compared to other forward models. Note, however, that due to the nature of the one way wave equations, the model does not consider backscatter or reflections. In this thesis, we only analyse relatively smooth corrosion defects, hence creating insufficient medium contrast for reflections, motivating the use of the recursive wave extrapolation method as a forward model in the inversion process for this thesis.

Because of the substantial computational expenses associated with computing synthetic data, the full nonlinear problem of waveform inversion typically adopts an iterative "descent" methodology. In this approach, the reduction of residuals is achieved through repeated calculation of a local gradient. Each iteration's gradient indicates the direction for minimising the objective function, in FWI, typically the L2 norm of data residuals. Prior to proceeding with the inversion of the linearised system, the sensitivity or Fréchet derivative matrix, which is the partial derivative of data concerning the model parameters, is explicitly determined.

1.3. Research gap

In the current tomography implementation at TNO, the derivative matrix is calculated with the use of a forward difference scheme. To that end, each parameter of interest has to be perturbed, and a new forward simulation has to run. Hence, constructing the sensitivity matrix becomes increasingly prohibitive to achieve high resolution imaging due to the high parameter count. TNO aims to develop a guided wave tomography technique that uses an adaptive meshing strategy to concentrate a large number of parameters around the defect shape. This approach, combined with a newly developed multi mode and frequency strategy tailored to address non uniqueness, improves the quality of the resulting

images. However, running a simulation for each parameter is impractical, even with the relatively inexpensive FX method.

Therefore, this thesis aims to implement an alternative gradient calculation strategy called the adjoint state method [30]. This methodology entails computing the gradient of the misfit function without explicitly constructing the sensitivity matrix. Contrarily, the computational effort required for gradient calculation in the adjoint method is similar to a single forward model. An important research question is whether this method produces misfit gradients of accuracy comparable to the state of the art tomography algorithm. To the author's knowledge, no prior research has been conducted on this topic.

1.4. Conclusion

The integrity of piping systems across industries is critical for operational safety and reliability. Various factors such as corrosion, erosion, and manufacturing defects can compromise these systems, potentially leading to catastrophic failures. Non destructive testing techniques, including ultrasonic wave testing, play a crucial role in assessing and monitoring the integrity of piping structures.

Guided wave tomography has emerged as a promising approach for inspecting long lengths of piping, offering advantages such as efficient monitoring in areas with restricted access. However, creating accurate thickness maps using GWT presents challenges, particularly in modeling the wave propagation.

Current methodologies for thickness mapping, including ray theory and Born approximations, have limitations in addressing diffraction effects and achieving high resolution imaging. Full waveform inversion techniques, originally developed in geophysics, offer a promising solution by iteratively reconstructing defect profiles based on numerical simulations.

Recent advancements in FWI include the use of frequency space domain methods for fast forward modeling. These developments aim to improve the efficiency and accuracy of thickness mapping in GWT, particularly in handling complex defect geometries and reducing computational costs.

The proposed research aims to address the limitations of current gradient calculation methods in GWT by implementing the adjoint state method. By bypassing the explicit construction of sensitivity matrices, this approach promises to reduce computational overhead while potentially maintaining accuracy. The effectiveness of this method, particularly in conjunction with one way forward models, will be the main subject of investigation for the master thesis.

2

Methodology

A guided wave is a type of wave whose energy is confined between boundaries and propagates along a waveguide [10]. Unlike bulk waves, which propagate freely without confinement to specific boundaries, guided waves are continuously reflected between the structure's boundaries. These waves are generated by an excitation at a specific location in the structure, producing various wave modes, each characterized by unique properties and suited to specific applications. These modes differ in their displacement patterns, wave velocities, and dispersion characteristics.

2.1. Guided waves

Lamb waves are a type of guided elastic wave that propagates in a thin plate or shell like structures, such as solid materials like metals, composite materials, or biological tissues. They were first described in a free plate problem (as seen in Figure 2.1) by the British mathematician Horace Lamb in 1917 [31].




Figure 2.1: Free plate problem.

These guided waves can propagate long distances without significant attenuation and interact strongly with thickness variations. The propagation characteristics of Lamb waves depend on various parameters, including the thickness and material properties of the structure, as well as the frequency and the mode of excitation. They exhibit dispersion, meaning that different frequency components of the wave propagate at different velocities, leading to waveform distortion over long distances [32]. These waves consist out of symmetric (S) and antisymmetric (A) modes, which describe the relative motion of particles on opposite sides of the plate's midplane, as indicated by the orange arrow in Figure 2.2, both having their own dispersive characteristics.




Figure 2.2: Propagation of Lamb waves in a plate, illustrating the different characteristics of symmetrical and asymmetrical wave modes.

Another wave mode that exist is the shear horizontal (SH) wave mode. The particle motion is parallel to the plane of the surface and orthogonal to the direction of propagation. This transverse motion results in in-plane displacement, simplifying their dynamic behavior compared to other guided waves. These waves are particularly notable for their simplicity compared to other guided wave types, as they lack coupling with longitudinal or flexural motions in isotropic materials, making them relatively easier to analyze and interpret.

To understand and predict the behavior of these wave modes, it is essential to delve into the dispersion equations governing wave propagation. Dispersion equations provide a quantitative relationship between key wave parameters: frequency, wavenumber, and phase velocity. Which encapsulate the influence of boundary conditions and material properties on wave behavior. These equations arise from solving the wave equation under the stress free boundary conditions of a finite structure, for the derivation the reader is referred to [33].

Solving the dispersion equations demonstrates the existence of multiple ultrasonic guided wave modes within a plate. Multiple solutions emerge for specific combinations of frequency and plate thickness. Therefore only numerical simulation methods can be used to solve these equations. Using a MATLAB based software, 'Dispersion Calculator' (DC) [34], the dispersion curves for multiple wave modes can be obtained, which exist in a certain material. This is then calculated, however, in the free case where the plate is assumed to be surrounded by a vacuum so that no energy can leave the plate.

Analogues to the free plate problem, one can derive the case with a liquid loading of water on the plate [35]. Liquid loading is relevant as pipelines are typically filled with liquids such as water or oil. One characteristic of Lamb waves is that they exhibit out of plane displacement [31]. Consequently, they have the potential to lose acoustic energy into the liquid medium, thereby leading to a decay or complete attenuation of the desired guided wave amplitude before it reaches the receiver. Furthermore, if this energy couples with the liquid, it can create a pathway for extra wave transmission between the source and receiver, introducing undesired coherent noise into the data.

In Figure 2.3, the phase velocities are depicted for the symmetric, anti symmetric and shear horizontal wave modes in a steel plate with material properties $E = 200 \text{ GPa}$, $\nu = 0.3$ and $\rho = 7850 \text{ kg/m}^3$. Here A_0 , S_0 and SH_0 are the anti symmetric, symmetric and shear horizontal fundamental wave modes, respectively. Furthermore, A_n , S_n and SH_n represent the higher order wave modes, with n denoting their order. The horizontal axis represents the product of frequency and the plate thickness, while the vertical axis denotes the phase velocity. The DC software is utilised for generating the graph illustrating the attenuation characteristics of wave modes within a steel plate. Figure 2.4 illustrates the attenuation behaviour of these various wave modes.

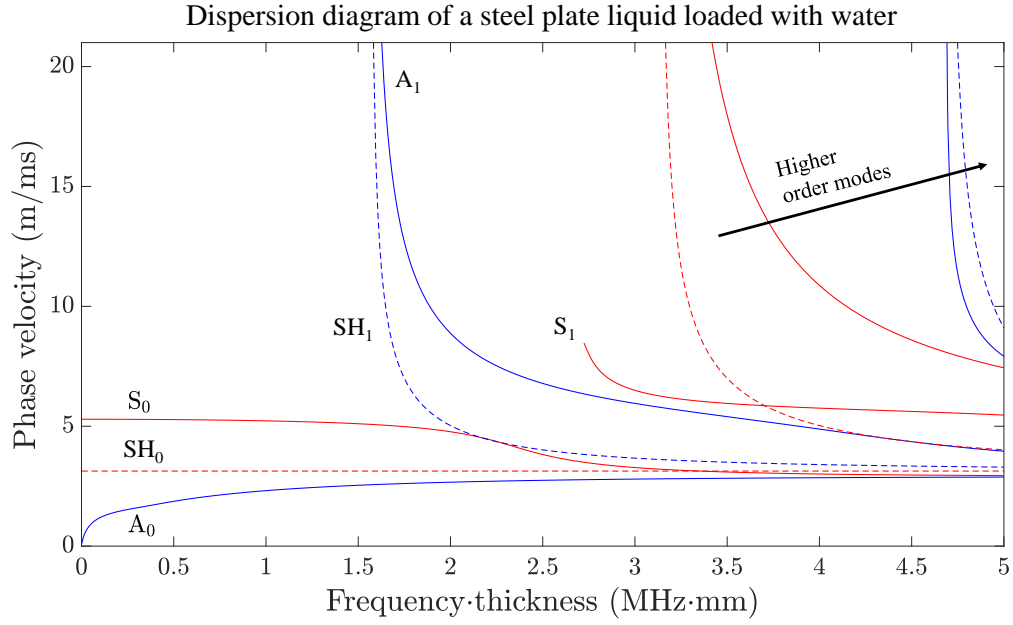


Figure 2.3: Dispersion curves for different modes in a steel plate with liquid loading of water [34].

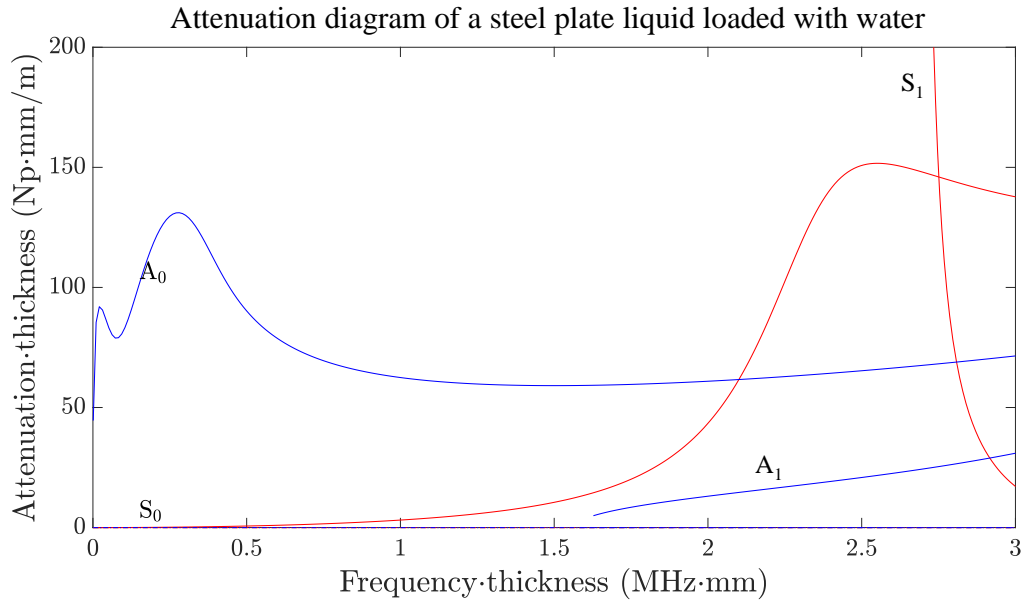


Figure 2.4: Attenuation for different wave modes in a steel plate with liquid loading of water. Np is the natural logarithm of the ratio of the initial amplitude to the final amplitude [34].

2.2. Wave mode selection

One of the fundamental aspects of the design of a guided wave tomography system is the choice of wave mode, which strongly affects the sensitivity of the inversion algorithm to defect interactions. Using the dispersion curves and assuming a fixed frequency, a relationship between the inverted phase velocity and the thickness of the plate can be established. When a wave mode displays high sensitivity to thickness variation, it can be utilized for thickness estimation. This dependency is represented by the dispersion curve's gradient, a steep gradient correlates with substantial velocity shifts relative to thickness. Although higher order modes are generally more sensitive than the fundamental wave modes, for certain frequencies, higher order wave modes co exist with other wave modes, as can be seen in Figure 2.3. Wave mode separation techniques do exist, Jidong Hou et al. [36] present a method using arrival time estimation to separate the multiple lamb wave modes. Although a smooth distribution of arrival time for the fundamental wave modes and A_1 wave mode is obtained, the mechanism behind the method is adding extra complexity and noise, especially for wave modes of higher order than A_1 , when compared to the utilisation of the fundamental wave modes S_0 and A_0 . Furthermore, each higher order wave mode has a cutoff frequency for which they don't exist for specific frequency thickness values (e.g. around 1.6 MHz·mm for the SH_1 mode). Therefore, when the frequency is fixed, below a certain thickness, these waves will be unable to propagate through the medium. The fundamental shear horizontal wave mode SH_0 , is not explored due to the non dispersive nature of the mode. However, the utilisation of the first higher order shear horizontal mode SH_1 has been explored, enabling an improvement in resolution compared to the fundamental modes [33].

Looking at Figure 2.3 we find that that A_0 is much more sensitive than S_0 at low frequency thickness values, which due to practical limits on the operating frequencies [37] is advantageous over S_0 which is more sensitive at frequency thickness values of around 1.5 MHz·mm. As can be seen in Figure 2.4, when a liquid boundary is present, such as in an oil pipe, the wave attenuation for A_0 is very high and therefore not well suited for liquid loaded pipes. From Figure 2.4, it can be found that the S_0 wave mode is less sensitive for attenuation at low frequency thickness values, hence allowing signals to be propagated and properly received at the receiver array. This thesis is part of a research project that explores the joint inversion of fundamental wave modes and shear horizontal wave modes. Within the scope of this thesis, we initially focus on the S_0 wave mode. This is because it has low attenuation and relatively high sensitivity at low frequency thickness values. However, numerical simulations incorporating multiple wave modes within a single simulation will be conducted aswell in this thesis to investigate their impact on mitigating the ill posedness of the problem. In these simulations, the A_0 wave mode will be examined alongside the S_0 wave mode.

2.3. Forward modeling

Forward modeling serves as a computational technique aimed at simulating the propagation behavior of guided waves within a given structure. The process of forward modeling typically involves solving the governing wave equations, such as the linear elastic wave equation or acoustic wave equation, within the structure of interest. These equations govern the propagation of guided waves, taking into account factors like material elasticity, geometry, boundary conditions, excitation sources, and defects [38].

The main model of acoustics in the frequency domain can be derived based on the Helmholtz equation. The Helmholtz equation is more efficient to solve than the time-dependent wave equation because it is time-independent, allowing for a static, frequency-domain analysis and reducing computational complexity. Assuming a constant density ρ , the Fourier transformed pressure field $P(\vec{r}, \omega)$ obeys the Helmholtz equation and is given as:

$$\nabla^2 P(\vec{r}, \omega) + \frac{\omega^2}{v_p^2} P(\vec{r}, \omega) = 0, \quad (2.1)$$

Where ω is the frequency, v_p the phase velocity and \vec{r} equals $\sqrt{x^2 + y^2 + z^2}$.

2.3.1. One-way wavefield extrapolation

While the elastodynamic wave equation offers a detailed description of the interactions of guided waves with defects, its computational requirements frequently pose practical limitations for real-world scenarios. To address this challenge, numerous contemporary implementations use 2D acoustic forward models as a viable alternative. The finite difference method applies to both the acoustic and elastic two-way wave equations. However, two-way wave-equation techniques often require substantial computational resources. Alternatively, methods based on one-way wave equations typically offer greater computational efficiency and straightforward implementation compared to their two-way counterparts. A method operating in the frequency space domain (FX) has been utilized in [39, 40]. A one-way wavefield can be propagated in a homogeneous media using an extrapolation operator in the wavenumber domain [41]. The acoustic approximation is made by utilizing the phase velocity for a given frequency in an acoustic model.

In order to derive the FX method, we follow the outline provided previously in the work of Hassefras et al. [14]. The FX method is based upon Huygens' principle, which states that the propagation of a wave within a medium can be comprehended through the collective contributions of all secondary sources distributed along its wavefront. This phenomenon is mathematically formalized in the Kirchhoff-Helmholtz integral. Based on this (for a detailed derivation, the reader is referred to [14]), we obtain the general formulation of the Rayleigh II integral, which is formulated as:

$$P(\vec{r}_A, \omega) = \int_S P(\vec{r}, \omega) \frac{\partial G}{\partial n} dS, \quad (2.2)$$

where $P(\vec{r}_A, \omega)$ is the Fourier transform of the pressure field, ω is the angular frequency, and G is the Greens's function. The anti-causal [42] Green's function G describes the wave field due to a point source excitation:

$$G = \frac{1}{4\pi\Delta r} \exp^{-j\frac{\omega}{c}\Delta r}, \quad (2.3)$$

where $\Delta r = |\vec{r} - \vec{r}_A|$ the distance between an observation and source location at A . Substituting Eq. 2.3 into 2.2 gives the Rayleigh II integral for wave field extrapolation:

$$P(\vec{r}_A, \omega) = \frac{z_A - z_0}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} P(\vec{r}, \omega) \frac{1 + j\frac{\omega}{c}\Delta r}{\Delta r^3} \exp^{-j\frac{\omega}{c}\Delta r} dx dy. \quad (2.4)$$

For each propagation step through the medium z_i , eq. 2.4 can now be rewritten as a convolution operation in Cartesian coordinates as [43]:

$$P(x, y, z_{i+1}, \omega) = W(x, y, z_{i+1}, z_i, \omega) * P(x, y, z_i, \omega). \quad (2.5)$$

Here, $*$ denotes the convolution operator and:

$$W(x, y, z_{i+1}, z_i, \omega) = \frac{\partial G(x, y, z_{i+1}|z = z_i, \omega)}{\partial z}, \quad (2.6)$$

where $W(z_{i+1}, z_i)$ is a propagation operator that describes propagation in the propagation direction z , from level z_i to level z_{i+1} .

The convolution in the space-frequency domain corresponds to a multiplication in the wavenumber-frequency domain given as:

$$\tilde{P}(k_x, k_y, z_{i+1}, \omega) = \tilde{W}(k_x, k_y, z_{i+1}, z_i, \omega) \tilde{P}(k_x, k_y, z_i, \omega), \quad (2.7)$$

where k_x and k_y are the lateral wavenumbers. The tilde symbol (\sim) indicates the wavenumber-frequency domain. For a homogeneous medium, the extrapolation operator \tilde{W} has the form [41]:

$$\tilde{W}(k_x, k_y, z_{i+1}, z_i, \omega) = \exp(-jk_z \Delta z), \quad (2.8)$$




Figure 2.5: Recursive Wavefield Extrapolation, as per [44]. A different operator is used to extrapolate the pressure field at every lateral position with varying velocities. [14]

with

$$k_z = \begin{cases} \sqrt{k^2 - (k_x^2 + k_y^2)}, & \text{for } k_x^2 + k_y^2 \leq k^2, \\ -j\sqrt{(k_x^2 + k_y^2) - k^2}, & \text{for } k_x^2 + k_y^2 > k^2, \end{cases} \quad (2.9)$$

Note, however, that this operator is only valid for a homogeneous medium, as local homogeneity is assumed for each grid point in this procedure. Furthermore, in a 2D medium, we find that $k_y = 0$. Figure 2.5 shows the basic concepts of a recursive wavefield extrapolation with space-variant spatial convolution operators. Using the whole length of the extrapolation operator (which theoretically goes to infinity) is computationally expensive. For that reason, the operator can be truncated to a finite number of grid points in the space domain, sacrificing resolution.

2.4. Full waveform inversion

Full-waveform inversion (FWI) in guided wave tomography is an advanced technique that reconstructs the spatial distribution of material properties or defect characteristics within a structure using the complete wavefield information. Unlike the ray tracing and diffraction methods for tomography, FWI leverages the entirety of the recorded wavefield, to iteratively refine the model until it accurately reproduces the observed data. An overview of a FWI scheme used in this thesis for a single frequency and single mode can be found in Figure 2.6. The reconstruction's multi-resolution characteristic is achieved by iterating over multiple discrete frequencies, starting with lower frequencies and progressively advancing to higher ones. After each iteration at a particular frequency, a new model parameter for the wall thickness is computed. It is set as the initial model for the subsequent iteration. The iterative process continues until a predefined accuracy or minimum update size is met, or the maximum iteration number is reached. In each iteration, a residual is calculated to set up an objective cost function of the model parameter m . More details of the inversion steps will be discussed in sections 2.5 and 2.6.

