

Seismic Broadband Full Waveform Inversion

by shot/receiver refocusing

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus Prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op vrijdag 18 januari 2013 om 10.00
uur
door

Peter Rudolf HAFFINGER

Master of Science in Applied Geophysics,
Technische Universiteit Delft, Nederland,
Eidgenössische Technische Hochschule Zürich, Zwitserland,
Rheinisch-Westfälische Technische Hochschule Aachen, Duitsland,
geboren te Trier, Duitsland.

Dit proefschrift is goedgekeurd door de promotoren:

Prof. dr. ir. A. Gisolf

Prof. dr. ir. P.M. van den Berg

Samenstelling promotiecommissie:

Rector Magnificus	voorzitter
Prof. dr. ir. A. Gisolf	Technische Universiteit Delft, promotor
Prof. dr. ir. P.M. van den Berg	Technische Universiteit Delft, promotor
Prof. dr. H.P. Urbach	Technische Universiteit Delft
Prof. dr. J.O.A. Robertsson	Eidgenössische Technische Hochschule Zürich
Dr. ir D.J. Verschuur	Technische Universiteit Delft
Dr. ir R.F. Remis	Technische Universiteit Delft
Dr. ir. A. Abubakar	Schlumberger-Doll Research
Prof. dr. ir. L.J. van Vliet	Technische Universiteit Delft, reservelid

SUPPORT

The research for this thesis was financially supported by the DELPHI Consortium.

ISBN 978-90-8891-559-8

Copyright © 2012 by P.R.Haffinger

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilised in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without the prior permission of the author.

Typeset by the author with the L^AT_EX 2_ε Documentation System.

Published by Uitgeverij BOXPress, 's-Hertogenbosch, The Netherlands

Printed by Proefschriftmaken.nl

Contents

1	Introduction	1
1.1	Future challenges of the upstream geo-energy	1
1.2	Acoustic (seismic) imaging	2
1.3	The seismic value chain	3
1.3.1	Seismic data acquisition	4
1.3.2	Standard processing	4
1.3.3	Interpretation	4
1.3.4	Reservoir characterisation	5
1.4	Imaging versus inversion	6
1.5	Historical and recent developments	8
1.6	Thesis overview	10
2	Theoretical formulation in 2-D	13
2.1	Integral formulation of the wave equation	14
2.2	Alternating solution of the nonlinear inverse problem	16
2.3	Linear inversion of the data equation	17
2.3.1	Minimisation problem	17
2.3.2	Conjugate-gradient scheme	18
2.3.3	Multiplicative regularisation	19
2.3.4	The importance of regularisation: A synthetic example	22
2.3.5	Convergence of multiplicative regularisation	24
2.4	Nonlinear field update based on the domain equation	25
2.5	Comments on background models	29
2.6	Outlook	30
3	Localised linear full waveform inversion	31
3.1	Computational requirements of a CG scheme	32
3.2	Localisation by source/receiver-refocusing	35

3.2.1	Backpropagation of the seismic surface data	37
3.2.2	Backpropagation of Green's functions and total fields	42
3.3	Numerical example of local full waveform inversion	45
4	Nonlinear total field estimation	51
4.1	Synthetic example I: The Delphi temple	52
4.2	Synthetic example II: 2-D layered geology	55
5	Combining localised full waveform inversion and nonlinear total field estimation	63
5.1	Subsurface model and synthetic seismic dataset	64
5.2	Initialisation of the scheme: Linear inversion	65
5.3	First nonlinear field update	69
5.4	Overburden and inter domain multiples	71
5.4.1	Classification of primaries and multiples	71
5.4.2	Subtraction of complimentary data as exact solution to localisation	75
5.5	Fully iterated nonlinear inversion results	77
6	Reservoir-oriented nonlinear full waveform inversion	83
6.1	Acoustic 1.5-D nonlinear full waveform inversion	84
6.1.1	A numerical example of acoustic 1.5-D full waveform inversion	87
6.2	Input for reservoir-oriented inversion from surface data	94
6.2.1	Reservoir-oriented inversion by refocusing (redatuming, double backpropagation, downward projection, etc.)	94
6.2.2	Reservoir-oriented inversion by demigration	95
6.3	Reservoir-oriented elastic full waveform inversion: A real data example	96
6.3.1	The elastic scattering framework	96
6.3.2	A field-data example from the Middle East	98
7	Conclusions and recommendations	111
7.1	Conclusions	111
7.2	Recommendations	113
A	Derivation of the gradients	115
	Bibliography	117
	Summary	123
	Samenvatting	125
	Acknowledgements	127

Curriculum Vitae

Introduction

1.1 Future challenges of the upstream geo-energy

One of the most important future challenges for human civilisation will be to meet the rapidly increasing global energy demand. Despite a steady growth in using alternative energy sources, hydrocarbons will need to provide the major contribution to the global energy mix for several decades. Expectations are that the major hydrocarbon reservoirs were already found and it can be observed that their production rates are declining. This development creates a major challenge for the upstream geo-energy industry to increase their production rates while at the same time replacing existing reserves (see figure 1.1).

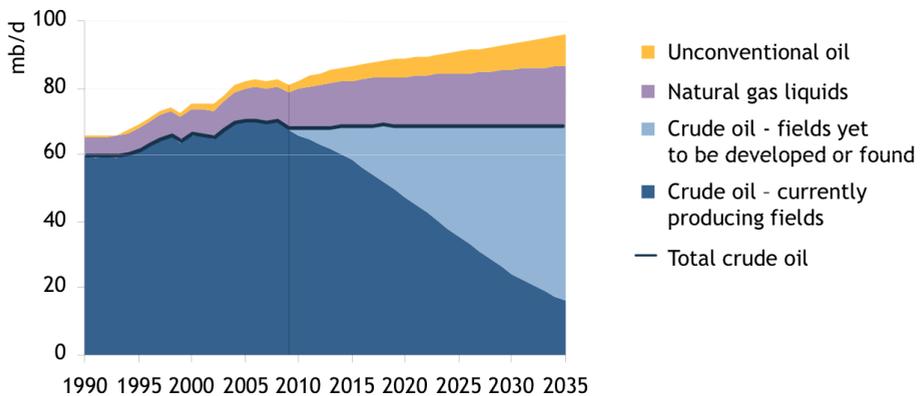


Figure 1.1: As conventional oil production declines, the gap between supply and demand would have to be filled by oil fields yet to be recovered. Graph: World Energy Outlook 2010, International Energy Agency.

The point is that the real energy issue is not about availability, but about affordability. As long as we can produce energy at an affordable price, the earth provides sufficient sources to fulfil the human energy demand. Interestingly, this statements holds for fossil fuels as well as for renewable energy sources as wind, solar or geothermal.

For the oil exploration industry the challenges are to produce the existing reservoirs in an efficient manner, but also to find new oil reservoirs to replace the current reserves. To fulfil this task, oil companies need to have a good understanding of the earth's subsurface, specifically the geology and the types of rocks down to a depth of several kilometres below the surface. As a result it turns out that the oil companies that have access to the best geo-imaging technologies are the most competitive ones. This observation is the motivation for this thesis: to improve the existing geo-imaging technology to increase the business opportunities of the upstream geo-energy industry while at the same time reducing its operational costs. Consequently, during a transition phase we will be able to provide fossil fuels at an affordable price until renewables can reach the capacity to deliver the dominant part of the global energy mix.

1.2 Acoustic (seismic) imaging

In the oil exploration industry it is crucial to have reliable images of the earth's subsurface in order to find and locate potential oil and gas reservoirs. It must be understood that the formation and accumulation of hydrocarbons as oil and gas occur only in very specific geological settings. Therefore, good structural as well as quantitative imaging of the geology and the rock formations in an area that is expected to contain oil reservoirs is absolutely crucial to the success of any oil company. One does not want to miss out on opportunities of a significant discovery. Also, one does not want to spend millions of dollars on the development of a potential field that eventually turns out to be unprofitable, or the other way round, underdevelop a field that has a higher potential.

A well known option to obtain detailed information on the properties of the subsurface is by drilling a well several kilometers deep into the earth. Although this method delivers in principle the ground truth, it can be extremely expensive and it provides information only along the well trajectory. In the upstream geo-energy industry images of an extended subsurface volume are required for appraisal and development decisions. The technical challenge then becomes to image the interior of a certain object, in our case the earth's subsurface, without actually opening or destroying it. It should be noted that this challenge exists in various application fields, among which medical imaging or non-destructive testing.

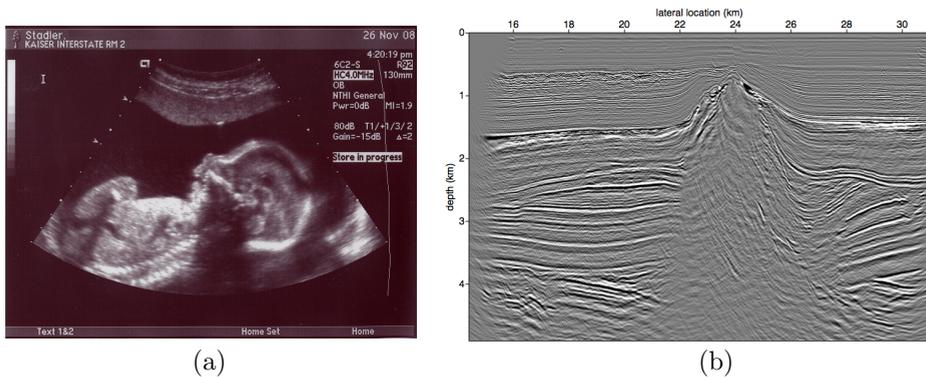


Figure 1.2: *Ultrasonic image of a growing baby in the belly of a pregnant women (a) and an image of the earth’s subsurface as it is used in the oil exploration industry (b). Interestingly, both images are obtained with similar technologies but the scale of the experiments obviously differ significantly.*

The task can be achieved by illuminating the interior of the object with e.g. acoustic or electromagnetic waves, while receivers placed around the object measure the response. From these measurements an image of the properties of the object’s interior can be derived. A well-known example of this procedure is ultrasonic images of a pregnant woman’s belly to determine the position and size of a growing foetus (see figure 1.2 (a)). Here an operator puts an ultrasonic (acoustic) probe on the belly and by measuring the reflections an image of the abdominal cavity is obtained. The same concept is widely used in the oil and gas exploration industry to obtain images of the subsurface (see figure 1.2 (b)). A typical experiment involves sending acoustic waves (sound) into the earth where they are reflected at the boundaries of different geological layers. The reflected waves can then be measured at the earth’s surface and images of the subsurface can be derived from these measurements.

1.3 The seismic value chain

The entire process of finding and eventually producing an oil field with the help of seismic imaging technology includes various steps of which each has a specific purpose. Initially seismic data must be acquired by sending acoustic wave fields into the subsurface and by measuring the response. These datasets can then be processed to obtain a structural subsurface image of the area under consideration. By interpreting the results it should ideally be possible to identify rock formations that due to their specific environment of deposition and structure potentially contain oil or gas. Once a discovery has been made an appraisal phase follows, in which more dedicated seismic is acquired and more wells are drilled. Finally,



Figure 1.3: *The process of finding an oil reservoir with the help of seismic imaging includes several phases which are described by the seismic value chain.*

a decision is made if the oil field can and will be produced economically. More sophisticated methods, to obtain also quantitative information about the target, might be applied during the reservoir characterisation phase. In the following we describe each phase of the seismic value chain (see figure 1.3) in more detail.

1.3.1 Seismic data acquisition

The very first step in seismic exploration is the acquisition of seismic data. Data can be measured on land (onshore), or in a marine environment (offshore) while the process itself is in principle the same. A seismic source is used to generate an acoustic wave that propagates into the earth's subsurface. The wavefield gets then reflected at the boundaries between geological layers and starts propagating back to the surface again where it is measured by a network of detectors. The earth response for each shot is measured for several (7-8) seconds before the next source is activated at a different surface location and a new measurement is started. By repeating this process several hundreds to thousands of times, a huge volume of seismic data can be acquired. Acquisition of a marine seismic dataset is sketched in figure 1.4.

1.3.2 Standard processing

Seismic data acquisition is the very first step in seismic exploration, but the raw data that are measured with detectors along the surface is not of much use for direct interpretation. To translate the seismic measurements into an image of the earth's subsurface, seismic data processing needs to be applied. The outcome of this process is a structural image of the subsurface showing the geology. This information is essential because hydrocarbon reservoirs are only found in specific geological areas. An example of a high resolution structural subsurface image is given in figure 1.5. As we can see, standard processing can provide a high resolution structural image.

1.3.3 Interpretation

Interpretation of structural images allows us to derive geological scenarios that might have happened over millions of years. Because oil and gas can develop and accumulate in very specific geological settings only, seismic interpretation

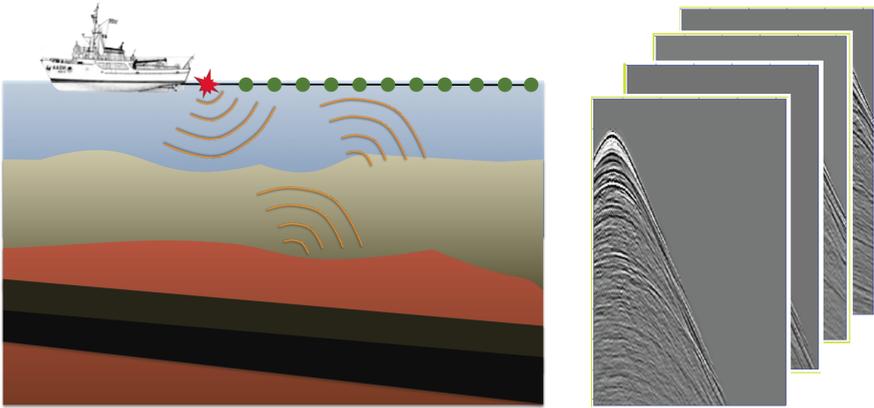


Figure 1.4: *Acquisition of seismic data: Acoustic waves are sent into the earth's subsurface where they are reflected at the boundaries between different geological layers. By measuring the response at the earth surface by a network of detectors a seismic shot gather is obtained. Repeating this process several hundreds to thousands of times results in a complete seismic survey.*

will usually result in identifying potential reservoirs. This can be considered as the most important phase in seismic exploration, because for most seismic datasets the result of this phase will be the final product of the seismic value chain. Still, structural images do not tell us what rock types, e.g. clay, salt, carbonate, sand etc., are present. The actual rock properties are of special interest because they give us more insight if hydrocarbons can be expected and in what quantities, in the area under investigation. Rock properties can be recovered by full waveform inversion but the process is computationally expensive and is therefore not routinely applied in the exploration industry.

1.3.4 Reservoir characterisation

After a potential reservoir has been identified on a structural image, more specialised techniques are usually applied to the target area. Because the reservoir is only a subset of the entire subsurface, also computationally expensive schemes become feasible, which might have been impossible to compute during the stage of standard processing. Here, inversion techniques, including full waveform inversion play a fundamental role because they have the potential to provide the rock properties, instead of structural information only. It has to be understood that oil bearing layers can be as thin as several meters and accurate high resolution characterisation is the key to successful production of a reservoir.

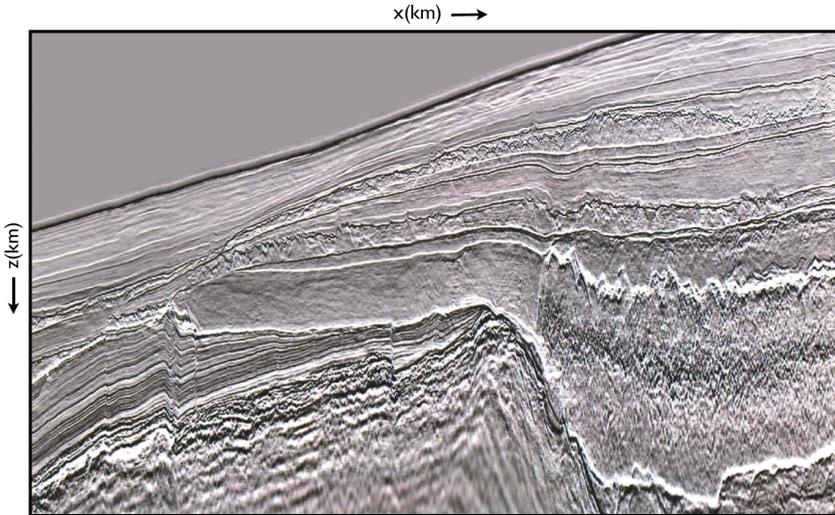


Figure 1.5: *Structural image of the subsurface obtained by seismic data processing. The geology of the area was reconstructed with a high resolution. Based on this result a potential oil reservoir will be identified. It should be realised that a structural image does not provide the rock properties yet. This information can in principle be obtained by full waveform inversion but the process is computationally rather expensive. (Image is property of PGS)*

1.4 Imaging versus inversion

We have seen that the seismic value chain contains different parts of which each one has a specific role. During the standard processing stage, imaging techniques are usually used to get structural information. Reservoir characterisation can then be performed with the help of inversion techniques to provide the rock properties. In this section we will discuss the differences between imaging and inversion and what limitations both techniques have.

Let us start by seismic imaging. Here seismic data are backpropagated to a certain point in the subsurface while a simulated source pulse is forward propagated to the same location. Next the imaging condition is applied. If at the point that is considered a geological boundary was present, this process should result in an image amplitude. If at this location no boundary was located, the imaging condition is supposed to give a zero result. This procedure is then repeated for all sources and every subsurface location and by summing all the results a structural image can be obtained. The process is schematically shown in figure 1.6. It is important to realise that the imaging condition is an intrinsically linear concept. Here it is assumed that a wavefield propagates downward until it reaches a geolo-

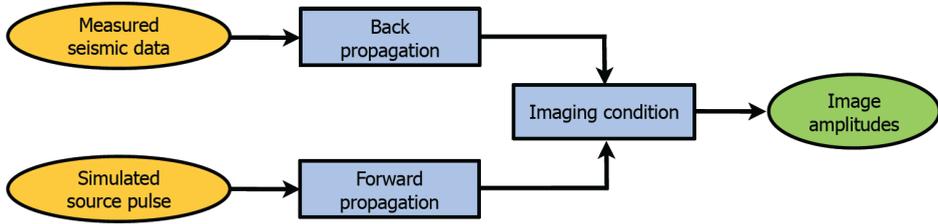


Figure 1.6: *Schematic description of seismic imaging.*

gical boundary. After being reflected the wavefield propagates upwards where it is measured at the acquisition surface. It should be understood that this separation into downward and upward propagating wavefields is artificial and introduces artefacts, especially in areas where no reflectors were present.

Seismic imaging can be performed for each source and each subsurface point independently. Afterwards the individual results are combined to get an image of the entire subsurface, using the illumination of all seismic sources. This is very advantageous from a computational point and if the process is implemented in a parallel manner, today's large computer clusters can deliver a high resolution image as it was shown in figure 1.5.

At the same time it should be realised that current seismic imaging techniques are just feedforward processes. Seismic data and a source pulse are input with image amplitudes being the output. There is no feedback which would allow us to detect inconsistencies between the obtained structural image and the measured seismic data.

The opposite situation can be found if inversion techniques are considered. Here the aim is to find the media properties on a dense subsurface grid, instead of image amplitudes. Synthetic seismic data are forward modelled and compared with the measured dataset. The differences are then used to update the full set of media properties and the process is repeated until a satisfactory match between modelled and measured dataset is obtained. In figure 1.7 a flow chart for the process is shown. It can be seen that, in contrast to imaging techniques, a closed loop exists, meaning that feedback from the obtained media properties is used to verify any inconsistency between the measured data and the inversion result.

The drawback of inversion is the need of rather large computer systems to make it feasible on a scale as required in the field of hydrocarbon exploration. This statement especially refers to the processes in the two blue boxes in figure 1.7. Forward

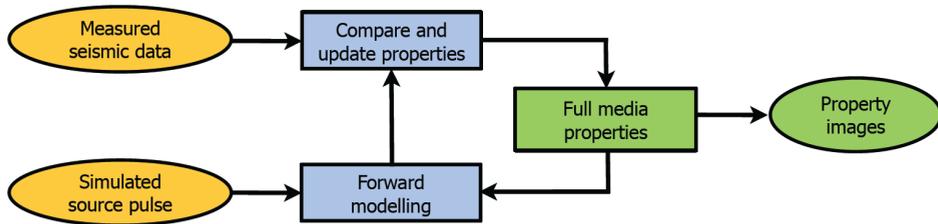


Figure 1.7: *Schematic description of seismic inversion.*

modelling of seismic wavefields is known to be expensive in terms of the number of calculations that must be performed. However, because it is highly parallelisable, with existing large computer clusters this process can be made feasible. In the upper blue box in figure 1.7 the difference between modelled and measured data is used to update the full media properties. This process requires large amounts of memory and we will see in chapter 3 that today’s computer systems cannot provide sufficient memory to make inversion feasible with a resolution that is competitive to global seismic imaging.

By global we do not refer to the entire earth but to the part of the subsurface that is supposed to be imaged by seismic techniques. The extension of a global domain in this context is usually several kilometres in width and depth. Later we introduce a local domain, which is smaller and has an extension of several hundreds of meters in each dimension.

This brings us to the scientific motivation of this thesis. We will develop a full waveform inversion scheme that has the potential to provide property images of the subsurface with a resolution that goes beyond the resolution of today’s imaging technologies and at the same time produces property images rather than just structural images.

Before we go into the detail of the method we will review previous scientific work that was done in the field of full waveform inversion in the upstream oil industry.

1.5 Historical and recent developments

It was already realised more than three decades ago that full waveform inversion is in principle a suitable tool to derive the properties of the subsurface from seismic data. First attempts at seismic applications were made by several authors starting in the late 70’s among which Bamberger et al. (1977, 1982), Symes (1981),

Lailly (1983) and Tarantola (1984, 1987).

The initial work on full waveform inversion has been developed further and today gradient based methods (Mora, 1987; Crase et al., 1990) and Gauss-Newton optimisation schemes (Pratt et al., 1998; Epanomeritakis et al., 2008; Fichtner and Trampert, 2011; Métivier et al., 2012) are most commonly used to recover the subsurface properties from seismic measurements. An excellent overview of existing full waveform inversion work was given by Virieux and Operto (2009). A drawback of most approaches is their large memory requirement and the necessity to solve many full forward problems during the process. If 2-D or 3-D full waveform inversion is aimed for, even currently available huge computer clusters cannot provide the capacity that would be required if the full seismic bandwidth of the measured data is to be used. By full bandwidth is meant spectra with maximum frequencies between 50 – 100 Hz.

To overcome this limitation, recent attempts have been made by encoding sources and receivers (Krebs et al., 2009; Abubakar et al., 2011; Routh et al., 2011), or by using compressive sensing, if the system to be solved is sparse in some sense (Moghaddam and Herrmann, 2010). Also multi scale approaches could be deployed to decrease the problem size while at the same time incorporating higher frequencies (Bunks et al., 1995; Brossier et al., 2009). Interestingly, the seismic community seems to focus on reduction of the data space while reduction of the model space is not really investigated yet. Despite all efforts to reduce the amount of data that is simultaneously needed, maximum frequencies used for inversion (approx. 20 Hz in 2-D and several Hz in 3-D) are still far away from the maximum frequencies measured in seismic data. Consequently, the role of full waveform inversion is still limited to providing background velocity models that can be used as input to a conventional imaging scheme (Plessix et al., 2010; Prioux et al., 2010; Vigh et al., 2011). This means that the final output of the seismic value chain remains a structural image that will be used for interpretation.

In this thesis we describe a gradient based full waveform inversion scheme with the potential to increase significantly the upper frequency limit that can be taken into account for full waveform inversion. We discuss mainly two new aspects: reduction of the memory requirements by localising the inversion and iterative calculation of the wave fields in the inversion domain to avoid solving the full forward problem many times. The purpose of the present work is not to speed-up waveform inversion but to reduce its computational complexity, especially in terms of simultaneous memory requirements. As a result, broadband full waveform inversion in 2-D and eventually 3-D should become feasible with a very fine subsurface discretisation (≈ 5 m) while inverting frequencies of 50 Hz and above. Consequently, full waveform inversion could become feasible to provide a resolution that is even better than conventional structural imaging. Because full waveform inversion provides structural as well as quantitative information, the

results would become the final output of the seismic value chain. In the following we give an overview of the thesis and discuss the relationship with the work done by other authors.

1.6 Thesis overview

Chapter 2 - We start by giving a theoretical formulation of 2-D full waveform inversion as it will be described in this thesis. This includes introduction of the integral representation of the wave equation which will be needed for forward modelling of seismic wavefields during the inversion process. Furthermore, we discuss a multiplicatively regularised conjugate gradient scheme to recover the subsurface properties from the data.

Chapter 3 - It was mentioned that full waveform inversion is computationally expensive. In this chapter we describe an approach to split up a large scale inversion problem into many small scale, and highly parallelisable, local linear inversions. In a first step seismic data are backpropagated into the subsurface (Wapenaar and Berkhout, 1989) which allows us to invert for the medium properties of a local domain only, while the rest of the subsurface can be disregarded for the moment. It should be realised that domain decomposition methods are generally used to speed-up the forward problem (Bamberger et al., 1997). This is different to our approach where we use the concept of splitting up the inverse problem into many local domains. The method combines redatuming (Berryhill, 1984; Shtivelman and Canning, 1988; Mulder, 2005) and target-oriented inversion approaches (Valenciano et al., 2006; Staal et al., 2010). It will be shown that the local results can be combined to a high resolution global property model.

Chapter 4 - In the approach discussed, so far local inversions are solved by a linear CG scheme. To incorporate the nonlinear relationship between the measured seismic data and the subsurface properties we describe a total wave field estimation scheme based on the integral representation of the wave equation. This means that nonlinear effects, e.g. multiple scattering, transmission and true traveltimes in the inverted medium, will be used to obtain improved property models, rather than being regarded as noise.

Chapter 5 - In this chapter we combine localised linear full waveform inversion with the previously discussed nonlinear total wave field estimation technique. In this way a high-resolution global property model can be obtained while fully honouring the nonlinear relationship between the data and the subsurface properties. We will see that a resolution can be obtained that goes even beyond current imaging techniques. Therefore, also improved structural resolution can be achieved by nonlinear full waveform inversion that cannot be obtained by standard imaging.

Chapter 6 - Here we apply nonlinear full waveform inversion for the purpose of reservoir characterisation to a real dataset from the Middle East. Although the main part of this thesis deals with acoustic 2-D inversion, we use a local 1.5-D assumption in this chapter. Assuming a horizontally layered subsurface under a local target point is a well known and robust approximation that is used in many quantitative interpretation tools today. As a result of this assumption it is then feasible to extend the full waveform inversion to the elastic framework and recover the elastic properties over a reservoir sequence.

Chapter 7 - Finally we summarise the developments of this thesis and give recommendations for future work.

Theoretical formulation in 2-D

Full waveform inversion is in principle a suitable tool to retrieve the properties of the subsurface from seismic measurements. The process involves simulation of seismic measurements assuming a known subsurface model and from the difference between synthetic and real seismic datasets an update of the subsurface model is supposed to be obtained. This implies that a theoretical framework is needed to describe the propagation of seismic wavefields through the subsurface. To this purpose we give a mathematical description of the inversion scheme as will be discussed in this thesis in the current chapter.

We have mentioned in the introduction that inversion techniques are mostly applied during the reservoir characterisation phase. Reservoir-oriented full waveform inversion is commonly based on the 1.5-D assumption. In this part of the thesis we will be dealing with 2-D inversion, taking into account two-dimensional wave propagation and subsurface models. The increase in dimensionality has its price, because it makes the computation of the scheme significantly more expensive. To demonstrate the feasibility of the proposed inversion scheme, we compromise on the set of parameters that is supposed to be recovered by the current implementation. We do this by limiting ourselves to the acoustic case where the density is assumed to be known and constant. Although extension to 2-D elastic and ultimately to 3-D elastic inversion is the aim of future research, we start with a single parameter inversion to prove the concept.

Before going into the equations we present a summary of the most important variables that will be used throughout this thesis, in table 1. It should be realised that the vector $\vec{x} = (x, z)$ is now two-dimensional, which is in contrast to the more traditional reservoir oriented case, where all variables depend on a local z -coordinate only.

Table 2.1: *Most important variables used throughout the thesis.*

ω	=	temporal angular frequency,
\mathcal{D}	=	scattering domain (subsurface),
\mathcal{S}	=	surface domain,
\vec{x}_s	=	source location along \mathcal{S} ,
\vec{x}_r	=	receiver location along \mathcal{S} ,
\vec{x}, \vec{x}'	=	scattering locations in \mathcal{D} ,
$c_0(\vec{x})$	=	acoustic wave velocity of the background medium,
$c(\vec{x})$	=	acoustic wave velocity of the true medium,
ρ_0	=	constant mass density of the background medium,
$\chi(\vec{x})$	=	contrast function,
\mathcal{G}	=	Green's function in background medium,
p_{data}	=	measured scattered field data,
p_0	=	background field at each point in the subsurface,
p_{tot}	=	total field at each point in the subsurface.

2.1 Integral formulation of the wave equation

Any full waveform inversion scheme relies on a forward modelling engine to simulate synthetic wavefields as they would be measured above a known subsurface model. This is needed to compare the modelled wavefields with the measured data, from the difference of which an update of the subsurface model may be obtained. Modelling seismic data requires solving the wave equation, acoustic or elastic, in a numerical sense. Several numerical methods exist to perform this task, the best researched and most widely used ones being based on finite-difference (FD) techniques (Mitchell and Griffiths, 1980). The medium of interest, in our case the subsurface, and the wave equation have to be discretised and can then subsequently be solved based on the differences between neighbouring grid-cells. In spite of its frequent application, FD has a lot of drawbacks among which its instability for high frequencies or large scatterers, its being subject to dispersion, and the fact that the method is computationally rather expensive. A reason for the notably high computational requirements is that FD techniques inherently solve for the exact wavefield at any point in the medium, taking an infinite order of scattering into account.

To overcome these problems, we use an integral formulation of the wave equation to describe the propagation of waves through the subsurface. In contrast to the previously discussed FD method, where differences were used between adjacent grid-points, the integral formulation sums the contributions of all grid-points. This makes it much more stable from a numerical point of view and consequently the sampling criteria can be reduced by a factor of two, compared to FD tech-

niques. Additionally, the integral formulation builds up the wavefields iteratively by adding single orders of scattering. In this way, the proposed inversion approach does not require us to run an expensive full FD solution after each iteration of the inversion process.

In the following we introduce the integral formulation of the wave equation and we show that properties of the subsurface can be derived from measured data through inversion using a conjugate gradient scheme. As mentioned before, our ultimate goal is to develop an elastic nonlinear full waveform inversion, but for the sake of practicality we limit ourselves to the acoustic 2-D case with constant density, in this thesis. We strongly believe that based on the concepts that will be described here, future extension to 2-D elastic and eventually 3-D full waveform inversion could become feasible within the next decade. This would be a major breakthrough for the oil exploration industry and could potentially change the way we process and interpret seismic data, in a fundamental way.

For the acoustic case with constant density, the wave equation can be described by a set of two coupled equations, being the data equation and the object equation. The data equation describes the seismic dataset, mostly measured along the earth's surface, in terms of a total field p_{tot} at each grid point in the subsurface, the contrast function χ and the Green's functions \mathcal{G} in a background medium:

$$p_{data}(\vec{x}_r, \vec{x}_s, \omega) = \int_{\vec{x} \in \mathcal{D}} \mathcal{G}(\vec{x}_r, \vec{x}, \omega) \chi(\vec{x}) p_{tot}(\vec{x}, \vec{x}_s, \omega) d\vec{x}. \quad (2.1)$$

Reading from right to left, equation (2.1) can be understood as follows. A source transmits a wavefield that propagates to every point in the subsurface. Note that this wavefield p_{tot} is generally quite complex because it takes the interaction of all scatterers in the subsurface already into account. This wavefield is creating secondary sources in all points where the contrast function χ is non-zero. The secondary sources transmit energy through a smooth background medium to the receivers, represented by the Green's function \mathcal{G} in equation (2.1). The measured seismic data at every receiver are then a summation of all secondary sources. It should be mentioned that direct waves, including ground roll and surface waves are supposed to be removed from the measured data to obtain p_{data} . The contrast function, which is in real data applications generally unknown and should be recovered by inversion of seismic data, is defined as

$$\chi(\vec{x}) = 1 - \left[\frac{c_0(\vec{x})}{c(\vec{x})} \right]^2. \quad (2.2)$$

It depends on the difference between a known background medium $c_0(\vec{x})$ and the true, but unknown, subsurface model $c(\vec{x})$. The total field on the right-hand side in equation (2.1) is dependent on the contrast χ , because it contains the interaction between all subsurface scatterers. Then it follows that there is a nonlinear

relationship between the subsurface properties and the measured seismic data.

If this dependency of the total field on the contrast function is neglected, the data equation can be linearised. This is called the Born-approximation and allows us to obtain approximate subsurface properties by straightforward linear inversion. In the Born-approximation it is assumed that the true total field can be approximated by a simple, non-scattering background field, hence $p_{tot} \approx p_0$.

On the other hand, if the exact total field would be known, the exact contrast function could be obtained again by a simple linear inversion. This is a very interesting observation and it tells us that full nonlinear inversion can be achieved by iterative linear inversions, whereby in every iteration more nonlinearity is incorporated in the total field. That is where the domain equation comes into play, because it describes the total field at each grid point in the subsurface in terms of the incident field and the Green's functions in the background, and the contrast function:

$$p_{tot}(\vec{x}, \vec{x}_s, \omega) = p_0(\vec{x}, \vec{x}_s, \omega) + \int_{\vec{x} \in \mathcal{D}} \mathcal{G}(\vec{x}, \vec{x}', \omega) \chi(\vec{x}') p_{tot}(\vec{x}', \vec{x}_s, \omega) d\vec{x}'. \quad (2.3)$$

For convenience we write can write equation (2.3) in operator notation

$$p_0(\vec{x}, \vec{x}_s, \omega) = [\mathcal{L} p_{tot}](\vec{x}, \vec{x}_s, \omega), \quad (2.4)$$

where the operator \mathcal{L} is defined as

$$[\mathcal{L} p_{tot}](\vec{x}, \vec{x}_s, \omega) = p_{tot}(\vec{x}, \vec{x}_s, \omega) - \int_{\vec{x} \in \mathcal{D}} \mathcal{G}(\vec{x}, \vec{x}', \omega) \chi(\vec{x}') p_{tot}(\vec{x}', \vec{x}_s, \omega) d\vec{x}'. \quad (2.5)$$

If the contrast function is known, the true total field, including the complete propagation in the true subsurface, can be obtained by solving equation (2.3) numerically. Substituting the obtained total field into equation (2.1) then yields seismic data containing the full nonlinear relationship with the subsurface properties. The problem in seismic inversion is that the contrast function is supposed to be found from the seismic data and is generally unknown beforehand. The challenge in nonlinear full waveform inversion is to find the total field at each point in the subsurface as well as the unknown contrast function.

2.2 Alternating solution of the nonlinear inverse problem

Nonlinear inversion requires solving both the data equation and the domain equation. This task can in principle be performed simultaneously but it would be computationally very expensive and rather complicated to do so. That is why in our approach we solve the coupled set of equations separately, in an alternating

manner. In the first step we use the Born approximation to obtain approximate properties of the subsurface by a simple linear inversion. Note that this procedure only involves the data equation. The obtained property model can then be used to update the total field as described by the domain equation. Substitution of the updated total field into the data equation and repeating the linear inversion is supposed to lead to an improved subsurface model. This iterative process is then repeated until the subsurface properties and the updated total field do not change any more. A requirement to obtain a nonlinear inversion result, is that this procedure converges. In the remainder of this chapter all steps that need to be taken to obtain a convergent system are discussed.

2.3 Linear inversion of the data equation

First, we focus on obtaining the contrast function, the subsurface properties, from linear inversion of the data equation. This is an iterative procedure during which seismic data are modelled based on a given subsurface model and the obtained synthetic data are compared to the measured data. The total field in the data equation is fixed at the currently best estimate of the total field. The objective is to find a parameterisation of the subsurface yielding a synthetic dataset that minimises the difference between the synthetic data and the field data. If the two datasets are nearly identical, the synthetic subsurface model is assumed to be a very good approximation of the real subsurface.

2.3.1 Minimisation problem

In our inversion scheme we try to minimise the L^2 -norm of the difference between the measured data p_{data} and the modelled data that is obtained using the currently best estimate of the contrast function and the fixed total field. We define an operator $[\mathcal{K}\chi]$ as the short-hand notation for the integral operator of equation (2.1),

$$[\mathcal{K}\chi](\vec{x}_r, \vec{x}_s, \omega) = \int_{\vec{x} \in \mathcal{D}} \mathcal{G}(\vec{x}_r, \vec{x}, \omega) p_{tot}(\vec{x}, \vec{x}_s, \omega) \chi(\vec{x}) d\vec{x}, \quad (2.6)$$

where the function $\mathcal{G} p_{tot}$ acts as the kernel of the integral operator. The residual between measured and modelled data is obtained as

$$r(\vec{x}_r, \vec{x}_s, \omega) = p_{data}(\vec{x}_r, \vec{x}_s, \omega) - [\mathcal{K}\chi](\vec{x}_r, \vec{x}_s, \omega). \quad (2.7)$$

The error functional F that should be minimised becomes:

$$F^{data} = \frac{\iiint |r(\vec{x}_r, \vec{x}_s, \omega)|^2 d\vec{x}_s d\vec{x}_r d\omega}{\iiint |p_{data}(\vec{x}_r, \vec{x}_s, \omega)|^2 d\vec{x}_s d\vec{x}_r d\omega}. \quad (2.8)$$

It should be realised that all integrals have to be replaced by discrete summations during implementation of the scheme. Next, we will minimise the error functional, defined in equation (2.8), with the help of a conjugate gradient scheme.

2.3.2 Conjugate-gradient scheme

To establish the CG scheme we define the error functional, given in equation (2.8), after each iteration n as

$$\begin{aligned} F_n^{data} &= \eta \iiint |r_n(\vec{x}_r, \vec{x}_s, \omega)|^2 d\vec{x}_s d\vec{x}_r d\omega \\ &= \eta \iiint |p_{data}(\vec{x}_r, \vec{x}_s, \omega) - [\mathcal{K}\chi^{(n)}](\vec{x}_r, \vec{x}_s, \omega)|^2 d\vec{x}_s d\vec{x}_r d\omega. \end{aligned} \quad (2.9)$$

Here we use the scaling factor

$$\eta = \left(\iiint |p_{data}(\vec{x}_r, \vec{x}_s, \omega)|^2 d\vec{x}_s d\vec{x}_r d\omega \right)^{-1}, \quad (2.10)$$

which makes the cost functional equal to one for a vanishing contrast.

