

Quantitative Transmission Tomography for Non-Destructive Imaging of Stored Grain and Biological Tissue

Max Hughson

A Thesis submitted to the Faculty of Graduate Studies of
The University of Manitoba
in partial fulfillment of the requirements of the degree of

MASTER OF SCIENCE

Department of Electrical and Computer Engineering
University of Manitoba
Winnipeg

Copyright © 2021

Abstract

This thesis presents the development, testing, and verification of transmission tomography software for the applications of small-scale biological imaging and large-scale stored-grain imaging. Transmission tomography is a non-invasive technique for producing quantitative images of an object's physical properties. This is done by interrogating the object with waves and measuring the resulting response at a set of measurement positions. The properties of the received waves are analyzed, and used to calculate the properties of the object which is interrogated.

First, a mathematical formulation of transmission tomography is presented and explained. The formulation is then used to build a numerical model of the physical properties which dictate wave transmission in two systems of linear algebraic equations. The algorithms that are required for building the numerical model are then explained.

The rest of this thesis is devoted to a series of experiments. These experiments show the usefulness of transmission tomography in some particular applications. They also motivate the incremental development of features of the transmission tomography algorithm that was developed for this thesis.

The first experiment uses acoustic data to perform two-dimensional transmission tomography. The first experiment shows that two-dimensional transmission tomography of acoustic data produces useful images of an object's wave speed. The successful two-dimensional experiments motivate the two following experiments, which use electromagnetic data to perform three-dimensional transmission tomography. In order to perform three-dimensional tomography of electromagnetic data, a modification is made to the numerical model. The successful three-dimensional transmission tomography algorithm is then applied to track the volume of grain stored in a bin. The final experiment uses acoustic data and the modified numerical model to perform three-dimensional transmission tomography to identify two distinct spherical objects which are submerged in a liquid medium.

Contributions

The primary contribution of this thesis is an algorithm which uses time-domain scan data to perform non-destructive, quantitative imaging.

Contributions to State of Knowledge

The key contributions made to the general state of knowledge are summarized here:

1. The development and implementation of a ray-based tomography algorithm which operates on either acoustic or electromagnetic data to quantitatively image an object's properties, which are represented with a finite basis
2. The optimization of algorithms which process time-domain signals to extract features which are useful to the tomography algorithm
3. The incorporation of both a pulse basis and a polynomial basis in the tomography algorithm, and analysis of their impact on image quality
4. The verification of the utility of the tomography algorithm, including:
 - Monitoring the volume of grain stored in a metallic silo, using electromagnetic measurements
 - Identifying regions of fibroglandular tissue embedded in a region of fat, in a tissue-mimicking acoustic phantom

Computer Program

The algorithm described in this work was implemented as a Matlab program. The program consists of computer code to carry out the following tasks:

- Reading user-specified configuration information from text files
- Reading scan data and extracting features from scans, according to the configuration

- Calculating line-integrals of basis functions along straight paths
- Solving matrix equations, according to the configuration
- Writing visualization files to the file system
- Packaging imaging results in a convenient data structure

Various scripts were written in order to automate the imaging workflows which produce the results presented in this thesis. These scripts are accessible on a code repository which is hosted at the University of Manitoba. These scripts perform the following tasks:

- Translating path data and scan data into a format that is readable by the code described above
- Processing the scan data to aid the feature extraction described above
- Calling the code described above to generate imaging results
- Comparing the imaging results to expected results, via quantitative error analysis
- Writing the results of the error analysis to files which can be included in this thesis

The computer code for the program is the author's own work, except for the following components:

`partitions` This function [1] calculates and enumerates the unique ways to partition an integer into a sum containing a given number of terms.

`writeVTK` This function [2] creates `vtk`-formatted files which can be viewed in ParaView [3].

The electromagnetic simulations which are used in Chapter 4, Chapter 5, and Chapter 6 make use of some computer code written by Hannah Fogel. Hannah is also responsible for maintaining and operating the hardware which is described in Chapter 3 and Chapter 7.

Acknowledgements

I thank my advisors, Dr. Joe LoVetri and Dr. Ian Jeffrey, for their support and guidance. You have helped me grow professionally and personally.

I thank my examining committee for spending some of your valuable time to review and evaluate my thesis.

I also thank all the members of the Electromagnetic Imaging Laboratory. Our group is unique, and I am lucky to have had the opportunity to work with all of you. I am particularly grateful to my friend and colleague, Hannah Fogel. Your work on modifying and operating the ultrasound imaging chambers provided the inspiration to begin this research, and my research would not have been possible without your technical support. Your time-domain electromagnetic simulation scripts were also indispensable for my research.

Finally, I thank my partner, Holly Macduff for your love, patience and support throughout all of my studies.

This work was enabled by financial support from the Natural Sciences and Engineering Research Council of Canada, the University of Manitoba, and AGCO Winnipeg.

Contents

List of Tables	xi
List of Figures	xiii
1 Introduction	1
1.1 Relevant History	2
1.2 Grain-Bin Imaging	4
1.3 Ultrasound Tomography	5
1.4 Transmission Tomography	6
2 Transmission Tomography	7
2.1 General Formulation	7
2.2 Wave Propagation Model	9
2.2.1 Total-Field and Scattered-Field Imaging	11
2.2.2 Discretizing the Model	15
2.3 Feature Extraction Details	22
2.3.1 Time-Of-Arrival Determination	22
2.3.2 Pulse Power Determination	29
2.3.3 Ignoring Data	31
2.4 Extracting Auxiliary Material Properties	32
2.4.1 Acoustic Properties	33
2.4.2 Electromagnetic Properties	33
3 Two-Dimensional Ultrasound Imaging on Unstructured Meshes	35
3.1 Introduction	35

3.2	Procedure	36
3.2.1	Data Processing Details	38
3.2.2	Development	39
3.2.3	Testing	40
3.3	Results	41
3.4	Discussion	42
4	Three-Dimensional Pulse-Basis Transmission Tomography in Grain Bins	45
4.1	Introduction	45
4.2	Procedure	47
4.2.1	Data Processing Details	51
4.3	Results	51
4.4	Discussion	54
5	Three-Dimensional Polynomial-Basis Transmission Tomography in Grain Bins	57
5.1	Introduction	57
5.2	Procedure	58
5.2.1	Data Processing Details	61
5.3	Results	61
5.4	Discussion	63
6	Monitoring Grain-Bin Fill Volume	65
6.1	Introduction	65
6.2	Procedure	66
6.2.1	Data Generation	66
6.2.2	Volume Extraction	67
6.2.3	Data Processing Details	67
6.3	Results	67
6.4	Discussion	68
7	Biological Imaging in Three Dimensions	71
7.1	Introduction	71

7.2	Procedure	72
7.2.1	Data Processing Details	75
7.3	Results	76
7.4	Discussion	76
8	Discussion and Conclusions	81
8.1	Future Work	84
8.1.1	Basis Degree Selection	84
8.1.2	Diffraction-Aware Solver	85
8.1.3	Iterative Polynomial Basis Refinement	85
8.1.4	On-Line Imaging	86
8.1.5	Regularization Study of Pulse-Basis Imaging	86
9	Bibliography	89
A	Line-Integrals Of Pulse Basis Functions	95
B	Line-Integrals Of Polynomial Basis Functions	99
B.1	Symbolic Representation	99
B.2	Computation	101
C	Details for Extracting Auxiliary Material Properties	103
C.1	Calculating the Complex Wave Number	104
C.2	Acoustic Properties	105
C.3	Electromagnetic Properties	106

List of Tables

2.1	Time-Of-Arrival of Sample Scan	25
3.1	Imaging metrics for two-dimensional ultrasound experiment . .	42
4.1	Expected material properties of grain reconstruction	50
4.2	Accuracy of reconstruction on meshes of varying density . . .	54
5.1	Accuracy of reconstruction on bases of varying degree	63
7.1	Expected material properties for 3D ultrasound experiment . .	74
7.2	Accuracy of ultrasound reconstruction on bases of varying degree	79

List of Figures

2.1	Physical setup of a general transmission tomography system, showing a single transmitter-receiver path	8
2.2	Example of a transmission tomography scan, as measured by receiver R_i	10
2.3	Flow of information in scattered-field and total-field formulations.	16
2.4	Over-determined and under-determined matrix systems. The shape of \mathbf{L} depends on the number of paths and the number of basis functions. \underline{p} is a vector of properties, and \underline{f} is a vector of features.	19
2.5	Physical causes of observed phenomena in time-domain scans .	24
2.6	Sample scan for illustrating algorithms for determining time or arrival	25
2.7	Time-of-arrival determination via peak detection	26
2.8	Time-of-arrival determination via modified energy ratio	27
2.9	Time-of-arrival determination via variance score	30
3.1	System schematic. Red lines indicate digital control connections, with the arrow head pointing to the controlled device. Green lines indicate digital signal connections and blue lines indicate analog signal connections, with arrows pointing in the direction of signal flow. Only 3 of the 64 transducer supports are shown, for visual simplicity.	36
3.2	Ring of ultrasound transducers, submerged in water	37
3.3	Horizontal and vertical scale of homogeneous gelatin phantom	38

3.4	Pulse-basis reconstructions of gelatin phantoms	43
4.1	Synthetic grain distribution for grain-bin imaging experiment .	49
4.2	Error and interrogation ratio with respect to mesh character- istic length	52
4.3	Pulse-basis reconstructions of wave speed from meep-generated electromagnetic data	53
5.1	Evaluation domain for the error metric. The region shown here represents the cylindrical portion of the grain bin. The upper and lower cones are omitted for visual clarity. The set of paths fails to interrogate the outer regions of the cylinder. .	59
5.2	Synthetic grain distribution for grain-bin imaging experiment .	61
5.3	Polynomial basis reconstructions. Colour indicates wave-speed, from 1×10^8 to $3 \times 10^8 \text{ m s}^{-1}$	62
5.4	Isosurfaces of polynomial basis reconstructions. Colour indi- cates wave-speed, from 1×10^8 to $3 \times 10^8 \text{ m s}^{-1}$	62
5.5	Error metrics with respect to basis degree	63
6.1	Schematic of grain distribution for volume-tracking experiment	66
6.2	Extracted volume plotted against grain height	68
6.3	Grain volume calculation, shown for three different fill levels .	69
7.1	Images of the three-dimensional ultrasound imaging chamber .	72
7.2	Known speed profile with transducer locations. Colour indi- cates wave-speed, ranging from 1400 m s^{-1} to 1600 m s^{-1}	73
7.3	Ultrasound signal pre-processing. The transmitter-receiver path is shown on the left, and the corresponding truncated signal is shown on the right	74
7.4	Face-view of sound-speed reconstructions. Colour indicates wave-speed, ranging from 1300 m s^{-1} to 1800 m s^{-1}	77
7.5	Contoured view of sound-speed reconstructions. Colour indi- cates wave-speed, ranging from 1300 m s^{-1} to 1800 m s^{-1}	78
7.6	Error metrics with respect to basis degree in three-dimensional ultrasound experiment	79

A.1 Calculating a surface intersection point on one face of a tetrahedral element 97

Glossary

attenuation Physical property which determines power lost by a wave during propagation, expressed in units of $\text{dB cm}^{-1} \text{MHz}^{-1}$. 9

decay Scaling of a pulse's amplitude due to attenuation of the transmission medium. 9, 20, 45

delay Deviation of a pulse's measured arrival time from its expected arrival time due to the slowness of the transmission medium. 9, 20, 45

feature extraction The process of determining the delay and decay of a particular scan. 20

inverse speed Inverse of wave speed. *see* slowness

object of interest Object embedded in the transmission medium, whose physical properties may differ from the transmission medium. 7

pulse Causal wave produced by a transmitting element. 7

receiving element Device which measures waves. 7

slowness Physical property which determines wave propagation speed in a medium, expressed in units of s m^{-1} . 9

transmission medium Region of space defined by spatially-varying physical properties which influence wave propagation. xiii, 7

transmitting element Device which produces a wave in a transmission medium. 7

Acronyms

CGLS conjugate-gradient least-squares. 19, 20, 36, 38, 86

CSI contrast-source inversion. 4, 34, 44

CURE computed ultrasound risk evaluation. 26

EIL Electromagnetic Imaging Laboratory. 1–4, 30, 33, 44, 72, 82

FDTD finite-difference time-domain. 45

MER modified energy ratio. 21, 24, 27, 75

MUBI multimodal ultrasound breast imaging. 23, 24

MWI microwave imaging. 2, 4, 5, 44, 45

RF radio-frequency. 2, 4, 44

TOA time of arrival. 11, 20–24, 26, 27, 29, 49, 74, 75

VNA vector network analyzer. 3

WFI wave-field imaging. 4, 5, 30

Chapter 1

Introduction

Humans are curious creatures, and engineers are especially curious creatures. Many engineers have felt the inclination to take something apart and put it back together again, just to see what's inside.

That is a fine technique for uncovering the secrets of a toaster or figuring out how the wipers go back and forth on your car, but 'taking it apart' is often not safe, and it may even be impossible! For example, in medicine, tissue biopsy is generally performed *after* attempts have been made at tissue imaging, via some modality such as X-ray, not the other way around. In other words, we often prefer non-destructive imaging techniques over other invasive techniques.

This thesis presents the development and testing of a non-invasive imaging algorithm which makes use of transmission tomography. The research undertaken in this thesis was motivated by the author's experience with imaging algorithms that are in use at the Electromagnetic Imaging Laboratory (EIL). The transmission tomography algorithm that is presented here complements the algorithms currently employed in the EIL. Specifically, the goal of the work described in this thesis is to create an imaging tool with the following features:

- Minimal reliance on prior information
- Robustness to modelling error

- Fast execution time

The importance and relevance of these features are explained in the following sections.

The primary use of the algorithm presented here is for the monitoring of stored agricultural commodities, such as corn or wheat in a metal silo, via radio-frequency (RF) electrical measurements [4]. Exploratory experiments are also documented which use transmission tomography in a medical ultrasonic imaging context. The experiments presented in this thesis all follow a similar workflow:

1. A set of data is generated by interrogating an object with pulsed waves and measuring the resulting wave patterns at a set of points. This step is performed synthetically in Chapter 4, Chapter 5, and Chapter 6. This step is performed using real hardware in Chapter 3 and Chapter 7.
2. Transmission tomography is used to generate maps of the material properties of the object from the set of data.
3. The maps of material properties are compared to the known material properties in order to evaluate the accuracy of the transmission tomography step.

The experiments show that meets the criteria from above, when used to image either electromagnetic data from a grain bin or acoustic data from an ultrasound imaging system.

1.1 Relevant History of Algorithm Development in the Electromagnetic Imaging Laboratory

The EIL is a research laboratory located in the Department of Electrical and Computer Engineering at the University of Manitoba. The EIL began

as a research lab focused on ground-penetrating radar, and has since grown to study several techniques to “image objects obscured from normal perception” [5]. In recent years, researchers at the EIL and elsewhere, have spent significant energy contributing to the field of microwave imaging (MWI) [6]. In particular, researchers at the EIL have developed and studied full-wave, non-linear imaging techniques [7].

This research into full-wave, non-linear imaging has produced impressive algorithms and systems, however, these algorithms have particular idiosyncratic drawbacks which have stood in the way of full system implementations.

For example, full-wave imaging studies often rely on *prior information* [8]. Imaging studies often need accurate information about the shape and bulk physical properties of the imaging target. Successful imaging of breast tissue via MWI relies on an accurate definition of the boundary of the breast, as well as an accurate model of the sub-cutaneous fat layer. Such prior information is frequently unavailable.

Full-wave, non-linear MWI is also highly sensitive to *modelling error*. MWI works by computationally modelling the wave phenomena in some imaging domain, and tweaking the physical properties in the imaging domain until the modelled electromagnetic fields match a set of real-world electromagnetic measurements. A modelling error is any difference between the computational model of the wave phenomena and the actual wave phenomena.

An example of modelling error is ignoring mutual antenna coupling in imaging studies. It is computationally expensive to accurately model antennas. Antennas are often represented with simple, less-accurate models, in order to lessen that computational burden. Transmitting antennas are often modelled as imposed current point-sources, while receiving antennas are often modelled as *probes* which do not disturb the field that they measure. In reality, any co-resident antennas in an imaging chamber will mutually couple, leading to a disturbance of the electromagnetic fields in the imaging chamber which will not be captured in the computational wave model. The frequency of operation can be chosen to minimize this effect [9], which unfortunately reduces the set of measurements available to the imaging algorithm.

Another source of modelling error is the exclusion of other measurement apparatus from the computational wave model. Antennas are often represented as either point-sources or probes, but in reality, antennas must be attached to a measurement device, such as a vector network analyzer (VNA), via a set of cables and switches. A VNA reports its measurements as S-parameters, but inversion algorithms often operate on either electric or magnetic *field* values. The process of developing a suitable translation between real-world measurements and their representations in the computational domain is known as *calibration*. Calibration often requires taking measurements with some reference object [10], which may be impossible in certain imaging scenarios.

The work in this thesis was undertaken after studying and using several of the state-of-the-art tools which have been developed in the EIL. Broadly speaking, the work undertaken in this thesis is an attempt to develop an imaging tool which does not suffer from the drawbacks mentioned above. It will be shown that those drawbacks can be avoided, at the cost of image resolution. The primary imaging application explored in this thesis is grain-bin imaging, and exploratory experiments are made using ultrasound hardware. The main issues that this thesis seeks to solve are briefly highlighted here.

1.2 Grain-Bin Imaging

The main application of the algorithm developed in this thesis is the monitoring of grain via RF measurements taken from a set of stationary antennas within storage silos. A quantitative imaging technique which uses full-wave, non-linear contrast-source inversion (CSI) [11] to detect localized grain spoilage in metallic bins has been shown to be successful [12]. In order to image the grain using CSI, however, it is currently necessary to estimate the shape of the heap of grain within the storage bin, and use the shape of the heap as prior information. Imaging via CSI is also computationally expensive, potentially requiring several hours of computation time [13].

The transmission tomography algorithm presented in this thesis can tolerate an inaccurate physical model of the grain bin. Transmission tomography

is tolerant to modelling errors because transmission tomography does not use a complete model of wave propagation phenomena. It is therefore hypothesized that transmission tomography could be used as a grain imaging algorithm which is more robust than CSI, but with lower resolution than CSI.

1.3 Ultrasound Tomography

Some of the tools for MWI which have been developed at the EIL have recently been applied to acoustic imaging [14], however, synthetic acoustic imaging studies have shown that successful imaging requires an amount and quality of data which is currently unattainable in a real-world acoustic imaging device [15]. The difficulty in performing wave-field imaging (WFI) with acoustic hardware lies in the operating wavelength. Commercial ultrasound transducers operate in the megahertz range, and ultrasound imaging is often used for targets whose physical properties are similar to those of water. The speed of sound in water is roughly 1500 m s^{-1} , so the operating wavelength of an acoustic WFI system would be roughly 1.5 mm. In contrast, the operating wavelength of a typical biomedical MWI system is roughly 10 cm.

In WFI it is necessary that transducers' positions are known with sub-wavelength certainty, in order to place time-harmonic field measurements on the correct *Riemann sheet* [16]. The small wavelength that is used in acoustic imaging is troublesome because the wavelength of operation can be similar to the size of the transducers, which makes transducer localization difficult.

Transmission tomography, as will be explained in Chapter 2, operates on time-of-flight measurements. Time-of-flight measurements are unaffected by phase, so the issue of placing transducers on the correct Riemann sheet does not exist in transmission tomography.

Another requirement in WFI is that the fields of interest be represented with sufficient spatial resolution. At high frequencies, this can lead to storage issues on the computer performing the imaging. A rule of thumb for representing fields with a pulse basis is that the basis points should be spaced at one tenth of the wavelength of operation [17]. Representing an ultra-

sound pressure wave with a pulse basis function would require pulse basis points with less than 0.15 mm spacing. Representing a 10 cm square region with such spacing would require more than 400×10^3 basis points. If the wavelength of operation is small, then the amount of resources required to represent the fields can be a computational burden. In transmission tomography, the interrogating waves do not need to be represented in this way, which alleviates this burden of storage.

1.4 Transmission Tomography

As the two previous sections have shown, transmission tomography has the following features:

- Tolerance to modelling errors
- Small computational burden

These features make transmission tomography attractive as a fast and robust imaging technique. The details of transmission tomography will be explained in Chapter 2, and its features will be demonstrated. The rest of this thesis documents the development and testing of a transmission tomography algorithm, which will be applied in a variety of contexts, in order to create low-resolution images of targets.

Chapter 2

Transmission Tomography

Transmission tomography is a non-destructive imaging technique which produces quantitative spatial maps of the physical properties of an object, given time-domain scan data. The goal of transmission tomography is to determine the physical properties of a target by modelling the way a causal, time-domain wave pulse changes as it propagates through the target [18] [19].

2.1 General Formulation

Transmission tomography requires, at the bare minimum, a transmitting element and a receiving element which are not co-located in space. Refer to Figure 2.1 for a visual representation of the key elements of a transmission tomography system. The transmitting element emits a signal, which will be referred to as a pulse. The pulse then propagates outward, and is measured by the receiving element. As the pulse propagates outward from the transmitter, it naturally decays as its wave front expands. The pulse is further modified by the transmission medium and the object of interest. The exact nature of the decay and modification is complex. For example, electromagnetic wave propagation is described by Maxwell's equations [20], while acoustic wave propagation in water can be described by a coupled set of partial differential equations which relate pressure variations to fluid velocity [21].

Transmission tomography simplifies the complex modification of the trans-

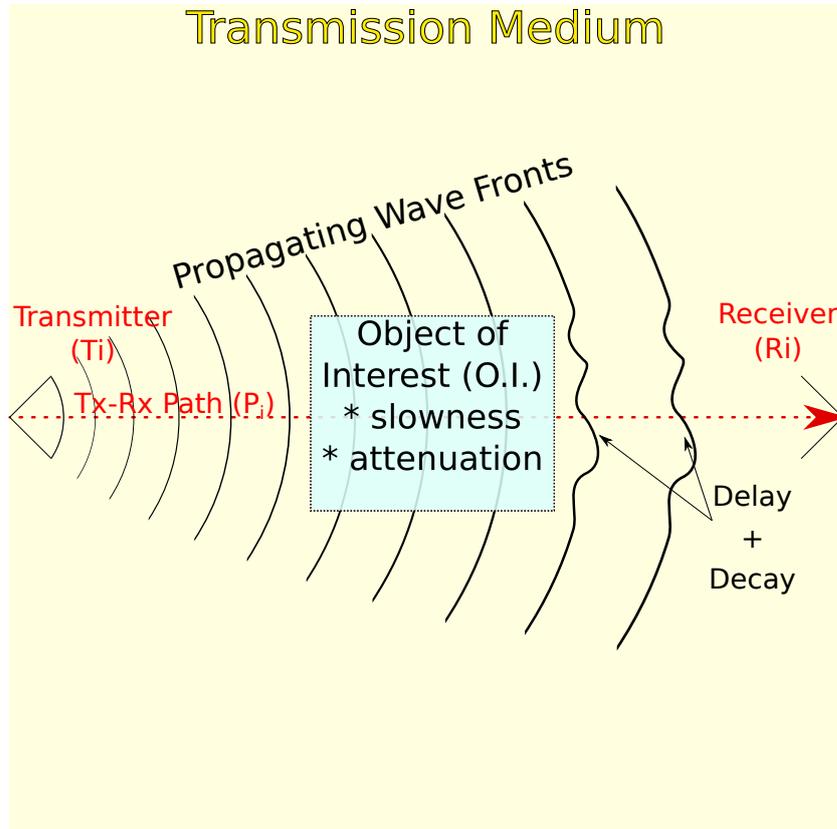


Figure 2.1: Physical setup of a general transmission tomography system, showing a single transmitter-receiver path

mitted pulse by assuming that the modification depends only on the material properties along the straight path from the transmitter to the receiver and the attenuation due to the expansion of the wave front. For a given transmitter and receiver, the propagating wave front is treated as a ray which travels from the transmitter to the receiver, ignoring the effects of refraction and diffraction. This assumption will be referred to as the *ray assumption*. The decay due to the expansion of the wave front will be referred to as *geometric attenuation*. Geometric attenuation is simple to model, for example, waves generated by an acoustic point-source experience $\frac{1}{r}$ geometric attenuation of their power, where r is the distance between the wave front and the source.

As the ray propagates from the transmitting element to the receiving element, it encounters regions of varying speed and attenuation. For example,

an acoustic pulse may propagate through a sample of bone, which is set in a region of pure water. In this case, the ray originates in a region of water, where the pulse propagates at roughly 1500 m s^{-1} , and experiences near-zero attenuation. In bone, the pulse would travel at roughly 4080 m s^{-1} and experience attenuation of $20 \text{ dB cm}^{-1} \text{ MHz}^{-1}$ [22].

The physical properties of the media through which the pulse propagates will impact the delay and the decay of the pulse, as viewed by the receiving element. In particular, the slowness of the media will impact the pulse's delay, and the attenuation of the media will impact the pulse's decay. Transmission tomography uses many pulses' delays and decays to extract information about the slowness and attenuation of the object of interest.

The ray assumption implies that transmission tomography can only provide useful information about the material properties *within* a region whose boundary is determined by the transmitter-receiver paths that are used for imaging. The only material properties that are visible to the transmission tomography algorithm are the properties between the transmitters and receivers. This point is explained more deeply in Chapter 5.

