Imperial College
London

Department of Earth Science
and Engineering

High-resolution imaging beneath the Santorini volcano

by

Kajetan Marcin Chrapkiewicz

Supervised by Prof. Joanna Morgan and Dr. Michele Paulatto

Submitted for the degree of Doctor of Philosophy
and the Diploma of Imperial College London

August 2021

Preface

Abstract

Volcanoes are surface expressions of much deeper magmatic systems, inaccessible to direct observation. Constraining the geometry and physical properties of these systems, in particular detecting high melt fraction (magma) reservoirs, is key to managing a volcanic hazard and understanding fundamental processes that lead to the formation of continents. Unfortunately, unambiguous evidence of magma reservoirs has not yet been provided due to the limited resolving power of the geophysical methods used so far.

Here, a high-resolution imaging technique called full-waveform inversion was applied to study the magmatic system beneath the Santorini volcanic field, one of the most volcanically and seismically active regions of Europe. Quality-controlled inversion of 3d wide-angle, multi-azimuth ocean-bottom seismic data revealed a previously undetected high melt fraction reservoir 3 km beneath the Kolumbo volcano, a centre of microseismic and hydrothermal activity of the field.

To enable the above method to handle land data, two major algorithmic improvements were added to the high-performance inversion code. First, to simulate instrument response of land seismometers, a pressure-velocity conversion has been implemented in a way that ensures reciprocity of the discretised 2nd-order acoustic wave equation. Second, the immersed-boundary method, originally developed for computational fluid dynamics, was implemented to simulate the wave-scattering off the irregular topography of the Santorini caldera. These advancements can be readily used to provide a higher-resolution image of the melt reservoir beneath the Santorini caldera already detected by means of travel-time tomography.

Key words: magma plumbing systems · continental back-arc volcanism · Santorini volcanic field · high-resolution imaging · seismic full-waveform modelling and inversion · finite-difference method

Statement of originality

I herewith certify that all material in this thesis that is not my own work has been duly acknowledged.

Copyright declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

I would like to express deep gratitude to my supervisors, Dr. Michele Paulatto and Prof. Joanna Morgan. They proved to be incredibly friendly, patient and supportive people. They gave me freedom to pursue my research ideas, yet they were always ready to lend me a hand. Without their assistance, this thesis would never have been submitted.

I also thank all the members of the PROTEUS team, especially its leader Prof. Emilie Hooft, Prof. Doug Toomey, Ben Heath (now a PhD), Brennah McVey (all from UO) and the Greek members, Prof. Evi Nomikou and Prof. Costas Papazachos, for sharing their knowledge during many Skype meetings. I had a pleasure to meet all these amazing people in person on the 2018 field trip to Santorini. I benefited a lot from the notes of Emilie and the way she has managed the team in general.

Furthermore, I want to thank Prof. Mike Warner of Imperial for being an inspirational team leader who creates a motivating, and at the same time uplifting, atmosphere of a scientific enterprise. Mike knows how to quickly find working solutions to seemingly insurmountable problems. I really enjoyed the company and learnt a lot from other Fullwave group members too, in no particular order: Nuno, Lluís, Tanya, George, Oscar, Mel, Rich, Yao, Matt, Yamina, Nik, Tenice, Jack, Tim, Gang. Special words of appreciation must go to Adrian Umpleby for his wisdom, patience and readiness to answer code-related questions of mine and others. I wish the PhD students who came after me, George, Debby and Lisa all the best with their research too. Prof. Saskia Goes and members of the PMAC research group had a positive influence on my stay at Imperial.

I also want to thank my supervisors for allowing me to take part in the field-data acqusition for the NZ3D experiment in New Zealand. One of the experiment's PIs, Dr. Rebecca Bell is yet another inspiring researcher full of positive energy from whom I learnt a lot.

Finally, I would like to thank my family, especially my parents and brothers, for all the love and support throughout my entire life, especially in the last period of writing up this dissertation. My late grandma helped me during the PhD more than she could have imagined.

Academic record

Talks

- *Cities on Volcanoes*, accepted, conference postponed;
- *European Geophysical Union General Assembly*, online, 2020;
- *American Geophysical Union Fall Meeting* Washington D.C., 2019;
- *Seminar of Institute of Geophysics* (invited), University of Warsaw, 2019;
- *FULLWAVE Consortium Steering Group Meeting*, Imperial College London, 2019;
- *FULLWAVE Consortium Steering Group Meeting*, Imperial College London, 2018;
- *British Geophysical Association Post-Graduate Research in Progress Conference*, Cardiff University, 2018;
- *ESE Department Student Conference*, Imperial College London, 2018.

Selected courses

- *Navigating the IODP*[1] *Proposal System for PhD and early-career scientists* 24-26 September 2018, National Oceanography Centre, Southampton, UK;
- *Shared-Memory Programming with OpenMP* 12-13 June 2017, Imperial College London;
- *ARCHER Driving Test* (passed online);
- *Guide to HPC Service at Imperial part 2 – OpenMP, MPI and hybrid jobs*, 25 May 2017, Imperial College London;
- *Guide to HPC Service at Imperial part 1 – from serial jobs to data parallelism*, 25 May 2017, Imperial College London

Field work

Around two-month field work consisting in deployment, maintanance and collecting land geophones in New Zealand, as a part of the NZ3D experiment.

Teaching

Serving as a graduate teaching assistant for:

- *Mathematical methods I* (three semesters);
- *Applied geophysics* (three semesters);
- *Advanced applied geophysics* (three semesters);
- *Volcanism and Internal processes* (two semesters);

[1] International Ocean Drilling Program

- *Geophysical inversion* (one semester);
- *Continental tectonics* (one semester);
- *Applied Computational Science & Engineering* (two workshops).

Supervision

Co-supervising *Computing the elastic properties of porous materials using FFT homogenization* MSc project of Chloe Spurling.

Thesis structure

The dissertation consists of three main parts:

- **Part** I provides the geological motivation of studying volcanoes, introduces the study area, the experimental setup and lastly the imaging method.
- **Part** II describes, after necessary theoretical background, the methodology developed to simulate land and/or marine multi-component data. Part of this work was presented at the *American Geophysical Union Fall Meeting 2019*.
- **Part** III describes the inversion of ocean-bottom field data and provides interpretation of its outcome. Some of these results were presented at the *European Geophysical Union General Assembly 2020*.

The first part provides a background of the project. The remaining two parts present the original results obtained by the author. Ch. 7 describes the preparation of the input used in ch. 8 as well as in the field-data examples of part II. Each of these chapters include its own suite of appendices (instead of one large at the end of the thesis) for easier referencing.

Note on notation

The indicial notation and Einstein's summation convention are used whenever they are thought to help the reader by keeping the equations concise and not too overwhelming. According to the summation convention, index appearing exactly twice in a monomial expression is called *dummy* and indicates that the expression in fact represents a sum of monomials over a full range of that index, e.g.:

$$\tau_{ij}\delta_{ij} \equiv \sum_{i=1}^{3}\sum_{j=1}^{3}\tau_{ij}\delta_{ij},$$

where δ_{ij} is the Kronecker delta. Indices that do not repeat in such expressions are called *free* and must appear exactly once on each side of an (in)equality. This way, a system of three equations can be written concisely as:

$$\rho\dot{v}_i = \tau_{ij,j} + f_i,$$

with i and j being a free and dummy index, respectively, τ_{ij} is a second-order tensor and a comma, as in $\tau_{ij,j}$, is a shorthand notation for differentiation:

$$\tau_{ij,j} \equiv \partial_j\tau_{ij} \equiv \frac{\partial\tau_{ij}}{\partial x_j}.$$

The summation convention is assumed whenever indicial notation is used, unless stated otherwise.

Contents

Preface **2**

Contents **10**

List of Figures **12**

List of Tables **22**

I Introduction **25**

1 Volcanology background **26**
1.1 Basic concepts . . . 26
1.2 Seismic studies . . . 31
1.3 Open questions . . . 32

2 Study area **34**
2.1 Tectonic setting . . . 34
2.2 Volcanic activity . . . 37
2.3 Magmatic system . . . 40
2.4 Open questions . . . 40

3 PROTEUS experiment **41**
3.1 Acquisition . . . 41
3.2 Observed data . . . 47
3.3 Travel-time inversion . . . 48
3.4 Imaging targets . . . 49

4 Full-waveform inversion **51**
4.1 Introduction . . . 51
4.2 Inverse theory . . . 53
4.3 Scattering theory . . . 61
4.4 FWI in practice . . . 63

II Full-waveform modelling **68**

5 Modelling multi-component data **69**

5.1 Introduction . . . 70
5.2 Theory . . . 72
5.3 Method . . . 81
5.4 Numerical tests . . . 86
5.5 Field-data example . . . 92
5.6 Conclusion . . . 98

Appendices **100**
5.A Locating stations . . . 100
5.B Analytical solutions . . . 100
5.C Caldera phase-plots . . . 101
5.D Travel-time picking . . . 101

6 Modelling irregular boundaries **106**
6.1 Introduction . . . 106
6.2 Basic IB method . . . 114
6.3 Point-projection in 3d . . . 126
6.4 Source injection . . . 130
6.5 Numerical tests . . . 134
6.6 Field-data example . . . 137
6.7 Conclusion . . . 143

Appendices **147**
6.A Vacuum method . . . 147
6.B `IMMERSE` software . . . 147
6.C Convergence tests . . . 148

III Inversion of field data **151**

7 Input preparation **152**
7.1 Metadata . . . 152
7.2 Observed data . . . 153
7.3 Source wavelet . . . 157
7.4 Starting model . . . 165

Appendices **169**
7.A `Jupyter` notebooks . . . 169

8 Kolumbo volcano **170**
8.1 Introduction . . . 171
8.2 Data . . . 174
8.3 Method . . . 180
8.4 Results . . . 184
8.5 Discussion . . . 190
8.6 Conclusion . . . 198

Appendices **201**

8.A Final inversion results . . . 201
8.B Anomaly-recovery tests . . . 201
8.C Melt sensitivity analysis . . . 201

Concluding remarks **209**

Bibliography **212**

List of Figures

1.1 Three major domains of a magmatic system. The transition between rock and mush (melt-connectivity transition, MCT) is characterised by a sudden drop in viscosity and strength. Upon the mush-magma transition (MMT), melt starts to control the rheology of the system. Modified after Sparks et al. (2019). . . . 27
1.2 Back-arc volcanism in a continental-ocean collision zone. Modified after Stern and Dickinson (2010). . . . 28
1.3 Schematic depiction of a magma plumbing system. From Magee et al. (2018). . . . 29
1.4 Time- and length scales of crustal magmatic activity. From Edmonds et al. (2019). . . . 31
1.5 Map of all the volcanoes studied by means of the seismic tomography. The interactive visualisation is available at `https://kmch.github.io/`. The database has been compiled by Lees (2007), and largely extended by M. Paulatto (2019, personal communication). . . . 31

2.1 Simplified geological map of Santorini. After Druitt et al. (1999). . . . 35
2.2 Regional tectonics showing Santorini at the centre of the Hellenic Volcanic Arc. The dashed box shows the approximate location of the study area (see ch. 3). The dotted line surrounds the Cycladic Block (Attico-Cycladic Massif). The red stars denote volcanic centres. From Hooft et al. (2017). . . . 36
2.3 The Hellenic subduction zone as imaged by Bijwaard et al. (1998). From van Hinsbergen et al. (2005). . . . 37
2.4 Major faults around Santorini. Note that the E side of the caldera is much more fractured than the W side. From Karstens et al. (2020). . . . 38
2.5 Comparison of volume of magma The eruption sizes were drawn by solving for the radius of a sphere using the eruption volumes shown. Modified after https://www.usgs.gov /media /images/ (public domain). . . . 39

3.1 The US Research Vessel Marcus G. Langseth of 71.5 m length (overall), 17 m width (beam, moulded) and 5.9 m of maximum draft (LD-EO, 2021). The moveable frame with a winch is used to deploy and collect OBS instruments. Top and bottom-left picture – modified from LD-EO (2021), bottom-right – from Montes (2013). . . . 42

3.2 Seismic source (airgun array) used in the experiment. Top-right and bottom-left – 1d and 2d view of the airgun array, respectively from Tolstoy et al. (2009). Remainder – size of selected elements of the array compared to people. From Hooft (2015a,b). 43
3.3 Seismic source wavelet of the R/V Langseth array simulated numerically by Lamont-Doherty Earth Observatory shown in time and frequency domain (left and right, respectively). Modified after Hooft (2015a). 44
3.4 Acquisition geomertry of the PROTEUS experiment. Shot lines P3 and P2 are the ones between/at the end of some of the primary (P1) shot lines. Numbering of transit lines (T) starts from 50. From Hooft (2015a). 45
3.5 Acquisition geomertry in local coordinates overlain on topography/bathymetry map. Red, blue and green triangles denote the WHOI, SIO and land stations respectively. The local coordinate system is based on the Universal Transverse Mercator rotated about 25.5° clockwise from north. The location of centre of the caldera is ca. at $36.4°\mathrm{N}, 25.4°\mathrm{E}$. Modified after M. Paulatto (unpublished). 45
3.6 Two types of ocean-bottom seismometers deployed in the PROTEUS experiment, from (top) Woods Hole Oceanographic Institution and (bottom) Scripps Institute of Oceanography instrument-pool. Modified after WHOI (2021) and SIO (2021), respectively. 46
3.7 OBS coverage highlighted by straight lines connecting sources and receivers for offsets shorter than 20 km. Rectangles denote regions of the densest coverage. 50

4.1 Subsurface, travel-time tomography, migration and FWI. Modified after ExxonMobil (2021). 52
4.2 Idea of a forward and inverse seismic problem. The forward step (green arrow) is realised by nature when the waves propagate through the Earth. Scientists simulate it numerically to solve the inverse problem (gray arrow). 53
4.3 Monochromatic wavepath. From Virieux and Operto (2009). 62
4.4 Scattering experiment. From Operto et al. (2015). 63
4.5 Time-lapse snapshots of the forward and adjoint wavefields, as well as their product, both instantanuous and time-integrated (columns from left to right, respectively). The adjoint wavefield was scaled with $\frac{\partial \mathbf{A}}{\partial \mathbf{m}}$. Modified after Tromp et al. (2005). 66