Subsequently we want to find a sequence of contrast functions $\chi^{(n)}(\vec{x})$, $n = 1, 2, \dots$, in which the error functional decreases with increasing iterations. Therefore, after each iteration the contrast function is updated as follows:

$$\chi^{(n)}(\vec{x}) = \chi^{(n-1)}(\vec{x}) + \alpha_n \zeta_n(\vec{x}), \quad \vec{x} \in \mathcal{D}. \quad (2.11)$$

Here, the step size of the update is determined by the parameter α_n while the update directions are conjugate gradient directions given by

$$\zeta_1(\vec{x}) = g_1(\vec{x}), \quad \zeta_n(\vec{x}) = g_n(\vec{x}) + \gamma_n \zeta_{n-1}(\vec{x}), \quad n \geq 2. \quad (2.12)$$

The update direction g_n will be the gradient of the cost functional F_n^{data} with respect to the unknown contrast χ , evaluated at step $n - 1$:

$$g_n(\vec{x}) = \eta \Re([\mathcal{K}^* r_{n-1}](\vec{x})), \quad \vec{x} \in \mathcal{D}. \quad (2.13)$$

See Appendix A for a detailed derivation of the gradient. In the equation above, the star denotes complex conjugation of the transposed operator and \Re indicates that only the real part will be used. A detailed derivation of the gradient and more information on CG schemes can also be found in Shewchuk (1994). The adjoint operator $[\mathcal{K}^* r_{n-1}]$ can be seen as a backprojection operator that maps the residual between measured data and modelled data from the surface domain \mathcal{S} to its associated location in the scattering domain \mathcal{D} .

In our conjugate gradient scheme we make use of the Polak-Ribière direction (Polak and Ribière, 1969), by taking

$$\gamma_n = \frac{\int g_n(\vec{x})[g_n(\vec{x}) - g_{n-1}(\vec{x})]d\vec{x}}{\int g_{n-1}(\vec{x}) g_{n-1}(\vec{x})d\vec{x}}, \quad \vec{x} \in \mathcal{D}. \quad (2.14)$$

Finally, the update parameter $\alpha_n(\vec{x})$ is obtained by minimising the cost functional with respect to α_n . Substitution of equation (2.11) into equation (2.9) leads to

$$F_n^{data} = \eta \iiint |p_{data} - [\mathcal{K}(\chi^{(n-1)} + \alpha_n \zeta_n)]|^2 d\vec{x}_s d\vec{x}_r d\omega. \quad (2.15)$$

The optimal step size is now found from minimisation of equation (2.15) with respect to α_n , by setting the derivative equal to zero:

$$\frac{F_n^{data}}{\partial \alpha_n} = 2\eta \iiint \left\{ (d - [\mathcal{K}\chi^{(n-1)}]) [\mathcal{K}\zeta_n] - \alpha_n |\mathcal{K}\zeta_n|^2 \right\} d\vec{x}_s d\vec{x}_r d\omega = 0. \quad (2.16)$$

Consequently the optimal step size becomes:

$$\alpha_n = \frac{\Re \left(\iiint r_{n-1}^*(\vec{x}_r, \vec{x}_s, \omega) [\mathcal{K}\zeta_n](\vec{x}_r, \vec{x}_s, \omega) d\vec{x}_s d\vec{x}_r d\omega \right)}{\iiint |\mathcal{K}\zeta_n(\vec{x}_r, \vec{x}_s, \omega)|^2 d\vec{x}_s d\vec{x}_r d\omega}, \quad (2.17)$$

in which ζ_n was defined in equation (2.12).

To initialise the conjugate gradient scheme we assume $\chi^{(0)}(\vec{x}) = 0$, leading to $r_0(\vec{x}_r, \vec{x}_s, \omega) = p_{data}(\vec{x}_r, \vec{x}_s, \omega)$. In view of equation (2.13) we find the gradient as

$$g_1(\vec{x}) = \eta \Re([\mathcal{K}^* p_{data}](\vec{x})), \quad \vec{x} \in \mathcal{D}. \quad (2.18)$$

Equation (2.17) gives us the first update parameter:

$$\alpha_1 = \frac{\Re \left(\iiint p_{data}^*(\vec{x}_r, \vec{x}_s, \omega) [\mathcal{K}g_1](\vec{x}_r, \vec{x}_s, \omega) d\vec{x}_s d\vec{x}_r d\omega \right)}{\iiint |\mathcal{K}g_1(\vec{x}_r, \vec{x}_s, \omega)|^2 d\vec{x}_s d\vec{x}_r d\omega}. \quad (2.19)$$

The updated contrast function can then be determined as $\chi^{(1)}(\vec{x}) = \chi^{(0)}(\vec{x}) + \alpha_1 g_1(\vec{x})$.

2.3.3 Multiplicative regularisation

In the previous section we did not address the fact that measured seismic data always contains some form of noise. This requires that we stabilise the inversion process. If we would not do so the instability of the inversion would result in contrast functions that are totally out of range. In this section we extend the conjugate gradient scheme in a way that it contains a multiplicative regularisation factor [van den Berg et al. (2003) and Abubakar et al. (2004)].

The error functional F^{tot} becomes a product of the original error functional defined in equation (2.8) and a newly introduced regularisation factor F^{reg} :

$$F_n^{tot} = F_n^{data} F_n^{reg}. \quad (2.20)$$

Here the regularisation factor is defined as

$$F_n^{reg} = \left(\int_{\vec{x}} d\vec{x} \right)^{-1} \int_{\vec{x}} \frac{(|\nabla\chi^{(n)}|^2 + \delta_{n-1}^2)}{(|\nabla\chi^{(n-1)}|^2 + \delta_{n-1}^2)} d\vec{x}, \quad (2.21)$$

which, in the remainder of this thesis, is simplified by defining the weighting function

$$b_n^2 = \left(\int_{\vec{x}} d\vec{x} \right)^{-1} (|\nabla\chi^{(n-1)}|^2 + \delta_{n-1}^2)^{-1}. \quad (2.22)$$

The given definition of the regularisation factor has an edge preserving effect while a smooth solution would be asked for if the second integral in equation (2.21) would be replaced by separate integrals over numerator and denominator. A comparison of both implementations can be found in van den Berg et al. (2003). Note that the regularisation factor contains the spatial gradients of the contrast function. This constraint can be seen as *a priori* information that the contrast distribution of the subsurface has, at least locally, some kind of profile with restricted spatial variability.

In equations (2.21) and (2.22) we make use of the so-called steering parameter δ^2 . We note that both for $\chi^{(n)} = \chi^{(n-1)}$ or for a large steering parameter, the regularisation factor becomes equal to one and the total cost functional F_n^{tot} reduces to the original cost functional F_n^{data} . We choose this steering factor equal to the mean of the variation $|\nabla\chi^{(n-1)}|^2$, over the whole scattering domain D ,

$$\delta_n^2 = \left(\int_{\vec{x}} d\vec{x} \right)^{-1} \int_{\vec{x}} |\nabla\chi^{(n-1)}|^2 d\vec{x}. \quad (2.23)$$

Experience shows that results are very insensitive to the choice of this steering parameter and δ^2 could be varied between 0.5 and 2.0 of the mean variance without changing the inversion result notably. If lower values are taken for δ^2 the regularisation becomes dominant while higher values make the regularisation ineffective. In the whole iterative process, the steering parameter determines itself. At the end of the iterative scheme, when $\chi^{(n)} \approx \chi^{(n-1)}$, the regularisation factor becomes equal to one and we get back all the characteristics of the original cost functional. This autonomous decrease in regularisation for increasing iterations is contrary to an additive regularisation where the damping remains constant for all iterations. This means that with an additive regularisation, signal in the data which is below the level of the damping can never be explained, whereas multiplicative regularisation only rejects signal components that give rise to a high

spatial variation of the contrast.

The gradient of the regularisation functional with respect to the contrast χ (see Appendix A) becomes:

$$g_n^{reg} = \nabla \cdot (b_{n-1}^2 \nabla \chi^{(n-1)}). \quad (2.24)$$

The gradient of the total cost functional can be written as

$$g_n(\vec{x}) = \eta [\mathcal{K}^* r_{n-1}] (\vec{x}) F_{n-1}^{reg} + F_{n-1}^{data} \nabla \cdot (b_{n-1}^2 \nabla \chi^{(n-1)}). \quad (2.25)$$

We refer the interested reader to Abubakar et al. (2004) for more information about the gradient of the total cost functional. The remaining task for each iteration is to calculate the step size α_n that belongs to the update direction ζ_n .

Substitution of the update equation (2.11) for $\chi^{(n)}$ into the total cost functional F_n^{tot} results in the product of two second order polynomials:

$$F_n^{tot} = \underbrace{(A_2 \alpha_n^2 + A_1 \alpha_n + A_0)}_{F_n^{data}} \underbrace{(B_2 \alpha_n^2 + B_1 \alpha_n + B_0)}_{F_n^{reg}}, \quad (2.26)$$

in which the constants are given by

$$\begin{aligned} A_2 &= \eta \iiint |\mathcal{K} \zeta_n|^2 d\vec{x}_s d\vec{x}_r d\omega, \\ A_1 &= -2 \eta \Re \left(\iiint r_{n-1}^* [\mathcal{K} \zeta_n] d\vec{x}_s d\vec{x}_r d\omega \right), \hat{\mathbb{E}} \\ A_0 &= \eta \iiint |r_{n-1}|^2 d\vec{x}_s d\vec{x}_r d\omega = F_{n-1}^{data}, \hat{\mathbb{E}} \\ B_2 &= \|b_n \nabla \zeta_n\|_{\mathcal{D}}^2, \hat{\mathbb{E}} \\ B_1 &= 2 \langle b_n \nabla \chi^{(n-1)}, b_n \nabla \zeta_n \rangle_{\mathcal{D}}, \hat{\mathbb{E}} \\ B_0 &= \|b_n \nabla \chi^{(n-1)}\|_{\mathcal{D}}^2 + \delta_{n-1}^2 \|b_n\|_{\mathcal{D}}^2. \end{aligned} \quad (2.27)$$

For a given $\chi^{(n-1)}$ this polynomial should be minimised with respect to the step size α_n . For this purpose we determine the derivative of F_n^{tot} with respect to α_n and set the result equal to zero. The root of the obtained third-order polynomial is obtained analytically with the help of Cardano's formula [Abramowitz and Stegun (1970)].

In conclusion, we want to emphasise again that the regularisation scheme does not require any user interaction and that all the necessary parameters are determined during the process itself. Hence, the regularisation scheme can be seen as data-driven.

2.3.4 The importance of regularisation: A synthetic example

To show the indispensability of regularisation of the inversion process of seismic data, we perform an experiment. We model linear seismic data for 21 sources and 21 receivers with a source/receiver increment of 12.5 m. This acquisition set-up was placed 25 m above a scattering domain of 21 by 21 grid cells with a cell spacing of 5 m. Using a temporal sampling of 4 ms a total record length of 512 ms was computed. No tapering was applied. Note that linear data assumes that $p_{tot} = p_0$ in equation (2.1) and that multiple scattering effects are neglected. However, for the current exercise this is acceptable, because we want to focus on the necessity of regularising the inversion process when noise is present in the data.

In the following we invert the modelled data with and without regularisation, and for different noise levels. Note that the noise levels are always given with respect to the average energy of the data. The inversion results can be found in figure 2.1. We see that in the ideal case of linear data without any added noise, a linear inversion leads to a nearly perfect reconstruction of the initially unknown property model. It unfortunately this ideal scenario never occurs in real life. Usually data are nonlinear and contains a varying amount of random noise. By modelling linear data we separate these issues to investigate the effect of random noise on the inversion result only.

We started by adding only 1% of random noise to the data and a linear inversion without regularisation leads to a rather distorted property image. The reason is that the inverse problem is ill-posed, meaning that several property realisations exist that all match the measured data within the same misfit. The same holds if noise is present, meaning that the inversion result explains the data but also partly the random noise at the same time.

It needs to be mentioned that for the noise-free example we stop the inversion automatically after a limited number of iterations. If we would let the inversion run for a very large number of iterations, eventually the scheme might diverge even in the noise-free case. Still, even if no noise is added, several property realisations exist that explain the measured data within the precision of the computer system used. This numerical noise would eventually be picked up by the inversion leading to unreasonable results.

If we increase the noise level to 10% the effect gets even worse and the original properties cannot be reconstructed at all. Since random noise in the data introduces random noise in the inverted property model we can use a regularisation to tackle this problem. The multiplicative regularisation as it was described earlier aims at finding a property model that is rather continuous and without random spikes. Although the inversion results as shown in figure 2.1 (c) and (d) described the data, they would not be accepted by the regularisation because of their high

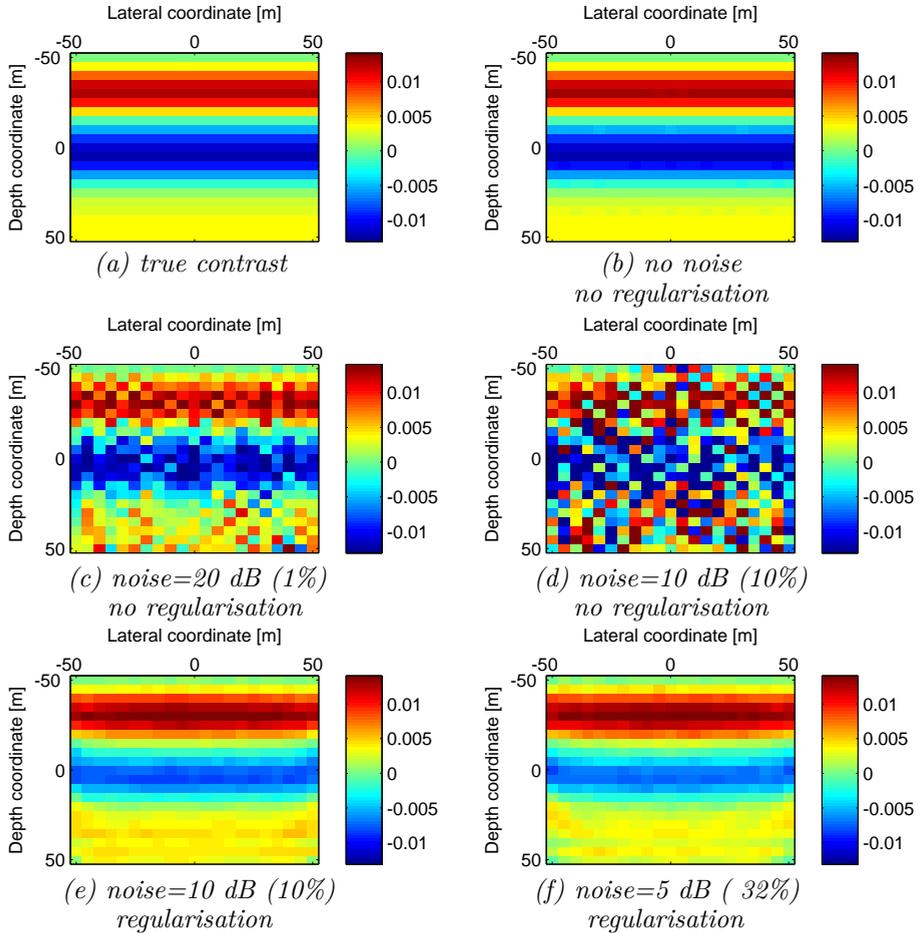


Figure 2.1: Inversion results for different noise levels and with or without regularisation of the inversion. In (a) the true contrast is shown. Linear and noise-free data can be inverted without regularisation and the original contrast is successfully reconstructed (b). As soon as noise is present in the data (c)+(d) the inversion process fails if no regularisation is included. Using the proposed regularisation scheme we are still able to reconstruct the original contrast even for a significant amount of noise in the data (e)+(f). Note that the noise levels are always given with respect to the average energy of the data.

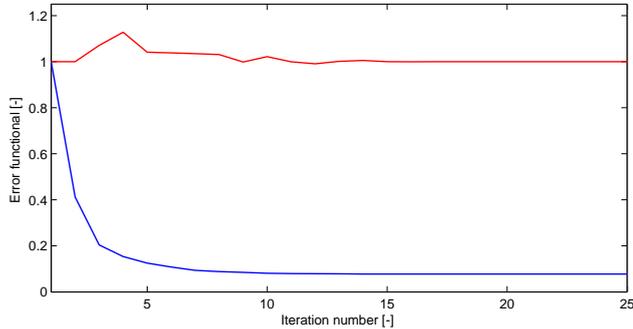


Figure 2.2: *Convergence of the data functional (blue) and the regularisation functional (red) as function of the iteration number. A remarkable characteristic of multiplicative regularisation is that it is only active in the beginning of the conjugate gradient scheme. With increasing iterations it converges to one, indicating that finally the original cost functional, based on the data misfit, will be minimised only.*

variation. It can be seen that if we repeat the inversion of data with 10% noise, but using now the multiplicative regularisation, a very good estimate of the true properties is obtained. The same even holds for 32% added noise, although such a large amount of random noise is generally not realistic to be found in field data. Nevertheless, it demonstrates the robustness of multiplicative regularisation and its ability to distinguish and reject noise from the data which contains information about the object to be recovered.

2.3.5 Convergence of multiplicative regularisation

Before concluding this section we would like to discuss the convergence behaviour of the multiplicatively regularised conjugate gradient scheme a bit more in detail. The total error functional that is supposed to be minimised and which was given in equation (2.20), consists of two parts: The data functional and the regularisation functional. To get a better understanding of how minimisation of these error functionals converges in the conjugate gradient scheme, we plot both factors as a function of the iteration number and for a noise level of 32% in figure 2.2.

First, it can be observed that the data functional converges nicely towards a specific value, related to the noise level in the data, while the regularisation functional is varying during the first couple of iterations but converges to one at a later stage. This behaviour can be explained in the following way. During the very first iterations, an inversion scheme has difficulties to distinguish between signal and noise in the measured seismic data. If no regularisation would be used, the outcome of full waveform inversion would be a contrast function that explains

the seismic signal as well as the noise. Unfortunately such synthetic subsurface models exist, although they might turn out to be quite meaningless from a geological point of view. We have gained this experience in the previous experiment where truly unregularised inversion of seismic data including variable amounts of noise lead to virtually random contrast functions.

The situation changes when multiplicative regularisation is used. The regularisation functional, as defined in equation (2.21), becomes dominant for large changes of the spatial gradient of the contrast function between two subsequent iterations. This behaviour occurs especially during early iterations when large spatial gradients, or differently expressed, spiky models are needed to explain the signal and the noise. Multiplicative regularisation rejects these spiky models, because it is designed to allow locally smooth models only. This has the effect that the best MR-CG can do to minimise the total cost functional, is to find a subsurface model that explains the signal only. Metaphorically speaking, the regularisation is supposed to guide the inversion scheme to a minimum that explains the data but not the noise. In figure 2.2 we can identify this guidance by the fact that regularisation functional is larger than one for the first several iterations. After the scheme converged to the desired minimum the regularisation term becomes one. This means that eventually the data functional will be minimised only while the regularisation prevents the scheme from picking up the random noise.

To support the current discussion we show the data residual corresponding to the inversion results displayed in figure 2.1 and for a single source in figure 2.3. It is interesting to see that in each scenario the measured data was fully explained by the conjugate gradient scheme but that the added random noise was rejected. Evaluation of the inversion results shows major differences between the cases when multiplicative regularisation was used and when the inversion was unregularised from the beginning.

2.4 Nonlinear field update based on the domain equation

So far we have described a multiplicatively regularised conjugate gradient scheme that allows us to recover the properties of the subsurface from measured seismic surface data. However, the CG scheme described performs a linear inversion only, meaning that the nonlinear relationship between the subsurface properties and the measured data is not accounted for. The nonlinearity in seismic data is represented by the coupling of the data equation and the domain equation. Solving the data equation by a conjugate gradient scheme is an iterative process during which the contrast function is updated until the measured data are explained. However, during this process it is neglected that every time when the contrast function changes also the total field in the data equation should be updated. This is why solving the data equation with a conjugate gradient scheme leads to

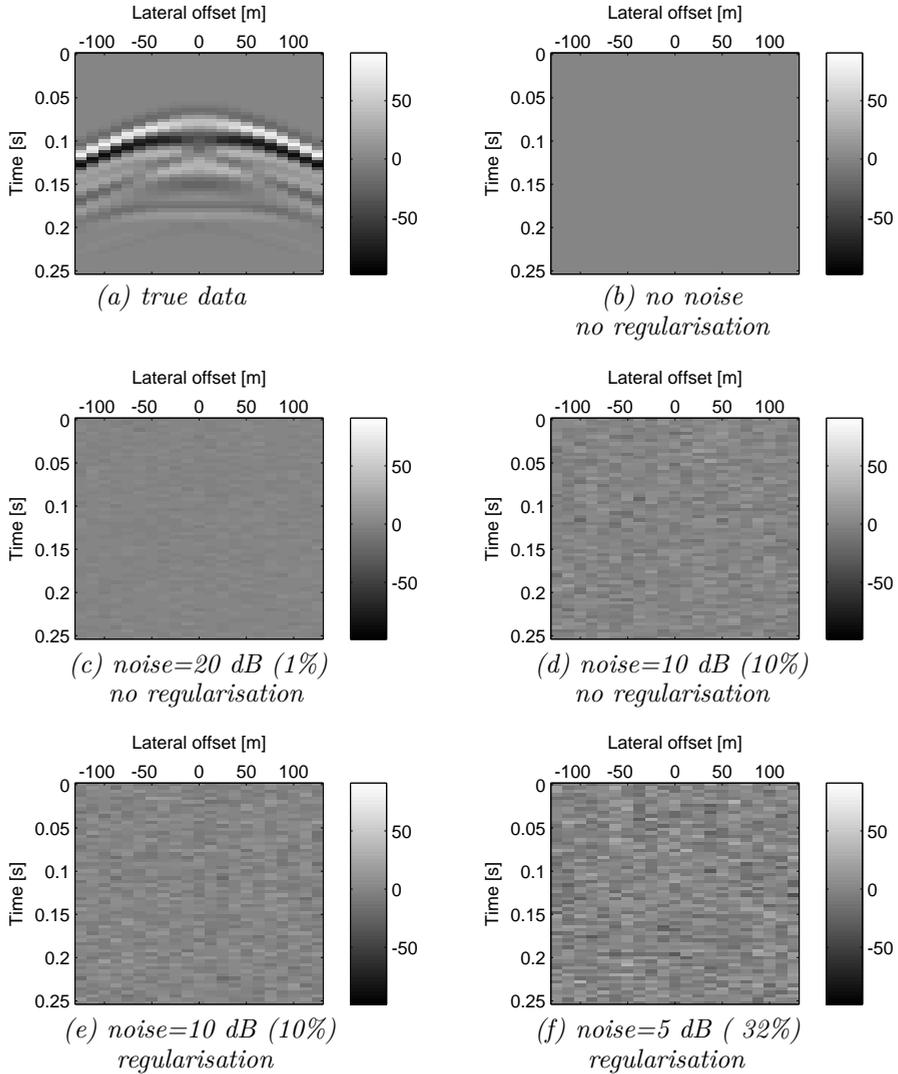


Figure 2.3: While the input data are given in (a), the residual energy for several experiments, differing in the amount of added noise and the use of regularisation or not, is shown in (b)-(f). Interestingly, the entire signal was explained by the inversion, independent of using regularisation or not. Still, the regularisation is necessary to avoid the inversion picking up the noise and previous results have shown that multiplicative regularisation is very effective in doing so.

an approximate solution of the subsurface parameter as long as the total field in the true medium is not known exactly. On the other hand it is essential to realise that if the total field would be known exactly, the relation between the subsurface properties and the measured seismic data becomes truly linear and a CG scheme should be able to recover the exact properties from the measured data.

In this section we explain how a total field update can be performed to incorporate nonlinearity in the inversion scheme. Since initially the subsurface properties are unknown, the very first inversion is always based on the assumption that wavefield propagation occurs in a smooth non-scattering background medium only. This means that the total field in equation (2.1) is approximated by the background field, hence $p_{tot} \approx p_0$. This is also called the Born-approximation and as long as it is used, nonlinear effects as multiple scattering and transmission in the subsurface are neglected. Furthermore, the inversion result will be based on the travel times in the background model and any possible differences between these travel times and the exact travel times in the true medium cannot be accounted for.

Assuming wave propagation in a smooth background medium that reasonably explains the kinematics of the subsurface under investigation will lead to an approximate inversion result. This initial inversion result, although coming from a linear data model, will show structural information that is expected to be similar to what would be shown by migration. This can be explained by the fact that existing migration schemes are based on a linear approximation as well. In Berkhout and Verschuur (2011) and Berkhout (2012) promising efforts were made to extend existing migration techniques by incorporating nonlinear effects as multiple scattering and transmission.

Using an approximate smooth property model of the subsurface immediately tells us that the wavefield propagation in such an approximate medium cannot be exact and that for exact results we should make the wavefields consistent with the currently best estimate of the true medium.

Since the output of linear full waveform inversion is a property model, we do not only obtain structural information but also a first order approximation of the subsurface properties at every grid point in the inversion domain. We can now use this property model to update the total field by solving the domain equation with in principle any suitable numerical method. We discussed earlier in this chapter why we do not use the most prominent modelling technique based on finite differences. Instead we iteratively build up the total fields as a sum of the background field and a number of basis functions:

$$p_{tot}^{(N)}(\vec{x}, \vec{x}_s, \omega) = p_0(\vec{x}, \vec{x}_s, \omega) + \sum_{n=1}^N \alpha_n^{(N)}(\vec{x}_s, \omega) \phi_n(\vec{x}, \vec{x}_s, \omega), \quad (2.28)$$

where $\alpha_n^{(N)}$ are frequency dependent weighting factors and depend on the number N of basis functions taken into account. If the weighting factors are all equal to one, equation (2.28) becomes the Neumann series which is known to be unstable for large contrast, high frequencies or large scattering domains. Note that N is equivalent to the iteration number and for each iteration one basis function is added and a full set of new weighting factors has to be calculated. The basis functions are defined as

$$\phi_n(\vec{x}, \vec{x}_s, \omega) = \int_{\vec{x}' \in \mathcal{D}} \mathcal{G}(\vec{x}, \vec{x}', \omega) \partial \mathcal{W}_n \, d\vec{x}', \quad (2.29)$$

and the incremental contrast sources $\partial \mathcal{W}_n$ are given by

$$\partial \mathcal{W}_1 = \chi^{(1)} p_{tot}^{(0)}, \quad \partial \mathcal{W}_n = \chi^{(n)} p_{tot}^{(n-1)} - \chi^{(n-1)} p_{tot}^{(n-2)}, \quad n > 1. \quad (2.30)$$

These contrast sources play a fundamental role in the integral formulation of the wave equation. By introducing them, every point in the subsurface becomes a scatterer that starts interacting with any other point. These point scatterers can be seen as secondary sources because they represent the scattering of the total field at every grid point which then becomes a new point source. It can be realised that for each iteration a higher order of scattering is included in the contrast sources and thereby in the basis functions. While \mathcal{W}_1 describes a first order interaction between two points in the subsurface, higher order terms are subsequently included in the contrast sources \mathcal{W}_n . It is important to realise that contrast sources are a purely mathematical concept and that they are no real physical sources. Nonetheless, it turns out that they are an exact way of describing the true wavefield propagation in the subsurface and that they are advantageous in the sense that they allow us to include higher order scattering terms iteratively. This would not be possible by finite-difference or finite-element methods, since they intrinsically solve for an infinite order of scattering.

It remains to determine the weighting factors $\alpha_n^{(N)}$ in order to allow us to update the total fields. To this purpose we substitute equation (2.28) into equation (2.3) and including the dependency of χ on the iteration number n leads to

$$\sum_{n=1}^N \alpha_n^{(N)}(\vec{x}_s, \omega) [\mathcal{L}_N \phi_n](\vec{x}, \vec{x}_s, \omega) = [(\mathcal{I} - \mathcal{L}_N) p_0](\vec{x}, \vec{x}', \omega), \quad (2.31)$$

where \mathcal{I} is the identity operator and where the operator $\mathcal{L}_q \phi_m$ was defined as:

$$[\mathcal{L}_q \phi_m](\vec{x}, \vec{x}_s, \omega) = \phi_m(\vec{x}, \vec{x}_s, \omega) - \int_{\vec{x}' \in \mathcal{D}} \mathcal{G}(\vec{x}, \vec{x}', \omega) \chi^{(q)}(\vec{x}') \phi_m(\vec{x}', \vec{x}_s, \omega) d\vec{x}'. \quad (2.32)$$

Equation (2.31) contains the unknown weighting factors $\alpha_n^{(N)}$, the basis functions ϕ_n and the background field p_0 . The unknown weighting factors can now be

calculated by minimising the difference between the left- and right-hand side of equation (2.31) in the least-squares sense. With $m = 1, 2, \dots, N$ this leads to the following $N \times M$ system of equations:

$$\begin{aligned} \sum_{n=1}^N \alpha_n^{(N)}(\vec{x}_s, \omega) \int_{\vec{x} \in \mathcal{D}} [\mathcal{L}_N \phi_m](\vec{x}, \vec{x}_s, \omega) [\mathcal{L}_N \phi_m]^*(\vec{x}, \vec{x}_s, \omega) d\vec{x} \\ = \int_{\vec{x} \in \mathcal{D}} [(\mathcal{I} - \mathcal{L}_N)p_0](\vec{x}, \vec{x}', \omega) [\mathcal{L}_N \phi_m]^*(\vec{x}, \vec{x}_s, \omega) d\vec{x}. \end{aligned} \quad (2.33)$$

Least-squares estimation of α ensures that the error in the total field update as described in equation (2.28) cannot diverge. This is an improvement of the proposed updating procedure compared to the straight-forward Neumann series or finite-difference methods which can suffer from instability issues, especially at higher frequencies. It is important to realise that each time a new basis function is added to the total field, N is increased by one and a complete new set of optimised $\alpha_n^{(N)}$ is calculated. Note that if the contrast is not changing, i.e. $\chi(\vec{x}) = \chi^{(n)}(\vec{x})$, the scheme becomes equivalent to the Krylov subspace method. A detailed description of the Krylov method can be found in Kleinman and van den Berg (1991). Then convergence under certain weak conditions of the operator can be proved. By iteratively performing total field updates, the data model used for inversion becomes more and more accurate. Finally, the total field at each grid point includes all multiple scattering and transmission that occur in the inverted subsurface model. Even the travel times will be honoured by the inversion if the kinematic behaviour of the inverted property model is different from the initial smooth background. Although having a reasonable background model to start with is still important, we will see later that the nonlinearity allows us to recover also spatial wave numbers of the initially unknown subsurface model that are outside the equivalent temporal bandwidth that is present in the seismic data.

2.5 Comments on background models

We discuss a full waveform inversion technique that is based on the integral formulation of the wave equation. This requires the computation of wavefields and Green's functions in smooth non-scattering background models. Current implementation of the 2-D scheme as will be demonstrated in this thesis, allows the use of homogeneous background models only. In a homogeneous background Green's functions can be calculated analytically, simplifying calculation of the entire scheme. We are fully aware that for real data applications the use of smooth inhomogeneous backgrounds is inevitable. This can be a rather expensive task, and effective modelling of Green's functions in inhomogeneous media is necessary. An interesting effort was made by Fokkema and van den Berg (2012) for medium stretching to effectively make it homogeneous. Application of this method

to a smooth background medium would then allow direct use of our current implementation. Another promising method to incorporate a smooth background medium is given by Abubakar et al. (2009). The WKB-approximation can also be used if smooth background media should be considered (Bremmer, 1951).

2.6 Outlook

In principal, alternating linear inversion by a CG scheme followed by nonlinear global field updates could now be used to obtain high-resolution property models by nonlinear full waveform inversion. Unfortunately, full waveform inversion in 2-D and especially in 3-D requires a huge amount of computer resources. To be more precise, the conjugate gradient scheme needs a rather large amount of physical memory because it has to store the kernel \mathcal{K} to model data based on a given property model. The size of this kernel grows rapidly with the number of samples in the input data and the number of unknowns in the discretised inversion domain. With large seismic datasets and the huge amount of unknowns in a finely discretised subsurface model, which is needed if high-resolution is aimed for, the size of the kernel becomes so large that the inversion is not feasible on today's computer systems. This statement surprisingly even is true if we consider huge computer clusters as they are used in the industry today. In the next chapter we will investigate the computational requirements of the CG scheme a bit more in detail and demonstrate a solution how these requirements can be reduced significantly. As a result, high-resolution full waveform inversion while using the full bandwidth of the seismic data should become feasible.

Localised linear full waveform inversion

In this chapter we will investigate the computational resources that are needed to perform a linear inversion of the data equation using a CG scheme. Here linear refers to the fact that the total field in equation (2.1) is kept fixed during the inversion. In chapter 4 we will discuss a total field estimation procedure to fully incorporate the nonlinear relationship between the measured data and the subsurface properties. Still, inversion of the data equation will always be performed using a linear CG scheme. In this chapter we use the background field as approximation of the total field.

We will see that current computer systems cannot provide the capacity to perform full waveform inversion on a high resolution grid. This also implies that the oil exploration industry cannot use the full bandwidth of the acquired signal, because on a relatively coarse grid, only the low temporal frequencies can be inverted. We will demonstrate how this limitation can be overcome by localising the inversion process. The scheme is based on backpropagation of sources and receivers into the subsurface and subsequent time windowing to separate the part of the backpropagated data that belongs to a local domain from the part that is associated with the surrounding. A local domain is only a subset of the entire subsurface, and supposing we can neglect the surrounding of the local domain for the moment, this inversion problem is much smaller. We can use the reduction of the problem size and invest the available resources in a much finer sampling of the inversion domain, which at the same time allows us to incorporate much higher frequencies than is possible in other schemes. The local inversion results can then be combined to a global property model yielding a high resolution property model of the subsurface. In figure 3.1 a flow chart visualises which building blocks of this

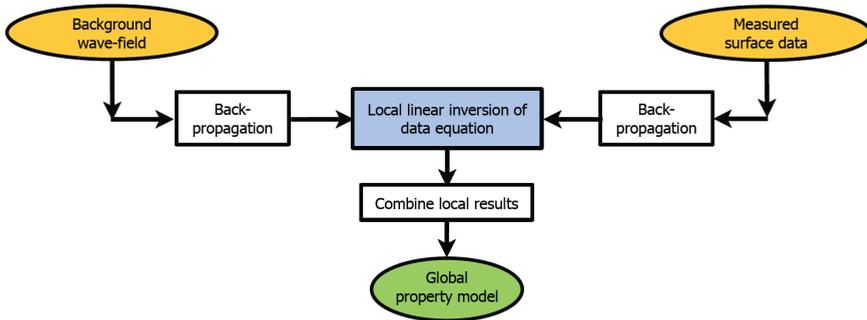


Figure 3.1: *By backpropagation of sources and receivers, a new acquisition is simulated at any arbitrary location in the subsurface. If only a limited time window of the back-propagated data is used, local inversion on a high resolution grid can be performed while using the full seismic frequency spectrum. To obtain a global property model, several local inversion results need to be combined. This concept is highly suitable for parallel implementation on a multi core computer cluster because the local inversions are fully independent of each other.*

linear inversion scheme will be discussed. Note that the processes in the white boxes of figure 3.1 are new aspects that will be introduced in this chapter.

3.1 Computational requirements of a CG scheme

The major computational requirements of a conjugate gradient scheme are related to the modelling kernel, containing the total fields from each source to any point in the subsurface combined with the Green's functions between every grid point and the receivers. The kernel \mathcal{K} was first introduced in equation (2.6) and can explicitly be written as

$$\mathcal{K}(\vec{x}_r, \vec{x}', \vec{x}_s, \omega) = \mathcal{G}(\vec{x}_r, \vec{x}', \omega) p_{tot}(\vec{x}', \vec{x}_s, \omega). \quad (3.1)$$

The kernel contains the total fields and the Green's functions before multiplication with the contrast function and before integration over the subsurface domain, it is a four-dimensional function. It depends on the frequency, the source locations, the grid locations in the scattering object, and finally the receiver locations. Let us investigate how large this kernel would be if we want to obtain a 3-D high resolution subsurface model from a seismic dataset that covers the entire surface above this domain.

A realistic seismic survey could easily cover a surface domain of $5 \times 5 \text{ km}^2$ with a source and receiver sampling of 25 m in each lateral direction. Note that in real life full spatial coverage is hardly available and especially the sources are often

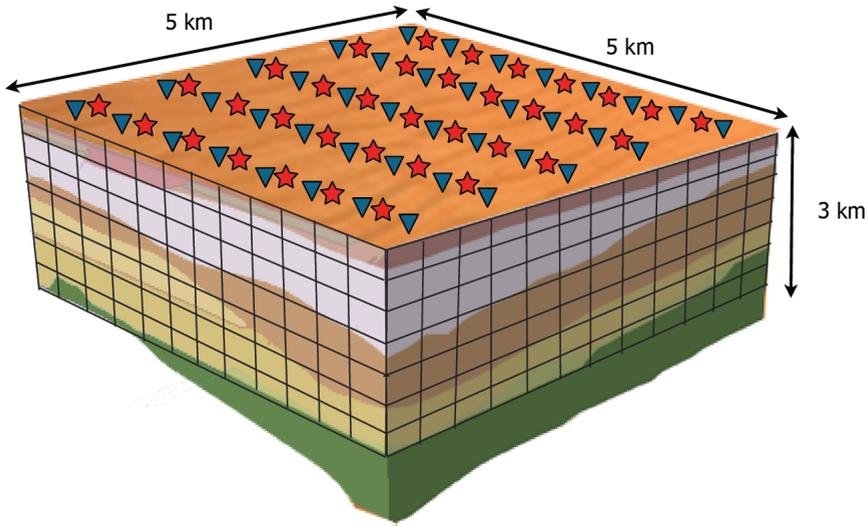


Figure 3.2: *In an ideal case for 3-D inversion, the full surface would be sampled with sources and receivers. From these measurements, the properties of a discretised subsurface model are supposed to be recovered.*

sparingly sampled in one lateral direction, but for the sake of simplicity we assume the perfect situation. This means that in total the dataset consists of 40000 sources and the same number of receivers. Each receiver measures 2048 ms of data, being equivalent to 512 data points if a sampling of 4 ms is used. From this seismic dataset we want, in principle, to recover the full set of elastic properties on a high resolution grid, but we limit ourselves to the acoustic velocity only. If this parameter should be recovered down to a depth of 3 km, the 3-D inversion domain has the extension of $5 \times 5 \times 3 \text{ km}^3$. Our ambition is to get the properties of this cube on a $5 \times 5 \times 5 \text{ m}^3$ grid, resulting in 6×10^8 unknown velocity values. A cartoon of a fully covered surface acquisition including a discretised subsurface model is shown in figure 3.2.