The following sections in this chapter describe a numerical implementation of transmission tomography. Section 2.2 describes a fundamental approximation which enables transmission tomography. The concepts of total-field and scattered-field imaging are described in Subsection 2.2.1. Then, Subsection 2.2.2 explains how transmission tomography is implemented in a computer program. Section 2.3 describes methods for determining the delay and decay of a pulse. Finally, Section 2.4 provides equations for converting slowness and attenuation to other physical properties.

2.2 Wave Propagation Model

Suppose an object of interest is interrogated by many pulses, each of which is generated by a transmitting element at some position and measured by a receiving element at some other position. Let i index the unique transmitter-receiver configurations. Refer to Figure 2.1 for a visual depiction of a transmitter-receiver configuration. Let T_i and R_i be the transmitter and receiver for the

configuration indexed by i , respectively. Let P_i be the straight path from T_i to R_i . An example of the signal measured by R_i is shown in Figure 2.2.

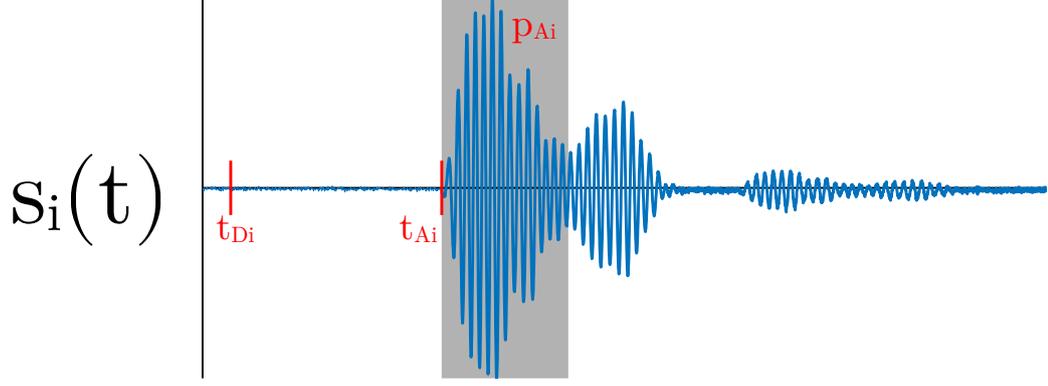


Figure 2.2: Example of a transmission tomography scan, as measured by receiver R_i

The transmitter, T_i , sends out a pulse at time t_{Di} , with some power p_{Di} . The pulse is assumed to be an oscillating function with centre frequency f_c , in order to align the transmission tomography presented in this thesis with standard acoustic imaging techniques, however, it is not strictly necessary to assume that the pulse has a centre frequency, as will be discussed in Section 2.3. The time of arrival (TOA) of the pulse, t_{Ai} , depends on the pulse's time of transmission, t_{Di} , and the slowness, c^{-1} , of the medium along the path. This dependence is expressed in the following equation:

$$t_{Ai} = t_{Di} + \int_{P_i} c^{-1}(\vec{r}(l)) dl \quad (2.1)$$

The power of the received pulse, p_{Ai} , depends on the transmitted power, p_{Di} , the distance-dependent geometric attenuation, $g(P_i)$, and the attenuation, α along the path. This dependence is expressed in the following equation:

$$p_{Ai} = p_{Di} \times g(P_i) \times 10^{\frac{-100f_c}{1 \times 10^6} \int_{P_i} \alpha(\vec{r}(l)) dl} \quad (2.2)$$

The exponent with a base of ten, and the factors of 100 and 1×10^6 in Equation 2.2 are present because α is expressed in units of $\text{dB cm}^{-1} \text{MHz}^{-1}$. This equation can be presented in a simpler way by taking logarithms, as in

the following equation:

$$\frac{-100f_c}{1 \times 10^6} \int_{P_i} \alpha(\vec{r}(l)) dl = 10 \log_{10} \left(\frac{p_{Ai}}{p_{Di} \times g(P_i)} \right) \quad (2.3)$$

The integrals in Equation 2.1 and Equation 2.3 can be isolated to produce the following set of equations:

$$\int_{P_i} c^{-1}(\vec{r}(l)) dl = t_{Ai} - t_{Di} \quad (2.4)$$

$$\int_{P_i} \alpha(\vec{r}(l)) dl = \frac{-1 \times 10^6}{100f_c} \left(10 \log_{10} \left(\frac{p_{Ai}}{p_{Di} \times g(P_i)} \right) \right) \quad (2.5)$$

The quantities under the integrals are the material properties. The quantities t_{Ai} and p_{Ai} will be referred to as *features* of the measured signal, $s_i(t)$. The time-delay of the received pulse is determined by the slowness of the propagation medium, along the straight-line path from the transmitter to the receiver. Likewise, the decay of the signal is determined by the attenuation of the propagation medium, along the same straight-line path.

If α was to be expressed in units of Np m^{-1} , the term on the right-hand side of Equation 2.5 would simplify to the following equation:

$$\int_{P_i} \alpha(\vec{r}(l)) dl = -\ln \left(\frac{p_{Ai}}{p_{Di} \times g(P_i)} \right) \quad (2.6)$$

The attenuation equation is expressed in the form of Equation 2.4 in order to align the work in this thesis with standard acoustic imaging techniques.

2.2.1 Total-Field and Scattered-Field Imaging

Consider the quantities present in Equation 2.4 and in Equation 2.5. The path, P_i , is known, since the path depends only on the location of the transmitting element and the receiving element. The geometric attenuation, $g(P_i)$, is known, since P_i is known. The centre frequency, f_c is also known. The material properties ($c^{-1}(\vec{r})$, and $\alpha(\vec{r})$) are unknown, that is, they are the quantities that will be reconstructed. The quantities t_{Ai} , t_{Di} ,

p_{Ai} , and p_{Di} are ostensibly unknown. Two of these quantities, t_{Ai} and p_{Ai} , are features which can be extracted from the time domain scans. Considering Figure 2.2, it is clear that the information about the TOA of the pulse, as well as the power of the pulse, are present in the received time-domain pulse. In Section 2.3, some techniques for extracting features from the time-domain scans are discussed. The two remaining quantities, t_{Di} and p_{Di} , *cannot* be extracted from the received time domain pulse. In this section, it will be shown how transmission tomography can be performed when these parameters are unknown. It will also be shown how transmission tomography can be performed when these parameters *are* known.

The quantities, t_{Di} and p_{Di} , will be referred to as *compensation factors*. If the compensation factors are known, then transmission tomography is possible via *total-field imaging*, otherwise, imaging must be performed via *scattered-field imaging*. Total-field imaging relies on interrogating the imaging domain with pulses **once**, and reconstructing the physical properties of the imaging domain from that single-pass set of data. The process of total-field imaging is summarized as follows:

1. The imaging domain is interrogated with pulses from several transmitters.
2. The pulses are received and measured by several receivers per transmitter.
3. The absolute physical properties are calculated, using the following information:
 - The paths' geometry
 - The compensation factors
 - The pulse features

Scattered-field imaging relies on interrogating the imaging domain with pulses **twice**, and reconstructing the physical properties of the imaging domain from the difference between the two sets of data. One of the interrogation steps is performed with a *known* target in the imaging domain, while the

other interrogation step is performed with an *unknown* target in the imaging domain. The data produced by the known target are referred to as *incident* data, while the data from the unknown target are referred to as *total* data.

As an example, consider a medical imaging scenario, where an ultrasound array is being used to image a body part. The array and the body part may be submerged in water, which would act as a matching medium between the ultrasound transducer and the body part. In this example, the *incident* data would be the measurements taken with the array submerged in water, and with the body part absent. Then, a new set of measurements would be taken with the body part present. This new set of measurements would make up the *total* data.

The differences between the total and incident data are then used to calculate the difference between the known target and the unknown target. Crucially, scattered-field imaging can be performed when the compensation factors are unknown. The process of scattered-field imaging is summarized as follows:

1. A known target is placed in the imaging domain.
2. Incident data are collected.
3. An unknown target is placed in the imaging domain.
4. Total data are collected.
5. The relative differences in the physical properties are calculated, using the following information:
 - The paths' geometry
 - The features of the incident pulses
 - The features of the total pulses
6. The absolute physical properties are calculated by adding the relative property differences to the properties of the known target.

Suppose there is a set of *known* target properties which produce a corresponding set of data. These properties will be referred to as incident properties, and their corresponding data will be referred to as incident data. Suppose there is another set of *unknown* target properties which produce a corresponding set of data. These unknown properties will be referred to as total properties, and their corresponding data will be referred to as total data.

The incident data are related to the incident properties, and the total data are related to the total properties, according to the following equations, derived from Equation 2.4 and Equation 2.5:

$$\int_{P_i} c^{-1\text{INC}}(\vec{r}(l)) dl = t_{Ai}^{\text{INC}} - t_{Di} \quad (2.7)$$

$$\int_{P_i} \alpha^{\text{INC}}(\vec{r}(l)) dl = \frac{-1 \times 10^6}{100f_c} 10 \log_{10} \left(\frac{p_{Ai}^{\text{INC}}}{p_{Di} \times g(P_i)} \right) \quad (2.8)$$

$$\int_{P_i} c^{-1\text{TOT}}(\vec{r}(l)) dl = t_{Ai}^{\text{TOT}} - t_{Di} \quad (2.9)$$

$$\int_{P_i} \alpha^{\text{TOT}}(\vec{r}(l)) dl = \frac{-1 \times 10^6}{100f_c} 10 \log_{10} \left(\frac{p_{Ai}^{\text{TOT}}}{p_{Di} \times g(P_i)} \right) \quad (2.10)$$

The superscripts, **INC** and **TOT** are used to differentiate between incident and total quantities, respectively. The compensation factors (t_{Di} and $p_{Di} \times g(P_i)$) are identical for the incident and total data.

On their own, Equation 2.9 and Equation 2.10 are the continuous representation of the total-field formulation of transmission tomography. With total-field formulation, the material properties are recovered by finding a solution to Equation 2.9 and Equation 2.10, over each path indexed by i .

The scattered-field formulation of the transmission tomography system is attained by subtracting Equation 2.7 from Equation 2.9, and subtracting Equation 2.8 from Equation 2.10.

$$\int_{P_i} \left(c^{-1\text{TOT}}(\vec{r}(l)) - c^{-1\text{INC}}(\vec{r}(l)) \right) dl = t_{Ai}^{\text{TOT}} - t_{Ai}^{\text{INC}} \quad (2.11)$$

$$\int_{P_i} \left(\alpha^{\text{TOT}}(\vec{r}(l)) - \alpha^{\text{INC}}(\vec{r}(l)) \right) dl = \frac{-1 \times 10^6}{100f_c} 10 \log_{10} \left(\frac{p_{Ai}^{\text{TOT}}}{p_{Ai}^{\text{INC}}} \right) \quad (2.12)$$

The following difference terms are introduced to simplify the left-hand sides of Equation 2.11 and Equation 2.12:

$$\delta c^{-1}(\vec{r}) = c^{-1\text{TOT}}(\vec{r}) - c^{-1\text{INC}}(\vec{r}) \quad (2.13)$$

$$\delta\alpha(\vec{r}) = \alpha^{\text{TOT}}(\vec{r}) - \alpha^{\text{INC}}(\vec{r}) \quad (2.14)$$

Then, Equation 2.11 and Equation 2.12 can be rewritten using Equation 2.13 and Equation 2.14. These equations are the continuous representation of the scattered-field formulation of transmission tomography.

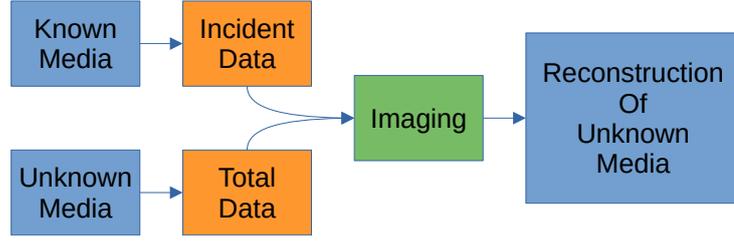
$$\int_{P_i} \delta c^{-1}(\vec{r}(l)) dl = t_{A_i}^{\text{TOT}} - t_{A_i}^{\text{INC}} \quad (2.15)$$

$$\int_{P_i} \delta\alpha(\vec{r}(l)) dl = \frac{-1 \times 10^6}{100f_c} 10 \log_{10} \left(\frac{p_{A_i}^{\text{TOT}}}{p_{A_i}^{\text{INC}}} \right) \quad (2.16)$$

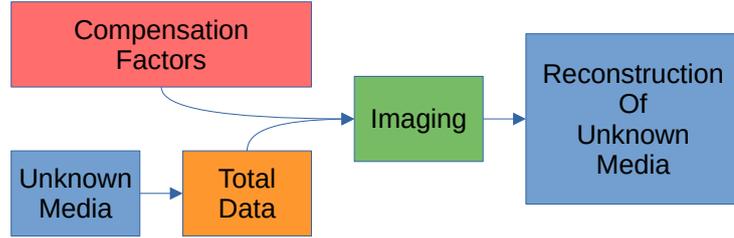
Subtracting the total-field equations from the incident-field equations eliminates the compensation factors, thus, the scattered-field formulation of transmission tomography permits imaging without compensation factors. The remaining right-hand-side terms only contain features of the time domain scans. The material properties are recovered by finding a solution to the above system, over each path indexed by i , then calculating $c^{-1\text{TOT}}(\vec{r})$ and $\alpha^{\text{TOT}}(\vec{r})$ from Equation 2.13 and Equation 2.14. Note that recovering the material properties from the scattered-field formulation requires knowledge of $c^{-1\text{INC}}(\vec{r})$ and $\alpha^{\text{INC}}(\vec{r})$. The flow of information in the total-field and scattered-field formulations is depicted in Figure 2.3. If the incident material properties are well-characterized, and scan data exist for both the incident field and the total field, then the scattered-field formulation can be used. If those conditions are not met, then the total-field formulation (Equation 2.9 and Equation 2.10) must be used, which requires that the compensation factors are known.

2.2.2 Discretizing the Model

In order to find a solution to the transmission tomography system, it is necessary to translate the continuous representations of the transmission to-



(a) Scattered-field formulation. Two sets of measurements must be taken to reconstruct the unknown target.



(b) Total-field formulation. One set of measurements is taken, and a set of compensation factors must be used in order to reconstruct the unknown target.

Figure 2.3: Flow of information in scattered-field and total-field formulations.

mography systems into forms which can be processed by a computer. In other words, the model must be discretized, both spatially and temporally. The model is spatially discretized by expressing the two material properties in some finite basis. The model is temporally discretized by representing the scans as sequences of samples, sampled at frequency f_s .

Spatial Discretization

Let $\{\phi_1(\vec{r}), \dots, \phi_N(\vec{r})\}$ be the chosen basis functions, and let $\{c_1^{-1}, \dots, c_N^{-1}\}$ and $\{\alpha_1, \dots, \alpha_N\}$ be the corresponding basis coefficients for $c^{-1}(\vec{r})$ and $\alpha(\vec{r})$, in a basis with N functions. Then, the material properties can be expressed

in the finite basis as follows:

$$c^{-1}(\vec{r}) = \sum_{j=1}^N c_j^{-1} \phi_j(\vec{r}) \quad (2.17)$$

$$\alpha(\vec{r}) = \sum_{j=1}^N \alpha_j \phi_j(\vec{r}) \quad (2.18)$$

By substituting Equation 2.17 and Equation 2.18 into Equation 2.4 and Equation 2.5, moving constant terms outside of the integral, and expanding the logarithm term, we obtain the following equations:

$$\sum_{j=1}^N c_j^{-1} \int_{P_i} \phi_j(\vec{r}(l)) dl = t_{Ai} - t_{Di} \quad (2.19)$$

$$\sum_{j=1}^N \alpha_j \int_{P_i} \phi_j(\vec{r}(l)) dl = \frac{-1 \times 10^6}{100f_c} (10 \log_{10}(p_{Ai}) - 10 \log_{10}(p_{Di} \times g(P_i))) \quad (2.20)$$

Performing this conversion for both the total-field and the incident field, we have the following equations:

$$\sum_{j=1}^N c_j^{-1 \text{TOT}} \int_{P_i} \phi_j(\vec{r}(l)) dl = t_{Ai}^{\text{TOT}} - t_{Di} \quad (2.21)$$

$$\sum_{j=1}^N \alpha_j^{\text{TOT}} \int_{P_i} \phi_j(\vec{r}(l)) dl = \frac{-1 \times 10^6}{100f_c} (10 \log_{10}(p_{Ai}^{\text{TOT}}) - 10 \log_{10}(p_{Di} \times g(P_i))) \quad (2.22)$$

$$\sum_{j=1}^N c_j^{-1 \text{INC}} \int_{P_i} \phi_j(\vec{r}(l)) dl = t_{Ai}^{\text{INC}} - t_{Di} \quad (2.23)$$

$$\sum_{j=1}^N \alpha_j^{\text{INC}} \int_{P_i} \phi_j(\vec{r}(l)) dl = \frac{-1 \times 10^6}{100f_c} (10 \log_{10}(p_{Ai}^{\text{INC}}) - 10 \log_{10}(p_{Di} \times g(P_i))) \quad (2.24)$$

These equations hold simultaneously for each path (P_i) indexed by i . The

four sets of simultaneous equations can then be expressed in matrix form.

$$\underline{\mathbf{L}}\underline{c}^{-1\text{TOT}} = \underline{t}_A^{\text{TOT}} - \underline{t}_D \quad (2.25)$$

$$\underline{\mathbf{L}}\underline{\alpha}^{\text{TOT}} = \underline{q}_A^{\text{TOT}} - \underline{q}_D \quad (2.26)$$

$$\underline{\mathbf{L}}\underline{c}^{-1\text{INC}} = \underline{t}_A^{\text{INC}} - \underline{t}_D \quad (2.27)$$

$$\underline{\mathbf{L}}\underline{\alpha}^{\text{INC}} = \underline{q}_A^{\text{INC}} - \underline{q}_D \quad (2.28)$$

Here, \mathbf{L} is a matrix such that $\mathbf{L}_{i,j} = \int_{P_i} \phi_j(\vec{r}(l)) dl$, and the underlined quantities are vectors created by the concatenation of the quantities in Equation 2.21, Equation 2.22, Equation 2.23, and Equation 2.24. The matrix, \mathbf{L} will be referred to as the *path integral matrix*. The cumbersome quantities from the right-hand-side of Equation 2.22 and Equation 2.24 are compressed by defining the following terms:

$$q_{Ai}^{\text{TOT}} = \frac{-1 \times 10^6}{100f_c} 10 \log_{10}(p_{Ai}^{\text{TOT}}) \quad (2.29)$$

$$q_{Ai}^{\text{INC}} = \frac{-1 \times 10^6}{100f_c} 10 \log_{10}(p_{Ai}^{\text{INC}}) \quad (2.30)$$

$$q_{Di} = \frac{-1 \times 10^6}{100f_c} 10 \log_{10}(p_{Di} \times g(P_i)) \quad (2.31)$$

The path integral matrix, \mathbf{L} is, in general, not a square matrix. The matrix may have more rows than columns, as in Figure 2.4a, or more columns than rows, as in Figure 2.4b. These two states are referred to as over-determined and under-determined, respectively. The number of rows in \mathbf{L} is the number paths along which the target was interrogated. The number of columns in \mathbf{L} is N , the number of functions in the chosen basis.

Taken together, Equation 2.25 and Equation 2.26, are the discrete representation of the total-field formulation of transmission tomography. The material properties can be recovered directly from the total-field formulation by calculating a solution to the two equations.

The discrete representation of the scattered-field formulation of transmission tomography is attained by subtracting Equation 2.27 from Equation 2.25 and subtracting Equation 2.28 from Equation 2.26. Let $\underline{\delta c}^{-1} =$

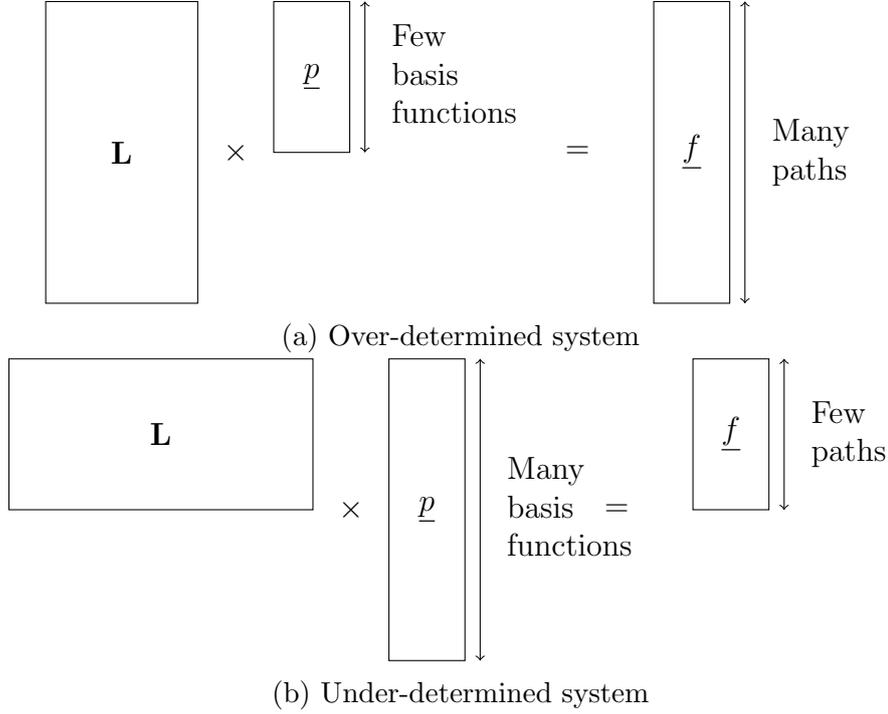


Figure 2.4: Over-determined and under-determined matrix systems. The shape of \mathbf{L} depends on the number of paths and the number of basis functions. \underline{p} is a vector of properties, and \underline{f} is a vector of features.

$\underline{c}^{-1\text{TOT}} - \underline{c}^{-1\text{INC}}$, and let $\underline{\delta\alpha} = \underline{\alpha}^{\text{TOT}} - \underline{\alpha}^{\text{INC}}$. The resulting equations are the following:

$$\mathbf{L}(\underline{\delta c}^{-1}) = \underline{t}_A^{\text{TOT}} - \underline{t}_A^{\text{INC}} \quad (2.32)$$

$$\mathbf{L}(\underline{\delta\alpha}) = \underline{q}_A^{\text{TOT}} - \underline{q}_A^{\text{INC}} \quad (2.33)$$

The only unknown quantities in Equation 2.32 and Equation 2.33 are the left-hand-side vectors, $\underline{\delta c}^{-1}$ and $\underline{\delta\alpha}$. The compensation factors have been eliminated by subtracting the incident equations from the total equations. Solving these two equations yields the vectors of physical property differences. The vectors of actual physical properties can be obtained by adding the

recovered differences to the known incident physical properties.

$$\underline{c}^{-1\text{TOT}} = \underline{c}^{-1\text{INC}} + \underline{\delta c}^{-1} \quad (2.34)$$

$$\underline{\alpha}^{\text{TOT}} = \underline{\alpha}^{\text{INC}} + \underline{\delta\alpha} \quad (2.35)$$

Finally, the vectors of property coefficients can be translated into spatial maps by expanding the physical properties in the chosen basis, using Equation 2.17 and Equation 2.18. This final translation step is the same for both total-field imaging and scattered-field imaging.

Solving Matrix Equations

Each of Equation 2.32, Equation 2.33, Equation 2.25, and Equation 2.26 use a matrix equation to relate a vector of properties to a vector of features via a path integral matrix. Let \underline{p} and \underline{f} represent the vector of properties and the vector of features in any of those equations.

$$\mathbf{L}\underline{p} = \underline{f} \quad (2.36)$$

There are several methods for calculating a \underline{p} which satisfies this equation for a given \mathbf{L} and \underline{f} . Choosing a method for solving Equation 2.36 depends largely on whether the equation is an over-determined or under-determined system. See Figure 2.4 for a graphical depiction of these two types of systems.

If Equation 2.36 is over-determined, then there may be no exact solution, that is, there is no \underline{p} which satisfies the equation. A common technique for calculating a solution to an over-determined system is least-squares analysis [23]. Since Equation 2.36 cannot be solved exactly, we instead quantify the error between $\mathbf{L}\underline{p}$ and \underline{f} , and minimize this error. The error is represented by $\mathcal{F}(\underline{p})$ in the following equation:

$$\mathcal{F}(\underline{p}) = \|\mathbf{L}\underline{p} - \underline{f}\|^2 \quad (2.37)$$

Then, we seek a value of \underline{p} which minimizes this error.

$$\underline{p} = \arg \min_{\underline{x}} \mathcal{F}(\underline{x}) \quad (2.38)$$

This optimal \underline{p} is calculated using the following equation [23]:

$$\underline{p} = (\mathbf{L}^T \mathbf{L})^{-1} (\mathbf{L}^T \underline{f}) \quad (2.39)$$

Several other techniques can be used to find solutions to over-determined matrix systems. Least-squares analysis is described here because it is the solution technique that will be used from Chapter 5 onward.

