Chair of Computational Modeling and Simulation
TUM School of Engineering and Design
Technical University of Munich

Frequency-domain Full Waveform Inversion in Non-destructive Testing

Scientific work to obtain the degree

Bachelor of Science (B.Sc.)

at the TUM School of Engineering and Design
of the Technical University of Munich.

Supervised by Tim Bürchner M.Sc.
Divya Shyam Singh M.Sc.
Chair of Computational Modeling and Simulation

Submitted by Aymen Hajri (03730058)
Helene-Mayer-Ring 7A/1011
D-80809 München
e-Mail: aymen.hajri@tum.de

Submitted on June 2, 2025

Abstract

In non-destructive testing (NDT) applications, Full Waveform Inversion (FWI) is a high-resolution imaging technique that uses complete ultrasonic wavefields to reconstruct material properties. Unlike most ultrasonic inspection approaches, FWI does not rely on primitive assumptions; rather, it reduces the difference between the measured and simulated wavefields, hence using full information for detailed material characterization. This thesis focuses on developing a two-dimensional frequency domain FWI framework. Wave propagation is modeled with the Helmholtz equation, while gradients for model updates are calculated using the adjoint-state method. BFGS optimization is used to ensure stable and efficient convergence. The Inversion is accomplished via multi-frequency techniques for increased robustness with two approaches: sequential frequency stepping from low to high frequencies and simultaneous inversion using all frequencies in each iteration. The work illustrates the capabilities of frequency domain FWI as a concrete and accurate method for ultrasonic inspection, demonstrating its potential for high-resolution material characterization in NDT applications.

Zusammenfassung

In der zerstörungsfreien Prüfung (ZfP) ist die Full Waveform Inversion (FWI) eine moderne Hochauflösungstechnik, die vollständige Ultraschallwellenfelder nutzt, um Materialeigenschaften präzise zu rekonstruieren. Im Gegensatz zu klassischen Ultraschallmethoden beruht FWI nicht auf vereinfachenden Annahmen. Stattdessen vergleicht sie gemessene und simulierte Wellenfelder und minimiert deren Unterschiede. So wird das volle Informationspotenzial der Daten genutzt, um ein besonders detailliertes Bild des Materials zu erzeugen.

Diese Arbeit entwickelt einen zweidimensionalen FWI-Ansatz im Frequenzbereich. Die Wellenausbreitung wird mithilfe der Helmholtz-Gleichung beschrieben, und die Gradienten zur Aktualisierung des Modells werden mit der Adjoint-State-Methode berechnet. Für eine stabile und effiziente Konvergenz kommt die BFGS-Optimierung zum Einsatz.

Die Inversion erfolgt mithilfe sogenannter Multi-Frequenz-Techniken, um die Robustheit zu erhöhen. Dabei werden zwei Strategien untersucht: zum einen die sequentielle Nutzung einzelner Frequenzen von tiefen zu hohen, zum anderen die gleichzeitige Verwendung aller Frequenzen in jeder Iteration.

Diese Arbeit zeigt, dass FWI im Frequenzbereich eine zuverlässige und präzise Methode zur Ultraschallprüfung ist und großes Potenzial für die hochauflösende Materialcharakterisierung in der ZfP bietet.

Contents

List of Abbreviations

2D	Two-Dimensional
BFGS	Broyden-Fletcher-Goldfarb-Shanno
CSC	Compressed Sparse Column
FWI	Full Waveform Inversion
L-BFGS	Limited-memory BFGS
LIL	List-of-Lists
LU	Lower-Upper (matrix factorization)
NDT	Non-Destructive Testing

Chapter 1

Introduction

1.1 Motivation and Application

First proposed in the 1980s by A. Tarantola and B. Valette [1] , Full Waveform Inversion (FWI) is a seismic imaging method that seeks to retrieve high-resolution information from seismic data using full waveform content. Unlike classical methods such as ray-based tomography or migration, which make critical physical oversimplifications regarding wave propagation, FWI utilizes the complete physics of the wave equation by iteratively updating subsurface models to reduce the difference between recorded and simulated wavefields [2].

The increasing feasibility of computational power and numerical methods has led to more simulation and waveform-based inversion techniques. In particular, academic and industrial applications of FWI have surged in recent years, owing to its effectiveness in regions that necessitate high-resolution subsurface models, such as hydrocarbon exploration, geotechnical engineering, and non-destructive testing. [3]

The numerical solution of the wave equation is one essential part of FWI, which can be approached in the time or frequency domain. For this thesis, the frequency-domain formulation was chosen because of its superior speed, lower memory capacity requirements, and better integration with multi-scale inversion methodologies [4]. Unlike time-domain simulations, frequency-domain techniques compute solutions to wave equations at discrete frequency intervals, thus enabling a more focused inversion process directed at the most informative sections of the spectrum. This aids in gradient computation by eliminating the need for reverse-time wavefield propagation and state-dependent wavefield memory, and it allows for tiered inversion methods that gradually shift to higher frequency bands for fine detail structural imaging. These beneficial aspects are crucial in ultrasonic nondestructive testing within small, confined volume spaces, dominated by limited scope, control, and complexity algorithms. [3], [4]

In the context of non-destructive testing (NDT), the development of effective frequency-domain FWI frameworks requires the use of a computational wave-based approach. This thesis follows the core principle of waveform inversion, which is building internal models by using full internal material characteristics like amplitude, phase, scattering, and multiple arrivals. Through modern computational facilities, simulation of wave propagation in complex heterogeneous materials is possible, thus making frequency-domain FWI a practically and accurately applicable tool for ultrasonic inspections in NDT [5].

The development of a two-dimensional frequency domain FWI framework is done in this thesis. It focuses on formulating the forward problem as a two-dimensional equation that utilizes the Helmholtz equation, and solves it numerically through finite difference discretization [6]. Gradient computation is done via the adjoint-state method [1], and optimization is executed with the quasi-Newton Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [7]. To enhance robustness and convergence, multi-frequency and multi-source information is added [8], [9]. The framework is tested with synthetic data, and complex subsurface velocity structures are successfully reconstructed [4].

1.2 Goal of the Thesis

This thesis aims to create and analyze the methodology of FWI in the frequency domain and its applicability towards refining velocity models for ultrasonic testing within the context of NDT [2], [5]. Emphasis is placed on developing a modular, self-sufficient inversion algorithm through the construction of a system that includes numerical simulation of acoustic wavefield, adjoint-state gradient computation [1], and iterative modifications of model building [6].

For this type of problem, the forward problem is modeled using the scalar wave equation, representing the propagation of waves through solid media relating to the NDT domain [2]. The computational model relies on finite difference techniques for solving the scalar Helmholtz equation within the frequency domain [9]. Each state of the system provides an estimate of the misfit function computed between observed and estimated data of ultrasonic wavefields and provides the gradient for the model update [6].

Multi-source and multi-frequency inversion are implemented to enhance the quality of the initial model and the final inversion results [3], [8]. This is done through the optimization technique based on the quasi-Newton BFGS approach, with proven efficiency in serial, highly dimensional inverse problems due to its fast convergence [7]. With these methods, sophisticated internal disruptions, velocity variances, and specific structures of the sample under test can be accurately determined.

Evaluating the algorithm is done through synthetic ultrasonic experiments, which check how accurately the system reconstructs the elastic property changes and verifies gradient computation under practical NDT scenarios [4], [10].

1.3 Overview

Chapter 2 presents the main concepts and theorems which serve as the basis of this thesis. It starts by describing the modeling of ultrasonic wave propagation in solids, beginning with the Helmholtz equation in the frequency domain. The chapter also covers the main ideas of FWI, such as the calculation of gradients using the adjoint-state method and model updates using a quasi-Newton approach for optimization. It also presents two important approximations for inversion: sequential and simultaneous use of frequencies, and explains how these are incorporated in the inversion process with BFGS optimization.

In Chapter 3, the focus shifts to the actual implementation of the FWI framework. It follows up with details on the structure of the code, which is implemented in Python, and on the solution of the Helmholtz equation by finite differences. The explanation includes the configuration of wave source and receiver locations, as well as the construction of the system matrix. Gradients are computed using the adjoint-state method and supplied to a BFGS optimizer. The chapter also describes how different configurations, like single vs. multiple sources and strategies for frequency use, are incorporated into the inversion loop.

Chapter 4 discusses the results from the synthetic tests aimed at evaluating the framework. Four different inversion setups are compared in terms of tuning accuracy, convergence speed, and numerical stability. This comparison illustrates the effects of using all frequencies at once versus adding them incrementally, as well as the impact of using single or multiple data sources on reconstruction quality.

Finally, Chapter 5 summarizes the results of this work and provides an outlook for potential future research. Suggested improvements include applying the framework to real experimental data, extending it to include more complex physical models, or integrating multiple inversion techniques to optimize both speed and accuracy.