Since the kernel is a four-dimensional function, its total size is determined by multiplying the size of each dimension. In our case the total size is equal to the number of data points times the number of unknowns. We do this calculation for our discussed 3-D example in table 3.1. It turns out that for the discussed 3-D inversion problem the kernel would have a size of $4 \cdot 10^{12} \text{ Gb}$. This multidimensional tensor has to be stored in the physical memory of the computer system, because it is used several times during each CG iteration. Realising that every inversion usually takes several hundred CG iterations, re-computation instead of storing

×	40000	number of sources
×	40000	number of receivers
×	512	measured time samples
×	$6 \cdot 10^8$	number of unknowns
<hr/>		
=	$5 \cdot 10^{20}$	samples of kernel
×	8 bits	per sample
<hr/>		
=	$4 \cdot 10^{12}$ Gb	total size of kernel

Table 3.1: *Calculation of the size of kernel \mathcal{K} for a realistic 3-D inversion problem aiming at recovering the acoustic velocity on a high resolution grid. For elastic inversion, the size would increase even more. It should be obvious that required memory to store the kernel is not available today and also will not be in the near future. This makes broadband high resolution full waveform inversion currently impossible.*

this huge kernel is not an option, unless extremely fast quantum computers are available. If we repeat the estimation of the kernel size for a 2-D acoustic inversion, assuming sources and receivers along a line only and skipping the third spatial dimension of the inversion domain, we end up with a size of 10^5 Gb.

The purpose of this exercise is simply to demonstrate that full waveform inversion is computationally an extremely expensive method and that current computer systems cannot provide the physical memory to deal with high resolution subsurface models by this technique. Due to the steadily increasing computer power, just in recent years full waveform inversion is being explored in the oil exploration industry as a tool to derive low resolution velocity models of the subsurface. To make this method feasible, only the low frequencies of the measured seismic data are used and inverted on a coarse grid. The outputs are smooth models that are mainly used as input to a migration scheme. Current schemes cannot deliver a true 2-D or 3-D solution that would be suitable for direct interpretation. This means that at the end of the seismic data processing chain, a structural image is used for interpretation and subsequent decision making. Although full waveform inversion is known to be able to improve these structural images we want to use this method for high resolution property estimation in the near future.

It needs to be mentioned that high resolution inversion of migrated data (pre-stack) is commercially available. If the 1.5-D assumption is made, the migrated data can be de-migrated to subsequently invert full waveform time-domain data.

Nonetheless, true 3-D or even 2-D full waveform inversion at reservoir resolution that does not make an implicit 1.5-D assumption is currently not on the market.

3.2 Localisation by source/receiver-refocusing

From the previous section it should be clear that it is inevitable to find a way of reducing the size of the inversion problem, if high-resolution full waveform inversion should be performed. This fact is certainly realised by the seismic community and major efforts are made by several researchers to achieve this task. Although good progress was made during recent years, most methods aim to reduce the amount of data that has to be inverted. In this thesis we demonstrate a new method that combines the entire surface acquisition smartly into a local backpropagated dataset at an arbitrary location in the subsurface. Beside the benefit of reducing the data volume significantly, it also allows us to invert for the properties of a limited subdomain of the entire subsurface only. Because for each local inversion we can neglect the rest of the subsurface, the size of the inversion domain is reduced and the overall problem size is reduced enormously.

To localise the inversion process we need to backpropagate the seismic surface data into the subsurface. Backpropagation, or inverse wavefield extrapolation as it is often called, is based on the Kirchhoff integral (Kirchhoff, 1883). The Kirchhoff integral states that from measurements of a wavefield, that obeys the scalar wave equation, along a closed surface, the wavefield can be predicted at any arbitrary location within this domain. Seismic measurements are generally carried out along the earth's surface only and the requirement of having measurements along a closed surface is never met. However, in case of one-sided illumination, Kirchhoff's integral can be re-written as an integral along an infinite plane surface only. This is called the Rayleigh integral (Rayleigh, 1896).

With the help of the Rayleigh integral, the backpropagation operators in the frequency domain can be defined as

$$[\mathcal{F}_r p_{data}](\vec{x}_{br}, \omega) = 2 \int_{\vec{x}_r \in \mathcal{S}} \frac{\partial}{\partial z_r} \mathcal{G}^*(\vec{x}_{br}, \vec{x}_r, \omega) p_{data}(\vec{x}_r, \omega) d\vec{x}_r, \quad (3.2)$$

$$[\mathcal{F}_s p_{data}](\vec{x}_{bs}, \omega) = 2 \int_{\vec{x}_s \in \mathcal{S}} \frac{\partial}{\partial z_s} \mathcal{G}^*(\vec{x}_{bs}, \vec{x}_s, \omega) p_{data}(\vec{x}_s, \omega) d\vec{x}_s, \quad (3.3)$$

where \mathcal{S} is in that infinite seismic acquisition plane and where \vec{x}_r and \vec{x}_s stand for the actual source and receiver positions, while \vec{x}_{br} and \vec{x}_{bs} denote the back-propagated source/receiver positions. The anti-causal Green's function between the locations of the backpropagated sources/receivers in the subsurface and the locations of the actual sources/receivers located along the surface is denoted by

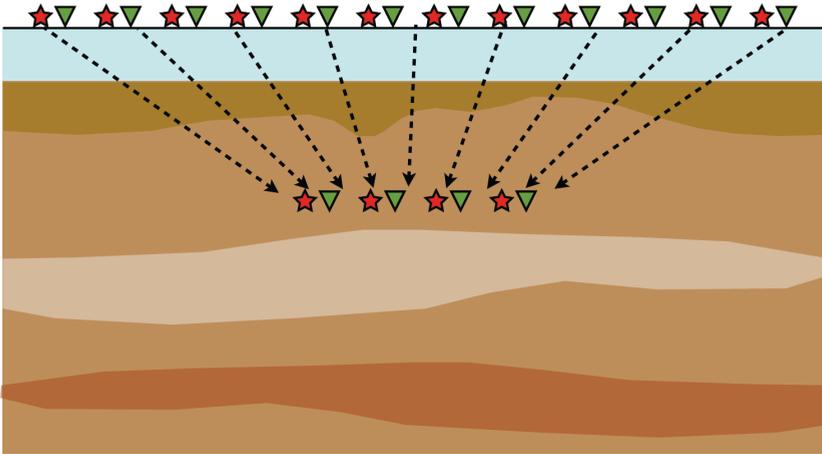


Figure 3.3: *By backpropagating the surface acquisition, a local dataset can be generated at an arbitrary location in the subsurface.*

\mathcal{G}^* . If the anti-causal Green's function is used, the process reduces the travel times of the seismic data and the procedure is referred to as inverse wavefield extrapolation. If the causal Green's function would be used, travel time would be added by the process and the procedure would be called forward wavefield extrapolation. In equations (3.2) and (3.3) the backpropagation operators are defined for sources and receivers separately. Backpropagation is applied first to the sources and a half-redatumed gather has a source located in the subsurface while the receivers are still along the surface. In a second step also the receivers are backpropagated in the same manner. It can be seen that if sources and receiver are positioned at the same locations along the surface, the operators become the same. By backpropagation we can generate a completely new acquisition at an arbitrary location in the subsurface from the seismic data measured at the surface. This procedure is sketched in figure 3.3.

Although we describe the method for surface seismic data, the same concept can in principle also be applied to other acquisition geometries, e.g. seismic well data. Still, a local backpropagated dataset has to be generated above the domain of interest and such a technique was recently proposed for VSP data by Soni et al. (2012).

In this thesis we use backpropagation not only to generate a seismic dataset in the subsurface but we also apply it to the total fields as part of the data equation. This means we apply a dual backpropagation (sources and receivers) to both sides of equation (2.1). So far we approximated the total field by the background field

and we will keep doing this in a synthetic experiment in this chapter. Nonetheless, the concept of localising full waveform inversion by backpropagation will be equally valid after applying the total field estimate, that will be discussed in the next chapter.

After carrying out the backpropagation we obtain the data equation with backpropagated sources located at \vec{x}_{bs} and backpropagated receivers located at \vec{x}_{br} :

$$\hat{p}_{sct}(\vec{x}_{br}, \vec{x}_{bs}, \omega) = \int_{\vec{x}' \in \mathcal{D}} \hat{\mathcal{G}}(\vec{x}_{br}, \vec{x}', \omega) \chi(\vec{x}') \hat{p}_{tot}(\vec{x}', \vec{x}_{bs}, \omega) d\vec{x}', \quad (3.4)$$

where $\hat{p}_{sct} = [\mathcal{F}_r[\mathcal{F}_s p_{data}]]$, $\hat{\mathcal{G}} = [\mathcal{F}_r \mathcal{G}]$ and $\hat{p}_{tot} = [\mathcal{F}_s p_{tot}]$. Assuming that the backpropagation operators are unique, equation (3.4) is equivalent to equation (2.1), with the difference that the backpropagated sources and receivers are located within the subsurface domain \mathcal{D} instead along the surface domain \mathcal{S} .

Applying the backpropagation to the left- and the right-hand side of the data equation has several beneficial implications. First, the whole procedure becomes effectively a preconditioning operation and no model error is made. The same reasoning tells us that we can replace the continuous integral in equations (3.2) and (3.3) by a discrete summation without introducing a model error. Since seismic data are inevitably measured at discrete locations this is desired and cancels out the acquisition imprint generally introduced by backpropagation of seismic data only. Second, although the backpropagation operators are defined based on the exact anti-causal Green's function \mathcal{G}^* , we can use an approximate Green's function, again without making any model error. The key still is to apply backpropagation to both sides of the data equation. In this way we overcome the need to know the exact Green's function between the true source/receiver location and the point in the subsurface to which backpropagation should be performed. Instead we can use operators based on travel times and approximate amplitudes while still keeping the data equation exact. It is well known that exact backpropagation is a nontrivial task and in this way this exact processing step is superfluous.

We assume that, to calculate the backpropagation operators, a background velocity model exists, that describes the kinematics of the subsurface under investigation reasonably well. This background model could come from standard velocity analysis or from full waveform inversion on a coarse grid, as routinely used by existing methods.

3.2.1 Backpropagation of the seismic surface data

In the following we investigate how the data and the total fields change when dual backpropagation is applied. For demonstration we make use of a horizont-

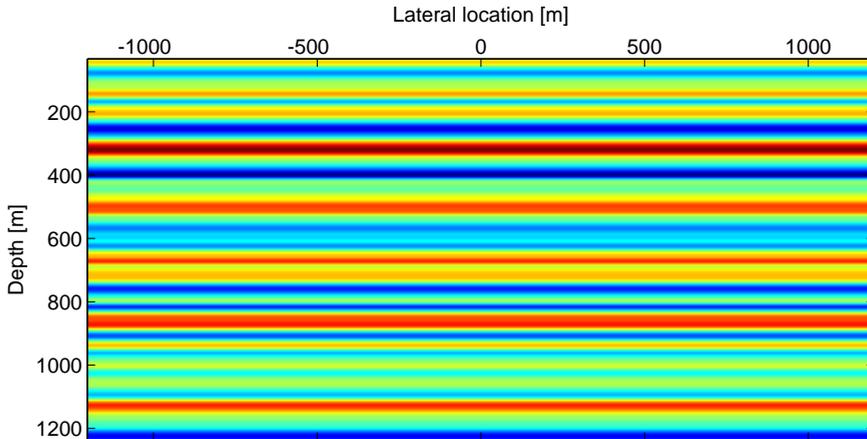


Figure 3.4: Contrast function used for modelling of linear data. A homogeneous background medium with $c=2000$ m/s and $\rho = 2000$ kg/m³ was used.

ally layered subsurface model above which linear primary data are generated by replacing p_{tot} by p_0 in equation (2.1). This means that no multiples are generated and that transmission effects due to scattering are neglected. At this point we want to exclude any nonlinear effects, in order to fully focus on the ability to localise seismic full waveform inversion by backpropagation. For modelling seismic data we assume a homogeneous background medium with $c = 2000$ m/s and $\rho = 2000$ kg/m³. The contrast function $\chi(\vec{x})$ is shown in figure 3.4.

Linear data was computed for 121 sources and 121 receivers located along the surface of the model. The source/receiver spacing was chosen to be 20 m while a total length of 512 samples with a 4 ms temporal sampling was acquired. The subsurface domain \mathcal{D} is defined on a 5 m grid in the vertical as well as in the horizontal direction. A selection of the surface data is displayed in figure 3.5.

In a next step we backpropagate the entire surface dataset to a depth level of $z = 600$ m while the operators are calculated in the homogeneous background medium. A subset of the backpropagated data is shown in figure 3.6. It has to be realised that all 121 sources and 121 receivers were combined to bring the dataset from the surface to the new depth level.

While seismic surface data are intrinsically causal and, consequently, has positive travel times only, backpropagated data can be divided into a causal and an anti-causal part. While the causal part consists of reflections that were generated below the depth level to which the data was backpropagated, the anti-causal part

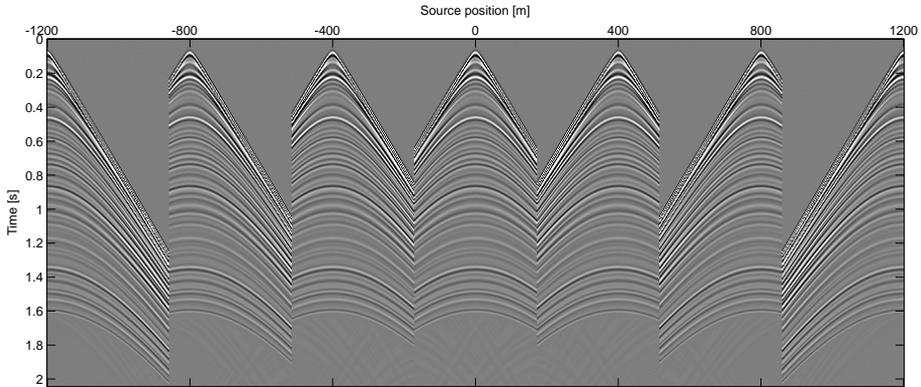


Figure 3.5: A shot selection of the linearly modelled primary data above the horizontally layer velocity model shown in figure 3.4. Seismic data as commonly acquired with sources and receivers along the surface contains causal events only. Hence, all reflections generated in the subsurface arrive at positive times.

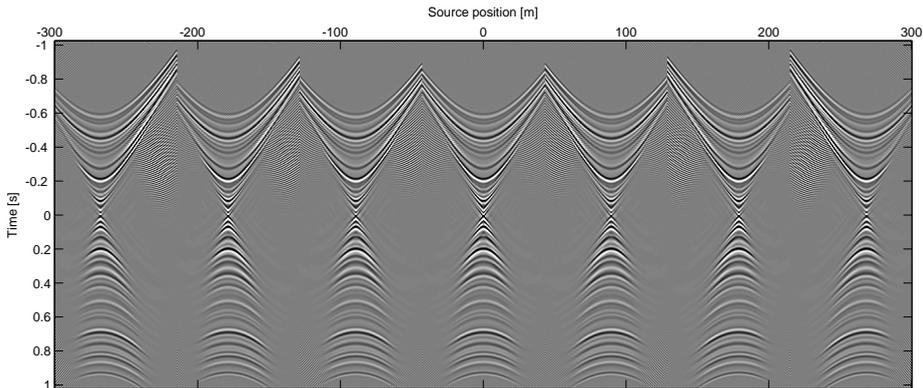


Figure 3.6: A selection of shots after sources and receivers were backpropagated to a depth level of $z=600$ m. It can be seen that the data after backpropagation can be divided into a causal and an anti-causal part. The causal part contains reflections from the subsurface below the depth level to which the data was backpropagated while the anti-causal part contains reflections from the subsurface above. It must be realised that this statement holds for linear data but in real life, multiples generated in the overburden may appear in the causal part as well. Because these multiples are nonlinear wave phenomena we will treat them later in this thesis.

shows reflection generated above this depth level. It should be mentioned that this statement holds only if seismic data are assumed to be truly linear, hence it consists of primaries only. If seismic data are nonlinear in the subsurface properties, multiples can be generated in the overburden between the target depth level and the surface, and these multiples might still show up in the causal part after backpropagation. Although at this point we are not concerned with nonlinearity, we will treat this special case at the very end of this thesis, after nonlinear total field estimation has been discussed.

If we assume for the moment that the linear assumption holds, the fact that backpropagated data can be split into a causal and an anti-causal part can be utilised to localise seismic full waveform inversion. If we use the causal part of the backpropagated data we can invert for parameters of the subsurface below the target depth level only while the overburden can be neglected, and vice versa, if the anti-causal part would be used. As a result we do not need to invert for the entire subsurface at once but we can reduce the size of the inversion domain by using a specific part of the backpropagated data.

Our aim is to reduce the inversion domain to a local domain with extensions of several hundreds of meters in each direction. So far we have explained that by backpropagation and choosing the causal or the anti-causal part, the subsurface can be divided into a lower and an upper part. To achieve also a limitation of the inversion domain with respect to the lateral direction we can choose only a subset of the backpropagated data. Here we make use of the fact that if only limited offsets are used, only the part of the subsurface between the first and last source/receiver is illuminated. Again, this statement will not be fully valid if nonlinearity plays a role or if the subsurface is not truly horizontally layered, but we will treat these cases later.

Finally, to obtain a local inversion domain, we apply a time window to the data to limit the inversion domain to a local depth range. Here we make use of causality again because we know that within a limited time wavefields can propagate only a limited distance from the source into the medium and back to the receiver. The length of the time window is dependent on the maximum travel times in the background medium from the source to the boundaries of the local domain and back to the receivers. Obviously, if the size of the local domain is chosen to be extended, also the time window used for inversion has to be extended. The complete three-step approach to localise the full waveform inversion by backpropagation, subselection of laterally limited sources/receivers and time windowing is schematically shown in figure 3.7

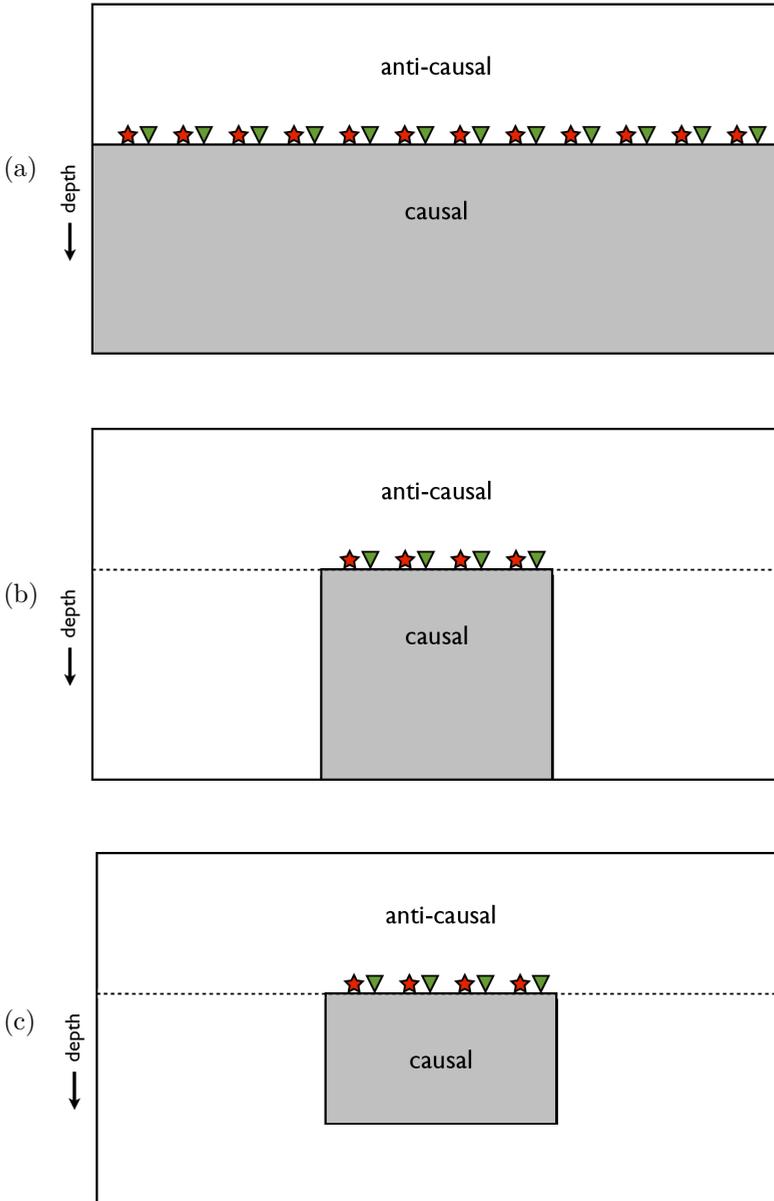


Figure 3.7: After backpropagation, the data can be split into a causal and an anti-causal part (a). By using only a local subset, a lateral limitation of the inversion domain is achieved (b). Finally a limitation in depth can be obtained by applying an appropriate time window dependent on the extension of the local domain (c).

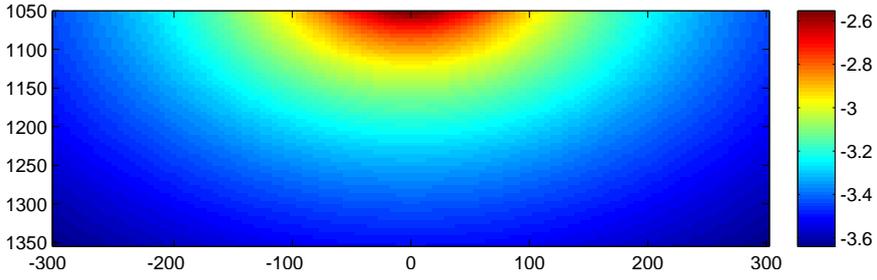
3.2.2 Backpropagation of Green's functions and total fields

We mentioned several times that it is important to backpropagate the data as well as the total fields and the Green's functions in equation (2.1). Furthermore we state that in this way the whole procedure becomes a preconditioning method and no model error is made. In the following we will investigate why this is a crucial step in localising full waveform inversion.

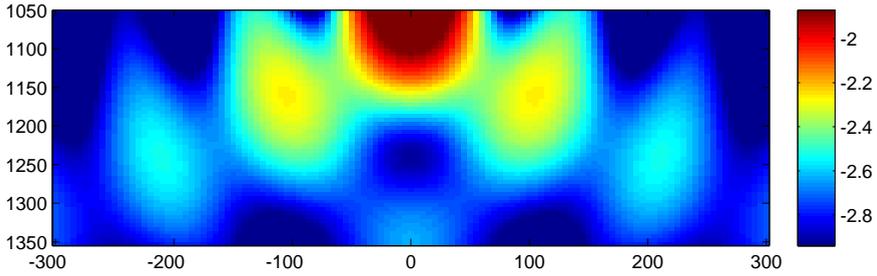
One could argue that it is sufficient to backpropagate the seismic surface data only and to replace the Green's functions and the total fields, originally connecting the depth domain to the acquisition surface, with their equivalents connecting the depth domain to the new acquisition level. In the case of a homogeneous background, Green's functions directly connecting the depth domain to the new acquisition level could be calculated analytically without the extra effort of backpropagating the right-hand side of the data equation. Actually that is the way we started this research but it turned out quickly that this approach will not deliver accurate results. The reason is that backpropagation introduces an acquisition imprint that depends on the source/receiver sampling and the maximum used offset. If we simply neglect this effect and use line sources and line receivers as the result of backpropagation, the data model is not accurate enough for full waveform inversion. We will demonstrate the influence of backpropagation on the Green's functions by an experiment.

We want to compare the total field for a line source directly located above a local domain and the total field in the same domain but for a backpropagated source. The line source response can be calculated analytically in a homogenous background but the backpropagated total field has to be generated numerically. To do so we start with an acquisition where line sources are located along the earth's surface. In a next step we calculate the backpropagation operators as defined in equations (3.2) and (3.3) for a target depth of $z = 1000$ m. Note that the total fields are only calculated for a local domain with $1050 \text{ m} < z < 1350 \text{ m}$ and $-300 \text{ m} < x < 300 \text{ m}$ because otherwise we would encounter the previously discussed memory limitations. The exercise is performed for two different frequencies of 10 Hz and 40 Hz. Because the Green's functions and the total fields are complex numbers in the frequency domain, we plot the logarithm of the absolute values in figures 3.8 and 3.9.

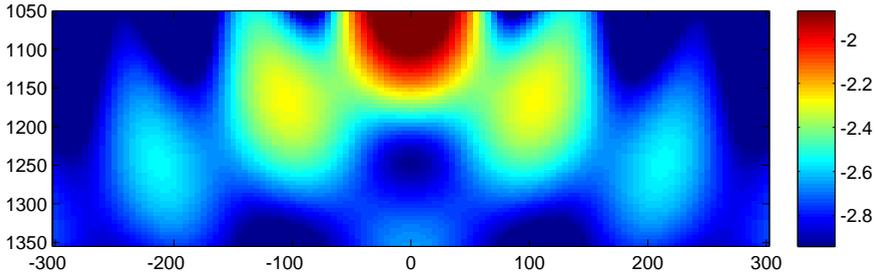
We can clearly observe a significant difference between the analytically calculated line source (a) and a source that was generated by backpropagation (b)-(d). While a line source in a homogeneous medium is simply a circularly symmetric function, the space dependency of the backpropagated source becomes rather complicated. This effect will intrinsically be present in the seismic data after backpropagation and using a line source in the data model of our inversion would be a strong approximation that does not lead to reliable results. The similarity



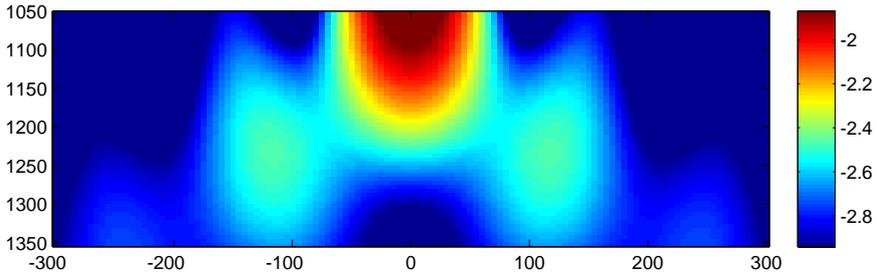
(a) analytic solution for line source



(b) backpropagated source from 401 sources with 25 m spacing at surface

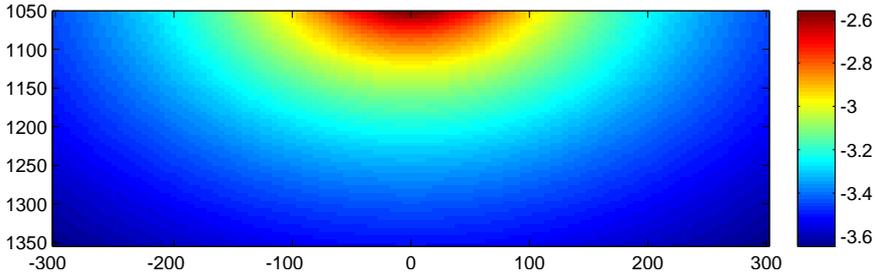


(c) backpropagated source from 201 sources with 50 m spacing at surface

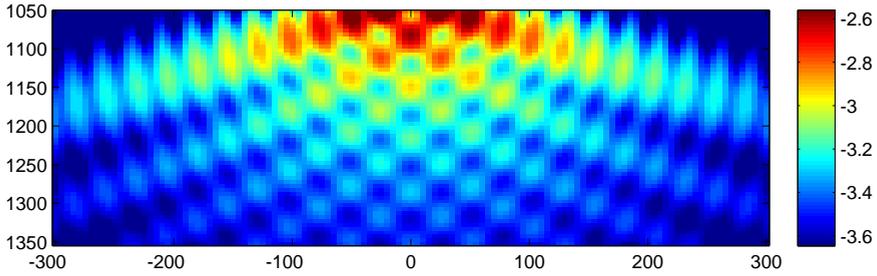


(d) backpropagated source from 201 sources with 25 m spacing at surface

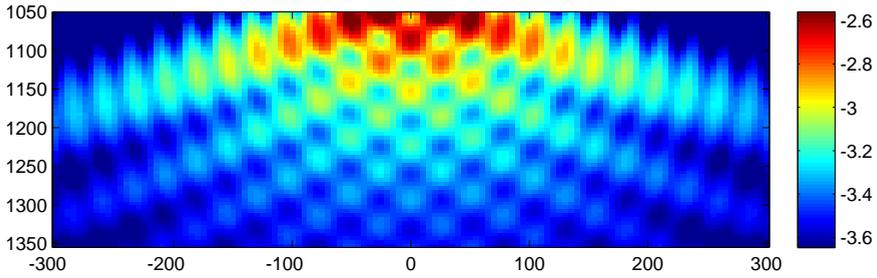
Figure 3.8: Comparison of line source with backpropagated source at $z = 1000$ m generated from surface acquisition. The experiment was performed for a single frequency of 10 Hz.



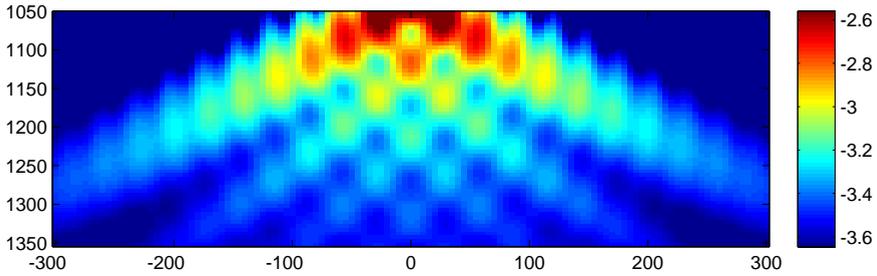
(a) analytic solution for line source



(b) backpropagated source from 401 sources with 25 m spacing at surface



(c) backpropagated source from 201 sources with 50 m spacing at surface



(d) backpropagated source from 201 sources with 25 m spacing at surface

Figure 3.9: Comparison of line source with backpropagated source at $z = 1000$ m generated from surface acquisition. The experiment was performed for a single frequency of 40 Hz.

between (b) and (c) shows that for the current example a source/receiver spacing of 50 m is obviously sufficient and a finer discretisation is not expected to result in any change. If the maximum offset of the surface acquisition is reduced (d), the imprint of the backpropagation becomes slightly different again. The maximum offset limits the highest angle that is still illuminated after backpropagation. Since a true line source in principle illuminates all angles equally well, a backpropagated source illuminates only a triangular area while for very high angles the source field becomes very low amplitude and does not contribute any more. This has to be kept in mind but the backpropagated data will immediately show which angles or offsets respectively can still be used for inversion.

It should be understood by now that the imprint that is introduced by backpropagation is strongly acquisition dependent. As proposed earlier in this section, by backpropagating both sides of equation (2.1) the same imprint will appear in the seismic data as well as in the modelling kernel used for inversion. This is the only way to ensure that an exact inversion result can be obtained, due to not making a model error.

The results also indicate that the effect is not only dependent on the acquisition geometry, i.e. spacing and maximum offset, but also on frequency. While for lower frequencies the spatial variations are much smoother, backpropagation introduces relatively rapid variations if applied to higher frequencies. Since for local inversion also time windowing needs to be applied, the backpropagation process has to be performed over the full bandwidth, because inverse Fourier transformation to the time domain is required. Afterwards, the inversion can be carried out in any desired domain, allowing the option to use only a few frequency components to reduce the problem size.

3.3 Numerical example of local full waveform inversion

Earlier in this chapter we have seen that the size of the kernel that needs to be stored during a conjugate gradient scheme grows rapidly with the size of the inversion domain and the amount of data that is involved. Consequently, a major challenge in inversion of seismic data is to overcome the limits of today's available computers by reducing the simultaneous memory requirements in a smart manner. In the following we demonstrate by a synthetic example how this goal can be achieved by localisation of the inversion. We make use of the seismic surface data that was linearly modelled above the horizontally layered subsurface model shown in figure 3.4.

The full surface acquisition consists of 121 sources shooting into 121 receivers each with a maximum trace length of 2048 ms (4 ms sampling). The inversion domain has an extension of 2400 m in the horizontal direction and 1200 m in

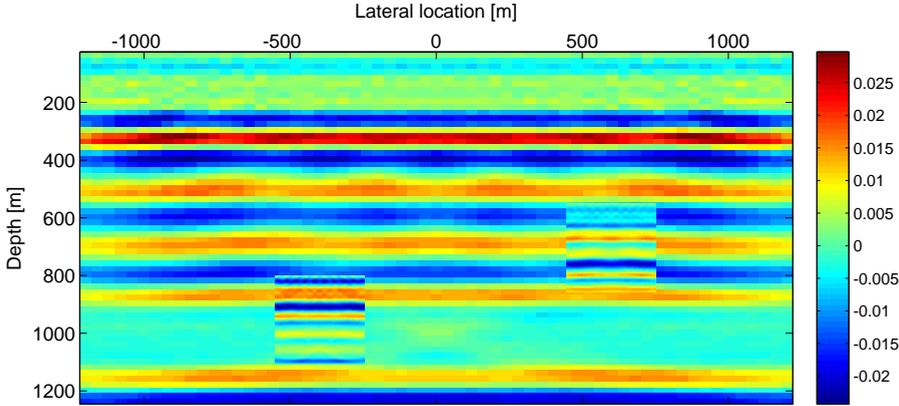


Figure 3.10: A global subsurface model on a coarse 20 m grid was obtained by direct inversion of surface data. After preconditioning the data in the time domain, local inversion on a fine 5 m grid was carried out for two different local domains. It becomes very clear how preconditioning and subsequent inversion on a fine grid improves the resolution and details present in the true model become visible. Note that the true model was given in figure 3.4.

depth. Taking into account that the subsurface parameter are supposed to be recovered on a 5 m grid, we end up with 481×241 unknown velocity values. Following the calculation that was made in table 3.1, we realise that the CG kernel would have a size of nearly 7000 Gb if the whole data should be inverted in one go. It is clear that with the CG scheme we cannot perform this task directly.

To demonstrate how data preconditioning in the time domain can be used to overcome this challenge we perform the following two experiments: First, we directly invert a subset of the surface data on a relatively coarse grid, to reduce the memory requirement. We use 21 sources and 21 receivers with a source/receiver increment of 120 m each. We limit the maximum length of each trace to 300 samples, which is sufficient to image the base of the inversion domain. In this first experiment, the inversion domain is sampled on a 20 m grid and due to the lower sampling we apply a high-cut-filter to the data to reduce the maximum frequencies to 18 Hz. This configuration represents quite common parameters in terms of data sampling but also in terms of resolution that is aimed for by conventional full waveform inversion schemes. Nonetheless, storage of the inversion kernel requires still 8 Gb of memory which can be performed on many standard computers nowadays. It is fully realised that the size of the inversion problem could be further reduced by other data reduction methods, but the purpose of this exercise is to demonstrate how by source/receiver refocusing (backpropagation), not only the data can be

reduced but also the inversion domain can be localised. After our localisation procedure any existing method to further bring down the problem size can still be applied and combined with our approach.

In a second experiment we localise the inversion process by preconditioning as described in this chapter. This allows us to perform linear inversion of all local domains independently, but on a much finer grid as would have been possible by direct inversion of surface data. By backpropagation we generate two local datasets at depth levels of $z_1 = 500$ m and $z_2 = 750$ m. If subsequently we apply a time window to the backpropagated data, we can perform local inversions for domains just below the backpropagated datasets. Note that the inversion domain starts 25 m below the newly generated acquisition. In this way we avoid instabilities that might occur very close to the sources and receivers. This also means that the very early causal part of the backpropagated data needs to be muted by time windowing before inversion.

Each local dataset consists of 11 sources and 11 receivers with a 30 m distance between each. A maximum trace length of 128 samples is sufficient to illuminate the entire local domain. Each local inversion domain consists of 121×61 grid cells but now sampled on a 5 m grid. It is interesting to realise that less than 1 Gb is needed to run the CG scheme within a local domain. This means that by localisation we can lower the memory requirements to an extent that local inversion on a fine grid becomes feasible on basically any existing computer system.

In figure 3.10 the results for inversion of surface data as well as for local inversion of two preconditioned datasets are displayed. It has to be mentioned that for the local results only the middle part of the inversion domain is displayed where illumination was good enough to obtain a reliable result.

Note that inversion of surface data on a coarse grid gives the low wavenumber part of the true model while detailed information of the subsurface properties cannot be recovered. Especially between 900 m and 1100 m depth the subsurface model obtained from global inversion seems to be homogeneous. After preconditioning the data and carrying out inversion over the two local domains, much more detail becomes visible, since inversion on a fine grid allows reconstruction of much higher wavenumber information. The same observation can be made in figure 3.11 where two single traces located at the centre of the local domains are shown.

So far it has been shown that by preconditioning the data in the time domain, a reduced dataset of the surface data can be generated that is dependent on a local domain only. It was demonstrated that the proposed method can be used to invert backpropagated data on a much finer grid than is possible with direct inversion of surface data, resulting in an increased resolution of the subsurface

properties. In the next part we demonstrate how those local inversion results can be combined to obtain a global property model. In figure 3.12 this is sketched by simply putting the local inversion results next to each other. However, we propose to combine the local results in an overlapping manner. The final properties are then obtained by simply calculating the average of the overlapping local domains. Averaging the local inversion results is expected to give more reliable property values and compensates inversion artefacts.

In this example we combine 25 local domains (5 domains in each dimension) by overlapping the single local domains in a way that each point in the global property model is covered 4 times. In figure 3.13 the true and the inverted property models are shown. Note that each local domain had an extension of 300 m by 250 m while the global model obtained by combining the local domain covers 900 m by 750 m. It is very interesting to note that by local inversion a global property model of any desired extension can be built up. The proposed method could even be used in a multiscale approach starting with low frequencies and a relatively coarse grid (Bunks et al., 1995). Subsequently, higher frequencies are included while making the local domains smaller but using a finer grid spacing. It is important to realise that in this way the simultaneous storage requirements of the computation are kept constant. Furthermore, the scheme is perfectly suitable for parallelisation since all local domains can be inverted independently.