If Equation 2.36 is under-determined, then there are several solutions, that is, there are several values of \underline{p} which satisfy the equation. A common technique for calculating a particular solution to an under-determined system of equations is the conjugate-gradient least-squares (CGLS) algorithm [23] [24]. This technique differs from least-squares analysis in that it is *iterative*. An initial guess for \underline{p} is imposed, then the value of \underline{p} is repeatedly updated until some stopping condition is met. The details of the CGLS algorithm will not be explained here. While there are many other techniques available for solving under-determined matrix systems, CGLS is mentioned here because it is used in Chapter 3 and Chapter 4.

Temporal Discretization

The time-domain scans, $s_i(t)$ are stored as sequences of samples, sampled at frequency f_s . It is equivalent to say the scan is sampled at sampling interval t_s , where $t_s = \frac{1}{f_s}$. The sampled version of the scan will be represented with square brackets, and indexed with n , according to the following relation:

$$s_i[n] = s_i(nt_s) \quad (2.40)$$

2.3 Feature Extraction Details

Transmission tomography operates on time-domain scans, in order to produce spatial maps of physical properties. As discussed in Section 2.2, the physical properties are chosen to match the delay and decay features of the set of measured time-domain pulses. The process of calculating the pulse features from the time-domain scans is *feature extraction*.

For each scan, its two features (TOA and power) are calculated in series. First, the TOA is calculated, then the power of the measured pulse is calculated. The TOA must be calculated first because the power calculation relies on having an accurate time window in which the pulse appears.

2.3.1 Time-Of-Arrival Determination

The pulses generally consist of several different important phenomena. The important phenomena are listed below, in order of appearance. This list is not exhaustive, since there may be infinite scattering events to consider, however, this list does contain the most common and important phenomena which appear in scans.

1. Noise due to cross-talk: This occurs when the generator-to-transmitter signal path and the receiver-to-measurement-device signal path get close to each other. This can happen if the pulse generator and the measurement device are the same device. Cross-talk can also be experienced when a switching device is used to multiplex the pulse generator and the measurement device through several transmitters and several receivers.
2. Low-power background noise: After the cross-talk has extinguished, the measurement device will see no signal except for background noise.
3. Fastest-path pulse: There are many possible paths along which the wave front will propagate on its journey from the transmitter to the receiver. The receiver is first excited by the portion of the wave front which travels along the path of least time.

4. Straight-path pulse: The portion of the wave front which travels along the straight path from the transmitter to the receiver arrives after the fastest-path pulse. The straight-line path is generally not the path of least-time from the transmitter to the receiver, in any non-homogeneous medium.
5. Slow pulses and reflections: The portions of the wave front which propagate along paths which are slower than the straight path will arrive at various points in time after the arrival of the straight-line pulse. Additionally, there may be late-arriving reflections due to the geometry of the imaging domain. For example, the imaging domain may be in some walled chamber.

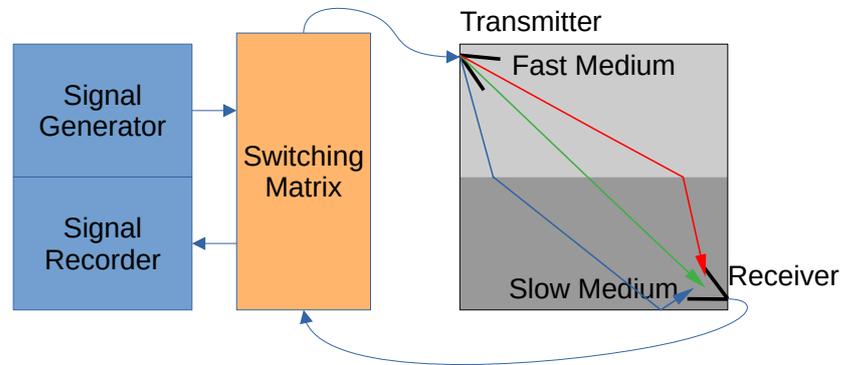
A visual representation of the cause of these phenomena is shown in Figure 2.5, for a simple imaging target. Note that the phenomena described above have intersecting support in time, that is, the phenomena blur into each other.

Three different TOA detection algorithms have been implemented in this thesis:

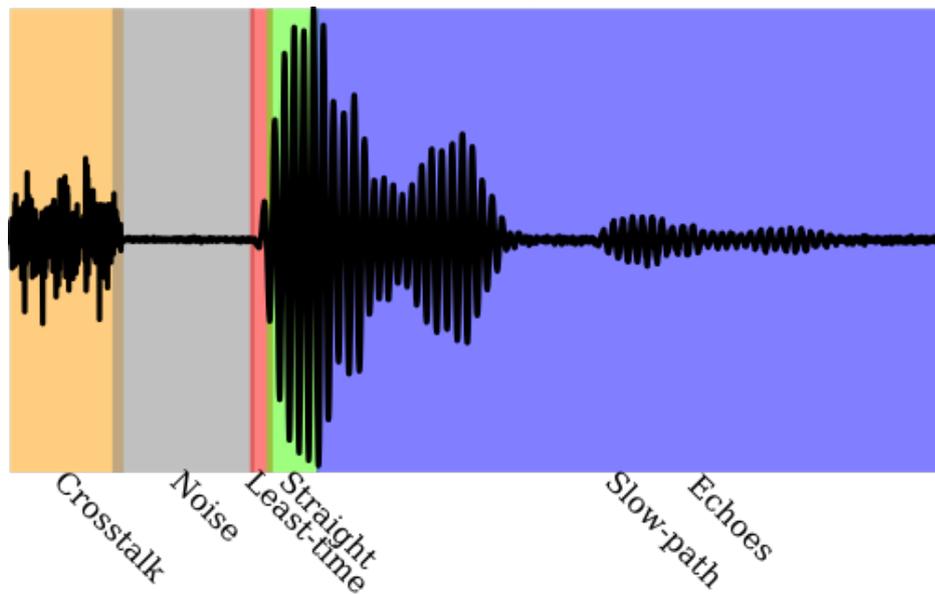
1. peak detection,
2. the modified energy ratio (MER), and
3. the variance score method.

Each of these algorithms has certain idiosyncratic benefits and drawbacks. These benefits and drawbacks mean that different sets of data require different algorithms for determining TOA. The different algorithms will be described here, with guidance as to where each algorithm is applicable. For each description, suppose we seek the TOA, t_{A_i} from the time-domain scan $s_i(t)$. A sample scan is shown in Figure 2.6. This scan is included as a visual aid for the explanation of the various TOA determination algorithms.

The sample scan's extracted TOA is shown in Table 2.1, for each of the three extraction methods, showing that they all agree, despite their different approaches. The experiments in this thesis make use of either the MER



(a) Schematic of sources of scan phenomena. The red path is the path of least-time. The green path is the straight-line path. The blue path is one of many slow paths.



(b) Scan phenomena. Sections are colour-coded to match the sources shown above.

Figure 2.5: Physical causes of observed phenomena in time-domain scans

method or the variance score method. Specifically, the experiments in Chapter 3 use the variance score method, the experiments in Chapter 4, Chapter 5, Chapter 6 and Chapter 7 all use the MER method. Further details will be provided in each chapter. The peak detection method is not used in this

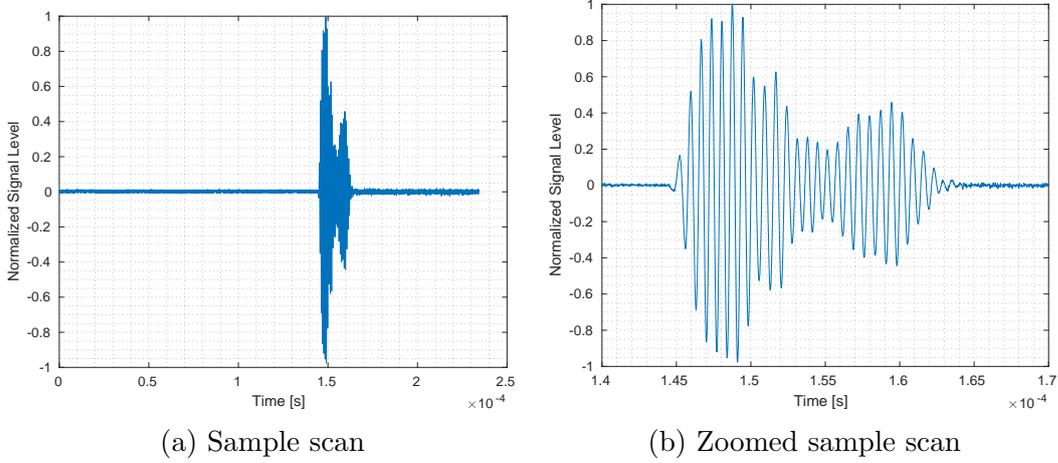


Figure 2.6: Sample scan for illustrating algorithms for determining time of arrival

Table 2.1: Time-Of-Arrival of Sample Scan

Method	Peak	MER	Variance Score
TOA [μs]	148.7	145.1	144.5

thesis, however, it is presented here since it has proven to be useful in other scenarios.

Peak

Peak detection, as a TOA algorithm, is a simple method. The TOA is calculated as follows:

$$t_{Ai} = \arg \max_t |s_i(t)| \quad (2.41)$$

That is, the TOA is taken to be the time where the scan reaches its peak absolute value. This algorithm is depicted in Figure 2.7, for the sample scan. This method has been shown to be effective on the multimodal ultrasound breast imaging (MUBI) data set [25], however, the peak detection method can fail if the time-domain scan contains a target with significant attenuation. The straight-line path may get attenuated enough that reflected and refracted pulses which pass around the target appear stronger than the straight-line pulse. Additionally, this method can fail due to strong cross-talk at the

beginning of the scan.

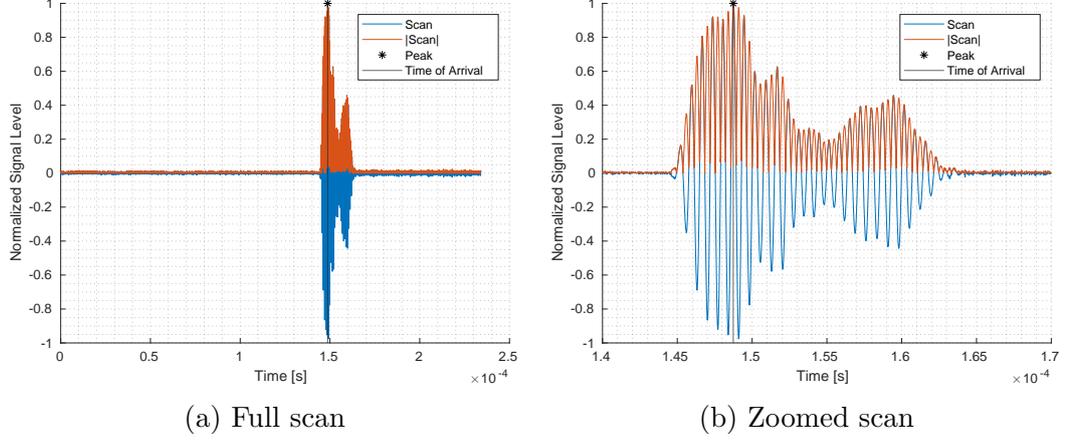


Figure 2.7: Time-of-arrival determination via peak detection

The negative impact of late-arriving reflection can be decreased by truncating the time-domain scans. If the minimum possible wave speed is known ahead of time, then it is possible to define a maximum possible time-of-flight (T_{MAX}) for each path with $T_{\text{MAX}} = \frac{\text{path distance}}{c_{\text{MIN}}}$. The scan may be safely truncated to T_{MAX} in order to ignore some late reflections, however, some late reflections may still be visible.

This method is implemented as a simple linear search through the scan, so its computational complexity is $\mathcal{O}(N_S)$, for a scan with N_S total samples.

Modified Energy Ratio

The MER method has also been shown to be effective on the MUBI data set [26]. The MER method finds a point in the scan where the power sharply increases. This method relies on calculating an energy ratio, $R(t)$.

$$R(t) = \frac{\int_t^{t+w} (s_i(\tau))^2 d\tau}{\int_{t-w}^t (s_i(\tau))^2 d\tau} \quad (2.42)$$

The ratio compares the energy in a window after t to the energy in a window before t . The parameter w is the size of the window, and it is typically taken to be 3 periods of the centre frequency of the scan. The ratio, $R(t)$, is then

modified via multiplication with the absolute value of the scan. The result, $M(t)$ is the *modified energy ratio*.

$$M(t) = |s(t)| R(t) \quad (2.43)$$

Finally, the TOA is chosen to be the maximizer of $M(t)$

$$t_{Ai} = \arg \max_t (M(t)) \quad (2.44)$$

Multiplying $R(t)$ by $|s(t)|$ helps ensure that if random noise produces a spurious jump in $R(t)$, then that spike is ignored. This algorithm is depicted in Figure 2.8, for the sample scan.

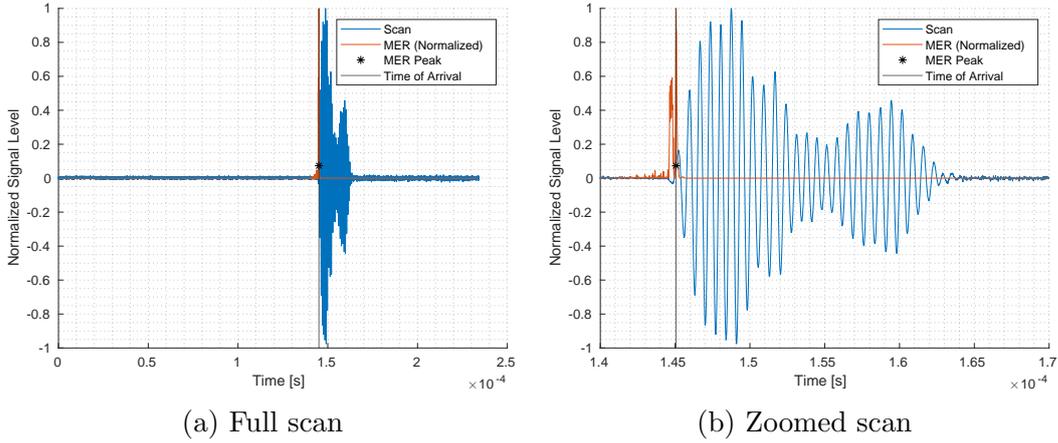


Figure 2.8: Time-of-arrival determination via modified energy ratio

This method is more resilient to cross-talk than the peak detection method, as long as cross-talk is the very first phenomenon encountered in the signal. In that case, cross-talk will appear in the numerator of $R(t)$ only if cross-talk appears in the denominator of $R(t)$, and no spurious spike in $R(t)$ will occur. This method is susceptible to strong, late-arriving reflections. If the difference between the arrival of the reflection and the arrival of the straight-line pulse is greater than w , and the late pulse has more power than the straight-line pulse, then that method will incorrectly assign t_{Ai} to the late pulse. Again, as was the case with the peak detection method, it may be possible to truncate the scans in order to decrease the negative impact of

some late reflections.

A naive implementation would achieve $\mathcal{O}(N_S w)$ computational complexity for this algorithm. The actual complexity can be reduced to $\mathcal{O}(N_S)$ by keeping track of the numerator and denominator terms for $R(t)$, and iteratively subtracting out old terms and adding in new terms as the windows move along the scan.

Variance Score

The method which will herein be referred to as the variance score method is adapted from a statistical TOA detection method, known as the AIC-picker, which has been shown to be effective on the computed ultrasound risk evaluation (CURE) dataset [27]. The variance score method seeks to divide the scan into two regions: the region before the arrival of the first pulse, and the region after the arrival of the first pulse. The time t_{Ai} is the point which separates these two regions. Suppose that the total time of the scan is T . The TOA, t_{Ai} , is chosen to be the minimizer of the score function, $S(t)$.

$$t_{Ai} = \arg \min_t S(t) \quad (2.45)$$

The score function is a ratio of two variances, $V_L(t)$, and $V_R(t)$. These are the left variance and the right variance, respectively. The left variance is the variance of s_i from 0 to t , while the right variance is the variance of s_i from t onward.

$$V_L(t) = \frac{1}{t} \int_0^t (s_i(\tau) - \bar{s}_L)^2 d\tau \quad (2.46)$$

$$V_R(t) = \frac{1}{T-t} \int_t^T (s_i(\tau) - \bar{s}_R)^2 d\tau \quad (2.47)$$

The left and right variance calculations require the calculation of the left and right means ($\bar{s}_L(t)$ and $\bar{s}_R(t)$), which are calculated with the following

equations:

$$\bar{s}_L(t) = \frac{1}{t} \int_0^t s_i(\tau) d\tau \quad (2.48)$$

$$\bar{s}_R(t) = \frac{1}{T-t} \int_t^T s_i(\tau) d\tau \quad (2.49)$$

This algorithm is depicted in Figure 2.9, for the sample scan. Specifically, all three steps of the variance score algorithm are shown. The left and right variances are shown in Figure 2.9a and Figure 2.9b. These figures highlight how the left and right variance can encode information about TOA. It is clear when the dividing time is near the TOA, the right-variance peaks. Then, as the dividing time moves beyond the TOA, the left variance increases sharply and the right variance decreases sharply. The ratio of these two values gives the variance score, shown in Figure 2.9c and Figure 2.9d. The minimum of the variance score is a good approximation of the TOA.

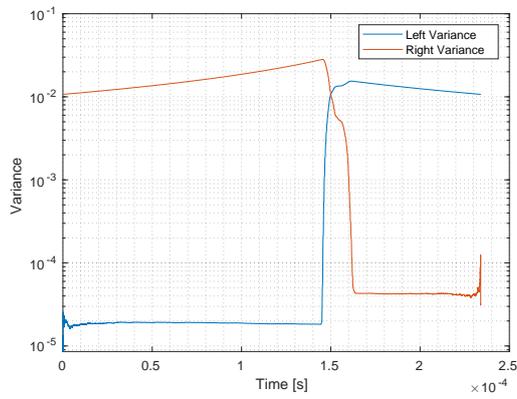
The variance score method excels at picking time-of-flight in the presence of noise [27], however, this method can experience the same failure modes as MER.

Care must be taken in the implementation of this TOA calculation method. A naive implementation would have $\mathcal{O}(N_S^2)$ computational complexity, as each of the N_S samples in time requires the calculation of the variance of the remaining $(N_S - 1)$ samples. Similar to the MER method, this computation can achieve $\mathcal{O}(N_S)$ complexity, by keeping track of the two mean values and the two variance values as the algorithm iterates through the scan.

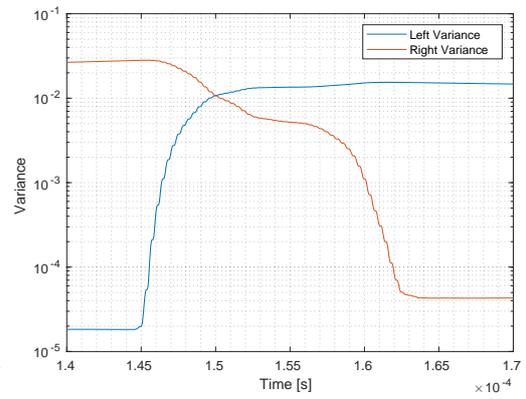
2.3.2 Pulse Power Determination

Once the TOA is determined, the power of the received pulse can then be calculated. The power is calculated using the following formula:

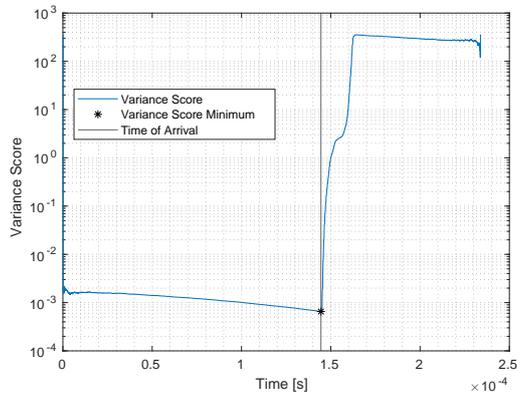
$$p_{Ai} = \frac{1}{w_p} \int_{t_{Ai}}^{t_{Ai}+w_p} (s_i(t))^2 dt \quad (2.50)$$



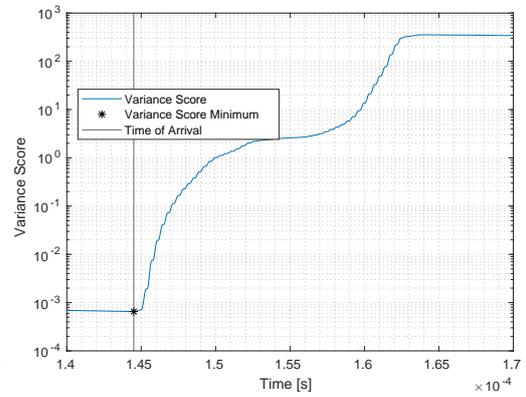
(a) Full view of left and right variance



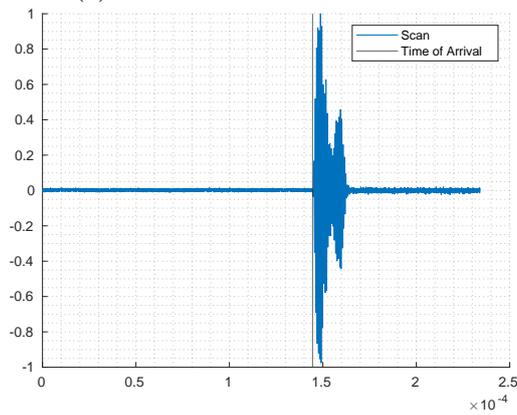
(b) Zoomed view of left and right variance



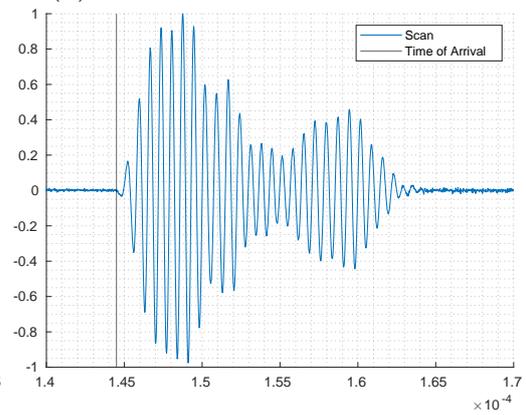
(c) Full view of variance score



(d) Zoomed view of variance score



(e) Full view of time of arrival via variance score



(f) Zoomed view of time of arrival via variance score

Figure 2.9: Time-of-arrival determination via variance score

The power is simply the average of the square of the received signal, over a period of time. The beginning of the period is taken to be the TOA. The parameter, w_p must be set prior to imaging, and is dependent on the shape of the transmitted pulse. For example, if the transmitted pulse is 5 wavelengths of an enveloped sinusoid with period t_{tran} , then a good value for w_p would be $5 \times t_{\text{tran}}$.

2.3.3 Ignoring Data

The feature extraction algorithms will occasionally fail, due to any of the causes listed above. The feature extraction step generally fails by identifying a wrong TOA for a particular pulse. If the TOA is identified incorrectly, then the pulse power calculation will also be wrong, since the power calculation relies on an accurate approximation of the envelope of the pulse. In this thesis, a data processing step takes place after all of the features are extracted for a particular imaging session.

The processing consists of identifying scans whose features are impossible, given some knowledge of the imaging domain. For example, suppose a pulse interrogates some object, immersed in a background with speed c_b , along a path of length l . Furthermore, suppose it is known that there are certain bounds, c_{MIN} and c_{MAX} on the expected speed of the material being imaged. For a scattered-field imaging scenario, the expected change in TOA could then be bounded, as follows:

$$l \left(\frac{1}{c_{\text{MAX}}} - \frac{1}{c_b} \right) \leq (t_{A_i}^{\text{TOT}} - t_{A_i}^{\text{INC}}) \leq l \left(\frac{1}{c_{\text{MIN}}} - \frac{1}{c_b} \right) \quad (2.51)$$

Likewise, given a background attenuation, α_b , and the bounds of α_{MIN} and α_{MAX} on the expected attenuation of the material being imaged, it is possible to put bounds on the relationship between the power of the total pulse and the power of the incident pulse. The bounds are calculated as follows:

$$l(\alpha_{\text{MIN}} - \alpha_b) \leq \frac{-1 \times 10^6}{100 f_c} 10 \log_{10} \left(\frac{p_{A_i}^{\text{TOT}}}{p_{A_i}^{\text{INC}}} \right) \leq l(\alpha_{\text{MAX}} - \alpha_b) \quad (2.52)$$

The criteria from Equation 2.51 and Equation 2.52 allow certain scans to be identified. Scans which fail either criteria are then ignored. Scans are ignored by simply removing their corresponding rows from the \mathbf{L} matrix, which was introduced in Equation 2.25 and Equation 2.26.

A median filter technique, presented in [27], was shown to be effective in identifying and remedying erroneous scans. That technique compares the features of the scan received by a particular receiver to the scans received by its neighbours. If the scan's calculated features differ from its neighbours' features by a certain amount, then the feature is replaced with the median of its neighbours' features. This technique is not implemented in this thesis, as the algorithm that was developed treats each transmitter-receiver path as an independent entity. It is not assumed that the transmitters or receivers are arranged in any regular order, so a regional median filter was not investigated.