Chapter 2

Theoretical Background

2.1 Ultrasonic Wave Propagation in Nondestructive Testing

Ultrasonic NDT utilizes elastic waves as probing instruments on the internal structure of solid materials. Often, transducers serve as sources of the waves that move through metals, composites, and other engineering materials. By measuring the wavefield response at the boundaries of the materials, the aim is to identify internal flaws, changes in the physical properties, or structural imbalances that exist within the material [3], [6].

The motion of waves in solids is defined in the context of the elastic wave equation, which is formed based on the principles of Newton's laws and other laws such as Hook's law [2], [11].

2.1.1 Wave Types in Ultrasonic Testing

Ultrasonic testing primarily involves bulk waves, namely compressional waves or P-waves, which propagate fastest and cause particle motion in the direction of travel, and shear waves or S-waves, which travel more slowly and move particles perpendicular to the direction of propagation, only in solids [2], [11].

While using FWI, the wave propagation is simplified as much as possible by making an acoustic approximation. In such cases, shear, mode conversion, and other effects are ignored. This is useful in balancing the amount of computation needed while maintaining important capabilities needed for imaging [6]. This is particularly true when the shear effects are weak or when there is limited data [4].

2.1.2 Frequency-Domain Modeling and the Helmholtz Equation

In NDT methods, the propagation of ultrasonic waves within solid structures is usually simulated in the frequency domain to improve computational speed and facilitate multi-frequency inversion techniques [2], [9]. The time-domain representation of the elastic wave equation can be transformed to its frequency-domain form, governed by the Helmholtz equation, which describes wavefields in steady state [6]:

$$(\omega^2 + c^2(x)\nabla^2)u(x,\omega) = f(x,\omega) \tag{2.1}$$

Where:

- $u(x, \omega)$: complex-valued displacement wavefield
- $f(x, \omega)$: frequency-domain source term
- ω: angular frequency
- $c(x)$: spatially varying wave speed in the material

This representation follows the parameterization from the numerical inversion implementation to solve for $c(x)$, which denotes wave speed. The equation is solved at discrete frequencies and source positions, where each solution corresponds to a medium response in steady state. These wavefields are then applied in the iterative process of the inversion to minimize the difference between simulated and observed data [8].

Unlike in seismic applications, where model boundaries must not induce artificial reflections, in ultrasonic testing, reflective boundaries are often relevant and physically present in small-scale test samples. Consequently, in this context, reflective boundary conditions are purposefully applied in order to emulate echo effects characteristic of ultrasonic NDT measurements [3].

FWI requires the simulation of accurate wavefields, and according to Fichtner [6], wavefield computation has a significant impact on the quality of the inversion. In ultrasonic testing, this is particularly important since there is a need to resolve small-scale features such as cracks or inclusions; the modeling of wave propagation and reflection needs to be performed with utmost accuracy.

To consider wave attenuation in practical materials, real-valued wave speed, $c(x)$, is transformed to a complex value:

$$c(x) = c_{\text{real}} + i \cdot \gamma \tag{2.2}$$

where γ is a small positive limit that dissipates energy. Such representation adds a system matrix of complex form in the Helmholtz equation, which, in turn, adds complexity to the simulation. It enables a more accurate modeling of the behavior of ultrasonic waves in attenuating media [2], [4].

2.2 Full Waveform Inversion Principles

FWI is an advanced imaging technique, or rather, a method that refines the spatial representation of geophysical features by reconstructing geophysical waveforms to fit the data. FWI is now commonly used in ultrasonic NDT, as it can help image subtle structural details such as cracks, inclusions, and complex internal defects [2], [5].

Some major differentiating attributes of FWI from classical imaging techniques include the use of approximate physics or the use of merit attributes from a partial waveform. Unlike these methods, FWI utilizes each recorded wavefield in its totality, using amplitude, phase, and interference of all frequencies from all sensors placed in each available position [2], [8], [10].

FWI can be expressed in both the time and frequency domains. The current paper implements the latter because of lower computational resources required for each frequency compared to the time domain, easier integration of multi-frequency inversion techniques, and less complex gradient formulation where wavefields are steady and harmonic [6], [9].

The inversion process comprises an iterative loop consisting of four basic steps. First, forward modeling is done to get the predicted wavefields by simulating the wave propagation through the current wavespeed [2]. Then the residual wavefield, which is the difference between the measured data and the calculated data, is computed, and the corresponding adjoint wave equation is solved in the state-based calculation [1], known as the adjoint-state calculation. After that, in gradient calculation, the misfit function's sensitivity concerning the model parameters and the forward adjoint is obtained with the help of the forward and adjoint wavefields [6]. Ultimately, in optimization, the model is updated by employing an optimization method, like the quasi-Newton BFGS, which attempts to decrease the error or residual [7]. In this way, FWI systematically tunes the model until there is a better match between the model's ultrasonic wavefield data and the recorded data.

2.2.1 Forward Modeling of the Inversion Problem's Wave Equation

As described in Equation (2.1), wave propagation in the frequency domain is governed by the Helmholtz equation. In the context of FWI, this equation forms the computational core of the forward modeling step. Upon discretization, it results in a complex-valued linear system:

$$A(c, \omega)\, u = f, \tag{2.3}$$

where $A(c, \omega) \in \mathbb{C}^{N \times N}$ is the system matrix that depends on the spatial distribution of the wave speed $c(x)$, $u \in \mathbb{C}^N$ is the complex wavefield at angular frequency ω, and $f \in \mathbb{C}^N$ is the source vector [6].

During the wave propagation phase of bounded solid specimens, ultrasonic NDT, as one of the measures done in ultrasonic inspection, reflections from material boundaries are intrinsic to the processes of measurements. Thus, it has been deliberated in this work that reflective boundary conditions need to be applied in order to replicate the actual behavior echoes in ultrasonic inspections [3].

For each frequency and source position, the forward problem is solved, generating synthetic wavefields which are compared to observed data [4].

2.2.2 Finite Difference Discretization of the Helmholtz Equation

2.2.2.1 Gradient Evaluation

The last step in the adjoint-state method consists of finding the gradient of the misfit function. This correlation is done between the adjoint wavefield and the Laplacian of the forward wavefield. The quantity obtained shows how local changes in wavespeed modulate the wavefield that is recorded [2], [6].

The Laplacian of the forward wavefield is given by [6]:

$$\Delta u = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}, \tag{2.4}$$

which can be computed using second-order central finite differences [6]:

$$\frac{\partial^2 u(i,j)}{\partial x^2} \approx \frac{u(i+1,j) - 2u(i,j) + u(i-1,j)}{\Delta x^2}, \tag{2.5}$$

$$\frac{\partial^2 u(i,j)}{\partial y^2} \approx \frac{u(i,j+1) - 2u(i,j) + u(i,j-1)}{\Delta y^2}. \tag{2.6}$$

This finite-difference approximations are used to compute Δu_i at every grid point i and then substitute into the gradient expressions from equations (2.11) and (2.12).

2.2.2.2 Discrete Helmholtz Equation Using Finite Differences

The discrete linear system outlined in Equation (2.3) is formed by implementing second-order finite-difference methods on the spatial derivatives in the Helmholtz equation. The Laplacian operator is evaluated using a five-point stencil on a uniform Cartesian mesh.

As previously defined in Equations (2.5) and (2.6), the second spatial derivatives are estimated using central differences. Substituting these expressions into the continuous Helmholtz equation gives the following discrete formulation at every grid point (i,j) [6]:

$$\omega^2 u_{i,j} + c_{i,j}^2 \left[\frac{u_{i+1,j} - 2u_{i,j} + u_{i-1,j}}{\Delta x^2} + \frac{u_{i,j+1} - 2u_{i,j} + u_{i,j-1}}{\Delta y^2} \right] = f_{i,j} \tag{2.7}$$

The distinction of the discrete Helmholtz operator $A(c,\omega)$ lies in the structure of this finite-difference expression. For each row, matrix A represents a particular grid point and contains five non-zero coefficients: one from the diagonal and four representing the off-diagonal terms related to its direct neighbors. The coefficients assigned to the diagonal result from the aggregation of all central difference contributions and the frequency-dependent term ω^2, while the off-diagonal terms represent the spatial coupling of neighboring grid points, scaled by the local squared velocity [6].

2.2.3 Calculate Adjoint-State Gradient and Misfit Function

The central goal of FWI is to iteratively reduce the difference between observed and simulated data. This is quantified using a misfit function, typically formulated as a least-squares objective [1], [2], [6].