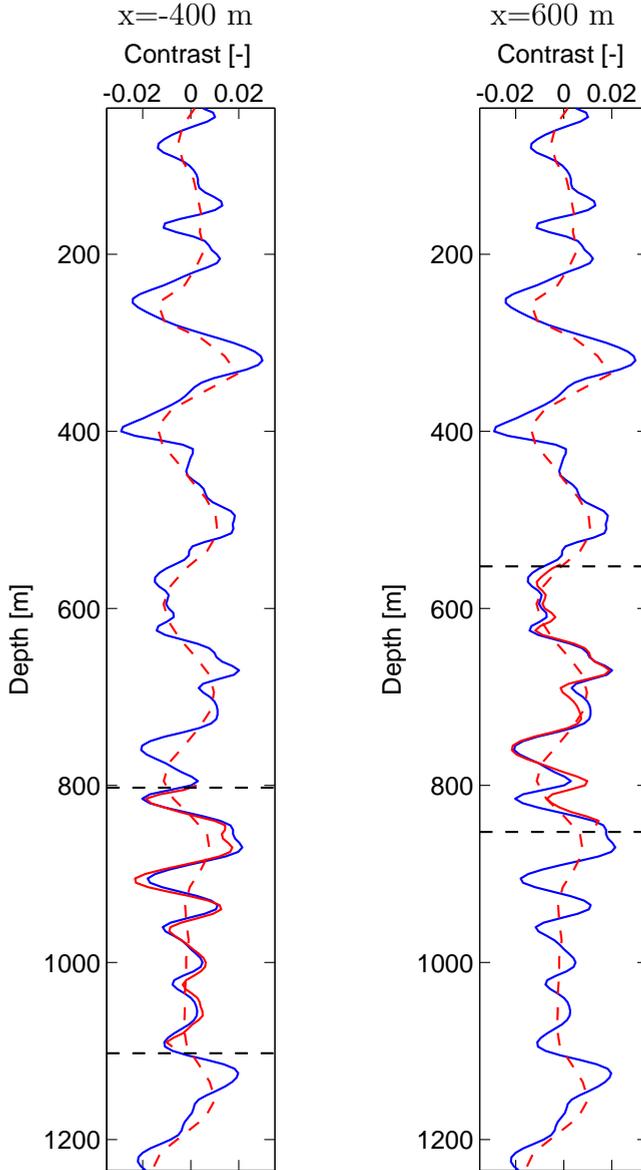


Figure 3.11: Inversion results for locations at the centre of each local domain. Inversion of surface data on a coarse grid (red dashed) results in the low frequency part of the true model (blue solid). After preconditioning and inversion on a fine grid the contrast within the full bandwidth of the use source wavelet could be recovered (red solid).

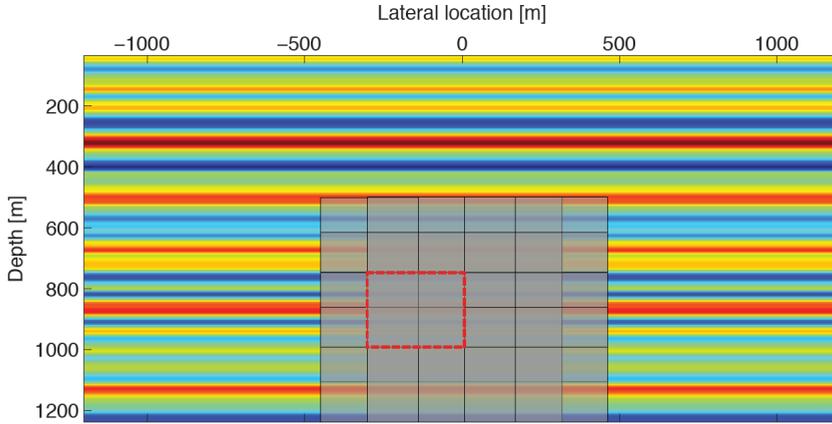


Figure 3.12: Combining local inversion results leads to a global property model. By this method a global property model on a fine grid and any desired size can be built up. Note that the red box indicates the extension of one local domain while due to the overlap it looks like more than 25 domains were combined.

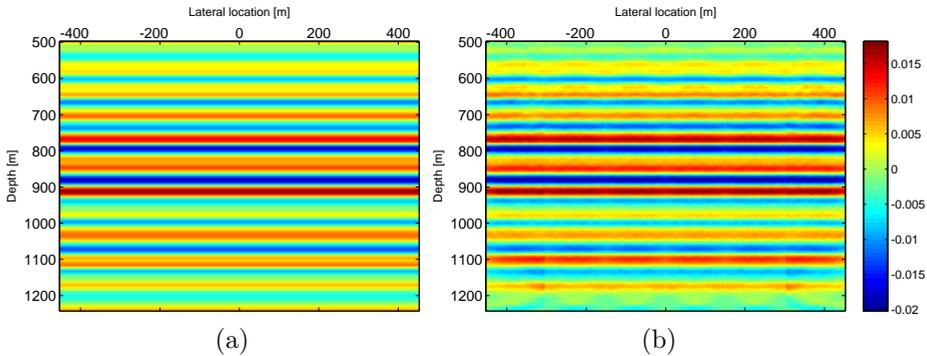


Figure 3.13: A global property model has been inverted by combining 25 local domains (five in each dimension). The true model (a) has been recovered over the whole domain very well (b). Even though in this example a global model of 900 m by 750 m has been built up there are no restrictions on the size for a global model.

Chapter 4

Nonlinear total field estimation

In the previous chapter we have shown that by backpropagating seismic sources and receivers into the subsurface we can localise the inversion process. As a consequence we can perform full waveform inversion on a much finer grid as would have been possible for direct inversion of surface data. Although in this way an increased resolution can be achieved, so far we demonstrated the technique only for linearly modelled data. It has to be understood that, although the approximation is nearly always made in quantitative seismic interpretation, real data are inherently nonlinear in the subsurface properties. For the previous exercise we modelled a linear dataset only to separate the localisation of the inversion from the challenge to incorporate nonlinearity in the scheme. In this chapter we will demonstrate how estimation of the total field, as opposed to the background field, allows us to estimate nonlinear effects like multiple scattering, transmission and also the true traveltimes in the inverted medium. We will see that if a nonlinear data model is used for inversion, the resulting subsurface models are far better than the results from linear inversion. This statement holds for the resolution of the structural image, but particularly so if one wants to recover quantitative information of the subsurface. In this chapter we will apply the total field estimate to a dataset that was measured above a domain that is equivalent in size to the local inversion domain from the previous chapter. In this way we can fully focus on the total field updates, whereas in the next chapter the combination of both techniques will be investigated. In figure 4.1, a flow chart visualises which building blocks of the nonlinear inversion scheme will be discussed in this chapter.

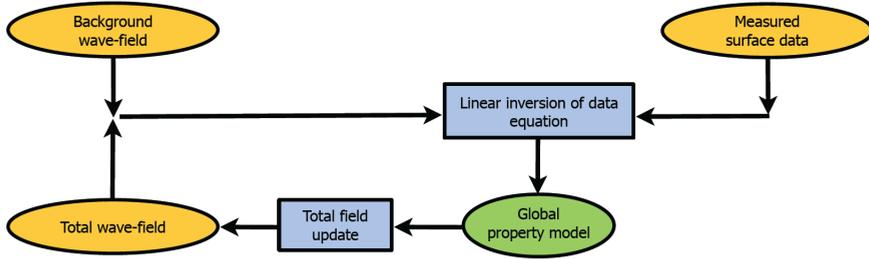


Figure 4.1: *In this chapter we discuss alternating linear inversion by a CG scheme followed by a nonlinear field update. It has to be realised that we apply the technique to a single domain which means that a local domain is equivalent to the global domain for this experiment. In this way we separate the localisation part to focus on the field update only while a combination of both methods will be investigated in the next chapter.*

4.1 Synthetic example I: The Delphi temple

The first synthetic example we use to demonstrate the added value of nonlinear inversion compared to linear inversion, is based on the logo of the Delphi research consortium. The logo includes a greek temple consisting of four major pillars under a triangularly shaped roof, as shown in figure 4.2. Additionally we added a layer below the temple to represent the basement. The whole structure is embedded in a homogeneous background with a constant acoustic velocity of $c = 2000$ m/s. The temple and the basement have a constant velocity of $c = 2218$ m/s, which can be translated to $\chi \approx 0.19$ as defined by equation (2.2). Although development of elastic 2-D inversion will be the subject of future research, current implementation of the scheme is based on the acoustic model with constant density. In the current experiment the density is chosen as $\rho = 2000$ kg/m³ and it is assumed to be known for the inversion. The corner frequencies of the wavelet that was used to generate the data were 7 – 12 – 45 – 55 Hz. The wavelet is shown in the time and frequency domain in figure 4.3. We model a synthetic dataset of 17 sources and 17 receivers with a 60 m spacing between each. Seismic data are generated by solving the exact integral formulation of the wave equation on a 5 m grid. In this way all nonlinear effects, including multiple scattering and transmission, are correctly modelled in the measured data that will be used for inversion later on. We add 3 % random noise to the data before inversion.

It is absolutely clear that the Delphi temple does not represent a very meaningful subsurface model from a geological point of view. Nevertheless, the Delphi temple is a very interesting model to evaluate the benefits of nonlinear inversion. The major challenge will be the reconstruction of the pillars since they form reflectors with a dip of 90°, partly hidden below the roof of the temple. It is well

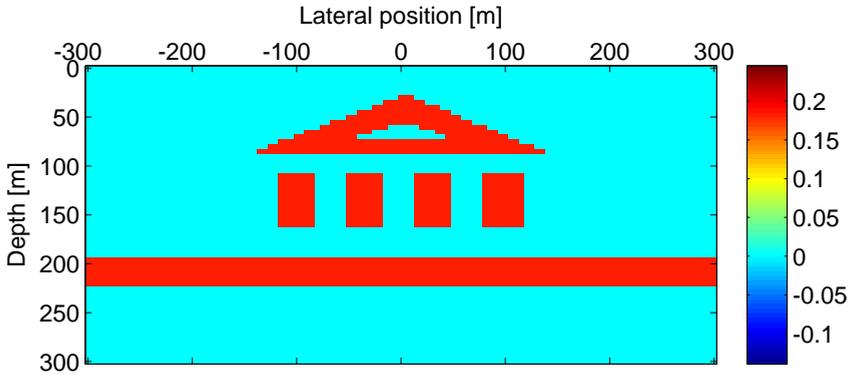


Figure 4.2: An acoustic model based on the Delphi temple was used to generate non-linear seismic data. The homogenous background has a constant acoustic velocity of $c_0 = 2000$ m/s. Embedded in this background is the Delphi temple with a constant velocity of $c = 2218$ m/s. This is equivalent to $\chi \approx 0.19$.

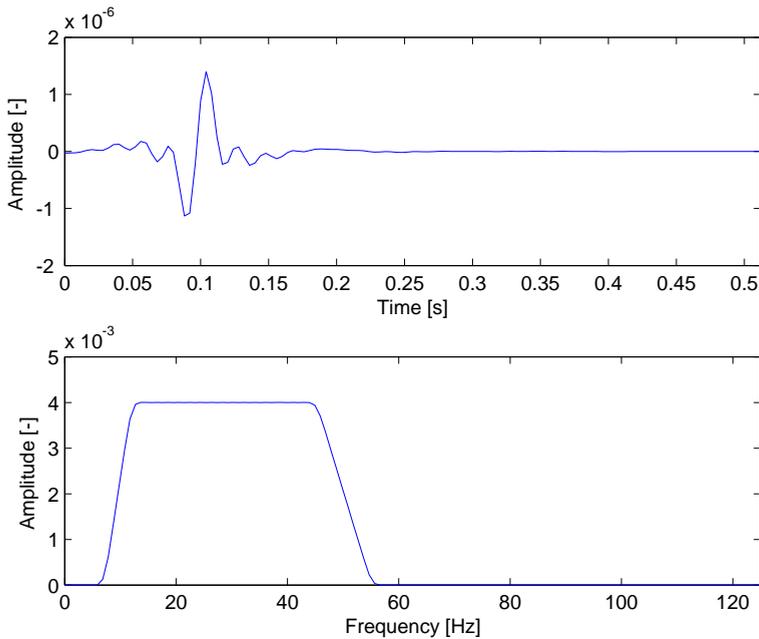


Figure 4.3: Wavelet used for modelling and inversion displayed in time (a) and frequency (b).

known that steep dips can hardly be reconstructed by methods that are based on primary illumination only, e.g. migration or linear inversion. That is why a comparison between linear and nonlinear inversion applied to data measured over the Delphi temple is of specific interest.

The inversion domain consists of 121 grid cells in the horizontal direction and 61 grid cells in the depth direction. Grid spacing was chosen to be 5 m in both spatial dimensions, resulting in a total size of the domain of 600 m by 300 m starting at a depth of 25 m. In the first iteration in figure 4.1 we make use of the Born approximation and we assume that the total wavefield in the domain \mathcal{D} is equal to the background field in the homogeneous embedding. Note that this approximation is not in agreement with the real measured data, since the measured data always contains nonlinear effects as transmission and multiple scattering. Furthermore, the wavefield always propagates in the true medium. Hence, if the kinematic information of the background model is not accurate it will leave an imprint on the linear inversion result.

Despite the fact that linear inversion is based on an approximation of the true wave propagation in the subsurface, it is expected to yield a reasonable result in terms of structure. This statement holds as long as the initial background model describes the kinematics reasonably well. We can then use the output of the linear inversion to update the total fields and make the wavefields more consistent with the inverted property model. In chapter 2 we described a method to build up iteratively the total fields at every grid point in the inversion domain. Instead of solving for the exact field during each iteration, we subsequently add basis functions where each of them includes a higher order of scattering. We would like to recapitulate that in the integral formulation of the wave equation, scattering caused by the contrast function is a mathematical phenomena that accounts for physical scattering at discontinuities, but also for transmission effects and travel-times as they occur in the medium defined by the inversion result. In the current experiment we performed 30 iterations of alternating linear inversions followed by nonlinear field updates, hence the total fields, that were used to obtain the final nonlinear inversion result, consisted of a sum of 30 basis functions. The obtained property models coming from linear inversion as well as nonlinear inversion for increasing orders of scattering are displayed in figures 4.4 and 4.5.

It can be observed that by linear inversion the roof of the temple and the baseline have been imaged reasonably well. However, linear inversion gave no indication of the presence of the four pillars. This can be explained by the fact that linear inversion allows for primaries only. By assuming that every wave is scattered only once in the subsurface there is no possibility for reflections from the pillars to reach the surface again. It should be clear that each wave that propagates from a source to a pillar and gets reflected will still propagate downwards. Only if we consider multiple reflections, those wavefields have a chance to get scattered

another time with the result that they can be detected by the receivers which are located along the surface. Note that a minimum of two orders of scattering are necessary to image a vertical reflector. In fact, many more orders of scattering are necessary to image the whole Delphi temple properly.

Beside the observation that the model was reconstructed much better from a structural point of view, it should also be appreciated that the quantitative properties have been recovered extremely well by nonlinear full waveform inversion. This is particularly interesting given the fact that the true contrast function that was supposed to be recovered has only positive values with respect to the homogeneous background model. This means that the contrast function contains low spatial wavenumbers which are outside the equivalent temporal frequency band of the seismic data. Linear inversion is not able to reconstruct information of the subsurface outside the seismic bandwidth and we can observe this in the inversion result as displayed in figure 4.4 (a). Although the positive contrast function has been inverted reasonably well along the roof and the basement, also negative contrasts are introduced by the linear inversion. These are artefacts that are generated due to the fact that linear inversion is not able to recover the very low spatial wavenumbers. By updating the total fields, nonlinear inversion allows to recover the full spectral bandwidth from a band limited seismic signal, effectively closing the gap between the seismic temporal bandwidth and the spatial bandwidth of the background.

4.2 Synthetic example II: 2-D layered geology

Next we apply the same scheme to a dataset that was acquired above a more geological model. The model is a subset of the very shallow part of the Marmousi2 model (Martin et al., 2006). In this experiment we used 11 sources and receivers with a 60 m spacing each. For inversion we used 50 frequencies equally sampled over the full bandwidth of the source signature. Again, the data was generated by solving the integral formulation of the wave equation on a 5 m grid using the same source wavelet as shown in figure 4.3. We added 3 % random noise to the data before inversion. The homogenous background model in which the Green's functions are calculated has again a constant acoustic velocity of $c = 2000$ m/s and a constant density of $\rho = 2000$ kg/m³. Starting with a linear inversion we subsequently run 20 alternating linear inversions followed by an equivalent number of field updates. The true and the inverted models coming from linear inversion as well as from fully iterated nonlinear inversion are shown in figure 4.6.

Evaluating the results shows that the differences between the linear and the nonlinear inversion result are much smaller than in the previous example. There are several reasons for this. One of them is that the model comes from the very shallow part of the Marmousi2 model and there the contrasts are not that large.

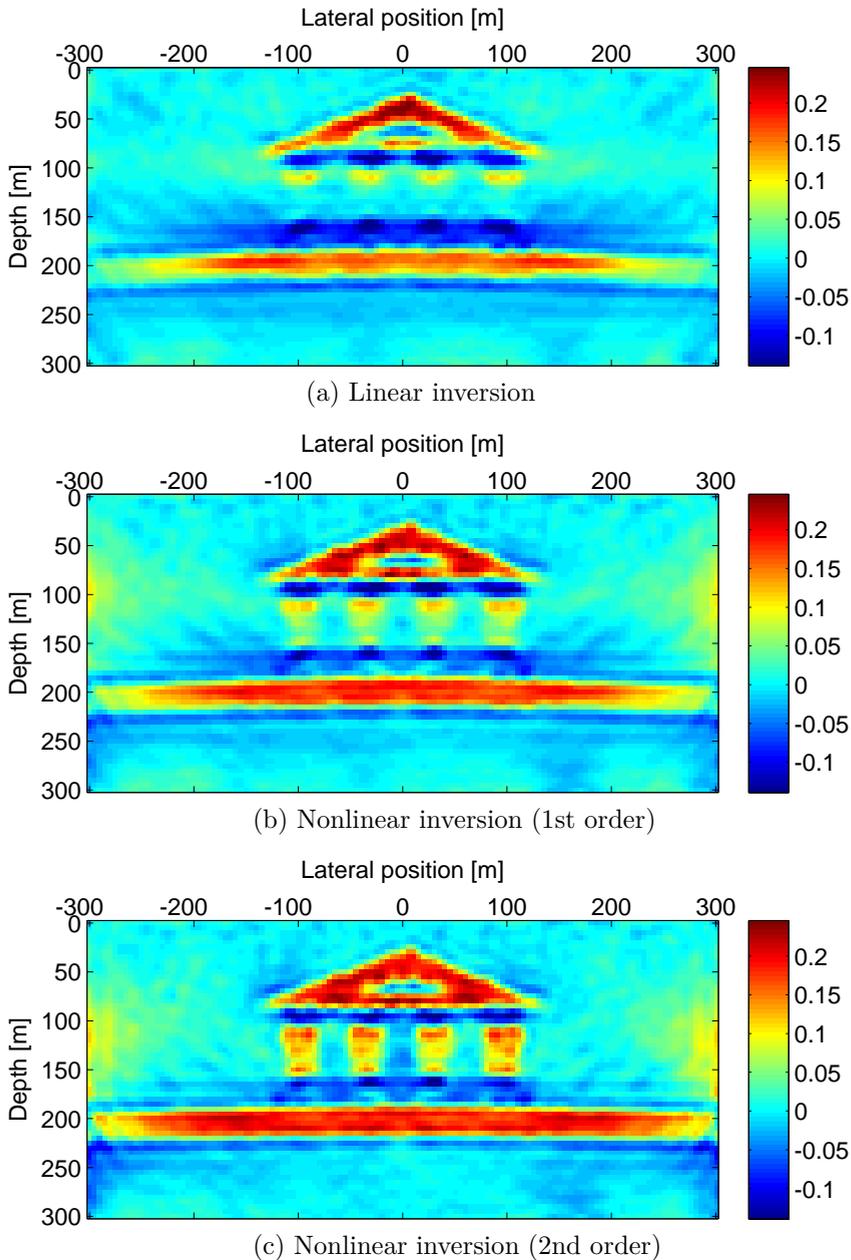
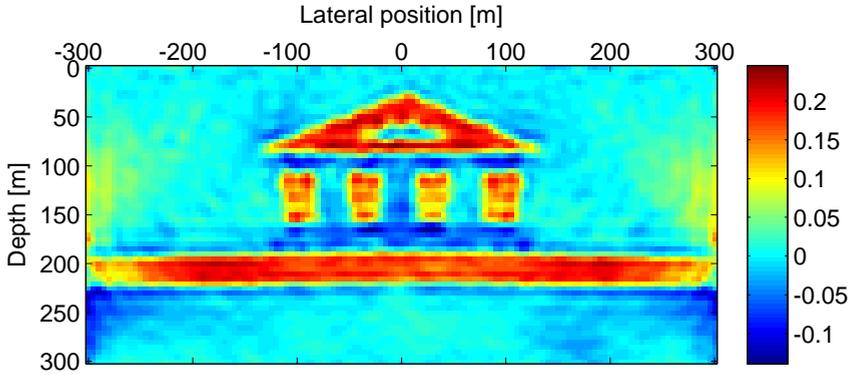
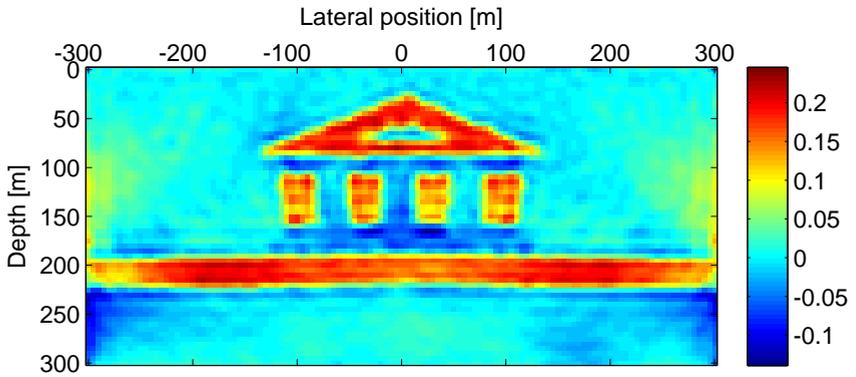


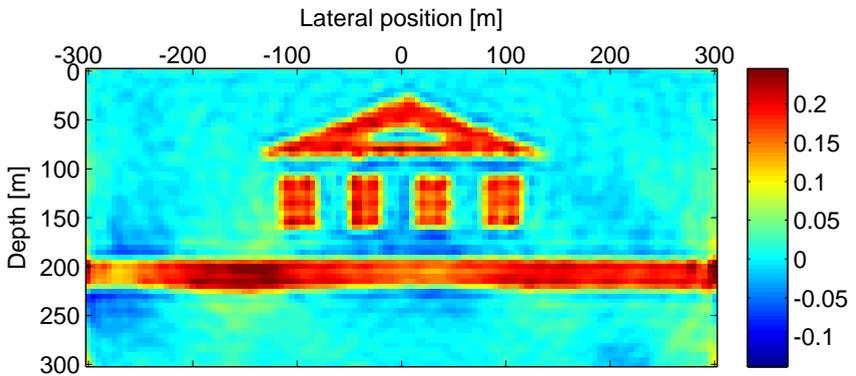
Figure 4.4: Inversion results for the Delphi temple. Starting with a linear inversion (a) up to thirty orders of scattering (b-f) were included for nonlinear inversion (see also next figure). Note how the pillars have been properly imaged by using the multiples in nonlinear inversion.



(d) Nonlinear inversion (3rd order)



(e) Nonlinear inversion (4th order)



(f) Nonlinear inversion (30th order)

Figure 4.5: Continuation of figure on the left page.

Smaller contrast immediately make the inversion problem that is being solved more linear and a linear technique is expected to lead to satisfactory results. Another reason is the fact that the model consist of more or less horizontally layers only. This implies that primary illumination should be sufficient to image the structure adequately. This is obviously very much different to the previous experiment in which the temple included vertical pillars which were impossible to recover by primary illumination only.

Looking at the inverted subsurface models, lets us discover differences. In general the image coming from nonlinear inversion looks sharper and less distorted compared to its linear equivalent. This holds especially for the very first 100 m. In the linear inversion result we can clearly identify an acquisition imprint as a result of the source receiver sampling. While inversion techniques are less sensitive to sampling criteria than direct imaging techniques, still a sufficient sampling is needed if a reliable model is expected to be obtained. The shallower the part that should be imaged, the finer the sampling should be. Our results give the impression that nonlinear inversion can get away with a coarser source and receiver sampling than linear inversion and indeed this is the case. To understand this point it has to be said that the very shallow part of the subsurface is not illuminated by primary reflections. If near surface imaging is the goal, usually refracted waves are used because they can propagate over long distances along the interfaces of the very shallow layers. Although primaries cannot be used to image the very shallow part, multiple reflections do illuminate this part of the subsurface and that is why nonlinear inversion is able to improve the image quality significantly. After 20 iterations the acquisition imprint is completely gone and also the very few meters in depth are recovered with a very high resolution.

It was mentioned already that linear inversion was able, as expected for the current model, to recover the structural information very well. If we investigate the absolute values of the inversion results we can observe differences again. To allow for a more detailed comparison in terms of quantitative properties we plot the contrast functions at $x=0$ m as function of depth in figure 4.7.

Although the structural images alone did not show major differences, plotting the inverted properties in the depth domain for a single lateral location reveals a notable discrepancy. While linear inversion led to a reasonable result, nonlinear inversion recovered the quantitative properties of the subsurface much more accurately. This could be achieved because by iteratively updating the total fields at every point in the subsurface, all nonlinear effects e.g. multiple scattering, transmission and updated traveltimes in the inverted medium were incorporated in the data model.

We conclude this section by transforming the results shown in figure 4.7 from depth to the spatial wavenumber domain. This allows us to relate the recovered

spatial wavenumbers to the temporal frequencies that were present in the measured seismic data. The results are displayed in figure 4.8. Again it becomes clear that nonlinear inversion is superior to linear inversion due to the reasons discussed. There is one last thing that has to be noted, which is unique for nonlinear inversion. The vertical dashed lines in figure 4.8 indicate the spectral bandwidth of the seismic data that was inverted. To calculate these values we simply take the minimum frequency (7 Hz) and the maximum frequency (55 Hz) in the seismic wavelet and transform them to spatial wavenumbers based on the homogeneous background velocity of 2000 m/s. While linear inversion showed a reasonable performance within the seismic bandwidth it failed to provide the subsurface properties accurately outside this bandwidth. Neither the very high wavenumbers nor the very low wavenumbers were recovered accurately. That is why in the seismic community there are long lasting discussions and a vast amount of research is performed on getting good background models for inversion. The argument is that wavenumbers below the low end of the seismic spectrum, have to be put into the background model because they cannot be recovered by the inversion. While it is certainly worthwhile to put effort in getting the 'best possible' background model, figure 4.8 demonstrates that the so-called spectral gap can be closed by nonlinear inversion and hence the dependency on the initial model is reduced. This is a unique characteristic of nonlinear full waveform inversion and as a result, absolute quantitative properties can be recovered by this method. Another valuable effect is that also the spatial position of discontinuities can be corrected by nonlinear inversion, although the kinematics of the initial background model might not have been entirely exact.

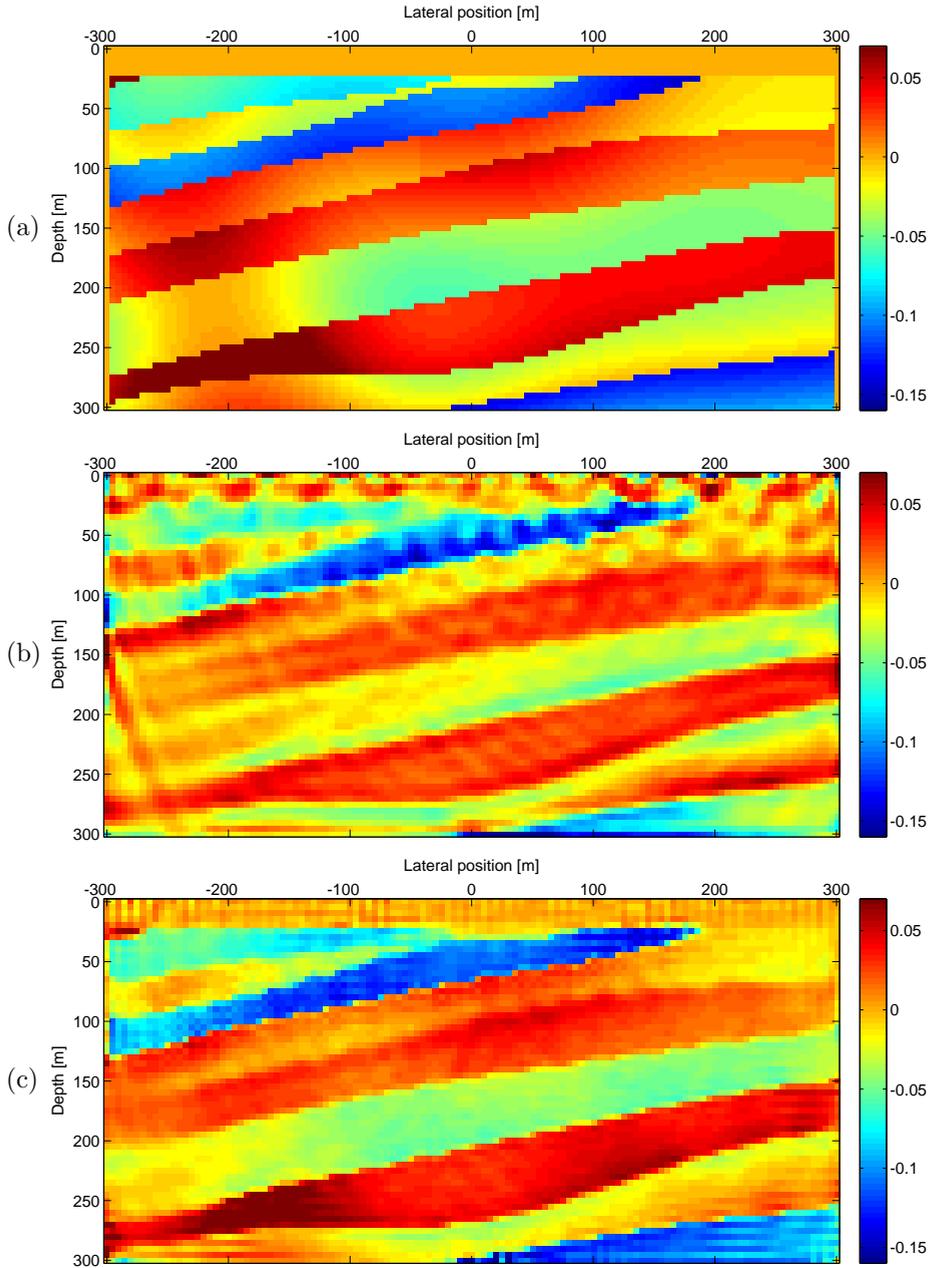


Figure 4.6: In (a) the true model is shown. Although from a structural point of view the results from linear (b) and nonlinear (c) inversion look similar, the absolute properties have been estimated much better by nonlinear inversion.

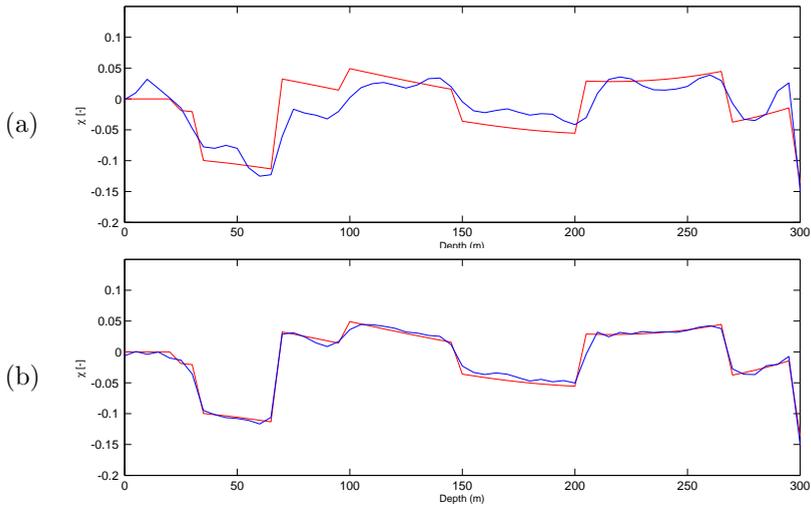


Figure 4.7: The true properties (red lines) at the centre of the model ($x=0$) have been estimated much more accurately by nonlinear inversion (b) than by linear inversion (a). This is due to the fact that nonlinear inversion takes multiple scattering into account while linear inversion is based on primaries only. The increased accuracy comes from the fact that multiples carry complimentary information compared to primaries.

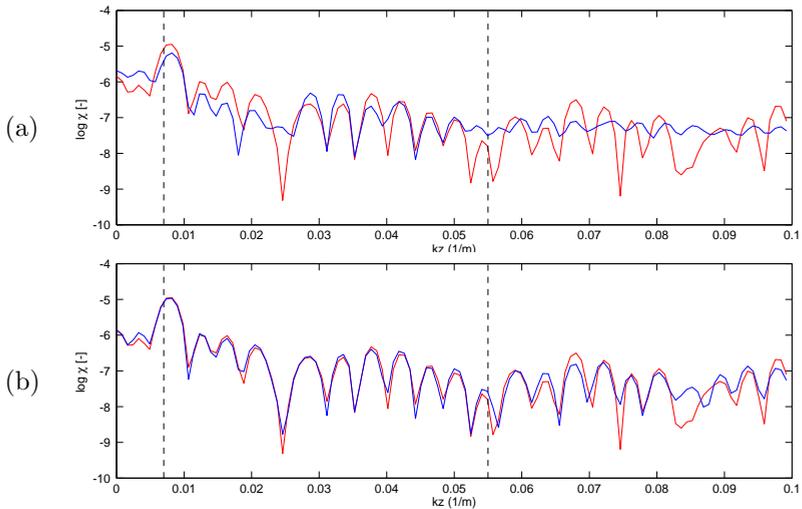


Figure 4.8: While linear inversion (a) can only predict the medium properties within the bandwidth of the seismic data (dashed vertical lines), nonlinear inversion (b) estimates the medium properties over the full wavenumber spectrum. Note that even the low and high wavenumbers which are outside the seismic bandwidth have been estimated very accurately.

Chapter 5

Combining localised full waveform inversion and nonlinear total field estimation

In the previous chapters 3 and 4 we discussed the two major building blocks of this thesis, separately. The first one being localised linear inversions and the second one describing how to update the total fields to incorporate the nonlinear relationship between the subsurface properties and the measured seismic data. What still remains is to demonstrate the combination of both methods and that will be the subject of the current chapter.

We have to recapitulate that in chapter 3 we successfully demonstrated that by backpropagation of sources and receivers into the subsurface, and subsequently choosing a limited time window only, the inversion process can be localised. Because a local domain is only a subset of the entire subsurface under consideration and only a spatially limited backpropagated acquisition was used, we were able to perform full waveform inversion on a much finer grid as would have been possible for direct inversion of surface data. Another consequence was that we could invert the full bandwidth of the measured seismic data (up to 55 Hz) while conventional inversion schemes usually only use lower frequencies (up to ≈ 20 Hz). However, so far we applied the technique to linearly modelled data, hence primary data only, and nonlinearity was supposed to be treated by the nonlinear field update. In the remainder of the thesis we show that localised inversion can be used in the same way to invert seismic data that contain nonlinear effects like multiple scattering and transmission, if combined with a global nonlinear field update.

The principle of single domain total field estimation was successfully demonstrated

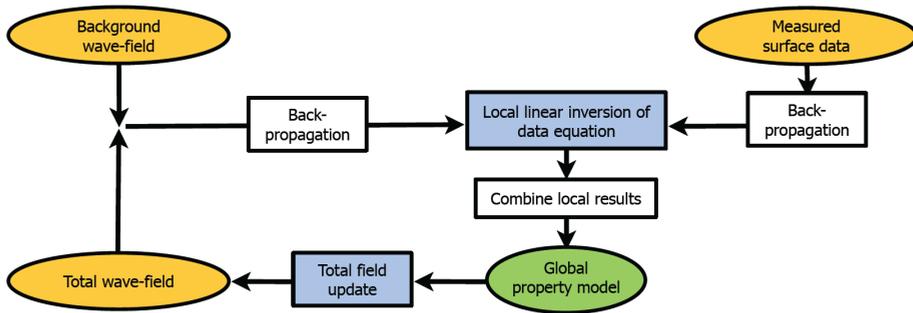


Figure 5.1: *This chapter is about combining the two techniques already discussed in chapters 3 and 4. Starting with approximate total fields, propagating in a background model only, localised linear inversions are carried out. Combining these local results leads to a global high resolution property model. In this global model the total fields are updated and made more consistent with the latest inversion result. In this way nonlinear effects as multiple scattering and transmission are incorporated. By alternating local inversions followed by global field updates a global high resolution nonlinear property model can be obtained.*

in chapter 4. By applying this technique to a single local domain only, we excluded the localisation process and we were able to ensure that the proposed total field updates worked reliably. Part of this chapter will be to apply the nonlinear field update to a property model that was obtained by combining several local inversion results. By alternatingly performing localised full waveform inversions followed by a total field update applied in a global domain covering the combined results, a global high resolution nonlinear property model can be recovered. In figure 5.1 we schematically show the steps involved in the new type of full waveform inversion as discussed in this thesis. Note that figure 5.1 is a combination of figure 3.1 and figure 4.1.