2.4 Extracting Auxiliary Material Properties

The transmission tomography tool described in this thesis reconstructs two physical properties: slowness, and attenuation. These material properties are diagnostically useful on their own, however, it is potentially necessary to translate these two properties to other properties. For example, the electromagnetic imaging tools currently in use at the EIL expect materials to be expressed in terms of complex-valued permittivity, $\tilde{\epsilon}$. The EIL also has several tools which perform time-harmonic WFI of acoustic fields. These tools expect materials to be expressed in terms of complex-valued compressibility, $\tilde{\kappa}$.

This section presents two techniques for translating from inverse wave speed and attenuation (c^{-1} and α) to either complex permittivity ($\tilde{\epsilon}$) or complex compressibility ($\tilde{\kappa}$). Depending on the imaging context, and the source of measured data, there are other properties which may be of interest. For example, in a geological imaging context, the diagnostic properties of interest may be the porosity and saturation of soil. The translations presented here are chosen specifically for their utility to the members of the EIL.

Detailed derivations of the equations in this section are shown in Ap-

pendix C.

2.4.1 Acoustic Properties

Given c^{-1} and α , in an ultrasound imaging context, the real and imaginary parts of the complex-valued compressibility of a material can be calculated. The following translation assumes that density (ρ) is constant:

$$\Re[\tilde{\kappa}] = \frac{(c^{-1})^2 - \left(\frac{\alpha \ln(10)}{4\pi \times 10^5}\right)^2}{\rho} \quad (2.53)$$

$$\Im[\tilde{\kappa}] = \frac{-\alpha c^{-1} \ln(10)}{2\pi \times 10^5 \rho} \quad (2.54)$$

2.4.2 Electromagnetic Properties

Given c^{-1} and α , in an ultrasound imaging context, the real and imaginary parts of the complex-valued permittivity of a material can be calculated. The following translation assumes that permeability (μ) is constant:

$$\Re[\tilde{\epsilon}] = \frac{(c^{-1})^2 - \left(\frac{\alpha \ln(10)}{4\pi \times 10^5}\right)^2}{\mu_0} \quad (2.55)$$

$$\Im[\tilde{\epsilon}] = \frac{-\alpha c^{-1} \ln(10)}{2\pi \times 10^5 \mu_0} \quad (2.56)$$

Chapter 3

Two-Dimensional Ultrasound Imaging on Unstructured Meshes

3.1 Introduction

Previous work in the EIL [14] has led to the construction of various ultrasound imaging systems. One such system, depicted in Figure 3.1, consists of a set of ultrasound transducers, which are set in a planar ring, and submerged in a bucket of water.

The system supports 64 piezoelectric transducers, which can act either as transmitters or receivers. In transmission mode, the transducers convert voltages to local pressure variations. In receiving mode, the transducers convert local pressure variations to voltages. To act as a transmitter, a transducer can be excited in one of two ways:

1. Excited by an impulse, and left to ring, or
2. Excited by a time-varying voltage source.

Impulse-excitation tends to provide more interrogation energy than exciting the transducer with a continuous signal, however, the frequency of the transducer's ringing cannot be controlled. The transducers resonate at roughly

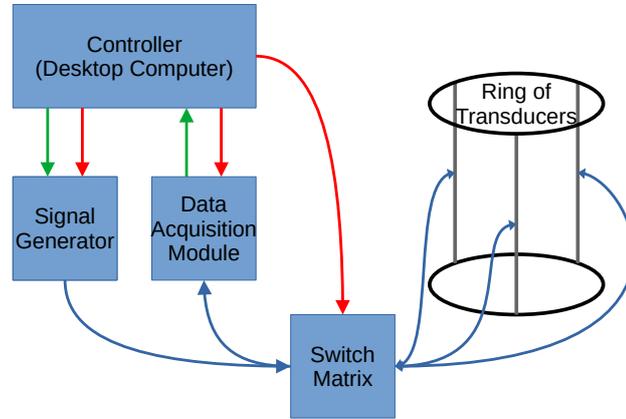


Figure 3.1: System schematic. Red lines indicate digital control connections, with the arrow head pointing to the controlled device. Green lines indicate digital signal connections and blue lines indicate analog signal connections, with arrows pointing in the direction of signal flow. Only 3 of the 64 transducer supports are shown, for visual simplicity.

1.4 MHz. A frequency of 1.4 MHz in water supports a wave with a wavelength of 1.1 mm. Previous work [14] sought to use the 1.4 MHz transducers to perform CSI in an imaging system similar to the one shown in Figure 3.2. The small wavelength made it difficult to perform full-wave imaging with the given transducers, however, the small wavelength is actually ideal for transmission tomography. Therefore, this bucket imaging system was chosen as a test-bed for initial transmission tomography tests, and the transducers were excited by a voltage impulse.

3.2 Procedure

The bucket holds 64 transducers, equally spaced on a ring of radius 11.4 cm. Three simple phantoms were developed and built by Hannah Fogel, a fellow student and member of the EIL. These phantoms were used to test the resolution of the tomography algorithm. The phantoms were made of a mixture of gelatin and water, which was allowed to set in a cylindrical vessel with a



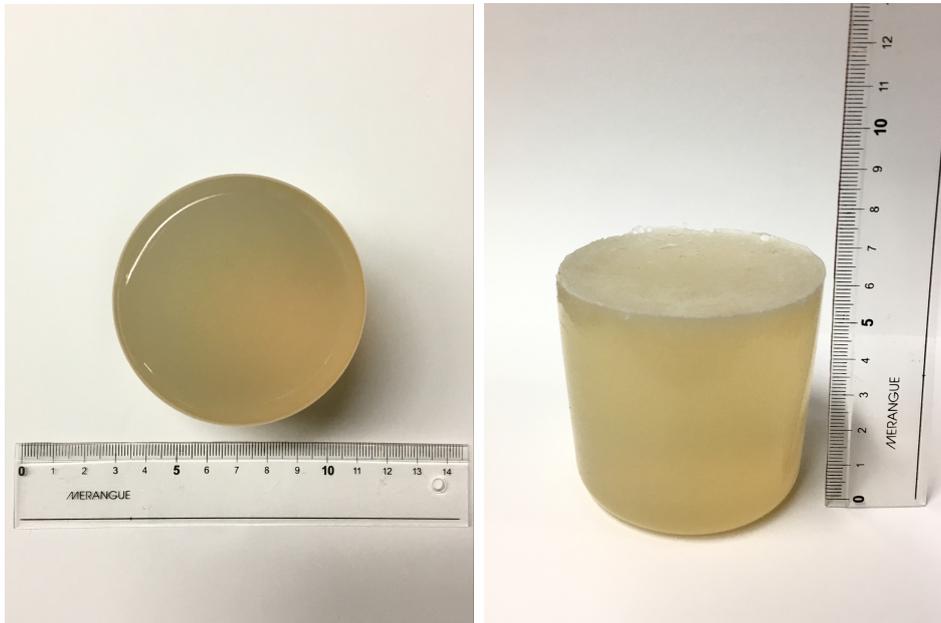
Figure 3.2: Ring of ultrasound transducers, submerged in water

diameter of 8 cm. One phantom, shown in Figure 3.3, was a simple monolithic region of the gelatin/water mixture. Two other phantoms were made by boring holes through homogeneous phantoms along the vertical axis. One phantom contained a single bore-hole, and the final phantom contained two bore-holes. All of the bore-holes had a diameter of 1 cm. The speed of sound in the gelatin was known to be $1.6 \times 10^3 \text{ m s}^{-1}$, and the material was known to be near-lossless.

The gelatin phantoms were interrogated by pulses from each of the 64 transducers. For each transmission, the non-transmitting transducers acted as receivers, generating a total of $64 \times 63 = 4032$ time-domain scans. In order to reduce random noise, each transmission event was repeated several times, and the resulting set of measurements were averaged. The data acquisition cards were set up with the following parameters:

- Transmitters excited by an impulse
- Sampling frequency of 50 MHz
- Averaging was enabled, with 50 recordings per scan

The 4032 time-domain scans were then inverted using the transmission tomography algorithm. The imaging domain was set to be a circular region



(a) Horizontal dimensions of homogeneous gelatin phantom

(b) Vertical dimension of homogeneous gelatin phantom

Figure 3.3: Horizontal and vertical scale of homogeneous gelatin phantom

with a radius of 11.4 cm. The imaging domain was meshed with a characteristic length of 5 mm, producing an unstructured mesh which contains 2135 triangular elements. The open-source software `gmsh` [28] was used to mesh the imaging domain. Each of the three gelatin phantoms was imaged on this mesh. The tomography equations were solved via CGLS, truncated at 10 iterations. The resulting two-dimensional images were then quantitatively evaluated for their reconstruction accuracy.

3.2.1 Data Processing Details

The imaging system used in this experiment produces strong crosstalk at the beginning of each scan, so each scan is preprocessed prior to imaging, in order to remove the crosstalk. This crosstalk decays earlier than the expected TOA for all of the scans, the preprocessing consists of deleting 300 entries from the beginning of each scan.

The TOA is extracted with the variance score method, and the window

size for the power calculation is set to six wavelengths, which is equivalent to 214 samples.

3.2.2 Development

It is necessary to choose a basis function for representing the material properties. A pulse basis was chosen for this test. A pulse basis is a set of functions, $\{\phi_1, \phi_2, \dots, \phi_N\}$, each of which has finite support. Each pulse basis function is unit-valued within its support, and zero-valued elsewhere. The pulse basis functions' supports are all mutually disjoint.

A pulse basis was chosen because it is easy to integrate pulse basis functions along a straight line. The integral, $\int_{P_i} \phi_j(\vec{r}(l)) dl$ is simply the length of the intersection of P_i with the support of $\phi_j(\vec{r})$.

Pulse bases are commonly supported by either square grids or unstructured meshes. Unstructured meshes permit more modelling freedom than square grids, so a mesh-based pulse basis was chosen for this experiment.

With a mesh-based pulse basis chosen, the matrix \mathbf{L} , from Equation 2.25 and Equation 2.26, is calculated by comparing each transmitter-receiver path to each mesh element, and calculating their intersection. Therefore, filling \mathbf{L} has a computation complexity of $\mathcal{O}(IJ)$, where I is the total number of transmitter-receiver paths, and J is the total number of basis functions. The number of mesh elements intersected by any path is small, when compared to the total number of mesh elements, therefore the resulting \mathbf{L} is a sparse matrix. Let l_c be the characteristic side-length of the mesh elements. As l_c decreases, the number of mesh elements intersected by each path grows proportionally to $\frac{1}{l_c}$, while the total number of mesh elements grows proportionally to $\left(\frac{1}{l_c}\right)^3$. The fill ratio of \mathbf{L} , then, is proportional to $\left(\frac{1}{l_c}\right)^2$.

Seeking direct solutions to Equation 2.25 and Equation 2.26 will produce solutions with high variations, which lead to non-physical images. This symptom has two causes:

1. Small errors which occur during feature extraction.
2. Mesh elements which are intersected by a small number of transmitter-

receiver paths (lonely mesh elements).

Suppose there is some path whose extracted features contain small errors, and which intersects a lonely mesh element. Each element which is intersected by the path will have its reconstructed physical properties impacted by the erroneous features, that is, the properties will change slightly in order to achieve the erroneous feature. The elements which will be modified the most are the elements that are intersected by the smallest number of paths. Therefore, if there are lonely mesh elements and erroneous features, then the lonely mesh elements' recovered physical properties will be excessively large.

This symptom can be alleviated by seeking *regularized* solutions to the tomography equations. The regularization technique employed in this test is the CGLS algorithm. This technique was chosen for its ease of implementation and its speed. For a problem with J basis functions and P different transmitter-receiver paths, solving the tomography equations with CGLS is a $\mathcal{O}(JP)$ operation. This is because each iteration of CGLS is a $\mathcal{O}(JP)$ operation, and the number of iterations is set independent of J and P and is determined by the desired spatial resolution. Matlab provides a built-in function, `lsqr()` which performs the CGLS algorithm. The built-in function was used for the inversions presented here, with appropriate options set to permit a fixed number of CGLS iterations.

3.2.3 Testing

The reconstructions for this experiment were evaluated by comparing the recovered sound-speed images to the known geometry of the gelatin phantoms. The attenuation images were not considered because the phantoms were known to be near-lossless. The following features of the reconstructed sound-speed images were quantitatively evaluated:

- Sound-speed reconstruction accuracy over the whole imaging domain
- Location of holes in the reconstructed sound-speed

The sound-speed reconstruction accuracy over the whole imaging domain was calculated using the following equation:

$$\mathcal{E}_c = \frac{\iint_D (c(\vec{r}) - c_{\text{EXP}}(\vec{r}))^2 d\vec{r}}{\iint_D (c_{\text{EXP}}(\vec{r}))^2 d\vec{r}} \quad (3.1)$$

Here, $c(\vec{r})$ is the reconstructed sound-speed, $c_{\text{EXP}}(\vec{r})$ is the expected sound-speed and D is the imaging domain.

The radius of the phantoms, their sound-speed, and the radius of the holes in the phantoms were accurately recorded prior to imaging. Some other geometric features of the gelatin phantoms were not accurately recorded, including the location of the phantoms within the imaging chamber and the locations of the holes that were placed in the gelatin phantoms. The location of the phantoms and the locations of their holes were calculated from the reconstructed images via the following procedure:

1. Mesh elements were classified, via a sound-speed threshold, as either gelatin or water.
2. The centre of the reconstructed phantom was set as the centre of the gelatin elements' bounding box.
3. The locations of the holes were determined by identifying elements within the phantom that are located in a low sound-speed neighbourhood. An element's neighbourhood is the set of elements whose centroid is within 5 mm of its centroid. Holes were assumed to be at the centroids of elements whose neighbourhoods' speed most closely matches the speed of water.

3.3 Results

The reconstructed sound-speed and attenuation images, for the three gelatin phantoms, are shown in Figure 3.4.

A quantitative evaluation of each phantom's reconstruction is shown in Table 3.1, which documents the calculated locations of the holes, and the

reconstruction accuracy over the entire imaging domain.

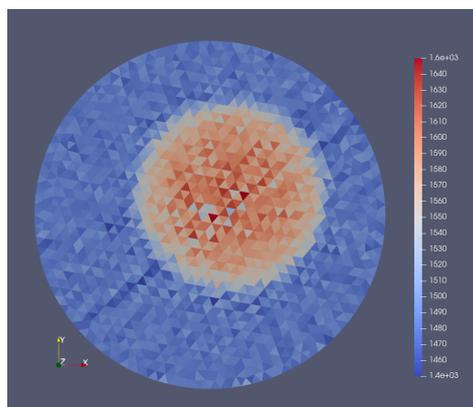
Table 3.1: Imaging metrics for two-dimensional ultrasound experiment

Phantom Name	Calculated Hole Location(s) [mm]	Expected Hole Location(s) [mm]	Domain Speed Error
Homogeneous	N/A	N/A	1.666×10^{-4}
One-Hole	(26.6, 2.88)	(25.1, 3.65)	2.378×10^{-4}
Two-Hole	(12.5, -7.34) (-15.5, 0.551)	(16.3, -5.14) (-14.8, 2.94)	3.029×10^{-4}

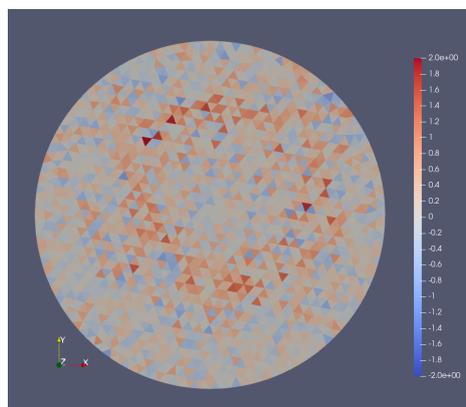
3.4 Discussion

The reconstructions in Figure 3.4 clearly show that it is possible to image the simple gelatin phantoms. It is clear that the sound-speed image in Figure 3.4a is the homogeneous gelatin phantom, the sound-speed image in Figure 3.4c is the single-hole phantom, and the sound-speed image in Figure 3.4e is the double-hole phantom. These images also show some straight-line artefacts. This type of sharp artefact is a motivator for the work in Chapter 5, which emerged as a method of smoothing out such imperfections. The attenuation images are not diagnostically helpful in this scenario. The attenuation images do present a faint outline of the gelatin phantoms, but these images are not clear enough to be useful for tissue identification.

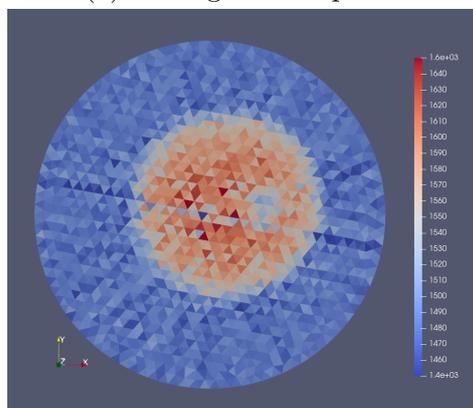
These time-of-flight tomography images demonstrate that it is possible to generate quantitative images of low-contrast targets, given pulse transmission data. These positive results motivate the development of a three-dimensional imaging framework, which will be presented in Chapter 4.



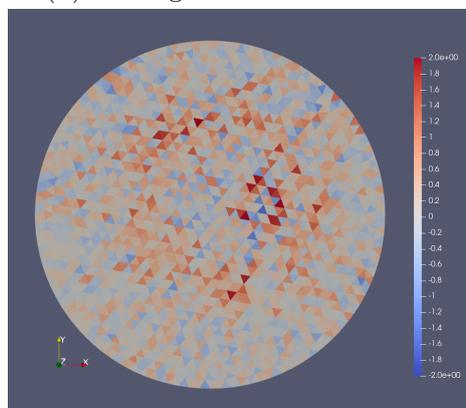
(a) Homogeneous Speed



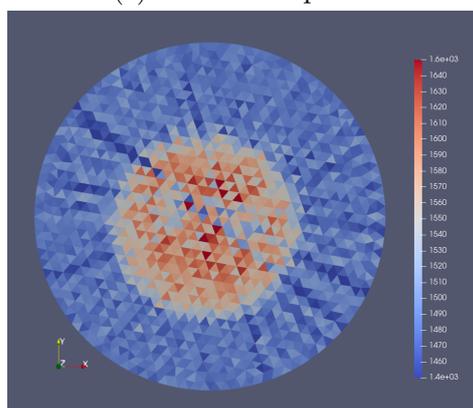
(b) Homogeneous Attenuation



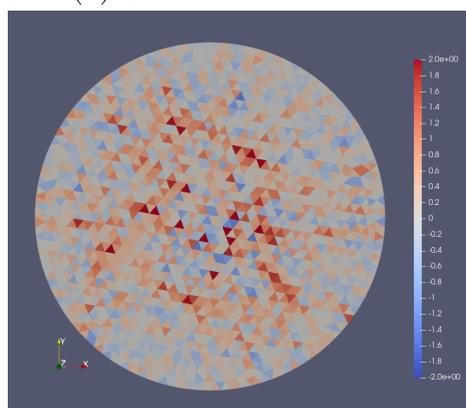
(c) One-Hole Speed



(d) One-Hole Attenuation



(e) Two-Hole Speed
 Column Colour Range:
 1400 to 1600 m s^{-1}



(f) Two-Hole Attenuation
 Column Colour Range:
 -2 to +2 $\text{dB cm}^{-1} \text{MHz}^{-1}$

Figure 3.4: Pulse-basis reconstructions of gelatin phantoms

Chapter 4

Three-Dimensional Pulse-Basis Transmission Tomography in Grain Bins

4.1 Introduction

The results presented in Chapter 3 demonstrated that transmission tomography is capable of imaging a geometrically simple target. A heap of grain, stored in a bin, is another geometrically simple target, however, a heap of grain is not a two-dimensional entity. Imaging a grain heap requires a three-dimensional imaging system. This chapter describes an exploratory experiment, which uses a mesh-based pulse-basis transmission tomography tool to calculate the spatially-varying wave speed and attenuation within a grain bin.

A tool which could extract wave-speed and attenuation from time-domain pulses in a grain bin would be useful for monitoring the status of stored grain. When grain is stored in a bin, it is desirable to know the moisture and temperature of the grain at each location within the bin. The grain's moisture is of economic interest to the farmer who stores the grain. Grain is traded by mass, so farmers would like to keep their grain as moist as possible so that they may maximize profit. A negative aspect of moist grain

is that moist grain is more likely to spoil than dry grain. When grain spoils, it generates heat, which results in a local spike in the grain's temperature. These phenomena give farmers a strong motivation to have detailed spatial information about the temperature and moisture of the grain in their bins. The agriculture company, AGCO Winnipeg, in partnership with the EIL, is developing tools that use non-linear, full-wave electromagnetic imaging to provide this information, via the following process:

1. Use the electromagnetic measurements to calculate the electrical properties (i.e. permittivity and conductivity) of the grain.
2. Translate the electrical properties to diagnostic properties which are of interest to farmers, such as moisture and temperature.

Electromagnetic measurements must be calibrated [10] when they are used as data in non-linear MWI, such as CSI. The goal of calibration is to negate the impact of all of the cabling and the switching network that are required when taking RF measurements. The process of calibration requires three entities:

1. Measurements taken while the imaging system is in a completely *known* state, called the calibration data
2. A computational or analytic model of the imaging system in the completely *known* state
3. A second set of measurements taken while the imaging system is in an unknown state

For example, in grain bin imaging, the calibration data may be collected from a completely empty bin. An analytic model does not exist for calculating radiation patterns in an arbitrary resonant enclosure, however, an empty grain bin can be modelled computationally. The calibration data, together with the 'empty bin' model, can then be used to calibrate another set of measurements, which may correspond to the bin filled with some amount of grain. The calibrated measurements can then be used to perform non-linear

MWI. A caveat of calibration is that the quality of the calibration process depends on the similarity between the object from which calibration data are collected and the object to be imaged [10].

Calibrating data for MWI in grain bins requires an accurate computational model of the bin and a set of measurements taken from a known distribution of grain. These two entities are then used to calibrate a second set of measurements, which are taken from a grain bin with some unknown distribution of grain. There are two problems with calibrating data for grain bins:

1. Data from the known distribution of grain may not be available, depending on how the grain bin is used.
2. Calibration may fail, due to excessive differences between the two sets of data. These differences may stem from the data acquisition system, whose characteristics change over time. These differences may also stem from a large differences between the shape of the two grain distributions.

Transmission tomography does not require such a calibration procedure, since the only properties of interest in the electrical measurements are delay and decay. The experiment presented in this chapter uses transmission tomography to quantitatively image a heap of grain, stored in a metallic chamber.

4.2 Procedure

This experiment relies on synthetic data. In Chapter 3, it was possible to experiment with real, measured data because the ground-truth properties of those targets were known. Ground-truth data for stored grain, however, are not easily accessible. The physical properties of grain which govern electromagnetic wave propagation are well-characterized [29], but a full-scale experiment which controls both the dielectric properties and the shape of stored grain is beyond the scope of this thesis. Using synthetic data allows

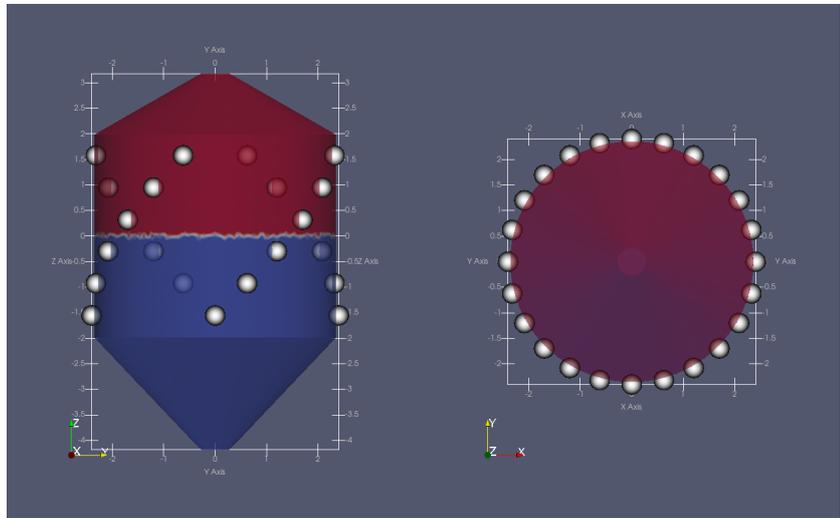
for numerical validation of imaging techniques, since the electrical properties of the target are known ahead of time.

Synthetic data were generated with the finite-difference time-domain (FDTD) method [30], using the open-source software `meep` [31], operating on a simulated model of a prototype grain imaging system located on-campus at the University of Manitoba. The bin consists of a cylindrical region, with a truncated cone region on top and a truncated cone region on the bottom. A total of 24 antennas are distributed in the cylindrical portion of the bin, in a quadruple-helix layout. The geometry of the bin is shown in Figure 4.1, along with the true wave-speed profile.