The goal of FWI is to find the wave speed model $\mathbf{c}^*$ that minimizes the misfit function Φ [6]:

$$\mathbf{c} = \arg\min_{m} \Phi. \tag{2.8}$$

$$\Phi(\mathbf{c}) = \frac{1}{2} \sum_{s=1}^{N_s} \sum_{r=1}^{N_r} \sum_{\omega \in \Omega} \| d_{\text{obs}}(s,r,\omega) - d_{\text{syn}}(s,r,m,\omega) \|_2^2 \tag{2.9}$$

where :

- $d_{\text{obs}}(s,r,w)$: Observed data for source s, receiver r, and frequency w.
- $d_{\text{syn}}(s,r,m,w)$: Synthetic (simulated) data, which depends on the model parameters m (and indirectly on wave speed c, since $c = \arg\min_m \Phi$).
- N_s: Number of sources.
- N_r: Number of receivers.
- Ω: Set of frequencies.
- $\|\cdot\|_2^2$: Squared L2-norm, meaning the squared difference between observed and synthetic data, summed over all components.

To compute the gradient of the misfit function efficiently, the adjoint-state method is used. This avoids the explicit computation of the Jacobian matrix $\partial d / \partial c$ and instead exploits the interaction between forward and adjoint wavefields [2].

The adjoint wavefield λ is obtained by solving the adjoint wave equation:

$$A(c,\omega)^{\dagger} \lambda = d_{\text{syn}} - d_{\text{obs}}, \tag{2.10}$$

where $A(c,\omega)^{\dagger}$ is the Hermitian transpose of the system matrix A, and the right-hand side is the data residual acting as a virtual source [6].

The gradient of the misfit function χ with respect to the real and imaginary parts of the complex wave speed c_i at grid point i is expressed as [6]:

$$\frac{\partial \chi}{\partial \operatorname{Re}(c_i)} = 4 \operatorname{Re}\left(c_i \lambda_i^* \Delta u_i\right), \tag{2.11}$$

$$\frac{\partial \chi}{\partial \operatorname{Im}(c_i)} = -4 \operatorname{Im}\left(c_i \lambda_i^* \Delta u_i\right), \tag{2.12}$$

Where :

- λ_i^*: complex conjugate of the adjoint wavefield at grid point i
- Δu_i: discrete Laplacian (second spatial derivative) of the forward wavefield at point i
- c_i: complex wave speed at point i
- $\mathrm{Re}\{\cdot\}$ and $\mathrm{Im}\{\cdot\}$: real and imaginary parts, respectively

This formulation links the forward and adjoint wavefields through the second spatial derivative operator, allowing the gradient to be computed efficiently without forming the full Jacobian matrix.

2.2.3.1 Residuals and Adjoint Source Vector

Assessing the data residuals, which determine the difference measuring gap between the simulated and observed wavefield data, is the starting point of the gradient calculation [5], [8].

$$\mathrm{r} = d_{\text{syn}} - d_{\text{obs}} \tag{2.13}$$

In the adjoint simulation, these residuals are treated as pseudo-point sources. Each residual is mapped onto the computational grid at its corresponding receiver location and used to construct the adjoint source vector r. So the residuals can be used directly without further scaling [1], [2].

2.2.3.2 Solving the Adjoint Equation

The adjoint wavefield λ is computed by solving a linear system based on the Hermitian transpose of the discretized Helmholtz operator [6]:

$$A(c,\omega)^{\dagger}\lambda = r, \tag{2.14}$$

where $A(c,\omega)^{\dagger}$ is the Hermitian transpose of the system matrix A, and the right-hand side is the data residual acting as a virtual source.

A sparse direct solver is typically used to solve this complex-valued system. Once the adjoint wavefield λ is obtained, it is reshaped into a 2D grid to match the spatial layout of the forward wavefield [2].

2.2.4 Optimization Strategy and Multi-Frequency Workflow

Once the gradient is evaluated, the next step is to update the model using an optimization algorithm. Due to the non-linear nature of the misfit function and the high dimensionality of the parameter space, quasi-Newton methods are commonly used. One widely used method is the Limited-memory BFGS (L-BFGS) algorithm, which computes model updates via [7], [10]:

$$m_{k+1} = m_k - \alpha_k H_k^{-1} \nabla \Phi(m_k) \tag{2.15}$$

To improve convergence behavior and reduce the risk of local minima, two frequency selection strategies are evaluated in this work. The first variant applies a multi-frequency stepping approach, beginning with low-frequency components to capture large-scale features and progressively introducing higher frequencies to resolve fine details. The second variant uses all frequency components simultaneously in each inversion iteration. This comparison allows for evaluating the influence of frequency scheduling on reconstruction quality and convergence. The impact of the two frequency-scheduling strategies on convergence and reconstruction quality is analysed in Chapter 4.

2.3 Methods of Optimization

This section presents the implementation details of the inversion process used in this thesis. The model is updated using the BFGS optimization algorithm, chosen for its favorable convergence properties in large-scale nonlinear problems.

2.3.1 Optimization Using BFGS and Quasi-Newton Methods

FWI is inherently a non-linear inverse problem with a highly complex solution space. The misfit function often presents a non-convex landscape riddled with sharp local minima and non-trivial curvature. First-order methods like steepest descent or conjugate gradient typically suffer from slow convergence or unstable updates when applied to such problems, especially when large models and datasets are involved. In such cases, quasi-Newton methods are often preferred, as they provide approximate second-order information while avoiding the high cost of computing and storing the full Hessian matrix [6], [7].

The BFGS algorithm is one of the most widely used quasi-Newton methods in numerical optimization. It incrementally builds an approximation to the inverse Hessian matrix by incorporating gradient and model differences from previous iterations. This matrix is then used to calculate a descent direction aligned with the local curvature of the objective function, allowing for more efficient convergence than first-order methods [7].

In this thesis, a full-memory implementation of BFGS is adopted. The scale of the problem, a two-dimensional synthetic acoustic FWI, is moderate in terms of computational demands, which makes full BFGS both feasible and beneficial. The richer curvature information retained by full BFGS leads to more accurate model updates. The update rule introduced in Section 2.2 forms the basis of the inversion process here.

The optimization is implemented using the BFGS option from the SciPy library, which makes use of the full-memory BFGS algorithm. This offers a consistent and extensively validated software for calculating search directions and carrying out line minimization [5].

In contrast, the L-BFGS variant is commonly used in large-scale three-dimensional FWI due to its reduced memory requirements [9], [10]. However, for the moderate-scale problems addressed in this work, full BFGS offers superior convergence characteristics.

Figure 2.1 shows that the BFGS algorithm achieves steady and stable convergence, which supports its suitability for this moderate-scale 2D inversion problem.

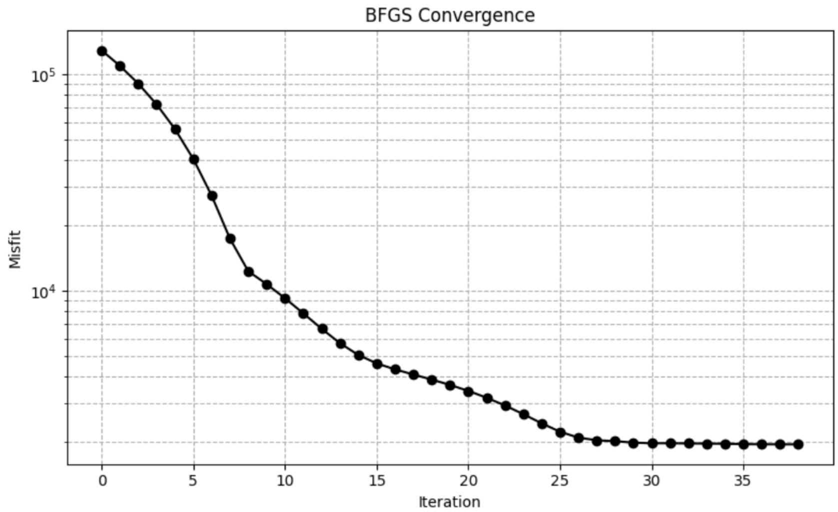

Figure 2.1: BFGS convergence curve showing consistent misfit reduction across iterations.

2.3.2 Multi-Frequency Inversion Strategy

The risk of convergence to local minima is one of the major challenges in FWI, especially when the starting model fails to adequately represent low-wavenumber features. A widely used technique to address this is frequency continuation, or multi-frequency inversion.

This strategy takes advantage of the fact that low-frequency wavefields predominantly contain long-wavelength energy, which describes the large-scale background structure of the model. These frequencies are less sensitive to smaller details, making them less prone to phase ambiguities and cycle skipping. Cycle skipping refers to situations where the inversion becomes trapped in the wrong phase due to large initial misfits between synthetic and observed data [2], [5], [8].

In this thesis, two multi-frequency methods are examined:

The initial approach takes on a traditional frequency stepping method, which begins with the lowest available frequencies for inversion. The frequencies are first used to build a smooth, large-scale model. Then, gradually higher frequencies are incorporated into the model, and with each cycle, finer details are added to the model. This approach follows a smooth-to-fine hierarchy that lowers non-linearity, increases stability, and improves convergence. For every frequency or band, multiple BFGS iterations are conducted, and for each, misfit reduction must reach a standstill. The model undergoes all the stages subsequently, ultimately leading to precise high-resolution reconstructions.