5.1 Amplitude spectrum of a direct wave (top) its time-derivative and (middle) direct wave (time-differentiated) with a ghost (bottom). In the last case, harmonic components of the two wave packets interfere leading to the notches in the spectrum. 71
5.2 Final models of Marmousi 2d recovered by FWI using different station coverage and sensor types. Benefits of inverting more than one component are subtle but noticable, especially for sparse acquisition. From Zheglova and Malcolm (2020). 71
5.3 Graphical depiction of the components of a moment tensor. Diagonal elements in the ellipse describe an explosion. Modified after (Shearer, 2009). 76
5.4 Schematic of reciprocal experiments involving monopole and dipole radiation patterns. Wavefield from a monopole source at location A recorded at location B by a dipole receiver (left) is equal to the wavefield from a dipole source at location B recorded at location A by a monopole receiver (right). The radiation pattern is denoted by the empty ('radiating out') and filled ('collecting in') shape (single/double circle for monopole/dipole) for the source and receiver, respectively. Bottom-row figures from Arntsen and Carcione (2000). 77

5.5 Schematic of the Dirac delta and its derivative (left, top and bottom) and their discrete versions – sinc function and its derivative (right, from Hicks, 2002). 82
5.6 Kaiser windows of the radius r equal to 3 grid nodes and the shape parameter α ranging from 0 (blue) to 100 (red). 83
5.7 Wavenumber spectrum of the band-limited discrete Dirac delta defined as a Kaiser-windowed sinc function (a). Case (b) same as (a) but with a larger radius; (c) same as (a) but with smaller offset from grid η (d) same as (b) but for smaller number of samples. The colormap of the lines encodes a window shape parameter α. The units of the horizontal axis (wavenumber) are scaled so that the value 1 corresponds to the Nyquist wavenumber. 83
5.8 Sinc function tapered with the Kaiser window defined with a 0th-order modified Bessel function of the first kind (left) in 1d (middle) and 2d (right). 83
5.9 Radiation from a monopole (left) and a dipole source (right) shown as 2d wavefield propagated in homogeneous model with a free surface at the top. Rows show different time-snapshots. Axes are scaled in metres. 85
5.10 Radiation from a monopole (top) and a dipole source (bottom) shown as x/y/z slices (from left to right) of 3d wavefield propagated from the dual-sensor OBS 166 in a realistic model of Christiana basin near Santorini. Axes scaled in grid nodes ($h = 50$ m). 85
5.11 Analysis of pressure-velocity reciprocity in 1d. Cases on the right and left are reciprocal. Each fig. shows source wavelet and recorded trace above/below the wavefield, respectively. Positive amplitude of traces/wavefield in blue. Dashed/solid black line – source/receiver location. Radiation-pattern symbols are explained in fig. 5.4. Constant v_P of 1 km/s with free surface at 0 and 2 km depth was used. Solver-code is supplied in list. 5.1. 87
5.12 Sinusoidal amplitude versus angle dependence (inset) of the data generated with a monopole source and recorded with the circular array of v_z receivers shown on the left (black star and coloured points, respectively). Each observed trace on the right corresponds to the receiver of the same colour shown on the left. The source function was a Ricker wavelet and the observed data consist of its derivative (fig. 5.22) convolved with a 2d Green's function (fig. 5.23). The lightblue trace recording the upgoing wave has a negative polarity and the orange one recording the downgoing wave – positive. The amplitudes of recorded the traces match the analytical solutions. 89
5.13 Reciprocity test in 2d: geometry (topmost), pressure source and pressure receiver (top), velocity source and pressure receiver (middle), pressure source and velocity receiver (bottom). X axis scale in time samples ($dt = 1$ ms). 91
5.14 Raw gradient calculated for a monopole (left) and dipole (right) source and a monopole receiver in a homogeneous of 2 km/s and absorbing boundaries. Starting model was 10% faster than the true one and the data residuals were used as adjoint sources. Source function – a Ricker wavelet of 4 Hz peak frequency. Domain size was 10×10 km sampled at 50 m; simulated time was 2 s sampled with 2000 time-steps. 92
5.15 Hydrophone data recorded by OBS 101 (pink dot in the top-left figure) for the shot line highlighted in red (top-left figure, other colours denote the shot-line numbers). The low-frequency noise dominates the data as shown in both time and frequency domain (top-right and bottom figure, respectively). 93
5.16 Geophone data recorded by the same OBS and shot line as in fig. 5.15, but with a much higher S/N ratio. 94

5.17 Starting v_P model (m/s) used in the field-data example. Top left – x slice, top right – z slice, bottom – y slice. Dashed white lines denote location of slices. Axes scaled in grid nodes ($dx = 50$ m), with variable aspect ratio. 95
5.18 Source wavelet used in the field-data example, here before and after the time-integration required to simulate the particle-velocity data. 95
5.19 Phase analysis of pressure and vertical component of particle-velocity data for the OBS station no. 157. The white star and coloured points denote horizontal locations of the reciprocal source and receivers, respectively. The colourmap the phase of data windowed in frequency at 3 Hz and time (1 s around the first arrivals), wrapped into $[-\pi, \pi)$ range. Left, middle and right panels show synthetic ϕ_{syn}, observed ϕ_{obs} and the residual phase, respectively. 96
5.20 Anomaly-recovery test using pressure vs. velocity data. The models recovered after 10 iterations at 3 Hz cut-off were calculated as the difference between the final and starting model (fig. 5.17, sliced in the same way). 97
5.21 Final P-wave velocity model (m/s) at Kolumbo (slices along the dashed black lines at z, y and x, respectively from top to bottom) recovered after 40 iterations using pressure (left) and velocity (right) data. Top-left – acquisition geometry overlain on bathymetry. 98
5.22 Gaussian function (negative) and its derivatives used to benchmark results of numerical tests (sect. 5.4). The second derivative is called the Ricker wavelet (e.g. Sheriff and Geldart, 1995). 101
5.23 Analytical solutions of acoustic wave equation in homogeneous space. Green's functions (left panel) convolved with the source wavelet (inset) give the observed traces (right panel). The Green's function depends on the dimensionality of the problem. Only in 3d the observed trace has the same shape as the source wavelet. In 1d it is its time integral and in 2d it has a non-trivial tail. Calculated based on Igel (2016) after derivation provided by Rienstra and Hirschberg (2019). 102
5.24 Phase analysis of pressure and vertical component of particle-velocity data for the OBS station no. 157. The white star and coloured points denote horizontal locations of the reciprocal source and receivers, respectively. The colourmap encodes the phase of data windowed in frequency at 3 Hz and time (1 s around the first arrivals), wrapped into $[-\pi, \pi)$ range. Left, middle and right panels show synthetic ϕ_{syn}, observed ϕ_{obs} and the residual phase, respectively. 103
5.25 Acquisition geometry of the PROTEUS experiment with the profile in the Christiana sedimentary basin (zoomed in the yellow box) used in the picking-error estimation. Only stations 181, 183, 184, 185, 186, 188, 189, 190, 192, 193 were used. 104
5.26 First arrivals picked for the stations in fig. 5.25. Travel-time of the arrival at station A generated by the source at station B should remain the same after interchanging A and B. . . 104
5.27 Left – schematic plot of shot and OBS location used for a geometrical correction of the travel-times in the reciprocity check. Right – reciprocity matrix with rows and columns representing respective receivers. Value at row i and column j stands for a delay of receiver i with respect to receiver j. Note that the matrix is by definition skew-symmetric. 105

6.1 Example of a pressure wavefield affected by numerical dispersion. Note that the error amplitude varies like $|\cos(2\theta)|$ where θ is the inclination from the grid. The vertical and horizontal directions suffer from the largest error, while the diagonal directions ($\theta = 45°$) are least affected. Parameters used in the simulation: $v_P = 1500$ m/s, $\Delta x = 50$ m, Ricker source with $f_{\text{peak}} = 3.5$ Hz containing some energy above 6 Hz; the spatial-sampling criterion is violated since the scheme requires at least 5 nodes per wavelength. 107

6.2 Effect of land topography on the full-waveform inversion of acoustic (top) and elastic (bottom) data. From left to right – true models, results of inversion using the actual and the flat free surface. Yellow and blue circles in the top-left figure denote sources and receivers respectively. Bottom row adapted from Nuber et al. (2016). 108

6.3 Effect of land topography on a seismic wavefield shown by 4-Hz pressure wavepaths calculated using (a) absorbing, and (b,c) free-surface boundary with (c) and without (b) time-windowing. Different code was used in (a) vs. (b,c) cases. Adapted from Bleibinhaus and Rondenay (2009). 109

6.4 Staircase approximation of the irregular free surface arising in vacuum/air-layer finite-difference methods. Modified after M. Paulatto (unpublished). 110

6.5 Idea of the immersed-boundary method. The physical wavefield propagating in the subsurface (interior grid nodes) is anti-symmetrically extrapolated above the free surface so that the non-physical wavefield at the ghost nodes imposes the zero-pressure boundary condition at the free surface. 111

6.6 Wavefield grid – inner grid (black) with the origin at node (1,1,1) and nx, ny,nz in size is extended by the left, right, front, back and bottom paddings (white with grey corners). . . 115

6.7 Orthogonal point-projection in 2d. Circles of different colors denote ghost nodes of different levels projected onto the surface Γ. From Zhao (2010). 118

6.8 Points used in the immersed-boundary method to extrapolate the wavefield along the normal vector (dashed line) across the free surface into the ghost node G. Here, the extrapolation is done in three steps: 1. values at fictitious nodes F_i are found by interpolation between auxiliary nodes A_{ij}; 2. value in R is found by interpolation between fictitious nodes and intersection I. 3. value in G is set equal to the value in R with the opposite sign. Modified after M. Paulatto (unpublished). 120

6.9 Variable orientation of the extrapolation stencil (fictitious and auxiliary nodes) for the free surface of a variable slope locally exceeding 45°. The same behaviour was implemented in 3d (right). 120

6.10 Exclusion of a fictitious node (grey rhombus) located less than half grid spacing from the free surface to avoid instability of some interpolators. The bottom-right inset adapted from Mulder (2017). 121

6.11 Polynomial (top-left) vs. radial-basis-function (IDW with $p = 2$) interpolation between four randomly generated data-points. The two-dimensional version of IDW is presented on the right. 123

6.12 Ghost wavefield before (left) and after (right) setting to zero during the vacuum step of the wavefield-extrapolation. 125

6.13 Geometry of the 3d orthogonal projection of a point (grey dot with black outline) onto a $z(x, y)$ surface (black curve). The part of the surface enclosed in the search region (dashed square) is denoted by the thick solid line (left). The XY projection of the search region is split into squares with corners being the grid nodes (middle). Each square corresponds to a surface patch generated by a bilinear interpolation between its corners (right). The shortest distance to each patch is found by the damped Newton optimisation. The scheme is readily generalisable to arbitrary surfaces. 127
6.14 Finely sampled surface patch (left), the corresponding objective function (middle), and the projection of the objective function onto the XY plane (right). The objective function expresses the distance between a point on the surface and the reference point (not shown). 128
6.15 Schematic of injecting a source close to a 1d free surface. Here, the method was generalised to 3d, requiring an iterative solution. From Hicks (2002). 132
6.16 Acquisition geometry of the experiments using the flat/rotated free surface: a single source and 101 receivers (yellow and blue dots, respectively); free surface at node $z = 0$. Relative positions and distances between points do not change under rotation of the frame of coordinates around its origin – this property was used to compare results for the flat and rotated free surface. 135
6.17 Shot gathers from the numerical test involving a horizontal planar free-surface using (left) the default `FULLWAVE3D` method and (right) the immersed-boundary method. The relative positions of sources, receivers and the free surface were the same in both cases. 135
6.18 Wavefield snaphots from the numerical test involving a planar free-surface tilted by 5° from the horizontal: (left) air-layer method; (right) immerse boundary method. 135
6.19 Wavefield snapshots from the numerical test involving a planar free-surface tilted by 75° from the horizontal: (top-left) air-layer method; (top-right) immerse boundary method. Note the spurious diffraction pattern resulting from the staircase-approximation for the air-layer method. Bottom – wave reflected from the free surface recorded at the central (51th) receiver for: (blue) air-layer method; (red) immersed-boundary method. 136
6.20 Wavefield snapshots from the numerical test involving a sinusoidal free surface with the wavelength of 10 grid nodes. 137
6.21 Wavefield snapshot after 3 s of the simulation using a 2d topographic surface cut horizontally through the center of the Santorini caldera, homogeneous model with $v_p = 1500$ m/s and Ricker source wavelet with $f_{\text{peak}} = 3$ Hz. 138
6.22 Numerical experiment involving a 3d fragment of the Santorini caldera: (left) three levels of ghost nodes coating the free surface; (top-right) wavefield snapshot sliced above the sea-level – the righthand part of the wavefield with the inverted polarity is the ghost wavefield; (middle-right) free surface topography measured in grid nodes below the top of the model; (bottom-right) data recorded by receivers located in the in-line direction across the center of the domain, below the sea level. 138
6.23 Wavefield snapshots of the simulation using a 3d topographic surface of the Santorini caldera. The 2d cross-sections were cut through the centre of the Nea Kameni island in the middle of the caldera. All boundaries are reflecting to test the stability of the scheme in the presence of strong multiples. 139
6.24 Acquisition geometry of the first field-data example. IDs of the OBS and the shot-line used in fig. 6.25 and 6.26 are highlighted in yellow. Modified after J. Morgan and M. Paulatto (unpublished). 140