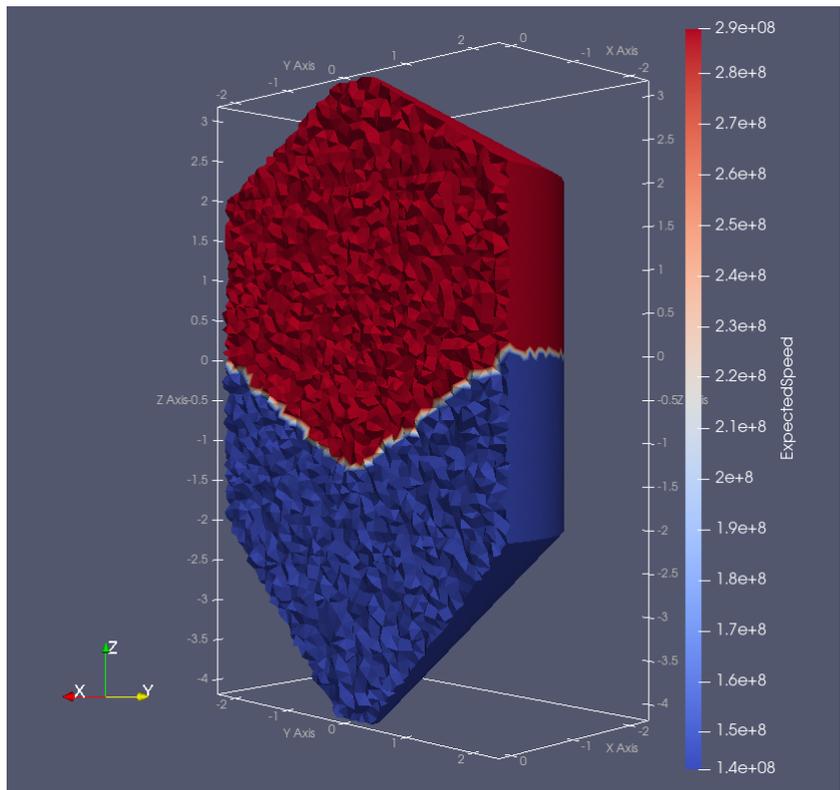
The model of the bin was separated into a grain region and an air region. The air region was modelled as free-space, with a complex relative permittivity of 1.0. The grain region was modelled by a non-magnetic material with a complex relative permittivity of $4.4 - 0.51i$, which is the approximate permittivity of corn with a moisture content of 17% [32]. The expected wave-speed in the grain was derived from the complex permittivity of the grain by solving Equation 2.55 for c . The expected wave-speed in the two simulated materials is shown in Table 4.1, along with the expected attenuation. The bin was filled halfway with grain, and the top surface of the grain was shaped in a negative cone, with a declination angle of 30 deg. This shape models a bin which is in the process of being emptied from a port at the bottom of the bin.

The scattered-field formulation was chosen for this experiment, and the incident medium was set as the bin filled with grain. Each of the 24 antennas was sequentially excited, while the other antennas acted as receivers, producing a set of 552 time-domain scans. The antennas were excited by an enveloped sinusoidal signal, with a centre frequency at 60 MHz, sampled at 9 GHz. Data were generated for both the incident medium and the grain distribution shown in Figure 4.1.

The synthetic time-domain data were used to image the wave-speed within the bin on meshes of varying density. Once again, the open-source software `gmsh` [28] was used to generate meshes. The density of the mesh was controlled by setting the *characteristic length*, l_c , in the meshing program. Five



(a) Antenna Positions - Front View (b) Antenna Positions - Top View



(c) Expected Speed Profile

Figure 4.1: Synthetic grain distribution for grain-bin imaging experiment

different characteristic lengths were studied, ranging from 0.1 m to 1.6 m. The wave-speed reconstruction accuracy, \mathcal{E}_c , was evaluated with the following formula:

$$\mathcal{E}_c = \frac{\iiint_{D'} (c(\vec{r}) - c_{\text{EXP}}(\vec{r}))^2 d\vec{r}}{\iiint_{D'} (c_{\text{EXP}}(\vec{r}))^2 d\vec{r}} \quad (4.1)$$

This equation is similar to Equation 3.1, with two modifications:

1. The surface integral is replaced by a volume integral, since this is a three-dimensional problem.
2. The domain of integration, D' , is modified to only include mesh elements which are intersected by a transmitter-receiver path.

Mesh elements which are not intersected by any paths have undefined speeds, so they will not be included in the error integral. The domain of integration is therefore different for each mesh. It is still reasonable to compare the errors from different meshes to each other, because the error integral is normalized by the integral in the denominator.

Table 4.1: Expected material properties of grain reconstruction

Material	$\Re[\varepsilon_r]$	$\Im[\varepsilon_r]$	c [m s ⁻¹]	α [dB cm ⁻¹ MHz ⁻¹]
Grain	4.4	-0.51	1.4251×10^8	2.2040×10^{-4}
Air	1.0	0	2.9979×10^8	0

An additional metric is used to measure how thoroughly the imaging mesh is interrogated by transmitter-receiver paths. This metric will be referred to as the interrogation ratio, and it is calculated with the following equation:

$$R_{\text{INT}} = \frac{M_{\text{INT}}}{M_{\text{MESH}}} \quad (4.2)$$

Here, M_{INT} is the number of mesh elements which are interrogated by at least one transmitter-receiver path, and M_{MESH} is the total number of mesh elements.

4.2.1 Data Processing Details

The scans' TOA were extracted with the MER method. The scans feature short arrival times, due to the use of a centre frequency of 60 MHz. Some scans' TOA do not satisfy the $3\text{-}\lambda$ requirement, which was discussed in Section 2.3. The expected TOA along the shortest transmitter-receiver path would be 1.89 ns, which is equivalent to 26 samples, however, the $3\text{-}\lambda$ requirement means that the MER algorithm expects at least 450 samples prior to the TOA. In order to meet this requirement, each scan is prefixed with some samples of white noise. The power of the noise is set to achieve a signal-to-noise ratio of 1000.

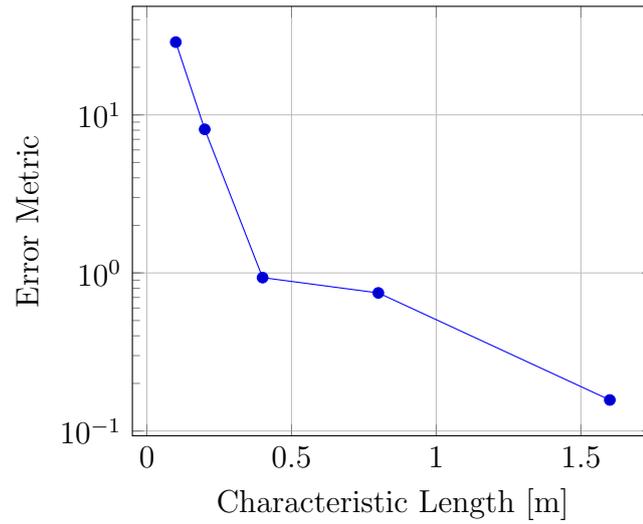
4.3 Results

The reconstructed wave-speed profiles are shown in Figure 4.3. Here, the mesh elements which are not interrogated by any transmitter-receiver paths are simply assigned the properties of the incident medium. Based on the antenna configuration shown in Figure 4.1, it is clear that the mesh elements in the upper and lower cones will not be intersected by any transmitter-receiver paths. For the denser meshes, there are mesh elements within the cylindrical portion of the grain bin which are not intersected by any transmitter-receiver paths, as is especially evident in Figure 4.3e. The mesh elements which are not interrogated by any transmitter-receiver paths are simply assigned the physical properties of the incident medium, which is grain.

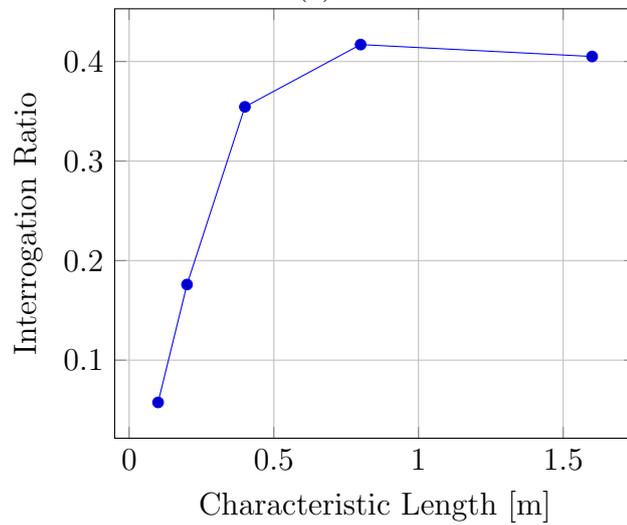
The mesh elements' speeds are depicted with a colour scale ranging from $1 \times 10^8 \text{ m s}^{-1}$ to $3 \times 10^8 \text{ m s}^{-1}$. Many of the elements' recovered speeds are actually outside of this range. The colour scale was restricted to those bounds in order to capture the range of the expected wave-speeds.

The error metric values for the different reconstructions are shown in Table 4.2. Note that the error metric is a normalized metric. An error of 0 would represent a perfect reconstruction, while an error of 1 would represent a reconstruction that assigns a speed of 0 m s^{-1} to each element. For reference, the error incurred by assigning the speed of grain to the entire evaluation

domain and the error incurred by assigning the speed of air to the entire evaluation domain are included in the table. The error metric values are also plotted in Figure 4.2, along with a plot of the interrogation ratio, R_{INT} .



(a)



(b)

Figure 4.2: Error and interrogation ratio with respect to mesh characteristic length

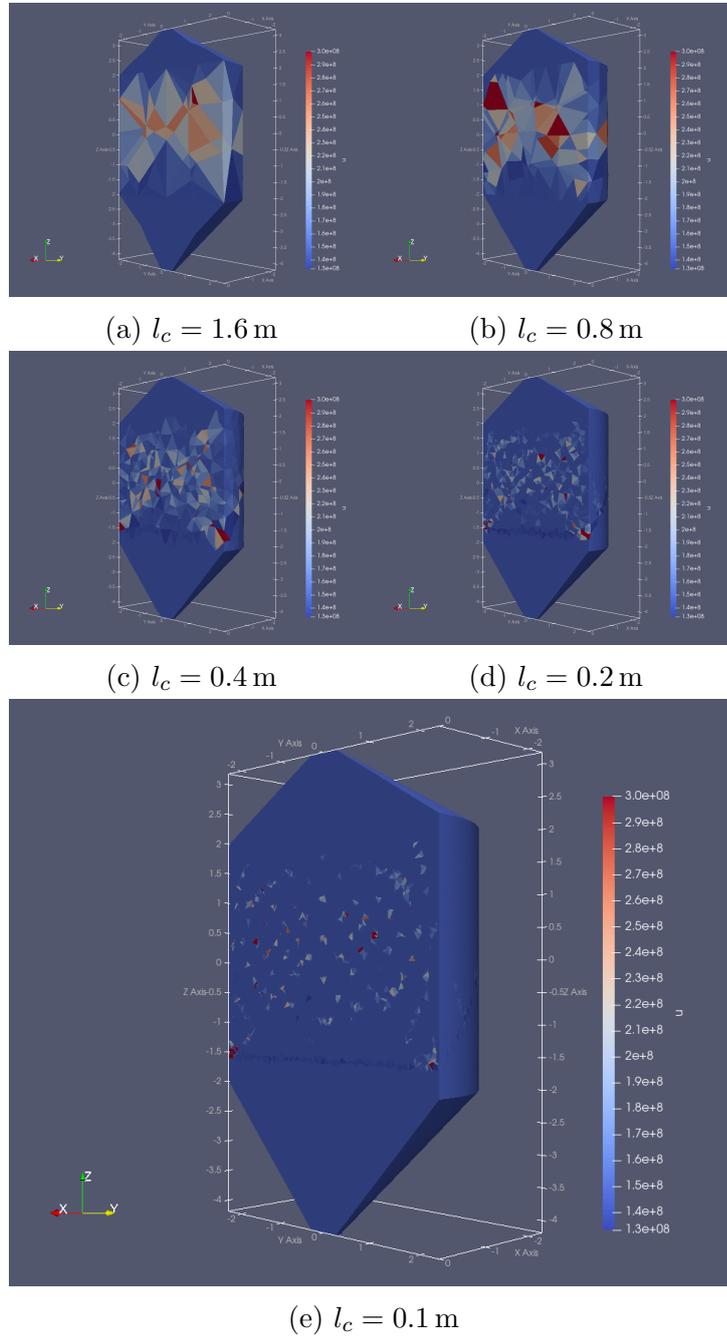


Figure 4.3: Pulse-basis reconstructions of wave speed from meep-generated electromagnetic data

Table 4.2: Accuracy of reconstruction on meshes of varying density

Characteristic Length [1 m]	Mesh Elements	Speed Error	Interrogation Ratio
1.60 m	437	0.157	0.405
0.800 m	1247	0.934	0.417
0.400 m	7299	0.934	0.354
0.200 m	43 782	8.10	0.176
0.100 m	362 764	28.9	0.0575

Reference	Speed Error
All Grain	0.116
All Air	0.249

4.4 Discussion

The images shown in Figure 4.3 do not resemble the expected wave-speed distribution shown in Figure 4.1. The error metrics shown in Table 4.2 are near 1.0 or greater than 1.0, except for the error associated with $l_c = 1.6$ m. Recall that an error of 1.0 is the error associated with assigning a uniform speed of 0 m s^{-1} . Looking at Figure 4.3a, it is clear that the result obtained from setting $l_c = 1.6$ m is not useful, despite its relatively acceptable error metric.

These results suggest that the three-dimensional transmission tomography tool is not capable of producing useful images of stored grain. The main issue with this type of imaging is that each basis function has finite support. The support of each basis function is the space occupied by the tetrahedral element to which that basis function is assigned. Each tetrahedral element is interrogated by subset of all of the pulse paths. Inversely, every tetrahedral element is *ignored* by a subset of all the pulse paths. The problem of mesh interrogation is highlighted in Figure 4.2b. For large characteristic lengths, the set of transmitter-receiver paths interrogates every mesh element which resides in the convex hull of the antennas. As the characteristic length of the mesh shrinks, though, the mesh interrogation ratio falls off sharply, which makes it difficult to draw physically coherent conclusions from the imaging results.

Consider the extreme case of an infinitely dense mesh, that is, consider a mesh with infinitely many tetrahedral elements, all of which are infinitesimally small. Furthermore, suppose that none of the pulse paths intersect each other, except at transmitter and receiver locations. In this case, each mesh element would be intersected by either 1 or 0 pulse paths, except for the elements containing the transmitter and receiver points. Along each path, there would be infinite degrees of freedom to satisfy the TOA and attenuation equations. Therefore, the TOA and attenuation equations would all be independent of each other. The resolution of the resulting images would be infinite, but there would be no spatial coupling of the recovered material properties, unless some additional regularization is imposed.

Consider the extreme case of a mesh which consists of a single tetrahedral element, which contains all of the pulse paths. In this case, the spatial coupling of the basis functions would be maximal, because the single tetrahedral element would be intersected by *all* of the pulse paths, however, the image resolution of the resulting images would be terrible, because a single material property would be assigned to the entire imaging domain.

It is necessary to strike a balance between these two extremes by choosing mesh elements which are big enough to be intersected by several pulse paths, while being small enough to offer diagnostically useful image resolution. One solution would be to increase the number and variety of pulse paths, in order to intersect more of the mesh elements, however, the target applications of agricultural and biological imaging, the number of pulse paths is set by hardware. These results show that in the target application of grain bin imaging, it is not possible to attain satisfactory image resolution, while simultaneously ensuring that each basis function is intersected by a sufficient number of straight-ray paths.

We seek to remedy this problem by using a more suitable basis, a basis which offers a higher degree of spatial coupling of the recovered material properties. A polynomial basis offers the desired spatial coupling of material properties. Polynomial functions have infinite support, so each basis function in a polynomial basis will be intersected by every pulse path. The development and testing of polynomial basis functions for grain bin imaging will be

described in Chapter 5.

Chapter 5

Three-Dimensional Polynomial-Basis Transmission Tomography in Grain Bins

5.1 Introduction

The pulse-basis reconstructions presented in Chapter 4 demonstrated that a pulse basis has some drawbacks which make it unsuitable for three-dimensional imaging. With a pulse basis, the basis functions' supports are all mutually disjoint. Spatial coupling of recovered material properties is only possible via mutual intersection of basis functions and transmitter-receiver paths. In the three-dimensional grain bin context, that amount of spatial coupling is insufficient for achieving images with diagnostically useful resolution.

This chapter presents transmission tomography with *polynomial* basis functions as a method of stabilizing the solution to the transmission tomography problem. It is known that any multivariate polynomial function can be reconstructed from a set of line-integrals of that function, as long as the number of line-integrals matches the function's degrees of freedom and the end-points of the line integrals lie on a convex surface [33] [34]. The shape of a pile of grain in a bin can be closely approximated with a low order polynomial function, therefore, it is hypothesized that representing the grain with

a polynomial basis will lead to stabler solutions than those achieved with a pulse basis.

5.2 Procedure

The experiment presented in this chapter is the same as the experiment presented in Chapter 4, except that a polynomial basis is used to represent the material properties. In Chapter 4, a series of meshes were tested, in order to examine the effect of the mesh density on the reconstruction accuracy. With a polynomial basis, there is no spatial density parameter that can be adjusted, instead, the degree of the polynomial basis will be parameter that will be examined. A total of six different bases will be examined, ranging from a basis of degree 1 to a basis of degree 6.

The error metric was modified to accommodate the change of basis. The support of the polynomial basis is naturally infinite, however, it would not be sensible to evaluate the reconstruction accuracy over an infinite domain. A natural evaluation domain would be the interior of the grain bin, however, the evaluation domain chosen for this experiment is even more restrictive than the interior of the grain bin.

The transmission tomography technique interrogates a region of space with transmitter-receiver paths in order to reconstruct a spatial map of certain physical properties. Several regions of the grain bin are not sufficiently interrogated by transmitter-receiver paths. The upper and lower cones of the bin are ignored, as are the portions of the main cylinder which are above or below the antenna locations, since there is no transmitter-receiver path which crosses through these regions. The portions of the main cylinder which are within the vertical range of the antennas and near the walls of the bin are also not sufficiently interrogated, due to the arrangement of the antennas on the wall of the bin. For these reasons, the evaluation domain was set to be the convex hull of the antenna positions. The convex hull of the antenna positions is shown in Figure 5.1, set within a cylinder that represents the cylindrical portion of the grain bin. The full set of transmitter-receiver paths, shown in Figure 5.1b fails to thoroughly interrogate certain regions

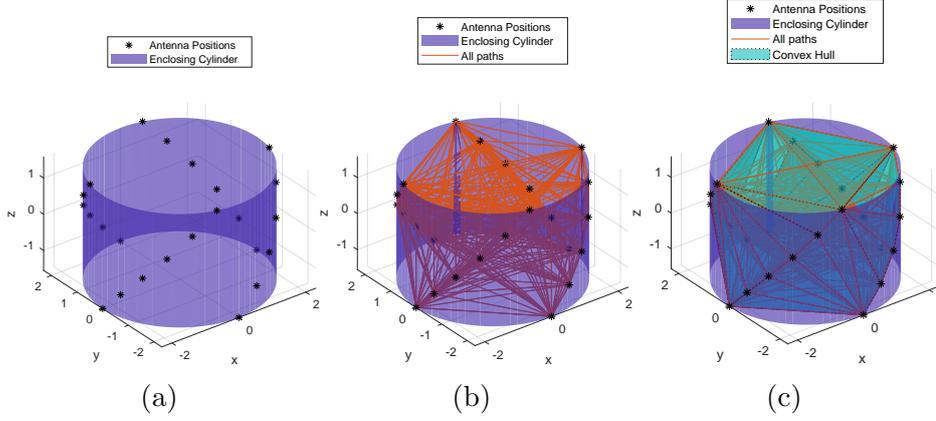


Figure 5.1: Evaluation domain for the error metric. The region shown here represents the cylindrical portion of the grain bin. The upper and lower cones are omitted for visual clarity. The set of paths fails to interrogate the outer regions of the cylinder.

of the bin's main cylinder, so the error metric is only evaluated within the convex hull of the antenna positions, which is shown in Figure 5.1c.

The visualizations presented in this chapter are also restricted to the convex hull of the antenna positions. The error metric, then, is expressed in the following equation, where H represents the convex hull of the antenna positions:

$$\mathcal{E}_c = \frac{\iiint_H (c(\vec{r}) - c_{\text{EXP}}(\vec{r}))^2 d\vec{r}}{\iiint_H (c_{\text{EXP}}(\vec{r}))^2 d\vec{r}} \quad (5.1)$$

An additional error, which will be referred to as the *best-fit error*, is calculated for the analysis of this experiment. The best-fit error, \mathcal{E}_{FIT} is calculated from the best-fit wave-speed, $c_{\text{FIT}}(\vec{r})$ as follows:

$$\mathcal{E}_{\text{FIT}} = \frac{\iiint_H (c_{\text{FIT}}(\vec{r}) - c_{\text{EXP}}(\vec{r}))^2 d\vec{r}}{\iiint_H (c_{\text{EXP}}(\vec{r}))^2 d\vec{r}} \quad (5.2)$$

The best-fit speed, $c_{\text{FIT}}(\vec{r})$ is the inverse of a corresponding best-fit slowness, $c_{\text{FIT}}^{-1}(\vec{r})$. The best-fit slowness, $c_{\text{FIT}}^{-1}(\vec{r})$ is the polynomial function which most closely matches the known slowness distribution, using a least-squares metric, at a set of test points. In this experiment, the set of test points is a uniformly-spaced set of points which lie within the region H . The metric

\mathcal{E}_{FIT} is included to give an indicator of the maximum accuracy that can be achieved by performing transmission tomography with a polynomial basis.

The number of basis functions for each degree is also reported, since the number of basis functions is the number of degrees of freedom that are available to the imaging algorithm. Consider Table 4.2, and the relationship between the mesh density and the number of basis functions. As the mesh density was increased, the number of basis functions increased, since each mesh element supports a unique basis function. In a polynomial basis, as the degree of the basis increases, the total number of basis functions increases. For example, consider the function f , represented with a three-dimensional polynomial basis of degree 1, using $\alpha_{i,j,k}$ as the coefficient of the $x^i y^j z^k$ monomial.

$$f(x, y, z) = \alpha_{0,0,0} + \alpha_{1,0,0}x + \alpha_{0,1,0}y + \alpha_{0,0,1}z \quad (5.3)$$

Thus, a polynomial basis of degree 1 requires 4 basis functions, and 4 basis coefficients. Expressing f with a basis of degree 2, we then have the following:

$$f(x, y, z) = \alpha_{0,0,0} + \alpha_{1,0,0}x + \alpha_{0,1,0}y + \alpha_{0,0,1}z + \alpha_{2,0,0}x^2 + \alpha_{0,2,0}y^2 + \alpha_{0,0,2}z^2 + \alpha_{0,1,1}yz + \alpha_{1,0,1}xz + \alpha_{0,1,1}xy \quad (5.4)$$

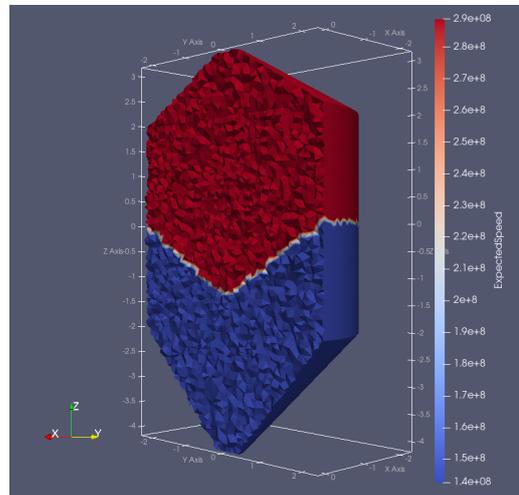
A polynomial basis of degree 2 requires 10 basis functions, and 10 basis coefficients. Suppose the polynomial basis is of degree D , and there are P paths which interrogate the target. In three spatial dimensions, there are $\frac{(D+1)(D+2)(D+3)}{6}$ total basis functions. This formula is equivalent to calculating the $(D+1)^{\text{th}}$ *tetrahedral* number, which is the three-dimensional extension of a *triangular number* [35]. The integrals of the polynomial basis functions along the given paths are performed analytically, so integrating a single basis function along a given path is a constant-time operation. The computational complexity of the basis function integral operation, then, is $\mathcal{O}(D^3P)$.

5.2.1 Data Processing Details

The scans for this experiment are the same as the scans used for the experiment in Chapter 4, so the scans are processed in the same way.

5.3 Results

The reconstructed wave-speed profiles are shown in Figure 5.3, and in Figure 5.4. Each of the visualizations shows the bin clipped along the y-z plane, viewed from the positive x direction, in order to show the reconstruction of the interior of the bin. The visualizations of Figure 5.3 are all shown with a colour scale ranging from $1 \times 10^8 \text{ m s}^{-1}$ to $3 \times 10^8 \text{ m s}^{-1}$, since that range covers the wave-speed values that are expected to be recovered. The visualizations in Figure 5.4 show clipped isosurfaces of the recovered wave-speed, in order to better show the three-dimensional shape of the reconstructions. The true grain distribution is also shown here, in Figure 5.2, for convenience.