The second strategy simultaneously employs all available frequencies throughout the inversion process. This approach is more direct because there is no logic to the frequency-staging process since every frequency will be applied during each iteration. This method helps accelerate the workflow towards the inversion and is used as a control to judge the impact staged frequency layering has on convergence and the quality of the model.

This configuration enables us to evaluate how frequency timing affects the performance of inversions. As seen in synthetic data studies and field-data experiments, the frequency-stepped approach is common practice [5]. Conversely, the full-frequency approach sheds light on practical tradeoffs associated with the method. Within the bounds of fully synthetic testing scenarios, all strategies can be deployed alongside clearly delineated bounds on frequency content, which allows for a thorough comparison.

2.3.3 Integration with the Inversion Workflow

The optimization method, which was detailed in Section 2.3, has the optimization integrated within the inversion framework presented in Section 2.2. Each frequency stage is executed in a canonical FWI loop of forward modelling, residual extraction, gradient extraction, and model updating. This allows the model to be progressively fine-tuned from low to high frequencies in a single-step procedure.

Chapter 3

Implementation

This chapter details the execution of the two-dimensional acoustic FWI framework in the frequency domain. It stems from the theoretical work done in Chapter 2. The inversion process is constructed using gradient-based optimization with the adjoint-state method:

Within the context of FWI, the objective of the method is to reconstruct high-resolution velocity models by minimizing the difference between the observed wavefield data and the one that is simulated. This chapter outlines the implementation of the inversion framework, focusing on how the simulation, optimization, and inversion components are structured and connected within the code.

3.1 Overview of the Framework

The Full Waveform Inversion FWI framework developed in this thesis is written entirely in Python. It addresses two-dimensional acoustic FWI problems in the frequency domain with gradient-based optimization using the adjoint-state method described in Chapter 2.2.

The codebase is structured into components: forward modeling, misfit computation, adjoint gradient evaluation, optimization, and visualization. Each module can be tested or modified independently.

The inversion loop implements the standard FWI workflow across discrete frequencies, as described in Chapter 2.2. The code is configurable via dedicated parameter files, allowing users to control grid size, frequency selection, source-receiver geometry, boundary conditions, and iteration limits.

Gradient validation tools verify the adjoint-state implementation using finite-difference checks. Visualization routines assist in analyzing velocity models, residuals, and gradients throughout the inversion process.

Python's ecosystem as NumPy, SciPy, and Matplotlib provides efficient handling of large numerical arrays, sparse solvers, and dynamic plotting. Although this version is limited to 2D acoustic media, the design is adaptable to more advanced physics (e.g., elastic or visco-acoustic formulations) and alternative optimization schemes (e.g., L-BFGS, stochastic methods).

The overall workflow of the inversion framework is summarized in Figure 3.1. It illustrates the core loop structure applied across multiple frequencies, including forward and adjoint simulations, gradient computation, and model updates using the BFGS algorithm. The sequential loop iterates over frequencies from low to high, while the simultaneous approach processes all frequencies in each iteration.

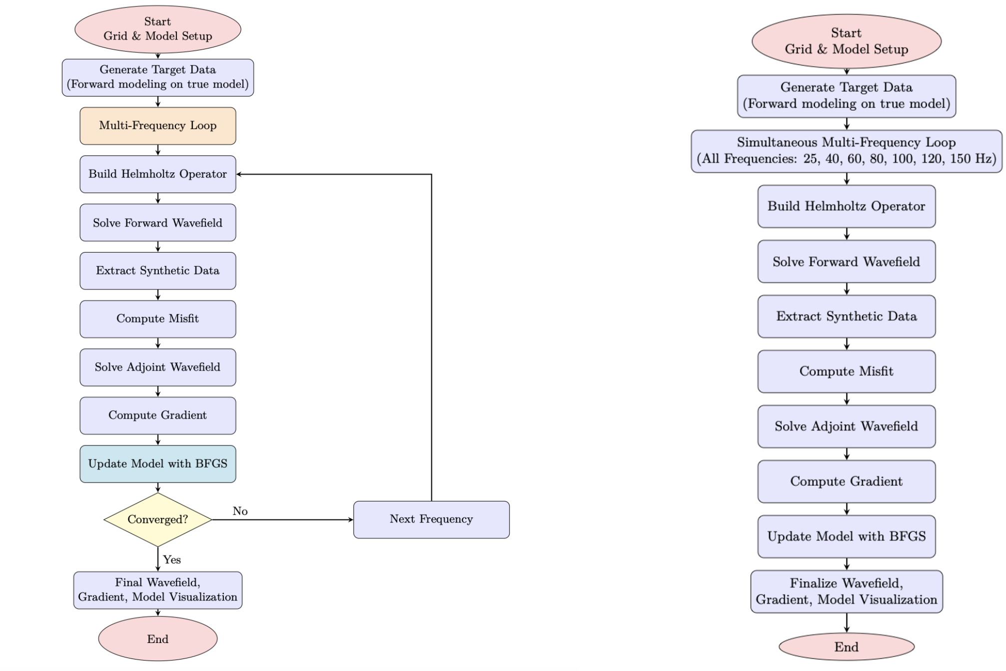

Figure 3.1: Flowchart of the FWI framework for sequential (left) and simultaneous (right) multi-frequency inversion. Both workflows involve forward and adjoint modeling, misfit and gradient computation, and model updating with BFGS.

3.2 Forward Solver: The Finite Difference Method of the Helmholtz Equation

Within the framework of frequency-domain FWI constructed in this study, the forward modeling step requires solving the Helmholtz equation for a specific frequency and velocity model. This chapter describes the implementation of this task through finite difference discretization techniques, which were previously introduced in Chapter 2 as part of the problem statement.

3.2.1 Numerical Formulation

The forward modeling problem addressed in this thesis is governed by the frequency-domain Helmholtz equation, provided again as Equation (2.1). This equation concerns the propagation of time-harmonically oscillating acoustic waves in a medium, which should be considered as the response of the medium to a periodic input at a particular frequency. As an elliptic partial differential equation, it is appropriate for frequency-domain simulations in which several predefined frequencies are computed directly [2], [11].

By partitioning the spatial domain into sections and implementing finite-difference methods on the differential operators, the Helmholtz equation can be expanded into a system of linear equations. This discrete form was presented as Equation (2.3) [6].

The discrete Helmholtz system $A(c,\omega)u = f$, introduced in Chapter 2, forms the core of the forward modeling step. Here, A is a sparse, complex-valued matrix assembled using second-order finite differences, u is the complex wavefield, and f is the source term. The matrix is constructed using a five-point stencil as described in Section 2.2.2, and supports efficient simulation across multiple frequencies [5], [9].

Figures 3.2 and 3.3 depict the velocity models corresponding to the inversion process.

The model for the target velocity distribution, shown in Figure 3.2, features a low-velocity square anomaly embedded in a background velocity of 3000 m/s. The perturbing anomaly is the target of reconstruction through the inversion process. Source and receiver locations are shown with blue dotted markers and are distributed on a circular array around the model domain.

On the other hand, Figure 3.3 displays the initial, homogeneous velocity model used to start the inversion, which lacks any internal anomaly. The absence of internal structure ensures that the inversion begins from a smooth, unbiased reference model. The target model presents a stark contrast to the initial model, underscoring the challenge faced by the inversion algorithm.

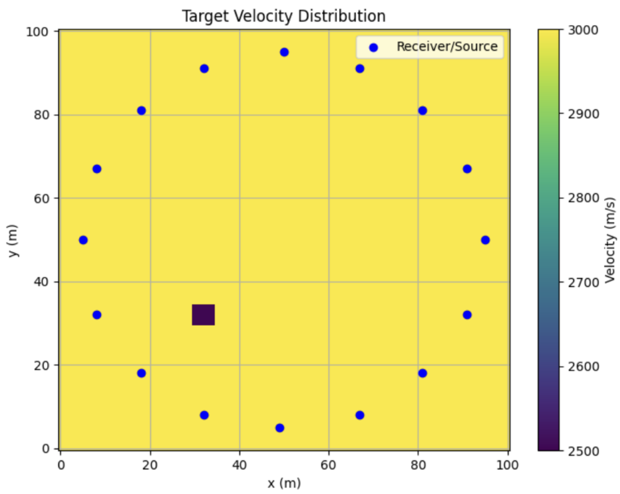

Figure 3.2: Target velocity model used in the inversion. The square anomaly represents a low-velocity inclusion within a homogeneous background of 3000 m/s. Blue dots indicate the receiver and source locations arranged on a circular array around the domain.