6.25 Wavefield snapshot from the first field-data example, superimposed over v_P model, sliced in the in-line direction through the source (OBS 160). Inset – data-derived source wavelet in the time (left) and frequency (right) domain. 140
6.26 Interleaved observed and synthetic data in the first field-data example shown for OBS 160 and shot line no. 26. Left – model with the flat free surface. Right – model with actual topography simulated using the immersed-boundary method. Vertical axis – no. of time-samples ($\mathrm{d}t$ was 0.0025 s). 141
6.27 Acquisition geometry (left) of the second field-data example (red rectangle), on land stations were analysed (red dots). Right – position of ghost nodes of teh first level. 142
6.28 Wavefield snapshots from the second field-data example sliced through stations 825, 829, 822 (from top to bottom) with YZ, XZ and XY planes (from left to right). 142
6.29 Flow of data in the `IMMERSE` software. The source-injection pre-calculator is implemented in the HicksPy code (sect. 5.6). 147
6.30 Acqusition geometry of the numerical experiments performed to study the convergence of the IB method. The cross-sections at $y = 500$ m are shown. The black star, circle and curve denote the source, receiver and irregular free surface, respectively. 148
6.31 Convergence of the results of the experiments from fig. 6.30 to the reference 'true' solutions. Left: traces observed at the receiver. Right: maximum error normalised to the coarsest-grid case. The 'ibm' and 'vac' stand for the immersed-boundary and vacuum method, respectively, with the grid interval indicated in the parentheses. 149

7.1 Low-pass Butterworth filter: (left) power spectrum of a filter with 4 poles and cut-off frequency of $\omega_c = 100$ rad/s denoted by the green vertical line. Note the semi-log scale and the flatness of the passband; (right) amplitude spectrum of Butterworth filters of 1 Hz cut-off frequency and different order n. 154
7.2 Effect of trace padding on filtering the field data (OBS 108, hydrophone channel, line 29, trace no. 250); purple – no padding, red – 5000 ms applied to both trace ends. 155
7.3 Data filtered (top) with and (bottom) without taper applied prior to the filtering. The difference indicated by the arrow is quite subtle. The Y axis in time-samples (0.0025 ms each). 155
7.4 Top and bottom mutes with different widths of the sine-squared taper starting at the 100th and 400th sample, respectively. 156
7.5 Different stages of processing the field data compared to the synthetic ones shown with a receiver gather (top panel with 0-7.5 s time on the Y axis) and its amplitude spectrum (bottom panel with 0-10 Hz frequency on the Y axis). Note the effect of muting on the amplitude spectrum. Arrivals before/after the direct wave highlighted in green/blue, respectively. Offsets shorter than 5 km excluded from the inversion highlighted in grey. 157
7.6 Flowchart of deriving the source wavelet from the field data using the Wiener filter. 158
7.7 Seismic phases and corresponding impulse response (Green's function). 159
7.8 Nearest-shot traces for all WHOI stations, unfiltered and aligned on the predicted water-wave arrival time;. The hydrophone component displays a variable shape of the negative pulse arriving right after the strongest positive pulse. Station IDs are shown next to the traces. Y axis in seconds. Modified after Heath (2019). 161
7.9 Map of the stations from which the traces were selected for stacking and wavelet extraction; axes scaled in local coordinates (in metres). For some stations more than one trace was used. Inset – offset distribution of the selected traces. 162

7.10 Selected traces (left) before and (right) after alignement using the cross-correlation lag. Different colors are used only to help discriminate between neigbhouring traces. X axis in seconds. 162
7.11 Stacked trace (using all the traces from 7.10) raw (left) and (right) bottom-muted with a 50-samples sine-squared taper at a time just before the arrival of the first multiple; the arrival time was calculated using 1500 m/s water column) and the coordinates of the receiver same as for the trace used for alignment. X axis in seconds. 162
7.12 Initial-guess (left) and the final (right) wavelet. The former was defined as a Kronecker delta at zero time filtered with a minimum-phase filter. X axis in seconds. 163
7.13 Observed (blue) vs. synthetic (red) data generated using (left) the initial-guess and (right) final wavelet. The traces on the right overlap almost entirely. X axis in seconds. 163
7.14 Wiener filter used to convert the synthetic into observed data in fig. 7.6b. X axis in time-samples ($dt = 5$ ms). 164
7.15 Observed vs. synthetic data calculated for OBS 104, line 37 using the final wavelet from fig. 7.12. Y axis in time-samples ($dt = 2.5$ ms). 164
7.16 Ingredients needed to build an accurate starting model for FWI. 165
7.17 Travel-time tomography model of Heath et al. (2019) sampled on a vertically-sheared 200x200x200 m grid. Minimum and maximum value of v_P is 1.69 and 9.22 km/s, respectively. 166
7.18 Water column velocity profiles (left, with v_P [m/s] and depth [m] on the horizontal and vertical axis respectively) determined from the bathythermographic (XBT) measurements at various locations in the study area (right, axes in geographical coordinates). Modified after Hooft (2015a). 166
7.19 Starting model of v_P (m/s) prepared from fig. 7.17 following alg. 7.5 (left) is characterised by a smaller topography imprint compared to the original version from Heath (2019). Axes scaled in local coordinates, in metres. 167
7.20 Starting model of v_P (m/s) overlain with the topography of the seabed and land surface (black line). Axes scaled in local coordinates, in metres. 168

8.1 Topographic map of the Aegean Sea: (left) regional tectonics from Huff et al. (2021), (right) zoom into the study area. 171
8.2 Hypothesised geological cross-section through the Hellenic Volcanic Arc, from the molassic back-arc Cretan Basin to the Cycladic Plateau near Ios. The country rock in the topmost 4 km beneath Kolumbo consists of interlaced metamorphic and non-metamorphic units. Modified after Kilias et al. (2013). 171
8.3 Seismic-reflection time-section through shallow sediments around Kolumbo. The K5 unit is associated with the most recent 1650 AD eruption. Note the bright spot at about 200 m below the crater floor (Hübscher et al., 2006). VC – volcanic cone. Modified from Hübscher et al. (2015). 173
8.4 P-wave velocity model (in m/s) across Kolumbo obtained by Heath et al. (2019), sliced at $y = 8.75$ km. It served as a starting model for this study. 173

8.5 Acquisition geometry of the domain of study (black rectangle) with the coverage density represented by lines between sources and receivers with offsets up to 20 km. The local system of coordinates was based on the Universal Transverse Mercator coordinates rotated 25.5° counterclockwise from north with the centre $(0, 0)$ m at ca. $36.4°\mathrm{N}, 25.4°\mathrm{E}$. Red and blue triangles denote OBSs from WHOI and SIO pools, respectively. The leading digit '4' is not shown in the station numbers. Shots are marked by black dots with line numbers in white squares. The large white rectangle denotes a potential extension of the domain for future inversions. 176

8.6 Histogram of water depths at shot positions in the study area. Most shots were fired in water of 250-450 m depth. 177

8.7 Histograms of source-receiver offsets for every OBS (semi-transparent colourscale) in the study area. 177

8.8 Phase plots (Shah et al., 2012a) used to determine the minimum invertible frequency. Each dot represents a phase residual (phase of synthetic data minus phase of observed data, wrapped into $[0, 2\pi)$ rad interval) of the signal extracted at 2, 3 or 4 Hz (from top to bottom, respectively). Colour scale from $-\pi$ (blue) to π rad (red), green denoting a zero residual, and thus good data fit. Incoherent colour variations indicate low S/N ratio. Note that the noise affects the long-offset data at each frequency. White star denotes the location of the OBS. Axes scaled in local coordinates in metres. 179

8.9 Maximum grid sampling (colour scale) for a numerical scheme accurate for at least 5 grid nodes per wavelength, as a function of the highest frequency of the source (X axis) and the lowest velocity present in the model (Y axis). 181

8.10 Phase residuals (left) at 3 Hz calculated for shots recorded by OBS 4176 and acquisition from fig. 8.5 with source-receiver offsets < 5 km excluded. The colour scale is the same as in fig. 8.13. The right figure shows the shots input to the inversion (dark blue) and shots excluded due to being noisy or cycle-skipped (grey). Axes scaled in grid nodes ($dx = 50$ m). 182

8.11 Anomaly recovered (final minus background model) after 60 iterations of inverting synthetic data generated from the same model as in fig. 5.20a. Top left – x slice, top right – z slice, bottom – y slice. Dashed white lines denote location of slices. Axes scaled in grid nodes ($dx = 50$ m). 183

8.12 Evolution of the objective function across iterations for each OBS (colour-coded). 185

8.13 Phase residuals at 3 Hz for four stations (from left to right: 4105, 4124, 4177, 4178) at iterations 0 (top) and 80 (bottom). The white star shows the OBS location, the coloured dots show shot locations. 186

8.14 Histograms of phase residual at 3 Hz for OBS 4137 at iteration 0 (red) and 80 (black, semi-transparent). 186

8.15 Data fit in the (left) time and (right) frequency domain for station 4137, and multiple shot-lines at the iteration 0, and 80 (top to bottom), filtered at high-cut frequency of 3 and 4.5 Hz, respectively. Each plot consists of interlaced blocks of synthetic and observed traces, 10 traces each, starting with synthetics. Data are normalised so that the maximum value of each trace is equal to 1. Each vertical arrow denotes 5 s time-window. 187

8.16 P-wave velocity model (left) and the anomaly relative to the starting model, in % (right), recovered after 20 iterations at the low-pass frequency of 3 Hz. 188

8.17 Slices through the final P-wave velocity model around Kolumbo: (top) NW-SE profile at $y = 8.75$ km; (middle) SW-NE profile at $x = 13.75$ km (shown as a perturbation of the starting model); (bottom) depth slice at $z = 2.75$ km. Dashed lines show locations of the slices. 189
8.18 Model recovered after inverting synthetic data generated from the starting model perturbed by a spherical LVZ added underneath Kolumbo. An anomaly as small as 200 m (FWHM) can be recovered, provided it is strong enough (here -50%). The final anomaly is smeared with a slightly under-recovered amplitude. 192
8.19 Partially molten olivine grain boundaries in an olivine-basalt aggregate. The light grey grains are basalt glass, the darker are olivines. Melt resides primarily in large pockets (black arrow), triple-junctions (white arrow) but also thin layers between grains (double white arrow), the latter two are poorly resolved in this back-scattered-electrons (BSE) image. Black cracks are a result of quenching. From Cmiral et al. (1998). 194
8.20 Flow of calculation of the effective P-wave velocity using the self-consistent scheme of Berryman (1980) and the Hashin-Shtrikman average (Hashin and Shtrikman, 1963). The input variables are in grey ellipses. The Voigt average is an arithmetic mean weighted by volume fractions (Mavko et al., 2009). The effective velocity is by far the most sensitive to pore-fluid bulk modulus and aspect-ratio highlighted in yellow. 196
8.21 P-wave velocity of a two-phase composite calculated with the self-consistent approximation (Berryman, 1980) in the low-frequency, equilibriated, limit starting with air-filled pores and subsequent fluid substitution according to Gassman's equation (Mavko et al., 2009). White area denotes Hashin-Shtrikman bounds (Hashin and Shtrikman, 1963). coloured lines correspond to different values of the pore aspect-ratio. Sensitivity to other parameters can be found in appendix 8.C. 197
8.22 Schematic model of Santorini and Kolumbo magmatic system(s). Not to scale. Modified after Rizzo et al. (2016). 197
8.23 Screenshots of a supplementary `Jupyter` notebook `Kolumbo_volcano.ipynb`. 200
8.24 Slices through the final model at $z = 2.75$ km b.s.l. The slice after 20 iterations is shown in fig. 8.16. 202
8.25 Slices through the final model at $x = 13.25$ km. 203
8.26 Slices through the final model at $y = 8.75$ km. 204
8.27 Final minus starting model (left) in the inversion of synthetic data generated from the 'true' model which was the starting/background model perturbed with a negative anomaly underneath Kolumbo. Right – true minus background model. The anomaly was a spherically-symmetric 3d Gaussian with a peak amplitude of -1% of the background model and a full-width at half-maximum (FWHM) of 800 m (right). 205
8.28 Same as fig. 8.27 but for anomaly with a peak amplitude of -5%. 206
8.29 Same as fig. 8.27 but for anomaly with a peak amplitude of -50%. 206
8.30 High versus low frequency limit (Mavko et al., 2009) used in the calculation of the effective parameters. All other figures of this appendix are for the low-frequency limit. Note that for a given v_P, melt fraction is lower in the low-frequency limit, and the difference increases with decreasing v_P. Blue and red curve denote the Hashin-Shtrikman lower and upper bound, respectively. 207

8.31 Analysis of the relationship between v_P and melt-fraction of a two-phase effective medium (Berryman, 1980) in terms of sensitivity to different properties of the solid and fluid phases. Each figure presents a single-parameter study (values in color, units in captions) with other parameters fixed. Area between Hashin-Shtrikman bounds denoted in white. 208

List of Tables

1.1 Seismic signature of mush zones vs. magma chambers. After M. Paulatto (unpublished). 32

4.1 Time-line of the FWI development in the context of the exploration industry until 2013. Later years have seen a flurry of incremental improvements and new applications of the method, see Wang et al. (2021) for a recent example of high-frequency elastic inversion. Great overview of the history of the method as well as the underlying theory is provided by Virieux and Operto (2009), Bleibinhaus (2014) and Virieux et al. (2014). The word 'spectacular' refers here to improvements in model resolution and fidelity relative to the travel-time tomography. 52

4.2 Formulations of seismic tomography in different Fourier domains of space and time variables. 61

4.3 Comparison of different inverse-scattering (imaging) methods. In linear methods the background velocity model, and thus the wave paths, remain constant. 61

5.1 Time-line of previous approaches to modelling and inversion of velocity/multi-component data. 72

6.1 Time-line of major developments of the immersed-boundary method for modelling seismic waves. 112

8.1 Time-line of major studies of the Kolumbo volcano. 172

8.2 Major-element composition of the K5 pumice (fig. 8.3), given on a volatile-free basis and normalised to 100 wt.%. The sum of major elements without water is 99.96 wt.%. The values are the average of NA007-058 NA007-071, NA007-073 NA007-075, NA007-085, NA007-090, NA014-019 samples reported in the supplement of Klaver et al. (2016). 193

List of Algorithms

3.1 Timeline of the acquisition and pre-processing of the PROTEUS data used in this study. 42
5.1 Quality control of the implemented particle-velocity simulator. 86
5.2 Estimation of travel-time picking error using the reciprocity principle. 103
6.1 Simulating an irregular free surface. 114
6.2 Orthogonal point-projection posed as a global optimisation problem. 118
6.3 Damped Newton method of finding the shortest distance from a surface patch. 128
6.4 Source injection at/near an irregular free surface. 132
7.1 Processing of the field data. 154
7.2 Filtering workflow that minimises wrap-around effects. Applied to each trace. 154
7.3 Muting workflow. 156
7.4 Source-wavelet derivation. See also the flowchart 7.6. 158
7.5 Embedding finely-sampled seabed on a coarse model grid. 167

Listings

5.1 Simple Python solver of the second-order 1d acoustic wave equation with first-order free-surface boundaries and first-order quiescent initial conditions. It uses the second-order approximation of the second derivative (both in time and space). Definitions of the source and receiver operators are provided in the appendix. Note how free-surface boundary conditions are imposed implicitly by preventing the first and the last grid node from being updated. 88

Part I

Introduction

1 Volcanology background

The importance of studying volcanoes is at least two-fold. First, the number of people potentially affected by volcanic activity is increasing with the ever-growing human population and its reliance on air transportation, especially vulnerable to explosive eruptions. Second, volcanoes are windows into fundamental processes that shape the Earth and other planets, in particular into the formation of the continental lithosphere.