(a) Expected Speed Profile

Figure 5.2: Synthetic grain distribution for grain-bin imaging experiment

The numerical error analysis for each basis is shown in Table 5.1. The reference errors for the all-grain and all-air case are included in this table, as

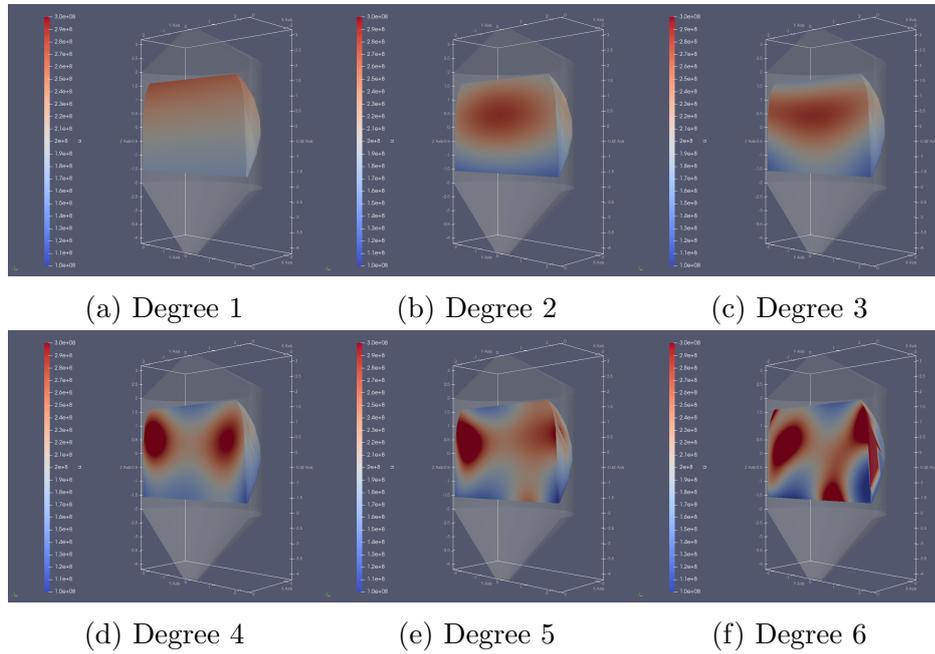


Figure 5.3: Polynomial basis reconstructions. Colour indicates wave-speed, from 1×10^8 to $3 \times 10^8 \text{m s}^{-1}$.

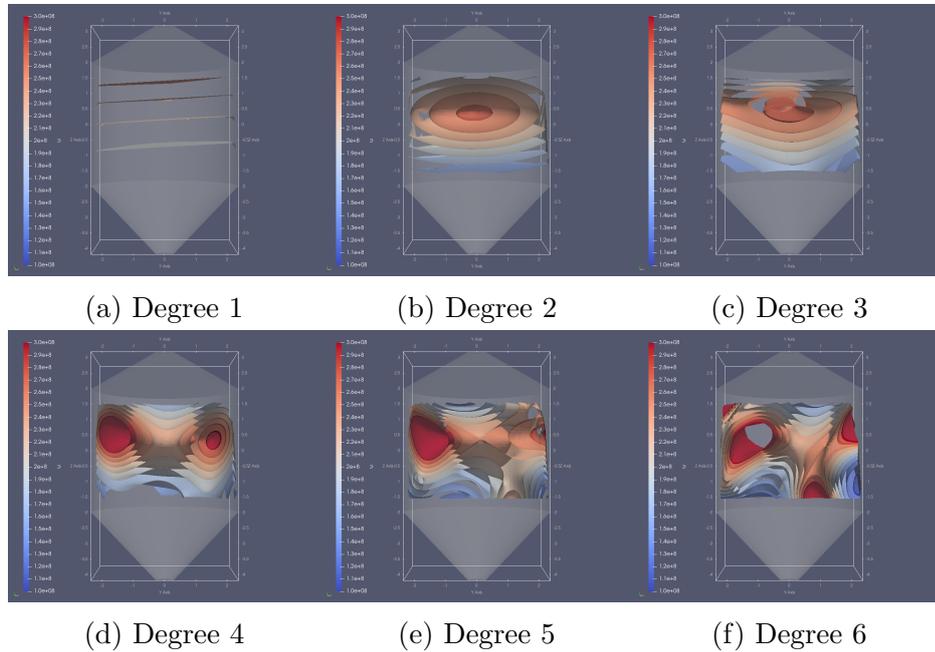


Figure 5.4: Isosurfaces of polynomial basis reconstructions. Colour indicates wave-speed, from 1×10^8 to $3 \times 10^8 \text{m s}^{-1}$.

are the best-fit errors, which are calculated according to Equation 5.2. These errors are plotted in Figure 5.5, using a logarithmic vertical axis.

Table 5.1: Accuracy of reconstruction on bases of varying degree

Basis Degree	Basis Functions	Speed Error	Best-Fit Error
1	4	0.0748	0.0633
2	10	0.0502	0.0306
3	20	0.0475	0.0253
4	35	0.0523	0.0266
5	56	0.0625	0.0189
6	84	0.132	0.0153
7	120	5850	0.0189

Reference	Speed Error
All Grain	0.249
All Air	0.115

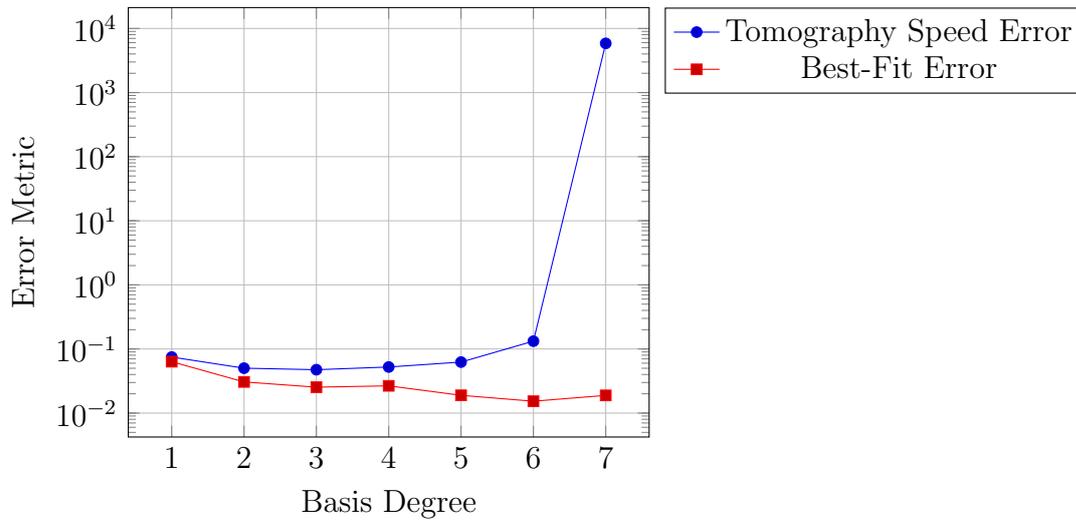


Figure 5.5: Error metrics with respect to basis degree

5.4 Discussion

Comparing the errors associated with polynomial-basis imaging in Table 5.1 with the errors associated with pulse-basis imaging in Table 4.2, it is clear

that using a polynomial basis for transmission tomography in a grain bin produces more accurate results than using a pulse basis for transmission tomography in a grain bin. The polynomial-basis images of the wave-speed in Figure 5.3 resemble the ground-truth grain distribution of Figure 5.2 more closely than the images of the pulse-basis wave-speed in Figure 4.3.

The most useful reconstruction of this set of data is the third-degree polynomial-basis reconstruction. This reconstruction achieves the lowest speed error, as can be seen in Table 5.1. The third-degree reconstruction is visually the best reconstruction, as well, as can be seen in Figure 5.3c. The negative-angled cone shape is clearly visible in the reconstructed wave-speed image.

The seventh-degree reconstruction is clearly an outlier in this experiment. For this set of data, attempting to perform transmission tomography with any degree higher than 7 results in a highly oscillatory wave-speed reconstruction. The challenge of choosing the optimal degree will be discussed in Chapter 6 and Chapter 8.

The results of the experiment presented in this chapter show that performing transmission tomography in a grain bin is possible when the material properties are represented with a polynomial basis. In contrast, the results shown in Chapter 4 show that transmission tomography in a grain bin is not possible with a pulse basis.

Chapter 6

Monitoring Grain-Bin Fill Volume

6.1 Introduction

The results presented in Chapter 5 show that transmission tomography can use electromagnetic measurements to produce an accurate three-dimensional map of the physical properties of a heap of grain, stored in a metal bin.

A spatial map of the wave speed of the grain heap is useful, however, a farmer may not be interested in this particular property. As was discussed in Chapter 4, a farmer is interested in the profit that can be made from the agricultural product that is stored in their bin. To that end, the farmer would want to know how much product they have.

This chapter presents an experiment which attempts to use the result of transmission tomography to monitor the volume of grain in a bin as it changes over time.

6.2 Procedure

6.2.1 Data Generation

For this experiment, transmission tomography was performed on several sets of synthetic data. Each set of data was generated using a different amount of grain in the simulated bin. The amount of grain in the simulated bin was controlled by setting the *grain height* parameter. The grain height is defined to be the distance from the bottom of the cylindrical portion of the grain bin to the vertical level of grain measured at the wall of the cylindrical portion of the grain bin. This parameter is shown schematically in Figure 6.1.

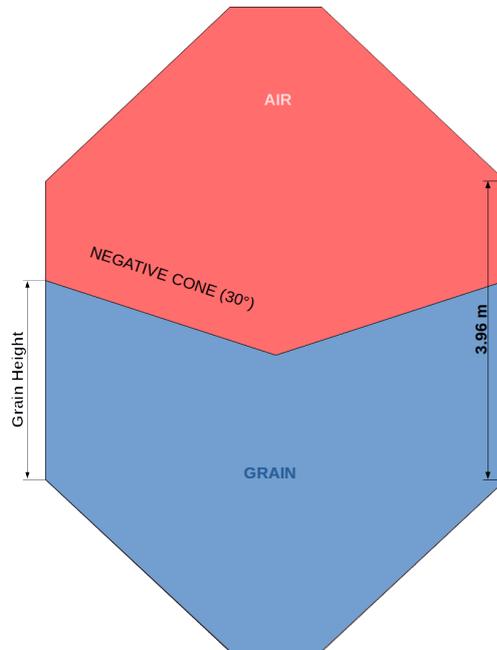


Figure 6.1: Schematic of grain distribution for volume-tracking experiment

Synthetic measurements were taken for grain heights ranging from 0 m to 3.96 m, in a bin whose main cylinder has a height of 3.96 m. For each grain height, the surface of the grain was shaped in a negative cone, with a declination angle of 30 deg, which matches the grain surface that was modelled in Chapter 4 and Chapter 5.

Similar to Chapter 4 and Chapter 5, the software meep was used to

generate the time-domain electromagnetic data.

6.2.2 Volume Extraction

Three-dimensional transmission tomography was performed on each set of data. Based on the results of Chapter 5, imaging was performed using a polynomial basis of degree 3. For each imaging result, a grain fill volume was calculated as follows:

1. The convex hull of the antenna positions was partitioned into a set of small, disjoint elements
2. The reconstructed speed profile was sampled at each element's centroid
3. Elements with a reconstructed speed of $2 \times 10^8 \text{ m s}^{-1}$ or lower were labelled as containing grain
4. The volume of the elements labelled with grain were summed

A similar process was used to calculate an *expected* fill volume. In order to calculate the expected fill volume, except elements were labelled based on the *position* of their centroid relative to the known surface of the grain.

6.2.3 Data Processing Details

The scans for this experiment are the same as the scans used for the experiment in Chapter 4, so the scans are processed in the same way.

6.3 Results

The calculated grain volume and the expected grain volume are plotted together in Figure 6.2. The volume extraction process is shown in Figure 6.3, for three different grain heights.

The left-most column of Figure 6.3 shows the expected wave-speed distribution, based on the known shape of the grain. The middle column shows

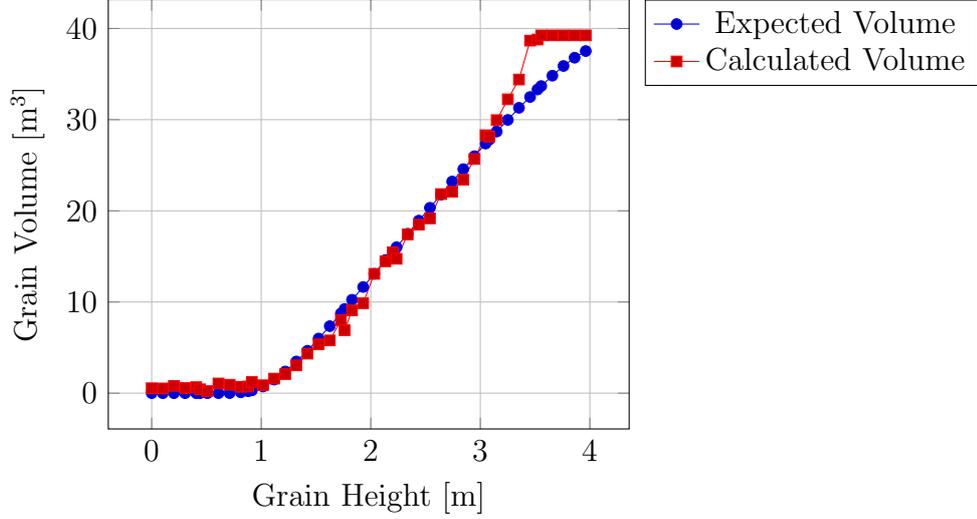


Figure 6.2: Extracted volume plotted against grain height

the wave-speed distribution that is calculated by the transmission tomography algorithm. The right-most column shows the speed distribution that is attained by partitioning the result of the middle column with a speed threshold of $2 \times 10^8 \text{ m s}^{-1}$. The top row corresponds to a low grain height, while the bottom row corresponds to a high grain height.

6.4 Discussion

The thresholding technique successfully tracks the volume of grain stored in the bin. Looking at Figure 6.2, the thresholding technique is most accurate for low to medium grain heights. The technique only appears to degrade for high grain heights, where the thresholding technique over-estimates the amount of grain stored in the bin.

The cause of the overshoot, which is observed on the right side of the graph in Figure 6.2, is the smoothing effect that is inherent in a low-order polynomial basis. Looking at the third row of Figure 6.3, the shape of the thresholded speed is not as sharply defined as the shape of the expected speed distribution, and this smoothing results in an apparent excess of grain.

Looking at the first row of Figure 6.3, it is clear that the thresholding

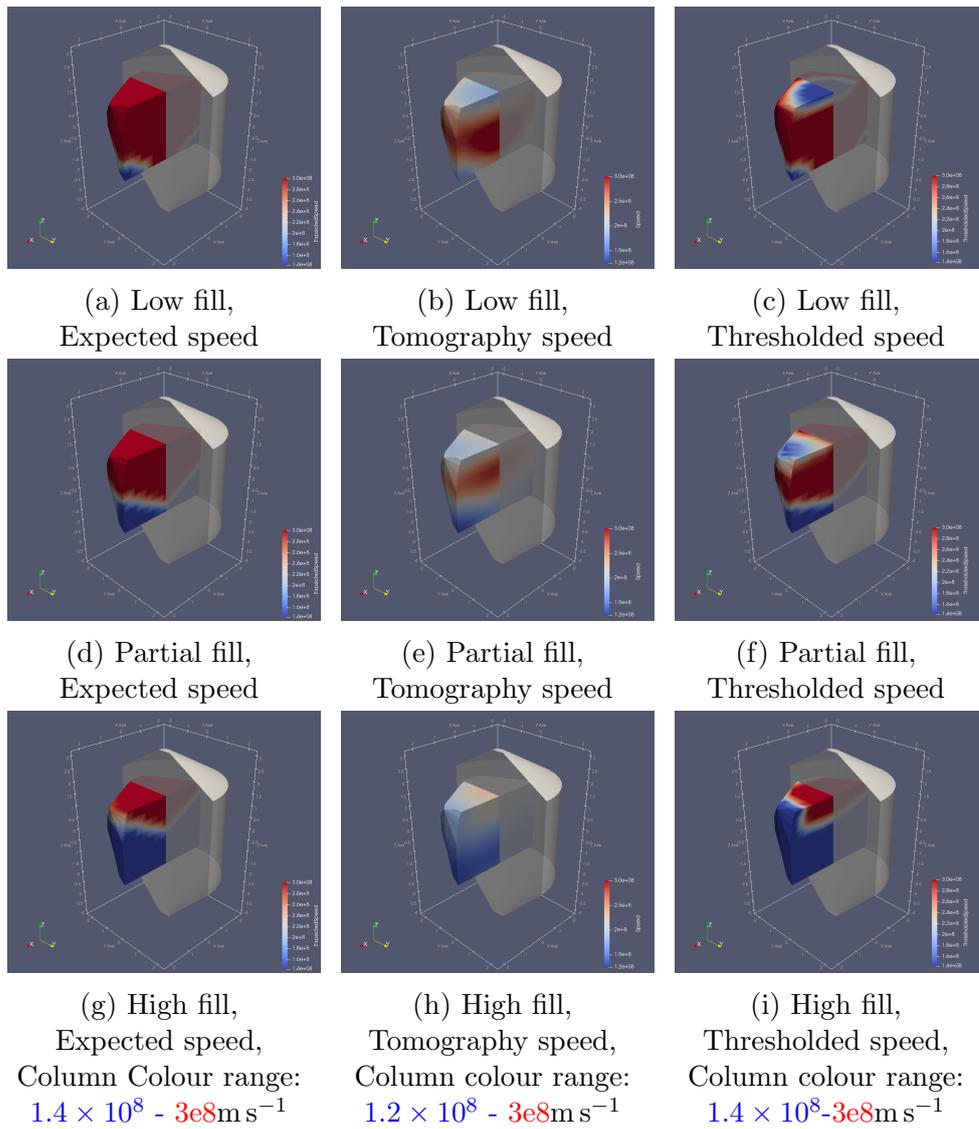


Figure 6.3: Grain volume calculation, shown for three different fill levels

technique introduces an artefact in the upper region of the grain bin. The thresholding technique interprets this artefact as a region of grain which is floating near the top of the grain bin, which is obviously unrealistic. Based on the data in Figure 6.2, it appears that this artefact does not have a significant negative impact on the calculation of the grain volume. This is because the volume of the artefact is actually quite small, despite its prominence in Figure 6.3. The artefact is so visible because it occupies a flat region at the *boundary* of the region being shown.

The polynomial basis degree was set to 3 because optimal results of the synthetic experiment shown in Chapter 5 were achieved at that degree. This insight may not be available in a real imaging scenario, when an actual heap of grain is being imaged. Having a robust method for choosing the optimal basis degree is necessary in order to apply the transmission tomography algorithm to arbitrary real-world data, however, developing such a method is beyond the scope of this thesis. A candidate method would be to perform an imaging study on synthetic data in order to choose the basis degree for imaging real measured data, which mimics the process used in this experiment, however, a thorough analysis of other options has not been performed.

The results of this chapter provide evidence that the transmission tomography algorithm can create diagnostically useful images from electromagnetic measurements from within a grain bin.

Chapter 7

Biological Imaging in Three Dimensions

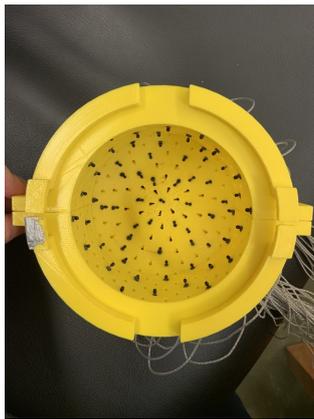
7.1 Introduction

It was shown in Chapter 3 that transmission tomography using pulse-basis functions on an unstructured mesh was capable of imaging low-contrast targets in two dimensions. Such an imaging formulation was unsuitable for the three-dimensional environment of grain-bin imaging, as was demonstrated in Chapter 4. Increasing the number and density of scans in the grain-bin is infeasible, and Chapter 5 presented a transmission tomography tool which uses a pulse basis to represent the material properties. Expressing the material properties in a global basis led to smoother reconstructions, which are diagnostically useful, and more accurate than the reconstructions which used a pulse basis. Smooth, spatially-coherent images are desirable in the scenario of grain-bin imaging because it is known that the material properties of grain in a bin are geometrically simple.

This chapter demonstrates transmission tomography in a small-scale, biological tissue imaging context. Specifically, the techniques described in this chapter make use of an ultrasound imaging chamber similar to the chamber described in Chapter 3. For this experiment, the chamber was configured as a three-dimensional imaging system. The 64 piezoelectric elements were

mounted in a shell consisting of a cylindrical portion capped on the bottom by a dome. The chamber is shown in Figure 7.1. The imaging chamber consists of a yellow plastic enclosure which houses the piezoelectric transducers. The transducers inserted through slots in the side of the enclosure, and held in place with rubber plugs, as shown in Figure 7.1c. The shape of the chamber was chosen to facilitate the imaging of breast tissue, which has been a research topic in the EIL in recent years.

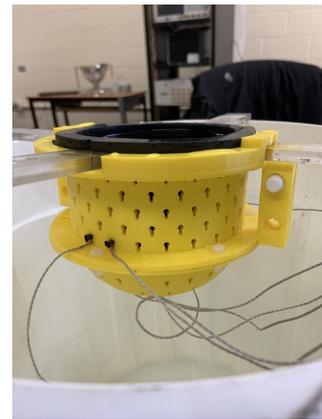
For this experiment, a water-tight cup, which can be seen in Figure 7.1c, sits inside the enclosure and rests on the piezoelectric transducers. The water-tight cup makes it easy to quickly create tissue-mimicking phantoms. This experiment makes use of phantoms, instead of actual tissue samples, so that the accuracy of the tomographic images can be accurately assessed.



(a) Enclosure populated with transducers



(b) Partially-assembled enclosure



(c) Enclosure with water-tight cup suspended in dry bucket

Figure 7.1: Images of the three-dimensional ultrasound imaging chamber

7.2 Procedure

A simple phantom was constructed by filling the cup with olive oil, and suspending two spheres of gelatin in the olive oil. The sound-speed profile of this phantom is shown in Figure 7.2, from several angles, with the transducers' positions shown as white dots. The olive oil and gelatin material was chosen

to approximate tissue found in a breast. The speed of sound was measured in a sample of each of the gelatin spheres, as well as in the olive oil, and in the water in which the system was submerged. Those measurements are shown in Table 7.1. Note that the two gelatin spheres were formulated to have different speeds. The sphere which sits higher in the shell has a higher sound speed than the sphere which sits lower in the shell.

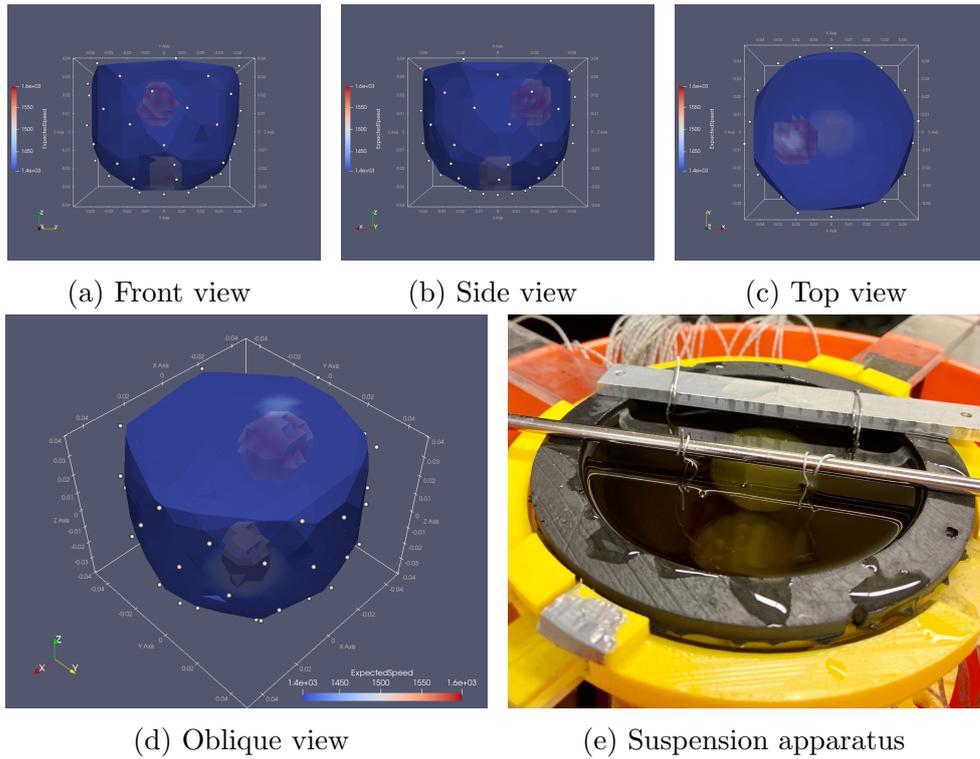


Figure 7.2: Known speed profile with transducer locations. Colour indicates wave-speed, ranging from 1400 m s^{-1} to 1600 m s^{-1} .

The attenuation of the materials could not be accurately measured in the lab, due to lack of precision power measurement tool, so the attenuation images will not be analyzed.