5.1 Subsurface model and synthetic seismic dataset

To demonstrate the methodology we use a subset of the Marmousi2 subsurface model. Again, modelling and inversion is performed in the acoustic approximation assuming a constant and known density of $\rho = 2000 \text{ kg/m}^3$. The subsurface model is discretised on a 5 m grid in the horizontal as well as in the depth direction. We choose a part of the Marmousi model that includes a steep fault, because it is very challenging from an imaging point of view, but also because this type of geology is very relevant for the exploration and production of hydrocarbon reservoirs. The exact contrast function for the acoustic velocity that was used to generate seismic data is shown in figure 5.2. Although for modelling we use a model that is 1500 m wide, we display the centre 1200 m in this chapter only,

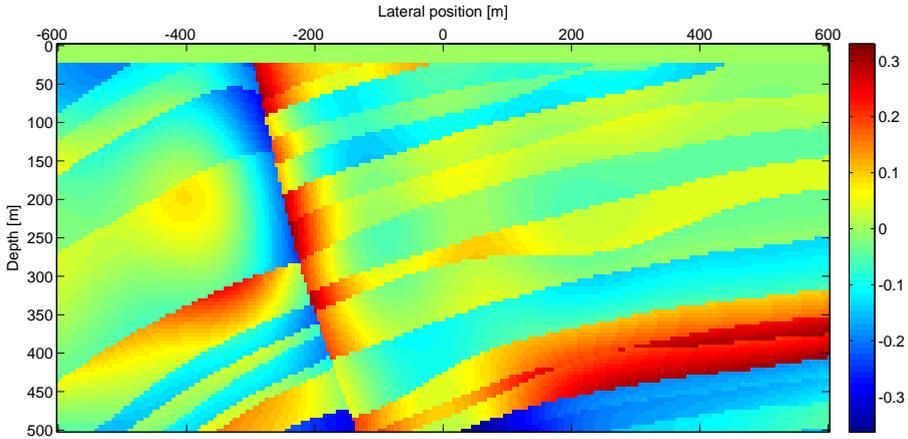


Figure 5.2: *Exact contrast function that was used to generate synthetic data with sources and receivers located along the surface.*

due to the fact that the sides are not illuminated by the limited acquisition and cannot be recovered satisfyingly by inversion. The contrast function is calculated with respect to the homogeneous background of $c_0 = 2000$ m/s that will later be used for inversion. We generate a synthetic dataset at the data acquisition surface with 51 sources and 51 receivers, with a source/receiver spacing of 30 m. Time sampling was chosen to be 4 ms and the maximum length of each recorded trace was 1024 ms. The seismic data were generated by solving the exact wave equation, meaning that all nonlinear effects like multiple scattering, transmission and traveltimes in the true medium were properly modelled. We used the same wavelet that was shown already in figure 4.3 containing frequencies between 7 Hz and 55 Hz. A selection of shot records can be found in figure 5.3.

5.2 Initialisation of the scheme: Linear inversion

We initialise the inversion process by performing localised linear inversion under the assumption that the total fields are equivalent to the background fields in the homogenous background model. Although we are fully aware that this is an approximation of the true wavefield propagation, in the first iteration that simple approximation must be made. We divide the global domain into 27 local domains consisting of nine lateral domains and three different depth levels. In the following we briefly explain the procedure of localising full waveform inversion by source/receiver backpropagation for two separate local domains.

If we consider a local domain that starts at the surface the situation becomes rather simple. A lateral limitation can immediately be achieved by choosing a

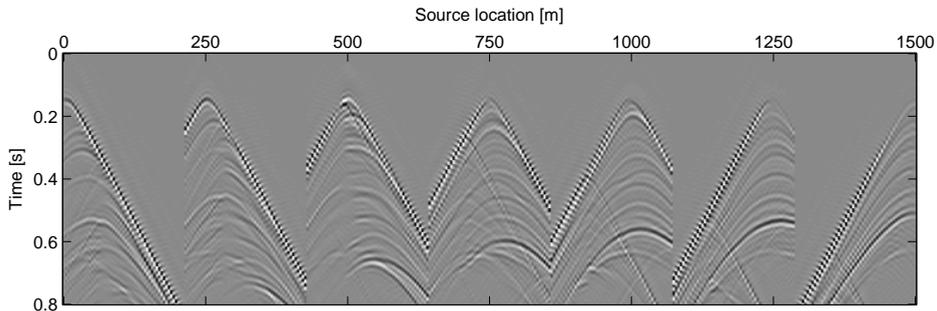


Figure 5.3: *Some shot records from the synthetic data that was generated with sources and receivers located along the surface of the model shown in figure 5.2. The complete dataset consists of 51 sources and 51 receivers with 30 m sampling each.*

laterally limited subset of the surface acquisition. By choosing a limited time window, and thereby making use of causality, also a limitation with respect to depth can be obtained. It can easily be understood that by repeating this process for a desired number of neighbouring local domains, and by combining the local results, a high resolution near-surface model can be obtained.

Since we aim at inverting an extended subsurface domain, we need to get high resolution results also for domains that do not start at the surface. To achieve this goal we backpropagate sources and receivers to a number of depth levels. After bringing the entire acquisition to a new depth level the previously described procedure of localisation can be carried out. Seismic data associated with two different local domains, one starting at $z = 0$ m and another one starting at $z = 300$ m, including the relevant time windows are shown in figure 5.4.

Although for every domain only a single (backpropagated) shot is shown, for inversion of each local domain we use 11 sources shooting into 21 receivers. The sources are separated by 60 m while the receiver spacing is 30 m. The inversion domains consist of 121×61 grid cells which results in 600×300 m² if we take the discretisation of 5 m into account. It is well known that at the sides of the inversion domain there is insufficient illumination to obtain a reliable result and that is why we keep only the middle parts of 300×300 m², to be combined to a global result. The properties within the two local domains that are associated with the data shown in figure 5.4, resulting from the first linear inversion are shown in figure 5.5, together with the global high resolution model that was obtained by combining 27 local inversion results.

The global result is obtained by simple averaging of the middle bits of all the local results while an overlap of two in both directions was used. In this way

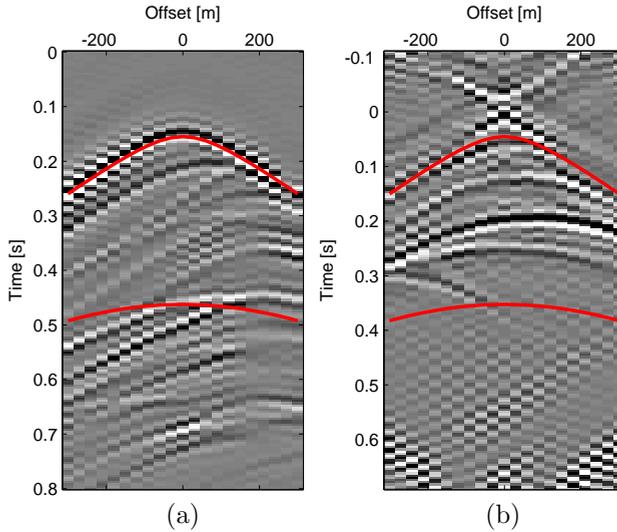
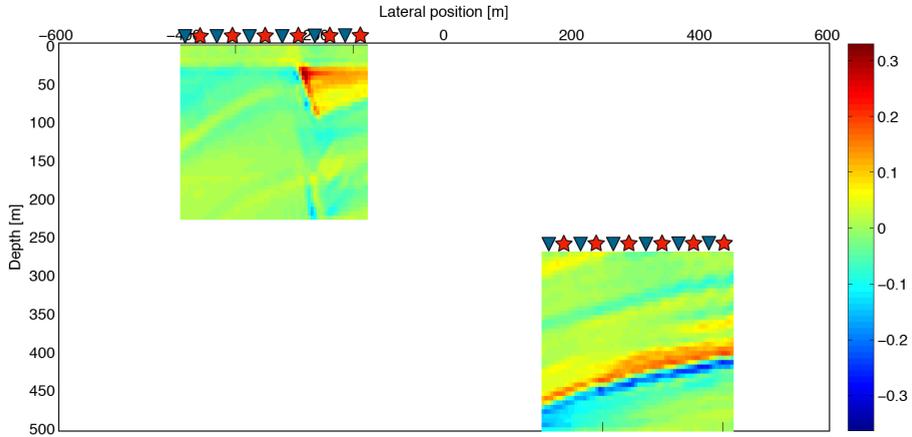


Figure 5.4: Shots corresponding to local inversion domains starting at $z = 0$ m (a) and $z = 300$ m (b) depth. Note that for local domains which do not start at the surface the data has to be backpropagated. For inversion, only the data within the red time window is used. This primary time window is dependent on the extension of the local domain.

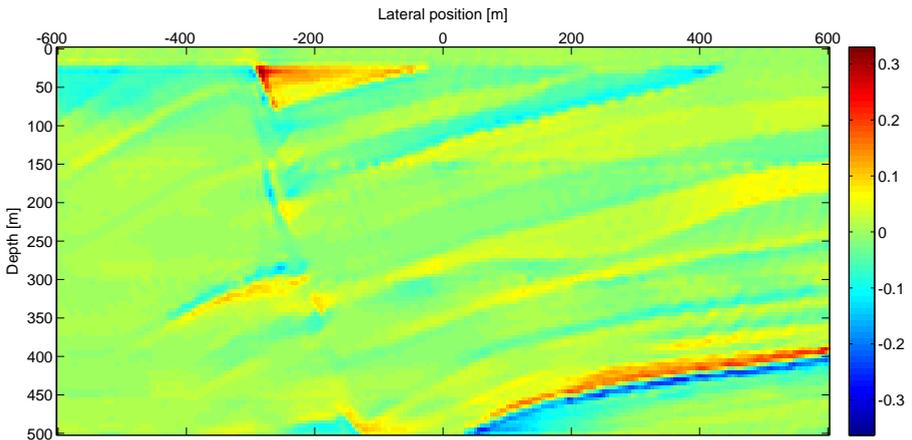
every subsurface grid point is covered four times, which is expected to improve the reliability of the inversion especially in the early iterations. We made this choice to be on the safe side, but some optimisation could be thought of, e.g. by reducing the overlap and/or the number of local inversions needed. In this way less computations would be needed and the scheme could be speeded up.

Evaluating the linear inversion result we make the observation that the layered structure of the subsurface was recovered reasonably well. All major layers are present in the inversion result and a rough idea of the geology could be obtained. In contrast, the steeply dipping fault is not imaged very well. This is expected, because very long offsets would be needed by a linear scheme to image structure with large dips. The problem of a linear imaging/inversion technique is the fact that it uses primaries only and these can hardly be reflected by steeply dipping geology in a way that they reach the surface again.

Think of a wavefield propagating from a source at the surface into the subsurface and approaching a truly vertical reflector. If we assume primaries only, the wavefield will inherently propagate further downwards after getting reflected. Although the event might be reflected multiple times afterwards and consequently could be measured by a receiver at the surface, a linear data model would not be able to explain this event and thereby the vertical subsurface structure. This is a



(a)



(b)

Figure 5.5: Local inversion results for two separate local domains (a) as well as the global result obtained by combining 27 local domains (b). Since the first inversion assumes the total fields in the subsurface to be the background fields in the background medium and because interaction between local domains was neglected an unsatisfactory result is obtained, when compared with the true profile in figure 5.2.

very figurative explanation for the fact that linear schemes have difficulties with steep dips and at the same time it serves as motivation to use a nonlinear scheme that is able to use multiple scattering.

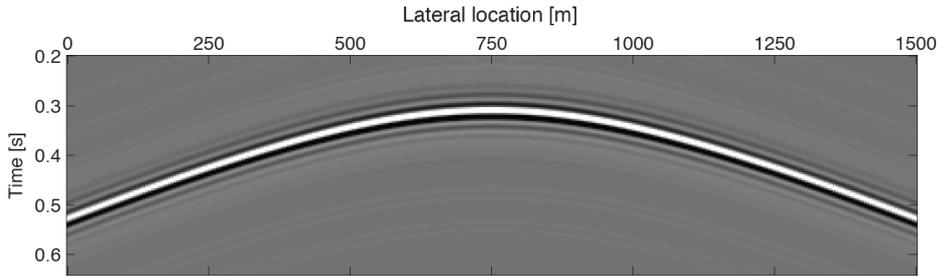
Some more insights can be gained by realising that the colour scale of figures 5.2 and 5.5 (b) are identical. The mismatch of the colours between the inversion result and the true model tells us that the properties were estimated unsatisfactorily by linear inversion. This drawback is related to the ability of linear inversion to recover the properties only within the bandwidth of the seismic data. Wavenumbers that are lower than their equivalent minimum temporal frequency in the inverted data have to be present in the background model. The low wavenumbers are essential if quantitative information of the subsurface is supposed to be recovered. Due to the spectral gap between the low-wavenumber background and the minimum temporal frequency in the measured data, linear inversion will not be able to recover quantitative properties.

5.3 First nonlinear field update

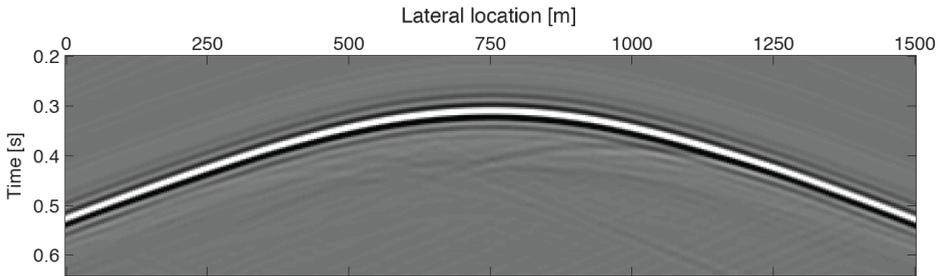
By now we understand the limitations of linear inversion, which are the reasons that the very first iteration did not lead to a very good estimate of the true subsurface model. However, we can use the currently available property model to update the total fields and make them more consistent with the latest inversion result.

We do this in the same manner as described in chapter 2 and as demonstrated in chapter 4. After the first iteration we add one basis function to the background field to bring in first order scattering in the domain equation. It has to be repeated that scattering in the integral representation of the wave equation is a mathematical term that accounts for multiple scattering but also for transmission and traveltimes in the inverted medium. Whenever we refer to an order of scattering it should not be confused with the order of multiples that is commonly used in the seismic community. For example, adding a first order basis function to the total field introduces an estimation of first order scattering. Still, it will take several orders of scattering in the mathematical sense to describe first order multiples in the measured data.

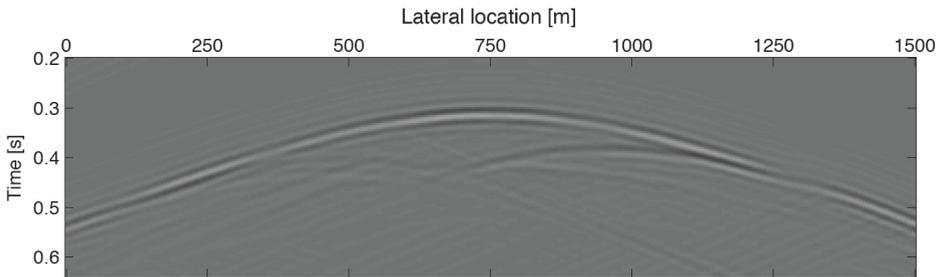
To see the effect of adding the first basis function, we plot the total field at a depth of $z = 400$ m and for a single source at $x = 0$ m in figure 5.6. It should be realised that the wavefields in the subsurface are not measured in the real world. These wavefields might be available locally by well seismic measurements, but certainly not at any location in the subsurface. Still, the domain equation allows us to calculate them from the inverted subsurface model. The background field propagates in the background medium only, which makes it a simple wavefield



(a)



(b)



(c)

Figure 5.6: *The background field for a source at $x = 0$ m and calculated at $z = 400$ m (a) is very simple since it propagated in the homogeneous background medium only. The total field was then updated by adding the first order basis function based on the inversion result from the first iteration (b). Plotting the difference in (c) shows that multiple scattering gets introduced but also that the wavefield changes along the first arrival. This is because the amplitudes as well as the traveltimes are changed by the total field update.*

that can be calculated rather easily. By updating the total field and adding a first order basis function we introduce multiple scattering but also transmission effects. Furthermore, scattering as described by the internal representation updates the traveltimes in the inverted medium. Now we basically could perform localised inversions followed by global field updates alternatingly, but the fact that we are introducing multiples in the total field creates problems for the time windowing in the localisation, which was based on primary arrivals only. This will be discussed and solved in the following section.

5.4 Overburden and inter domain multiples

In chapter 3 we discussed in detail that by backpropagation of sources and receivers into the subsurface, full waveform inversion can be limited to a local domain. We demonstrated the procedure for truly linear data that contained primary reflections only. At this earlier stage we also mentioned that backpropagation is a linear concept that cannot account for multiples that are generated in the overburden, i.e. the part between the target depth level and the acquisition surface. The purpose of this chapter is to combine the localised inversions with global nonlinear total field updates when the data to be inverted contains all kinds of multiples that are generated in the subsurface. This includes intra-domain multiple scattering as well as inter-domain multiple scattering. The global field update will handle all these effects automatically, but we have to understand which multiples can be used for local inversion and which have to be treated in a different way.

5.4.1 Classification of primaries and multiples

In principle we can classify four different types of events in the context of local full waveform inversion. The first class contains primaries and internal multiples that are scattered within the local domain under consideration only. A wavefield propagates from a source into the subsurface, it is scattered at least once within a specific domain, and consequently propagates back to the acquisition surface where it is detected by a receiver. It is important to realise that in the integral formulation, propagation back to the receiver occurs in the smooth background only, and no further scattering occurs outside the local domain. A schematic description of these Type-I events, as we will call them, is given in figure 5.7 (a). Along the same line, Type-II events get scattered at least once outside the local domain, followed by at least one order of scattering inside the local domain we are interested in. Although several Type-II events can be thought of, we display a single example in figure 5.7 (b). Type-I and Type-II events have in common that the last bounce before arriving at the acquisition surface occurs inside the local domain. This is in contrast to Type-III and Type-IV events that will be introduced next. Type-III events get reflected at least once within the local domain, they also get scattered at least once outside the local domain, but this

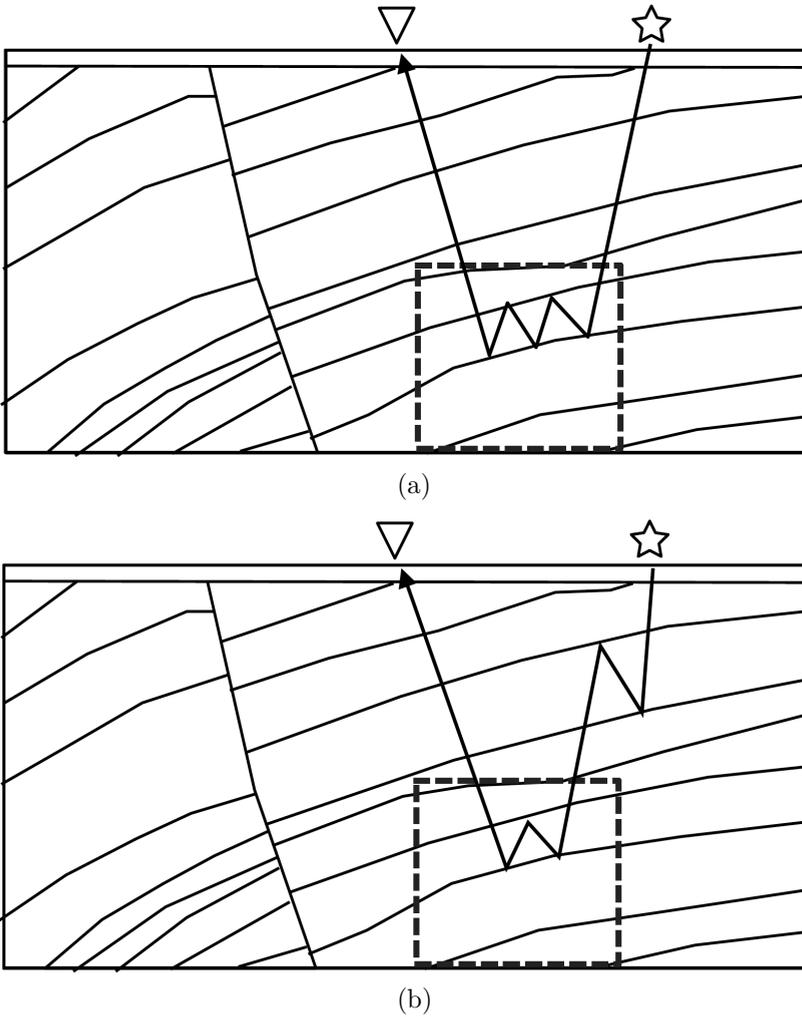
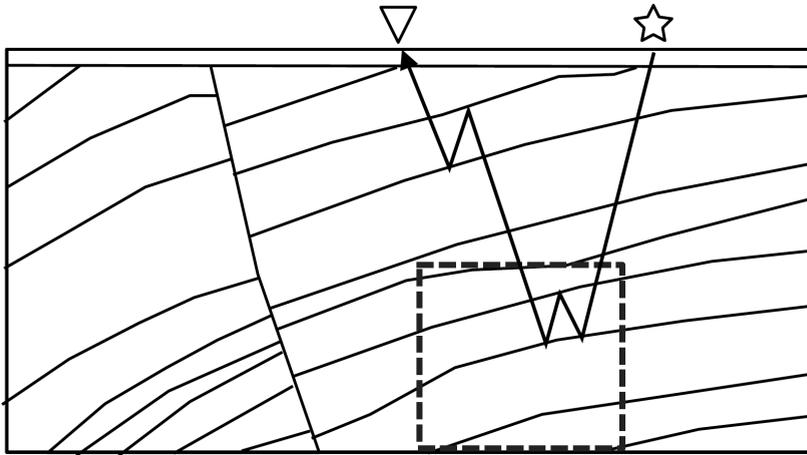
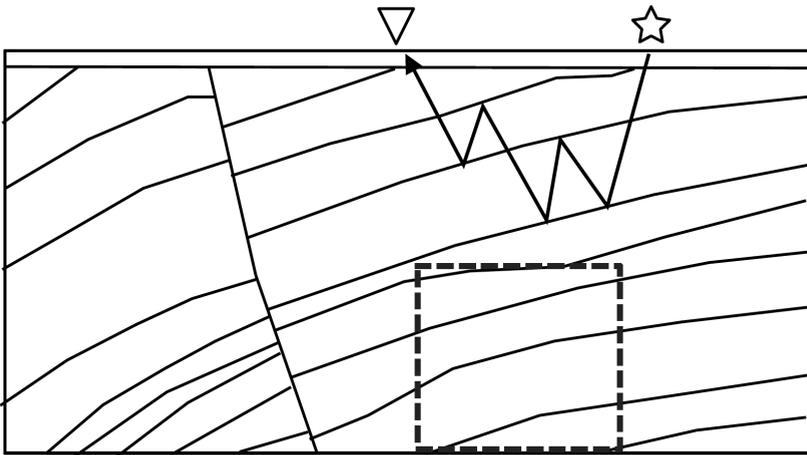


Figure 5.7: *Seismic data at the surface can be described as superposition of different reflections. We classify different reflections as Type-I (a) to Type-IV (d) events (see also next page) depending on whether scattering occurred inside and/or outside the local domain. Note that Type-I (a) and Type-II (b) events have the last reflection point inside the local domain while Type-III (c) and Type-IV (d) events are scattered finally outside the local domain.*



(c)



(d)

Figure 5.8: Continuation of figure on the left page.

time the last bounce before detection was outside the local domain. An example of a Type-III event is given in figure 5.8 (a). Seismic data also contains events that were scattered outside of the local domain only and hence never propagated through the area of interest. Think of a multiple that bounces in the overburden and gets measured at the acquisition surface but the wavefield never propagated deep enough to reach the local domain we want to invert for. These reflections are classified as Type-IV events and an example is given in figure 5.8 (b).

Time-wise, all four categories of multiple events could fall in the primary time window of the local domain under consideration. Although the integral representation does not intrinsically differentiate the defined classes of events we make this separation to understand better which reflections contribute to a specific local inversion. It is obvious that Type-I events are purely associated with a local domain and hence they fully contribute to local inversion. It has to be realised though that these events might occur outside the primary time window that is applied to the data after backpropagation. Type-I multiples with an arrival time larger than the primary time window will then be muted and consequently they cannot contribute to the inversion result. We call this the coda of the local domain.

With Type-II and Type-III events the situation gets more complicated. Here it is the last reflection point that determines if an event is contributing to the local domain inversion. Because Type-II events experience the last reflection inside the local domain before reaching the acquisition surface, they will add to the inversion result in the local domain. It is interesting that multiples generated by the geology outside the local domain increase the local illumination as long as their last reflection point is located inside the local domain and their arrival time falls into the time window used for inversion.

The same reasoning tells us that Type-III events do not contribute to the local domain that is indicated in figures 5.7 and 5.8 because their last reflection lies outside the local domain. Classification into different types of events is dependent on the specific domain one is interested in. While the indicated event in figure 5.8 (a) is a Type-III event for the outlined local domain, it will turn into a Type-II event if the local domain above would be considered. It is only the last reflection point that determines if an event contributes to a specific local domain or not.

Off course, by applying reciprocity, one can always turn a Type-II arrival into a Type-III arrival, and vice versa, without changing the information content. However, it is the way the data are organised that determines what information can be extracted from it.

The last class that we consider are Type-IV events and they can be referred to as overburden multiples. These events get scattered in the overburden only without ever reaching the local domain we want to invert. Consequently, these

multiples might have a travelttime larger than the beginning of the primary time window that is used for inversion. This means that Type-IV multiples can leak into the time window that is used for local inversion although they are not related to the local domain at all. Because backpropagation is a linear concept it cannot handle Type-IV multiples adequately and we need to find a way to deal with the undesired Type-III and Type-IV multiples in a different way.

5.4.2 Subtraction of complimentary data as exact solution to localisation

In the integral representation of the wave equation, seismic data in the data domain (acquisition surface) is modelled by summing the contributions from all grid points of the discretised subsurface. No distinction between primaries and different types of multiples is made while all wave propagation effects are fully incorporated in the total fields that are the source strengths of the secondary sources in all grid points. The integral formulation allows us to split the data equation, as it was defined in equation (2.1), into a part which is the contribution of the local domain only and its complimentary part. Mathematically the process can be described by splitting up the data equation:

$$\begin{aligned}
 p_{data}(\vec{x}_r, \vec{x}_s, \omega) &= \int_{\vec{x} \in \mathcal{D}_\cup} \mathcal{G}(\vec{x}_r, \vec{x}, \omega) \chi(\vec{x}) p_{tot}(\vec{x}, \vec{x}_s, \omega) d\vec{x} \\
 &+ \int_{\vec{x} \in \mathcal{D}_\cap} \mathcal{G}(\vec{x}_r, \vec{x}, \omega) \chi(\vec{x}) p_{tot}(\vec{x}, \vec{x}_s, \omega) d\vec{x}. \quad (5.1)
 \end{aligned}$$

where the first term on the right hand side is the contribution from the current local domain, indicated by \mathcal{D}_\cup and the second term is the contribution from the complement of the local domain, indicated by \mathcal{D}_\cap . The first term contains all the recently introduced Type-I and Type-II events, while the second term is the sum over the undesired Type-III and Type-IV events.

It should be noted that in the equation above sources and receivers are still located along the surface. Once we have a first estimate of all contrasts in the global domain, the complementary part for a specific local domain can be calculated and subtracted from the measured surface data in order to backpropagate only the part of the seismic data that is related to the local domain under consideration. This can be done because we iteratively build up knowledge of the contrast function as well as of the total field at every grid point. We define the surface data related to a specific local domain L , after subtracting the complementary part at

iteration n as

$$\begin{aligned}
 p_L^{(n)}(\vec{x}_r, \vec{x}_s, \omega) &= p_{data}(\vec{x}_r, \vec{x}_s, \omega) \\
 &- \int_{\vec{x} \in \mathcal{D}_\cap} \mathcal{G}(\vec{x}_r, \vec{x}, \omega) \chi_n(\vec{x}) p_{tot}^{(n-1)}(\vec{x}, \vec{x}_s, \omega) \, d\vec{x} \\
 &= \int_{\vec{x} \in \mathcal{D}_\cup} \mathcal{G}(\vec{x}_r, \vec{x}, \omega) \chi_n(\vec{x}) p_{tot}^{(n-1)}(\vec{x}, \vec{x}_s, \omega) \, d\vec{x}. \quad (5.2)
 \end{aligned}$$

This means that after the first iteration, equation (5.2) will be used and the result $p_L^{(n)}$ is backpropagated to localise the inversion process. In this way overburden multiples (Type-IV) and inter domain effects (Type-III) are removed from the data for a specific local domain in an iterative procedure. Furthermore, time windowing could finally be abandoned and multiples that have the final reflection point inside the local domain (Type-I and Type-II) but arrive outside the primary time window can be used and can therefore contribute to the inversion as well. Following this procedure an exact solution is obtained without making any approximations by the localisation. Backpropagation only serves the purpose of preconditioning the inversion problem and to enhance the sensitivity of the inverse scheme to a specific local domain.

It is interesting to realise that after a number of iterations, the backpropagation could be left out altogether. This means that the concept of backpropagation can be abandoned and local inversion could be performed from the acquisition surface directly, if the complementary part of the data was subtracted beforehand.

Before proceeding with the results of the fully iterated nonlinear inversion we demonstrate subtraction of the complementary data using the total fields and contrast function obtained by linear inversion. The measured data we show in figure 5.9 (a) belongs to a local domain that starts at $z = 300$ m depth and with the centre being located at $x = 150$ m. For the first iteration, the entire measured surface data was backpropagated to localise the inversion, because initially the contrast function is zero everywhere and we are not able to subtract contributions from the complement of the local domain.

After the first iteration, we can use the obtained contrast function and the initial total field $p_{tot} = p_0$ to estimate the part of the data that does not belong to the local domain, but comes from the overburden or the surrounding of the local domain. This can be done by modelling surface data while integrating over grid points in the domain \mathcal{D}_\cap (the complement of the local domain) only. Note that this is equivalent to the second term on the RHS of equation (5.1). The corresponding data are shown in figure 5.9 (b). The first estimate of the complementary data explains the surface data partly, but quite some residual energy remains. It should be understood by now that figure 5.9 (b) shows predicted Type-III and Type-IV

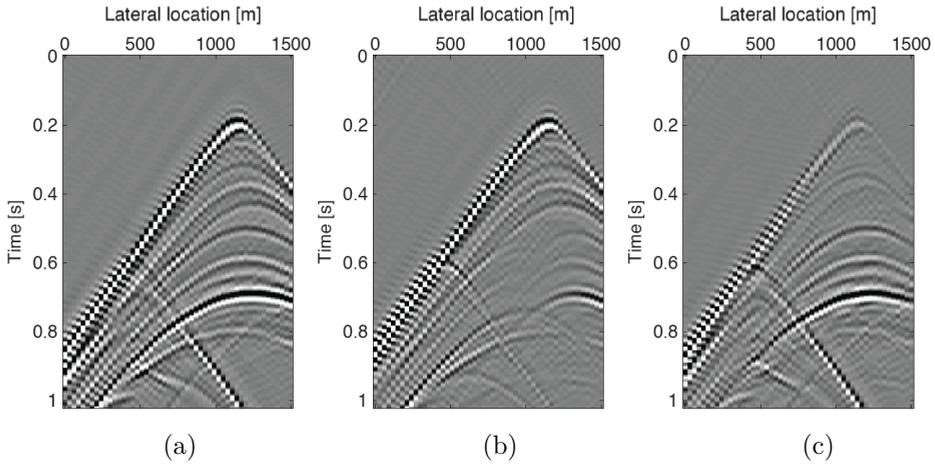


Figure 5.9: Surface data for a single shot at the surface, located just above a local domain starting at $z = 300$ m and with the centre being located at $x = 150$ m (a). The complementary data (b) obtained by integrating over grid points after the first iteration, but over the domain \mathcal{D}_\cap only. Difference between (a) and (b) which is the current best estimate of the data associated with the local domain \mathcal{D}_\cup , is shown in (c). Only the difference will be used for the next local inversion of this specific domain. Note that for other domains the complementary data will differ but the concept will be equally valid.

events while the residual data in figure 5.9 (c) should ideally contain Type-I and Type-II events only. Looking at the predicted complementary data it can be realised that reflection energy is distributed over the entire shot gather. This tells us that reflections generated outside the local domain overlap with energy generated inside the local domain. From surface data, a direct separation seems to be very hard to achieve and that is why we use backpropagation of sources and receivers as preconditioning. The results displayed earlier in this chapter confirm that this is a suitable method for localisation of full waveform inversion. The remaining data overlap due to multiple scattering and/or the fact that a reflection is not only generated at a single grid point but over an extended Fresnel zone, must be handled by estimation and subtraction of the complementary data for the local domain under consideration.

5.5 Fully iterated nonlinear inversion results

During the first iteration the inversion was based on the Born approximation, assuming that wavefields propagate in a smooth nonscattering background medium only. Localised inversions were obtained by backpropagating sources and receivers into the subsurface and by subsequent time windowing. Based on the linear inversion result a total field update was performed and the complementary data was estimated for each local domain. This three-step procedure of localised

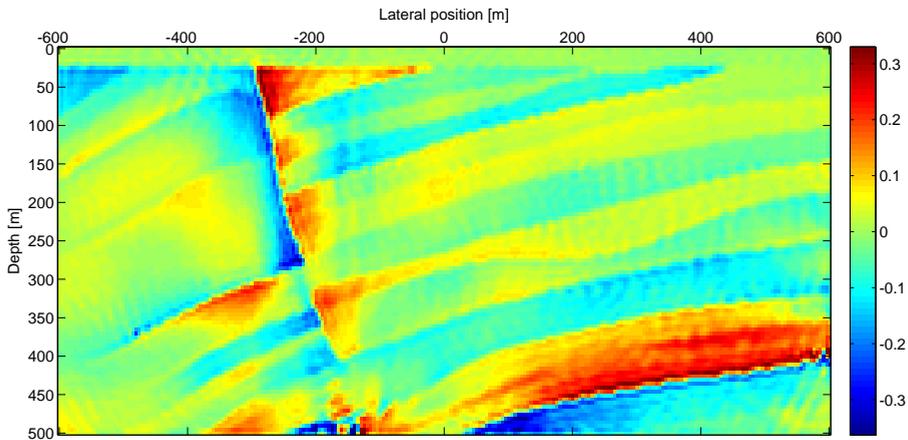


Figure 5.10: *Final contrast function coming from our nonlinear inversion scheme. Note how well the structure is reconstructed but also how the quantitative information of the subsurface is nicely estimated. The fault in the model is clearly visible, which demonstrates that the resolution of nonlinear inversion is far beyond what can be obtained by linear imaging/inversion techniques.*

inversions, global total field update and estimation and subtraction of the complementary part will be performed 35 times. Then we stop the procedure because the total field and the contrast function obtained by inversion do not significantly change any more.

In figure 5.10 we show the final subsurface model which is the result of the fully iterated nonlinear inversion scheme as it was described in this thesis. A comparison between linear inversion as shown in figure 5.5 (b) and nonlinear inversion clearly shows the improvement that can be obtained in terms of structure and properties. For this synthetic example and controlled experiment, nonlinear inversion was able to reconstruct the absolute properties of the subsurface very well. Interestingly, the fault in the model could be imaged by nonlinear inversion very accurately while it could not by linear inversion. Reasons why the illumination can be improved by using nonlinear effects like multiple scattering have been given earlier. Still, it is not only multiple scattering but the entire nonlinear relationship between the data and the subsurface properties that allows us to get a much better estimate of the subsurface. Obviously, the improved structural and quantitative image will have a major impact in quantitative interpretation of the seismic, and reservoir modelling or simulation.

In figure 5.11 the updated total field after the last iteration is shown for the same configuration as in figure 5.6. The final total field is a weighted sum of 35 basis functions and it clearly differs from the background field. A notable

amount of multiple scattering was estimated by the iterative procedure and the change along the first arrival indicates that also amplitudes and traveltimes were updated accordingly. It should be mentioned again that the field update must be performed for the original source geometry which is usually along the earth's surface. By estimating and using the nonlinear relationship between the seismic data and the subsurface properties, the obtained property model was improved significantly.

To conclude this chapter we display the results on estimation and subtraction of the complementary data in figure 5.12. It is obvious that after 35 iterations, the complementary data for a specific local domain is estimated very well. Wavefields reflected in the overburden are nearly perfectly estimated and subtracted out. Please note the event going obliquely through the data, which is generated by the fault. After 35 iterations it is nicely predicted and subtracted out because it is not generated within the local domain we want to invert.

It is worthwhile mentioning that by performing a global field update and subsequently subtracting the complementary data, an exact inversion scheme is obtained, in which the localisation takes only the role of preconditioning. It also becomes clear that if the complementary data can be estimated very accurately, time windowing the backpropagated data is in principle not necessary any more and even the coda that arrives outside the primary time window could be used for inversion. This is possible because the interference of the coda from a local domain with reflections from the subsurface below the local domain will be subtracted out as well.