Volcanic eruptions in which molten rocks emerge on the Earth's surface are expressions of much deeper magmatic processes. A thorough understanding of these processes is key to managing the volcanic risk and unlocking the inner workings of our planet.

Even though the last 30 years have seen a flurry of progress due to increased volume of available data as well as advances in instrumentation and computing (Cashman et al., 2013; Magee et al., 2018), our understanding of volcanoes is still fragmentary. In particular, despite being a cornerstone of petrological models, no magma chamber has ever been observed under a continental arc volcano directly. Models of generation of silicic magma at arc volcanoes as well pre-eruptive storage and mobilisation require better physical constraints on the geometry, melt content, and the interconnections of magma bodies in the crust (Cashman et al., 2017).

1.1 Basic concepts

The central concept in volcanology is that of a magmatic system (Sparks et al., 2019). It is where the magma is generated, stored, differentiated and finally makes its way to the surface in the phenomenon known as a volcanic eruption. Below, all these terms are given a precise meaning.

1.1.1 Definitions

Three domains of the subsurface relevant to volcanology can be distinguished (Sparks et al., 2019):

1. ***host rock*** (fig. 1.1a,b) – with either no melt (i.e. at a sub-solidus temperature) or with the unconnected melt of very low fraction present (implying a super-solidus temperature).

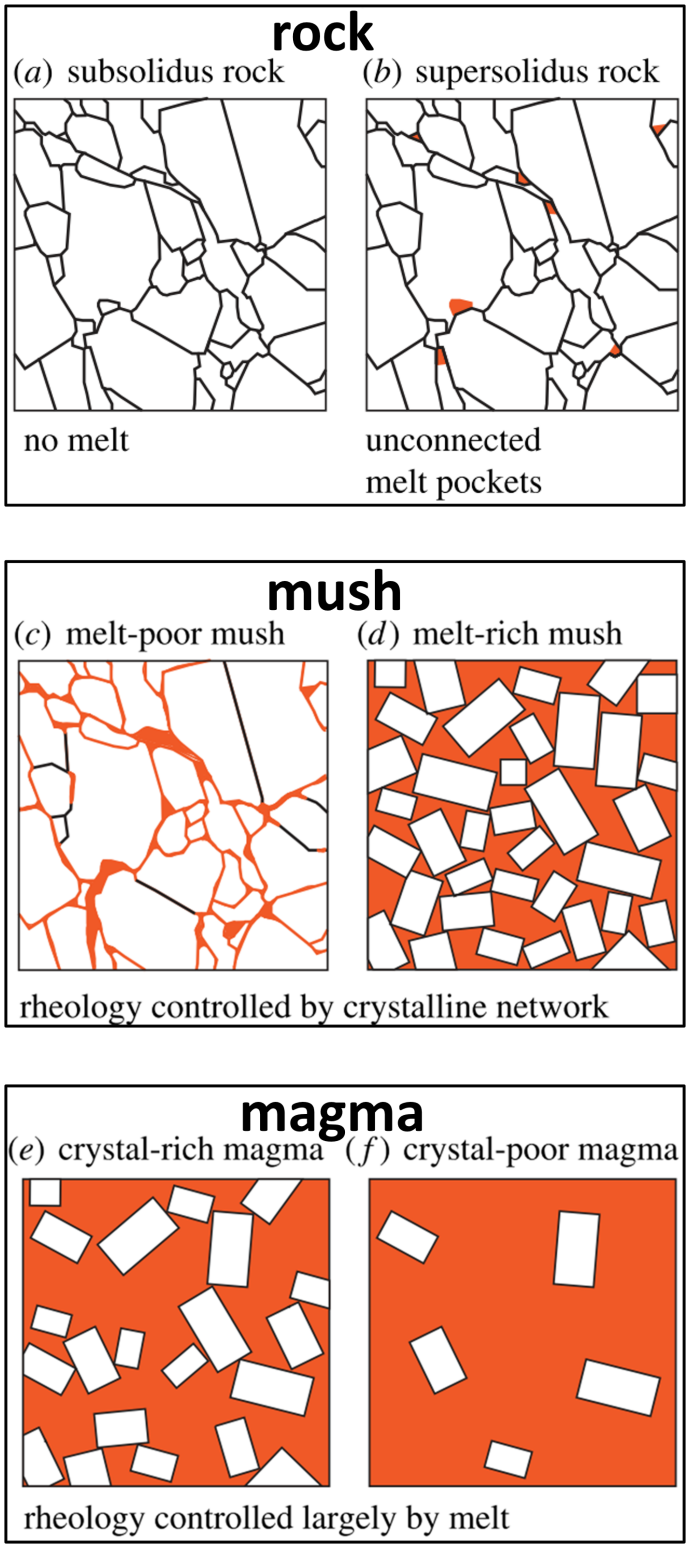


Figure 1.1: Three major domains of a magmatic system. The transition between rock and mush (melt-connectivity transition, MCT) is characterised by a sudden drop in viscosity and strength. Upon the mush-magma transition (MMT), melt starts to control the rheology of the system. Modified after Sparks et al. (2019).

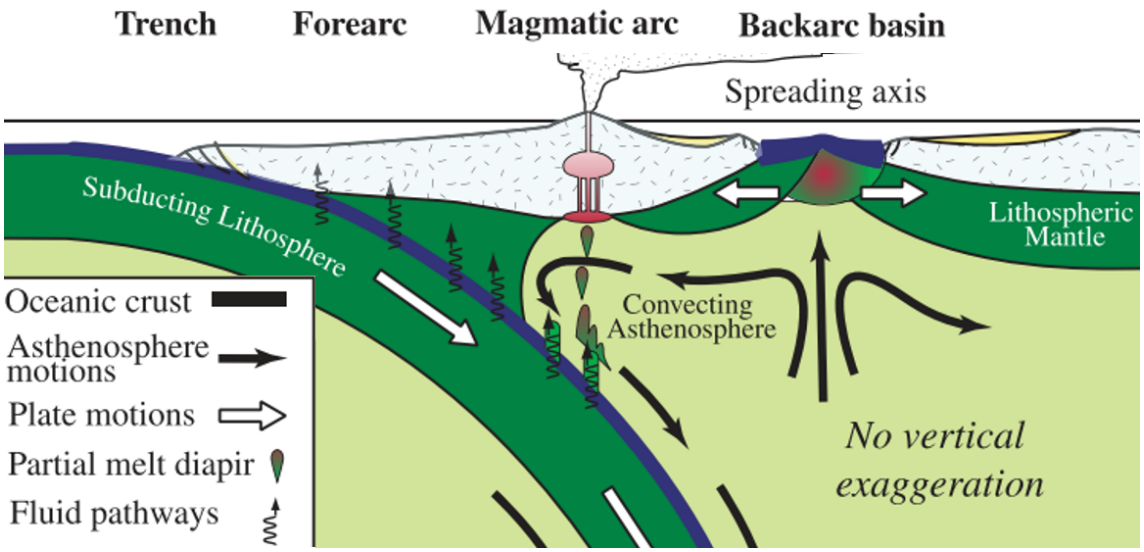


Figure 1.2: Back-arc volcanism in a continental-ocean collision zone. Modified after Stern and Dickinson (2010).

2. ***mush*** (fig. 1.1c,d) – mostly crystals, connected melt.
3. ***magma*** (fig. 1.1e,f) – mostly melt (high fraction).

There exist two sharp transitions of rheology and elastic properties between these domains. A step-change between the domain 1 and 2 called the *melt-connectivity transition* (MCT) is characterised by a rapid decrease of viscosity and strength. A jump between the domain 2 and 3 called *mush-magma transition* (MMT) changes the controlling factor for rheology from crystalline network to melt, respectively. Magma and mush domains consist of three phases: melt, volatiles (e.g. supercritical H_2O and CO_2) and crystals. These phases have very different physical properties which makes any attempt of modelling such a system challenging (Sparks et al., 2019).

Let us define a *magmatic system* as a volume of sub-solidus rocks + *melt reservoir*(s) that has contained magma at least at some point in the past, even if currently melt is stored as mush. The term *magma chamber* is used here to mean a storage place for magma, i.e. a zone with a high melt fraction. This is not synonymous with a melt reservoir, which is a broader rheological term that does not assume any time-scales or volumes, in particular does not require the chamber to be large and long-lived (Sparks et al., 2019).

1.1.2 Magma generation

Rock melting and further seperation of melt and crystals to generate magma occurs in a variety of tectonic settings – whenever there is a pressure decrease and/or enough heat and/or water supplied (Best, 2002). Here, only continental-ocean subduction zones and resultant continental arcs will be discussed.

When the slab reaches about 60-100 km below the surface (England and Katz, 2010; Perrin et al., 2018) its water-saturated sediments and altered (serpentinised) oceanic crust (basalt and gabbro) undergo metamorphic phase change into an eclogite (consisting mainly of garnet and pyroxene) which can be thought of as a high-pressure version of a basalt. This reaction releases bound water from hydrous minerals such as amphiboles, as they get transformed into the anhydrous garnet and pyroxene. The

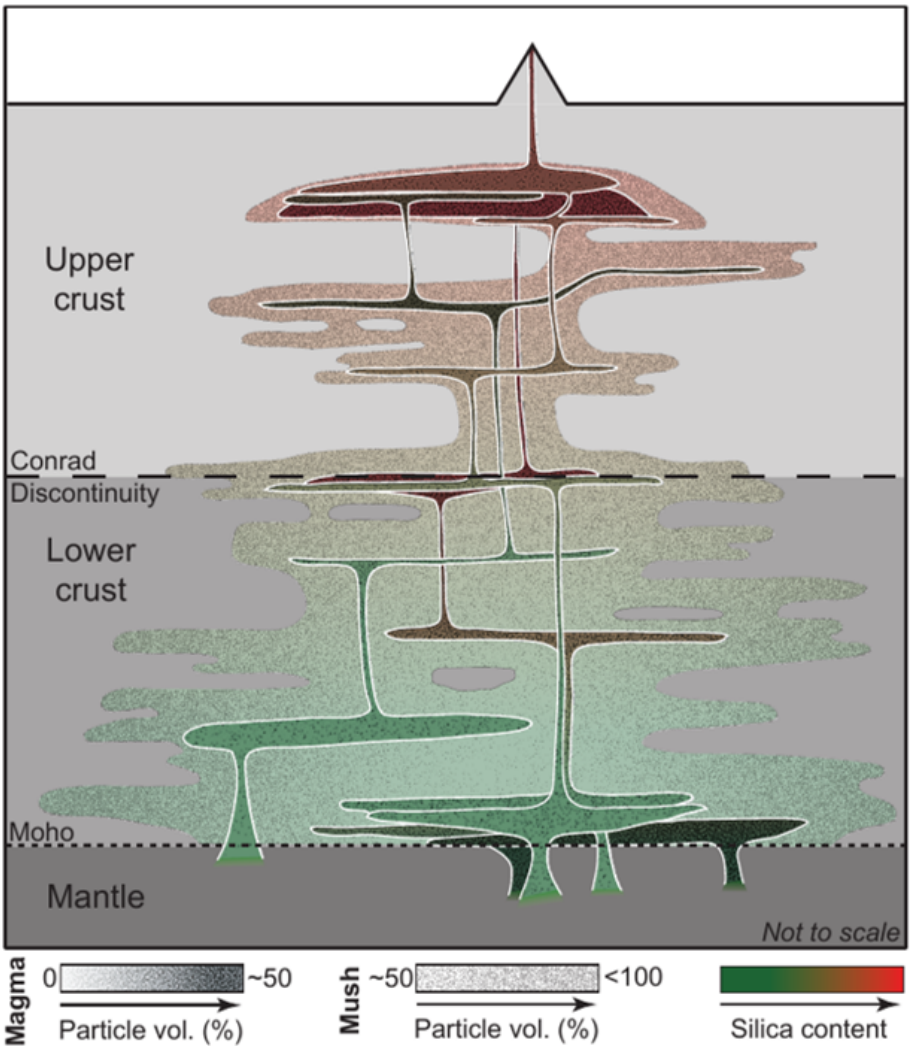


Figure 1.3: Schematic depiction of a magma plumbing system. From Magee et al. (2018).

largest contribution to the water release is likely to occur during the blueschist-eclogite transition at 300-500°C (Paulatto et al., 2017).

The water acts as a crystal network modifier breaking up Si-O polymers of SiO_4 tetrahedra and lowering the melting temperature. The H_2O-enriched (typically a few % of weight) mantle wedge starts to melt at much lower temperatures than the dry one (Best, 2002). It is worth emphasising that it is the mantle that contributes to melt-generation the most, not the slab, unless the latter is very young and hot. Heating the slab up to the required temperature would otherwise take too long (Perrin et al., 2018).