The imaging chamber was populated with 64 transducers, as is shown in Figure 7.1a. Similar to Chapter 3, the transducers were excited in sequence, while the non-excited transducers acted as receivers, producing 4032 time-domain scans. The scans were then imaged using the transmission tomography algorithm. Based on the findings of Chapter 5, a polynomial basis

was chosen to represent the material properties.

Table 7.1: Expected material properties for 3D ultrasound experiment

Material	Speed [m s^{-1}]
Olive Oil	1405
Gelatin (upper sphere)	1598
Gelatin (lower sphere)	1516
Water	1447

$$\mathcal{E}_c = \frac{\iiint_H (c(\vec{r}) - c_{\text{EXP}}(\vec{r}))^2 d\vec{r}}{\iiint_H (c_{\text{EXP}}(\vec{r}))^2 d\vec{r}} \quad (7.1)$$

The accuracy of the recovered sound-speed map was determined quantitatively using the error metric from Equation 7.1. This error metric is the same as the metric used in Chapter 5, and is shown again in this chapter for the reader's convenience. The function c_{EXP} was generated from accurate measurements of the position and size of the gelatin spheres, together with the values from Table 7.1.

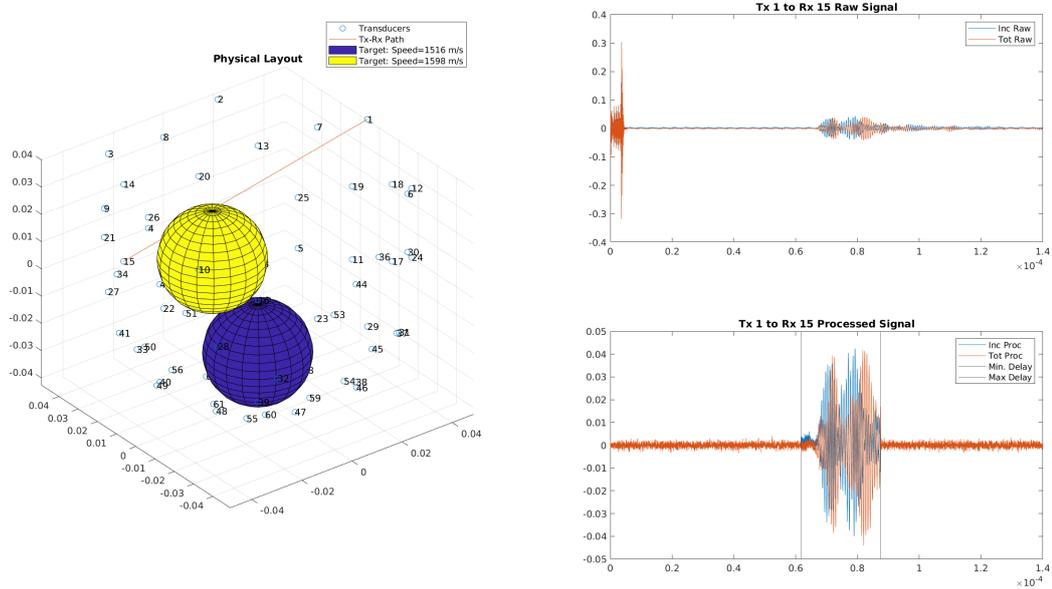


Figure 7.3: Ultrasound signal pre-processing. The transmitter-receiver path is shown on the left, and the corresponding truncated signal is shown on the right

Similar to the experiment in Chapter 5, the transmission tomography algorithm was run on polynomial bases of increasing degree, to examine the impact of increasing the number of degrees of freedom.

7.2.1 Data Processing Details

The rigid plastic enclosure presents an opportunity for the generation of reflections, which can impair TOA detection. In order to minimize the impact of reflections, the scans were processed prior to imaging as follows:

1. The length of the transmitter-receiver path was calculated
2. A maximum delay was determined by combining the contributions of the following effects:
 - Pulse propagation in the slowest expected medium
 - Pulse propagation in the switch and along cables
 - Calculating a time-window to suit the chosen TOA detection algorithm
3. The portion of the signal following the maximum delay was replaced with low-energy Gaussian noise
4. A minimum delay was determined by combining the contributions of the following effects:
 - Pulse propagation in the fastest expected medium
 - Calculating a time-window to suit the chosen TOA detection algorithm
5. The portion of the signal preceding the minimum delay was also replaced with low-energy Gaussian noise

A graphic representation of the data processing procedure is shown in Figure 7.3. The scans' TOA were then extracted with the MER method. Gaussian noise was chosen as the filler in order to prevent any divide-by-zero

errors in the TOA process. Both the MER and variance score methods rely on evaluating a ratio of terms which are derived from the scan signal. Using zeroes as the filler signal would introduce the opportunity for the ratio's denominator to be zero.

7.3 Results

Images of the recovered sound-speed are shown in Figure 7.4 and Figure 7.5, and the corresponding values of \mathcal{E}_c are documented in Table 7.2. Results are presented for polynomial basis degrees ranging from 1 to 9, as the imaging results for higher polynomial bases degrade quickly.

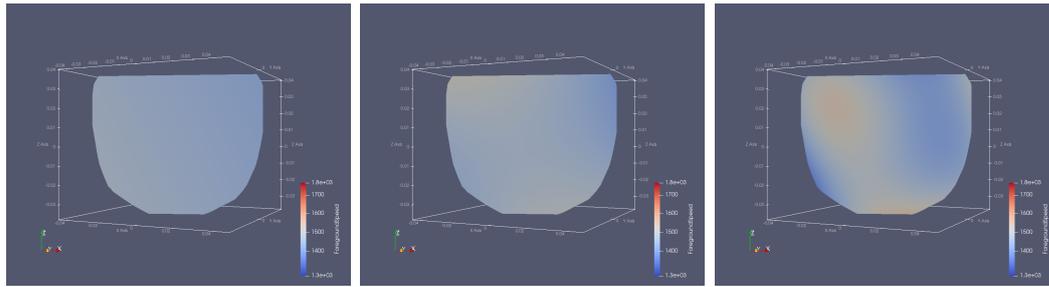
The images in Figure 7.4 show a clipped view of the three-dimensional reconstructed sound-speed. The plane of the clip is chosen to cut through the two gelatin spheres, to show the shape of the reconstructed sound-speed. These results are further processed by applying a contour filter, as can be seen in Figure 7.5. The contour filter highlights the three-dimensional features of the reconstructions.

The errors in Table 7.2 are reported along with the number of basis functions in the given basis, and the best-fit speed error. The number of basis functions indicate the number of degrees of freedom available to the imaging algorithm. Recall from Chapter 5 that the best-fit error represents speed error that is achieved when the slowness polynomial is the polynomial of best-fit for the *known* slowness distribution. The best-fit error gives an indicator of the best possible error that could be achieved with a polynomial basis.

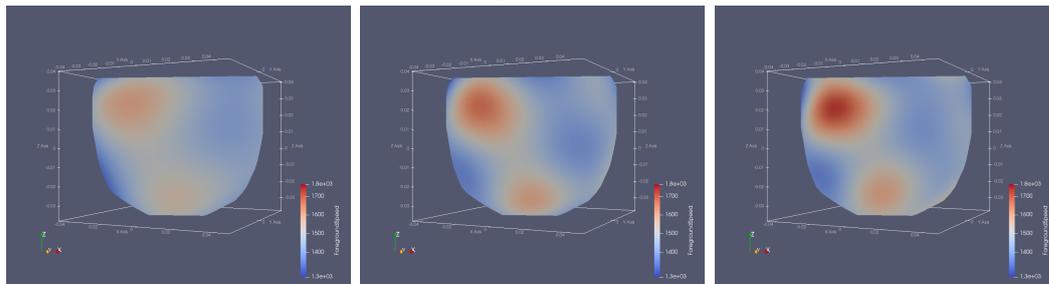
7.4 Discussion

It is clear in Figure 7.4 and Figure 7.5 that the transmission tomography algorithm recovers the gelatin targets.

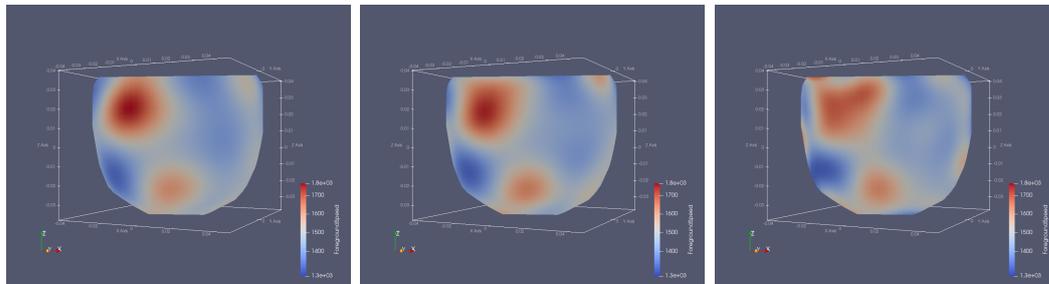
Among all of the sound-speed reconstructions, the highest accuracy is attained with a polynomial basis of degree 5, using 56 basis functions to



(a) Degrees 1 to 3

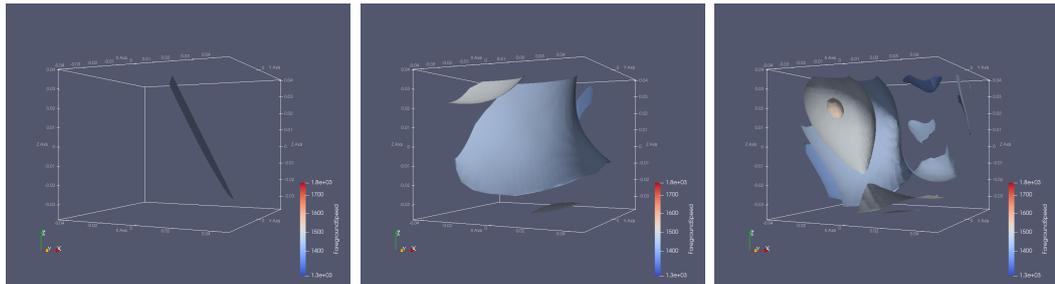


(b) Degrees 4 to 6

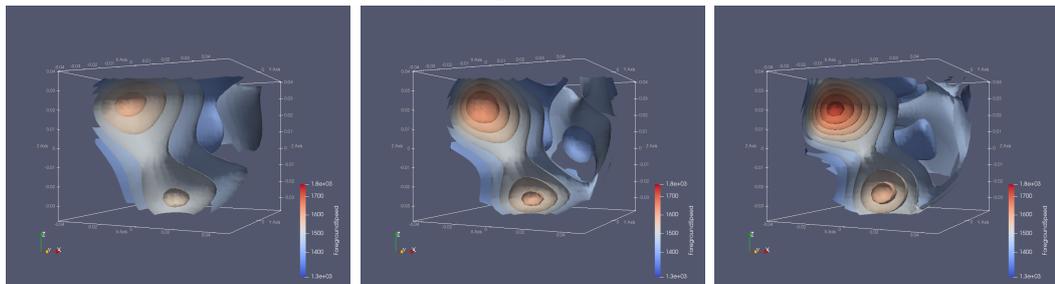


(c) Degrees 7 to 9

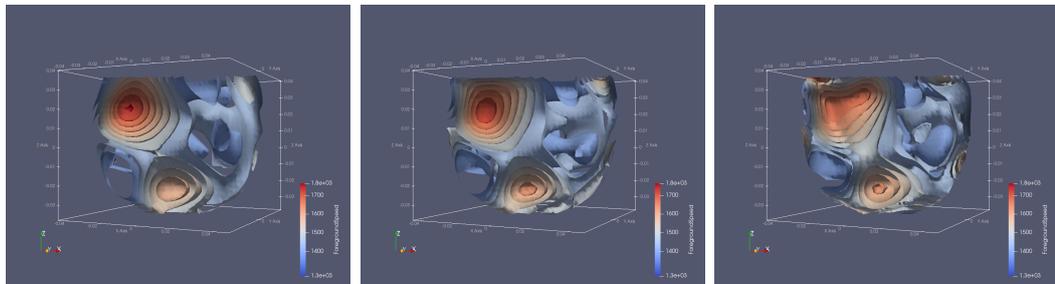
Figure 7.4: Face-view of sound-speed reconstructions. Colour indicates wave-speed, ranging from 1300 m s^{-1} to 1800 m s^{-1} .



(a) Degrees 1 to 3



(b) Degrees 4 to 6



(c) Degrees 7 to 9

Figure 7.5: Contoured view of sound-speed reconstructions. Colour indicates wave-speed, ranging from 1300 m s^{-1} to 1800 m s^{-1} .

Table 7.2: Accuracy of ultrasound reconstruction on bases of varying degree

Basis Degree	Basis Functions	Speed Error	Best-Fit Error
1	4	0.001 55	0.001 07
2	10	0.001 49	0.000 949
3	20	0.001 32	0.000 763
4	35	0.001 35	0.000 619
5	56	0.001 28	0.000 447
6	84	0.001 31	0.000 342
7	120	0.001 49	0.000 283
8	165	0.001 61	0.000 261
9	220	0.002 14	0.000 254

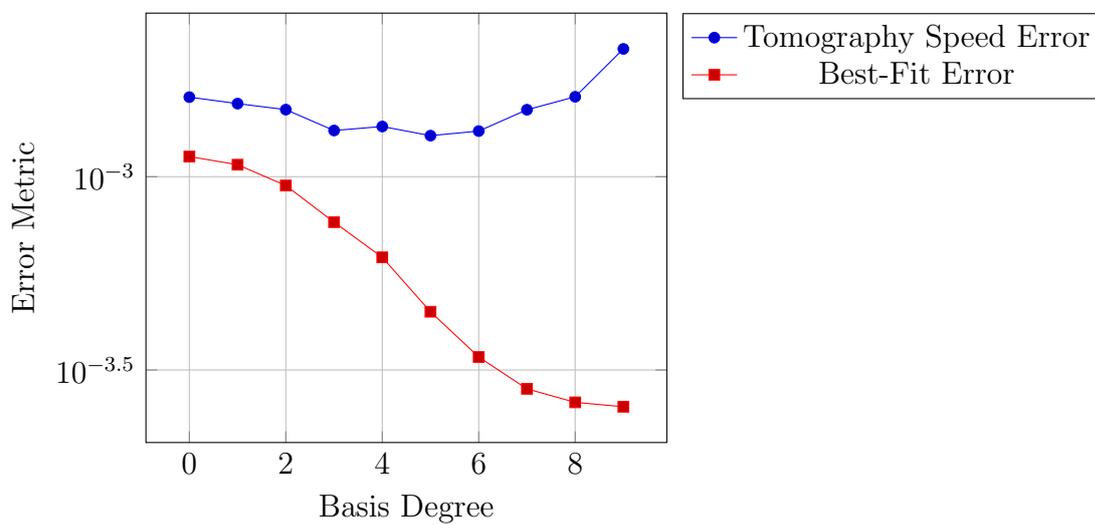


Figure 7.6: Error metrics with respect to basis degree in three-dimensional ultrasound experiment

represent the slowness function. The recovered speed of the upper sphere is clearly higher than the recovered speed of the lower sphere. At degree 5, the recovered speed of the upper sphere is roughly 1650 m s^{-1} , and the recovered speed of the lower sphere is roughly 1600 m s^{-1} . These recovered speeds are higher than the expected speeds of 1598 m s^{-1} and 1516 m s^{-1} , respectively, due to the inherent error of using a polynomial to represent a step discontinuity.

The errors in Table 7.2 show that the tomography algorithm is robust over a range of basis function orders. The reconstructions' speed error is smooth up to the degree-8 reconstruction, after which the speed error quickly rises.

The results of this experiment show that transmission tomography, using a polynomial basis, is capable of identifying objects using quantitative and qualitative analysis of the spatial sound-speed distribution. The stable, low-resolution results of this experiment suggest that transmission tomography, using a polynomial basis, could be a useful aid to non-linear inversion techniques. Non-linear inversion techniques benefit from the inclusion of *prior information* about the imaging target [8]. Without sufficient prior information, these techniques often get stuck in local minima within the optimization search space. Using the output from transmission tomography as an initial guess for a non-linear inversion algorithm could direct the non-linear algorithm away from local minima, toward an accurate result.

Chapter 8

Discussion and Conclusions

The work undertaken in this thesis sought to produce an imaging tool which has the following features:

- Minimal reliance on prior information
- Robustness to modelling error
- Fast execution time

These features were realized by implementing transmission tomography, particularly using a polynomial basis to represent the material properties of interest.

This thesis has presented a robust transmission tomography tool which can be used to generate low-resolution, quantitative property maps of imaging targets, given time-domain scan data. The only prior information that is required by the transmission tomography algorithm is the location of the transducers which produce and measure the interrogating waves. Additional information about the maximum and minimum values of the expected material properties *can* be exploited, but that information is not strictly necessary. The algorithm is fast due to the simple wave propagation model, which obviates the need for detailed and expensive modelling of coupled differential equations. The simple wave propagation model also makes transmission tomography robust to modelling errors, as transmission tomography does not

require sub-wavelength knowledge of transducer locations or perfect knowledge of the geometry of the imaging domain.

The algorithm developed for this thesis is agnostic of the particular wave propagation phenomena being studied, however, special attention was paid to the applications of grain-bin imaging and medical imaging. This special attention was warranted by recent work that has been undertaken in the EIL.

A general framework for transmission tomography was presented in Chapter 2. The transmission tomography framework relies on a simplified wave propagation model which ignores diffraction. It is assumed that a wave front propagates directly away from its source, at a speed which depends only on the material properties at the location of the wave front. As a consequence of this assumption, it can be assumed that the time required for a pulse to propagate from point A to point B in space depends *only* on the properties of the medium on the straight-line path which connects A and B. Furthermore, it is assumed that the the portion of the wave front's attenuation which is due to material losses depends only on the material parameters of the medium on the straight-line path which connects A to B.

Spatial and temporal discretization schemes were then introduced. These discretization schemes allow the simplified wave propagation model to be expressed numerically, in the form of a system of linear algebraic equations. The system of linear algebraic equations is then solved to produce an estimate of the material properties of the imaging target.

Imaging tests were then performed on simple targets, in Chapter 3, using a two-dimensional form of the transmission tomography framework. Acoustic data were used because of availability and ease of rapid testing. These simple tests were performed with pulse basis functions, which were chosen for their ease of implementation. The results of the two-dimensional test provided evidence that the development of a three-dimensional transmission tomography tool could work for grain imaging.

In Chapter 4, the two-dimensional transmission tomography tool was augmented to perform three-dimensional transmission tomography, using pulse basis functions. Three-dimensional tomography was performed on synthetic data which were generated by a full-wave, time-domain numerical simulation.

The results of the three-dimensional pulse basis experiment showed that the imaging tool was incapable of accurately imaging a heap of grain stored in a metal bin.

The results of Chapter 4 suggested that the transmission tomography tool must adapt in order to successfully image a heap of grain stored in a bin. To that end, the transmission tomography tool was further augmented to represent the material properties in a polynomial basis, instead of a pulse basis. A pulse basis permits high-resolution imaging, due to the mutually disjoint support of its basis functions, however, a sufficient amount of data must be supplied to the algorithm in order to achieve a high-resolution image. In the target application of grain-bin imaging, sufficient data are *not* available, due to hardware cost constraints. A polynomial basis features functions whose support is infinite. As a consequence, the amount of interaction between basis functions and interrogation paths is greater with a polynomial basis than with a pulse basis. With polynomial basis functions, the transmission tomography tool's accuracy in grain-bin imaging was improved to a degree that the transmission tomography tool is diagnostically useful for grain-bin imaging, as was demonstrated in Chapter 5.

The application of grain-bin imaging was explored further in Chapter 6. Transmission tomography was used to track the volume of grain stored in a bin, using the recovered speed distribution to partition the bin in regions of grain and air. The resulting estimate of the volume of grain in the bin was accurate for several different fill levels, ranging from a nearly empty bin to a full bin.

The three-dimensional transmission tomography tool was then used in an acoustic imaging context, as documented in Chapter 7. Similar to Chapter 3, some imaging tests were performed on geometrically simple targets. Qualitatively and quantitatively, it was shown that the transmission tomography tool was capable of identifying the rough location and properties of the targets.

The main goal of this thesis was to develop a fast and simple imaging tool which could be used primarily for grain-bin imaging. Experiments were performed which demonstrated that transmission tomography is a useful tool

for fast, low-resolution imaging in grain bins. Some additional, exploratory experiments were performed in ultrasound imaging, which demonstrated that transmission tomography could provide valuable information in a field where full-wave inversion is still struggling to perform. The tool developed in conjunction with this thesis will be used in future grain-bin imaging work, likely in boundary detection.

8.1 Future Work

This thesis represents one particular journey down a branching path of research ideas. Along the way, several possible research paths were identified which could augment the algorithm presented in this thesis. Some of those paths are presented here.

8.1.1 Basis Degree Selection

In its current state, the transmission tomography algorithm lacks a robust method for choosing the optimal degree for the material properties' basis polynomials. The experiments in Chapter 5 and Chapter 7 presented the results of imaging using several different degrees, and selected the best result. Selecting the best result, however, was only enabled by the existence of ground-truth data.

In Chapter 6, the basis degree was selected based on the synthetic study in Chapter 5. Both experiments consisted of similar imaging targets, in the same bin, so the insight from the first experiment was applied to the second experiment.

In a real-world imaging scenario, there is no ground-truth which can be used to evaluate imaging results. Therefore, it will be necessary to have a robust method for choosing the basis degree to use for imaging. It is possible that the method used in Chapter 6 would suffice, however, the inconsistency between synthetic data and real data may derail that method. Developing and testing some degree-picker methods could make the transmission tomography algorithm more broadly applicable than it currently is.

8.1.2 Diffraction-Aware Solver

The physical model of wave propagation that is employed in this thesis relies on a key assumption. It is assumed that the waves propagate without diffraction, that is, the pulse which propagates from a transmitter to a receiver is affected only by the material properties on the straight-line path between the transmitter and the receiver. That fundamental assumption is not true, however, it permits the development of a fast, simple imaging tool, and the imaging results align with reality well enough to justify the assumption.

Diffraction can be accounted for via an iterative imaging technique [36]. This technique augments the straight-path transmission tomography that was employed in this thesis by using the result of transmission tomography to modify the straight-path pulse propagation assumption. Once a wave-speed image is generated, that wave-speed image is used to find a new set of interrogation paths which minimize time of flight. Then, the new interrogation paths are used to update the wave-speed image. These two steps are repeated iteratively until a stopping condition is met.

This iterative technique was not explored in this thesis because the calculation of optimal paths, as well as the solution of many matrix systems, could slow the imaging algorithm down. The potential to improve image resolution, however, is tantalizing, and it may be fruitful to add this feature the transmission tomography algorithm.

8.1.3 Iterative Polynomial Basis Refinement

In this thesis, the basis coefficients of the material property functions are recovered by solving a matrix equation. The basis coefficients are all recovered at the same time, and their values are chosen to satisfy the matrix equation by some measure.

It was observed in Chapter 5 and Chapter 7 that the reconstructions tend to degrade for basis degrees above some threshold. This degradation may be due to an excess of degrees of freedom, that is, the imaging algorithm is allowed to tweak too many variables at once in order to satisfy the tomography matrix equations. It may be advantageous to calculate the ba-

sis coefficients in sequence, and treat the already-calculated coefficients as known values. For example, imaging would first be performed with a basis of degree zero. Then, the zero-order terms in the basis would be held constant while imaging is performed with additional linear terms in the basis. Then, the zero-order and linear terms would be held constant while imaging is performed with additional quadratic terms. Applying this process iteratively may allow imaging to succeed with bases of higher degree.

8.1.4 On-Line Imaging

All of the imaging performed in this thesis was performed off-line. In general, the off-line imaging workflow was the following:

1. For each transmitter, in turn
 - Excite the transmitter, and
 - Measure a received signal at each receiver
2. Store all of the scan data in a computer file
3. Read the data from the file
4. Perform transmission tomography
5. Store tomography results in a computer file

For grain-bin imaging, the imaging target is a heap of grain, which could be scanned continuously to create a never-ending stream of imaging data. The basis coefficients could be updated with each new scan, to present a living image of the target.

8.1.5 Regularization Study of Pulse-Basis Imaging

It was shown in Chapter 4 that representing material properties with a pulse basis leads to unacceptable imaging results. This was remedied in Chapter 5 by using a different set of basis functions. Another research would be to seek a method of regularizing the transmission tomography matrix system

which makes pulse-basis imaging acceptable. The only regularization which was applied to pulse-basis imaging in this thesis was the CGLS algorithm. Several other regularization techniques could be explored such as:

- Tikhonov regularization
- Penalizing property differences between neighbouring mesh elements
- Using higher-order representation of material properties within mesh elements
- Employing iterative matrix solvers other than CGLS

Perhaps one of these techniques could permit high-resolution imaging, while still maintaining the other attractive features of transmission tomography.