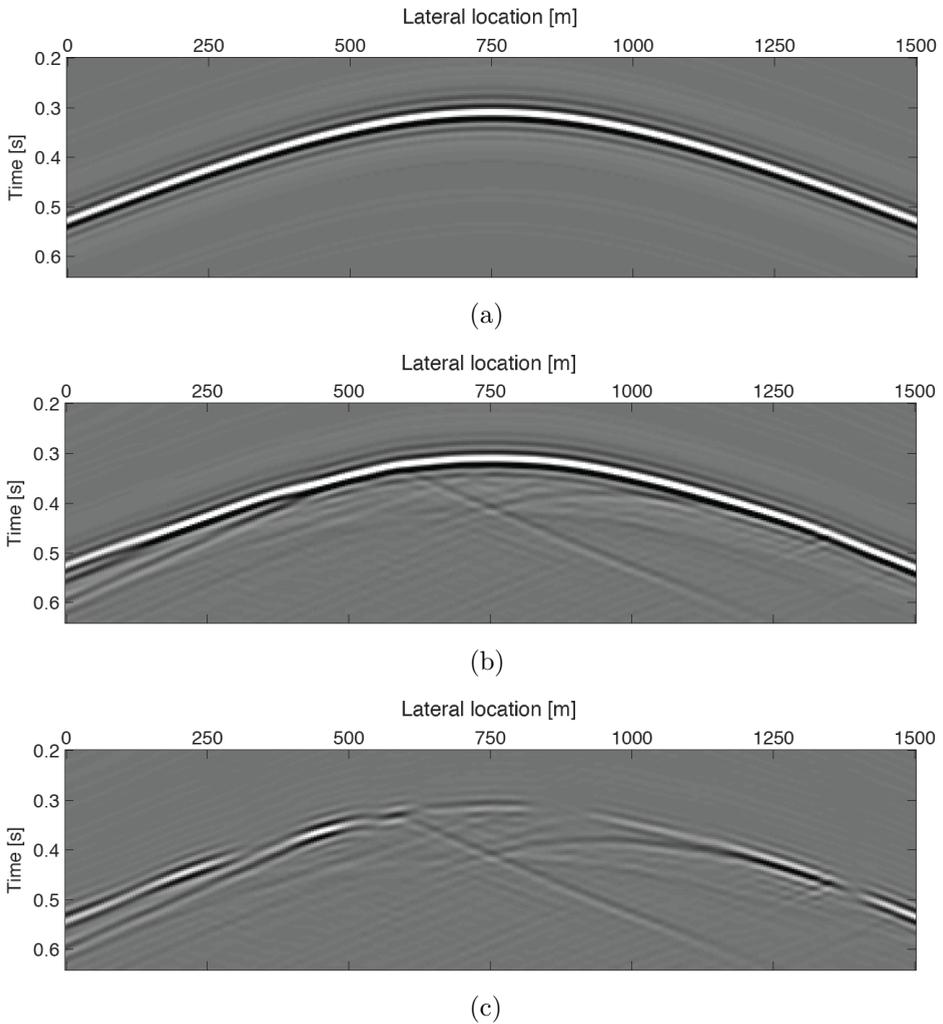


Figure 5.11: *The background field at $z = 400$ m (a) is very simple since it propagated in the homogeneous background medium only. After 35 iterations of alternating local inversions followed by a global field update (b), the total field becomes much more complex since it describes the true wavefield propagation in the best estimate of the medium. Updating the total field, estimates multiple scattering as well as the correct traveltimes and amplitudes in the inverted subsurface model. This becomes clear by plotting the difference in (c).*

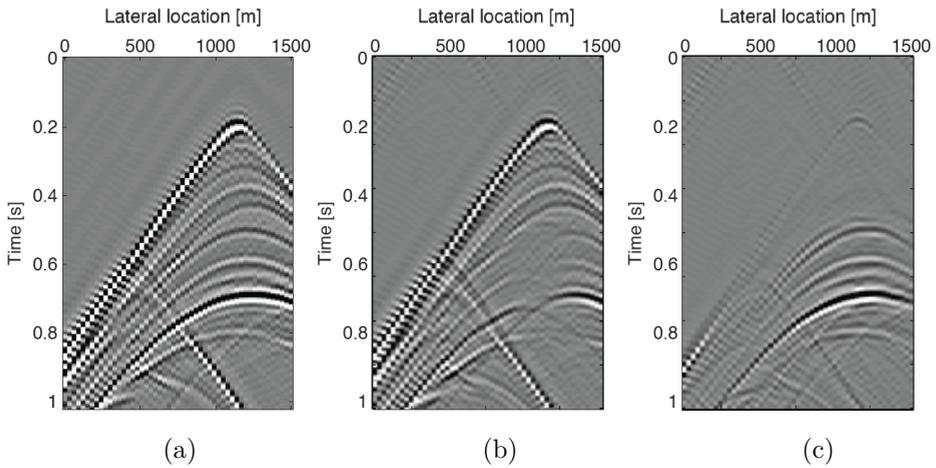


Figure 5.12: Same data as shown in figure 5.9 but now after 35 iterations. Note how well the complementary data (b) matches the measured data (a) in the parts that do not belong to the local domain. The difference (c) that is used for subsequent inversions of the particular local domain only contains events that are related to the local domain. Note how the linear event, associated with the fault in the subsurface, was nicely predicted and removed after 35 iterations. This is a desired result since the fault lies outside the local domain under consideration and reflections from this part of the subsurface should not be taken into account by local inversion.

Chapter 6

Reservoir-oriented nonlinear full waveform inversion

We have seen in the previous chapters that full waveform inversion is a good tool to derive high resolution property models of the subsurface from surface seismic data. On the other hand we realised that significant computational resources must be available to make the scheme feasible on a large scale. This is one of the reasons why application of full waveform inversion is still limited in seismic exploration although the method was discussed already in the early 1980's in this context.

One part of the seismic value chain as introduced in figure 1.3, in which full waveform inversion can have a major impact, is reservoir characterisation. This specific step is applied after a reservoir has been identified by the means of structural imaging. Because a reservoir is only a subdomain of the entire subsurface, the process becomes much more feasible from a computational point of view. This statement especially holds if the commonly used 1.5-D approximation is made. In this case the data model assumes a locally plane layered subsurface structure, which turns out to be a very robust approximation in many reservoir settings. The layering is supposed to be parallel to a target horizon, but the property values for all lateral locations are independent of each other. All layer thicknesses are equal to the vertical discretisation step (typically 5 m for acoustic data, 3-5 m for elastic data). The layered reservoir sequence to be inverted for typically has a thickness of approximately 500 m.

In this chapter we will describe nonlinear full waveform inversion as it can be directly used for reservoir characterisation. Although the principle is very similar to the scheme that was discussed in this thesis so far, some changes had to be

made. As already indicated we make a step back in terms of multi-dimensionality. While previous chapters were about 2-D full waveform inversion, we will use the local 1.5-D assumption in this chapter. The resulting scheme will then be applied to individual CMP gathers after the surface data has been redatumed to a target boundary at the top of the reservoir sequence. Although we give a full mathematical formulation of 1.5-D acoustic inversion only, aiming at compressional velocity and density, we eventually apply elastic full waveform inversion to a field dataset from the Middle East. For the sake of brevity and clarity, the full elastic equations are only shown in symbolic form. Another extension that will be made in this chapter is the usage of smooth inhomogenous background models, which is essential for realistic representation of the property models.

6.1 Acoustic 1.5-D nonlinear full waveform inversion

Any inversion algorithm relies on a data model to predict synthetic data from a known set of subsurface properties. Like in the previous chapters, our implementation of this so-called forward problem is based on the scattering approach. In the acoustic approximation the parameters that determine the propagation of wavefields are the compressibility κ and the mass density ρ . Instead of directly inverting for these parameters they can be reduced to contrasts which are defined with regard to a background medium that is given by κ_0 and ρ_0 :

$$\chi_\kappa(z) = \frac{\kappa(z) - \kappa_0(z)}{\kappa_0(z)}, \quad \chi_\rho(z) = \frac{\rho(z) - \rho_0(z)}{\rho_0(z)}. \quad (6.1)$$

In equation (6.1) we made use of the 1.5-D approximation already because the contrasts χ are functions of the local depth variable z only. The background models are smooth and represent the best *a-priori* information that is available before the inversion. It should be realised that the definition of the contrast function is based on the difference between the background medium and the true properties. This is in contrast to the definition of the reflectivity that is based on the difference between the true properties across a boundary.

In this chapter, seismic data will be mainly displayed in the plane-wave domain (τ, p) . This is also the domain in which full waveform inversion is applied. Nonetheless, derivation of the scheme will be carried out in the frequency domain and the equivalent time-domain expressions can be obtained by an inverse Fourier transform.

The Green's functions are defined in the smooth background medium with velocity $c_0 = 1/\sqrt{\kappa_0\rho_0}$ and are calculated in the WKBJ approximation as discussed

by Robinson (1986). In the frequency domain we have:

$$\mathcal{G}(z, z', p) = -\frac{1}{2}\gamma_0(z') \left[\frac{\gamma_0(z')\rho_0(z)}{\rho_0(z')\gamma_0(z)} \right]^{1/2} e^{-i\omega \int_z^{z'} \sqrt{1-c_0^2(z'')p^2/c_0(z'')} dz''}, \quad (6.2)$$

where $\gamma_0(z) = (\omega^2/c_0^2(z) + \omega^2 p^2)^{1/2}$ and p is the horizontal slowness (rayparameter). Let us further define

$$\begin{aligned} \mathcal{G}_{11}(z, z') &= \mathcal{G}(z, z'), \\ \mathcal{G}_{12}(z, z') &= \mathcal{G}(z, z')\text{sign}(z - z'), \\ \mathcal{G}_{21}(z, z') &= \mathcal{G}(z, z')\text{sign}(z - z'), \\ \mathcal{G}_{22}(z, z') &= \mathcal{G}(z, z'). \end{aligned} \quad (6.3)$$

Next we introduce two equations that describe the measured seismic data as well as the wavefields at every grid point in the object under investigation, hence in our case at every vertical location in the local plane layered domain. The data equation in the rayparameter-frequency (p, ω) domain is given by

$$\begin{aligned} P_{data}(p, \omega) &= \int_{\mathcal{D}} [\mathcal{G}_{11}(0, z, p, \omega)\chi_1(z)P_{tot}(z, p, \omega) \\ &\quad + \mathcal{G}_{12}(0, z, p, \omega)\chi_2(z)V_{tot}(z, p, \omega)] dz, \end{aligned} \quad (6.4)$$

where the upper-case field quantities indicate Fourier transforms. Equation (6.4) describes the seismic data $P_{data}(p, \omega)$ at $z = 0$ i.e. in terms of the contrast properties χ , the Green's functions \mathcal{G} and the total fields P_{tot} and V_{tot} in the subsurface domain \mathcal{D} . The contrasts χ_1 and χ_2 are combinations of the contrasts defined in equation 6.1 and can be calculated as

$$\chi_1(z) = \frac{1}{1 - c_0^2(z)p^2} \left[\chi_\kappa(z) + c_0^2(z)p^2 \frac{1}{1 + \chi_\rho(z)} \chi_\rho(z) \right], \quad \chi_2(z) = \chi_\rho(z). \quad (6.5)$$

In equation (6.5), the factor $1/(1 + \chi_\rho(z))$ is linearised by using the contrast from the previous iteration ($\chi_\rho^{(0)} = 0$). The total fields in the subsurface are given by the object equations:

$$\begin{aligned} P_{tot}(z, p) &= P_0(z, p) + \int_{\mathcal{D}} [\mathcal{G}_{11}(z, z', p)\chi_1(z')P_{tot}(z', p) \\ &\quad + \mathcal{G}_{12}(z, z', p)\chi_2(z')V_{tot}(z', p)] dz', \end{aligned} \quad (6.6)$$

$$\begin{aligned} V_{tot}(z, p) &= V_0(z, p) + \int_{\mathcal{D}} [\mathcal{G}_{21}(z, z', p)\chi_1(z')P_{tot}(z', p) \\ &\quad + \mathcal{G}_{22}(z, z', p)\chi_2(z')V_{tot}(z', p)] dz'. \end{aligned} \quad (6.7)$$

Nonlinear full waveform inversion requires that equations (6.4), (6.6) and (6.7) are satisfied. In principle one could solve for both acoustic parameters simultaneously, but this would be a very expensive task. To solve the nonlinear problem in a more efficient way we solve the data equation and the domain equation in an alternating iterative manner. The procedure is equivalent to the one described in chapter 2 with the difference that now we need to find two parameters.

A well known method to solve equations (6.6) and (6.7) is to substitute iteratively the total fields from the left-hand side into the integral on the right-hand side. The resulting sequence of estimates is called the Neumann series, which is straightforward, but has the disadvantage that it might not converge for large contrasts, high frequencies or large inversion domains. For this reason we update the total fields in an optimised and intrinsically stable manner. The scheme is based on the Krylov subspace expansion where the total fields can be written as

$$P_{tot}^{(N)}(z, p) = P_0(z, p) + \sum_{n=1}^N \alpha_N^{(n)}(p) \phi^{(n)}(z, p) \quad (6.8)$$

$$V_{tot}^{(N)}(z, p) = V_0(z, p) + \sum_{n=1}^N \alpha_N^{(n)}(p) \psi^{(n)}(z, p), \quad (6.9)$$

with basis functions $\phi^{(n)}$ and $\psi^{(n)}$ defined as

$$\phi^{(n)}(z, p) = \int_{\mathcal{D}} [\mathcal{G}_{11}(z, z', p) \partial \mathcal{W}_P^{(n)}(z', p) + \mathcal{G}_{12}(z, z', p) \partial \mathcal{W}_V^{(n)}(z', p)] dz' \quad (6.10)$$

$$\psi^{(n)}(z, p) = \int_{\mathcal{D}} [\mathcal{G}_{21}(z, z', p) \partial \mathcal{W}_P^{(n)}(z', p) + \mathcal{G}_{22}(z, z', p) \partial \mathcal{W}_V^{(n)}(z', p)] dz'. \quad (6.11)$$

In our notation we will omit the space and slowness arguments for the sake of simplicity in the remainder of this section. The difference between the contrast sources of the last two iterations can then be calculated as

$$\begin{aligned} \partial \mathcal{W}_P^{(1)} &= \chi_1^{(1)} P_0, & \partial \mathcal{W}_P^{(n)} &= \chi_1^{(n)} P^{(n-1)} - \chi_1^{(n-1)} P^{(n-2)}, & n > 1, \\ \partial \mathcal{W}_V^{(1)} &= \chi_2^{(1)} V_0, & \partial \mathcal{W}_V^{(n)} &= \chi_2^{(n)} V^{(n-1)} - \chi_2^{(n-1)} V^{(n-2)}, & n > 1. \end{aligned} \quad (6.12)$$

The weighting factors α depend on frequency and iteration N , which is equivalent to the total number of N basis function taken into account. This means that after each iteration one basis function is added, but an entirely new set of weighting factors is determined. It is important to realise that the same weighting factors are used to update both total fields. Although separate weights could be found for both total fields, our research indicated that a single set is probably sufficient. This is also why estimation of the weighting factors is only based on optimising equation (6.6) and equation (6.7) is not used in this process. Substitution of

equations (6.8) and (6.9) into the total pressure field as described in equation (6.6) yields

$$\begin{aligned} \sum_{n=1}^N \alpha_N^{(n)} \phi^{(n)} &= \int_{\mathcal{D}} [\mathcal{G}_{11} \chi_1^{(n)} P_0 + \mathcal{G}_{12} \chi_2^{(n)} V_0] dz' \\ &+ \int_{\mathcal{D}} [\mathcal{G}_{11} \chi_1^{(n)} \sum_{n=1}^N \alpha_N^{(n)} \phi^{(n)} + \mathcal{G}_{12} \chi_2^{(n)} \sum_{n=1}^N \alpha_N^{(n)} \psi^{(n)}] dz'. \end{aligned} \quad (6.13)$$

Note that we only use the object equation for the pressure field in this optimisation procedure. Both equations (6.6) and (6.7) could in principle be used but our research indicated that this is superfluous. Equation (6.13) contains the unknown weighting factors $\alpha_N^{(n)}$. If we define the operators $\mathcal{L}_1^{(q)} \phi^{(m)}$ and $\mathcal{L}_2^{(q)} \psi^{(m)}$ as

$$[\mathcal{L}_1^{(q)} \phi^{(m)}] = \phi^{(m)} - \int_{\mathcal{D}} \mathcal{G}_{11} \chi_1^{(q)} \phi^{(m)} dz', \quad (6.14)$$

$$[\mathcal{L}_2^{(q)} \psi^{(m)}] = - \int_{\mathcal{D}} \mathcal{G}_{12} \chi_2^{(q)} \psi^{(m)} dz', \quad (6.15)$$

the weighting factors can be found by least-squares minimisation of the difference between the left- and right-hand side of equation (6.13). With $m = 1, 2, \dots, N$ this leads to the following linear $N \times M$ system of equations (c.f. equation (2.33)):

$$\begin{aligned} \sum_{n=1}^N \alpha_N^{(n)} \int_{\mathcal{D}} ([\mathcal{L}_1^{(N)} \phi^{(n)}] + [\mathcal{L}_2^{(N)} \psi^{(n)}]) ([\mathcal{L}_1^{(N)} \phi^{(n)}] + [\mathcal{L}_2^{(N)} \psi^{(n)}])^* dz \\ = \int_{\mathcal{D}} \left([(\mathcal{I} - \mathcal{L}_1^{(N)}) P_0] - [\mathcal{L}_2^{(N)} V_0] \right) ([\mathcal{L}_1^{(N)} \phi^{(n)}] + [\mathcal{L}_2^{(N)} \psi^{(n)}])^* dz, \end{aligned} \quad (6.16)$$

where the identity operator is denoted by \mathcal{I} .

6.1.1 A numerical example of acoustic 1.5-D full waveform inversion

In this section we will demonstrate full waveform inversion as described here, on a synthetic acoustic dataset. Using the Kennett algorithm (Kennett, 1983), one CMP gather is modelled in the plane-wave domain (τ, p) . Keep in mind that this is also the domain in which the inversion is actually carried out. In this experiment the synthetic seismic data to be inverted is generated directly at the top of the local domain. To generate the data, directly in the plane-wave domain, we use the Kennett method, which solves the full acoustic wave equation, which means that the modelled data contains the full nonlinear relationship between the subsurface properties and the measurements. Multiple scattering, transmission effects and exact traveltimes in the true medium are properly modelled in this way.

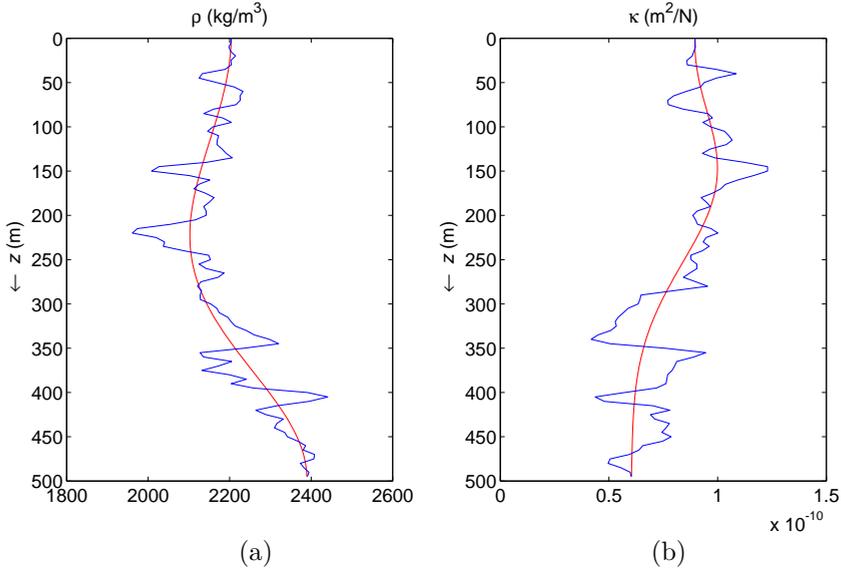


Figure 6.1: True density (a) and compressibility (b) that were used to generate synthetic data by Kennett modelling. The background properties ρ_0 and κ_0 are very smooth versions of the true properties and displayed as red lines.

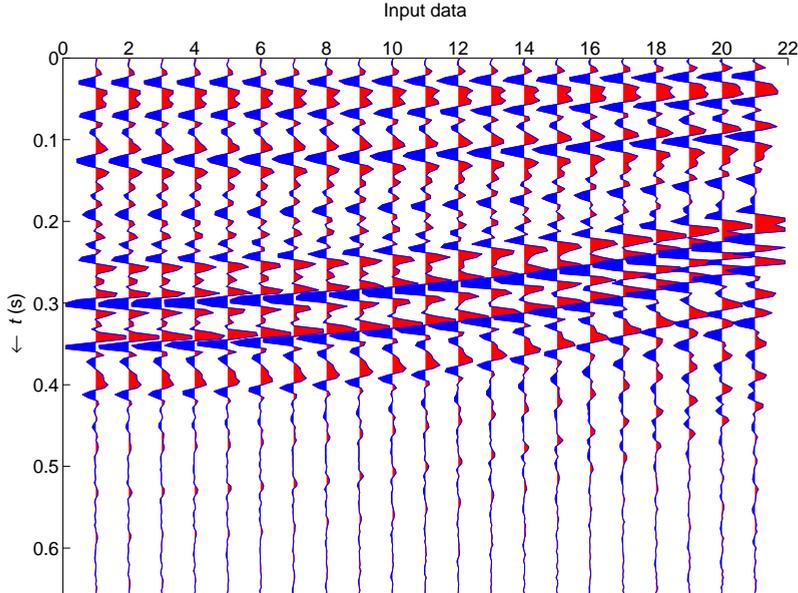


Figure 6.2: Acoustic data that was modelled in the (τ, p) -domain by the Kennett method, using the properties shown in figure 6.1.

In figure 6.1 the acoustic properties are shown that were used to generate the data shown in figure 6.2. Before inversion 5% random noise was added to the data to simulate the fact that noise is always present in real data. In the first iteration we invert the dataset by assuming that the total fields propagate only in the background media that are indicated by the red lines in figure 6.1. This approximation leads to a linear inversion result that gives a first estimate of the subsurface properties. The contrasts χ_ρ and χ_κ as defined in equation (6.1) and coming from the initial linear inversion are plotted in figure 6.3. Note that the results are spatially filtered back to the seismic bandwidth because linear inversion is not able to close the spectral gap. Although over the seismic bandwidth linear inversion leads to a reasonable result for the compressibility, the density was recovered rather poorly. This can be understood by realising that a linear data model was used to invert nonlinear seismic data and this obviously leaves an imprint on the inversion result. This is already a first indication that nonlinear inversion has to be performed if quantitative information on the subsurface is to be obtained.

In figure 6.4, the data residual after the initial linear inversion is displayed. It can be observed that the data within the primary time window is fully explained. Nonetheless, all events that scattered multiple times and consequently arrive outside of the primary time window cannot be explained by linear inversion. The reason is that the so-called coda can never be matched by a linear data model. The coda builds up by multiple scattering which is fundamentally a nonlinear phenomenon. On the other hand, it has to be understood that also multiple scattering is present within the primary time window itself. This nonlinear effects have been almost fully explained by the linear inversion although in principle it should not. The inconsistency between the used data model and the measured seismic data leaves an imprint on the inversion result, as can clearly be seen in figure 6.3.

Since we got an initial estimate of the subsurface properties, however imperfect, we can use this knowledge to update the total fields in the subsurface and thereby incorporate a first order of nonlinearity in the modelled data. We do this by adding a basis function that is the best estimate of the first order scattering as it was described in the previous section. It should be mentioned again that every basis function adds one order of scattering only, before another inversion leads to updated contrasts. This alternating procedure is repeated 60 times until the contrasts and the total fields do not change any more and convergence of the data equation and the object equation is achieved. The result of nonlinear inversion, after being spatially filtered back to the seismic bandwidth, is given in figure 6.5. This spatial filter is derived from the seismic wavelet spectrum, with the help of the average background velocity. The raw results over the full spatial bandwidth are shown in figure 6.6.

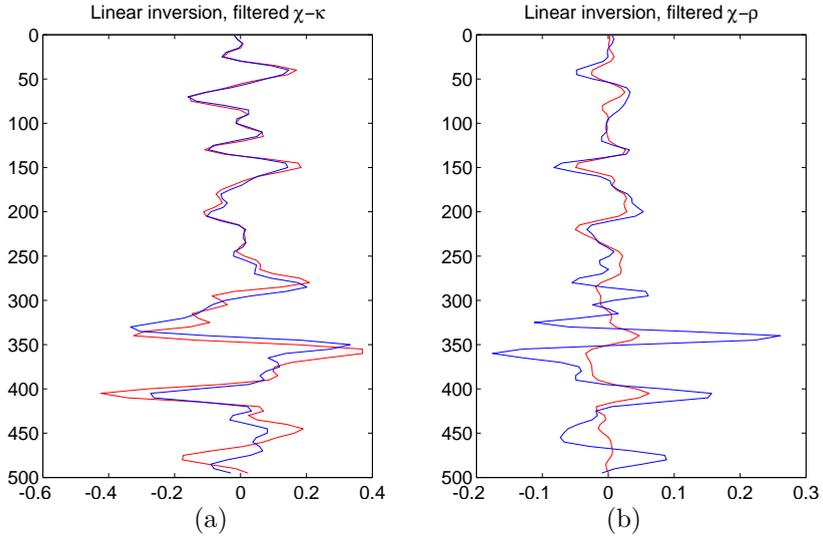


Figure 6.3: First estimate of the compressibility (a) and the density (b) coming from the initial linear inversion in which the total fields are assumed to propagate in the known smooth background media only. The results have been spatially filtered back to the seismic bandwidth.

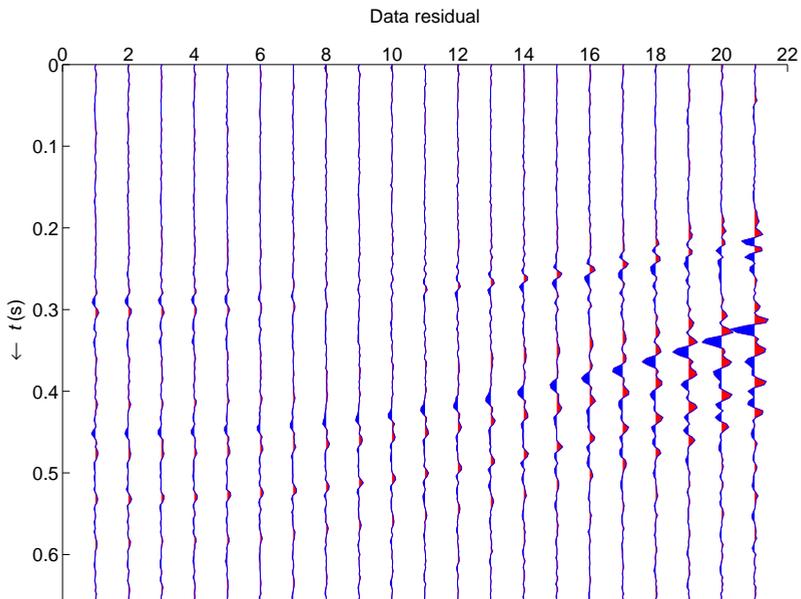


Figure 6.4: Data residual after linear inversion.

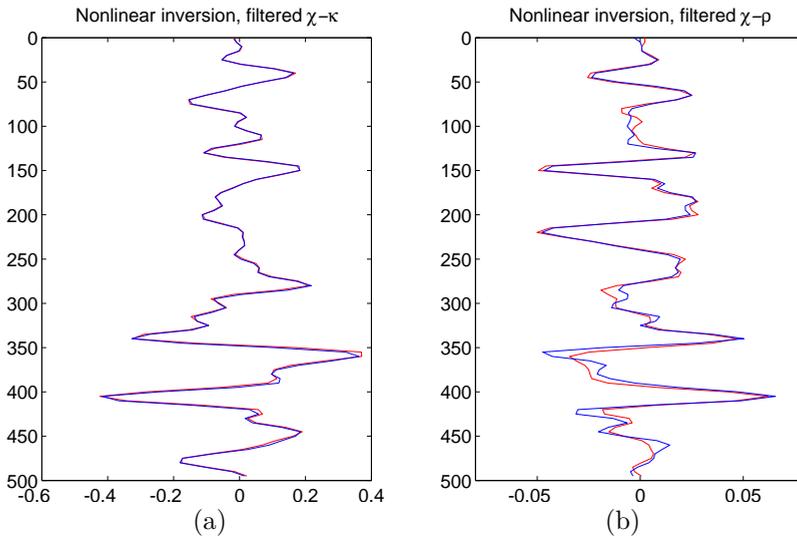


Figure 6.5: Final estimate of the contrasts in compressibility (a) and density (b) coming from nonlinear inversion in which the total fields are iteratively updated based on the best known properties. The results have been spatially filtered back to the seismic bandwidth. Note the improvement over the linear inversion results that were shown in figure 6.3.

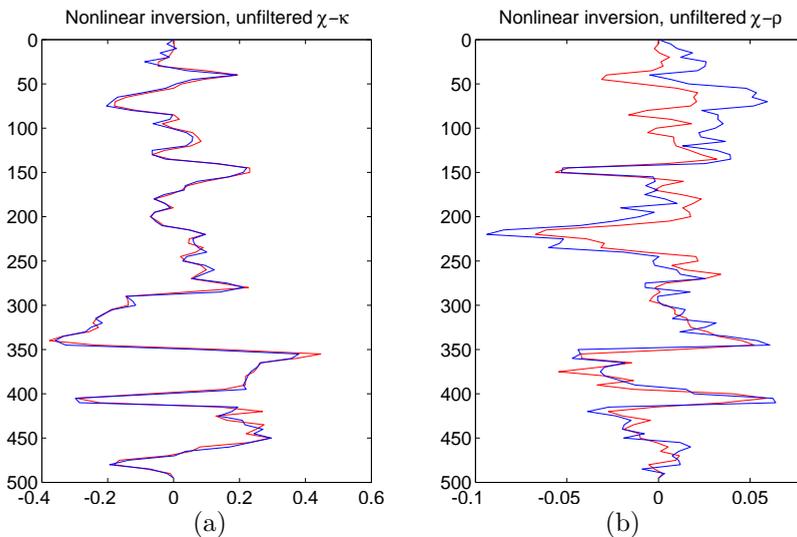


Figure 6.6: Final estimate of the contrasts in compressibility (a) and density (b) coming from nonlinear inversion but now without any spatial filter applied. Nonlinear inversion very accurately estimated quantitative properties, which cannot be achieved by linear inversion.

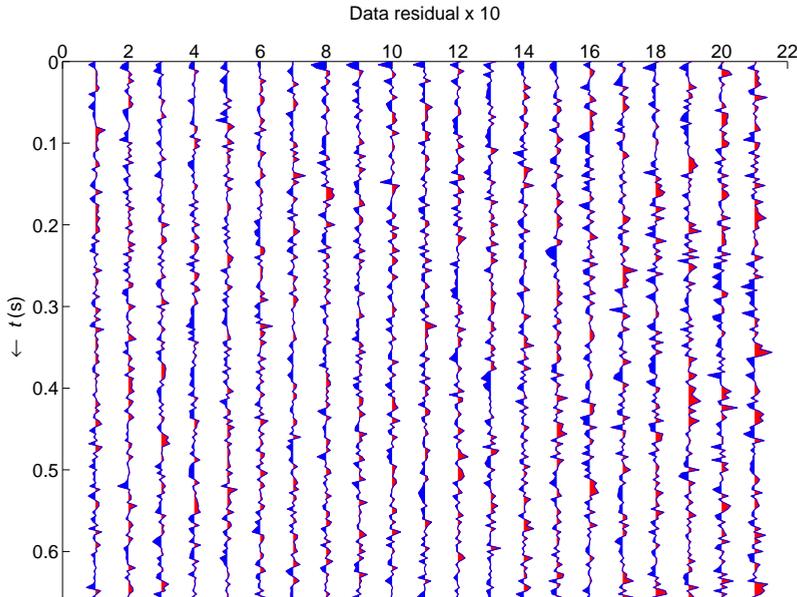


Figure 6.7: *Data residual after nonlinear inversion and enhanced by a factor of ten. The residual contains only the random noise that was added before the inversion and no indication of any data structure is left.*

A comparison between the true and the inverted properties from nonlinear inversion shows an outstanding match. It can be said that nonlinear inversion was able to recover the compressibility as well as the density nearly perfectly although 5% of random noise was added to the data before inversion. This statement holds for the case that results are spatially filtered back to the seismic bandwidth but also if no filter is applied. This already indicates that nonlinear inversion is able to provide quantitative properties. We will elaborate on this observation later in this chapter.

Looking at the data residual, given in figure 6.7, underlines this observation because also the data are fully explained by nonlinear inversion. Although the residual is scaled by a factor of ten, there is absolutely no data structure left in the residual. In contrast to linear inversion the data are fully explained over the entire time interval and also the coda is completely described. The remaining residual is just the 5% of random noise that was added to the data.

To evaluate the nonlinear inversion result even further, we make a comparison between true and inverted properties in the spatial-wavenumber domain. This allows us to relate the spatial bandwidth of the inversion result to the temporal

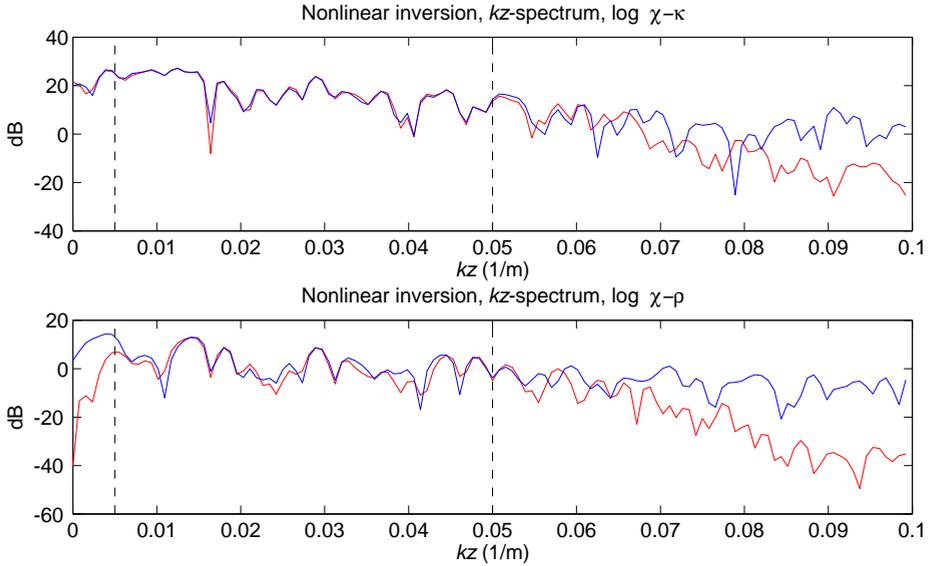


Figure 6.8: Comparison between true (blue) and inverted (red) properties in the spatial-wavenumber domain. The upper plot shows the contrasts in compressibility while the density contrasts are displayed in the lower plot. The dashed lines are a projection of the temporal bandwidth of the inverted data.

bandwidth of the measured data. For this purpose we apply a spatial Fourier transformation to convert the properties from depth to wavenumber. The comparison is then displayed in figure 6.8. It can be observed that over the seismic bandwidth, compressibility and density were estimated very accurately in both cases. Interestingly, the compressibility was even predicted very well outside the seismic bandwidth. This statement holds especially for the very low wavenumbers but also at the high end wavenumbers were recovered that were not present in the inverted seismic data. The low-wavenumber density components beyond seismic bandwidth were not recovered by nonlinear inversion, suggesting that this information is not in the data, or masked by noise.

The contrast between the DC component and the lowest frequency in the data are often referred to as the spectral gap. It covers information that is too high frequency to be captured by the kinematics, but still below the frequency spectrum of the data. While conventional linear schemes cannot recover this information, nonlinear inversion obviously can, at least for the compressibility. This is a great benefit of nonlinear inversion, while linear inversion will only be able to get at best a reasonable result within the bandwidth of the seismic data, assuming the problem is sufficiently linear. For the results shown in figure 6.3 this is clearly not the case.

6.2 Input for reservoir-oriented inversion from surface data

So far we have described and successfully demonstrated a 1.5-D nonlinear full waveform inversion scheme that was able to recover the properties at one CMP location as function of depth from measured seismic data. In principle this allows us to apply the method in a reservoir-oriented mode to get information about an extended area by repeating these inversions for a number of CMPs. However, since seismic data are usually measured with an acquisition geometry along the earth's surface, the data has to be processed to be suitable for reservoir-oriented inversion. In other words, it has to be ensured that the dataset that is inverted, contains information about the reservoir only and reflections from the overburden should not be taken into account for this process.

If reservoir-oriented inversion is supposed to be carried out, in principle the most suitable acquisition would be to locate sources and receivers as close to the reservoir as possible. This kind of seismic data are called proximity measurements and they require the presence of various horizontal wells above the reservoir in which the sources and receivers can be positioned. Although this technique is currently under development it is by far not common practice yet and other alternatives have to be thought of to convert the measured surface data into a suitable input for reservoir-oriented inversion.

6.2.1 Reservoir-oriented inversion by refocusing (redatuming, double back-propagation, downward projection, etc.)

For deeper targets it is necessary to re-position seismic sources and receivers from the actual location where the measurements were performed to a target boundary just above the objective sequence. An example where the seismic data are backpropagated into the subsurface, is the near-surface problem (Berryhill, 1984; Shtivelman and Canning, 1988; Haffinger and Verschuur, 2012). Here a virtual dataset is created along a boundary below the near surface, to remove the distortions that are generated in the very shallow part of the subsurface.

In chapter 3 we explained how measured seismic surface data can be backpropagated into the subsurface. Physical measurements with sources and receivers along the surface are transformed into a virtual dataset that should equal seismic data that was measured just above a target interval. Note that in this context the terms backpropagation, refocusing and re-datuming might be used equivalently. The concept of backpropagation for reservoir-oriented inversion is sketched in figure 6.9.

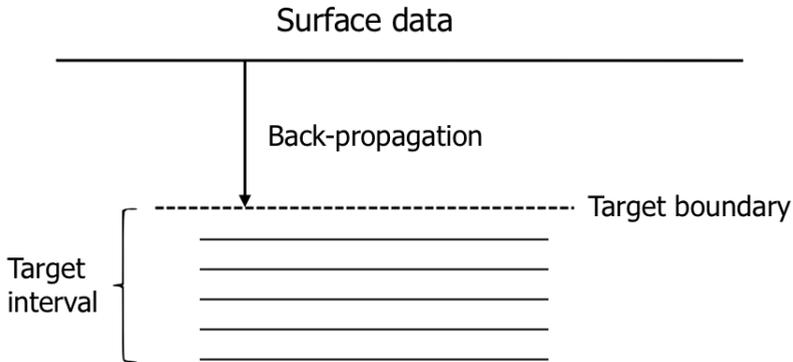


Figure 6.9: *By backpropagation surface data can be directly brought to a target boundary in the subsurface. Subsequent inversion will yield the properties of a target interval below the target boundary. Note that the target interval is only a limited subset of the full subsurface.*

6.2.2 Reservoir-oriented inversion by demigration

Previously we discussed that backpropagation can be used to transform a seismic dataset that was acquired along the earth's surface into a dataset that is suitable for reservoir-oriented inversion. Next we propose demigration as a tool to get a dataset as input to a reservoir-oriented scheme. Demigration has the huge advantage that it uses migrated angle stacks as starting point. It must be realised that oil companies or their contractors put a major effort in processing seismic data with the aim to obtain a high-quality migrated angle stack. By using these angle stacks we can build on all processing steps that were already applied previously. Migration of seismic data involves backpropagation of seismic sources and receivers into the subsurface where these backpropagated datasets are cross-correlated to obtain an image amplitude. A detailed description of seismic migration was given by Berkhout (1982). The image amplitudes as a result of migration are a function of space only because only the zero-leg of the cross-correlation is used, which removes any traveltimes. By using an available velocity model these images can be converted to pseudo time but the result is still fundamentally different from true time-domain data. If time-domain data at a boundary is desired, the traveltimes, which were removed by the migration process, have to be restored. This can be achieved by a so-called demigration. In principle demigration is only defined in a plane layered model where plane waves can simply be timeshifted. This is precisely the local model we adapted for target-oriented inversion. The demigration process for reservoir-oriented inversion is sketched in figure 6.10. The time shifts are calculated in the background and are angle (rayparameter) dependent.