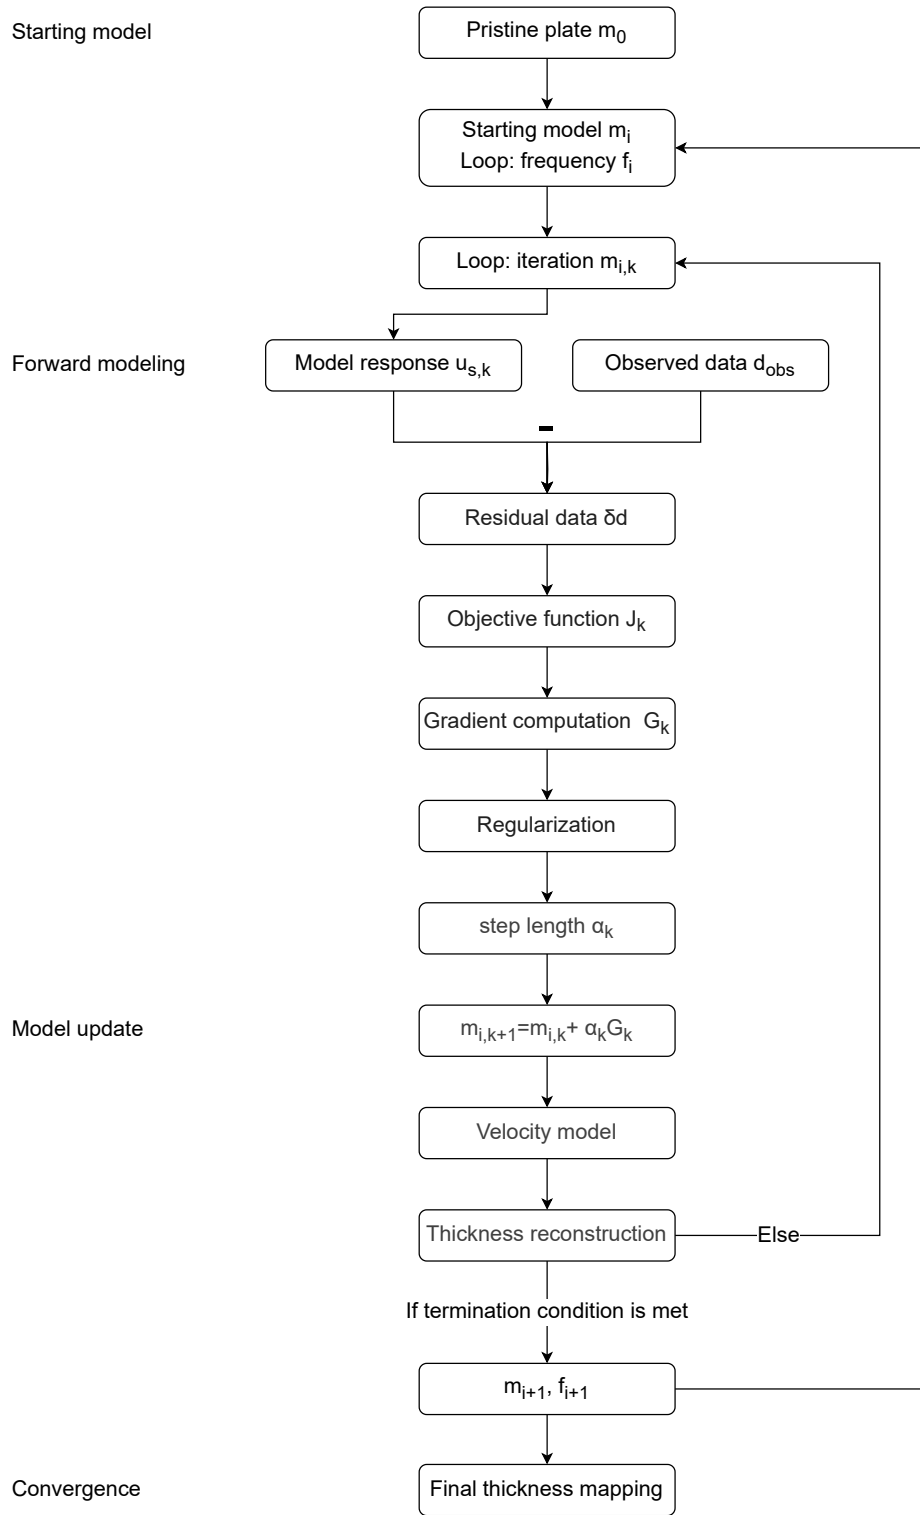


Figure 2.6: FWI scheme of a multi frequency and single mode inversion

2.5. Gradient optimization

The goal of using an FWI multi parameter simulation is to minimize the cost function in Equation 2.10, which the misfit between the observed data d_{obs} and synthetic data u_s . Here, m is the model parameter, d_{obs} is recorded at the receiver position r due to the source f_s , and u_s , a complex vector of N elements corresponding to the pressure field to the shot at s . The efficiency of the method greatly depends on the accuracy and efficiency of the computation of this gradient. A common way of defining the objective function is by using the least square norm, given as:

$$J(m) = \frac{1}{2} \sum_{\omega} \sum_{s,r} ||d_{obs}(x_r, x_s, \omega, m) u_s(x_r, x_s, \omega)||^2, \quad (2.10)$$

In this thesis, m is considered to be the wall thickness. x_s and x_r denote the source and receiver locations, respectively. The gradient of the functional can be obtained with the Fréchet derivatives of the state variables. The Fréchet derivatives are the derivatives of the state variables with respect to the model parameter. This results in the Jacobian or sensitivity matrix.

2.5.1. Perturbation method

A straightforward and well known approach for evaluating the sensitivity of the functional involves employing finite differences. Using forward differences of first order, one can approximate the gradient by:

$$\frac{\partial J}{\partial m} \approx \frac{J(m + \Delta m) - J(m)}{\Delta m},$$

Δm is a small step size. In the context of full waveform inversion, the quantity of design variables m can become very large, given that the amount of parameters depend on the mesh size. In that case, the gradient can be approximated as follows:

$$\frac{\partial J(m)}{\partial m_i} \approx \frac{J(m_1, \dots, m_i + \Delta m, m_{i+1}, \dots, m_N) - J(m_1, \dots, m_i, \dots, m_N)}{\Delta m}, i = 1, \dots, N \quad (2.11)$$

However, this would require N evaluations of J . As every evaluation of J corresponds to one additional solution of the (forward model) state equation, this clearly becomes impractical for large N . Hence, computing the Fréchet derivatives can be resource intensive [45, 46].

In order to reduce the computational load a series of points on the plate's surface can strategically be selected to serve as representative locations for the parameterization of corrosion defects or thickness reductions. These points are specifically chosen to cover regions where corrosion is most likely to occur, ensuring that the coverage sufficiently exceeds the spatial extent of the defects. These evaluation points are sampled coarser than the computational mesh for forward modelling and an interpolation step is performed from the parameterisation to the forward model mesh to obtain a final image. This strategy of a sub mesh of evaluation points to describe the geometric thickness is also referred to as geometrical full waveform inversion, as in [47].

2.5.2. Adjoint state method

If non-linear optimization methods such as for example the non-linear conjugate gradient method, are used, it may suffice to compute only the sensitivity of the objective function instead of calculating all the Fréchet derivatives. A method based on the adjoint state can be used to efficiently calculate the sensitivity field of the functional. The approach using the adjoint fields is independent of the number of model parameters, making it, in contrast to the perturbation method, well suited for inversion problems with a large amount of model parameters.

In order to derive this method we follow [48] and start by stating the Helmholtz equation here again:

$$\left[\nabla^2 + \frac{\omega^2}{v^2(x)} \right] u(x, x_s, \omega) = -f(\omega) \delta(x - x_s), \quad (2.12)$$

Which we can rewrite in the general simple form of

$$A(x, \omega, m) u(x, x_s, m) = f(\omega). \quad (2.13)$$

The gradient of the objective function can be found via the chain rule (for full derivation please see Appendix A):

$$\nabla_m J(x) = \frac{\delta J}{\delta m(x)} = \Re [\Gamma(x)^T \delta d^*(x, x_s, \omega)] . \quad (2.14)$$

From our objective function we find that $\delta d = d_{obs}(x_r, x_s, \omega) - u(x_r, x_s, \omega)$ which is the data misfit, \Re denotes the real part and Γ is the Fréchet derivative matrix given as:

$$\Gamma = \frac{\delta u}{\delta m_i(x)} . \quad (2.15)$$

Where δu represents the perturbation in the wavefield due to a small perturbation in the model parameters. Instead of directly calculating the Fréchet matrix we can define a more efficient solution by taking the derivative of equation 2.13 with respect to the model parameter m :

$$A(x, \omega, m) \frac{\delta u}{\delta m_i(x)} = - \frac{\delta A}{\delta m(x)} u_s(x, x_s, \omega), \quad (2.16)$$

which can be rewritten as:

$$\frac{\delta u}{\delta m_i(x)} = -A(x, \omega, m)^{-1} \frac{\delta A}{\delta m_i(x)} u_s(x, x_s, \omega). \quad (2.17)$$

We can now substitute equation 2.17 into equation 2.14 to obtain the final form for the gradient of our objective function:

$$\frac{\delta J}{\delta m} = \Re \left[-u_s^T(x, x_s, \omega) \frac{\delta A^T}{\delta m_i} A^{-1}(x, \omega, m) \delta d^* \right] . \quad (2.18)$$

Since we have an objective function on the receiver axis for every source position and every frequency, and given that in the construction of the adjoint sensitivity kernel phase velocity is used as model parameter we find that:

$$\frac{\delta A}{\delta m} = \frac{\delta \left[\nabla^2 + \frac{\omega^2}{v^2(x)} \right]}{\delta v(x)} = - \frac{2\omega^2}{v_p^3(x)}, \quad (2.19)$$

hence we find the final form for construction of the sensitivity kernel as:

$$\frac{\delta J}{\delta m} = 2\Re \left[\sum_{\omega} \sum_s u_s(x, x_s, \omega) \frac{\omega^2}{v_p^3(x)} \lambda_s^*(x, \omega) \right] . \quad (2.20)$$

$\lambda_s^*(x, \omega)$ denotes the adjoint field, which is back propagation of the misfit. So the source term is the the misfit between the simulated data u_s (which is a pristine plate for the first iteration) and the observed data u_{obs} at the receiver line, as can be seen in Figure 2.7. Here Figure 2.7c is the source perturbation (the misfit between Figure 2.7a and 2.7b) that is injected at the receiver line and back propagated in order to construct the adjoint field $\lambda_s(x, \omega)$ for a single source.

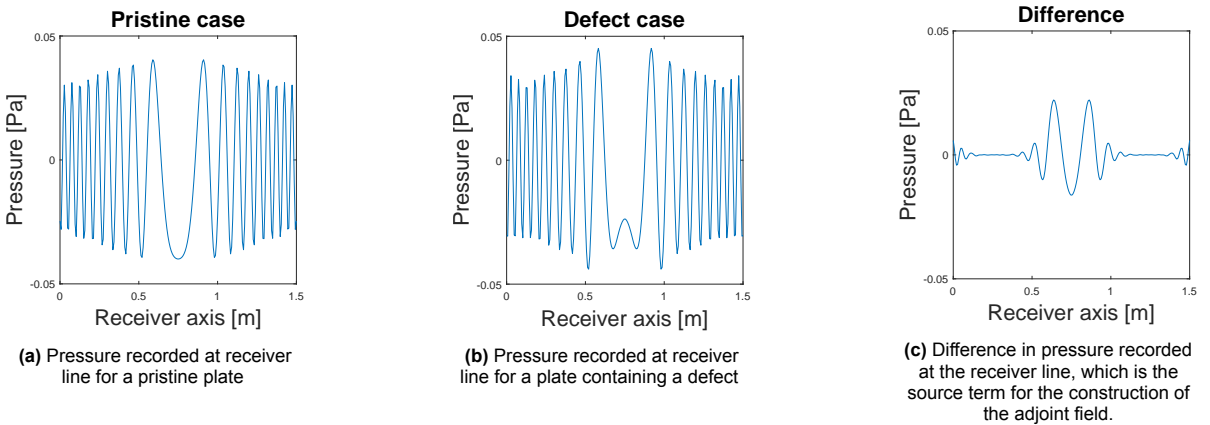


Figure 2.7: Comparison of pressure recordings for 2 different scenarios for a center located defect and $\pm 1N$ helical path

The construction of the adjoint sensitivity kernel for a single frequency is displayed in Figure 2.8. Observing the computation of cumulative interactions between forward and adjoint fields yields insight into what is commonly referred to as a sensitivity kernel. This kernel encapsulates the sensitivity of the cost function to perturbations in model parameters. Specifically, the adjoint field sets out the location of defects, serving as their secondary source. At the same time, the forward solution traverses through these defects, resulting in an interaction between the two wave fields at the perturbation location. Therefore, the adjoint field concentrates its energy predominantly at the site of the defect, as a peak in amplitude while diminishing elsewhere, as can be seen in Figure 2.8d.

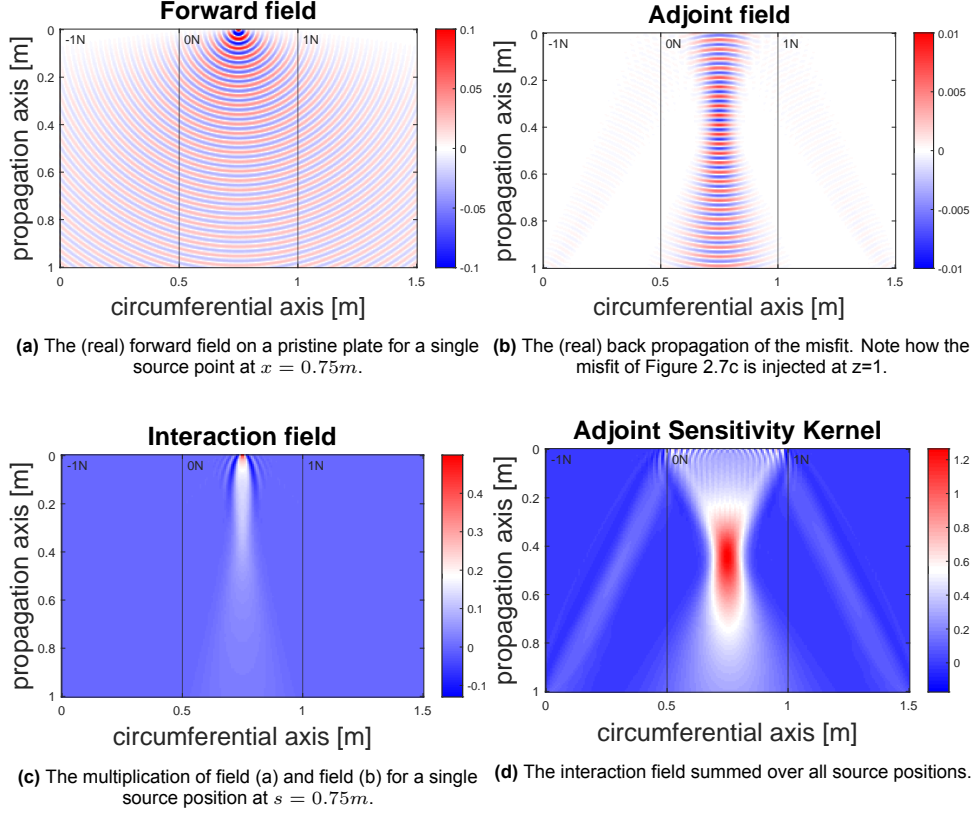


Figure 2.8: The construction of an adjoint sensitivity kernel for a single frequency using $\pm 1N$ helical paths.

2.5.3. Limited view

The number and spatial distribution of sources and receivers significantly influence the detection of the position and shape of an unknown defect [49]. Since the kernel is constructed as a sum over the number of sources and receivers, a higher transducer count enhances the solution's ability to distinguish the defect from its surroundings due to a stronger contrast visible in the sensitivity kernel. Additionally, strategically placing sources around the defect improves its shape description since shooting at the defect from different angles enriches the sensitivity kernel and enabling more precise characterization of the defect, which in turn reduces the ill posedness of the problem. However practical limitations, such as mainly the geometry of the pipe and cost considerations, impose constraints on sensor deployment.

However a limited number of sources on the circumferential axis can in many applications constrain the ability to effectively image defects. As briefly discussed in Section 1.1, helical paths are considered as a means to enhance coverage by interacting with the defect from multiple angles. By utilizing these paths, the defect is exposed to a wider range of orientations, increasing the shape reconstruction of the defect by the sensitivity kernel. This can be seen in Figure 2.9, where Figure 2.9a uses $\pm 1N$ helical paths and Figure 2.9b uses $\pm 3N$ helical paths for a centered Gaussian defect shape. As can be seen in the figure, using a single helical path makes the defect in the sensitivity kernel much more elongated while using 3 helical paths encapsulates the round shape better. This approach helps mitigate the challenges associated with ill posed optimization problems, as the increased number of angles improves

the stability and accuracy of the solutions. Consequently, the use of helical paths proves valuable in improving defect detection and characterization in scenarios with restricted source locations.

It is important to recognize that maximum interaction between wave fields does not necessarily occur at the defect location but rather at the primary source location, as can be seen in Figure 2.8c. This is due to the higher amplitude of the primary source compared to the secondary source. When the spatial sampling of the transducers is below the nyquist limit, notable aliasing contributions can be observed at the transducer locations. This will be explained in more detail in Section 3.1.3. Furthermore a trace of nonzero components from the source to receiver locations via the defect is created, which are unwanted as it does not physically represent wall thickness loss. These effects can be partly managed by applying a mask to suppress contributions near the transducer locations or by utilizing a Tukey window centered in the region of energy focus, as shown in Figure 2.9c.

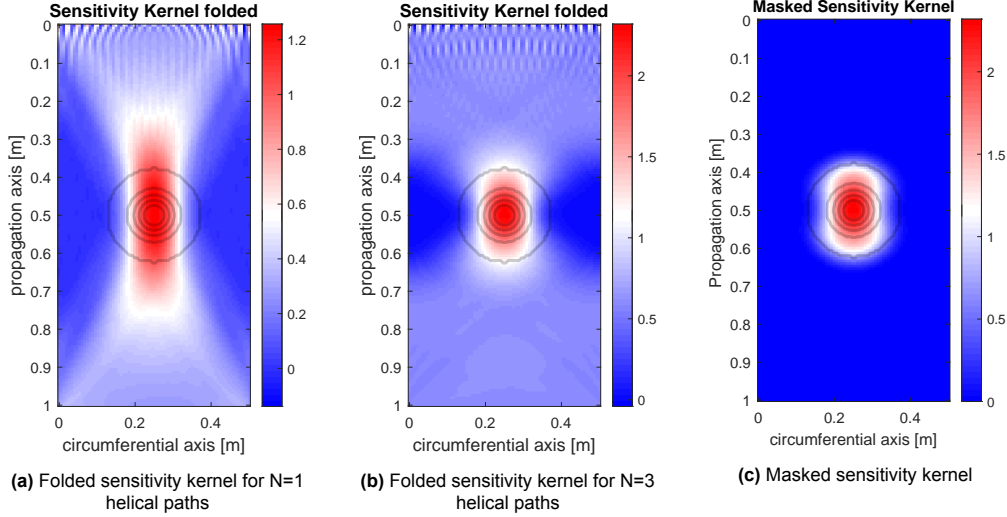


Figure 2.9: Sensitivity kernels for different helical path configurations and the application of a masking technique. The defect is displayed as contours.

2.6. Constrained non linear optimization methods

Unconstrained nonlinear optimization methods aim to find the minimum or maximum of an objective function without any explicit constraints on the variables. One of the most simple algorithm's is the Steepest Descent. The gradient of the objective function ∇J always points in the direction of the steepest ascent, perpendicular to the contour lines of the objective function. The parameter update is then given as:

$$m_{k+1} = m_k + \alpha d_k, \quad \text{with } d_k = -\nabla J(m_k) \quad (2.21)$$

Where m is the model parameter, and α the step size, often determined via a line search. Even though the negative gradient makes sure that the logarithm finds the steepest decrease of the objective function in a small region around the current solution, it is not necessarily the optimal choice globally since its always perpendicular to the contour lines. Therefore, convergence to a local minimum can require many iterations, as illustrated in Figure 2.10.




Figure 2.10: Illustration of the steepest descent method on an arbitrary objective landscape. The search direction follows the negative gradient, which is always perpendicular to the contour lines. This can lead to slow convergence, requiring many iterations to reach a local minimum.

In order to improve the convergence rate, Newton's method improves upon the steepest descent method by incorporating second-order curvature information, leading to faster and more stable convergence. For Newton's method the Hessian matrix H is required, which captures the local curvature of the function:

$$m_{k+1} = m_k - \alpha H^{-1} \nabla J(x_k). \quad (2.22)$$

By considering second-order information, Newton's method adapts the step direction more effectively, avoiding the zigzagging behavior of steepest descent and often achieving quadratic convergence near a local minimum. However, it requires computing and inverting the Hessian. The Hessian is given as

$$H = \Gamma_f(m)^T \Gamma_f(m) \quad (2.23)$$

Where $\mathcal{J}_f(m)$ is the Jacobian matrix. Since the Hessian consists of second-order partial derivatives, its computation requires calculating the Fréchet derivatives, making the method computationally expensive, particularly when applied in a large parameter optimization problem. Given that this thesis focuses on exploring optimization methods that use the adjoint state method and therefore do not rely on Fréchet derivatives, Newton's method is not considered in this thesis.

A commonly used approach that improves on the convergence rate of the steepest descent, but avoids the calculation of the Hessian, is the non linear Conjugate Gradient (CG) method. The CG method is based upon the steepest descent but extends it by ensuring that each successive search direction is not only a descent direction, but also conjugate to all previous directions. Therefore it avoids previously optimized directions, leading to more direct convergence to the minimum, as visualized in Figure 2.11.

The nonlinear conjugate gradient method is used in this thesis due to its ease of implementation and its ability to efficiently handle large parameter spaces without requiring explicit second-order information. Therefore in the following section the algorithm is further explored in more detail.




Figure 2.11: Illustration of the nonlinear Conjugate Gradient (CG) method on objective landscape. Unlike steepest descent, CG ensures that each search direction is conjugate to the previous ones, preventing redundant steps and improving convergence efficiency.

2.6.1. Conjugate Gradient

The outline of the algorithm can be found in the pseudo code 1. In the context of adjoint-based optimization, the search direction at each iteration is defined as the negative of the (adjoint) sensitivity kernel since the initial step the sensitivity kernel is directly used as gradient. Since the sensitivity kernel is a sum over the sources and scaled by the frequency squared, the amplitude is very much dependent on the frequency and amount of sources used in the system. This can lead to very high amplitudes in the sensitivity kernel that are multiple factors higher than the nominal thickness of 10 mm. Hence the sensitivity kernel is scaled to a range where the maximum of the kernel is within the bound of 1 and -1 mm, ensuring a numerical stable update. At iterations, the conjugacy coefficient β is used to ensure that the new search direction maintains conjugacy with previous directions.