1.1.3 Magma storage

There has been a long-standing debate about the physical state and spatial distribution of the melt stored in the crust (Cashman et al., 2017). Most systems are likely to consist of all three domains (1.1) at once, with melt reservoirs either connected or separated by a sub-solidus rock (Sparks et al., 2019). An emerging paradigm is the one of magmatic systems being vertically extensive, transcrustal, interconnected networks of magma conduits and magma/mush reservoirs (e.g. Cashman et al., 2017) referred to as *magma plumbing systems* (fig. 1.3). In such a sytem, parts of the erupted magma may

be stored cold (below the solidus) and later rejuvenated (e.g. Cooper and Kent, 2014) with melt-rich magma chambers being only short-lived features.

It is worth noting that magma chambers beneath mid-ocean ridges are *not* a subject of similar controversy and have been detected by means of reflection seismics at East Pacific Rise (Detrick et al., 1987; Kent et al., 2000), Lau back-arc spreading center of the Tonga trench (Collier and Sinha, 1990) and Mid-Atlantic Ridge (Singh et al., 2006), which represent slow, intermediate and fast-spreading centres, respectively.

1.1.4 Magma differentiation

As the magma ascends from its source region to intercrustal reservoirs, a place where it stalls, it undergoes compositional changes (differentiation). These changes are caused by:

- ***fractional crystallisation*** – sometimes referred as AFC (Assimilation, Fractional Crystallisation). In this old model, magma differentiation happens in mostly molten magma chambers by fractional crystallisation and melting and assimilation of host rock. As the more refractory minerals crystallise, they fall to the bottom of the chamber and form layered mafic and ultramafic cumulates. As this process progresses, the residual melt becomes more and more silicic;
- ***mixing/mingling*** of magmas from other reservoirs, often when a hotter mafic magma intrudes into cooler more felsic magma and often triggers an eruption (Sparks et al., 1977); surface studies suggest mixing of magmatic products is a ubiquitous process (Edmonds et al., 2019);
- ***assimilation*** (via mixing) of partial melts of the crust – this process is characteristic of continental settings (Solano et al., 2012);
- ***reactive flow*** – another mechanism that has gained traction in the last couple of decades is differentiation by reactive flow through the crystal mush (Jackson et al., 2018).

The differentiation leads to magmas that are more felsic (SiO-enriched) than the parental mantle melt. In this way a plumbing system may be thought of as a filter that transforms the mantle into the continental rocks.

Continental arcs display a great diversity of magma compositions from mafic basalt to felsic rhyolite with intermediate rocks (dacite and andesite) being dominant. Commonly a single eruption brings to the surface more than one type of magma, often in a form of mafic enclaves created by rapid crystallisation of mafic magma when brought in contact with felsic magma (Martin et al., 2006).

1.1.5 Magma mobilisation

The formation of mobile (eruptible) magmas is postulated by Sparks et al. (2019) to be governed by four main mechanisms:

- direct intrusion at a rate fast enough to prevent crystallisation,
- reactive flow through the mush with melt extraction induced e.g. by compaction, volatile activity or compositional convection (Jackson et al., 2018),
- 'defrosting' by the heat or volatiles, both of which are most likely supplied by magma recharge (Burgisser and Bergantz, 2011) and can be thought of as an indirect influence of an intrusion,
- merging of smaller magma bodies formed by the previous three mechanisms.

The first two are very slow, whereas the last two are relatively fast processes (fig. 1.4). There is no reason to think that all these mechanisms cannot operate in tandem.

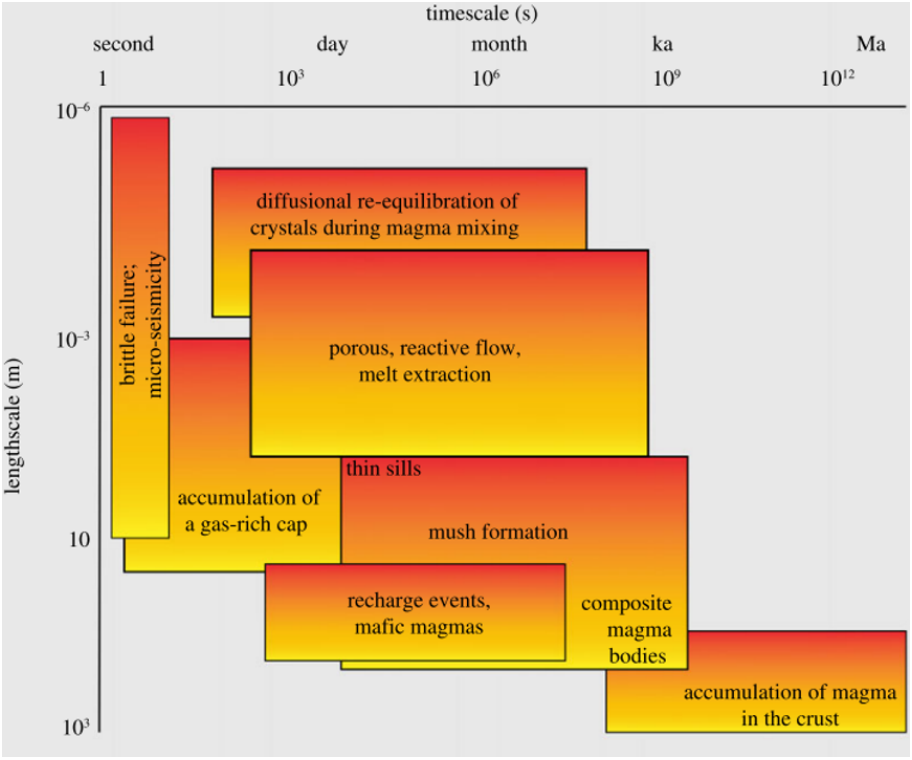


Figure 1.4: Time- and length scales of crustal magmatic activity. From Edmonds et al. (2019).

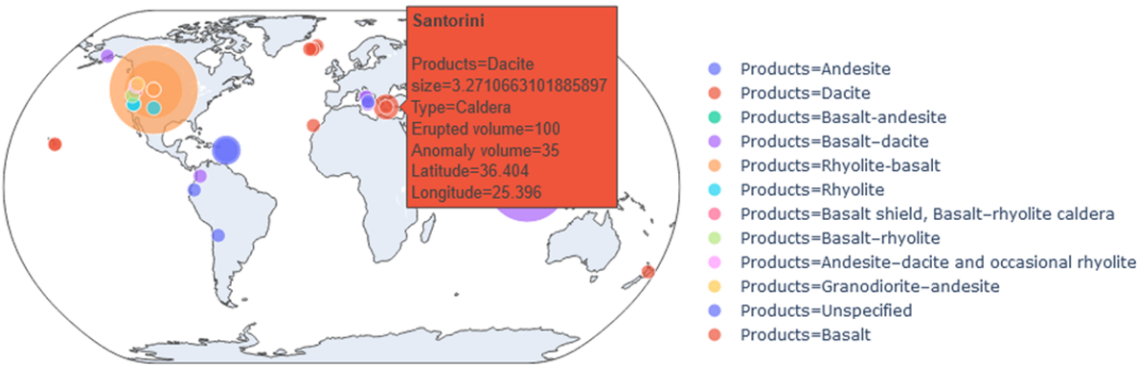


Figure 1.5: Map of all the volcanoes studied by means of the seismic tomography. The interactive visualisation is available at `https://kmch.github.io/`. The database has been compiled by Lees (2007), and largely extended by M. Paulatto (2019, personal communication).

1.2 Seismic studies

Volcanology is an interdisciplinary field of research taking advantage of various specialised scientific disciplines. One of the most important, petrology, is particularly useful in providing constraints on pressure and temperature conditions, as well as the duration and sequence-order of thermodynamic processes of the past eruptions (e.g. Druitt et al., 2012). These tools, however, do not give much insight into geometry (shapes and volumes) of the the magmatic system in its current state. On the other hand, geophysical methods can be used to locate changes in physical parameters related to the current structure beneath the volcano. For example, low-velocity anomalies detected by seismic travel-time tomography can be associated with a high melt content.

mush zone	**magma chamber**
-10% to -30% velocity anomalies	-30% to -60% velocity anomalies
High Vp/Vs ratio	Vp/Vs$\rightarrow \infty$, Vs=0 (S-wave shadow zones)
Anisotropy	No anisotropy
No reflections or diffractions	Strong reflections and diffractions

Table 1.1: Seismic signature of mush zones vs. magma chambers. After M. Paulatto (unpublished).

Need for high resolution f. A snapshot of an interactive visualisation that compiles the results of most volcano-targeted tomographic studies conducted to date, is shown in fig. 1.5.

1.3 Open questions

1. How is melt distributed in the crust?

What are the depths and types of reservoirs (mush/magma)? The models range from single sills to large mush regions (Cadoux et al., 2014; Druitt et al., 2016). A multi-level system is likely to be a common feature for all arc volcanoes of similar crustal thickness. Related questions concern:

- What is the depth, geometry and melt content of magma reservoirs distributed throughout the crust and at the Moho?
- Are these magma bodies connected by dike systems or by vertical, crystal-rich complexes?
- Is the mush present in the upper crust?
- What are the pathways that link deep sources of melt and the active volcanoes?

2. How long do shallow magma chambers live?

Caldera-forming eruptions imply that large magma bodies do exist, however large melt-dominated zones have not yet been observed directly (Lees, 2007). Lack of direct evidence for them may be related to the resolution limit of the seismic methods used so far (see sect. 1.2) or their current absence beneath studied volcanoes may suggest that magma bodies are transient, short-lived features features. This is supported by the observation that phenocrysts get suspended in melt for a relatively short time (up to a few centuries) before eruption (Druitt et al., 2012) and the spread of their ages suggests that some of them are stored cool/cold, i.e. close to/below the solidus for a much longer period before being remobilised (Jackson et al., 2018). Other relevant questions to answer include:

- Is a shallow magma chamber a pre-requisite for an eruption?
- Space problem – how is the room for a magma chamber made?

3. Where does the differentiation occur?

Three hypotheses can be formulated:

- Crystallisation of mafic melts occurs in shallow crustal magma chambers (Gelman et al., 2013);
- Magma evolves continuously as it resides in, and moves through, multiple levels of magma reservoirs (e.g. Tarasewicz et al., 2012). Jackson et al. (2018) propose that magma storage and differentiation occrus primarily in mush zones by reactive flow, rathern in magma chambers by fractional crystallisation;
- Differentiation and/or mixing with melts of surrounding rock occurs almost entirely in the mid-lower crust in the regions called deep crustal hot zones, DCHZ (Annen et al., 2006) or melting-

assimilation-storage-hybridisation (MASH) zones (Hildreth and Moorbath, 1988). The only process that occurs in the upper crust is accumulation which may result in solidification and/or eruption.

4. How can magma layers coalesce so rapidly?

Separate magma bodies can coalesce rapidly prior to or during eruptions (Druitt et al., 2012). Understanding this process is likely to be a major direction of future research (Sparks et al., 2019) and can aid in the volcanic eruption forecasting (Brenguier et al., 2008).

2

Study area

SANTORINI and other islands of the Cyclades archipelago are popular tourist destinations. Santorini itself, with its dramatic views and idyllic architecture, is visited annually by a hundred times more tourists per unit area than the country's average.

The mechanisms responsible for the region's natural appeal are also a source of a serious natural hazard. The Aegean sea tops the list of the most seismically active parts of Europe (Papazachos et al., 2000). Its long tectonic history is reflected in a complex network of faults, many of which are active today. In 1956 the Santorini-Amorgos fault zone hosted a devastating M7.5 earthquake. Santorini itself, with the last eruption 5 years before the Santori-Amorgos event, has been actively recharging its magma reservoir after the Late Bronze Age eruption – the largest explosion in the past 10,000 years of the Earth's history. With a record of several violent caldera-forming eruptions, it remains one of the most hazardous volcanoes in the world capable of Phreatoplinian eruptions generating a tsunami and ejecting pyroclastic material more than 25 km into the atmosphere. Less than 10 km from Santorini lies Kolumbo, another volcano that experienced a caldera-forming eruption in 1650 AD. The inevitable renewal of volcanic and/or tectonic activity is going to have a severe economic impact and may affect the region by tsunami generation and air-traffic disruption.

To credibly assess the volcanic risk, a better understanding of the underlying magmatic processes is needed. The evaluation of the earthquake hazard will benefit from a delineation of the major transcrustal faults by seismic methods. This chapter summarises the current knowledge of the study area with the focus on the Santorini volcano. More details on the Kolumbo volcano are given in ch. 8.

2.1 Tectonic setting

The Santorini volcanic field[1] (Sigurdsson et al., 2006) comprises the inactive volcanic island of Christiana, the Santorini caldera and more than twenty smaller submarine volcanoes extending about 20 kilometres to the northeast of Santorini, the largest of which is Kolumbo. The Santorini caldera forms

[1] Also known as the Christiana-Santorini-Kolumbo volcanic field (Nomikou et al., 2019).