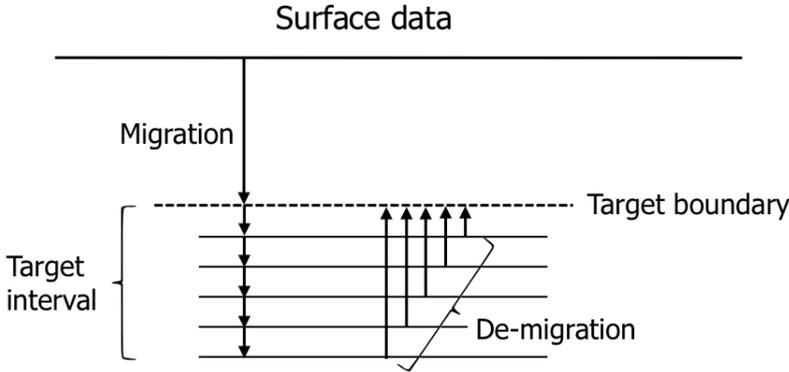


Figure 6.10: *Migration brings the data to every depth level in the target interval. Time-domain data at a target boundary above this interval can be generated by demigration. Demigration effectively restores the traveltimes between the target boundary and the other depth levels.*

6.3 Reservoir-oriented elastic full waveform inversion: A real data example

Earlier in this chapter nonlinear acoustic full waveform inversion has been described for a reservoir-oriented setup. Furthermore, it was explained how surface data can be transformed into a dataset that is suitable for reservoir-oriented inversion. In this section we apply this concept including the extension to elastic full waveform inversion to a real dataset from the Middle East.

In the following we will give a brief description of the elastic framework for full waveform inversion. No purpose is served by filling many pages with a full mathematical description of all elastic full waveform operators. The main purpose of this section is to demonstrate that target-oriented elastic inversion is feasible.

6.3.1 The elastic scattering framework

Our elastic full waveform inversion is again based on the scattering formulation and the equations for the full 3-D problem can be found in Yang et al. (2008), who based their equations mainly on the book of de Hoop (1995). Elastic inversion aims at recovering the Lamé parameters λ and μ . If the elastic problem is described in terms of the scattering formulation these stiffness parameters are transformed to compliances as:

$$K = \frac{1}{\lambda + 3\mu} \quad M = \frac{1}{\mu}, \quad (6.17)$$

where K is the well known compressibility, equal to one over the bulk modulus. These properties are related to the compressional velocity v_p , shear velocity v_s and density ρ as

$$v_p = \sqrt{\frac{1}{\rho} \left(\frac{1}{K} + \frac{4}{3M} \right)}, \quad v_s = \sqrt{\frac{1}{\rho M}}. \quad (6.18)$$

Instead of inverting directly for K , M and ρ , we solve for contrast functions depending on the backgrounds K_0 , M_0 and ρ_0 and defined as:

$$\chi_K(z) = \frac{K(z) - K_0(z)}{K_0(z)}, \quad \chi_M(z) = \frac{M(z) - M_0(z)}{M_0(z)}, \quad \chi_\rho(z) = \frac{\rho(z) - \rho_0(z)}{\rho_0(z)}. \quad (6.19)$$

Similar to the acoustic case, the backgrounds are smooth media in which the WKBJ approximation can be used to calculate the Green's functions. In the 1.5-D elastic case we need to solve four object equations, related to two elements of the stress tensors τ_{zz} , τ_{zx} and two elements of the particle velocity vector v_z and v_x . In principle there would be six object equations but since we use the 1.5-D approximation, the remaining stress tensors τ_{xx} , τ_{yy} can be expressed in terms of the other four, whereas τ_{yx} and τ_{yz} are zero. The full system of object equations then becomes:

$$\tau_{zz} = \tau_{zz,0} + \mathcal{L}_{11}\tau_{zz} + \mathcal{L}_{12}\tau_{zx} + \mathcal{L}_{13}v_z + \mathcal{L}_{14}v_x \quad (6.20)$$

$$\tau_{zx} = \tau_{zx,0} + \mathcal{L}_{21}\tau_{zz} + \mathcal{L}_{22}\tau_{zx} + \mathcal{L}_{23}v_z + \mathcal{L}_{24}v_x \quad (6.21)$$

$$v_z = v_{z,0} + \mathcal{L}_{31}\tau_{zz} + \mathcal{L}_{32}\tau_{zx} + \mathcal{L}_{33}v_z + \mathcal{L}_{34}v_x \quad (6.22)$$

$$v_x = v_{x,0} + \mathcal{L}_{41}\tau_{zz} + \mathcal{L}_{42}\tau_{zx} + \mathcal{L}_{43}v_z + \mathcal{L}_{44}v_x. \quad (6.23)$$

In the above equation, the integral operators \mathcal{L}_{ij} contain the Green's function from any point in the domain to any other point. The 16 operators \mathcal{L}_{ij} contain combinations of the 24 Green's functions that relate the six different source elements to the four field quantities. To make the integral formulation of the elastic wave equation complete, we need the two data equations that relate the total fields to the measured datasets PP and PS if available. Instead of using PP and PS directly the measured data can be expressed in the scalar potentials Φ and Ψ as:

$$\Phi = \mathcal{K}_K^\Phi \chi_K + \mathcal{K}_M^\Phi \chi_M + \mathcal{K}_\rho^\Phi \chi_\rho \quad (6.24)$$

and

$$\Psi = \mathcal{K}_K^\Psi \chi_K + \mathcal{K}_M^\Psi \chi_M + \mathcal{K}_\rho^\Psi \chi_\rho, \quad (6.25)$$

respectively. The operators \mathcal{K} , which relate the contrasts to the measured data, contain the total fields. With the same reasoning as already discussed in the acoustic case, we solve the domain equations and the data equations in an alternating manner to obtain a nonlinear inversion result.

6.3.2 A field-data example from the Middle East

To demonstrate the feasibility of our elastic inversion approach we applied it to a field dataset from Saudi-Arabia. The data used in this case study was kindly provided by the national oil company Saudi-Aramco. The data was measured with an ocean bottom cable (OBC) configuration where the receivers lie on the sea bottom and a seismic vessel, towing a source, shoots while sailing across or along these receiver lines. In total a 3-D narrow azimuth survey was acquired and every shot was measured by four receiver lines. Although the pressure as well as the three components of the particle velocity were measured, for inversion only the pressure component after PZ-summation (Loewenthal et al., 1985; Hoffe et al., 2000) was used. In this way the receiver ghost can be attenuated (Barr, 1997). Beside the seismic data, logged well information was available that is used to calibrate the data at the target level and to determine an angle-dependent wavelet as input for the inversion. Although not discussed specifically from a theoretical point of view, estimation of an angle-dependent wavelet allows us to calibrate the data with respect to angle dependent transmission effects from the overburden, in an optimal way. A single shot gather and a stack of the surface data is displayed in figure 6.11.

In this case study we go the demigration route to get input data that is suitable for target-oriented inversion. We were provided with eight migrated angle stacks. Several preprocessing steps were applied to the data before migration, including removal of the mud roll and surface related multiple elimination. Removal of surface-related multiples is of special interest because this type of multiples is not generated within the reservoir sequence and therefore cannot be predicted by our inversion scheme. An example demigrated angle stack is shown in figure 6.12. The section was dipfiltered after demigration to suppress high frequency noise.

The result of migration are reflection amplitudes which are each associated with a certain point in the subsurface. Using a velocity model, e.g. velocities from semblance analysis, a spatial image can then be converted to two-way vertical traveltimes. To obtain true time-domain data from a migrated dataset, a demigration has to be applied to restore the traveltimes that were previously removed by the migration process. In figure 6.13 a single CMP gather is shown before and after demigration. After migration, but before demigration, the angle gathers are flat. The different traveltimes, dependent on the angle of incidence, are removed by migration because the image amplitudes are all related to the same depth point. By demigration the true traveltimes of the data can be restored, which results in a real time-domain dataset at the target boundary. The data shown is related to a CMP location close to the well and obviously this process is carried out for all CMPs. In this way a full time-domain (τ, p) dataset at the target boundary can be generated and used as input to our nonlinear inversion scheme.

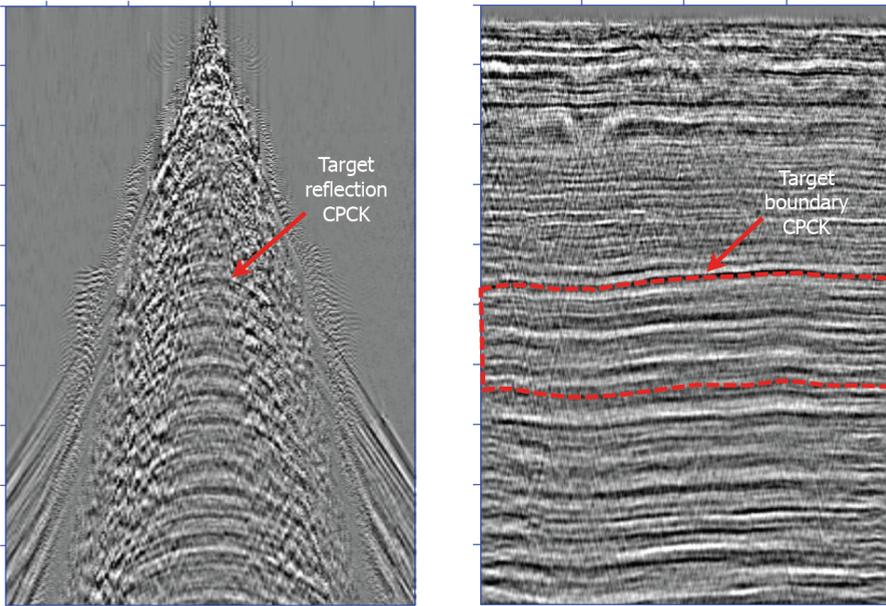


Figure 6.11: Field data used to demonstrate reservoir-oriented nonlinear full waveform inversion. On the left a shot gather is displayed while on the right a stack of the surface data is shown. The arrows indicate the reflections from the target boundary while the red box renders the reservoir area for which the elastic parameter will be determined.

Like in every other inversion scheme, our algorithm needs knowledge of the wavelet for calculation of the incident field. Wavelet estimation and calibration of the data to account for transmission effects in the overburden can be carried out with the help of the available well information. We were provided a full suite of logs including, ρ , v_p and v_s . The logged properties at the well location are shown in figure 6.14. What stands out are the extremely strong velocity contrasts in the upper part of the section as well as the very high P-velocities reaching a maximum of 5000 m/s. The strong variations in the logged properties are related to a very soft shale, sandwiched between carbonates and sand. These strong reflectors are expected to generate notable internal multiples, which would make the dataset very nonlinear.

To estimate the wavelet we model broadband (0-5-90-100 Hz) synthetic data based on the logged properties. For generating the synthetics we use the Kennett algorithm which is an exact implementation of the wave equation in strongly inhomogeneous layered media. A wavelet can then be estimated by least-squares matching of the broadband synthetics and the real data for a CMP located at the

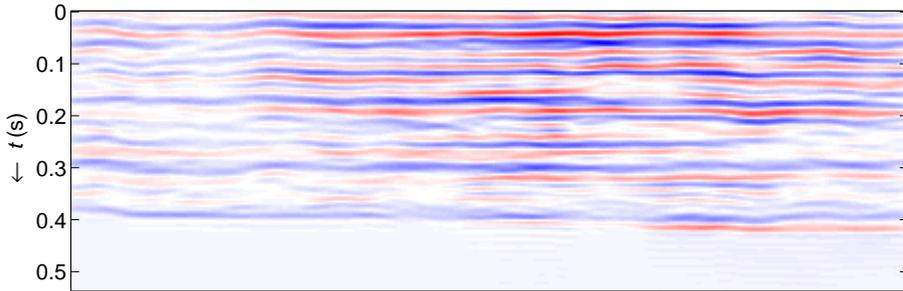


Figure 6.12: *An example demigrated angle stack for the highest available p -value. Note that $t = 0$ corresponds to the target boundary as indicated in figure 6.11.*

well location. Matching seismic to well synthetics is a well known procedure and an example is given by White (2003).

In figure 6.15 we display the broadband well-synthetics, the seismic data for a CMP location very close to the well, and the result of the seismic-to-well match. The seismic-to-well match was performed for each angle while the matching filter, hence the wavelet, was kept as short as possible in time. In this way a p -dependent wavelet was estimated, taking the transmission in the overburden for each different angle of incidence into account.

In figure 6.16 the estimated wavelets are shown. Although the estimated wavelets are quite acceptable, it becomes clear that the seismic-to-well match is rather poor for this dataset. While the match is reasonable for the very shallow part, where the very strong reflectors appear, the lower part is hardly matched at all. An evident reason is, that least-squares matching is dominated by the strong events which are generated in the shallow part. Another reason might be that amplitudes were not fully preserved during preprocessing. This could be a result of applying a gain to the data to balance the amplitudes over a certain interval. Also an imperfect PZ-summation and overburden multiples could be reasons for a poor seismic-to-well match. On the other hand, since the CMP location of the real data does not exactly coincide with the location of the well, some mismatch, although on a much smaller order, could be expected. For the moment there is no other choice than continuing the study while keeping the poor seismic-to-well match.

Another required input to our inversion algorithm are background models for K , M and ρ in which the Green's functions and the incident fields are calculated. From the true logs we extract very smooth background models by applying a high-cut filter with corner frequencies of $0 - 0 - 2.7 - 5$ Hz. Note that we translate these temporal frequencies into spatial wavenumbers based on the average velocity that appears over the objective section. The low-wavenumber backgrounds

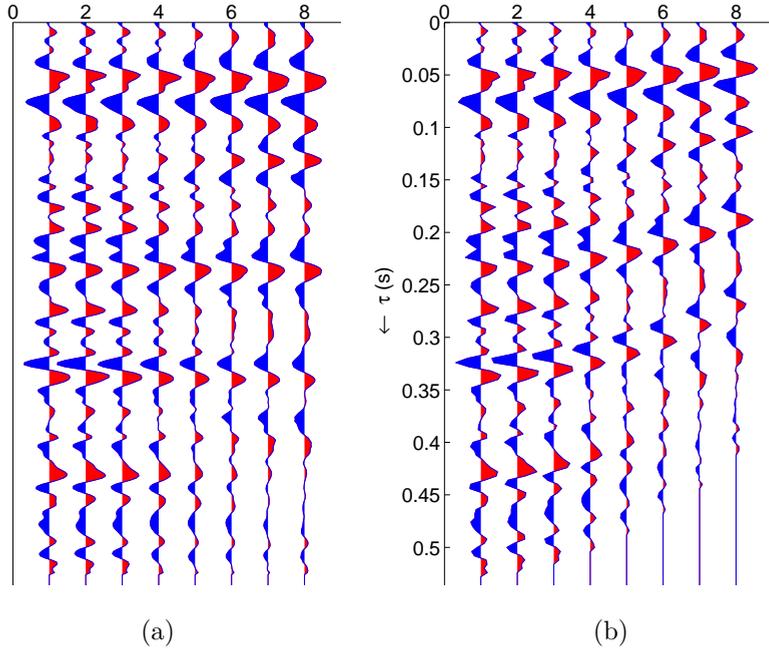


Figure 6.13: A single gather, for a CMP close to the well location, before (a) and after (b) demigration. While before demigration, related events have equal pseudo traveltimes, the true traveltimes are restored by the demigration process. Traveltimes are dependent on slowness, while for smaller p -values arrival times increase.

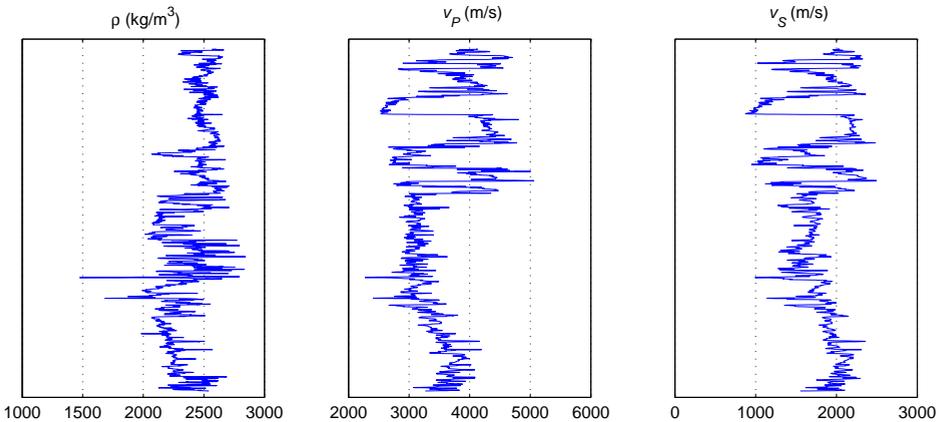


Figure 6.14: Logged properties which will be used for estimation of a wavelet, calibration of overburden effects and extraction of smooth background models. The property logs have a 0.5 ft sampling which is equivalent to 0.1524 m.

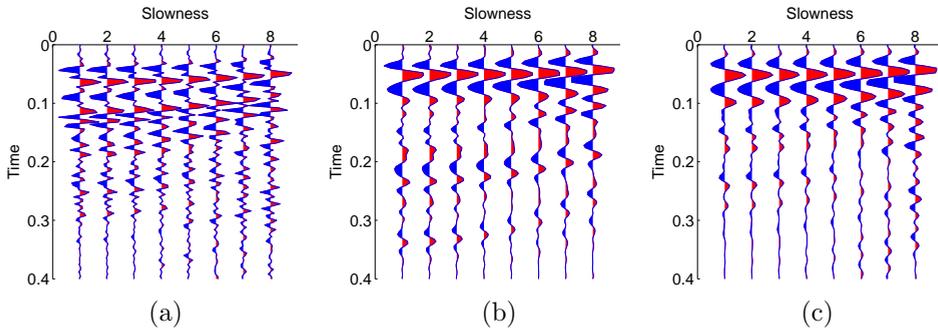


Figure 6.15: Based on the logged properties broadband synthetics (a) were generated using Kennett modelling. By least-squares matching the synthetics to the real data (b) for a CMP at the well location a wavelet can be estimated. The matched synthetics are shown in (c).

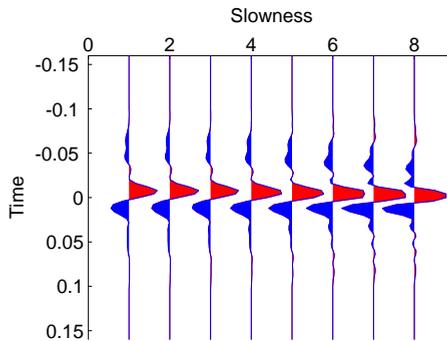


Figure 6.16: Angle dependent wavelets as a result of the seismic-to-well match.

are very smooth and do not contain any structural *a-priori* information. We keep the backgrounds constant for all CMPs. When more wells would be available, a low wavenumber background model could be created by interpolation. Other possibilities for background estimation as a function of CMP location and depth could be investigated. These might be based on analysing the kinematics of the data to obtain backgrounds for the propagation velocities as well as gravity methods for extraction of variable density backgrounds. However, this is beyond the scope of this thesis. To allow a fair comparison between logged properties and inverted properties we apply a high-cut filter to the log data. In this way the resolution becomes similar to what can be expected from full waveform inversion on a 5 m grid. The high-cut filtered logs including the backgrounds are shown in figure 6.17.

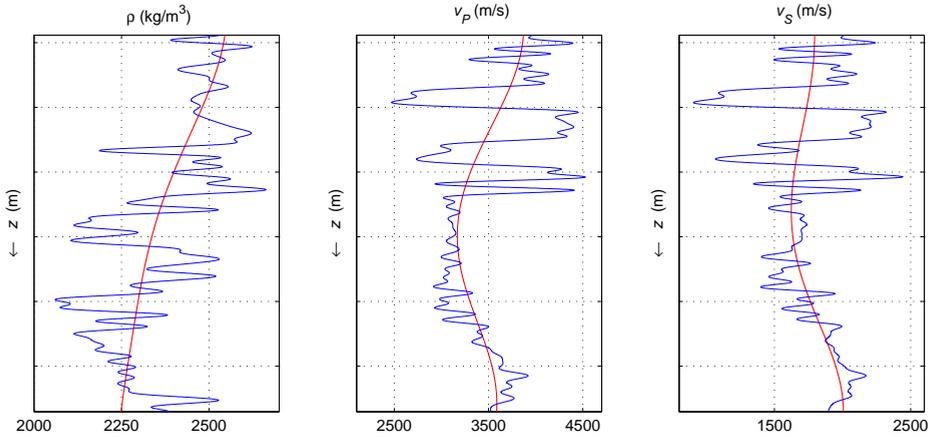


Figure 6.17: *Since inversion results will be on a 5 m grid we apply a high-cut filter to the logged properties to allow a fair comparison. The red lines indicate the backgrounds which are very low-wavenumber equivalents of the well. Note that the backgrounds are very smooth and do not contain any structural a-priori information.*

With the backgrounds obtained we are now able to apply nonlinear full waveform inversion to our demigrated dataset. Since for this case study only PP-data are available, we will invert for M and K only while using the density from the log as a fixed property during the iterations. The poor seismic-to-well match already indicated an insufficient data quality to support inversion for all three elastic parameter. If data quality is not good enough to reliably estimate all three elastic parameters, ρ needs to be skipped first, followed by v_s in the case of very low data quality.

Our inversion approach updates the contrasts and the total fields in an alternating manner. The scheme is initialised with a total field propagating in the background media only. This is the so-called Born approximation which ignores nonlinear effects as multiple scattering, transmission and the traveltimes in the true medium. The results of linear inversion are displayed in figure 6.18.

The property estimates χ_K and χ_M as a result of linear inversion come out better than expected. This is surprising because the strong contrasts in the shallow part of the target section were expected to generate a rather nonlinear dataset. However, analysing the linear results a bit more in detail, some interesting observations can be made. The inversion result for χ_K is considerably better than the result for χ_M , which is rather unstable. This is related to the fact that the parameter K is mainly determined by the kinematics of the data while M is more related to the angle dependency of the amplitudes. Obviously, the kinematics constrain linear inversion much better than the amplitudes do. It would be possible to improve

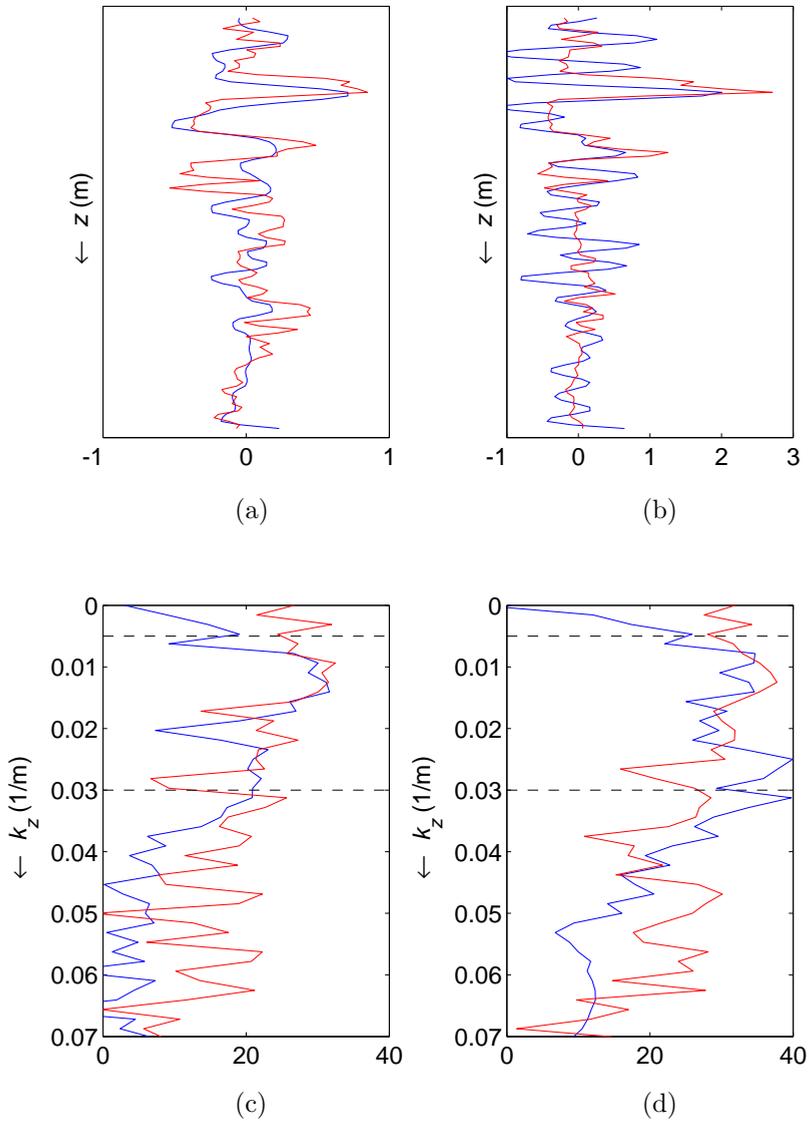


Figure 6.18: Linear inversion results after the first iteration. The true contrasts at the well location are displayed in red and the inversion results are shown in blue: (a) χ_K in depth domain, (b) χ_M in depth domain, (c) χ_K in spatial-wavenumber domain, (d) χ_M in spatial-wavenumber domain.

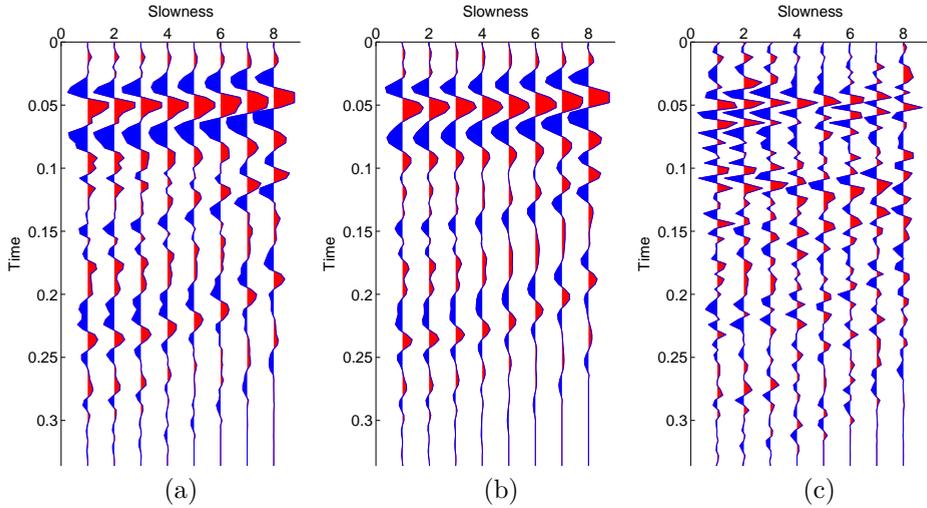


Figure 6.19: The demigrated data at the well location (a) is explained satisfyingly by the predicted data coming from linear inversion (b). The residual (c) is blown up by a factor of five and shows some strong events. Taking the poor seismic-to-well match into account, this could be expected.

the result by applying some damping to the inversion. However, this requires the choice of a damping parameter which makes the result inevitably dependent on the user. This is in strong contrast to our nonlinear inversion scheme where there are virtually no user parameters that directly affect the result. Furthermore, χ_M reaches values smaller than -1 which is not allowed from a physical point of view.

By applying a spatial Fourier transform to the results, we can plot the properties in the wavenumber domain. This allows us to judge the obtained properties in terms of spectral bandwidth and in relation to the temporal frequencies that are present in the seismic data. Again, conversion from temporal frequency to spatial wavenumber is based on the average velocity over the objective sequence. The seismic bandwidth in plots 6.18 (c)-(d) is indicated by the dashed black lines. It can be observed that linear inversion leads to plausible results within the seismic bandwidth while outside the bandwidth the results are very poor. For the low wavenumbers this is the so-called spectral gap between the bandwidth of the background model and the seismic bandwidth. This gap cannot be closed by linear inversion and, in that case, would have to be filled by *a-priori* information that is available for inversion.

In figure 6.19 we compare the measured data at the well location with the predicted data coming from linear inversion. The input data are explained well by the predicted data although the residual, blown up by a factor of five, shows quite

strong events. Because the seismic-to-well match was very poor, this is not very surprising.

In a next step we use the inverted properties to update the total fields at every point in the inversion domain according to the description that was given earlier. We repeat the process 15 times, meaning that the inversion and the field update are performed 15 times. Finally, the nonlinear inversion result is based on a total field that consists of 15 basis functions, taking the nonlinear relationship between the properties and the seismic data into account, to the 15th order of scattering. The results of nonlinear inversion are displayed in figure 6.20.

Nonlinear inversion utilises the full nonlinear relationship between the subsurface properties and the seismic data. This can be clearly observed in figure 6.20 because the results are significantly improved in comparison with the linear inversion results. The contrasts χ_K and χ_M follow the true properties much closer, particularly for χ_M . An interesting observation can be made by plotting the results in the spatial-wavenumber domain. Nonlinear inversion is able to estimate the properties far beyond the wavenumber equivalent of the frequency spectrum of the seismic data. The low wavenumbers are estimated nearly down to the DC component and the spectral gap has been entirely closed. Furthermore, the high wavenumbers were estimated far beyond the highest temporal frequency in the input data. It should be realised that closing the seismic gap allows us to obtain absolute properties of the subsurface and quantitative imaging becomes possible. This is of special importance for reservoir characterisation where quantitative properties are needed for good decision making.

As previously, a good match between input data and inverted data is obtained (see figure 6.21), although the overall residual increased. This is a very intriguing result that goes against intuition, because more accurate properties would raise the expectation of a better data fit. That the opposite is the case can be understood in the following way. Nonlinear inversion, together with our regularisation scheme that minimises the variation in the parameters, has the potential to reject noise in the data that cannot be explained by the nonlinear data model. Seismic events that do not fulfil the wave equation will not be explained and end up in the residual. This makes nonlinear inversion very much superior to linear inversion that cannot distinguish between signal and noise. The beauty of nonlinear inversion is that the process itself will tell what part belongs to the data and what can be considered as noise, in the sense that it does not obey the wave equation in the local domain.

Finally we apply elastic full waveform inversion along a line, consisting of 101 CMPs with 25 m separation each. Although inversion of each CMP is based on the 1.5-D assumption, by combining the results a full 2-D image is obtained. In figures 6.22 and 6.23 we plot the results of linear and nonlinear inversion for K

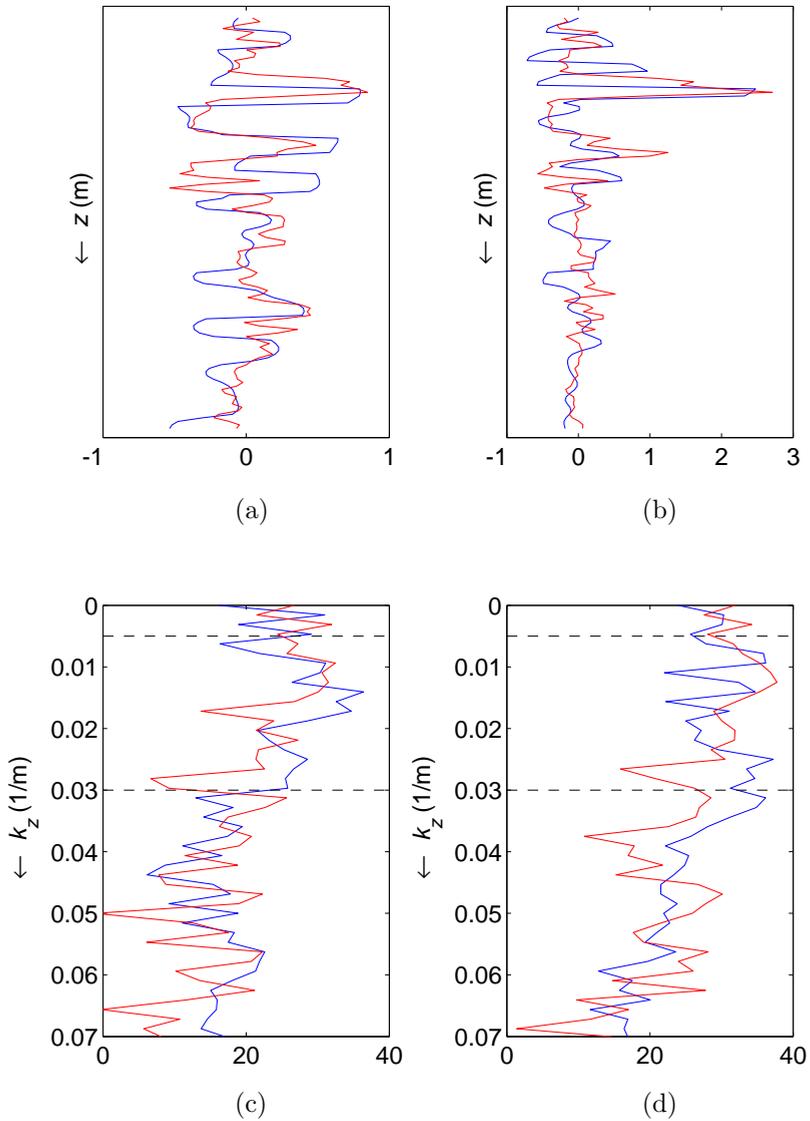


Figure 6.20: Nonlinear inversion results after 15 iterations. The true contrasts at the well location are displayed in red and the inversion results are shown in blue: (a) χ_K in depth domain, (b) χ_M in depth domain, (c) χ_K in spatial-wavenumber domain, (d) χ_M in spatial-wavenumber domain.

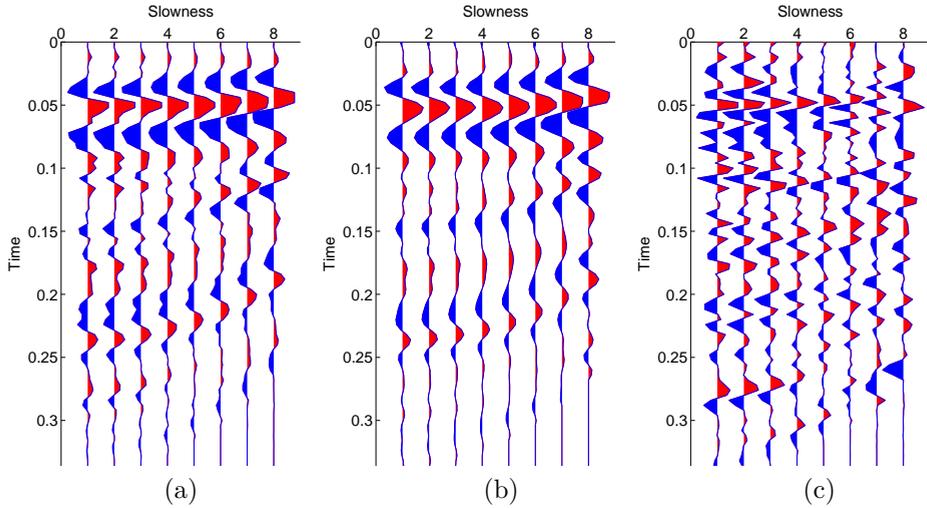
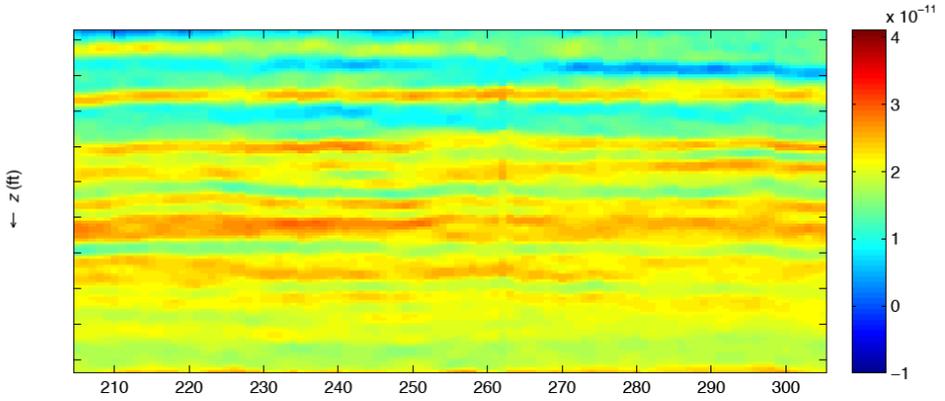


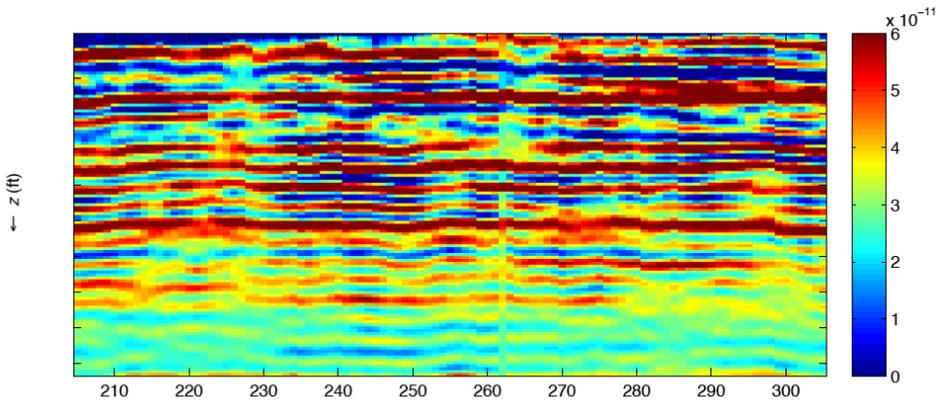
Figure 6.21: Shown is the demigrated data at the well location (a), predicted data from nonlinear inversion (b) and the residual (c) blown up by a factor of five. Interestingly, the residual from nonlinear inversion is larger than the one from linear inversion. This indicates the capability of nonlinear inversion to reject noise in the data that does not obey the nonlinear data model.

and M . Again, it can be observed that the difference between linear and nonlinear inversion is smaller for K as it is for M . Because a nonlinear data model explains the amplitudes much better than a linear data model, a significantly better subsurface model in terms of M can be obtained by this method.

For some lateral locations a slight instability for the DC component of the inversion result can be observed. This could easily be removed by filtering the very low spatial wavenumbers (up to 1 Hz). The filtered part should then be replaced by the background model. In this way no significant spectral gap would be introduced but a laterally much more consistent result is obtained.



(a)



(b)

Figure 6.22: Linear inversion results for absolute K (a) and absolute M (b). The 2-D image is composed of 101 CMPs. While M was recovered reasonably well, the result for Q is very unstable.