Algorithm 1 Nonlinear Conjugate Gradient Method

```

1: Input: Initial guess  $x_0$ , minimal update amplitude  $\epsilon$ 
2: Compute gradient  $g_0 = \nabla J(m_0)$ 
3: Scale gradient:  $abs(g_0) < (1 \cdot 10^{-3})$ 
4: Set search direction  $d_0 = -g_0$ 
5:  $k \leftarrow 0$ 
6: while  $\alpha_k d_k > \epsilon$  do
7:   Initialise  $\alpha_{0,k}$ 
8:   Compute step size  $\alpha_k$  using line search &  $\alpha_{0,k}$ 
9:   Update solution:  $m_{k+1} = m_k + \alpha_k d_k$ 
10:  Compute new gradient:  $g_{k+1} = \nabla J(x_{k+1})$ 
11:  Scale gradient:  $g_{k+1} \leftarrow scale(g_{k+1})$ 
12:  Compute conjugate coefficient  $\beta_k$ 
13:  Update search direction:  $d_{k+1} = -g_{k+1} + \beta_k d_k$ 
14:   $k \leftarrow k + 1$ 
15: end while
16: Output: Approximate solution  $m_k$ 

```

An important component of the CG method is the line search, which determines the step size α with which the search direction is scaled, ensuring the prevention of large overshoots aswell as a faster convergence rate. The algorithm can be found pseudo code 2. α is computed using the strong Wolfe conditions, which ensure sufficient decrease and curvature of the objective function (line 5 and 10 in alg. 2 respectively). The zoom function, which can be found in pseudo code 3 is the part of the line search algorithm that refines α to satisfy the strong Wolfe conditions when an acceptable step has not been found in the main line search loop. Within the interval $[\alpha_{low}, \alpha_{high}]$ the function iteratively updates α and re-evaluates the objective function and gradient. Once the conditions are satisfied, the algorithm returns the refined step size α . For methods that do not generate well scaled search directions, it is important to use current information to make an initial guess for $\alpha_{0,k}$ [50]. Therefore the first-order change in the function at iterate m_k will be similar to that at the previous step. The initial step length at each iteration $\alpha_{0,k}$ is chosen such that:

$$\alpha_{0,k} = \begin{cases} 1, & \text{if iteration} = 1, \\ \alpha_{k-1} \frac{\nabla f_{k-1}^T d_{k-1}}{\nabla f_k^T d_k} & \text{otherwise} \end{cases}$$

For updating the conjugacy coefficient β , the Polak-Ribière (PR) method is a widely used approach in conjugate gradient optimization [50]. It refines earlier methods by incorporating additional gradient information, enhancing stability and performance, particularly in complex optimization landscapes. Unlike basic formulations that may struggle with stagnation in curved objective functions, the PR method adjusts β based on the difference between successive gradients. It is defined as:

$$\beta_{PR} = \frac{\nabla J_k^T (\nabla J_k - \nabla J_{k-1})}{\nabla J_{k-1}^T \nabla J_{k-1}} \quad (2.24)$$

This formulation allows the PR method to adapt more effectively to changes in the optimization landscape, making it particularly robust in ill-conditioned problems. Numerical experience suggests that it often outperforms other approaches, offering improved convergence behavior and efficiency in practical applications.

Algorithm 2 Line Search with Strong Wolfe Conditions

```

1: Input: Objective function  $J$ , gradient  $\nabla J$ , initial step size  $\alpha$ , direction  $d$ , parameters  $c_1, c_2$ 
2: Compute initial function value  $J_0 = J(m)$  and gradient  $g_0 = \nabla J(m)^T d$ 
3: for  $i = 1$  to max iterations do
4:   Evaluate  $J(m + \alpha d)$  and  $\nabla J(m + \alpha d)$ 
5:   if  $J(m + \alpha d) > J_0 + c_1 \alpha g_0$  or ( $i > 1$  and  $J(m + \alpha d) \geq J(x_{extprev})$ ) then
6:     Call Zoom( $\alpha_{extlow}, \alpha$ )
7:     Return  $\alpha$ 
8:   end if
9:   Compute  $g = \nabla J(m + \alpha d)^T d$ 
10:  if  $|g| \leq -c_2 g_0$  then
11:    Return  $\alpha$ 
12:  end if
13:  if  $g \geq 0$  then
14:    Call Zoom( $\alpha, \alpha_{extlow}$ )
15:    Return  $\alpha$ 
16:  end if
17:  Update  $\alpha_{extlow} = \alpha$ , increase  $\alpha$ 
18: end for
19: Return  $\alpha$  (warning: max iterations reached)

```

Algorithm 3 Zoom Function for Line Search

```

1: Input: Bracketing values  $\alpha_{extlow}, \alpha_{exthigh}$ , function  $J$ , gradient  $\nabla J$ , parameters  $c_1, c_2$ 
2: for  $j = 1$  to max iterations do
3:    $\alpha = 0.8\alpha$ 
4:   Evaluate  $J(m + \alpha d)$  and  $\nabla J(m + \alpha d)$ 
5:   if  $J(m + \alpha d) > J_0 + c_1 \alpha g_0$  or  $J(m + \alpha d) \geq J(x_{low})$  then
6:     Update  $\alpha_{high} = \alpha$ 
7:   else
8:     Compute  $g = \nabla J(m + \alpha d)^T d$ 
9:     if  $|g| \leq -c_2 g_0$  then
10:      Return  $\alpha$ 
11:    end if
12:    if  $g(\alpha_{high} - \alpha_{low}) \geq 0$  then
13:      Update  $\alpha_{high} = \alpha_{low}$ 
14:    end if
15:    Update  $\alpha_{low} = \alpha$ 
16:  end if
17: end for
18: Return  $\alpha$  (warning: zoom did not converge)

```

3

Numerical simulations

This chapter evaluates the performance of the nonlinear conjugate gradient optimization algorithm, in combination with the adjoint state method, as introduced in Section 2.6.1. The objective is to assess its effectiveness in full waveform inversion by examining key parameters, selecting an appropriate transducer configuration, and determining the most suitable modes and frequencies for the analysis.

The investigation begins with a simplified defect, modeled as a Gaussian shaped pit. Starting with a simplified defect allows for a clear and controlled validation of the algorithm, as the smooth and well defined nature of the Gaussian defect ensures that inversion errors can be more easily attributed to methodological performance rather than complex geometric effects and hence helps in systematically identifying and understanding the influence of different parameters in the simulation settings.

Following this, the analysis is extended to a more complex defect. Unlike the idealized Gaussian shape, real world corrosion induced defects often exhibit irregular geometries characterized by sharp edges, asymmetry, and rough surfaces due to the stochastic nature of material degradation. To better represent such complexities, an analysis of the so called 'Echo' defect [33] is introduced, which is reconstructed from laser scans of an actual corrosion defect provided by industry.

For the analyses in Sections 3.2 and 3.3, an inverse crime is intentionally carried out. This refers to the case where synthetic data used for inversion are generated using the same numerical model as the one employed in the inversion process. This approach provides an validation of the algorithm's consistency and accuracy. It enables a controlled assessment without uncertainties from model mismatches, isolates discrepancies to numerical accuracy and algorithm performance, and verifies the correctness of the adjoint state implementation. However, it does not reflect real world conditions, where discrepancies between the forward model and the physical system introduce additional challenges.

To bridge this gap, in Section 3.4, we introduce noise into the synthetic data. This step is important as real experimental measurements are inherently affected by noise due to sensor limitations, environmental influences, and numerical discretization errors. By incorporating noise, we assess the algorithm's robustness and its ability to reconstruct defects under more realistic conditions.

3.1. Model configuration

A well-defined model configuration is important for obtaining reliable and meaningful results in the experiments. This section details the key aspects of the model setup, including the selection of operating frequencies, the configuration of transducers, and strategies to mitigate non-uniqueness in the obtained solutions. Additionally, we outline performance indicators used to assess the model's effectiveness.

3.1.1. Plate and defect configuration

To test the algorithm, a 2D model of an unfolded steel pipe with a length L_{pipe} of 1 m, and a diameter D_{pipe} of 0.159 m is used. When unfolded, this diameter corresponds to a circumferential length of 0.5 m. Furthermore, 3 helical paths will be included in the simulation. The nominal thickness of the pipe is 10

mm. It must be noted that during manufacturing, slight variations in pipe or plate thickness can occur, with deviations of up to 10% [33] not being unusual. In real world scenarios these thickness variations will result in corresponding variations in wave velocity. Another factor that may influence wave velocity is material anisotropy, which can arise from uneven rolling during manufacturing. However in this thesis, in order to focus on the validation of the adjoint method methodology, the material is assumed to be isotropic with a nominal thickness.

To model a Gaussian-shaped defect in a plate, a spatially varying function is used to define the defect profile. The defect is centered in the plate. To validate the inversion methodology, an initial defect is introduced with a width w_{gauss} of 0.25 m and a depth d_{max} of 2.5 mm. The chosen dimensions ensure a clear defect response while remaining practical for analysis. A 0.25 m width spans half the circumferential width, allowing wave interactions to be observed in both defected and undisturbed regions. The 2.5 mm depth, moderate relative to the 10 mm plate thickness, is sufficient to affect wave propagation without introducing excessive distortions.

The width w_{gauss} of the Gaussian defect is determined using the relationship

$$w_{\text{gauss}} = \frac{w}{2\sqrt{-\log(2\sigma)}} \quad (3.1)$$

In this context, σ represents the standard deviation of the Gaussian function, which controls the spread of the defect, the factor $\sqrt{-\log(2\sigma)}$ ensures that the Gaussian function reaches 95.45% of its maximum at a distance of $w_{\text{gauss}}/2$ from the center. The defect profile is then described by

$$D(x, z) = 1 - d_{\text{max}} \exp\left(-\frac{2((z - z_0)^2 + (x - x_0)^2)}{w_{\text{gauss}}^2}\right) \frac{1}{D_{\text{norm}}} \quad (3.2)$$

where d_{max} is the maximum defect depth, and (x_0, z_0) represents the center position of the defect. The normalization factor D_{norm} ensures that the maximum defect depth is preserved, independent of the spatial distribution described by the Gaussian function. The exponential function is scaled correctly and is defined as:

$$D_{\text{norm}} = \max\left(\exp\left(-\frac{2((z - z_0)^2 + (x - x_0)^2)}{w_{\text{gauss}}^2}\right)\right). \quad (3.3)$$

The resulting plate can be seen in Figure 3.1a, where the defect is represented as a reduction in wall thickness, serving as the model parameter to be updated during the inversion process. Via the dispersion curve, a known frequency and wave mode, we can map the spatial defect to its associated phase velocity. This is illustrated using the S_0 mode at a frequency of 130 kHz, which can be seen in Figure 3.1b.




Figure 3.1: Visualization of a 2D Gaussian-shaped defect. (a) The defect is characterized as a reduction in wall thickness, serving as the model parameter to be updated in the inversion process. (b) Calculated phase velocity using the S_0 mode dispersion curve at 130 kHz.

In this chapter, the plate and defect configuration outlined in this section will be used as the default, and all simulations will be conducted with these settings in a noise free environment unless specified otherwise.

3.1.2. Frequency selection

Selecting an appropriate frequency (range) is critical for the effectiveness of the inversion. The choice of frequency influences both the sensitivity to defects and the practical limitations imposed by attenuation. This section explores the trade offs involved in frequency selection, and by analyzing the wavenumber domain representation we establish a strategy for optimizing frequency selection to enhance the defect imaging accuracy.

As discussed in Section 2.2 due to its low attenuation characteristics we initially consider the S_0 mode. The dispersion curve for this mode indicates that, for a 10 mm steel plate, frequencies within the range of approximately 100 kHz to 300 kHz are within the sensitive region for this mode. Intuitively one might suggest that picking the highest possible sensitivity is desirable which, as can be seen in Figure 3.2 can be achieved by picking high frequencies. Although higher frequencies, with their shorter wavelengths, offer finer spatial resolution, they also experience higher attenuation during propagation, as shown in Figure 2.4, which reduces their effectiveness over longer piping sections.




Figure 3.2: Frequencies and its associated location on the dispersion curve for 10 mm plate containing a defect with a maximum depth of 2.5 mm (red) and pristine case (blue). As the frequency increases, the phase velocity difference between the pristine and defective cases also increases, enhancing sensitivity.

To achieve effective imaging of a defect, it is crucial to consider the minimum required spatial resolution. In real world applications, the defect is typically unknown, making the required spatial resolution also uncertain. Therefore, a strategic approach is necessary, such as initially using low frequencies to locate the defect's general position, followed by iterative refinement with higher frequencies to resolve finer details.

We can determine the minimal spatial resolution by performing a 2D Fourier transform of the spatial defect into the wavenumber domain. The spatial resolution refers to the smallest detectable feature size and is inversely proportional to the wavenumber. Specifically, the spatial resolution Δ at a given wavenumber k is defined as:

$$\Delta = \frac{2\pi}{k} \quad (3.4)$$

where k represents the radial wavenumber. As the wavenumber increases, the resolution improves, allowing for the detection of finer spatial features in the original data.

In diffraction tomography, the maximum wavenumber components are constrained by the Ewald limiting circle in k space, where the zero frequency component is centered in the spectrum [33]. Low spatial frequency (or wavenumber) components are located at the center of k space, whereas high frequency components appear toward the outer regions. This is visualized in appendix B, where a defect is smoothed by filtering out the high frequency components in the wavenumber domain. As shown in Figure 3.3, the Gaussian defect exhibits smooth edges, resulting in the absence of high frequency components in the image.




Figure 3.3: wavenumber domain representation of the Gaussian defect.




Figure 3.4: Wavelength vs frequency plot for a plate with nominal thickness of 10 mm for the S_0 wave mode. A frequency of 130 kHz has a wavelength of 0.04 m.

As discussed in Section 1.2, the theoretical resolution limit is governed by the diffraction limit and given as $\Delta_{max} = \lambda/2$. However in realistic (elastic) modeling data the resolution limit drops to around 1.5 to 2λ [20]. Therefore, to make a conservative estimation, we find that the minimum required wavelength would be approximately $0.06/2 = 0.04$ m. From Figure 3.4, based on the dispersion curve's phase velocity relation with the wavelength $\lambda = v/f$, we find that this corresponds to a minimal frequency of 130 kHz.

Another consideration that needs to be taken into account is phase cycle skipping. The initial misfit plot, which compares the observed signal with the calculated signal on a pristine plate, reveals cycle skipping effects around 210 kHz, as shown in Figure 3.5. Cycle skipping is a common numerical challenge in guided wave tomography, where the observed and predicted waveforms differ by an integer multiple of the wave period. This arises because phase differences between the two signals are multiples of 2π , leading the inversion algorithm to converge to an incorrect solution. Although 210 kHz is not considered in this case due to its relatively high attenuation, cycle skipping can still occur at lower frequencies when

sensitivity is high, especially when the difference in wall thickness between the defect and the pristine case is very large. While various techniques exist to mitigate cycle skipping, as discussed in Section 2.2, this thesis focuses on selecting frequencies and defects that naturally avoid cycle skipping issues, reducing the need for additional computational complexity and overhead.

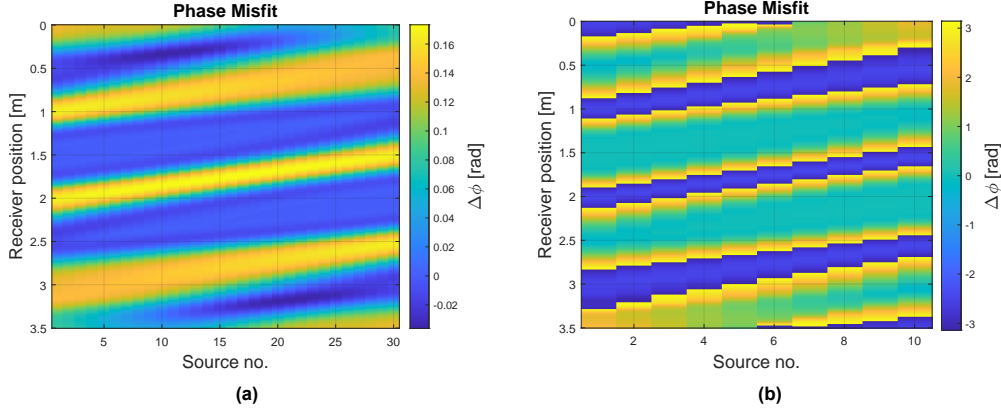


Figure 3.5: Phase misfit at varying frequencies: (a) 130 kHz and (b) 210 kHz, which illustrates the transition from π (yellow) to π (blue) known as cycle skipping. Each column represents the receiver axis for a single source. Due to the interaction with the defect, a phase difference is observed. As the source location moves from left to right along the plate, the corresponding arrival location on the receivers shifts accordingly.

We determined that each defect requires a minimum spatial resolution, which sets a lower bound on the frequency. In addition to spatial resolution, cycle skipping and attenuation need to be considered. There is however another factor to consider. One aspect that remains unaddressed but influences the choice of frequencies is the number of transducers, which will be discussed in the following section.

3.1.3. Transducer configuration

The positioning of sources and receivers plays a fundamental role in guided wave tomography, directly impacting the accuracy and resolution of defect imaging. Proper placement ensures that waves propagate through the structure in a way that maximizes coverage for tomographic reconstruction. The optimal configuration balances uniform spatial sampling around the defect, to minimize imaging artifacts and improve defect characterization. Unfortunately, as discussed in Section 2.5.3, guided wave tomography in pipeline inspection using only circumferential transducer rings suffers from a limited view problem, as sources and receivers can only be placed along the circumferential axis of the pipe. This constraint means that wave propagation is only observed from a restricted set of angles. In order to maximize the amount of angles the sources are always evenly distributed over the circumferential axis.

Sources

The number and distribution of sources directly affect the accuracy of the sensitivity kernel. If the number of sources is too low, resulting in large spacing between sources, aliasing effects can occur when performing the summation in Equation 2.20. This can introduce artificial artifacts in the sensitivity kernel, as can be seen in Figure 3.7a. These aliasing effects occur due to the way the sensitivity kernel is constructed. Recall from Section 2.5.2 that for each source position we calculate the interaction field, which is the cross correlation of the adjoint field and the forward field. The interaction field is depicted again in Figure 3.6a.

With sufficiently small source spacing, the combined interaction fields of all sources cause destructive interference at the localized high-amplitude regions. Thereby not interfering with the focusing around the defect. On the other hand, when the source spacing is too large, the interaction fields do not have enough coverage to properly cancel out these high amplitude regions near the sources. This is illustrated in Figures 3.6b and 3.6c. In both cases, 10 sources are used, one set placed centrally with sufficiently small spacing, which avoids aliasing, and another set evenly distributed, which results in aliasing artifacts. In this scenario with 10 sources, it is important to note that although aliasing artifacts are present, the focusing effect created by the evenly spaced sources significantly outweighs the placing the sources sufficiently close to each other, preventing aliasing but losing the focusing effect.

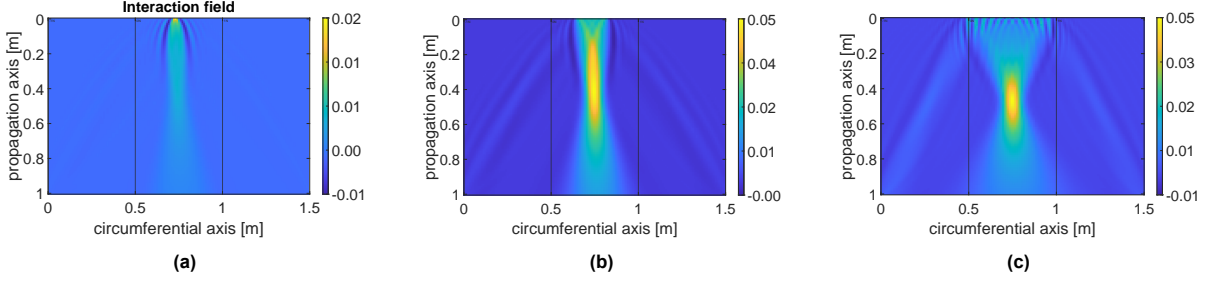


Figure 3.6: (a) Interaction field. (b) Kernel with 10 sources that are sufficiently spaced in the center, reducing aliasing effects but losing the focusing effect accordingly. (c) Kernel with 10 sources that are evenly spaced at the circumferential axis, inducing aliasing effects but keeping the focusing effect.

In order to minimize the effects the spatial sampling of the sources should be at least below the Nyquist limit of $\lambda/2$, as demonstrated in Figure 3.7a-c, where the aliasing artifacts become visible for a source spacing of λ , and disappear for a source spacing meeting the Nyquist limit. These aliasing artifacts act as noise, distorting the kernel, which can be seen in Figure 3.7d, which shows the aliasing contribution to the initial update step.

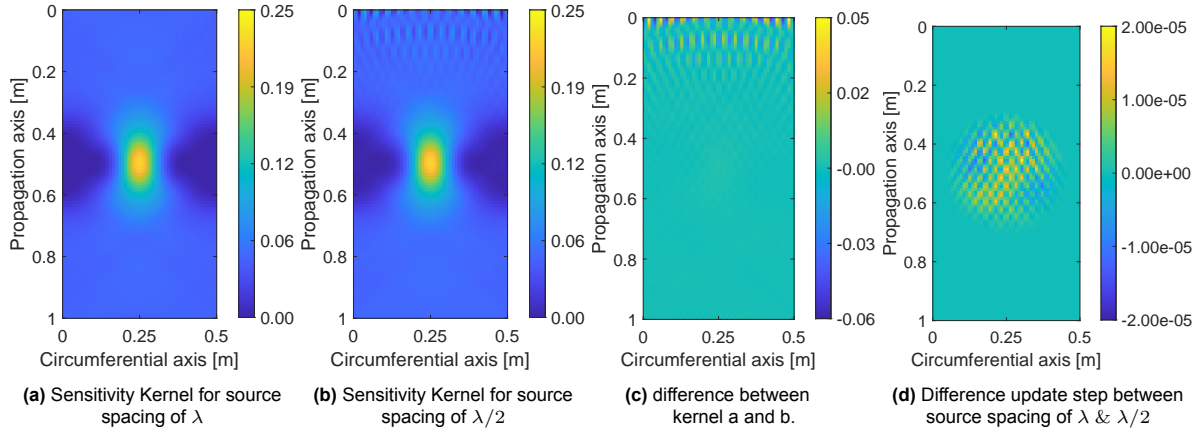



Figure 3.7: The effect of source spacing on aliasing in the adjoint sensitivity kernel and the influence of aliasing on the update step (αd_k)


Receivers

Not only the number of sources play an important role, but also the number of receivers determine the accuracy of the sensitivity kernel. When a sufficiently dense receiver array is used, the misfit is captured with adequate spatial resolution, ensuring a better reconstruction of the adjoint field, which is constructed from the back propagated misfit. In this case, the aliasing effects are mitigated because the high frequency components introduced by undersampling are minimized.

When the receiver spacing exceeds the Nyquist limit, the back propagated misfit exhibits artificial oscillations, as can be seen in Figure 3.8. Similar to insufficient source spacing, this effect in turn introduce high frequency noise into the sensitivity kernel, as can be seen in Figure 3.9. To minimize these aliasing effects, the spatial sampling of the receivers must be atleast above the Nyquist limit of $\lambda/2$.



(a) Adjoint field for a single source and receiver spacing of λ .



(b) Adjoint field for a single source and receiver spacing of $\lambda/2$.

Figure 3.8: The adjoint field for a single source positioned at the center of the circumference. It becomes evident that when the receiver spacing falls below the Nyquist criterion for spatial sampling, aliasing artifacts emerge in the adjoint field.