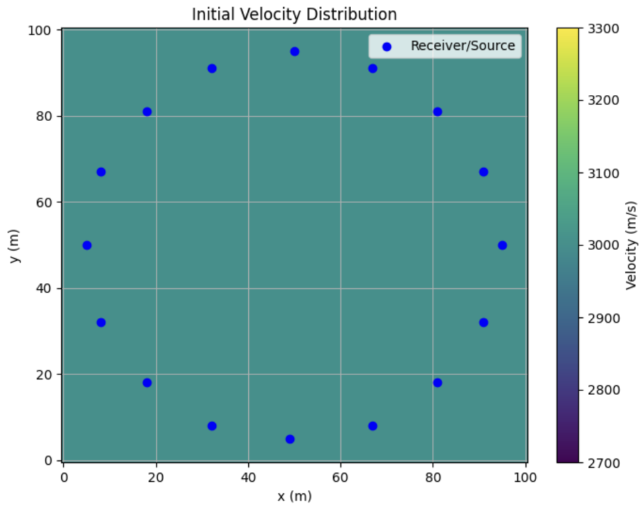

Figure 3.3: Initial velocity model used to initialize the Full Waveform Inversion. The model is homogeneous and lacks the internal anomaly, providing a smooth starting point for the inversion.

3.2.2 Discrete Matrix Construction

To program the discrete Helmholtz operator, each grid point (i, j) is translated into a one-dimensional linear index according to

$$\mathtt{idx} = i \cdot n_y + j \tag{3.1}$$

This mapping allows the global system matrix to be stored and manipulated using sparse matrix formats. Initially, the matrix is assembled using the List-of-Lists (LIL) representation, which enables efficient insertion of non-zero entries. Once the assembly is complete, it is converted to Compressed Sparse Column (CSC) format to improve memory efficiency and solver performance [2], [11].

The practical method for constructing the system matrix is outlined in Algorithm 1. In this procedure, local wave speeds are computed, diagonal and off-diagonal entries are assigned based on the finite-difference stencil, and neighbor connectivity is managed using index offsets [9].

Algorithm 1 Discrete construction of the Helmholtz system matrix

1: **for** each grid point (i, j) in the computational domain **do**
2: Compute linear index: $\texttt{idx} \leftarrow i \cdot n_y + j$
3: Compute squared wave speed: $c^2 \leftarrow c[i, j]^2$
4: Set diagonal entry:
$$A[\texttt{idx}, \texttt{idx}] \leftarrow \omega^2 - 2 \cdot c^2 \cdot \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} \right)$$
5: Set horizontal off-diagonals:
$$A[\texttt{idx}, \texttt{idx} \pm 1] \leftarrow \frac{c^2}{\Delta x^2}$$
6: Set vertical off-diagonals:
$$A[\texttt{idx}, \texttt{idx} \pm ny] \leftarrow \frac{c^2}{\Delta y^2}$$
7: **end for**
8: Convert matrix from LIL to CSC format

3.2.3 Numerical Implementation

3.2.3.1 Boundary Conditions

To limit the scope of the computations and eliminate nonphysical reflections, boundaries are set. This implementation uses homogeneous Dirichlet conditions on all sides, meaning the wavefield is nullified on the outer edges of the grid. This approach reduces the computational effort by simplifying the system matrix [2], [11].

Though this boundary condition is useful, it poses a severe constraint. The zero wave boundary condition means there will be no reliable field values in the vicinity of the boundary pores. To mitigate that, it does not affect the inversion, all sources and receivers are placed far enough from the boundaries. This alternative is sensible for synthetic tests where such control over the layout is available.

More sophisticated boundary treatments, like the PMLs, would provide greater freedom with regard to receiver placement, as well as increased absorption. The chosen method, however, strikes a reasonable compromise between ease of use and accuracy, given the scope of the experiments conducted.

3.2.3.2 Source Term and Receiver Modeling

The source term component is assumed to be located at the source position with a very high amplitude delta function and is modeled as put into the grid point closest to a given source location, used as a monitor. This simulates an idealistic case of a frequency-domain impulsive source [8].

Receiver positions are defined as part of a group of grid points that are located in the form of a circle or have circular coordinates centered around the domain center. After computing the provided system, the values for the wavefield at these receiver points are evaluated to create synthetic data that will be later utilized in misfit estimation and inversion.

3.2.3.3 Solver Strategy

When the matrix A is accompanied by the source vector f, the forward problem is simplified to determining the solution of the sparse linear system posed in Equation (2.3). Depending on the number of sources and performance estimations, two solver strategies from SciPy are implemented.

In the case of rather basic problems or single-source cases, a direct approach solver is adopted. Here, we "sparsely solve" the system using:

$$u = \text{spsolve}(A, f) \tag{3.2}$$

In contrast, in scenarios where several right-hand sides are present for example, in multi-source or simultaneous source simulations, the solution is developed more efficiently. The matrix A is LU-factorized, and a reusable solver object is created using:

$$\text{LU} = \text{factorized}(A) \tag{3.3}$$

Subsequent wavefield computations are then performed with reduced computational cost by:

$$u_k = \text{LU}(f_k) \tag{3.4}$$

The index k denotes the source number. This process is beneficial in frequency-domain FWI applications involving multiple frequencies and sources, as it uses factorization to improve performance, solving the same matrix repeatedly while the model and frequency remain constant [2], [9].

3.3 Adjoint-State Gradient Calculation

With the use of the adjoint-state technique in FWI, gradient computation becomes efficient, as it necessitates only a single wavefield simulation per source, both forward and adjoint. This efficiency greatly surpasses that of direct finite-difference methods, which incur a mounting computational cost stemming from the need for distinct forward simulations for each parameter in the model [6].

Recalling what was discussed in Chapter 2, the metric of evaluation, the misfit function defined in Equation (2.9), quantifies the difference between data that was observed and data that was simulated. Its gradient with respect to speed c_i, which is taken to be a complex number, is derived from Equations (2.11) and (2.12), treating the constituents separately as real and imaginary parts.

The adjoint wavefield λ is solved from the linear system in Equation (2.10) by employing the data residuals as fictitious sources. These sources are placed at the receiver positions and retro-propagated through the medium to calculate the adjoint field.

Figure 3.4 shows the gradient magnitude, which attains its highest sensitivity in perceiving changes in the wave speed. The regions in Figure 3.4 that have the highest sensitivity are those where the misfit function unilaterally changes and are therefore marked in bright colors. The blue markers correspond to the locations of the sources and the receivers.

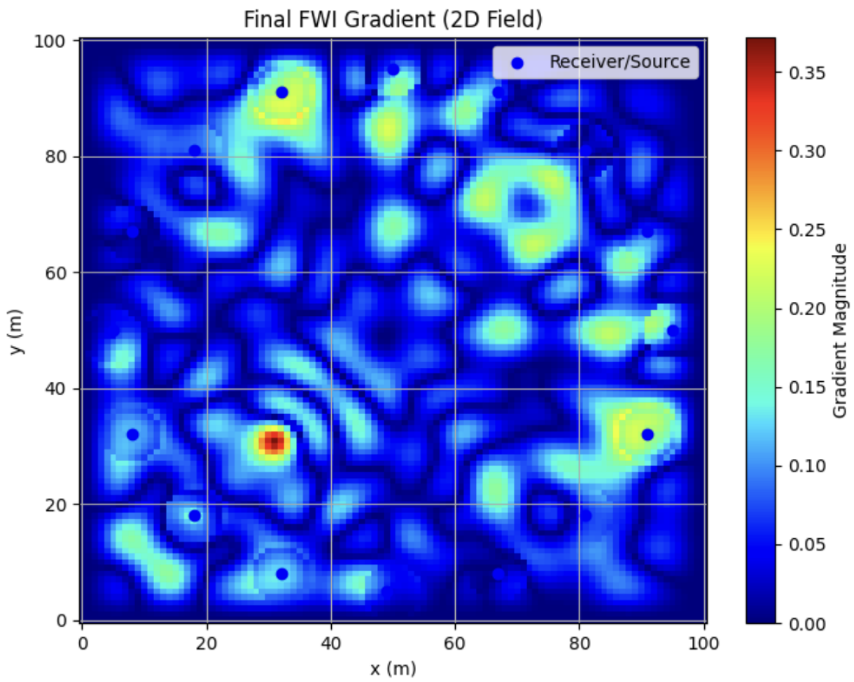

Figure 3.4: Final gradient magnitude distribution computed from the adjoint-state method. Bright areas indicate higher sensitivity of the misfit to local model changes. Blue dots denote source and receiver positions.

This part explains the steps involved in implementing gradient computation, including defining residuals on the grid, solving the adjoint equation with the Hermitian transpose of the system matrix, and calculating the gradient by correlating the adjoint wavefield with the Laplacian of the forward wavefield. The gradient obtained is input into the BFGS optimization routine for model updating within the inversion loop.