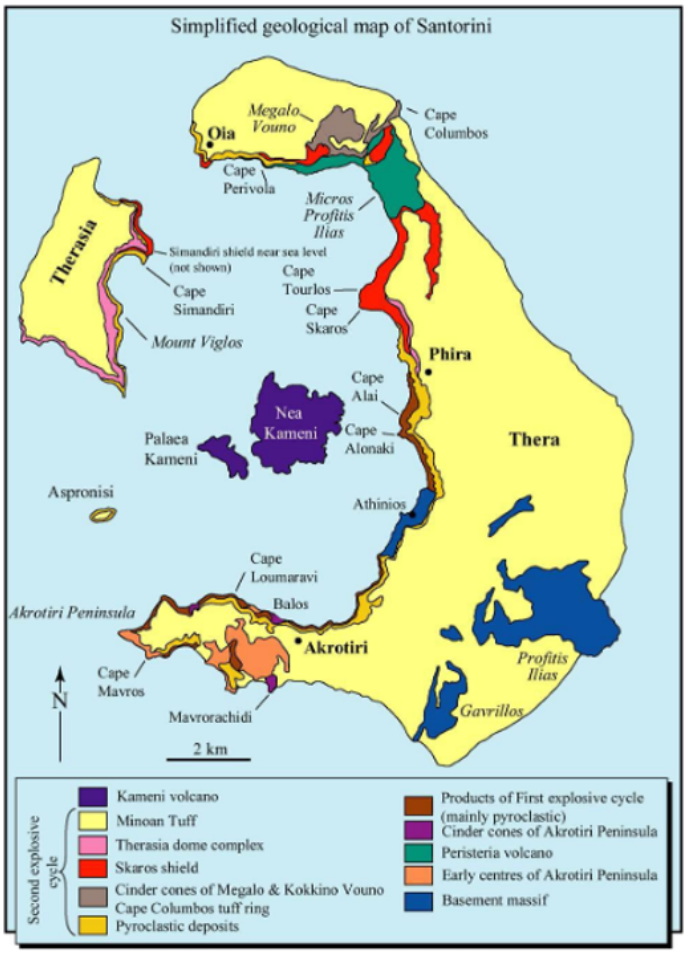


Figure 2.1: Simplified geological map of Santorini. After Druitt et al. (1999).

a small archipelago consisting of five islands: Thera, Therasia, Aspronisi forming the fragmented caldera rim and the two central islands of Palaea Kameni and Nea Kameni islands that are a subaerial expression of a lava shield, a living proof of a dome-forming phase of the volcanic cycle (fig. 2.1).

2.1.1 Regional scale

The Santorini volcanic field is located at the southeastern edge of the Cycladic block[2], a part of the Aegean microplate (fig. 2.2). The two major driving forces of the Aegean tectonics are:

- subduction of the oceanic Ionian plate beneath the continental Aegean plate inducing extensional stress in two directions: N-S extension due to a slab roll-back and trench-parallel extension close to the trench due to its curvature.
- westward extrusion of the Anatolian plate along the North Anatolian Fault with dextral (right-lateral) strike-slip motion (McKenzie, 2020).

The subducting slab has been imaged by seismic travel-time (Spakman et al., 1988; Bijwaard et al., 1998; Piromallo and Morelli, 2003) and full-waveform inversion (Fichtner et al., 2013). It extends to the lower mantle where it is probably anchored, given the longevity of this subduction zone (Jolivet and Brun, 2010).

[2]Also called the Central Aseismic Plateau or Attico-Cycladic Massif.

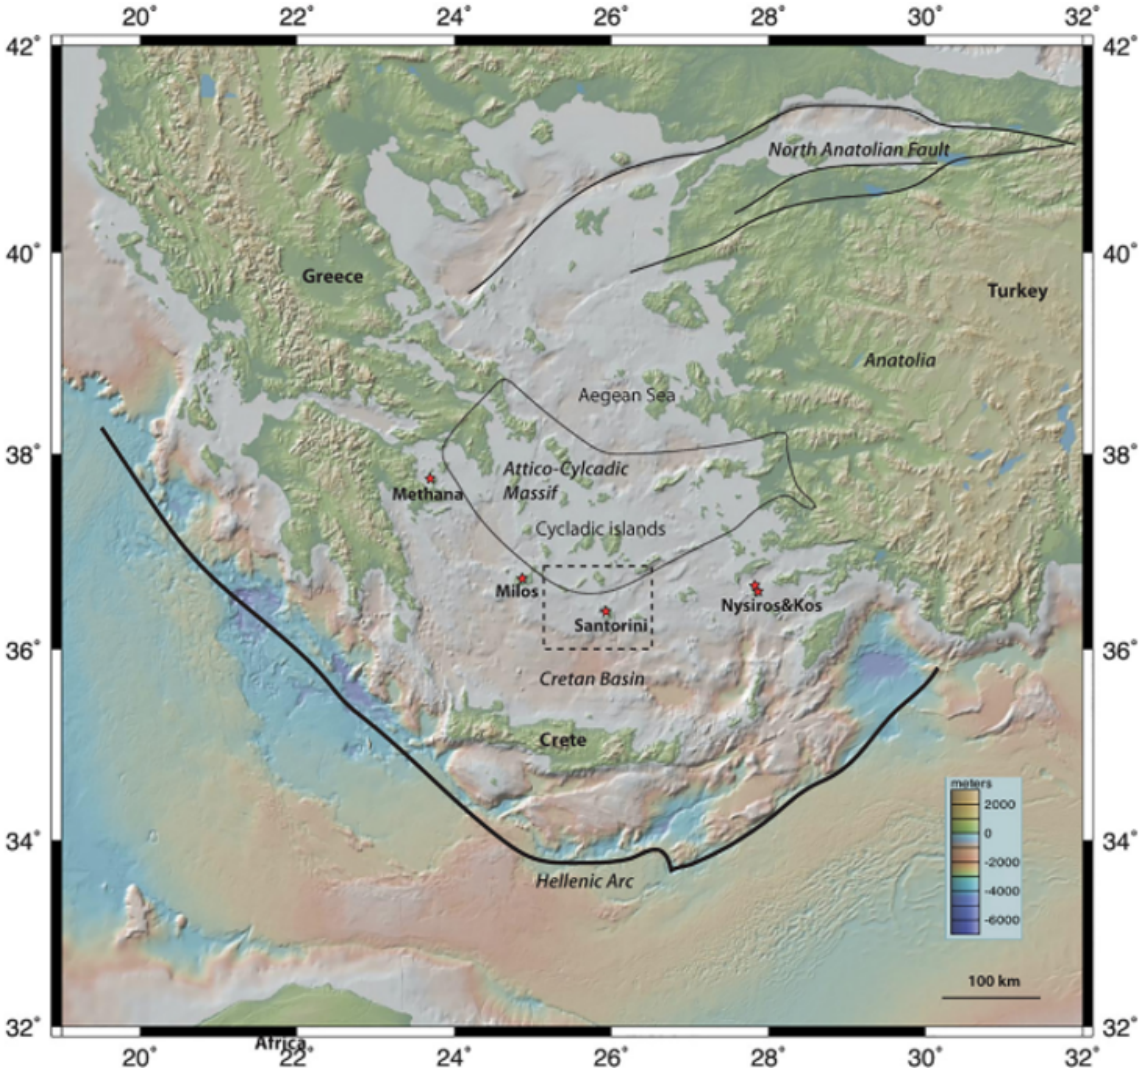


Figure 2.2: Regional tectonics showing Santorini at the centre of the Hellenic Volcanic Arc. The dashed box shows the approximate location of the study area (see ch. 3). The dotted line surrounds the Cycladic Block (Attico-Cycladic Massif). The red stars denote volcanic centres. From Hooft et al. (2017).

The Aegean Sea is a flooded, back-arc basin formed by the slab roll-back and subsequent southward retreat of the Hellenic trench and arc. It has a thinned crust of 20-30 km as shown by gravity and seismic data (Tirel et al., 2004). The current-arc volcanoes (Methana, Milos, Santorini, Nysiros, Kos and others) lie in the pull-apart basin along major strike-slip faults (fig. 2.2).

The regional metamorphism occurred in two stages reaching high-grade blueschist and eclogite facies:

1. high pressure, low temperature – directly related to the collision;
2. low pressure, high temperature – caused by thinning of the lithosphere due to the slab roll-back (Jolivet et al., 2013).

Blocks of metamorphic rocks have been exhumed during Miocene and younger extension phases (Piper and Perissoratis, 2003) and now compose most of the islands in the region, e.g. Ios, and Amorgos. Faults defining metamorphic blocks now aligned along NE-SW trends are no longer perpendicular to N-S extension, indicating a regional rotation of strain, most likely during the Miocene (Piper et al., 2007).

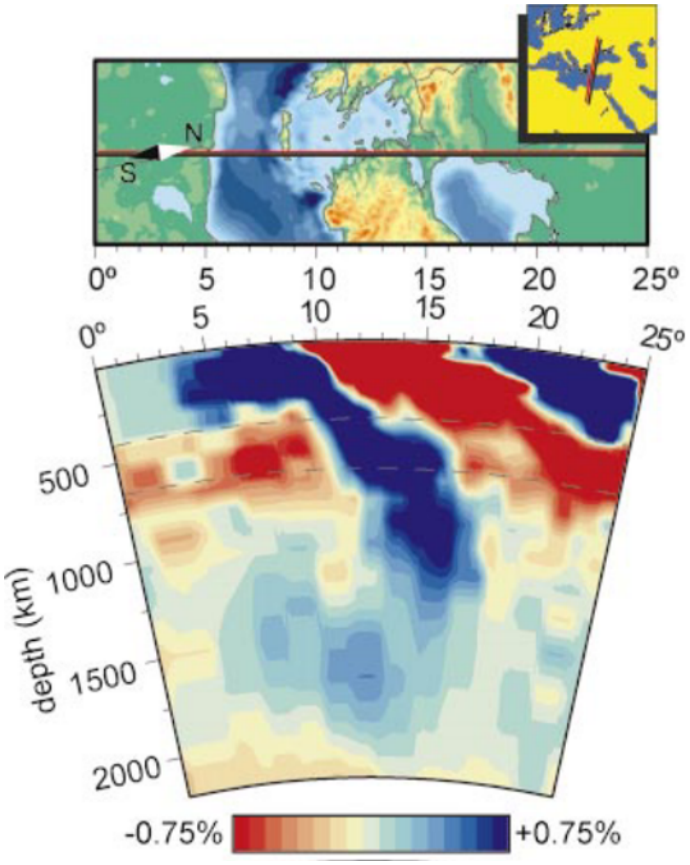


Figure 2.3: The Hellenic subduction zone as imaged by Bijwaard et al. (1998). From van Hinsbergen et al. (2005).

2.1.2 Local scale

The Santorini field lies within a N50°E trending rift zone at a high angle to the volcanic arc (Nomikou et al., 2013). Estimates of the present-day crustal thickness under Santorini vary from 15 km (fig. 4 of Bohnhoff et al., 2001) to 25 km (fig. 5 of Karagianni et al., 2005).

The main structures are two Anhydros and Anaphi basins (grabens) separated by the Santorini-Amorgos Ridge (horst) created by the normal-to-transtensional Santorini-Amorgos Fault Zone (fig. 2.4). This fault zone is suspected to be pathways for magma (Piper et al., 2007). East of Santorini extensive faulting is observed, indicated also by lowered P-wave velocities on trave-time tomograms. An important tectonic feature is the Santorini-Amorgos Basin undergoing right-lateral transtension. It divides the Hellenic Arc into a seismically active East and relatively quiet West (Bohnhoff et al., 2006). The metamorphised sediments outcrop in the SW part of Santorini (Druitt et al., 1999). Because metamorphic lithics are found in the Santorini ignimbrites a fractured metamorphic block most likely underlies the rest of the island too (Druitt et al., 1999). The two main tectono-magmatic lineaments are the Kameni and Kolumbo line (Heath et al., 2019).

2.2 Volcanic activity

In general, two alternating phases of a volcano's evolution can be distinguished – effusive, dome-forming and explosive, sometimes involving a caldera-collapse. The current island arc's volcanism

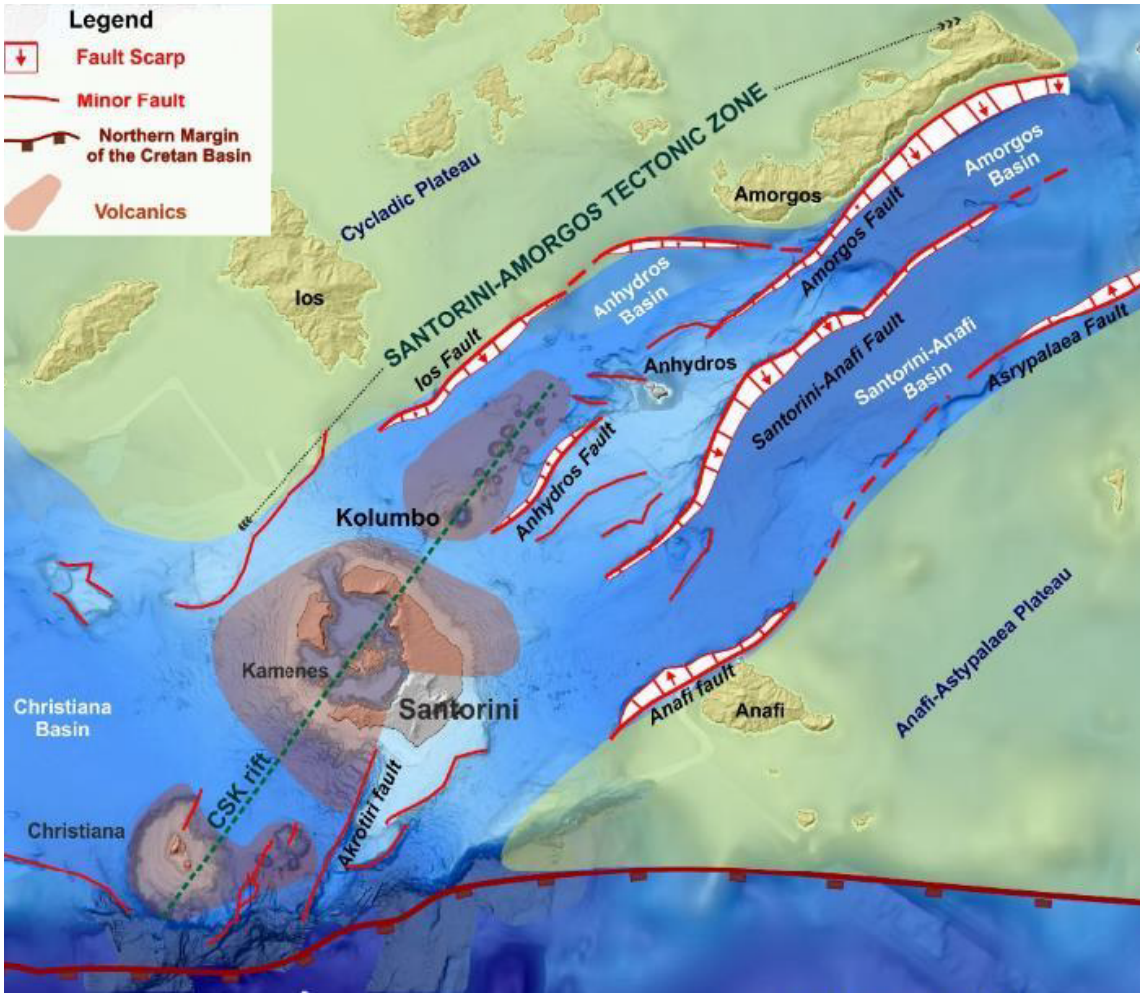


Figure 2.4: Major faults around Santorini. Note that the E side of the caldera is much more fractured than the W side. From Karstens et al. (2020).