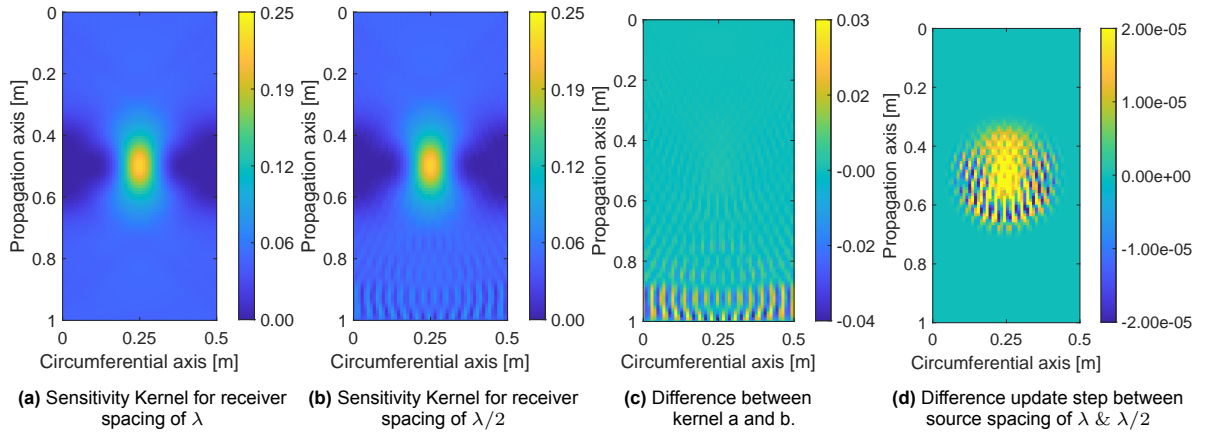


Figure 3.9: The effect of receiver spacing on aliasing in the adjoint sensitivity kernel and the influence of aliasing on the update step

At the TNO lab and in most real world scenarios, the maximum number of physical transducers that can fit around a pipe with circumferential of 0.5 meter is around 24. This means that for our default case the distance between each source is $0.5\text{m}/24 = 0.021\text{m}$. Consequently, the maximum frequency that can be used is the frequency that has the minimum wavelength λ_{min} of 0.042 mm. Based on the dispersion curve of the S_0 mode, for a 10 mm thick plate the maximum frequency that can be used without any aliasing effects occurring is slightly above 123 kHz.

Although aliasing effects reduce the accuracy of the sensitivity kernel, they do not necessarily cause instabilities in the inversion. Instead, they primarily degrade the overall performance of the algorithm capability to have a proper shape estimation. The impact of aliasing on the algorithm will be further examined in Section 3.2.1. Furthermore in Sections 3.2.4 and 3.3.4 parameter sweeps on the number of sources and receivers will be conducted.

Unless stated otherwise, the default for the simulations in this chapter use 24 sources and 24 receivers. The sources are located at the position where the length of the pipe $z=0$, and the receivers at $z=L$, both placed on the circumferential axis, as can be seen Figure 3.10, in which 12 sources and receivers are plotted for visualization purpose.




Figure 3.10: Default transducer layout. For better visualization 12 sources and 12 receivers are used in this figure.

3.1.4. Reducing non-uniqueness

Accurate reconstruction of wall thickness variations is challenging due to the inherent non-uniqueness of the inverse problem. One key issue suffering from this is model parameter trade off, where errors in one region of the model are compensated for by errors in another. In our context of wall thickness inversion, this means the algorithm might artificially add wall thickness in one area while instead of removing wall thickness in another, leading to a solution that still satisfies the phase match, but does not accurately reflect the true defect geometry. To mitigate this, additional independent constraints are required. Two possible strategies for improving inversion accuracy are the use of multiple frequencies and multiple wave modes.

Multiple frequencies

Although the spatial resolution at 130 kHz might be sufficient for detecting and estimating the overall shape of a defect, certain fine details, such as the minimal depth error, may still be inaccurate due to a lack of sensitivity and constraints. This can be particularly problematic when using a single frequency, as phase shifts at a given frequency can not provide enough independent constraints to distinguish between different thickness distributions. By using multiple frequencies, additional independent constraints are introduced into the inversion process, as visualized in Figure 3.11 and Figure 3.12. Here the increased sensitivity of the 150 kHz compared to the 130 kHz, next to an increase in amplitude, becomes visible as at distances far away from the source (top right for source = 72) the defect artifact is better defined compared to the 130 kHz case. Note that the misfit for each frequency is scaled to the same order as the maximum frequency to ensure that constraints imposed by lower frequencies remain sufficiently influential. The inversion is performed in a sequential manner, starting with a single low frequency and gradually adding higher frequencies in subsequent iterations.

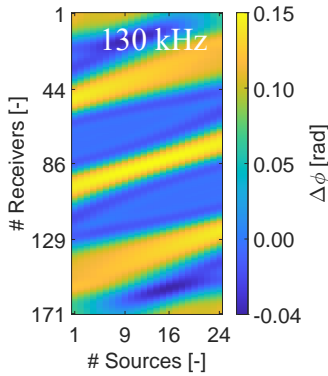


Figure 3.11: Misfit using a single frequency.




Figure 3.12: Misfit using multiple frequencies, further constraining the simulation.

Multiple wave modes

Incorporating multiple wave modes, such as the A_0 or SH_1 modes can further constrain the inversion, similar to using multiple frequencies. Different modes interact uniquely with the defect, possibly providing complementary information that a single mode alone cannot capture.

For example, the A_0 mode is highly sensitive to thickness variations at low frequencies and higher order shear horizontal modes like SH_1 exhibit low attenuation and strong sensitivity, making them advantageous in certain conditions. Remember however that the A_0 mode suffers from high attenuation in liquid filled pipelines, limiting its use to gas filled pipes. Different wave modes fundamentally differ in their interaction with defects. For example, in contrast to the S_0 mode, for the A_0 mode the phase velocity goes down with a decrease of wall thickness, resulting in negative phase shift, as can be seen in Figure 3.13. This could offer greater independence between constraints that reduces the non-uniqueness of the problem, and thereby attribute to a better inversion result.

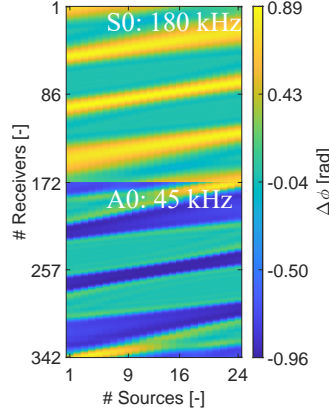


Figure 3.13: Misfit using 2 both S_0 and A_0 modes in a single inversion.

3.1.5. Performance indication

To assess the performance of the inversion result, two metrics were introduced [47]. The first error metric e_1 , referred to as shape error, evaluates the quality of the shape reconstruction.

The shape error is calculated as:

$$e_1 = \sqrt{\frac{\int \int [T(x, z) - T^*(x, z)]^2 W(x, z) dx dz}{T_0^2 \int \int W(x, z) dx dz}}, \quad (3.5)$$

where T_0 represents the nominal thickness, which is 10 mm for this thesis, $T(x)$ denotes the thickness map of the true defect model, and $T^*(x)$ refers to the reconstructed thickness map. The parameter W is a windowing function centered around the defect, with a radius equal to the defect width. This ensures that the evaluation focuses on the defected region and avoids misleading results caused by unaffected areas of the plate.

A lower value of e_1 indicates that the reconstructed defect shape better matches the true defect shape, meaning a more accurate inversion result with $e_1 = 0$ being a perfect reconstruction. Conversely, a higher value of e_1 suggests a larger deviation from the true shape, implying that the reconstruction does not accurately capture the defect geometry.

Since a pipeline is only as strong as its weakest point, small pit-like corrosion defects can be critical. Therefore, one of the key parameters of interest is the deepest point of the defect. To evaluate the deepest point a second error metric e_2 is introduced, which measures the accuracy of the minimum depth error, hence referred to as depth error.

The depth error is calculated as:

$$e_2 = \frac{|T_m - T_m^*|}{T_0}. \quad (3.6)$$

where T_m and T_m^* represent the minimum thickness of the true model and the synthetic model, respectively.

A lower value of e_2 means that the reconstructed minimum defect depth is close to the actual minimum defect depth, indicating high accuracy in capturing the deepest point of the defect. In contrast, a higher value of e_2 means that the reconstructed defect depth deviates from the true defect depth.

3.2. Inversion results for the Gaussian defect

A comparison of three single-frequency inversions at 130, 150, and 180 kHz is conducted to analyze the impact of different frequencies. The defect is discretized by the wavelength of the frequency that is used. In order to make a fair comparison for each frequency a mesh is used with the size of $\lambda_{high}/4$, where λ_{high} is the wavelength of the highest frequency. Following this, an inversion incorporating multiple frequencies is conducted to evaluate the effect of further constraining the problem through the use of multiple frequencies. Furthermore an inversion is done using both the S_0 and the A_0 modes, evaluating the influence of multimodal wave propagation on the inversion process. Finally, 2 parameter sweeps are conducted to explore the algorithm's sensitivity and robustness to different experimental conditions. One sweep varies the number of transducers to analyze the consequences of aliasing on shape reconstruction, while another investigates the effect of different defect configurations, providing insights into the algorithm's robustness across diverse defect geometries.

3.2.1. Single frequency inversion

Figure 3.14a shows the initial phase misfit and sensitivity kernel for the 130 kHz inversion.

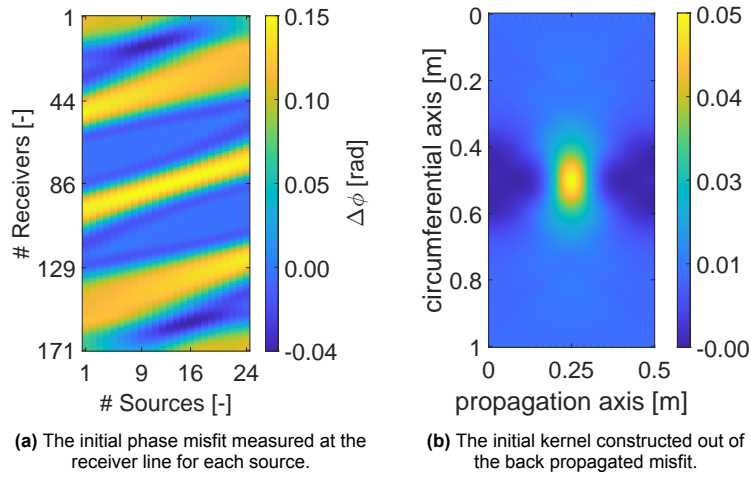


Figure 3.14: Wall thickness, phase misfit and adjoint sensitivity kernel on iteration 0

The thickness reconstruction is shown in Figure 3.15b. The imaging process reveals a clear presence of the defect, with the circumference and shape reconstruction being well-defined in axial direction. However due to the inherent limitations in the field of view, the resulting image shows elongation of the defect, which is consistent with expectations based on the sensitivity kernel which also has the shaded side areas with limited coverage. An alternating pattern of positive and negative errors can be seen in Figure 3.15c. The alternating pattern arises because the algorithm adds and removes wall thickness in unintended locations. However, due to the non-uniqueness of the problem, these changes still lead to a reduction in the misfit, allowing the pattern to be preserved. Lastly, the Tukey window also creates artificial wall thickness variations at its borders due to the gradient of the Tukey window. Both effects cause the misfit to be illegitimately reduced, which in turn leads to depth underestimation.

The convergence plot, which can be found in Figure 3.16, demonstrates a relatively stable and well-behaved optimization process, with the misfit ratio progressively decreasing towards a solution. Notably, at the second iteration the misfit does not decrease. This can be attributed to an imposed constraint on the maximum allowable change in α , which is implemented to prevent excessively large updates that lead to numerical instability or divergence, particularly in the early iterations when the algorithm is gradients are very strong.




Figure 3.15: Inversion result for 130 kHz. The colorbars indicate wall thickness in mm.




Figure 3.16: Convergence plot of the inversion where χ_0/χ_i is the misfit ratio between the misfit iteration and the initial misfit calculated in the pristine case. χ_i is the root mean square (RMS) of the misfit vector.

In an attempt to improve the performance, the frequency is increased to 150 and 180 kHz. The resulting thickness reconstructions are shown in Figure 3.17 and Figure 3.18. Higher frequencies have shorter wavelengths, allowing them to detect finer details of the defect. As a result, the slopes near the defect edges are expected to be better captured, slightly reducing the interchange of wall thickness and minimizing the alternating error pattern. However, when looking at the reconstructed thickness results for 150 and 180 kHz, shown in Figure 3.17 and Figure 3.18, we observe that while both frequencies improve over 130 kHz in terms of error metrics, the inversion at 180 kHz shows a slight increase in error over 150 kHz.




Figure 3.17: Inversion result for 150 kHz. The colorbars indicate wall thickness in mm.




Figure 3.18: Inversion result for 180 kHz. The colorbars indicate wall thickness in mm.

Since we are operating above the Nyquist limit for the transducer spacing, the contribution of aliasing is inspected for the 180 kHz inversion. Figure 3.19 compares the first inversion update with an inversion that uses sufficient transducer sampling. Since the update step is in the order of 10^{-3} mm, the aliasing contribution to wall loss is negligible. Furthermore, Figures 3.20 and 3.21 show the sensitivity kernel in the wavenumber domain before regularization at the first and final iteration. From these plots we find that in the first iteration the update step is dominated by low frequency components, corresponding to an initial more coarse reconstruction of the defect. Although at around $k_x = \pm 50$ the aliasing introduces some high frequency noise components, they are considered to be negligible. Through the iterations we find that the aliasing effects slightly increase but stay negligible. In the final iteration, the sensitivity kernel captures some high-frequency components that are near $\lambda = 27$ mm, which is the wavelength for 180 kHz. However, low-frequency components continue to dominate even at this stage of the inversion, suggesting that the advantage of using 180 kHz for a smooth, noise free, defect remains limited.




Figure 3.19: Difference between the initial update step for a non aliased inversion and the inversion case of Figure 3.18. Colorbars indicate wall thickness in mm.




Figure 3.20: Wavenumber domain representation of the Kernel in the first iteration for 180 kHz inversion with a transducer sampling of λ .




Figure 3.21: Wavenumber domain representation of the Kernel in the last iteration for 180 kHz inversion with a transducer sampling of λ .

In order to get a better understanding of the influence of the frequency on the behavior of the algorithm,

inversions are done for both S_0 and A_0 in a wavelength range of 25 to 40 cm, corresponding to a frequency range of 125 to 190 kHz for the S_0 mode, and 45 to 87 kHz for the A_0 mode. The corresponding error metrics are presented in Figures 3.22a and 3.22b. However, these figures do not indicate a clear correlation between frequency and inversion performance.

Numerical experiments reveal that the inversion outcome is sensitive to the scaling of the stepsize α . Since the gaussian defects exhibit smooth, well-defined gradients, the optimization takes a highly effective first step. This rapid initial improvement causes the residual gradient norm $\|g\|$ to shrink significantly, leading to a small denominator in the step size update formula and, consequently, a much smaller α in subsequent iterations. As a result, the optimization slows down prematurely. Therefore for each iteration the growth of the initial $\alpha_{0,k}$ is restricted to a maximum of 20%. When varying this grow restriction, from which the results can be found in Appendix C, variations in step sizes lead to differences in the inversions outcome, indicating that there is no single minima.

Another possible explanation is that for all inversions conducted in this section a noise free case is considered. The absence of noise and smooth defect configuration could make the problem similarly posed for each frequency, and since the differences in errors across frequencies are small, the inversion converges near the same minimum, with slight variations possibly due to the line search. If noise were introduced into the system, we expect increased variability in the inversion results since high frequencies can handle noise better. In section 3.4, further investigation is conducted to draw conclusions on the influence of noise on the inversion process.

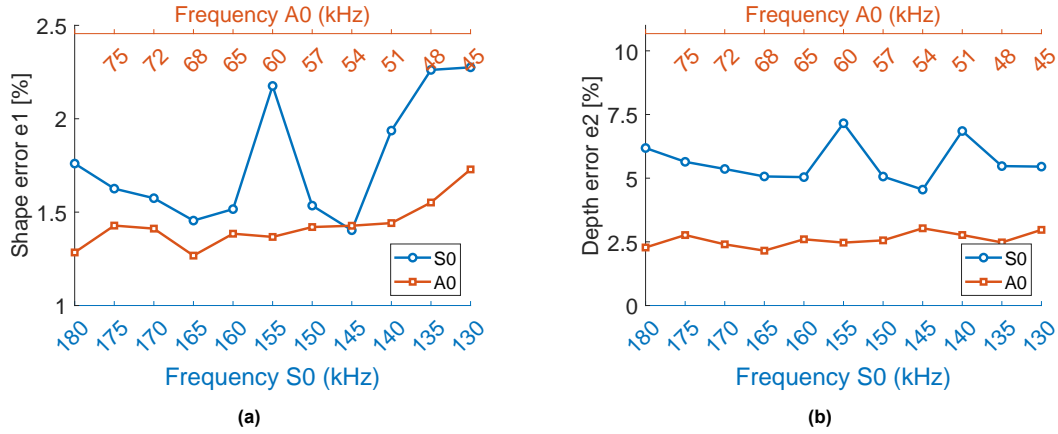


Figure 3.22: (a) Shape error for multiple wavelengths evaluated as single frequency inversions. (b) Depth error for multiple wavelengths evaluated as single frequency inversions.

3.2.2. Multi frequency inversion

From the single frequency inversions we find that increasing the frequency does not directly lead to better inversion performance. However as discussed in section 3.1.4 we aim to further improve the performance by constraining the problem using multiple frequencies that are within the usable frequency range. An inversion is initially conducted using frequency continuation with 130 and 150 kHz (130/150 kHz). The resulting thickness reconstruction is shown in Figure 3.23. It is found that both the shape and minimum depth error decreased relative to both the 130 and 150 kHz single frequency inversion. The shape error compared to the 130 kHz inversion improved from 2.28% to 1.68% by adding 150 kHz consecutive to the inversion. When comparing Figure 3.15c and Figure 3.23c, the increase in performance is mainly achieved by the reduction of the alternating error pattern, indicating that in the algorithm the extra constraint misfit does not allow the interchange of wall thickness.

Additionally, an inversion is performed using 150/180 kHz, as well as one using 130/180 kHz. The inversion results can be found in Figure 3.23 and Figure 3.25. When combining multiple frequencies, it is observed that selecting frequencies that are as far apart as possible generally leads to better performance. This is because each frequency constrains the problem in a different way, capturing complementary information about the defect. Therefore a wider separation between frequencies ensures a more diverse set of constraints, improving the overall inversion accuracy.




Figure 3.23: Inversion result for multi frequency inversion using 130/150 kHz. The colorbars indicate wall thickness in mm.




Figure 3.24: Inversion result for multi frequency inversion using 150/180 kHz. The colorbars indicate wall thickness in mm.




Figure 3.25: Inversion result for multi frequency inversion using 130/180 kHz. The colorbars indicate wall thickness in mm.

To further investigate whether the addition of more frequencies improves the inversion performance, we extend the analysis by incorporating a third frequency. A new inversion is conducted with three frequencies: 130/150/180 kHz. The inversion result can be seen in Figure 3.26.

When comparing the results in table 3.1, compared to the inversions performed with 130/150 kHz and 150/180 kHz, we observe an improvement in performance using a third frequency. However, when compared to the 130/180 kHz inversion, the performance goes down. Examining the convergence plot for the triple frequency inversion in Figure 3.27, we see that for the first two frequencies, there is a significant reduction in the misfit. Here, it is important to note that for each inversion, the misfit ratio is relative to its first iteration, causing it to reset to 1 at the first iteration for each frequency. The use of the third frequency, 180 kHz, does not lead to a significant reduction in the misfit. Since at this stage of the inversion the defect is already relatively well defined by the previous two frequencies, the sensitivity is not large enough to make proper distinctions in phase velocity. As a result, the overall performance does not surpass that of the 130/180 kHz inversion, where 180 kHz still retains sufficient sensitivity after the 130 kHz inversion to meaningfully contribute to the defect imaging. This suggests that simply adding more frequencies does not always enhance inversion performance. Rather, the chosen frequencies must effectively complement each other to provide meaningful constraints. It should be noted that for the Gaussian defect considered in this study, the lower frequencies already capture the defect configuration relatively well. For more complex defect geometries, introducing a third frequency might be more significant for the inversion quality.




Figure 3.26: Inversion result for multi frequency inversion using 130/150/180 kHz. The colorbars indicate wall thickness in mm.

f [kHz]	e_1 [%]	e_2 [%]
130	2.28	5.45
150	1.74	6.15
180	1.79	6.41
130/150	1.68	5.55
150/180	1.50	4.79
130/180	1.39	4.25
130/150/180	1.46	4.84

Table 3.1: Error metrics e_1 and e_2 for different single and multi-frequency inversions.




Figure 3.27: Convergence plot for an inversion using 130/150/180 kHz. Note that the misfit ratio is depicted on the y-axis. The misfit is reset to 1 at each first iteration of each frequency.

3.2.3. Multi mode inversion

To further improve inversion performance, an alternative approach is to incorporate multiple wave modes. The sensitivity kernels for the S_0 and A_0 modes in Figures 3.28a and 3.28b reveal how each mode responds uniquely to the defect. In order to make sure that both sensitivity kernels have sufficiently influence they are scaled to the same magnitude. To obtain the final sensitivity kernel, both mode kernels are summed, resulting in the final sensitivity kernel, as visualized in Figure 3.28c.

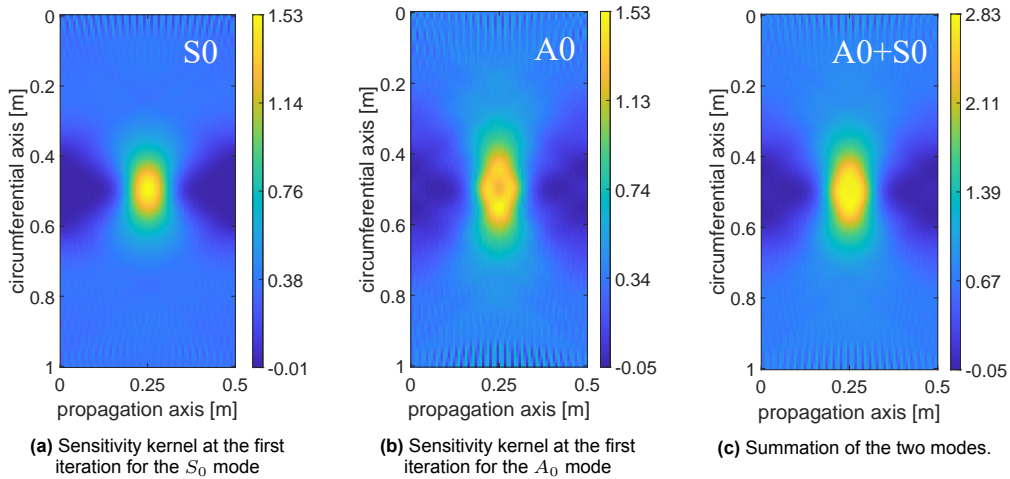


Figure 3.28: Construction of the sensitivity kernel for multiple modes.