3.4 Inversion Strategy

This part details the inversion framework within the context of FWI, which incorporates forward modeling, adjoint-based gradient calculation, iterative model updating with the BFGS algorithm, and model optimization. In the implementation version, two design decisions are made which include the frequency component treatment and source position arrangement. As for frequency scheduling, two variants are considered. The first variant employs sequential multi-frequency band assimilation where lower frequency ranges are introduced first followed by higher frequencies which aids in coarse-to-fine model reconstruction [5], [8]. The second variant applies all frequencies in a single iteration to every inversion during each iteration to speed up convergence. As for source configuration, the framework analyzes single-source configurations versus multi-source scenarios where each simulation step uses one or multiple sources [9].

These strategies aim to reconstruct high-resolution velocity models by minimizing the difference between observed and simulated wavefield data. The consecutive sub-chapters explain each of the strategies' reasoning and describe their execution on a more detailed level.

3.4.1 Multi-Frequency Approach

Here we explain the strategy of guiding the inversion from coarse to fine resolution using multi-frequency strategy. This sequential technique adds frequencies stepwise, beginning with low-frequency components, and later integrating higher-frequency bands. Low-frequency components enable large-scale structural recovery first during the inversion process, while higher-frequency components enable more detailed recovery afterwards [5], [8]. The frequency bands assigned for sequential inversion are 25 Hz, 40 Hz, 60 Hz, 80 Hz, 100 Hz, 120 Hz, and 150 Hz.

During each frequency stage, inversion is done following a standard loop, where the Helmholtz equation is solved via finite-difference discretization, synthetic wavefields are created and matched with observed data, and the misfit is computed. Using residuals, the gradient is calculated based on the adjoint-state method [2], [6]. To prevent instabilities around the source, a masking technique is used in the gradient field. BFGS Optimization algorithm is employed for updating the model and iterates until a convergence threshold or a predefined maximum iteration count is reached (for instance, 100 iterations).

In contrast, the simultaneous method adds all frequency components in one inversion step. A global loss function combines the misfit from all frequencies and calculates a composite gradient. This approach optimizes convergence without cycling through iterative frequency steps. On the downside, it demands higher computational effort because the solution for multiple frequencies is computed at the same time [9].

3.4.2 Handling Multi-Source Data

The enhancement of the multi-source approach lies in its use of data at multiple source positions which improves the robustness of the inversion. In the configuration used, it is formed by a circular array of 16 source-receiver pairs that are placed at a radius of 45 m from the center of the computational domain (see Figure 3.2). All receivers are recording synthetic wavefield data from each source which uses a high amplitude delta function (1.0×10^{10}). Each source's misfit function includes the contributions of all source-receiver pairs and all frequencies as outlined in Equation (2.9).

To balance computational efficiency with other factors such as accuracy, the system matrix $A(m, \omega)$ is LU-factorized once per frequency, and the factorization is reused for all source terms. This reduces the overall cost of solving the linear system (Equation (2.3)) multiple times as the matrix decomposition does not need to be recomputed for each individual source [9], [10].

3.4.3 Optimization Loop

As highlighted in Section 2.3.1, the optimization loop combine forward modeling, computation of an adjoint-state gradient, and model updates via the BFGS algorithm. This optimization is performed through iterative steps with the intent of minimizing the difference between observed and synthetic wavefield data.

The previously described inversion loop is implemented here using a homogeneous initial velocity model and iteratively refined using the BFGS optimizer.

Then, in accordance with Equations (2.11) and (2.12), the gradient of the misfit function is calculated using the adjoint-state method. This gradient is useful because it provides directional sensitivity information about each component of the velocity model. Afterward, using BFGS, a model to optimize the problem and provide the solution is obtained, which considers information of the previous iteration of the model along other auxiliary gradient and curvature data in order to speed up the convergence [2], [7].

During the inversion cycle, intermediate results like updated velocity models, computed wavefields, and gradient fields are saved for subsequent analysis and visualization. This step-by-step monitoring aids in evaluating the convergence behavior and helps in determining the effectiveness of inversion processes, as shown in Figures 3.2, 3.3, 3.4, and 2.1.

Depending on the setup, the loop can operate in a sequential mode, processing one frequency at a time, or simultaneously, where all frequencies and sources are integrated into a single set. In the latter approach, computational performance is usually improved through LU factorization for systems with repetitive solutions [9].

3.4.4 Overview of Strategy Differences

This research examines four distinct inversion strategies: single-source with sequential multi-frequency, single-source with simultaneous multi-frequency, multi-source with sequential multi-frequency, and multi-source with simultaneous multi-frequency.

Chapter 4 contains a detailed analysis of the strategies based on their performance, convergence trends, and reconstruction quality.

Chapter 4

Results and Discussion

4.1 Overview of the Approaches under Study

The previous chapter presented four inversion strategies with each having a given source configuration and frequency scheduling arrangement. Each of these methods combines a single- or multi-source setup with either sequential or simultaneous multi-frequency inversion. Each configuration offers different advantages and disadvantages with respect to computational cost, reconstruction quality, and convergence behavior [5], [6], [8].

This chapter analyzes the performance of these strategies based on results from inversion experiments. The analysis is arranged into two primary parts to provide a clear focus and meaningful comparison. First, the analysis focuses on assessing single-source strategies to evaluate how the order of frequency application impacts inversion results, independent of source coverage. This isolatory approach aids in assessing the influence of frequency scheduling [9].

The second part of the study concentrates on assessing multi-source configurations. These arrangements, which integrate several sources at once, generally provide more extensive spatial coverage of the area in question.

Each inversion experiment is performed on a (2D) spatial domain of $100\,\mathrm{m} \times 100\,\mathrm{m}$. This area contains a uniform Cartesian grid consisting of 101×101 points, which can be discretized as $\Delta x = \Delta y = 1\,\mathrm{m}$. The initial wave speed model is set to be homogeneous at $3000\,\mathrm{m/s}$, except for a localized anomaly around the center of the domain, where wave speed is reduced to $2500\,\mathrm{m/s}$ within a 5×5 sub-region. The model is further subjected to constant attenuation of 50 applied throughout.

The source in the model emits wavefields at the following seven frequencies: $25, 40, 60, 80,$ $100, 120,$ and $150\,\mathrm{Hz}$. 16 receivers are arranged circularly with a radius of $45\,\mathrm{m}$, positioned at the center of the domain.

Figure 4.1 shows the final velocity models obtained using two different inversion strategies. On the left is the result from the sequential single-frequency approach, and on the right is the result from the simultaneous multi-frequency method. Both methods are able to reconstruct the velocity anomaly well, capturing its shape and strength. In both cases, the velocity inside the anomaly reaches about 2800 m/s.

Looking a bit closer, there are some subtle differences. The sequential method tends to produce a slightly sharper outline of the anomaly, while the simultaneous approach shows smoother transitions around the edges. Still, both results are quite similar in terms of locating the anomaly and recovering the correct contrast.

For a fair comparison, it is important to consider not only the quality of the inversion results, but also the computational cost. The sequential method processes one frequency at a time, which means running separate inversions for each. In contrast, the simultaneous approach includes all frequencies in every iteration. This leads to a major difference in runtime: the sequential method took around 2 hours and 36 minutes, while the simultaneous strategy completed in just 22 minutes.

This difference is particularly significant for single-source strategies. Since the sequential method moves through the frequency spectrum step by step, it naturally requires more iterations. The simultaneous strategy avoids that by using all available frequency information in each update, reusing wavefields efficiently. Although this can slightly smooth out the anomaly boundary, the time savings are substantial and the accuracy of the final result remains high.

To conclude, while both strategies deliver similar results in terms of reconstruction quality, the simultaneous method clearly stands out when it comes to computational efficiency. This makes it a highly practical choice, especially when working with limited time or resources.

4.2 Convergence Behavior

The sequential strategy uses a coarse-to-fine paradigm, where the inversion begins with low-frequency components first [5]. This strategy captures the large-velocity features of the model and then shifts focus to smaller details. Usually, Such a strategy ensures that the updates are smoother while avoiding cycle skipping, particularly when the starting point is a homogeneous initial model. The inversion process itself tends to be more stable, as the low-frequency components that are employed earlier in the sequence correspond well with the major, uniform characteristics of the initial model [6]. In general, however, it tends to have a higher number of total iterations because every frequency tier requires a separate optimization loop.

In contrast to the sequential strategy, the simultaneous approach combines all frequency elements within a single iteration loop. This fundamentally alters the gradient as it uses the entire spectral content at once. The increase in information contained in this gradient allows the misfit function to be reduced more quickly and stably than with update steps using incremental frequency updates. [9]. As shown in Figure 4.1, tthe simultaneous approach achieves a normalized misfit of 0.3 by iteration 40 and 0.18 by iteration 60. In contrast, the sequential method reaches misfit values of only 0.65 and 0.45 at the same iterations. respectively, After 100 iterations, the final misfit for the simultaneous strategy is 0.07 while the sequential method lags at 0.25. The simultaneous approach converges faster within 100 iterations.