Chapter 9

Bibliography

- [1] Partitions of an integer: List all partitions of an integer. <https://www.mathworks.com/matlabcentral/fileexchange/12009-partitions-of-an-integer>. Accessed: 2021-08-23.
- [2] writevtk: Exports triangulations and triangulated surfaces to vtk file format. <https://www.mathworks.com/matlabcentral/fileexchange/25784-writevtk>. Accessed: 2021-08-23.
- [3] Amy Henderson. *ParaView Guide, A Parallel Visualization Application*. Kitware Inc., 2007.
- [4] Mohammad Asefi, Ian Jeffrey, Joe LoVetri, Colin Gilmore, Paul Card, and Jitendra Paliwal. Grain bin monitoring via electromagnetic imaging. *Computers and Electronics in Agriculture*, 119:133–141, 2015.
- [5] Electromagnetic imaging laboratory overview. <https://home.cc.umanitoba.ca/~lovetrij/EMILab/research.html>. Accessed: 2021-02-18.
- [6] Amer Zakaria, Anastasia Baran, and Joe LoVetri. Estimation and use of prior information in fem-csi for biomedical microwave tomography. *IEEE Antennas and Wireless Propagation Letters*, 11:1606–1609, 2012.

- [7] Amer Zakaria, Ian Jeffrey, and Joe LoVetri. Full-vectorial parallel finite-element contrast source inversion method. *Progress In Electromagnetics Research*, 142:463–483, 2013.
- [8] Anastasia Baran. *Microwave Breast Imaging Techniques in Two and Three Dimensions*. Ph. d. thesis, University of Manitoba, 2016.
- [9] Colin Gilmore, Puyan Mojabi, Amer Zakaria, Majid Ostadrahimi, Cameron Kaye, Sima Noghianian, Lotfollah Shafai, Stephen Pistorius, and Joe LoVetri. A wideband microwave tomography system with a novel frequency selection procedure. *IEEE Transactions on Biomedical Engineering*, 57(4):894–904, 2009.
- [10] M. Ostadrahimi, P. Mojabi, C. Gilmore, A. Zakaria, S. Noghianian, S. Pistorius, and J. LoVetri. Analysis of incident field modeling and incident/scattered field calibration techniques in microwave tomography. *IEEE Antennas and Wireless Propagation Letters*, 10:900–903, 2011.
- [11] Peter M Van Den Berg and Ralph E Kleinman. A contrast source inversion method. *Inverse problems*, 13(6):1607, 1997.
- [12] Colin Gilmore, Mohammad Asefi, Jitendra Paliwal, and Joe LoVetri. Industrial scale electromagnetic grain bin monitoring. *Computers and Electronics in Agriculture*, 136:210–220, 2017.
- [13] Nicholas Geddert. An electromagnetic hybridizable discontinuous galerkin method forward solver with high-order geometry for inverse problems. M. sc. thesis, University of Manitoba (Canada), 2020.
- [14] G. Faucher. Calibration of an ultrasound tomography system for medical imaging with 2d contrast-source inversion. M. sc. thesis, University of Manitoba, 2013.
- [15] Pedram Mojabi. Ultrasound tomography: An inverse scattering approach. M. sc. thesis, University of Manitoba (Canada), 2014.

- [16] Tomasz M. Grzegorzcyk, Paul M. Meaney, Soon Ik Jeon, Shireen D. Geimer, and Keith D. Paulsen. Importance of phase unwrapping for the reconstruction of microwave tomographic images. *Biomed. Opt. Express*, 2(2):315–330, Feb 2011.
- [17] Ergün Simsek, Jianguo Liu, and Qing Huo Liu. A spectral integral method and hybrid sim/fem for layered media. *IEEE transactions on microwave theory and techniques*, 54(11):3878–3884, 2006.
- [18] Nebojsa Duric, Peter Littrup, Lou Poulo, Alex Babkin, Roman Pevzner, Earle Holsapple, Olsi Rama, and Carri Glide. Detection of breast cancer with ultrasound tomography: First results with the computed ultrasound risk evaluation (cure) prototype. *Medical physics*, 34(2):773–785, 2007.
- [19] Jonathan B Ajo-Franklin, Burke J Minsley, and Thomas M Daley. Applying compactness constraints to differential travelttime tomography. *Geophysics*, 72(4):R67–R75, 2007.
- [20] James Clerk Maxwell. Viii. a dynamical theory of the electromagnetic field. *Philosophical transactions of the Royal Society of London*, (155):459–512, 1865.
- [21] Allan D Pierce. *Acoustics: an introduction to its physical principles and applications*. Springer, 2019.
- [22] Jerry L Prince and Jonathan M Links. *Medical imaging signals and systems*. Pearson Prentice Hall Upper Saddle River, 2006.
- [23] Edwin KP Chong and Stanislaw H Zak. *An introduction to optimization*. John Wiley & Sons, 2004.
- [24] Richard Barrett, Michael Berry, Tony F Chan, James Demmel, June Donato, Jack Dongarra, Victor Eijkhout, Roldan Pozo, Charles Romine, and Henk Van der Vorst. *Templates for the solution of linear systems: building blocks for iterative methods*. SIAM, 1994.

- [25] M Pérez-Liva, JL Herraiz, N González-Salido, L Medina-Valdés, J Camacho, C Fritch, and JM Udías. Ultrasound computed tomography for quantitative breast imaging. In *2016 Global Medical Engineering Physics Exchanges/Pan American Health Care Exchanges (GMEPE/PAHCE)*, pages 1–6. IEEE, 2016.
- [26] Pedram Mojabi and Joe LoVetri. Experimental evaluation of composite tissue-type ultrasound and microwave imaging. *IEEE Journal on Multiscale and Multiphysics Computational Techniques*, 4:119–132, 2019.
- [27] Cuiping Li, Lianjie Huang, Nebojsa Duric, Haijiang Zhang, and Charlotte Rowe. An improved automatic time-of-flight picker for medical ultrasound tomography. *Ultrasonics*, 49(1):61 – 72, 2009.
- [28] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering*, 79(11):1309–1331, 2009.
- [29] Stuart Nelson. *Dielectric properties of agricultural materials and their applications*. Academic Press, 2015.
- [30] Allen Taflov and Susan C Hagness. *Computational electrodynamics: the finite-difference time-domain method*. Artech house, 2005.
- [31] Ardavan F Oskooi, David Roundy, Mihai Ibanescu, Peter Bermel, John D Joannopoulos, and Steven G Johnson. Meep: A flexible free-software package for electromagnetic simulations by the fdtd method. *Computer Physics Communications*, 181(3):687–702, 2010.
- [32] KC Lawrence, SO Nelson, and AW Kraszewski. Temperature dependence of the dielectric properties of wheat. *Transactions of the ASAE*, 33(2):535–0540, 1990.
- [33] B Bojanov and IK Georgieva. Interpolation by bivariate polynomials based on radon projections. *Studia Mathematica*, 162:141–160, 2004.

- [34] Hakop A Hakopian. Multivariate divided differences and multivariate interpolation of lagrange and hermite type. *Journal of Approximation Theory*, 34(3):286–305, 1982.
- [35] Charles Hutton. *A Philosophical and Mathematical Dictionary Containing... Memoirs of the Lives and Writings of the Most Eminent Authors*, volume 1. the Author, 1815.
- [36] Mailyñ Perez-Liva, José Manuel Udías, Jorge Camacho, Elena Merčep, Xosé Luís Deán-Ben, Daniel Razansky, and Joaquín L. Herraiz. Speed of sound ultrasound transmission tomography image reconstruction based on bézier curves. *Ultrasonics*, 103:106097, 2020.
- [37] The MathWorks Inc. *MATLAB Function Reference*. The MathWorks Inc., 2021.

Appendix A

Line-Integrals Of Pulse Basis Functions

Integrating pulse basis functions along straight paths is the same as determining the length of the intersection between the elements which support the basis functions and the straight paths. For this thesis, the elements which support the basis functions are either tetrahedra in a three-dimensional tetrahedral mesh, or triangles in a two-dimensional triangular mesh. The process of determining the length of the intersection between elements and paths in three dimensions is described here. The process of integrating a two-dimensional function is sufficiently similar that it is not shown.

Let A and B be the start and end point of the path. Let P , Q , R , and S be the points which define the tetrahedral element. The following conditions are short-circuit checks which can quickly tell if there is *no* intersection between the path and the element:

- The x-components of each of $\{A, B\}$ is greater than the x-component of each of $\{P, Q, R, S\}$
- The y-components of each of $\{A, B\}$ is greater than the y-component of each of $\{P, Q, R, S\}$
- The z-components of each of $\{A, B\}$ is greater than the z-component of each of $\{P, Q, R, S\}$

- The x-components of each of $\{A, B\}$ is less than the x-component of each of $\{P, Q, R, S\}$
- The y-components of each of $\{A, B\}$ is less than the y-component of each of $\{P, Q, R, S\}$
- The z-components of each of $\{A, B\}$ is less than the z-component of each of $\{P, Q, R, S\}$

If any of the above conditions are true, then the intersection between the path and the element is zero. If none of the above conditions are true, then the intersection must be determined using the following procedure:

- For each face of the tetrahedral element
 - Calculate where the path and the plane defined by the face intersect
 - Check whether the intersection point occurs within the face or outside the face
 - If the intersection point occurs outside the face, then ignore it
- If two valid intersection points were found, then calculate the distance between those points
- If one valid intersection point was found
 - One of $\{A, B\}$ is located inside the tetrahedral element
 - Calculate the distance between the interior point and the valid intersection point
- If no valid intersection point was found
 - Check whether both of $\{A, B\}$ are inside the tetrahedral element
 - If both A and B are within the element, calculate the distance between A and B
 - If both A and B are outside the element, then the path-element intersection is zero

Some linear algebra is required to find out where a path and a plane intersect. An example calculation will be shown for just one of the faces of a tetrahedral element, in particular, the face defined by the points $\{P, Q, R\}$. Such an intersection is shown in Figure A.1. Let C be the point where the

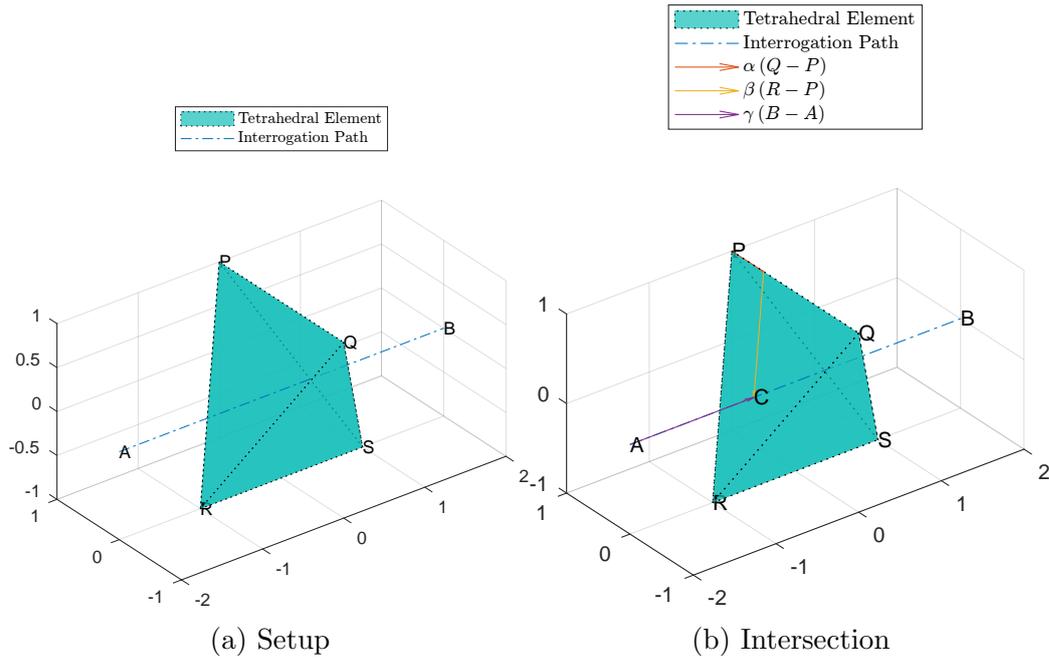


Figure A.1: Calculating a surface intersection point on one face of a tetrahedral element

path intersects the plane defined by $\{P, Q, R\}$. The point C can be located by starting at A and travelling some distance, γ , along \overline{AB} .

$$C = A + \gamma(B - A) \quad (\text{A.1})$$

The point C can also be located by starting at P , then travelling some distance, α , along \overline{PQ} , and then travelling another distance, β , along \overline{PR} .

$$C = P + \alpha(Q - P) + \beta(R - P) \quad (\text{A.2})$$

Equate the two expressions to produce an equation with three unknowns, α ,

β , and γ .

$$P + \alpha(Q - P) + \beta(R - P) = A + \gamma(B - A) \quad (\text{A.3})$$

Isolate P and A .

$$\alpha(Q - P) + \beta(R - P) + \gamma(A - B) = A - P \quad (\text{A.4})$$

The points are all three-dimensional entities, so Equation A.4 is actually three equations. The three unknowns can be recovered from these three equations. Represent Equation A.4 in matrix form, assuming each point is a column vector.

$$\begin{bmatrix} A - P \end{bmatrix} = \begin{bmatrix} (Q - P) & (R - P) & (A - B) \end{bmatrix} \times \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} \quad (\text{A.5})$$

Solving this equation yields the distances α , β , and γ . These distances, as they pertain to finding C , are shown in Figure A.1b.

Whether C exists on the line segment between A and B depends on the value of γ . C exists on that line segment only if $0 \leq \gamma \leq 1$

Whether C is contained in the face defined by $\{P, Q, R\}$ depends on the values of α and β . C is contained in the face only if $0 \leq (\alpha + \beta) \leq 1$ and $0 \leq \alpha$ and $0 \leq \beta$.

Appendix B

Line-Integrals Of Polynomial Basis Functions

Polynomial basis functions have infinite support, so integrating a polynomial basis function along a path is not as simple as calculating the intersection of a geometric element and a line, as is the case for pulse basis functions. Performing this integration requires some calculus and algebra which allows the integrand to be expressed as a polynomial of a single variable, t . Once the integrand is expressed in terms of a single variable, a computational method is required to evaluate the integral.

The process of integrating a three-dimensional polynomial basis function along a straight line is described here. The process of integrating a two-dimensional function is sufficiently similar that it is not shown.

B.1 Symbolic Representation

Let $f(x, y, z)$ be a single polynomial basis function. Let i , j , and k be the exponents of x , y , and z , respectively.

$$f(x, y, z) = x^i y^j z^k \tag{B.1}$$

100 APPENDIX B. LINE-INTEGRALS OF POLYNOMIAL FUNCTIONS

Let $I_{i,j,k}$ be the integral of f along some straight path which is parameterized by the variable t , which ranges from 0 to S .

$$I_{i,j,k} = \int_{t=0}^S x(t)^i y(t)^j z(t)^k dl(t) \quad (\text{B.2})$$

The variable, t , is simply a parameterizing variable, not a length. That is why the infinitesimal length, dl , is expressed as a function of t . Given dx , dy , and dz , the following equation can be used to express dl :

$$dl = \sqrt{dx^2 + dy^2 + dz^2} \quad (\text{B.3})$$

The functions $x(t)$, $y(t)$, and $z(t)$ are all linear functions of t . Let (x_0, y_0, z_0) be the starting point of the straight path, and let \hat{k} be a unit vector along the straight path.

$$x(t) = x_0 + \hat{k}_x t \quad (\text{B.4})$$

$$y(t) = y_0 + \hat{k}_y t \quad (\text{B.5})$$

$$z(t) = z_0 + \hat{k}_z t \quad (\text{B.6})$$

These equations allow dx , dy , and dz to be expressed in terms of dt .

$$dx = \hat{k}_x dt \quad (\text{B.7})$$

$$dy = \hat{k}_y dt \quad (\text{B.8})$$

$$dz = \hat{k}_z dt \quad (\text{B.9})$$

Now, dl can be expressed in terms of dt .

$$dl = dt \sqrt{(\hat{k}_x^2 + \hat{k}_y^2 + \hat{k}_z^2)} \quad (\text{B.10})$$

Since \hat{k} is a unit vector, this equation simplifies.

$$dl = dt \quad (\text{B.11})$$

Incorporate the linear functions for x , y , and z into the integral, with $dl = dt$.

$$I_{i,j,k} = \int_{t=0}^{t=S} \left(x_0 + \hat{k}_x t\right)^i \left(y_0 + \hat{k}_y t\right)^j \left(z_0 + \hat{k}_z t\right)^k dt \quad (\text{B.12})$$

Due to the choice of \hat{k} , the value of S is the length of the straight path.

B.2 Computation

The integration is then solved computationally, using vectors of coefficients to represent polynomials, and using the `conv` (\cdot, \cdot) function in Matlab [37] to multiply those polynomials. Let $\mathcal{V}(\cdot)$ represent the transformation of a polynomial to its vector representation. For example, vector representation of the polynomial $x_0 + \hat{k}_x t$ is the following:

$$\mathcal{V}\left(x_0 + \hat{k}_x t\right) = \begin{bmatrix} x_0 \\ \hat{k}_x \end{bmatrix} \quad (\text{B.13})$$

Multiplying two polynomials is analogous to convolving their vector representations.

$$\mathcal{V}\left(\left(x_0 + \hat{k}_x t\right)^2\right) = \text{conv}\left(\begin{bmatrix} x_0 \\ \hat{k}_x \end{bmatrix}, \begin{bmatrix} x_0 \\ \hat{k}_x \end{bmatrix}\right) \quad (\text{B.14})$$

The coefficients of $\left(x_0 + \hat{k}_x t\right)^i$ are calculated by iterative convolution.

$$\mathcal{V}\left(\left(x_0 + \hat{k}_x t\right)^i\right) = \text{conv}\left(\begin{bmatrix} x_0 \\ \hat{k}_x \end{bmatrix}, \underbrace{\text{conv}\left(\begin{bmatrix} x_0 \\ \hat{k}_x \end{bmatrix}, \dots \text{conv}\left(\begin{bmatrix} x_0 \\ \hat{k}_x \end{bmatrix}, \begin{bmatrix} x_0 \\ \hat{k}_x \end{bmatrix}\right) \dots}_{(i-1) \text{ times}}\right)\right) \quad (\text{B.15})$$

The same process can be used to determine the coefficients of $\left(y_0 + \hat{k}_y t\right)^j$ and $\left(z_0 + \hat{k}_z t\right)^k$. Finally, the coefficients of the integrand of $I_{i,j,k}$ are calculated

via convolution.

$$\text{conv} \left(\mathcal{V} \left((x_0 + \hat{k}_x t)^i \right), \text{conv} \left(\mathcal{V} \left((y_0 + \hat{k}_y t)^j \right), \mathcal{V} \left((z_0 + \hat{k}_y t)^k \right) \right) \right) \quad (\text{B.16})$$

With the coefficients of the integrand known, the integral $I_{i,j,k}$ can be evaluated with the Matlab functions `polyint` and `polyval` [37].

Appendix C

Details for Extracting Auxiliary Material Properties

The algorithm described in this thesis deals with slowness and attenuation as the material properties of interest. This appendix provides derivations and instructions for translating slowness and attenuation to other material properties which govern wave propagation. The properties to be calculated depend on the imaging modality, i.e. acoustic or electromagnetic. For acoustic waves, slowness and attenuation will be translated to the real and imaginary parts of a complex compressibility, $\tilde{\kappa}$. For electromagnetic waves, slowness and attenuation will be translated to the real and imaginary parts of a complex permittivity, $\tilde{\epsilon}$. The translation for acoustic properties is shown in Section C.2, and the translation for electromagnetic properties is shown in Section C.3.

Both phenomena will use the following two-step procedure to translate slowness and attenuation to other properties:

1. Use slowness and attenuation to calculate a complex wave number, \tilde{k}
2. Use the complex wave number to calculate the modality-dependent properties.

In this appendix, the real part of quantity Q is represented as $\Re[Q]$, and the imaginary part of quantity Q is represented as $\Im[Q]$. The angular

frequency, ω is related to the actual frequency, f via the following relation:

$$\omega = 2\pi f \quad (\text{C.1})$$

C.1 Calculating the Complex Wave Number

The translation of material properties is enabled by examining the behaviour of a plane wave. It will be assumed that the plane wave propagates along \hat{z} , and that the plane wave is completely described by an initial value, u_0 and a complex wave number, \tilde{k} .

$$u(z) = u_0 \exp(-j\tilde{k}z) \quad (\text{C.2})$$

The term $\exp(-j\tilde{k}z)$ can be split into a term that controls magnitude and a term that controls phase.

$$u(z) = u_0 \exp(-j\Re[\tilde{k}]z) \exp(\Im[\tilde{k}]z) \quad (\text{C.3})$$

The term $\exp(\Im[\tilde{k}]z)$ controls the magnitude of $u(z)$, while the term $\exp(-j\Re[\tilde{k}]z)$ controls the phase.

The real part of \tilde{k} can be extracted from the wavelength, λ , which is dictated by the frequency of operation, f , and the slowness, c^{-1} .

$$\Re[\tilde{k}] = \frac{2\pi}{\lambda} = 2\pi f c^{-1} \quad (\text{C.4})$$

It will be useful later on to work with the term $\frac{\Re[\tilde{k}]}{\omega}$, so that term is calculated here.

$$\frac{\Re[\tilde{k}]}{\omega} = c^{-1} \quad (\text{C.5})$$

The imaginary part of \tilde{k} can be related to the attenuation by considering the magnitude of Equation C.3.

$$|u(z)| = |u_0| \exp(\Im[\tilde{k}]z) \quad (\text{C.6})$$

It is useful to consider this magnitude with respect to a reference point.

$$\frac{|u(z + \Delta z)|}{|u(z)|} = \exp\left(\Im\left[\tilde{k}\right] \Delta z\right) \quad (\text{C.7})$$

This ratio can also be determined by via the material's attenuation.

$$20 \log_{10}\left(\frac{|u(z + \Delta z)|}{|u(z)|}\right) = -\alpha \times (100\Delta z) \times \frac{f}{1 \times 10^6} \quad (\text{C.8})$$

Combine Equation C.7 and Equation C.8 to get the following equation:

$$20 \log_{10}\left(\exp\left(\Im\left[\tilde{k}\right] \Delta z\right)\right) = \frac{-\alpha \Delta z f}{1 \times 10^4} \quad (\text{C.9})$$

Reduce the logarithm and cancel terms to produce the following equation:

$$\Im\left[\tilde{k}\right] = \frac{-\alpha f \ln(10)}{2 \times 10^5} \quad (\text{C.10})$$

Again, this term will be expressed with respect to ω .

$$\frac{\Im\left[\tilde{k}\right]}{\omega} = \frac{-\alpha \ln(10)}{4\pi \times 10^5} \quad (\text{C.11})$$

C.2 Acoustic Properties

The wave number of an acoustic plane wave is given by the following equation:

$$\tilde{k}^2 = \omega^2 \rho \tilde{\kappa} \quad (\text{C.12})$$

Here, ρ is the medium's density, and $\tilde{\kappa}$ is the complex-valued compressibility. The real and imaginary parts of this equation reduce to the following two equations:

$$\Re\left[\tilde{k}\right]^2 - \Im\left[\tilde{k}\right]^2 = \omega^2 \rho \Re\left[\tilde{\kappa}\right] \quad (\text{C.13})$$

$$2\Re\left[\tilde{k}\right] \Im\left[\tilde{k}\right] = \omega^2 \rho \Im\left[\tilde{\kappa}\right] \quad (\text{C.14})$$

Solve for the real and imaginary parts of $\tilde{\kappa}$.

$$\Re[\tilde{\kappa}] = \frac{\Re[\tilde{k}]^2 - \Im[\tilde{k}]^2}{\omega^2 \rho} \quad (\text{C.15})$$

$$\Im[\tilde{\kappa}] = \frac{2\Re[\tilde{k}]\Im[\tilde{k}]}{\omega^2 \rho} \quad (\text{C.16})$$

Substitute in Equation C.5 and Equation C.11 to get rid of the wave number and introduce slowness and attenuation. The resulting equations constitute a translation from slowness and attenuation to complex compressibility.

$$\Re[\tilde{\kappa}] = \frac{(c^{-1})^2 - \left(\frac{\alpha \ln(10)}{4\pi \times 10^5}\right)^2}{\rho} \quad (\text{C.17})$$

$$\Im[\tilde{\kappa}] = \frac{-\alpha c^{-1} \ln(10)}{2\pi \times 10^5 \rho} \quad (\text{C.18})$$

C.3 Electromagnetic Properties

The wave number of an electric plane wave in a non-magnetic domain is given by the following:

$$\tilde{k}^2 = \omega^2 \mu_0 \tilde{\varepsilon} \quad (\text{C.19})$$

Here, μ_0 is the vacuum permeability, and $\tilde{\varepsilon}$ is the complex-valued permittivity. The complex permittivity is related to the dielectric permittivity, ε and conductivity, σ via the following equation:

$$\tilde{\varepsilon} = \varepsilon - j \frac{\sigma}{\omega} \quad (\text{C.20})$$

The form of Equation C.19 is the same as Equation C.12, with μ_0 taking the place of ρ and $\tilde{\varepsilon}$ taking the place of $\tilde{\kappa}$. It follows, then, that the real and

imaginary parts of $\tilde{\varepsilon}$ are given by the following equations:

$$\Re[\tilde{\varepsilon}] = \frac{(c^{-1})^2 - \left(\frac{\alpha \ln(10)}{4\pi \times 10^5}\right)^2}{\mu_0} \quad (\text{C.21})$$

$$\Im[\tilde{\varepsilon}] = \frac{-\alpha c^{-1} \ln(10)}{2\pi \times 10^5 \mu_0} \quad (\text{C.22})$$