The effectiveness of this approach is evaluated by testing different frequency pairings for the two modes using wavelengths that are similar to the multi frequency simulations, as can be seen in Table 3.2.

	S_0 mode	A_0 mode
$\lambda = 39.7 \text{ mm}$	130 kHz	45 kHz
$\lambda = 33.8 \text{ mm}$	150 kHz	58 kHz
$\lambda = 27.4 \text{ mm}$	180 kHz	79 kHz

Table 3.2: Wavelength and their corresponding frequencies for S_0 and A_0 modes.

From Tables 3.3 and 3.4 we find that, in contrast to the multi frequency simulation, higher-frequency components for the A_0 and S_0 modes improve inversion accuracy since the lowest errors are achieved with both highest frequencies.

	$f_{A_0} = 45 \text{ kHz}$	$f_{A_0} = 79 \text{ kHz}$		$f_{A_0} = 45 \text{ kHz}$	$f_{A_0} = 79 \text{ kHz}$
$f_{S_0} = 130 \text{ kHz}$	1.47	1.36	$f_{S_0} = 130 \text{ kHz}$	2.70	2.54
$f_{S_0} = 180 \text{ kHz}$	1.41	1.20	$f_{S_0} = 180 \text{ kHz}$	2.62	1.96

Table 3.3: Shape error in % for multi-modal inversion using both S_0 and A_0 modes.

Table 3.4: Depth error in % for multi-modal inversion using both S_0 and A_0 modes.

Unlike multi-frequency inversion where maximizing frequency separation reduced the ill posedness, the multi-mode inversion already enforces different constraints due to the fundamental differences between the wave modes. Therefore, maximizing the frequency spacing between modes does not necessarily yield the best results. Instead, using the highest available frequencies for both modes provides the greatest spatial resolution in each case, ultimately leading to better defect imaging.

In Figures 3.29c and 3.30c, which compare the inversion results to the true profile for $f_{S_0} = 180 \text{ kHz}$ combined with $f_{A_0} = 45 \text{ kHz}$ and $f_{A_0} = 79 \text{ kHz}$, it is observed that variations in wall thickness are reduced. This reduction leads to a decrease in shape error.




Figure 3.29: Inversion result for multi mode inversion using 45 kHz for the A_0 mode and 180 kHz for the S_0 mode. The colorbars indicate wall thickness in mm.




Figure 3.30: Inversion result for multi mode inversion using 78 kHz for the A_0 mode and 180 kHz for the S_0 mode. The colorbars indicate wall thickness in mm.

3.2.4. Parametric study

To systematically assess the performance and limitations of the inversion algorithm, a parametric study is conducted. This study investigates the influence of key parameters on the defect reconstruction process, focusing on variations in defect geometry and transducer configurations.

Different defect configurations

A parameter sweep is conducted using the default configuration settings and a frequency continuation with 130 kHz and 180 kHz. This parametric study varies the defect configuration by adjusting its width and depth in each inversion.

The defect width is tested at 4 different values: at the diffraction limit of half a wavelength, a single wavelength, which is below the realistic minimum resolution for elastic data, two wavelengths, which is slightly above the realistic minimum resolution for elastic data, and lastly a reasonably large defect of 250 mm is considered, representing 50 percent of the plate's circumference. Similarly, three different defect depths are considered, ranging from a shallow defect of 10% to a moderate depth of 25% and a critical depth of 50%. Since defects deeper than 50% are already considered highly critical, deeper depths are not investigated.

The results of the parameter sweep in Figure 3.32 demonstrates that the algorithm can effectively estimate defect shapes across various defect sizes. However for depth estimation we find that an increase in depth leads to higher errors. When a defect is very shallow, there is minimal shift along the frequency-thickness axis on the dispersion curves, as shown in Figure 3.31a. Therefore, detecting shallow defects requires selecting a sensitive region of the dispersion curve. Conversely, for very deep defects, as illustrated in Figure 3.31c, the phase velocity shift occurs in the less steep portion of the dispersion curve. As a result, the system's sensitivity decreases, making it difficult to distinguish between variations in this depth range. This increases the likelihood that the inversion algorithm accepts multiple solutions. Consequently, when inspecting deep defects using the S_0 mode, high frequencies, or the use of a more sensitive mode, are necessary to ensure sufficient sensitivity. Furthermore the results show that increasing the defect width leads to a reduction in shape error. This is expected because, for a Gaussian defect with a fixed depth, a greater width results in a shallower wall loss gradient. As a result, there are more points of contact for the propagating wave to interact with the defect, leading to more phase shift, and thereby improving the accuracy of the reconstructed defect shape.




Figure 3.31: Frequencies and its associated location on the dispersion curve for 10mm plate containing a defect (red) and pristine case (blue). As the frequency increases, the phase velocity difference between the pristine and defective cases also increases, enhancing sensitivity.

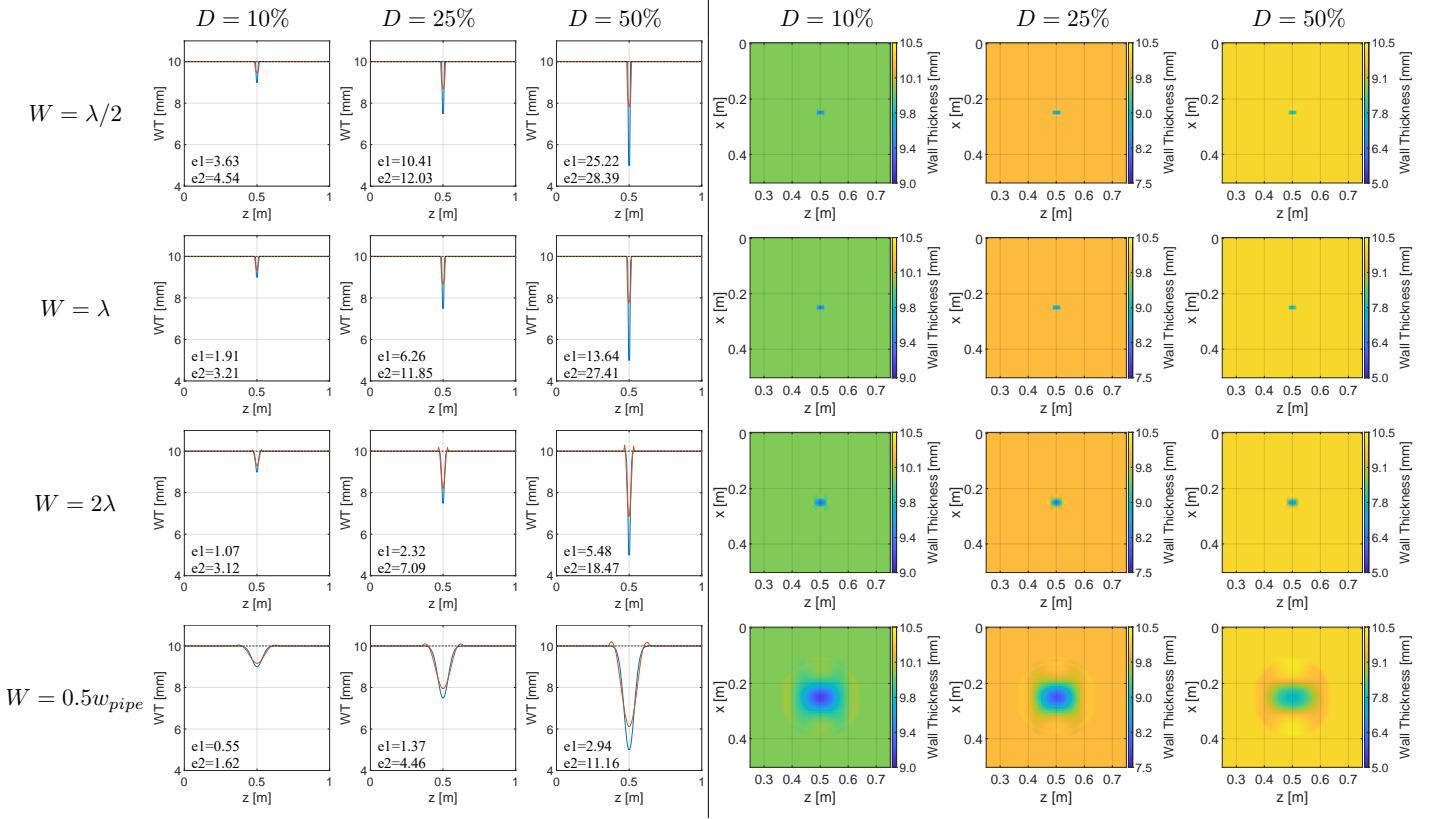


Figure 3.32: 130/180 kHz inversion results for different source (N_s) and receiver (N_r) counts. The left three columns show circumferential defect representations at the deepest point, the right three columns show the corresponding defect shape.

Different transducer configurations

To better understand the effect of transducer quantity, another parameter sweep is performed with the default defect configuration using 130/180 kHz. The results can be found in Figure 3.33. Across the entire range of transducer counts, we observe that the inversion stably converges, and the defect shape is well reconstructed. However, when the source count is reduced to three, aliasing effects become significant, introducing noticeable jittering artifacts in the inversion result. This arises from the construction of the sensitivity kernel, which is a summation over individual contributions. With fewer sources, the absolute amplitude of the kernel is lower, while the aliasing effects remain present, making the inversion more susceptible to these jittering effects, degrading the reconstruction quality.

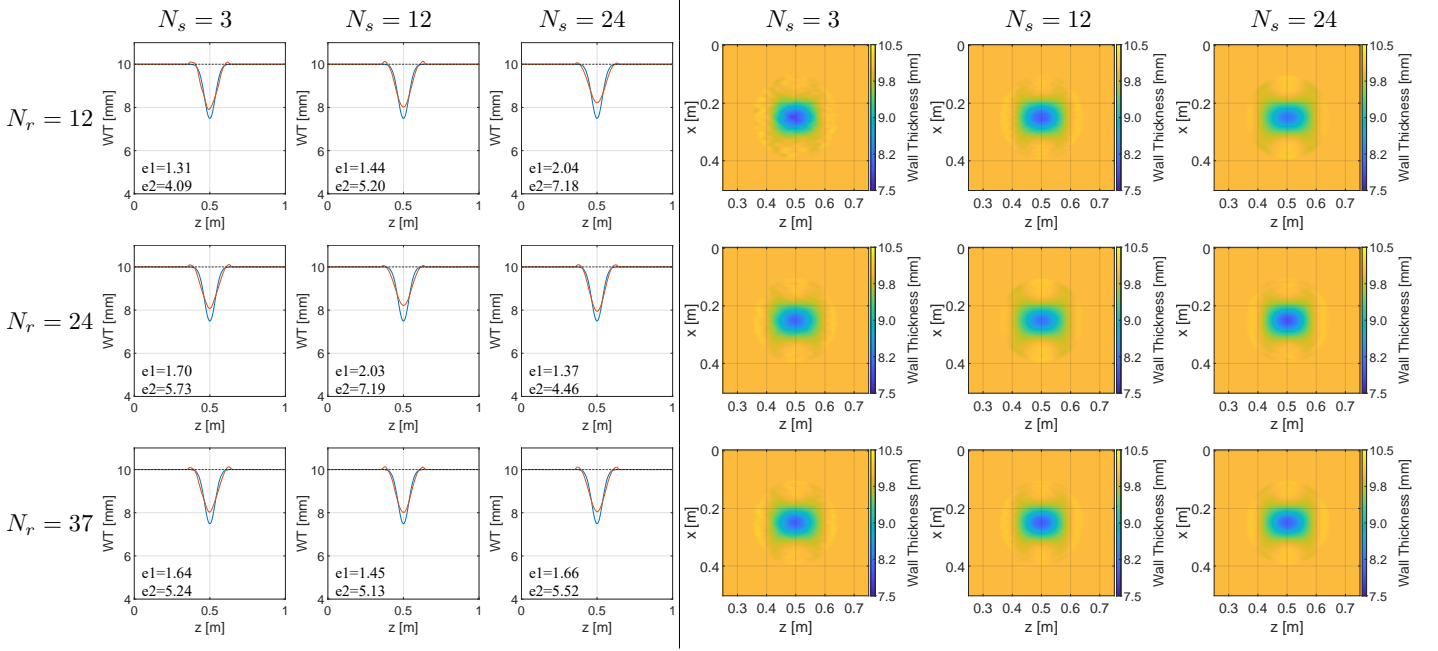


Figure 3.33: Defect configurations for different relative widths (W) and depths (D) for a nominal thickness of 10 mm and circumferential width of 0.5 m. The left three columns show circumferential defect representations at the deepest point, the right three columns show the corresponding defect shape.

3.3. Inversion results for the Echo defect

In real-world scenarios, corrosion-induced defects rarely exhibit smooth, Gaussian-like shapes. Instead, they often present irregular geometries characterized by sharp edges, asymmetry, and rough surfaces due to the stochastic nature of material degradation. As a result, studying idealized, smooth defects may not fully capture the structural and operational challenges posed by real corrosion damage. To address the gap, this section focuses on the investigation of the more complex 'Echo' defect in a noise free environment.

3.3.1. Model configuration

As mentioned in the introduction, the Echo defect is reconstructed from laser scans of an actual corrosion defect provided by industry. The defect exhibits pronounced asymmetry and non-uniform surface characteristics, as shown in Figure 3.34.




Figure 3.34: Wall thickness of the so-called Echo defect.

Similar to the analysis of Gaussian defects, we ensure a fair comparison across frequencies by reconstructing the defect on a mesh based on the wavelength at 180 kHz. This is particularly crucial for the Echo defect, as using a frequency-dependent parametrization could otherwise smooth out its sharp edges when a low frequency is used with large wavelengths. To evaluate the necessary spatial resolution, we once again perform a 2D Fourier transform of the defect, representing it in the wavenumber

domain. The resulting representation is shown in Figure 3.35. From this figure, it is found that the image contains relatively high amplitude values at a high radius from the center, indicating that the defect is characterized by high-frequency components which are associated with sharp edges. In order to accurately image this defect a wavelength of at least $2 \cdot 0.015 \text{ m} = 0.03 \text{ m}$ is required. From Figure 3.36 it becomes apparent that for the S_0 and A_0 modes the required frequency is 168 and 69 kHz, respectively.




Figure 3.35: Wavenumber domain representation of the Echo defect.




Figure 3.36: Wavelength vs frequency for different modes

3.3.2. Single frequency inversion

Since we aim to image sharp edges in the defect, ideally the inversion is performed at the maximum usable frequency where the attenuation is still sufficiently low. Therefore a maximum frequency of 180 kHz is used.

The inversion results for 130, 150 and 180 kHz is shown in Figures 3.40, 3.41 and 3.42, respectively. When analysing the FFT of the initial sensitivity kernel for 130 kHz, which can be seen in Figure 3.37, we find that at the initial iteration the sensitivity kernel only captures low frequency components and hence makes a coarse model update, as can be seen in Figure 3.39, which visualizes the wall thickness after the first iteration step. Throughout the inversion more and more high frequency components are introduced in the wavenumber domain up to the diffraction limit of $\lambda/2$ (which is 20mm for 130 kHz). As can be seen in Figure 3.38, which is the wavenumber domain of the sensitivity kernel in the final iteration. In the plot, the effect of the limited-view problem becomes apparent. While the wavenumber components in k_x direction are well-defined, those in the k_z direction remain uncaptured.




Figure 3.37: Wavenumber domain representation of the kernel at the first iteration at 130 kHz




Figure 3.38: Wavenumber domain representation of the kernel at the last iteration at 130 kHz




Figure 3.39: Wall thickness after first iteration step for a 130 kHz inversion.

From the difference plot in Figure 3.40c, it is evident that the shape error is primarily concentrated around the sharp edges of the defect that require a finer spatial resolution than 20 mm as these regions correspond the components that are not captured by the sensitivity kernel. Furthermore, depth estimation errors remain high across all inversions of the Echo defect due to physical limitations imposed by the diffraction limit. The pit, which has a pike-like shape that narrows with depth, can only be resolved down to a width of $\lambda/2$, leading to significant minimum depth estimation errors.




Figure 3.40: Inversion result for 130 kHz. The colorbars indicate wall thickness in mm.




Figure 3.41: Inversion result for 150 kHz. The colorbars indicate wall thickness in mm.




Figure 3.42: Inversion result for 180 kHz. The colorbars indicate wall thickness in mm.




Figure 3.43: Wavenumber domain representation of the kernel at the first iteration at 180 kHz




Figure 3.44: Wavenumber domain representation of the kernel at the last iteration at 180 kHz

When looking at the first and last wavenumber domain representation of the sensitivity kernels for the 180 kHz inversion, visualised in Figure 3.43 and Figure 3.44, a similar pattern is observed. At the initial iteration the update step is mainly dominated by low frequency components. Notably, in contrast to 130 kHz, some high frequency components are visible, which can be attributed to the aliasing effects attributing to the sensitivity kernel. Again, in the final iteration we find that frequency components are captured up to the diffraction limit which is 14 mm for 180 kHz.

To determine the maximum achieved resolution achieved, we analyze the wavenumber domain of the reconstructed defect for 130 kHz and 180 kHz. The decibel plot reveals a very strong dominance of low frequencies. To improve resolution assessment and facilitate comparison, the figures are presented using the 97th percentile. This means the highest 3% of values are excluded, emphasizing the range that contains 97% of the data. This approach helps reveal high-frequency components, as shown in Figure 3.45 and Figure 3.46 for 130 kHz and 180 kHz, respectively. From the figure it becomes apparent that some frequency components are captured at the diffraction limit $\lambda/2$, however their contribution is negligibly small. The significant contributions are primarily within 1 to 1.5λ (corresponding to 60 and 41 mm for 130 and 180 kHz respectively) and above. Although for the 130 kHz case there are significant attributions at a resolution of λ .




Figure 3.45: Wavenumber domain of inversion result for the Echo defect at 130 kHz. Plotted as the 97% percentile.




Figure 3.46: Wavenumber domain of inversion result for the Echo defect at 180 kHz. Plotted as the 97% percentile.

When performing inversions with a range of different frequencies, although small, there is a correlation between increasing the frequency and a reduction in shape error. This trend is observed for both the S_0 and A_0 modes, as shown in Figure 3.47. This outcome is expected for the complex Echo defect, as its reconstruction relies on high resolution to capture finer details, unlike the Gaussian defect with its smooth edges.

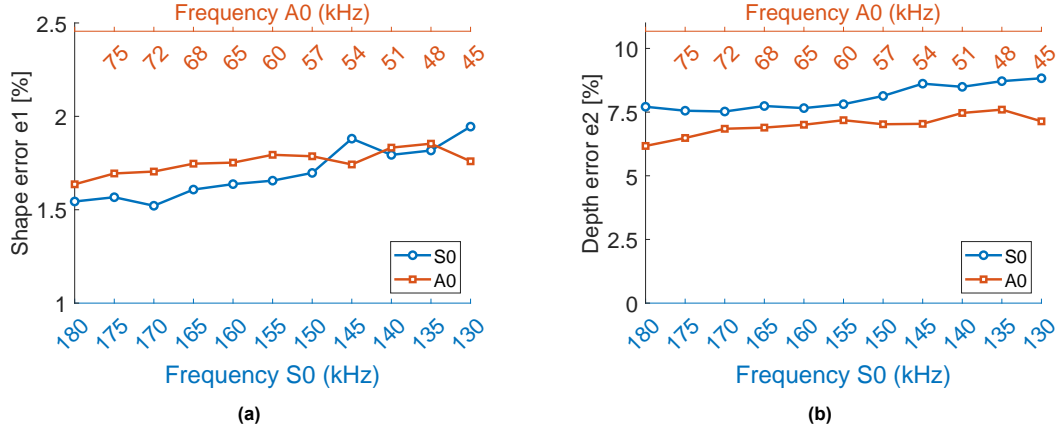


Figure 3.47: (a) Shape error for multiple wavelengths evaluated as single frequency inversions. (b) Depth error for multiple wavelengths evaluated as single frequency inversions.

3.3.3. Multi frequency inversion

Similarly to the Gaussian defect we aim to improve the performance by constraining the problem using multiple frequencies within the usable frequency range. An initial inversion is conducted using a two-frequency continuation approach with 130 and 150 kHz (130/150 kHz). The resulting thickness reconstruction is shown in Figure 3.48. Compared to the single-frequency results at 180 kHz we find that the shape error goes up from 1.72% to 1.82%. This indicates that the shape error is dominated by the sharp edges that cannot be imaged, as in our multi frequency inversion we miss out on the higher resolution provided by the 180 kHz wavelength.

Further inversions are performed using 150/180 kHz and 130/180 kHz. The results, shown in Figures 3.49 and 3.50, indicate that the 150/180 kHz inversion achieves the best reconstruction. From table 3.5 we find that, unlike the Gaussian defect case, where spacing frequencies further apart generally improved performance, for the Echo defect the higher frequency combination of 150/180 kHz outperforms the wide spacing of 130/180 kHz. This additionally suggests that for a defect containing rough surfaces, higher frequencies play a more dominant role the reconstruction.




Figure 3.48: Inversion result for 130/150 kHz. The colorbars indicate wall thickness in mm.




Figure 3.49: Inversion result for 130/180 kHz. The colorbars indicate wall thickness in mm.




Figure 3.50: Inversion result for 150/180 kHz. The colorbars indicate wall thickness in mm.

f [kHz]	e_1 [%]	e_2 [%]
130	1.94	8.83
150	1.89	7.91
180	1.72	7.48
130/150	1.82	8.11
130/180	1.78	8.16
150/180	1.67	7.20

Table 3.5: Error metrics e_1 and e_2 for different single and multi-frequency inversions.