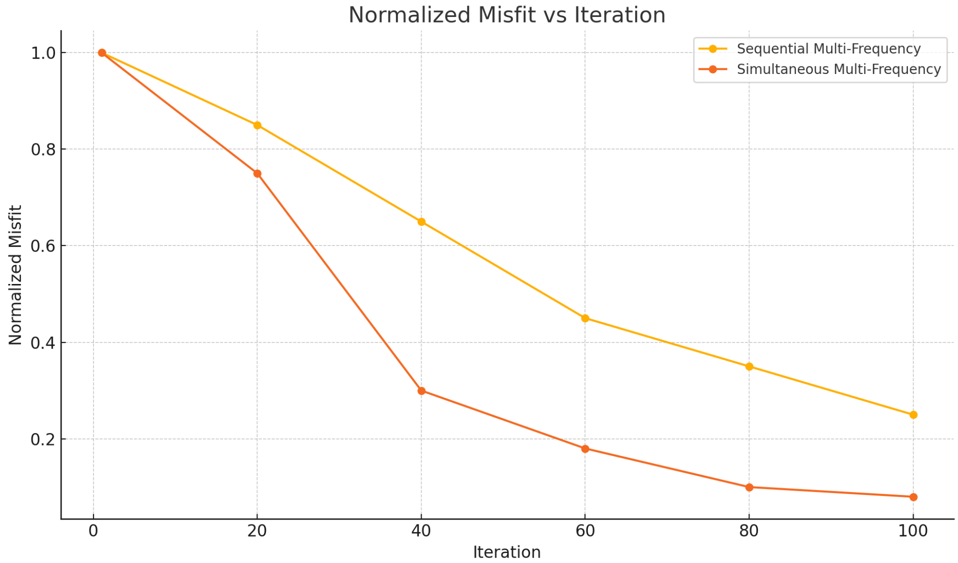

Figure 4.1: Normalized misfit vs. iteration for single-source strategies. The simultaneous approach achieves faster convergence, reaching a misfit of 0.07 by iteration 100, while the sequential strategy reaches 0.25.

4.3 Single-Source Strategies: Sequential vs. Simultaneous

Following the implementation detailed in Section 3.4.1, this section evaluates the effectiveness of single-source approaches relative to multi-frequency simultaneous and sequential scheduling. In the simultaneous approach, all frequency ranges are utilized at every step, whereas the sequential approach adds frequencies from low to high [5], [8].

Both methods aim to reduce the gap between the real and simulated wavefields; however, they differ regarding the techniques employed at the replacement phase during inversion rendering. Their performance is evaluated in this case based on convergence, reconstruction quality, and the stability of numerical calculations [2].

Figure 4.2 shows the final velocity models obtained using two different inversion strategies. On the left is the result from the sequential single-frequency approach, and on the right is the result from the simultaneous multi-frequency method. Both methods are able to reconstruct the velocity anomaly well, capturing its shape and strength. In both cases, the velocity inside the anomaly reaches about 2800 m/s.

Looking a bit closer, there are some subtle differences. The sequential method tends to produce a slightly sharper outline of the anomaly, while the simultaneous approach shows smoother transitions around the edges. Still, both results are quite similar in terms of locating the anomaly and recovering the correct contrast.

For a fair comparison, it is important to consider not only the quality of the inversion results, but also the computational cost. The sequential method processes one frequency at a time, which means running separate inversions for each. In contrast, the simultaneous approach includes all frequencies in every iteration. This leads to a major difference in runtime: the sequential method took around 2 hours and 36 minutes, while the simultaneous strategy completed in just 22 minutes.

This runtime difference is particularly significant for single-source strategies. Since the sequential method moves through the frequency spectrum step by step, it naturally requires more iterations. The simultaneous strategy avoids that by using all available frequency information in each update, reusing wavefields efficiently. Although this can slightly smooth out the anomaly boundary, the time savings are substantial and the accuracy of the final result remains high.

To conclude, while both strategies deliver similar results in terms of reconstruction quality, the simultaneous method clearly stands out when it comes to computational efficiency. This makes it a highly practical choice, especially when working with limited time or resources.

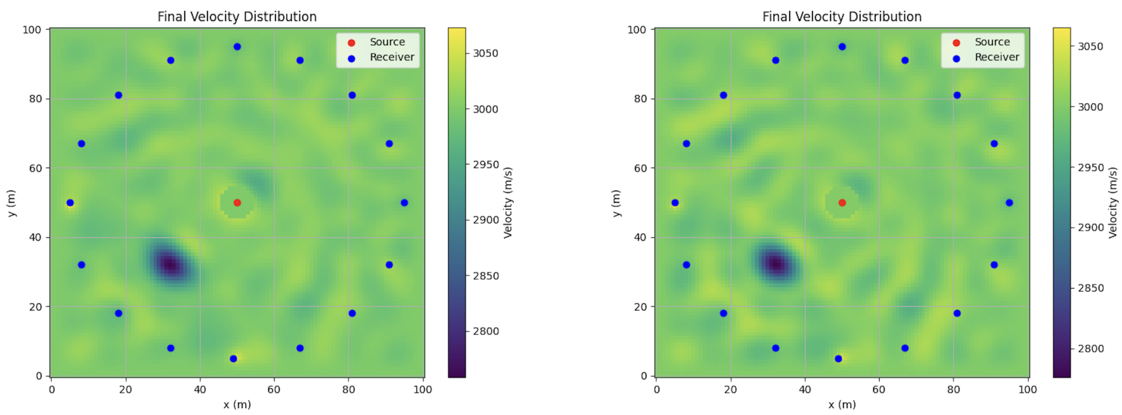

Figure 4.2: Final velocity models obtained using two different inversion strategies. The left image shows the result from the sequential single-frequency approach, while the right image shows the result from the simultaneous multi-frequency method.

4.4 Multi-Source Strategies: Sequential vs. Simultaneous

Now we will consider the multi-source sequential and simultaneous multi-frequency strategies with respect to reconstruction quality, convergence stability, and numerical artifacts. Both strategies employ a circular array of sensors and emitters which improves the spatial coverage of the wavefield compared to single source techniques [2], [5]. This illumination improves the results obtained from inversion, especially in the sharpness and fullness of the reconstructed anomaly [3], [9].

Figures 4.3 and 4.4 provide a more detailed view of the final recovered velocity profiles for the sequential and simultaneous multi-frequency strategies, respectively, in the multi-source setting. These plots show the real part of the velocity, averaged laterally across the anomaly region.

In Figure 4.3, corresponding to the sequential strategy, the recovered real velocity profile shows a sharp and relatively deep anomaly, closely matching the shape of the true model.

Figure 4.4 shows the same profile for the simultaneous strategy. The recovered velocity displays a broader anomaly with smoother transitions, as expected from the collective frequency treatment.

Building on these results, Figure 4.5 presents the full 2D velocity distributions for both strategies. The left panel shows the result from the sequential multi-frequency approach, while the right panel shows the simultaneous counterpart. Though the convergence behavior is generally similar to what was observed in the single-source case, some differences are noticeable.

Both strategies are successful in recovering the main features of the anomaly. However, there is a slight difference in the final recovered velocities. The sequential method reconstructs the anomaly with a minimum velocity of about 2720 m/s, whereas the simultaneous approach achieves a slightly lower value of approximately 2680 m/s. This suggests that the simultaneous strategy is marginally more accurate in terms of velocity recovery.

In terms of shape and resolution, the sequential result appears sharper and more focused, while the simultaneous model shows a smoother transition around the anomaly edges. This behavior aligns with the inherent nature of each method: sequential inversion builds frequency content progressively, while the simultaneous method handles all frequencies at once, which can lead to some smoothing.

On the computational side, both methods benefit from the factorization of the Helmholtz operator $\mathbf{L}$, which allows for wavefield reuse across multiple sources. This optimization significantly reduces the overall runtime in the multi-source setup. The sequential approach completed in approximately 34 minutes, while the simultaneous strategy finished in just 26 minutes. This faster runtime, combined with slightly better velocity recovery, gives the simultaneous method a practical edge.

Ideally, both methods should ideally be run for the same amount of time. It is quite possible that the simultaneous method would deliver even better results if it were allowed to run as long as the sequential one. Because of this, the current difference in performance should be viewed with some caution.

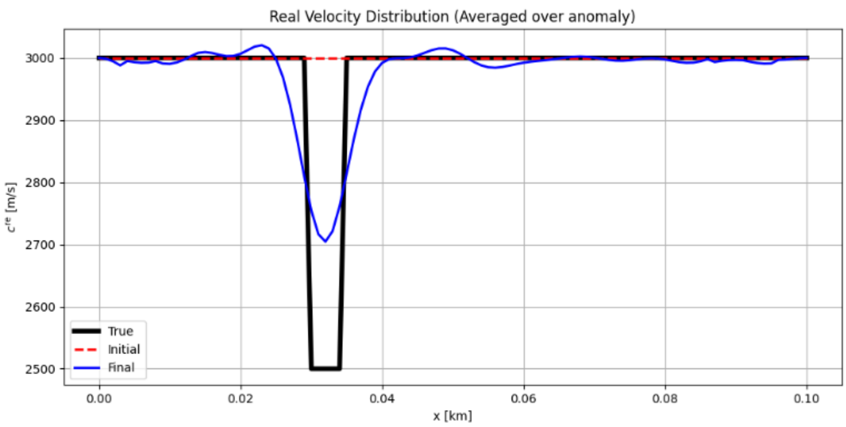

Figure 4.3: Final recovered real velocity profile using the sequential multi-frequency strategy.