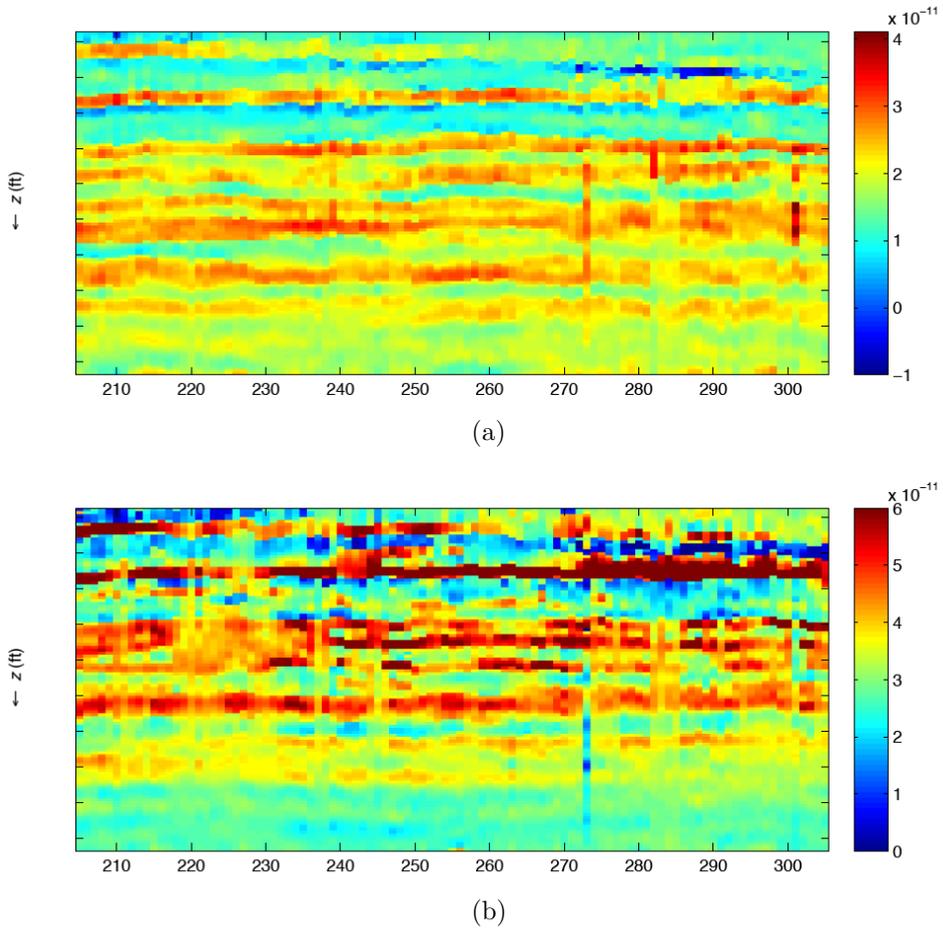


Figure 6.23: Nonlinear inversion results for absolute K (a) and absolute M (b). The 2-D image is composed of 101 CMPs. Compared to the linear inversion results, shown in figure 6.22, a much more stable result is obtained for M . Nonlinear inversion significantly improved the resolution which allows better definition of the structure, but also quantitatively the results are superior. This was also discussed and demonstrated in figure 6.20

Conclusions and recommendations

In this thesis we demonstrated a full waveform inversion approach to recover property models of the subsurface with a resolution that goes beyond the one of current imaging techniques. The current acoustic implementation inverts for the compressional velocity only while the density is assumed to be known and constant. Furthermore, in the 2-D case we used homogeneous background models so far because it allowed us to use analytic Green's functions instead of numerically calculating them. While in chapter 2 the problem was mathematically formulated in 2-D, we subsequently applied the discussed methodology to synthetic seismic datasets as well as to a field dataset from the Middle East.

7.1 Conclusions

In chapter 3 of the thesis we have shown that by backpropagation of seismic sources and receivers into the subsurface the inversion process can be localised. The major findings of this part are:

- By localising the inversion process, the memory requirements of full waveform inversion can be reduced significantly.
- Localised full waveform inversions can be performed on a significantly finer grid compared to direct inversion of surface data.
- Splitting up the large scale inversion problem into many small scale local inversion allows us to use the full seismic bandwidth (50 Hz and beyond).
- Local inversion results can be combined to a high resolution global property model.

In chapter 4 we illustrated that by performing a global total wavefield update after each iteration, to replace the initial assumption that the total wavefield is equal to the wavefield in the background, the full nonlinear relationship between the subsurface properties and the seismic data was utilised. The main outcomes of this chapter are:

- Quantitative broadband properties can be estimated from a bandlimited seismic measurement.
- Nonlinear full waveform inversion has the potential to image structural features that cannot be revealed by linear schemes, e.g. linear inversion or migration.
- The key to improved imaging is to include nonlinear effects like multiple scattering, transmission and the exact traveltimes in the inverted medium in the inversion process.
- Doing so we are able to image even vertical reflectors because multiple scattering introduces illumination that can never be explained by a linear data model.

In chapter 5 we combined localised full waveform inversion with the global total field update. Taking multiple scattering into account also implies that an interaction between the local domains occurs. During the first iteration the interaction between separate local domains has to be ignored due to the fact that the subsurface is still unknown. Because we iteratively estimate the subsurface properties and the wavefields at every location in the subsurface, we are able to compute the inter-domain scattering leading to the following discoveries:

- We can successfully subtract the predicted data, that does not belong to the domain under consideration, from the measured data before localisation.
- The imprint of overburden multiples can be removed while at the same time revealing the coda of a local domain that could potentially contribute to the inversion result at later iterations.
- By combining localised inversion, the nonlinear wavefield update and the subtraction of inter-domain effects, an exact inversion scheme is obtained without making any approximations.

In chapter 6 we used nonlinear full waveform inversion, based on the previously discussed concepts, for reservoir characterisation. In this case study a field dataset from the Middle East was used. Reservoir-oriented full waveform inversion was based on the 1.5-D assumption which is well known to be accurate in many reservoir studies. Although making a step back in terms of dimensionality of the inversion, we also made steps forward by using smooth inhomogeneous backgrounds and extending the scheme to the full elastic case. This first application

of the discussed methodology to real data gave us the following very valuable insights:

- Although the seismic-to-well match in this case study was very poor, nonlinear full waveform inversion turned out to be extremely robust.
- Nonlinear inversion increased the resolution and also performed better in estimating quantitative properties, compared to linear inversion.
- Nonlinear full waveform inversion seems to have the potential to reject noise in the data that cannot be explained by the nonlinear data model.

7.2 Recommendations

Although in this thesis we aimed to find answers to several scientific questions, at the end open questions remain. This is related to the fact that a research project is supposed to be carried out in a specified time frame, but also because researching brings up new topics that are interesting to be investigated in more detail. In the following we discuss some potential research that could help to improve the work presented, especially with focus on real data applications.

The major part of this thesis was about 2-D full waveform inversion. For demonstration purposes we limited ourselves to an acoustic implementation in which the density is assumed to be constant and known. Furthermore, we used homogeneous background models only because they significantly speed up the computation of the scheme. For real data applications these limitations are absolutely insufficient. If real data applications are aimed for, the implementation of inhomogeneous smooth background models is inevitable. This means that the Green's functions in the smooth background medium need to be computed numerically. This could be done by using the WKB approximation but a 2-D implementation is then required.

To demonstrate the improvements that can be achieved by nonlinear full waveform inversion on real data, we propose to focus on applications for the near-surface problem. Here, localisation by backpropagation is not necessary because only the first several hundred meters below the surface are of interest. A lateral limitation can immediately be achieved by selecting an appropriate subset of the full surface acquisition. By time windowing the data a limitation with respect to depth is obtained. In this way a high resolution nonlinear inversion result of the near-surface can be obtained. It should be realised that this application would be of great interest to the exploration industry as well. A successful application can help to avoid drilling hazards that often occur due to unknown features, e.g. gas pockets, in the near-surface. In addition, resolving the complexities of the shallow subsurface, would greatly improve imaging of the deeper data.

Next, quantitatively correct backpropagation should be incorporated in real data applications. Still, further research would be necessary on how accurately the incident wavefields should be estimated in the subsurface. Inverting the data on a 5 m subsurface grid does not solve for the transmission that happens on an extremely fine (centimetre) scale in the overburden. This could require the presence of well logs to calibrate the inversion process with respect to dispersion. It might also be that a more general transmission function can be found that is sufficient to explain this effect, but certainly more research is needed in this field.

The implementation of elastic full waveform inversion for reservoir characterisation was demonstrated on real data. However, here too a better understanding of several processes will be very helpful. A big disappointment in the case study was the extremely poor seismic-to-well match. If quantitative properties are supposed to be recovered one would hope to have a much better match between the measured data at the well location and synthetic data based on the logged properties in the well. One explanation for the poor match could be that reservoir-oriented inversion does not take the overburden between the surface and the reservoir into account. This means that multiples that are generated in this overburden can leak into the time window that is used for target-oriented inversion. Since these events were not generated within the target area they cannot be explained by the forward modelling engine. It should be realised that in the case study we did not apply the global scheme, but local inversions were combined with local total field updates over the reservoir sequence. One approach to tackle this problem are recent developments of migration techniques that can estimate and utilise multiple scattering. In this way, multiples generated in the overburden can be predicted and added to the incident field for full waveform inversion of the reservoir sequence below.

Appendix A

Derivation of the gradients

In this appendix we derive the gradients of data functional as well as of the regularisation functional with respect to χ_{n-1} . We start by the general definition of the Fréchet derivative as given by Kantorovich and Akilov (1982):

$$\partial F(a; b) = \lim_{\epsilon \rightarrow 0} \frac{F(a + \epsilon b) - F(a)}{\epsilon}, \quad \epsilon \in R. \quad (\text{A.1})$$

The Fréchet derivative measures the variation of the functional ∂F per variation ϵ along the direction b . First we derive the gradient of the data functional as it was defined in equation (2.8) with respect to the contrast function χ_{n-1} . Substitution of the data functional into equation (A.1) results in

$$\begin{aligned} & \partial F_{n-1}^{data}(\chi_{n-1}; v_n) \\ = & \lim_{\epsilon \rightarrow 0} \frac{\eta \iiint |p_{sct} - [\mathcal{K}(\chi_{n-1} + v_n)]|^2 - |p_{sct} - [\mathcal{K}\chi_{n-1}]|^2 d\vec{x}_s d\vec{x}_r d\omega}{\epsilon} \\ = & \lim_{\epsilon \rightarrow 0} \frac{\eta \iiint [\epsilon \mathcal{K} v_n]^2 - 2(p_{sct} - \mathcal{K}\chi_{n-1})\mathcal{K} \cdot v_n d\vec{x}_s d\vec{x}_r d\omega}{\epsilon} \\ = & -\eta \iiint 2(r_{n-1} \mathcal{K}) \cdot v_n d\vec{x}_s d\vec{x}_r d\omega \\ = & -2 \langle \eta \mathcal{K} r_{n-1}, v_n \rangle \end{aligned} \quad (\text{A.2})$$

Since the gradient provides the direction in which the variation of the functional is maximum, we may identify the gradient with respect to the Fréchet derivative, apart from the constant -2 as:

$$g_n^{data} = \eta \mathcal{K}^* r_{n-1}. \quad (\text{A.3})$$

In a similar way we derive now the gradient of the regularisation functional. We start again by substituting the regularisation functional as defined in equa-

tion (2.21) into equation (A.1):

$$\begin{aligned}
\partial F_{n-1}^{reg}(\chi_{n-1}; v_n) &= \lim_{\epsilon \rightarrow 0} \frac{F_{n-1}^{reg}(\chi_{n-1} + \epsilon v_n) - F_{n-1}^{reg}(\chi_{n-1})}{\epsilon} \\
&= \lim_{\epsilon \rightarrow 0} \frac{\int_{\vec{x} \in \mathcal{D}} [2\epsilon(\nabla\chi) \cdot (b_{n-1}^2 \nabla v) + \epsilon^2 |b_{n-1}^2 \nabla v|^2] d\vec{x}}{\epsilon} \\
&= 2 \int_{\vec{x} \in \mathcal{D}} (\nabla v_n) \cdot (b_{n-1}^2 \nabla \chi_{n-1}) d\vec{x}. \tag{A.4}
\end{aligned}$$

Using the relation $\nabla \cdot (v_n b_{n-1}^2 \nabla \chi_{n-1}) = (\nabla v_n) \cdot (b_{n-1}^2 \nabla \chi_{n-1}) + v_n \nabla \cdot (b_{n-1}^2 \nabla \chi_{n-1})$, we obtain

$$\begin{aligned}
\partial F_{n-1}^{reg}(\chi_{n-1}; v_n) &= 2 \int_{\vec{x} \in \mathcal{D}} [\nabla \cdot (v_n b_{n-1}^2 \nabla \chi_{n-1}) \\
&\quad - v_n \nabla \cdot (b_{n-1}^2 \nabla \chi_{n-1})] d\vec{x}. \tag{A.5}
\end{aligned}$$

Application of Gauss's theorem to equation (A.5) results in

$$\begin{aligned}
\partial F_{n-1}^{reg}(\chi_{n-1}; v_n) &= 2 \int_{\vec{x}' \in \partial \mathcal{D}} \vec{n} \cdot [v_n b_{n-1}^2 \nabla \chi_{n-1}] d\vec{x}' \\
&\quad - 2 \int_{\vec{x} \in \mathcal{D}} v_n \nabla \cdot (b_{n-1}^2 \nabla \chi_{n-1}) d\vec{x}, \tag{A.6}
\end{aligned}$$

where $\partial \mathcal{D}$ is the boundary of \mathcal{D} and \vec{n} is the normal to $\partial \mathcal{D}$. When we choose $\partial \mathcal{D}$ to lie outside the scattering domain where the contrast vanishes, the first integral in equation (A.6) vanishes. This can be justified because by definition the contrast outside the scattering domain is assumed to be zero. Hence, also $\nabla \chi$ becomes zero along $\partial \mathcal{D}$. We therefore can write the Fréchet derivative of the regularisation functional as an inner product:

$$\partial F_{n-1}^{reg}(\chi_{n-1}; v_n) = -2 \langle \nabla \cdot (b_{n-1}^2 \nabla \chi_{n-1}), v_n \rangle. \tag{A.7}$$

With the same reasoning as previously, we may identify the gradient with respect to the Fréchet derivative, apart from the constant -2 as:

$$g_n^{reg} = \frac{\partial F_{n-1}^{reg}}{\partial \chi_{n-1}} = \nabla \cdot (b_{n-1}^2 \nabla \chi_{n-1}). \tag{A.8}$$

Bibliography

- Abramowitz, M., and Stegun, I. A., 1970, Handbook of mathematical function; with formulas, graphs, and mathematical tables: Dover Publications, Inc.
- Abubakar, A., van den Berg, P. M., Habashy, T. M., and Braunisch, H., 2004, A multiplicative regularization approach for deblurring problems: IEEE Trans. Image Processing, Vol. 13, pages 1524–1532.
- Abubakar, A., Hu, W., Habashy, T., and van den Berg, P., 2009, Application of the finite-difference contrast-source inversion algorithm to seismic full-waveform data: Geophysics, **74**, no. 6, WCC163.
- Abubakar, A., Habashy, T. M., and Pan, G., 2011, Source-receiver compression approach for 3d full-waveform inversion with an iterative forward solver: 81st Ann. Internat. Mtg., Soc. Expl. Geophys., Expanded abstracts, 2411–2417.
- Bamberger, A., Chavent, G., and Lailly, P., 1977, Une application de la théorie du contrôle à un problème inverse sismique: Ann. Geophys., **33**, 183–200.
- Bamberger, A., Chavent, G., Hemons, C., and Lailly, P., 1982, Inversion of normal incidence seismograms: Geophysics, **47**, 757–770.
- Bamberger, A., Glowinski, R., and Tran, U. H., 1997, A domain decomposition method for the acoustic wave equation with discontinuous coefficients and grid change: SIAM J. Math. Anal., **34**, no. 2, 603–639.
- Barr, F. J., 1997, Dual sensor obc technology: The Leading Edge, **16**, 45.
- Berkhout, A. J., and Verschuur, D. J., 2011, Nonlinear, full waveform modeling in terms of forward multiple scattering: 73th Ann. Internat. Mtg., B017.
- Berkhout, A. J., 1982, Seismic migration, imaging of acoustic energy by wave field extrapolation, a: theoretical aspects: Elsevier.

- Berkhout, A. J., 2012, Combining full wavefield migration and full waveform inversion, a glance into the future of seismic imaging: *Geophysics*, **77**, no. 2, S43–S50.
- Berryhill, J. R., 1984, Wave equation datuming before stack (short note): *Geophysics*, **49**, no. 11, 2064–2066.
- Bremmer, H., 1951, The WKB approximation as a first term of a geometric-optical series:, *in* *The Theory of Electromagnetic Waves: A Symposium*, 169–179.
- Brossier, R., Operto, S., and Virieux, J., 2009, Seismic imaging of complex onshore structures by 2d elastic frequency-domain full-waveform inversion: *Geophysics*, **74**, WCC63–WCC76.
- Bunks, C., Saleck, F. M., Zaleski, S., and Chavent, G., 1995, Multiscale seismic waveform inversion: *Geophysics*, **60**, no. 5, 1457.
- Cruse, E., Pica, A., Noble, M., McDonald, J., and Tarantola, A., 1990, Robust elastic non-linear waveform inversion: Application to real data: *Geophysics*, **55**, 527–538.
- de Hoop, A. T., 1995, *The handbook of radiation and scattering of waves*: Academic Press, Delft.
- Epanomeritakis, I., AkDelik, V., Ghattas, O., and Bielak, J., 2008, A newton-cg method for large-scale three-dimensional elastic full waveform seismic inversion: *Inverse Problems*, **24**, 1–26.
- Fichtner, A., and Trampert, J., 2011, Hessian kernels of seismic data functionals based upon adjoint techniques: *Geophys. J. Int.*, **185**, 775–798.
- Fokkema, J. T., and van den Berg, P. M., 2012, Stretched backgrounds for acoustic scattering models: *Journal of Computational Physics*, **231**, 1728–1742.
- Haffinger, P., and Verschuur, D. J., 2012, Estimation and application of near-surface full waveform redatuming operators: *Geophys. Prosp.*, **60**, no. 2, 270–280.
- Hoffe, B. H., Lines, L. R., and Cary, P. W., 2000, Applications of obc recording: *The Leading Edge*, **19**, no. 4, 382.
- Kantorovich, L. V., and Akilov, G. P., 1982, *Functional analysis*: Pergamon Press, New York, USA.
- Kennett, B. L. N., 1983, *Seismic wave propagation in stratified media.*: Cambridge University Press.
- Kirchhoff, G., 1883, Zur theorie der lichtstrahlen: *Ann. d. Physik*, **18**, no. 2, 663ff.

- Kleinman, R. E., and van den Berg, P. M., 1991, Iterative methods for solving integral equations: *Radio Science*, **26**, no. 1, 175–181.
- Krebs, J., Anderson, J., Hinkley, D., Neelamani, R., Lee, S., Baumstein, A., and Lacasse, M., 2009, Fast full-wavefield seismic inversion using encoded sources: *Geophysics*, **74**, no. 6, WCC177–WCC188.
- Lailly, P., 1983, The seismic inverse problem as a sequence of before stack migrations, *in* Bednar, J. B., Redner, R., Robinson, E., and Weglein, A., Eds., *Conference on inverse scattering: theory and application: Society of Industrial and Applied Mathematics*.
- Loewenthal, D., Lee, S. S., and Gardner, G. H. F., 1985, Deterministic estimation of a wavelet using impedance type technique: *Geophys. Prosp.*, **33**, no. 7, 956–969.
- Martin, G. S., Wiley, R., and Marfurt, K. J., 2006, Marmousi2: An elastic upgrade for marmousi: *The Leading Edge*, **25**, 156–166.
- Métivier, L., Brossier, R., Virieux, J., and Operto, S., 2012, Toward gauss-newton and exact newton optimization for full waveform inversion: 74th Ann. Internat. Mtg., Eur. Ass. of Geosc. and Eng., Expanded abstracts, P016.
- Mitchell, A. R., and Griffiths, D. F., 1980, *The finite difference method in partial differential equations.*: J. Wiley & Sons, London.
- Moghaddam, P., and Herrmann, F., 2010, Randomized full-waveform inversion: a dimensionality-reduction approach: 80th Conference and Exhibition, Soc. Expl. Geophys., Expanded abstracts.
- Mora, P., 1987, Nonlinear two-dimensional elastic inversion of multi-offset seismic data: *Geophysics*, **52**, no. 9, 1211–1228.
- Mulder, W. A., 2005, Rigorous redatuming: *Geophys. J. Int.*, **161**, 401–415.
- Plessix, R. E., Baeten, G., de Maag, J. W., Klaassen, M., Rujie, Z., and Zhifei, T., 2010, Application of acoustic full waveform inversion to a low frequency large offset land data set: 80th Conference and Exhibition, Soc. Expl. Geophys., Expanded abstracts.
- Polak, E., and Ribière, G., 1969, Nonlinear two-dimensional elastic inversion of multi-offset seismic data: Note sur la convergence de méthodes de directions conjuguées, **16**, 35–43.
- Pratt, R. G., Shin, C., and Hicks, G. J., 1998, Gauss-newton and full newton methods in frequency-space seismic waveform inversion: *Geophys. J. Int.*, **133**, 341–362.

- Prieux, V., Brossier, R., Kommendal, J. H., Barkved, O. I., Operto, S., and Virieux, J., 2010, Application of 2d acoustic frequency-domain full-waveform inversion to obc wide-aperture data from the valhall field: 80th Conference and Exhibition, Soc. Expl. Geophys., Expanded abstracts.
- Rayleigh, J. W. S., 1896, The theory of sound (volume 2):.
- Robinson, E. A., 1986, Migration of seismic data by the wkbj method: Proceedings of the IEEE, **74**, no. 3, 428–439.
- Routh, P., Krebs, J., Lazaratos, S., Baumstein, A., Lee, S., Cha, Y. H., Chikichev, I., Downey, N., Hinkley, D., and Anderson, J., 2011, Encoded simultaneous source full-wavefield inversion for spectrally shaped marine streamer data: 81st Conference and Exhibition, Soc. Expl. Geophys., Expanded abstracts.
- Shewchuk, J., 1994, An introduction to the conjugate gradient method without the agonizing pain: <http://www.cs.cmu.edu/~jrs/jrspapers.html>.
- Shtivelman, V., and Canning, A., 1988, Datum correction by wave equation extrapolation: Geophysics, **53**, no. 10, 1311–1322.
- Soni, A. K., Wouters, W., and Verschuur, D. J., 2012, Target oriented VSP imaging - a sparse-inversion approach: 74th Ann. Internat. Mtg., Eur. Ass. of Geosc. and Eng., Expanded abstracts, P095.
- Staal, X. R., Verschuur, D. J., Gisolf, A., Tonellot, T., and Burnstad, R., 2010, Improved target-oriented linear full waveform inversion: 72th Ann. Internat. Mtg., Eur. Ass. of Geosc. and Eng., Expanded abstracts, P388.
- Symes, W., 1981, Stable solution of the inverse reflection problem for a stratified elastic medium: SIAM J. Math. Anal., **12**, no. 3.
- Tarantola, A., 1984, Inversion of seismic reflection data in the acoustic approximation: Geophysics, **49**, no. 8, 1259–1266.
- Tarantola, A., 1987, Inverse problem theory, methods for data fitting and model parameter estimation: Elsevier Science Publ. Co., Inc.
- Valenciano, A., Biondi, B., and Guitton, A., 2006, Target-oriented wave-equation inversion: Geophysics, **71**, A35–A38.
- van den Berg, P. M., Abubakar, A., and Fokkema, J. T. ., 2003, Multiplicative regularization for contrast profile inversion: Radio Sci., Vol. 38, No. 2, pages 23–1 – 23–10.
- Vigh, D., Kapoor, J., Moldoveanu, N., and Hongyan, L., 2011, Breakthrough acquisition and technologies for subsalt imaging: Geophysics, **76**, no. 5, WB41–WB51.

- Virieux, J., and Operto, S., 2009, An overview of full-waveform inversion in exploration geophysics: *Geophysics*, **74**, no. 6, WCC1.
- Wapenaar, C. P. A., and Berkhout, A. J., 1989, Elastic wave field extrapolation: redatuming of single- and multi-component seismic data: Elsevier Science Publ. Co., Inc.
- White, R. E., 2003, Tutorial: Good practice in well ties: *First Break*, **21**, no. 10, 75–83.
- Yang, J., Abubakar, A., van den Berg, P. M., Habashy, T. M., and Reitich, F., 2008, A cg-fft approach to the solution of a stress-velocity formulation of three-dimensional elastic scattering problems: *Journal of Computational Physics*, **227**, 10018–10039.

Summary

Full waveform inversion is a tool to obtain high-resolution property models of the subsurface from seismic data. However, the technique is computationally expensive and so far no multi-dimensional implementation exists to achieve a resolution that can directly be used for seismic interpretation and characterisation. In this thesis we discuss a method to overcome some of the current limitations of seismic full waveform inversion. The scheme consists of alternating local linear inversions followed by global nonlinear total field estimations.

By backpropagating sources and receivers from the acquisition surface into the subsurface, the inversion process can be localised. Inversion of backpropagated data to obtain the properties in a local domain only, is a much smaller problem compared to direct inversion of surface data, aiming at the properties of the full subsurface at once. As a result we can use the full frequency spectrum, up to 55 Hz and beyond, to obtain property models with a resolution that is comparable to the one obtained by conventional structural imaging. A high-resolution global subsurface model in terms of properties is then obtained by combining many of these local results. Because all local inversions are fully independent, the scheme is extremely suitable for parallelisation.

In the first iteration we assume that the wavefields are propagating in a smooth background model only. Obtaining a high-resolution property model by localised inversion, showing structural features of the subsurface, tells us that wavefield propagation in a smooth background is not a good approximation. This is why in a next step we perform a total wavefield estimation based on the latest combined local inversion results. In this way the full nonlinear relationship between the measured seismic data and the subsurface properties is utilised. We show how nonlinear wavefield effects as multiple scattering, transmission and true travel-times in the inverted medium, can properly be estimated and how they can be used to obtain significantly improved subsurface models. Examples are given where nonlinear full waveform inversion allows us to recover structural information, particularly very steep dipping geology, that can never be imaged by using

a linear data model.

It is demonstrated that by using the nonlinear relationship between the measured data and the obtained subsurface properties, the resolution can be increased even further. Full waveform inversion as proposed in this thesis, allows reconstruction of subsurface models with a spatial resolution that goes beyond the equivalent temporal frequencies in the measured seismic data. Consequently, estimation of absolute quantitative properties from a bandlimited seismic signal becomes possible.

It is discussed that every scatterer in the subsurface introduces a coda that will have an effect on the imaging of the deeper subsurface. Especially the near-surface can introduce nonlinearity, particularly in terms of multiple scattering, which makes imaging the deeper area of interest, e.g. a reservoir, difficult. In this thesis we show how these internal multiples cannot only be handled, but also how they can contribute to the inversion process by introducing enhanced illumination.

The method is first derived in 2-D under the acoustic approximation and assuming a constant and known density. Since current implementation of the scheme requires homogeneous background models as input to the inversion, in 2-D the method is demonstrated on synthetic datasets only. Field data application using a 2-D data model needs extension to inhomogeneous background models and to true elastic inversion, which is in principle straightforward, but has not been implemented yet.

In this thesis we perform these extensions for the 1.5-D case, assuming a horizontally layered subsurface per CMP location. Nonlinear elastic full waveform inversion, aiming at the recovery of three elastic parameters over a reservoir sequence, is applied to synthetic data as well as to a field dataset from the Middle East. It is shown that nonlinear full waveform inversion is suitable for reservoir characterisation while increased resolution and better results in terms of quantitative properties are obtained.

Samenvatting

Full waveform inversie is een techniek om met hoge resolutie modellen van de eigenschappen van de ondergrond te verkrijgen van seismische data. Echter, de techniek is rekenkundig duur en tot op heden bestaat er geen multidimensionale implementatie om een resolutie te bereiken die direct gebruikt kan worden voor seismische interpretatie en karakterisatie. In dit proefschrift wordt een methode besproken om de huidige limitaties van seismische full waveform inversie te doorbreken. Het schema bestaat uit afwisselend lokale lineaire inversies en globale niet-lineaire herberekening van het totale veld in het object.

Door inverse golfveldpropagatie van bronnen en ontvangers vanaf het acquisitie oppervlak naar deelgebieden in de ondergrond, kan de inversie worden gelokaliseerd. Inversie van terug gepropageerde data om de eigenschappen in een lokaal domein te verkrijgen, is een veel kleiner probleem dan directe inversie van de oppervlakte data, die alle eigenschappen van de hele ondergrond tegelijk bepaalt. Derhalve kan het gehele frequentie spectrum worden gebruikt, tot 55 Hz en verder, om eigenschap modellen te verkrijgen met een resolutie die vergelijkbaar is met de resolutie verkregen met conventionele structurele beeldvorming. Een globaal model met een hoge resolutie, in termen van eigenschappen, van de ondergrond wordt verkregen door het combineren van al deze lokale resultaten. Omdat alle lokale inversies volledig onafhankelijk zijn, is het schema extreem geschikt voor parallelisatie.

In de eerste iteratie nemen we alleen aan dat de golfvelden propageren in een 'smooth' achtergrond model. Het verkrijgen van een model met een hoge resolutie van eigenschappen door gelokaliseerde inversie, dat de structurele eigenschappen van de ondergrond laat zien, verteld ons dat golfveld propagatie in een 'smooth' achtergrond geen goede benadering is. Dit is de reden waarom in de volgende iteratie een schatting van het totaal golfveld wordt uitgevoerd, die is gebaseerd op de laatste gecombineerde lokale inversie resultaten. Op deze manier wordt de volledig niet-lineaire relatie tussen de gemeten seismische data en de eigenschappen van de ondergrond gebruikt. We laten zien hoe niet-lineaire golfveld effecten,

zoals meervoudige verstrooiing, transmissie en echte looptijden in het medium, kunnen worden geschat en hoe ze kunnen worden gebruikt om significant verbeterde modellen van de ondergrond te verkrijgen. Voorbeelden worden gegeven waar niet-lineaire full waveform inversie structurele informatie oplevert, in het bijzonder steile hellingen in de geologie, die nooit kunnen worden afgebeeld door gebruik te maken van een lineair data model.

Er wordt gedemonstreerd dat door gebruik te maken van de niet-lineaire relatie tussen de gemeten data en de verkregen eigenschappen van de ondergrond, de resolutie zelfs verder kan worden verbeterd. Full waveform inversie, zoals gepresenteerd in dit proefschrift, is in staat modellen van de ondergrond te reconstrueren met een spatiale resolutie die hoger is dan op grond van de temporele frequenties in de gemeten data verwacht had kunnen worden.

Er wordt vastgesteld dat elke verstrooier in de ondergrond een coda introduceert dat een effect heeft op de beeldvorming van de diepere ondergrond. In het bijzonder de ondergrond dichtbij het oppervlak kan niet-lineariteit introduceren, in het bijzonder door meervoudige verstrooiing, wat de beeldvorming van het diepere gebied van interesse bemoeilijkt. In dit proefschrift laten we zien hoe de meervoudige verstrooiing niet alleen wordt gebruikt, en hoe deze kan bijdragen tot het inversie proces door het introduceren van verbeterde illuminatie.

De methode is eerst afgeleid in 2-D onder de akoestische benadering met een constante en bekende massadichtheid. Omdat de huidige implementatie van het schema een homogeen achtergrond model nodig heeft, wordt de methode in 2-D alleen gedemonstreerd op synthetische data. De toepassing op echte data met een 2-D data model vereist uitbreiding tot inhomogene achtergrond modellen en echte elastische inversie, wat in principe meer van hetzelfde is, maar nog niet is geïmplementeerd.

In dit proefschrift breiden we de methode uit naar 1.5-D, er vanuit gaande dat de ondergrond vlak gelaagd is per CMP locatie. Niet-lineaire elastische full waveform inversie, gericht op het vinden van drie elastische parameters voor alle roosterpunten in een deelgebied, is toegepast op synthetische data, zowel als op een echte dataset uit het Midden-Oosten. Niet-lineaire full waveform inversie is geschikt voor reservoir karakterisatie met een verhoogde resolutie en betere resultaten in termen van kwantitatieve eigenschappen, dan verkregen zou worden met behulp van conventionele lineaire inversie.

Acknowledgements

During the last four years as PhD student I travelled a journey that will remain unforgettable for the rest of my life. Still, it is not the journey itself that makes it such an important period but all the different people I met and who greatly contributed to my scientific work but also to my personal development.

Foremost, I would like to express my sincere gratitude to my promoters Prof. Dries Gisolf and Prof. Peter van den Berg. It was a pleasure to work with you and I highly appreciate the many hours I spend in your offices discussing about inverse problems, regularisation, ill-posedness, nonlinear relationships but also about subjects that go far beyond the scientific part of my PhD.

During the first six month of my PhD I was among the lucky ones to be supervised by Dr. Eric Verschuur. It was a joy working with such an outstanding researcher and I truly enjoyed this experience.

I always considered being a DELPHI student as a privilege and I would like to express my appreciation to Prof. Berkhout for making this possible by directing this outstanding consortium over the last 30 years.

My defence committee is acknowledged for their time reading this thesis and for improving it by their valuable suggestions.

Saudi Aramco is gratefully acknowledged for their hospitality during my internship at the EXPEC Advanced Research Center in Dhahran, Saudi Arabia. I particularly recognise Dr. Panos Kelamis for hosting me in his research group and Dr. Thierry Tonellot for his scientific supervision during my visit.

I would like to show my appreciation to all sponsors of the DELPHI consortium for their financial support, but also for the invaluable comments and suggestions they made during the sponsor meetings.

Sincere thanks are given to Margaret for all the assistance she provided in any matter. Edo and Henry are highly appreciated for their technical support whenever and wherever it was needed. I would like to thank Jan for all his help whenever my Mac did not do what I wanted it to do and I truly appreciate all the support from Hanneke during the Delphi meetings I attended.

I am very grateful to my very good friends Araz, Hannes and Panos. I truly en-

joyed the time with you guys and all the unforgettable moments we shared during the last four years.

My thanks are offered to Alok, Dana, Justyna, Gabrio and Jiawen for the good time we had while sharing an office.

These thanks are naturally extended to the many DELPHI members I had the pleasure to meet, particularly Ahmad, Alex, Amarjeet, Daria, Gabriel, Gerrit, Gert-Jan, Jan-Willem, Marwan, Mikhail, Mohammed, Rolf, Tomo, Xander and Yimin. Those I unintentionally forgot to mention will hopefully forgive me.

Von tiefstem Herzen bedanke ich mich bei meiner Familie - meinen Eltern Brigitte und Paul sowie meiner Schwester Sabine - für ihre schier unendliche Unterstützung, Geduld und Liebe. Ihr seid es die mir Mut und Ehrgeiz aber auch Respekt und Dankbarkeit mit auf den Weg gegeben haben welche unersetzlich waren um meinen eigenen Weg zu finden. Ich werde euch ewig dankbar sein.

Zuletzt möchte ich meine unendliche Dankbarkeit gegenüber Linda zum Ausdruck bringen, dass sie während der letzten Jahre für mich zum wichtigsten Menschen geworden ist. Es gibt niemanden auf der Welt dem ich mich nur ansatzweise so verbunden fühle wie dir. Es ist unser gegenseitiges Vertrauen und unsere Zuneigung die mir jeden Tag die Kraft gibt neue Dinge anzugehen, stets gewiss eine Partnerin zu haben die zu mir steht. Danke das es dich gibt.

*Delft,
December 2012*

Peter Haffinger

Curriculum Vitae

Personal Information

- Full name: Peter Rudolf Haffinger
- Date of Birth: 01/05/1985
- Place of Birth: Trier, Germany
- Email: phaffinger@gmx.com

Education

- PhD student in Applied Physics, TU Delft, The Netherlands, 2008-2012
- Master of Science in Applied Geophysics, IDEA League (TU Delft, ETH Zurich, RWTH Aachen), 2006-2008
- Vordiplom (pre-diploma) in Physics, RWTH Aachen, Germany, 2004-2006
- Abitur (german secondary school), Friedrich-Spee-Gymnasium Trier, Germany, 1995-2004

Internships and Work Experience

- Delft University of Technology, The Netherlands, 2008-2012
- Internship at EXPEC Advanced Research Center, Saudi Aramco, Dhahran, Saudi-Arabia, 10/2011-12/2011
- Internship with Western Geco (Multicomponent DP group), Gatwick, Great Britain, 01/2008-07/2008
- Geophysical field work for archeological prospecting, ETH Zurich, Switzerland, 04/2007

Publications

- Haffinger P., Gisolf A. and van den Berg P.M., 2012, Towards high resolution quantitative subsurface models by full waveform inversion: *Geophysical Journal International*, Revision submitted.
- Vrolijk J.W., Haffinger P. and Verschuur D.J., 2012, Multi-datum based estimation of near-surface full-waveform redatuming operators: *Journal of Applied Geophysics*, Vol. 82, 30-45.
- Haffinger P. and Verschuur D.J., 2012, Estimation and application of near-surface full waveform redatuming operators: *Geophysical Prospecting*, Vol. 60, No. 2, 270-280.
- Haffinger P., Gisolf A. and van den Berg P.M., 2012, Nonlinear full waveform inversion for high resolution quantitative property estimation: *Extended Abstracts of the 82nd SEG Annual International Meeting in Las Vegas*.
- Haffinger P., Gisolf A. and van den Berg P.M., 2012, Broadband Full Waveform Inversion by Shot/Receiver Re-focusing: *Extended Abstracts of the 74th EAGE Conference & Exhibition in Copenhagen*.
- Haffinger P., Gisolf A. and van den Berg P.M., 2011, Towards broadband non-linear full-waveform inversion with the help of shot/receiver refocussing: *Extended Abstracts of the 81st SEG Annual International Meeting in San Antonio*.
- Vrolijk J.W., Haffinger P. and Verschuur D.J., 2011, Multi-datum based estimation of near-surface full-waveform redatuming operators: *Extended Abstracts of the 81st SEG Annual International Meeting in San Antonio*.
- Haffinger P. and Gisolf A., 2011, Full Bandwidth 2-D Full-waveform Inversion by Domain Decomposition: Presented at the SEG Summer Research Workshop in Quebec City.
- Haffinger P., van den Berg P.M. and Gisolf A., 2011, Localised Linear 2-D Inversion by Data Pre-conditioning in the Time-domain: *Extended Abstracts of the 73rd EAGE Conference & Exhibition in Vienna*.
- Haffinger P. and Verschuur D.J., 2010, Estimation and application of full waveform redatuming operators: *Extended Abstracts of the 72nd EAGE Conference & Exhibition in Barcelona*.