3.3.4. Parametric study

Similar to the Gaussian defect study, a parameter sweep is conducted for the Echo defect using the default configuration settings and a frequency continuation at 130 kHz and 180 kHz. However it is important to note that the absence of noise in this study significantly influences the results. As the noise free environment allows the inversion process to over-perform compared to scenarios that include real world constraints such as the sensor noise.

Different defect configurations

When creating the observed data set the defect is discretized on the grid used for forward simulation, leading to pixelation effects when the defect width is very small, as can be seen in Figure 3.51. As a result, the defect shape becomes highly distorted at these small widths, forming a highly pixelated version of the intended defect geometry.

The results of the parameter sweep indicate similar trends to those observed in the Gaussian defect study. However, already at $W = \lambda$ inversion struggles as defect depth increases due to a lack of sensitivity. When the defect width is less than half a wavelength, the available wavefield does not have sufficient spatial resolution to distinguish these fine-scale features. As a result, the inversion process attempts to recover details that cannot be resolved, leading to significant overshooting. Beyond the diffraction limit, the inversion algorithm can reconstruct the defect, even though the resolution achieved is lower than the realistic resolution of 1.5λ . This is because the inversion does not have to deal with discrepancies between the model used for inversion and the actual wave behavior in experiments which normally degrade the resolution. Due to the discretization process, the maximum defect depth scores remain low for small widths since the pit inside the defect is not accurately modeled. At larger defect widths, this becomes more apparent, as the formation of the pit significantly increases the depth error. Overall, the inversion process effectively estimates defect shape across various defect sizes but the influence of discretization must be carefully considered, particularly when interpreting results for narrow and deep defects.

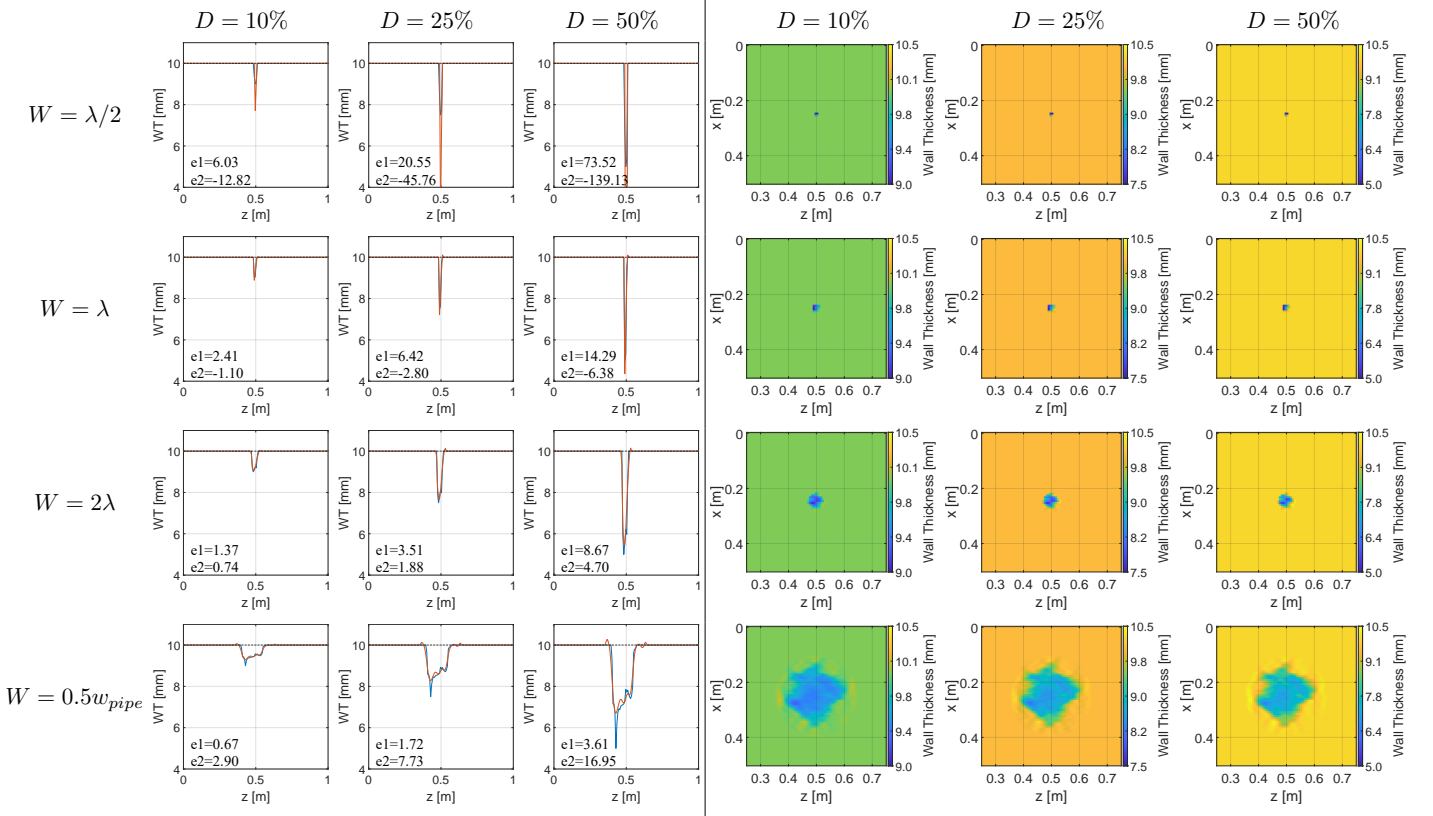


Figure 3.51: Defect configurations for different relative widths (W) and depths (D) for a nominal thickness of 10 mm and circumferential width of 0.5 m. The left three columns show circumferential defect representations at the deepest point, the right three columns show the corresponding defect shape.

Different transducer configurations

Also a parameter sweep is conducted using the default Echo defect configuration and a frequency continuation at 150 kHz and 180 kHz. The conclusions drawn from this parameter sweep are similar to those observed in the Gaussian defect case. Across the range of transducer counts, the inversion process remains stable, and the defect shape is effectively reconstructed. However at low transducer counts such in the case of 3 sources and 12 receivers, the jitter caused by the aliasing effects start to become problematic since the deepest point is not clearly distinguishable to the rest of the image.

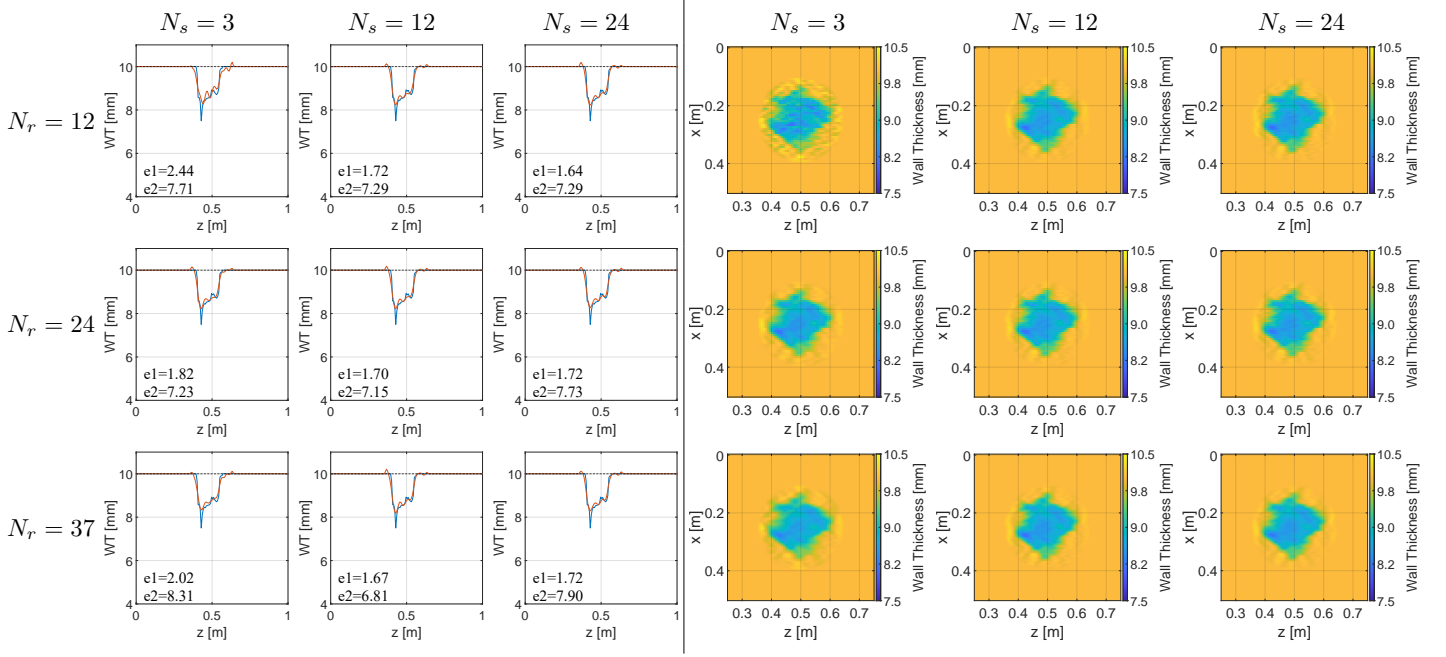


Figure 3.52: circumferential defect representations at the deepest point and shape reconstruction with a varying number of sources (N_s) and receivers (N_r). Conducted using frequency continuation simulation using 130 and 150 kHz, which expects aliasing effects below 30 sources and receivers.

3.4. Inversion performance under noise

In real-world applications, measurement data is always affected by noise components due to various factors such as sensor imperfections, environmental conditions, and signal processing limitations. Therefore to make a more realistic assessment of the inversion performance, in this section we again analyze the Echo defect but now with noise added to the observed data.

Since we found that the resolution we achieve in the noise free case is around 1.5λ , in order to accurately image this defect a wavelength of at least $0.015/1.5 \text{ m} = 0.01 \text{ m}$ is required. From Figure 3.53 it becomes apparent that for the S_0 and A_0 modes the required frequency is 320 and 278 kHz, respectively. For liquid loaded pipes, based on the attenuation curve in Figure 2.4, using a frequency this high is not feasible for both modes. The shear horizontal wave modes (SH_1) can offer a suitable alternative as it experiences significantly less attenuation from liquid loading. This reduction in attenuation is due to the absence of vertical, out of plane displacement in the shear horizontal propagation. However the required frequency is around 442 kHz, which is not feasible due to practical transducer limitations. Although a reconstruction that perfectly describes the details of the defect might not be feasible due to a lack of resolution, a usable reconstruction indicating the presence and shape of the defect up to the resolution limit, which still gives valuable information about the structural health of the pipeline.




Figure 3.53: Wavelength vs frequency for different modes

Unless stated otherwise the signal to noise ratio (SNR) is generated as broadband Gaussian noise in the time domain at 30 dB, meaning it initially has a flat spectral distribution before processing. The noise is transformed into the frequency domain using a Fourier transformation, where its power is computed and scaled at the center frequency of 155 kHz to the desired SNR. The noise degrades the quality of the misfit, as can be seen in Figure 3.54 and Figure 3.55, where we see the same misfit for the previous case of a noise free misfit, and the misfit in which noise is added.




Figure 3.54: Misfit for the Echo defect at 130 kHz without noise.

Figure 3.55: Misfit for the Echo defect at 130 kHz with noise.

Figure 3.56: Misfit for the Echo defect at 180 kHz with noise.

When comparing the misfit plots for 130 kHz and 180 kHz with noise in Figure 3.55 and Figure 3.56, we observe that higher frequencies are less influenced by the noise. This is due to the increased sensitivity at high frequencies. Therefore the misfit pattern is being more in contrast with the noise.

The inversion result for 180 kHz including noise is shown in Figure 3.57. When compared to Figure 3.42, we find that at 180 kHz noise case the algorithm is still capable of reconstructing the defect without any significant loss of performance. Based on the misfit plot this is also not unexpected since the defect artifacts are still clearly distinguishable and not significantly distorted.




Figure 3.57: Inversion result for 180 kHz with noise (SNR = 30 dB). The colorbars indicate wall thickness in mm.

If we further decrease the SNR to 20 dB, we observe a significant distortion of the misfit, as visualized in Figure 3.58. This degradation is also reflected in the error metrics, as can be seen in Figure 3.59, where the shape error increases from 1,71% to 2,13%. Since the depth error remains relatively stable (changing from 7,52% to 7,35%), we can attribute the increase in shape error primarily to noise artifacts introduced in the sensitivity kernel, as the noise is prominently present in the measured misfit.




Figure 3.58: (a) 180 kHz noise free case and (b) 180 kHz case with a SNR 20 dB.




Figure 3.59: Inversion result for 180 kHz with noise (SNR = 20 dB). The colorbars indicate wall thickness in mm.

Due to the ability of high frequencies to better handle noise, we now observe a more clear correlation between the algorithm's performance and frequency, as can be seen in Figure 3.60. The overall shape error decreases in a linear fashion as frequency increases, while the maximum depth error keeps fluctuating. Besides the influence of the algorithm's line search, the fluctuation in depth error can be attributed to the random nature of the noise introduced at each frequency. Although the SNR remains fixed at 30 for all inversions, the specific realization of noise varies, affecting certain frequency cases more than others. This variability can lead to inconsistencies in the reconstructed depth, sometimes causing an underestimation of the deepest point. Additionally, noise artifacts may introduce jitter, which

can result in a local deep point within the defect region that does not necessarily correspond to the actual deepest location. This effect could also contribute to the outliers observed in the shape error at 145 and 155 kHz. However, further investigation is required to better quantify the impact of this randomness on inversion results.




Figure 3.60: (a) Shape error for multiple wavelengths evaluated as single frequency inversions with SNR = 30 dB. (b) Depth error for multiple wavelengths evaluated as single frequency inversions also with SNR = 30 dB.

Similarly to the previous sections we conduct two parameter sweeps for the noise case with a SNR of 30 dB and a frequency continuation of 150/180 kHz. The results can be found in Figure 3.61 and Figure 3.32. As expected for the width of $0.5w$ we find that there is hardly any performance loss due to the noise since in the misfit for 180 kHz, the defect pattern is clearly distinguishable. However at small defect widths we observe that the method already starts breaking down at a width of 2λ . Looking at the misfit for this inversion we find that for this width the noise does play a significant role as the defect is hardly visible in the sensitivity kernel, as can be seen in Figure 3.62 where the initial and third kernel is plotted for a defect width of 2λ . Due to the lack of sensitivity the amplitude of the focusing does not have a clear contrast to the background. This is further confirmed when looking at the convergence plot in Figure 3.63, where it becomes apparent that a steady state is being reached while the misfit is hardly reduced.

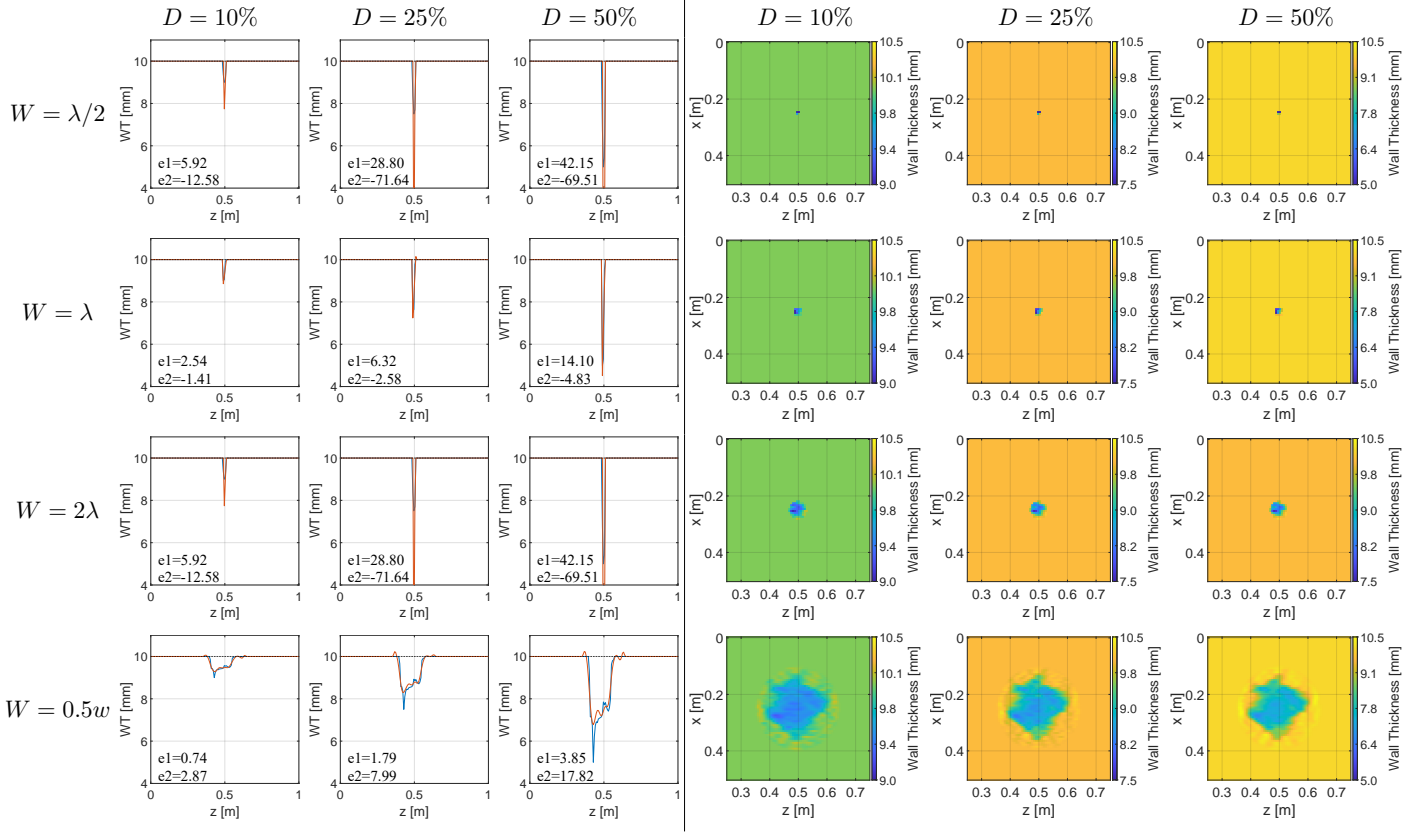


Figure 3.61: Defect configurations for different relative widths (W) and depths (D) for a nominal thickness of 10 mm and circumferential width of 0.5 m. The left three columns show circumferential defect representations at the deepest point, the right three columns show the corresponding defect shape.

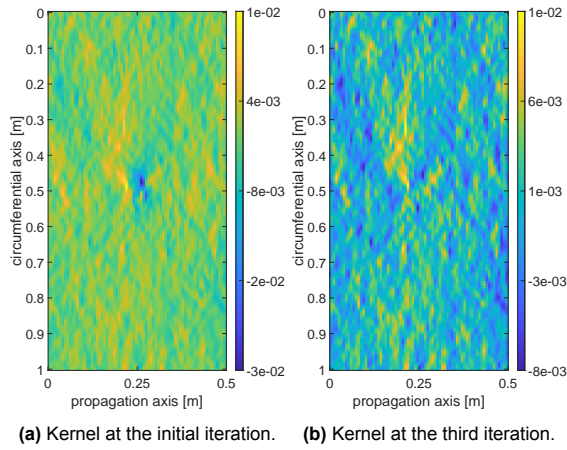


Figure 3.62: Sensitivity kernels at different iterations for a defect with a width of 2λ and a depth of $0.25T$ using 150/180 kHz and 24 sources/receivers.




Figure 3.63: Convergence plot for the inversion of a defect with a width of 2λ and a depth of $0.25T$ using 150/180 kHz.

In order to increase the contrast of the defect focusing and the background, either the sensitivity needs to be increased (e.g. by increasing the frequency) or the number of sources should be increased since the summation over sources gives an increase amplitude of the focusing compared to the background noise, as can be seen in Figure 3.64 where a spatial sampling of $\lambda/4$ for both sources and receivers is used for 180 kHz, resulting in 73 sources and 73 receivers. As can be seen in Figure 3.65 with the increased sensor count an effective image can be created.

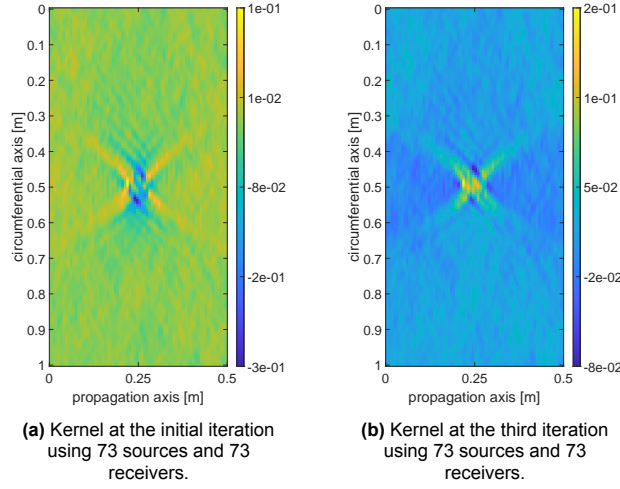


Figure 3.64: 2 Sensitivity kernels at different iterations in an inversion for a defect with a width of 2λ and a depth of $0.5T$ using 150/180 kHz and 73 sources and 73 receivers.




Figure 3.65: Inversion result for defect of width of 2λ and a depth of $0.5T$ using 150/180 kHz and 73 sources and 73 receivers. The colorbars indicate wall thickness in mm.

The effect of varying SNR levels on the reconstruction quality is further investigated by performing inversions with 150 kHz/180 kHz with SNR values ranging from 25 to 5 dB. The results, shown in Figure 3.66, indicate that the overall defect shape is relatively well preserved down to an SNR of 15 dB. Below this level, significant degradation of the defect shape occurs, leading to unrecognizable defect shapes. While the general outline remains identifiable at an SNR of 15 dB and higher, finer defect details are increasingly disturbed, with the deepest region of the defect becoming indistinguishable from surrounding noise artifacts at 20 dB and below. Furthermore, when the SNR exceeds 15 dB, noise artifacts starts to form around the defect, indicating that the inversion algorithm has difficulty constraining the reconstruction as result of the noise.