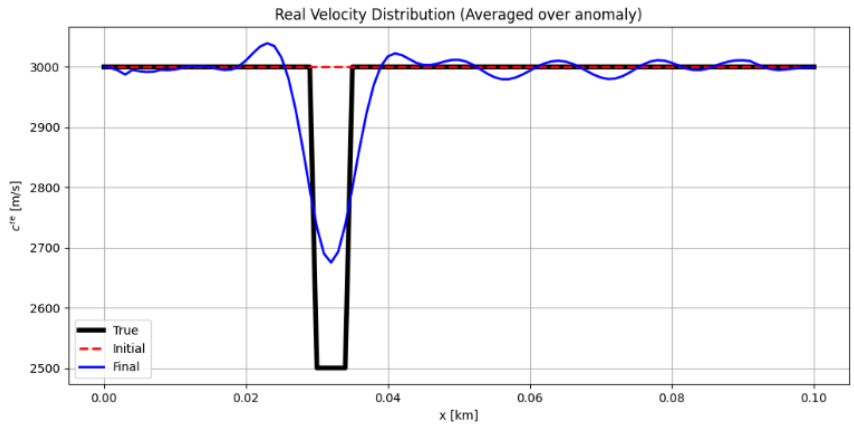

Figure 4.4: Final recovered real velocity profile using the simultaneous multi-frequency strategy.

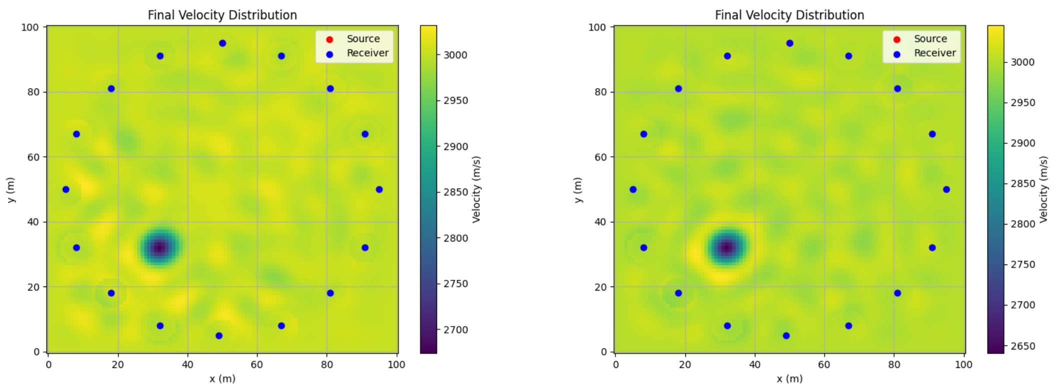

Figure 4.5: Final velocity distributions for multi-source strategies. Left: Sequential multi-frequency. Right: Simultaneous multi-frequency.

Chapter 5

Conclusion and Outlook

5.1 Conclusion

In this research, the investigation is focused on the application of frequency-domain FWI in ultrasonic NDT. Starting with the modeling of ultrasonic wave propagation, the Helmholtz equation, and the FWI optimizations, this work reviews the two central strategies for frequency injection, which are sequential and simultaneous multi-frequency inversion. Their implementations are studied in both single-source and multi-source configurations to evaluate their potential in recovering low-velocity velocity anomalies in a synthetic model.

Within the constructed framework, a Helmholtz problem solver, an adjoint-state gradient calculation implementer, and a Quasi-Newton model updater were developed in Python. This system interfaces with sequential and simultaneous multi-frequency techniques which increases the versatility for testing inversion configurations.

The sequential strategy seemed to produce more defined sharper anomaly boundaries, probably due to the construction of frequency information accumulation processes. This effort, however, notably increased computation duration, particularly in the single-source scenario. Even though the final images were climatically satisfying and well-defined, the achieved velocity values were slightly lower in accuracy than the simultaneous approach.

However, the simultaneous approach converged much faster, especially when multiple sources were employed since reusing the LU factorization provided great accelerative value. Although tere was a drop in accuracy capturing edge boundaries, the models were closer to the true value and captured them in significantly less time.

For a fair comparison, both strategies should really be tested under the same time limits. If given the same runtime as the sequential method, the simultaneous approach would likely deliver even better results. That makes it the more efficient and practical option in many full-waveform inversion scenarios.

The thesis shows how selecting a combination of frequency scheduling and source configuration can improve FWI performance dramatically. The multi-source setup has better spatial coverage as well as stronger gradients. The choice of sequential versus simultaneous strategies depends on speed or accuracy objectives for the inversion.

5.2 Outlook

This thesis' framework and findings can be extended in several ways. One promising avenue is to replace the current acoustic model with an elastic one to better capture wave physics relevant to NDT in industrial settings. Also, implementing adaptive frequency selection could improve convergence speed and resolution, compared to fixed schemes.

Future work could incorporate the inversion framework along with ultrasonic experimental data to enable validation within practical scenarios. These efforts would need to solve problems such as sufficient noise, source wavelet estimation, and calibration of boundary conditions. Adding real-time capabilities for simultaneous approaches, particularly those using LU optimization, could enable industrial applicability.

Conclusively, hybrid inversion approaches that integrate the resilience of sequential frequency scheduling with the rapid execution of simultaneous updates may provide an optimal solution. These could consist of low-frequency sequential early updates followed by rapid refinements in later stages. This hybrid approach exemplifies advancement towards improving the applicability of frequency domain FWI in NDT.

References

[1] A. Tarantola and B. Valette, "Inverse problems = quest for information," *Journal of Geophysics*, vol. 50, pp. 159–170, 1982. [Online]. Available: https://journal.geophysicsjournal.com/JofG/article/view/28/16#nav-menu.

[2] A. Fichtner, *Full Seismic Waveform Modelling and Inversion* (Advances in Geophysical and Environmental Mechanics and Mathematics). Berlin Heidelberg: Springer, 2011, ISBN: 978-3-642-15806-3. DOI: 10.1007/978-3-642-15807-0.

[3] D. Köhn, T. Meier, M. Fehr, D. D. Nil, and M. Auras, "Application of 2d elastic rayleigh waveform inversion to ultrasonic laboratory and field data," *Near Surface Geophysics*, vol. 14, no. 5, pp. 461–476, 2016. DOI: 10.3997/1873-0604.2016027. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.3997/1873-0604.2016027.

[4] J. Rao, "Quantitative imaging using ultrasonic tomography based on full waveform inversion," Ph.D. dissertation, Nanyang Technological University, 2018. DOI: 10.32657/10356/75928. [Online]. Available: http://hdl.handle.net/10356/75928.

[5] J. Virieux and S. Operto, "An overview of full waveform inversion in exploration geophysics," *Geophysics*, vol. 74, no. 6, WCC127–WCC152, 2009. DOI: 10.1190/1.3238367. [Online]. Available: https://doi.org/10.1190/1.3238367.

[6] A. Fichtner, "Lecture notes on inverse theory," *Cambridge Open Engage*, 2021. DOI: 10.33774/coe-2021-qpq2j.

[7] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd. Springer, 2006. DOI: 10.1007/978-0-387-40065-5.

[8] R. G. Pratt, "Seismic waveform inversion in the frequency domain, part 1: Theory and verification in a physical scale model," *Geophysics*, vol. 64, no. 3, pp. 888–901, 1999. DOI: 10.1190/1.1444597. [Online]. Available: https://library.seg.org/doi/abs/10.1190/1.1444597.

[9] Y. Tao and M. K. Sen, "Frequency-domain full waveform inversion with plane-wave data," *Geophysical Journal International*, vol. 195, no. 1, pp. 642–658, 2013. [Online]. Available: https://doi.org/10.1190/geo2012-0267.1.

[10] T. Bürchner, P. Kopp, S. Kollmannsberger, and E. Rank, "Immersed boundary parametrizations for full waveform inversion," *Computer Methods in Applied Mechanics and Engineering*, vol. 406, p. 115 893, 2023. DOI: 10.1016/j.cma.2023.115893. [Online]. Available: https://doi.org/10.1016/j.cma.2023.115893.

[11] H. Igel, *Computational Seismology: A Practical Introduction*. Oxford University Press, 2017, ISBN: 978-0-19-871740-9.

Declaration

I hereby affirm that I have independently written the thesis submitted by me and have not used any sources or aids other than those indicated.

Location, Date, Signature