875 began 3-4 My ago (Druitt et al., 1999). The two long-term volcanic cycles with clear trends in magma composition can be observed. They both start with mafic-to-intermediate eruptions and terminated with silicic (rhyodacite, or even rhyolite in the case of Kolumbo) eruptions involving a caldera collapse. In total 12 major explosive eruptions of which at least nine were Plinian (Druitt et al., 1999), and at least 3 large lava shields can be distinguished at the Santorini caldera.

880 2.2.1 The 1600 BC eruption

A major caldera-forming eruption happened at Santorini around 1600 BC, called the Late Bronze Age (LBA) eruption. It was a series of large Plinian eruptions with a total volume of 60–70 km^3 (fig. 2.5) dense-rock equivalent of magma (Karátson et al., 2018). It discharged three magma types, two andesitic and one rhyodacitic (Druitt, 2014). The zoning of plagioclase crystals suggests that the upper
885 magma chamber had grown substantially a few decades before the eruption (Druitt et al., 2012).

2.2.2 The 2011-12 unrest

The Santorini volcano is currently in a dome-forming phase (Nomikou et al., 2014b). Since 1570 AD five subaerial eruptions in 14- up to 160-year intervals happened, leading to the growth of the Nea

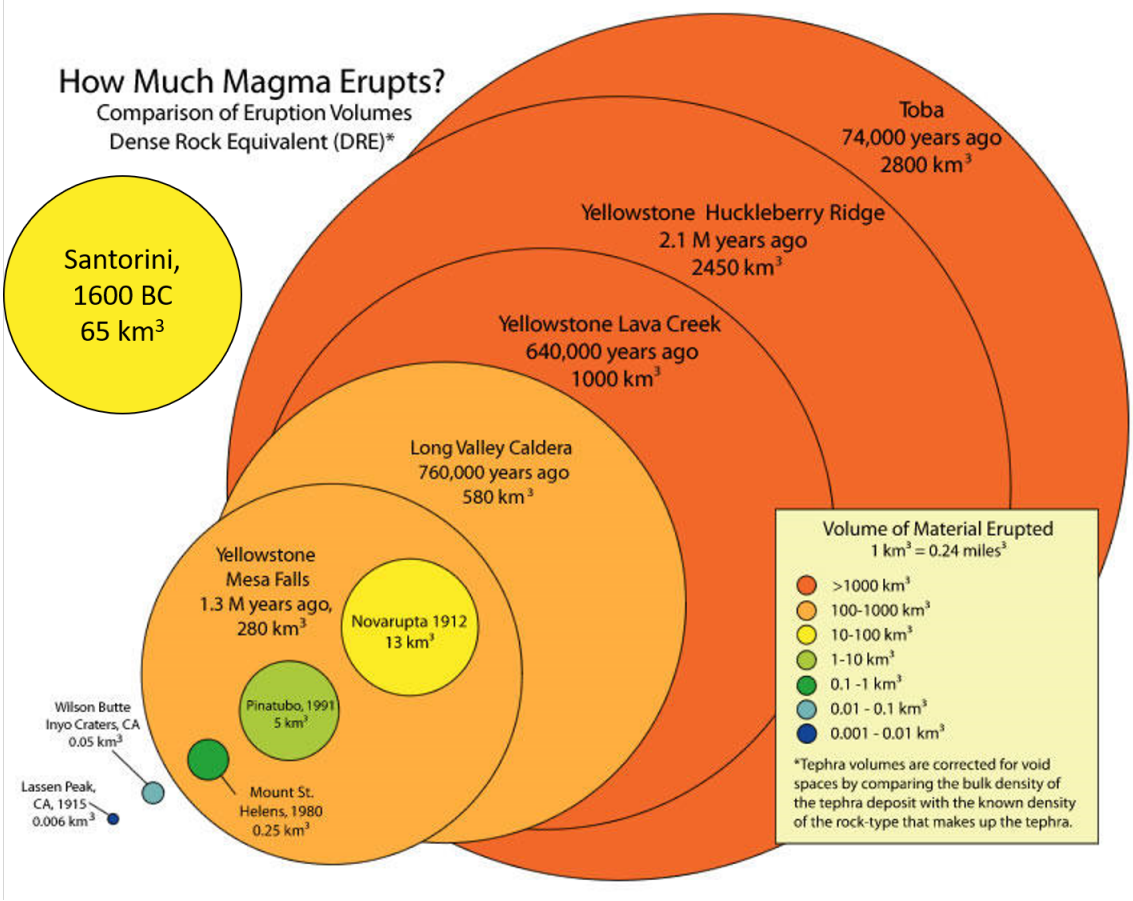


Figure 2.5: Comparison of volume of magma The eruption sizes were drawn by solving for the radius of a sphere using the eruption volumes shown. Modified after https://www.usgs.gov /media /images/ (public domain).

Kameni islands with the last two eruptions occurring in 1939-1941 and 1950 (Druitt et al., 1999). The Santorini volcano was quiescent for more than 60 years (unlike Kolumbo, see ch. 8), until 2011-2012 when the seismic crisis (intense earthquake swarm) accompanied by very high-rate inflation (150-180 mm/yr). The swarm with tens of recorded earthquakes (M1.0-3.2) per day was centreed beneath the Kameni islands and extended in the NE direction along a sub-vertical fault between 1 and 6 km depth (Newman et al., 2012). The total uplift of around 10 cm observed by GPS and InSAR measurements was radially symmetric and the modelling showed that it is consistent with a Mogi source located 3.3-6.3 km beneath the submerged northern section of the caldera, offset from the seismic swarm (Parks et al., 2012; Newman et al., 2012). The prevailing interpretation is that a large eruptible volume of magma has accumulated in a shallow chamber (Parks et al., 2015). The total volume of two pulses of the magma injected from lower to uppermost magma chamber was estimated to be 10-20 mln m^3 which is 10-50% of the volumes of recorded eruptions Parks et al. (2012). This was the only significant intrusion since 1955 (15 years after the last eruption) and is thought to likely be a beginning of a new episode of activity Parks et al. (2012).

2.3 Magmatic system

The multi-modal volcanism of the Santorini volcano indicates a presence of a multi-layer plumbing system (Druitt et al., 1999). Petrologic studies indicate that all four caldera-forming eruptions were fed from a similar depth of about 8 km (Cadoux et al., 2014; Fabbro et al., 2017). Phase equilibria studies suggest a presence of another magma body at ~15 km depth (Cadoux et al., 2014). Combined with a petrological observation that a rapid injection of 15% of erupted volume magma less than 100 years happened before the LBA eruption, it supports a view that a shallow magma chamber is charged not only by continuous, low-volume influxes but also by episodic, rapid and high-volume ones (Druitt et al., 2012). The 2011–12 event involved the intrusion of two magma pulses at about 4 km depth, resembling the behaviour of the volcano before the Plinian eruption.

Local earthquake tomography does not resolve the seismic structure directly beneath Santorini below 2-3 km depth, but does detect the presence of a low-velocity zone ca. 6 km beneath Kolumbo suggesting its possible connection to Santorini (Dimitriadis et al., 2009, 2010).

2.4 Open questions

Some of the questions remaining to be answered are listed below:

1. What is the (more) precise volume and melt fraction of the shallow reservoir beneath Santorini recharged in 2011-12?
2. Is there a shallow melt reservoir beneath Kolumbo?
3. Is there a connection between Santorini and Kolumbo? If yes, at what depth?
4. Do sills of mafic magma currently pond in the lower crust?
5. What is the nature of the connections between all levels of the crust?

This list is complemented by the more general questions provided in sect. 1.3.

3

PROTEUS experiment

MANY of the grand questions of volcanology are related to the continental arc volcanism. Answering these questions require geometric constraints that can only be provided by high-resolution seismic imaging techniques. Petrology provides barely any geometrical constraints. Depths of reservoirs are known with large uncertainty and lateral variations as well as connectivity are not known at all. In order to illuminate the deep roots of an arc volcano, large-aperture acquisition is needed. The high resolution and large aperture imply a requirement of collecting data using a dense array of seismic stations deployed over a large area. For any land volcano such acquisition would be prohibitively expensive, not to mention that the land seismic datasets are notoriously difficult to analyse. Marine seismic acquisition is favored as it is not affected by the land topography, human activity or restrictions regarding the source of seismic energy to be used.

Santorini is an outstanding natural laboratory for studying continental arc volcanism not only because of its particularly well recognised geology and volcanic history. As a semi-submerged system it lends itself to dense 3d marine-land seismic studies and promises to address the questions of sect. 1.3 and 2.4. Such a dataset was acquired in 2015 in the framework of the PROTEUS[1] experiment discussed below. Complete analysis of the PROTEUS dataset will make Santorini the best studied volcano in the world.

3.1 Acquisition

Acquisition of all the data analysed in this work took place from 19 November 2015 to 15 December 2015 during the scientific cruise on board of the R/V Marcus G. Langseth. The National Science Foundation (NSF) of United States of America covered most (~3 million USD) of the expedition costs with contributions from Imperial College London (partial funding), Aristotle University of Thessaloniki (a military helicopter) and University of Athens (other resources).

The acquired data consisted of actives-source marine and land seismic data, bathymetric data collected using Simrad Kongsberg EM122 12 kHz multibeam echo sounder, reflectivity of the uppermost

[1] Which stands for 'Plumbing Reservoirs Of The Earth Under Santorini'.

Algorithm 3.1 Timeline of the acquisition and pre-processing of the PROTEUS data used in this study.

1. Deployment of 65 land stations on the islands of Thera, Nea Kameni and Palea Kameni.
2. Departure of R/V M. G. Langseth from Piraeus, Greece, 17 Nov 2015 for the study area (1 day).
3. Deployment of 91 ocean bottom seismometers (5 days).
4. Acquiring active-source marine-land seismic as well as bathymetric and bathythermographic data (11 days).
5. Recovery of ocean bottom seismometers recovered (5 days, recovery rate 90/91).
6. Arrival of R/V M. G. Langseth in Heraklion, Greece, 15 Dec 2015.
7. Producing SEGY files of all OBS data.
8. Data archiving in a NSF-supported database and publishing the cruise report.

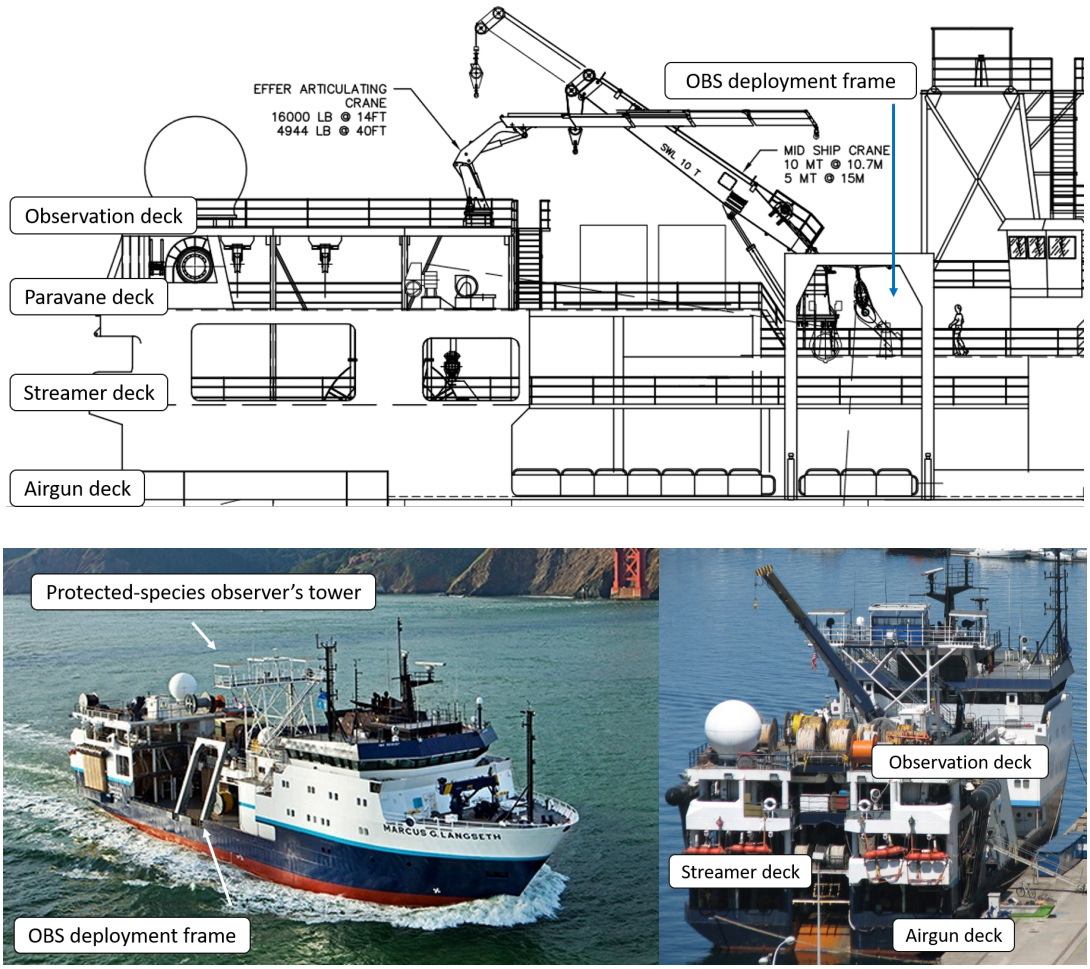


Figure 3.1: The US Research Vessel Marcus G. Langseth of 71.5 m length (overall), 17 m width (beam, moulded) and 5.9 m of maximum draft (LD-EO, 2021). The moveable frame with a winch is used to deploy and collect OBS instruments. Top and bottom-left picture – modified from LD-EO (2021), bottom-right – from Montes (2013).