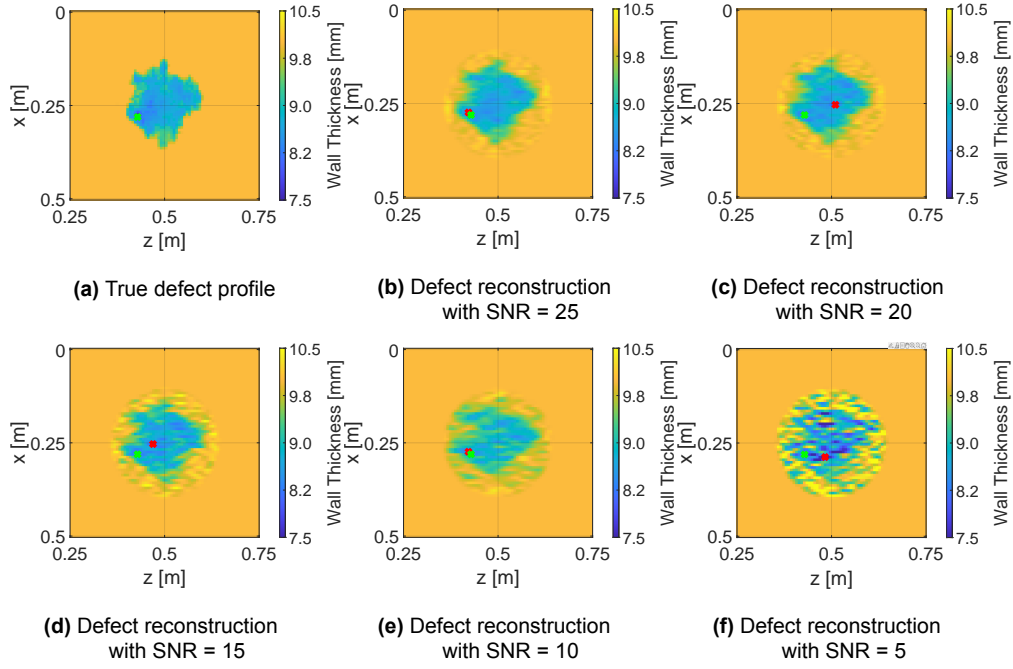


Figure 3.66: Echo defect reconstruction using 150/180 kHz under different noise levels for the S_0 mode. The green cross indicates the true maximum depth location. The red cross indicates the reconstructed maximum depth location.

Although 150/180 kHz at an SNR of 30 is relatively unaffected by noise, the transducer sweep with noise on the Echo defect yields results similar to the noise-free case. However, the noise introduces additional background interference, adding to the existing background noise caused by the aliasing effects. This reduces contrast with the sensitivity kernel, leading to additional jittering effects and incorrect minimal depth locations. Notably, at lower frequencies the distortions becomes more pronounced due to reduced defect sensitivity. Consequently, jittering could occur even at higher transducer counts compared to the 150/180 kHz case.

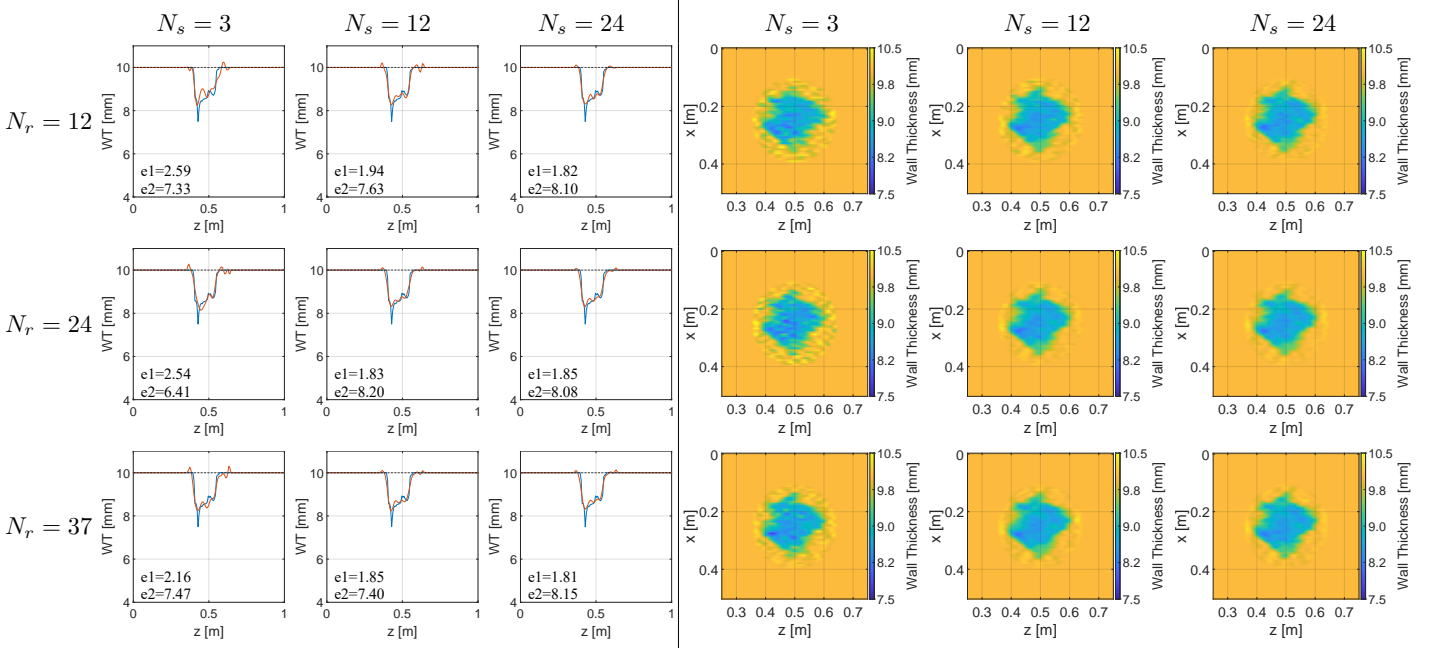


Figure 3.67: Transducer sweep results for different configurations of N_s and N_r under noisy conditions.

4

Discussion

The defect reconstruction is influenced by modeling choices, numerical approximations, and regularization techniques. Achieving accurate reconstructions requires balancing theoretical performance with practical applicability, especially when dealing with noise, model discrepancies, and limited data availability. This chapter examines key factors that impact inversion performance, including the role of inverse crime, the importance of proper kernel scaling, and the application of regularization techniques to improve stability and robustness.

First, we discuss the inverse crime. While this approach eliminates modeling inconsistencies and provides an idealized test environment, it does not account for real-world uncertainties such as material property variations and numerical approximations, leading to overly optimistic results that may not translate well to practical applications.

Next, we explore the kernel scaling, a crucial step in ensuring stable update steps during inversion. Without proper scaling, the inversion process suffers from instability, with updates that are too large causing divergence. Or, the update steps are too small, leading to slow convergence.

Finally, we address the role of regularization in the defect reconstruction, which helps mitigate the ill-posed nature of the problem by constraining the solution space. Techniques such as the application of a Tukey window and threshold constraints can improve inversion accuracy. However, these methods can also introduce application limitations which need to be carefully considered.

4.1. Inverse crime

By using the same numerical model for both data generation and inversion, discrepancies due to modeling errors, numerical approximations, or unknown physical effects are eliminated. This leads to faster convergence and high reconstruction accuracy. The inverse crime offers several benefits: it provides an idealized test scenario to verify the correctness of the inversion algorithm, ensures that observed errors stem from the inversion method itself rather than modeling inconsistencies, and enables controlled assessments of the adjoint methodology. Although the inverse crime is very valuable for proving the fundamental effectiveness of the adjoint method, in practical applications, model discrepancies exist, and the inversion may not perform as well. Key considerations include:

- The inversion results may appear more accurate than they would be with real experimental data, where modeling errors and noise play a significant role.
- The method's ability to handle realistic conditions is not tested, potentially leading to unexpected challenges in real-world applications.
- Without modeling errors, the inversion may appear robust even under conditions where it would typically struggle, such as limited transducer count.

While the inverse crime is useful for testing inversion techniques, assessing performance under more realistic conditions is crucial. Future research should introduce modeling errors, such as discrepancies

in material properties affecting phase velocity calculations, numerical approximations, discretization errors, or incorrect transducer characteristics. These factors can significantly impact the accuracy and robustness of the inversion. To mitigate this, future work should consider using a different forward model for inversion than for data generation or incorporating measured data, such as the 3D elastodynamic wave equations.

4.2. Scaling and Regularization

This section discusses two key aspects of scaling and preconditioning in the current inversion scheme. First, we address kernel scaling, which ensures that sensitivity kernels produce updates of appropriate magnitude. Furthermore we discuss the preconditioning using a Tukey window, which helps localize the inversion process around the defect while minimizing interference from irrelevant regions.

4.2.1. Kernel scaling

The sensitivity kernel is not naturally scaled within the range of a feasible update step. If left unscaled, the magnitude of the updates could be in the meter range, leading to instability in the inversion. In the current implementation the sensitivity kernel undergoes rescaling to maintain update values between -1 mm and +1 mm to preserve meaningful update ranges. Setting the scaling limit of the sensitivity kernel to high values such as ± 3 mm or ± 5 mm results in larger update steps which could speed up convergence for large defects. However this approach presents risking inefficient oscillations or divergence when updates become too forceful, especially for smaller defects. When the scaling factor is excessively small it causes the convergence to occur at a slow pace. There is no a-priori knowledge of the defect configuration which prevents the scaling parameter from being tuned. In the case where the initial scaling is way off the actual defect sizing, the step size parameter α functions as a safety mechanism to further scale the update step. If the sensitivity kernel scaling is too large, the line search will naturally select a smaller α to prevent overshooting. Conversely, if the update step is too small, the line search may allow for a larger α to speed up convergence while keeping updates within a stable range. Furthermore after a few iterations the sensitivity kernel is naturally scaled within the scaling bounds do to the diminishing sensitivity as the defect reconstruction is better shaped every iteration, effectively making the scaling factor 1 for later iterations.

4.2.2. Regularization

Since inverse problems in defect reconstruction are often ill-posed, regularization techniques have appeared to be essential for the inversion process, as can be seen in the unregulated inversion in Figure 4.2a. They impose constraints that limit the solution space, thereby enhancing robustness and reducing non-uniqueness. In our inversion we make use of a Tukey window to spatially constrain the reconstruction to the defect region. Furthermore threshold forging is introduced for noise cases to prevent artificial wall growth beyond the nominal thickness. Both approaches improve inversion performance but also introduce challenges.

Tukey window

The current inversion scheme applies a Tukey window at the defect site to suppress artifacts in regions of the plate that are not near the defect, reducing the ill-posedness of the problem. By aligning the window size with the defect dimensions, the method optimally reduces unnecessary degrees of freedom.

While the Tukey window improves inversion stability and efficiency, it also has drawbacks. The current implemented round shaped Tukey window is only effective for localized defects, making it unsuitable for larger defects, such as long axial grooves. The current implementation relies on prior knowledge of the defect, which isn't available in real-world applications. In order to mitigate this, a Tukey window can be designed that is placed based on the focusing in the sensitivity kernel. A threshold value can be set in the kernels amplitude to determine the Tukey window size. Numerical experiments show that the sensitivity kernel in the initial iteration captures the rough defect shape of non circular defect shapes well, as visualized in Appendix D. Also for different pipe configurations a clear contrast remains in the sensitivity kernel for the default case. Therefore, a method utilizing a threshold to determine the Tukey window is not restricted to circular defects or short pipe lengths.

Furthermore, the Tukey window creates an artificial gradient around the defect, which is directly printed on the sensitivity kernel. Therefore the Tukey window artificially creates wall thickness growth, contributing to an increase in shape error. Here, a trade-off needs to be made between the rate of decay of the Tukey window, and the increase of null space.

In order to mitigate this effect another regularization step can be applied to the problem which prevents the wall thickness from exceeding the nominal thickness value after the update step. This can be achieved by applying a threshold that sets all contributions exceeding the nominal thickness to zero, which can be illustrated well for the noise free Gaussian defect in Figure 4.1. Using the method the non-uniqueness of the problem is reduced, as the redistribution of wall thickness is constrained below the nominal thickness. Furthermore this is particularly beneficial for the adjoint method, as the limited view can create shadow regions adjacent to the defect, leading to artificial wall thickness growth. These unintended variations result in compensatory wall losses throughout the pipe, ultimately reducing the phase misfit. However, this step can only be done after the update step, as otherwise an overshoots in the reconstruction cannot be corrected anymore. As can be seen in Figure 4.2c, using this method alone does not significantly improve the inversion result and thus could be used in combination with the Tukey window.

Both regularization methods can be considered valid in different scenarios depending on defect characteristics. However, they also introduce the risk of over performance, where the algorithm achieves artificially high accuracy under controlled conditions that may not translate well to real-world applications. For example if the assumed nominal thickness is not accurate (e.g., due to manufacturing variations where the nominal pipe thickness can vary by up to $\pm 0.75\%$ [51]), applying this threshold may systematically obstruct the reconstruction. Careful consideration of these limitations is necessary to ensure the robustness and practical applicability of the inversion approach in real defect reconstruction scenarios.




Figure 4.1: Visualization of threshold forging using through cut at the center of a gaussian defect case. No wall growth is allowed above the nominal thickness, and hence is set back to 0, as indicated by the red arrows.

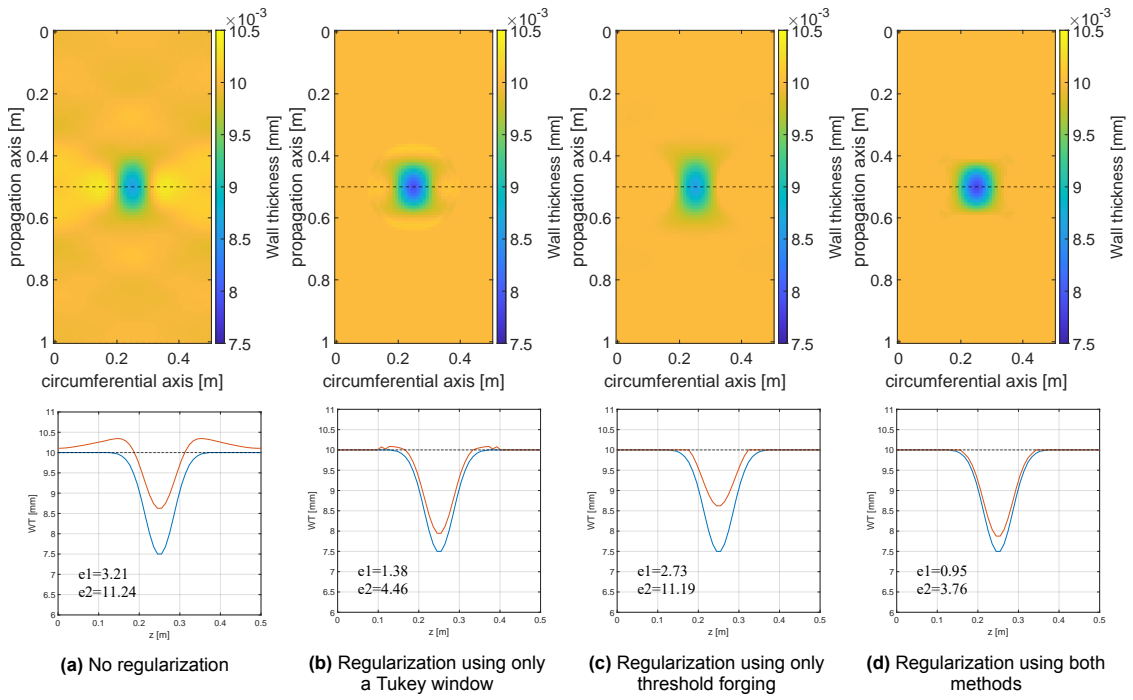


Figure 4.2: Inversion for a single frequency of 130 kHz using different regularization methods.

5

Conclusion

In this thesis, the adjoint method has proven effective for localizing defects in a pipe using forward modeling in the FX domain. The method works well in both noise-free and noisy environments, provided there is sufficient contrast between the defect focusing and its background.

If the sensitivity is insufficient, additional sensors can be employed to create the focusing effect in the sensitivity kernel, or higher frequencies should be used due to their increased sensitivity. For defects that lead to high sensitivity, using only a few transducers remains feasible, however for a low contrast defect case noise plays a more significant role and additional transducers are required. Although insufficient transducer sampling produces aliasing effects when the spacing is larger than $\lambda/2$, their effect on the inversion performance is found to be minimal.

Since the adjoint method back-propagates the misfit measured at the receiver line, the method is vulnerable to sensor measurement noise. For the intended use case of the S_0 mode with 24 transducers and a maximum frequency of 180 kHz, the signal to noise ratio needs to be atleast 25 before reconstruction quality is significantly reduced.

During the initial stages of the reconstruction process, for both the simple Gaussian and the more complex defect, low-frequency components dominate, resulting in a rough initial image. High-frequency components refine the reconstitution at later iterations. Wavenumbers up to the diffraction limit are included in the subsequent iterations but have weak amplitude. The algorithm achieves a resolution of approximately 1 to 1.5λ in noise-free scenarios.

The method demonstrates convergence without regularization but achieves significantly better performance through application of a Tukey window that limits the null space and decreases the problems ill-posedness. However in the current method a-priori information is used to effectively model the Tukey window. To overcome this constraint a windowing technique could be created that uses iterative sizing throughout the inversion process.

Similar to the perturbation gradient method employed at TNO, the use of the adjoint method results in wall thickness interchanges which cause an alternating error pattern that is an over and underestimation of wall thickness. Using multiple frequency continuation or multiple modes in a single inversion further constrains the solution space which helps to reduce wall thickness interchange. Through numerical experiments, it is observed that the choice of frequency is crucial for overall performance. Frequencies should be chosen complementarily to enhance inversion performance. While low frequencies provide a rough reconstruction of the defect, the subsequent lack of sensitivity at high frequencies prevents the capture of fine details in the reconstruction.

For the Echo defect, the S_0 mode proves unsuitable for perfect reconstruction, as errors are primarily concentrated at sharp edges that cannot be modeled due to limited resolution. However, the presence of the Echo defect is clearly depicted, and a good shape reconstruction is obtained but without details that are smaller than a width of approximately 1 to 1.5λ . Increasing resolution by using higher frequencies is, similar to the A_0 mode, not feasible for the S_0 mode due to attenuation. Different modes such

as the SH_1 mode should be explored to increase resolution due to their low attenuation characteristics.

To conclude, the adjoint method offers a promising approach for gradient calculation in guided wave tomography. While it requires regularization steps, which can impose limitations in certain use cases, its ability to efficiently compute gradients and locate defects makes it a promising tool for optimizing defect reconstruction. However experimental validation using measured data or 3D-elastodynamic data is essential to bridge the gap between simulation and real world applications, ensuring the method's robustness and practical feasibility.

5.1. Recommendations and further research

Building upon the findings of this study, several recommendations for improving reconstruction accuracy are proposed in this section.

5.1.1. Hessian-Based Optimization and trust region

To further improve reconstruction accuracy and stability, an alternative optimization method based on a Hessian approximation should be explored. Instead of directly computing the Hessian using Fréchet derivatives, quasi-Newton methods, such as BFGS (Broyden–Fletcher–Goldfarb–Shanno), provide an efficient alternative [50].

In the BFGS method an approximation of the inverse Hessian is done using only gradient information, hence the sensitivity kernel can be used for that. With the extra Hessian information improved inversion performance is expected since better step directions can be made in the optimization process. Furthermore this potentially help constraining the problem and thus mitigating the effects of wall thickness interchange. Future work should investigate the integration of BFGS or limited-memory BFGS (L-BFGS) into the inversion framework to enhance reconstruction quality and computational efficiency. However the gradient information needs to be carefully scaled in order to be effective. Similar to the non linear Conjugate Gradient method, an integral part of the BFGS method is calculating the step size.

In our current algorithm, a line search method is used to determine step sizes in the optimization process. While line the search ensures progress toward the optimum by selecting a step size that sufficiently decreases the objective function, it is computational expensive in complex large scale non linear optimization regions. An alternative approach is the trust-region method [50], where the update step is constrained within a region around the current iteration where the model is expected to be 'reliable'. Trust-region methods generally have improved stability in ill-conditioned problems. However they can also perform bad if the gradient information is not well scaled.

Therefore future work should explore the integration of trust-region strategies alongside quasi-Newton methods like BFGS to increase robustness, particularly in noise cases.

5.1.2. Validation on Real-World and 3D Elastodynamic Data

To ensure the effectiveness of the proposed inversion techniques and to prove it's promising use in real world applications, testing on elastodynamic or real-world measured data is necessary. Currently the artificial data provides a solid foundation for developing and adjoint based reconstruction algorithms, but real experimental data introduces additional complexities, such as sensor noise and material irregularities. Comparing results from real-world ultrasonic measurements will help evaluate the practical feasibility and robustness of the inversion approach.

Furthermore, future research should investigate the extension of the inversion to three dimensional elastodynamic data. Incorporating 3D simulations will allow a more accurate assessment of defect characterization, and can possibly be extended to more complex geometries like a curve. By validating the methodology on both real-world experimental data and full 3D elastodynamic simulations, the usability of the adjoint method in the reconstruction can be further established.


5.1.3. Validation on different configurations

This thesis has been mainly focused around single center circular defects, as mention before, pipelines can develop various defects during their lifespan, including corrosion due to chemical or electrochemical reactions, or gouging caused by impacts from foreign (third party) objects, which can appear as ridges

and/or grooves that contribute to stress concentration. Dents result from localized deformation, leading to a reduction in diameter at a specific section. Additionally, combined defects can occur, such as dents with cracks, dents with gouges, or a combination of cracks and dents within corroded areas [52].

As already mentioned, the adjoint sensitivity kernel is definitely not limited to the defect shapes used in this thesis. In future research these more 'exotic' defect configurations should be examined to further argument on the feasibility of the adjoint in real world applications. It is likely that different defect geometries will require more complex regularization strategies for effective inversion. For instance, the current circular Tukey window, which works well for single-center circular defects, is not be suitable for a groove defect. Instead, alternative regularization approaches, such as the earlier adaptive windowing technique should be explored to better accommodate the specific geometry and characteristics of each defect type.

Furthermore, in this thesis, we solely focused on a pipe with a length of 1 meter and a circumference of 0.5 meters. However, from short numerical experiments the initial kernel showed also promising results for mapping defects in longer pipe lengths or different circumferences. Evaluating the performance of the adjoint sensitivity kernel in these cases is important to understanding its robustness and practical applicability. In particular, when pipelines are very short or the defect is very close to the transducers, aliasing effects might play a more significant role, potentially affecting the accuracy of the inversion. Future research should investigate these scenarios, ensuring reliable defect characterization across a wider range of pipeline configurations.



Derivation of the Gradient of the Objective Function

We derive the gradient of the objective function $J(m)$, defined as:

$$J(m) = \frac{1}{2} \sum_{\omega} \sum_{s,r} \|d_{obs}(x_r, x_s, \omega) - u(x_r, x_s, \omega)\|^2, \quad (\text{A.1})$$

with respect to the model parameter $m(x)$.

The objective function can be expanded explicitly as:

$$J(m) = \frac{1}{2} \sum_{\omega} \sum_{s,r} (d_{obs}(x_r, x_s, \omega) - u(x_r, x_s, \omega))^* (d_{obs}(x_r, x_s, \omega) - u(x_r, x_s, \omega)), \quad (\text{A.2})$$

where $*$ denotes the complex conjugate. Letting $\delta u = d_{obs}(x_r, x_s, \omega) - u(x_r, x_s, \omega)$ represent the data misfit, this becomes:

$$J(m) = \frac{1}{2} \sum_{\omega} \sum_{s,r} \delta u^*(x_r, x_s, \omega) \delta u(x_r, x_s, \omega). \quad (\text{A.3})$$

To compute the gradient $\nabla_m J(m)$, we utilize the chain rule. The derivative of J with respect to $u(x_r, x_s, \omega)$ is given by:

$$\frac{\delta J}{\delta u(x_r, x_s, \omega)} = -\delta u^*(x_r, x_s, \omega). \quad (\text{A.4})$$

The gradient can then be expressed as:

$$\nabla_m J(m) = - \sum_{\omega} \sum_{s,r} \Re \left[\delta u^*(x_r, x_s, \omega) \frac{\delta u(x_r, x_s, \omega)}{\delta m(x)} \right]. \quad (\text{A.5})$$

The relationship between the perturbation in the wavefield u and the perturbation in the model m is described by the Fréchet derivative Δ :

$$\Gamma = \frac{\delta u}{\delta m_i(x)}. \quad (\text{A.6})$$

Substituting Γ into the gradient expression, we obtain:

$$\nabla_m J(m) = \Re \left[\sum_{\omega} \sum_{s,r} \delta u^*(x_r, x_s, \omega) \Gamma(x) \right]. \quad (\text{A.7})$$

B

Smoothing defect shapes by wavenumber domain filtering

It is possible to filter out the high-frequency components in the wavenumber domain by applying a windowing function in k -space, as can be seen in Figure B.1. By suppressing these high frequencies and performing an inverse 2D Fourier transform, the reconstructed defect is smoothed, effectively removing the sharp edges introduced by the corrosion. While this process retains the overall structure of the original image, it results in a noticeably blurrier reconstruction, as the fine details of the defect are intentionally filtered out to reduce imaging artifacts.

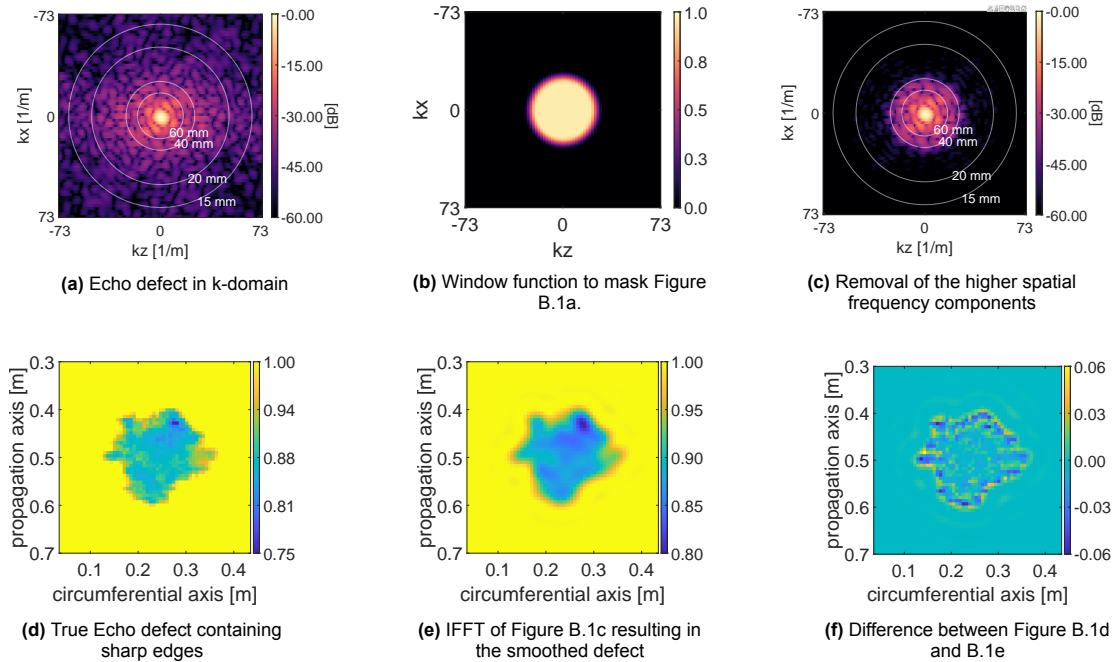



Figure B.1: Removal of sharp edges in spatial domain of the true defect by masking high frequency components in the k -domain.



Varying the $\alpha_{0,k}$ maximal growth rate

In this appendix, we examine the effect of varying the maximal allowable growth rate, $\max \delta \alpha_{0,k}$. This parameter serves as an upper bound on the initial step size $\alpha_{0,k}$ before the line search begins. Specifically, when we update the step size using

$$\alpha_{\text{new}} = \alpha \frac{g_{\text{prev}}^T P_{\text{prev}}}{g^T P} \quad (\text{C.1})$$

on each iteration, the step size α decreases rapidly for Gaussian defects, while it remains relatively stable for more complex defects. Gaussian defects exhibit smooth, well-defined gradients, allowing the optimization to take a highly effective first step. This rapid initial improvement causes the residual gradient norm $\|g\|$ to shrink significantly, leading to a small denominator in the step size update formula and a very small α . As a result, the optimization holds after the second or third iteration prematurely due to a update step that is in the magnitude of micrometers. This is why in Figure C.5 and Figure C.6 the error remains constant for the higher frequencies, as the first update step is the same due to the scaling factor on the sensitivity kernel.

The following figures present the shape and depth errors for multiple wavelengths, evaluated as single-frequency simulations for the S_0 mode. The results are shown for different values of $\max \delta \alpha_{0,k}$, illustrating its impact on the solution.




Figure C.1: Shape error for multiple wavelengths ($\max \delta \alpha_{0,k} = 0.5$).




Figure C.2: Depth error for multiple wavelengths ($\max \delta \alpha_{0,k} = 0.5$).




Figure C.3: Shape error for multiple wavelengths ($\max \delta\alpha_{0,k} = 0.75$).




Figure C.4: Depth error for multiple wavelengths ($\max \delta\alpha_{0,k} = 0.75$).




Figure C.5: Shape error for multiple wavelengths ($\max \delta\alpha_{0,k} = 1.0$).





Figure C.6: Depth error for multiple wavelengths ($\max \delta\alpha_{0,k} = 1.0$).

D


Initial kernels of non circular defects

While circular defects have been extensively studied in this thesis, real-world corrosion patterns often have more complex geometries. This appendix presents the initial sensitivity kernels corresponding to non-circular defect configurations, which can be found in Figure D.1, providing insight into how different shapes influence the sensitivity distribution at the first iteration of the inversion process. Additionally, in Figure D.2, the effect of pipe length on the sensitivity kernel is investigated by considering different pipe lengths.


The selected defect geometries include a rounded groove, a donut-shaped defect, and a double pit configuration. Each case is examined using a 130 kHz excitation frequency with a default transducer configuration.




(a) Groove defect configuration



(b) Sensitivity kernel for the groove defect configuration using 130 kHz



(c) Donut defect configuration



(d) Sensitivity kernel for the donut defect configuration using 130 kHz

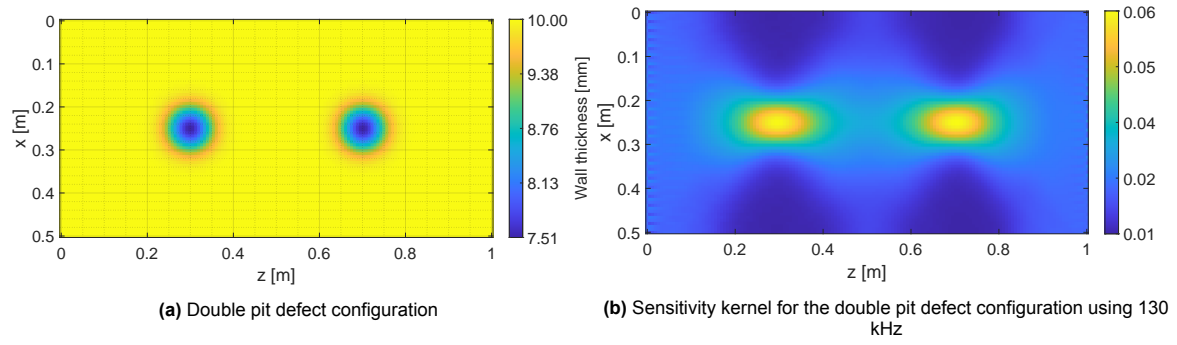


Figure D.1: Comparison of different defect configurations and their corresponding sensitivity kernels at the first iteration using 130 kHz excitation with the default transducer configuration.




Figure D.2: Initial sensitivity kernel for the Gaussian default case for different pipe lengths using 130 kHz

References

- [1] U.S. Department of Justice - Office of Public Affairs. *BP Exploration Alaska to pay \$25 million penalty for Alaskan North Slope oil spill*. 2014. URL: <https://www.justice.gov/opa/pr/bp-exploration-alaska-pay-25-million-penalty-alaskan-north-slope-oil-spill> (visited on 04/03/2014).
- [2] G. Pluvinage. "Pipe-defect assessment based on the limit analysis, failure-assessment diagram, and subcritical crack growth". In: *Materials Science* 42.1 (Jan. 2006), pp. 127–139. ISSN: 1573-885X. DOI: 10.1007/s11003-006-0065-x. URL: <https://doi.org/10.1007/s11003-006-0065-x>.
- [3] Naoya Kasai et al. "Evaluation of back-side flaws of the bottom plates of an oil-storage tank by the RFECT". In: *NDT & E International* 41.7 (2008), pp. 525–529.
- [4] David M Amos. "Magnetic flux leakage as applied to aboveground storage tank flat bottom tank floor inspections". In: *Materials Evaluation* 54.1 (1996).
- [5] Naoya KASAI, Kazuyoshi SEKINE, and Hiroaki MARUYAMA. "Non-destructive evaluation method for far-side corrosion type flaws in oil storage tank bottom floors using the magnetic flux leakage technique". In: *Journal of the Japan Petroleum Institute* 46.2 (2003), pp. 126–132.
- [6] T Warren Liao and Yueming Li. "An automated radiographic NDT system for weld inspection: Part II—Flaw detection". In: *Ndt & E International* 31.3 (1998), pp. 183–192.
- [7] Mridul Gupta et al. "Advances in applications of Non-Destructive Testing (NDT): A review". In: *Advances in Materials and Processing Technologies* 8.2 (2022), pp. 2286–2307.
- [8] Jitesh Jodhani et al. "Ultrasonic non-destructive evaluation of composites: A review". In: *Materials Today: Proceedings* 78 (2023). 3rd Biennial International Conference on Future Learning Aspects of Mechanical Engineering (FLAME 2022), pp. 627–632. ISSN: 2214-7853. DOI: <https://doi.org/10.1016/j.matpr.2022.12.055>. URL: <https://www.sciencedirect.com/science/article/pii/S2214785322074296>.
- [9] P. Huthwaite. "Evaluation of inversion approaches for guided wave thickness mapping". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 470.2166 (2014), p. 20140063. DOI: 10.1098/rspa.2014.0063. eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.2014.0063>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rspa.2014.0063>.
- [10] Jing Rao. "Quantitative imaging using ultrasonic tomography based on full waveform inversion". en. PhD thesis. Nanyang Technological University, 2018. DOI: 10.32657/10356/75928. URL: <http://hdl.handle.net/10356/75928> (visited on 11/22/2022).
- [11] P. Huthwaite. "Evaluation of inversion approaches for guided wave thickness mapping". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 470.2166 (2014), p. 20140063. DOI: 10.1098/rspa.2014.0063. eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.2014.0063>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rspa.2014.0063>.
- [12] Alexander Velichko and Paul D Wilcox. "Excitation and scattering of guided waves: Relationships between solutions for plates and pipes". In: *The Journal of the Acoustical Society of America* 125.6 (2009), pp. 3623–3631.
- [13] C.L. Willey et al. "Guided wave tomography of pipes with high-order helical modes". In: *NDT & E International* 65 (2014), pp. 8–21. ISSN: 0963-8695. DOI: <https://doi.org/10.1016/j.ndteint.2014.03.010>. URL: <https://www.sciencedirect.com/science/article/pii/S0963869514000449>.

- [14] Emiel Haseffras, Arno Volker, and Martin Verweij. "Efficient Guided Wave Modelling for Corrosion Mapping via One-Way Wavefield Extrapolation." In: *Sensors* (May 2024). ISSN: 1424-8220.
- [15] J Wong, P Hurley, and GF West. "Crosshole seismology and seismic imaging in crystalline rocks". In: *Geophysical research letters* 10.8 (1983), pp. 686–689.
- [16] ND Bregman, RC Bailey, and CH Chapman. "Crosshole seismic tomography". In: *Geophysics* 54.2 (1989), pp. 200–215.
- [17] Pierre Belanger and Peter Cawley. "Feasibility of low frequency straight-ray guided wave tomography". In: *NDT & E International* 42.2 (2009), pp. 113–119.
- [18] PR Williamson. "A guide to the limits of resolution imposed by scattering in ray tomography". In: *Geophysics* 56.2 (1991), pp. 202–207.
- [19] Pierre Belanger, Peter Cawley, and Francesco Simonetti. "Guided wave diffraction tomography within the born approximation". In: *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 57.6 (2010), pp. 1405–1418.
- [20] P Huthwaite and F Simonetti. "High-resolution imaging without iteration: A fast and robust method for breast ultrasound tomography". In: *The Journal of the Acoustical Society of America* 130.3 (2011), pp. 1721–1734.
- [21] Jean Virieux and Stéphane Operto. "An overview of full-waveform inversion in exploration geophysics". In: *Geophysics* 74.6 (2009), WCC1–WCC26.
- [22] Albert Tarantola. "A strategy for nonlinear elastic inversion of seismic reflection data". In: *Geophysics* 51.10 (1986), pp. 1893–1903.
- [23] R Gerhard Pratt and Michael H Worthington. "INVERSE THEORY APPLIED TO MULTI-SOURCE CROSS-HOLE TOMOGRAPHY. PART 1: ACOUSTIC WAVE-EQUATION METHOD 1". In: *Geophysical prospecting* 38.3 (1990), pp. 287–310.
- [24] Jean Virieux. "P-SV wave propagation in heterogeneous media: Velocity-stress finite-difference method". en. In: *GEOPHYSICS* 51.4 (Apr. 1986), pp. 889–901. ISSN: 0016-8033, 1942-2156. DOI: 10.1190/1.1442147. URL: <https://library.seg.org/doi/10.1190/1.1442147> (visited on 03/04/2024).
- [25] Jean Virieux. "SH-wave propagation in heterogeneous media : Velocity-stress finite-difference method". en. In: *Society of Exploration Geophysicists* (1984).
- [26] Kurt J Marfurt. "Accuracy of finite-difference and finite-element modeling of the scalar and elastic wave equations". In: *Geophysics* 49.5 (1984), pp. 533–549.
- [27] Gabrio Rizzuti et al. "A Non-Linear Full-Wavefield Inversion Method for Ultrasound Imaging in Non-Destructive Testing". In: *GEOPHYSICS* (July 2017). DOI: 10.13140/RG.2.2.22953.29284.
- [28] R Brossier, Jean Virieux, and Stéphane Operto. "Parsimonious finite-volume frequency-domain method for 2-DP–SV-wave modelling". In: *Geophysical Journal International* 175.2 (2008), pp. 541–559.
- [29] Michel Bouchon, Michel Campillo, and Stephane Gaffet. "A boundary integral equation-discrete wavenumber representation method to study wave propagation in multilayered media having irregular interfaces". In: *Geophysics* 54.9 (1989), pp. 1134–1140.
- [30] Guy Chavent. "Identification of functional parameters in partial differential equations". In: *Joint Automatic Control Conference*. 12. 1974, pp. 155–156.
- [31] Horace Lamb. "On waves in an elastic plate". In: *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character* 93.648 (1917), pp. 114–128.
- [32] Kranthi Peddeti and Sridhar Santhanam. "Dispersion curves for Lamb wave propagation in pre-stressed plates using a semi-analytical finite element analysis". In: *The Journal of the Acoustical Society of America* 143.2 (2018), pp. 829–840.
- [33] Andreas Armin Ernst Zimmermann. "Shear-horizontal guided wave tomography". en-US-GB. Accepted: 2022-05-19T11:14:17Z Publisher: Imperial College London. PhD thesis. Imperial College London, Sept. 2021. URL: <http://spiral.imperial.ac.uk/handle/10044/1/96897> (visited on 06/21/2023).

- [34] Armin Huber, researcher at the Center for Lightweight Production Technology (ZLP) of the German Aerospace Center (DLR). *Dispersion Calculator*. Version V1. 2018. URL: https://www.dlr.de/zlp/en/desktopdefault.aspx/tabid-14332/24874_read-61142/.
- [35] Joseph L. Rose. *Ultrasonic Guided Waves in Solid Media*. en. 1st ed. Can also be found under ISBN 978-1-107-27361-0. Cambridge University Press, Aug. 2014. ISBN: 978-1-107-04895-9. DOI: 10.1017/CB09781107273610. URL: <https://www.cambridge.org/core/product/identifier/9781107273610/type/book> (visited on 11/22/2022).
- [36] Jidong Hou, Kevin R Leonard, and Mark K Hinders. "Automatic multi-mode Lamb wave arrival time extraction for improved tomographic reconstruction". In: *Inverse Problems* 20.6 (2004), p. 1873.
- [37] Peter Huthwaite et al. "Mode selection for corrosion detection in pipes and vessels via guided wave tomography". In: *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 60.6 (2013), pp. 1165–1177.
- [38] Erik Luiten. "Wall Thickness Tomography in Complex Geometries with Ultrasonic Guided Waves". PhD thesis. Delft: TU Delft, Apr. 2009.
- [39] Augustinus Johannes Berkhout. *Seismic Migration: Imaging of Acoustic Energy by Wave Field Extrapolation...: Imaging of Acoustic Energy by Wave Field Extrapolation*. Elsevier, 2012.
- [40] Dave Hale. "Stable explicit depth extrapolation of seismic wavefields". In: *Geophysics* 56.11 (1991), pp. 1770–1777.
- [41] Jenő Gazdag. "Wave equation migration with the phase-shift method". In: *Geophysics* 43.7 (1978), pp. 1342–1351.
- [42] Niels Portzgen, Dries Gisolf, and Gerrit Blacquiere. "Inverse wave field extrapolation: a different NDI approach to imaging defects". In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 54.1 (2007), pp. 118–127. DOI: 10.1109/TUFFC.2007.217.
- [43] Edith Jacoba van Veldhuizen. "Integrated approach to 3-D seismic acquisition geometry analysis: Emphasizing the influence of the inhomogeneous subsurface". In: *Geology, Engineering* (2006).
- [44] Jan W Thorbecke, Kees Wapenaar, and Gerd Swinnen. "Design of one-way wavefield extrapolation operators, using smooth functions in WLSQ optimization". In: *Geophysics* 69.4 (2004), pp. 1037–1045.
- [45] Robert Seidl. "Full waveform inversion for ultrasonic nondestructive testing". PhD thesis. Technische Universität München, 2018.
- [46] R.-E. Plessix. "A review of the adjoint-state method for computing the gradient of a functional with geophysical applications". In: *Geophysical Journal International* 167.2 (Nov. 2006), pp. 495–503. ISSN: 0956-540X. DOI: 10.1111/j.1365-246X.2006.02978.x. URL: <https://doi.org/10.1111/j.1365-246X.2006.02978.x> (visited on 11/24/2022).
- [47] Peng Zuo and Peter Huthwaite. "Quantitative mapping of thickness variations along a ray path using geometrical full waveform inversion and guided wave mode conversion". en. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 478.2257 (Jan. 2022), p. 20210602. ISSN: 1364-5021, 1471-2946. DOI: 10.1098/rspa.2021.0602. URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.2021.0602> (visited on 11/22/2022).
- [48] Hussain Hammad and Gary Margrave. "Waveform tomography for areas of complex near surface". In: *CREWES Research Report* (Nov. 2009). DOI: 10.13140/RG.2.2.28186.24005.
- [49] Andreas Fichtner. *Full Seismic Waveform Modelling and Inversion*. en. Advances in Geophysical and Environmental Mechanics and Mathematics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011. ISBN: 978-3-642-15806-3. DOI: 10.1007/978-3-642-15807-0. URL: <https://link.springer.com/10.1007/978-3-642-15807-0> (visited on 03/08/2024).
- [50] Jorge Nocedal and Stephen J. Wright. *Numerical optimization*. 2. ed. Springer series in operations research and financial engineering. New York, NY: Springer, 2006. XXII, 664. ISBN: 978-0-387-30303-1. URL: http://gso.gbv.de/DB=2.1/CMD?ACT=SRCHA&SRT=YOP&IKT=1016&TRM=ppn+502988711&sourceid=fbw_bibsonomy.
- [51] Union Steels. *API 5L Pipe Size Tolerances*. Accessed: 2025-02-20. 2025. URL: <https://www.union-steels.com/newsdetail/api-5l-pipe-size-tolerances.html>.

- [52] S. Vishnuvardhan, A. Ramachandra Murthy, and Abhishek Choudhary. "A review on pipeline failures, defects in pipelines and their assessment and fatigue life prediction methods". In: *International Journal of Pressure Vessels and Piping* 201 (2023), p. 104853. ISSN: 0308-0161. DOI: <https://doi.org/10.1016/j.ijpvp.2022.104853>. URL: <https://www.sciencedirect.com/science/article/pii/S0308016122002381>.

This page is left blank intentionally