100 m of sediments using Knudsen 3260 3.5 kHz echosounder, temperature-derived water velocity using expandable bathythermograph (XBT), gravity, magnetic field, GPS, weather, and sea state (Hooft, 2015b; Hooft et al., 2017).

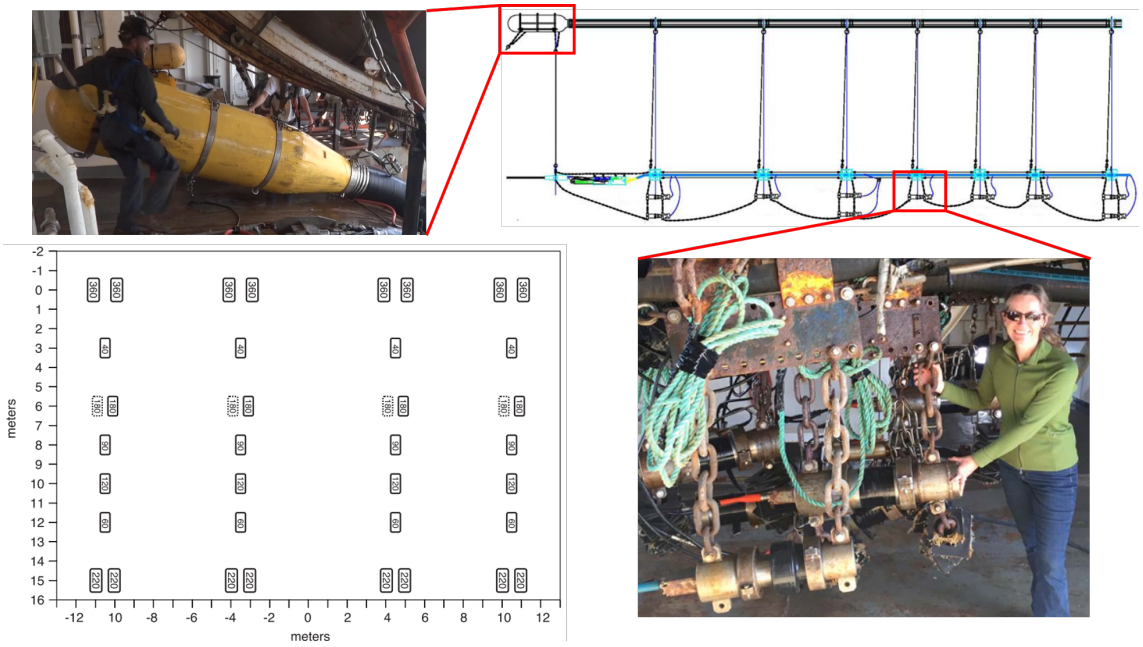


Figure 3.2: Seismic source (airgun array) used in the experiment. Top-right and bottom-left – 1d and 2d view of the airgun array, respectively from Tolstoy et al. (2009). Remainder – size of selected elements of the array compared to people. From Hooft (2015a,b).

3.1.1 Research vessel

The main ingredient of the experimental apparatus was the NSF-owned Research Vessel (R/V) Marcus G. Langseth, arguably the best equipped academic seismic vessel in the world. It is managed by the Lamont-Doherty Earth Observatory with a home port in New York. Its main capabilities (fig. 3.1) include towing streamers and airgun arrays (sect. 3.1.2) and deploying ocean-bottom seismometers.

The crew of 52 people consisted of 3 prinicipal investigators (PIs) – Emilie Hooft, Douglas Toomey (both University of Oregon) and Paraskevi Nomikou (University of Athens), 19 permanent staff members, 16 science technicians from WHOI and SIO (see 3.1.3), and 14 other academics more or less involved in the project, mainly doctoral students from the USA, UK and Greece.

3.1.2 Seismic source

The PROTEUS experiment belongs to a class of active-source seismic experiments. It means that the seismic waves recorded by receivers are generated on purpose instead of being naturally occuring earthquakes.

In the experiment, an airgun source, the most common marine source, was used. An airgun is a large container loaded with compressed air which is released in water to generate a submarine explosion that further propagates through water as a sound wave and penetrates under the seabed sampling the elastic properties of the subsurface before arriving at a receiver.

There are numerous considerations to take into acount when designing an airgun source, including a total energy released, its directivity and oscillations of air bubbles, among other factors (Sheriff and Geldart, 1995). To address these issues, typically an array of airguns is used. The R/V Langseth uses four identical, separately-towed linear arrays (strings). Each strings consists of ten airguns – three double-gun clusters and four single guns. The presence of clusters improve the array tuning by producing a

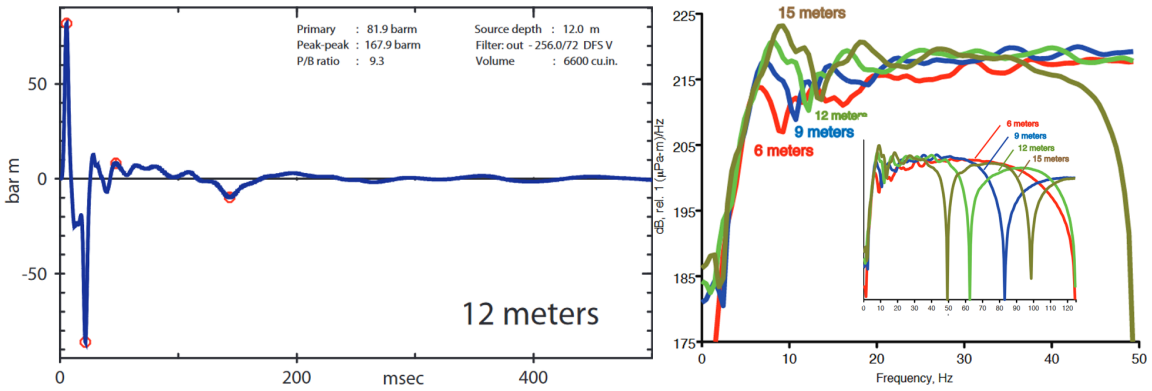


Figure 3.3: Seismic source wavelet of the R/V Langseth array simulated numerically by Lamont-Doherty Earth Observatory shown in time and frequency domain (left and right, respectively). Modified after Hooft (2015a).

larger residual air bubble with slower and weaker reverberations (Tolstoy et al., 2009; Diebold et al., 2010). One of the guns in the central cluster of each string serves as a spare part and is turned off (fig. 3.2).

Each string was towed at a depth of 12 m below the sea surface. All active guns were fired simultaneously giving a 36-element airgun array with 6600 cu. in. of a total volume. A total number of 14,320 shots was generated along 2500 km of track line at regular intervals resulting in an average shot spacing of 144 to 165 m (fig. 3.5). The optimal shot spacing is a trade-off between dense data coverage and overlapping of the shots.

The far-field source time-function for the 12-m towing depth estimated by the Lamont-Doherty Observatory is shown in fig. 3.3 (left). The power spectrum for several towing depths is given in fig. 3.3 (right). A careful spectral analysis of the Lamont source and water waves recorded by the OBSs revealed differences (probably resulting from the local wave propagation effects present only in the field recordings) that prevent it from being used in this study (Papazachos, 2017). Instead, a source function was extracted from the field data as described in sect. 7.3.

3.1.3 Ocean-bottom seismometers

All 91[2] ocean-bottom seismometers (OBSs) used in the experiment (fig. 3.5) belonged to the Ocean Bottom Seismograph Instrument Pool (OBSIP), of which 61 were designed and operated by the Scripps Institution of Oceanography (SIO) and 30 by the Woods Hole Oceanographic Institution. All but one (S116, where S stands for SIO) were recovered from the seafloor using an acoustic signal to cut off their anchors. One had a memory card error (S152) giving 89 stations that could be further analysed. The stations maintained by the WHOI yielded overall better data quality.

Sensor types Each OBS was a short-period instrument recording four-component data – scalar pressure and three components of particle velocity. These components will be denoted throughout as p,v_z,v_y,v_x or HY, VZ, VY, VX or 4,1,2,3 or 0,1,2,3, depending on the context. The pressure was recorded by High Tech HTI-90-U hydrophones. Its amplitude response to hydrostatic-pressure variation was flat between 6-80 Hz and dropped by two-thirds at 3 Hz. The particle velocity was recorded by a three-axis geophone – Sercel L-28-3D or Geospace GS-11D for SIO and WHOI, respectively. Both

[2]The station numbers go from 101 to 193, however stations 111 and 159 were never deployed.

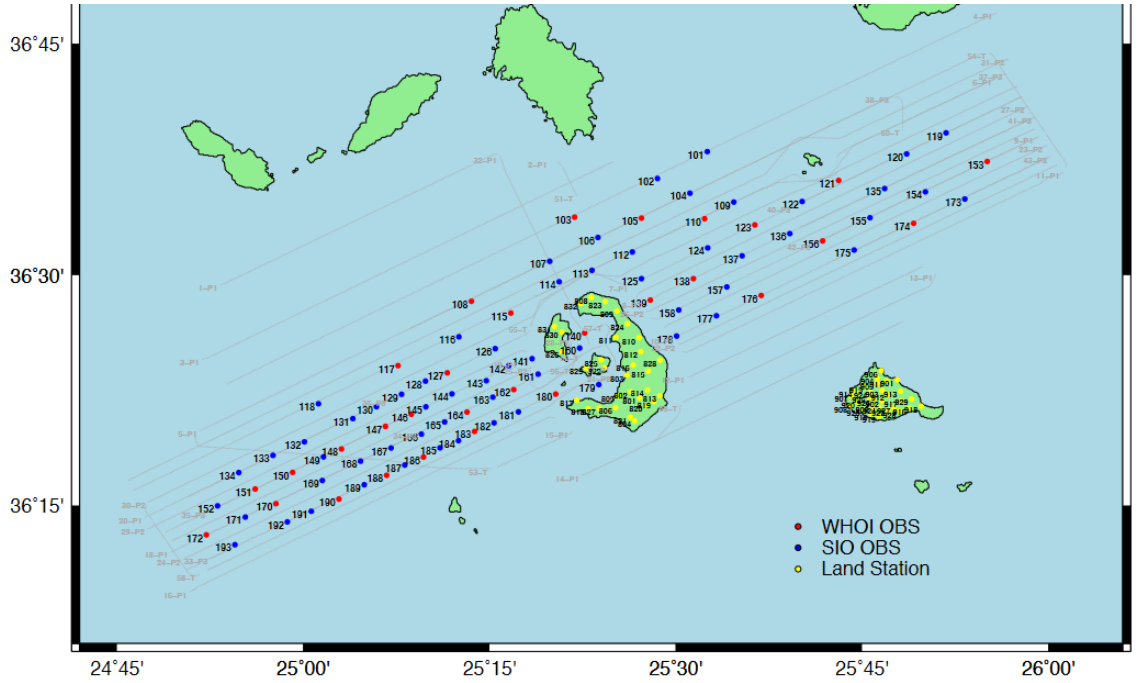


Figure 3.4: Acquisition geomertry of the PROTEUS experiment. Shot lines P3 and P2 are the ones between/at the end of some of the primary (P1) shot lines. Numbering of transit lines (T) starts from 50. From Hooft (2015a).

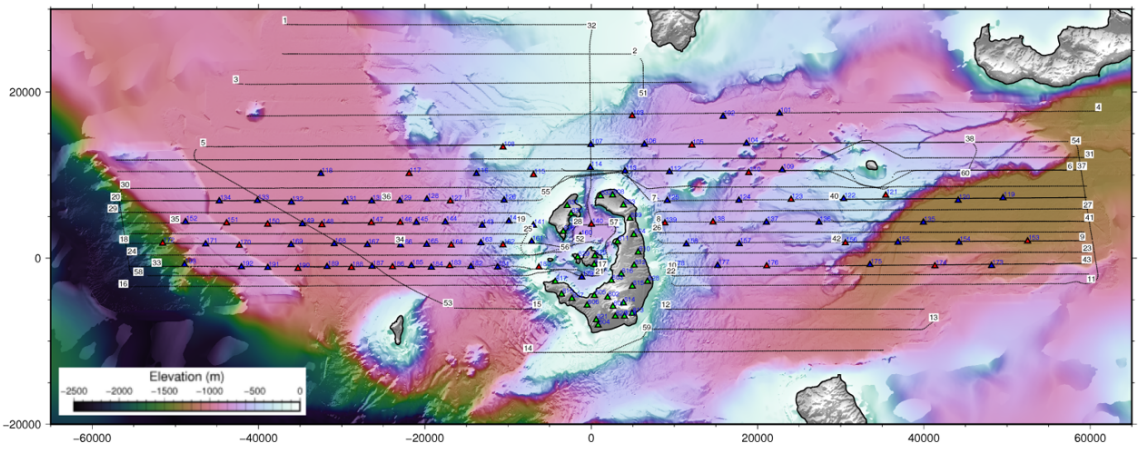


Figure 3.5: Acquisition geomertry in local coordinates overlain on topography/bathymetry map. Red, blue and green triangles denote the WHOI, SIO and land stations respectively. The local coordinate system is based on the Universal Transverse Mercator rotated about 25.5° clockwise from north. The location of centre of the caldera is ca. at $36.4°\mathrm{N}, 25.4°\mathrm{E}$. Modified after M. Paulatto (unpublished).

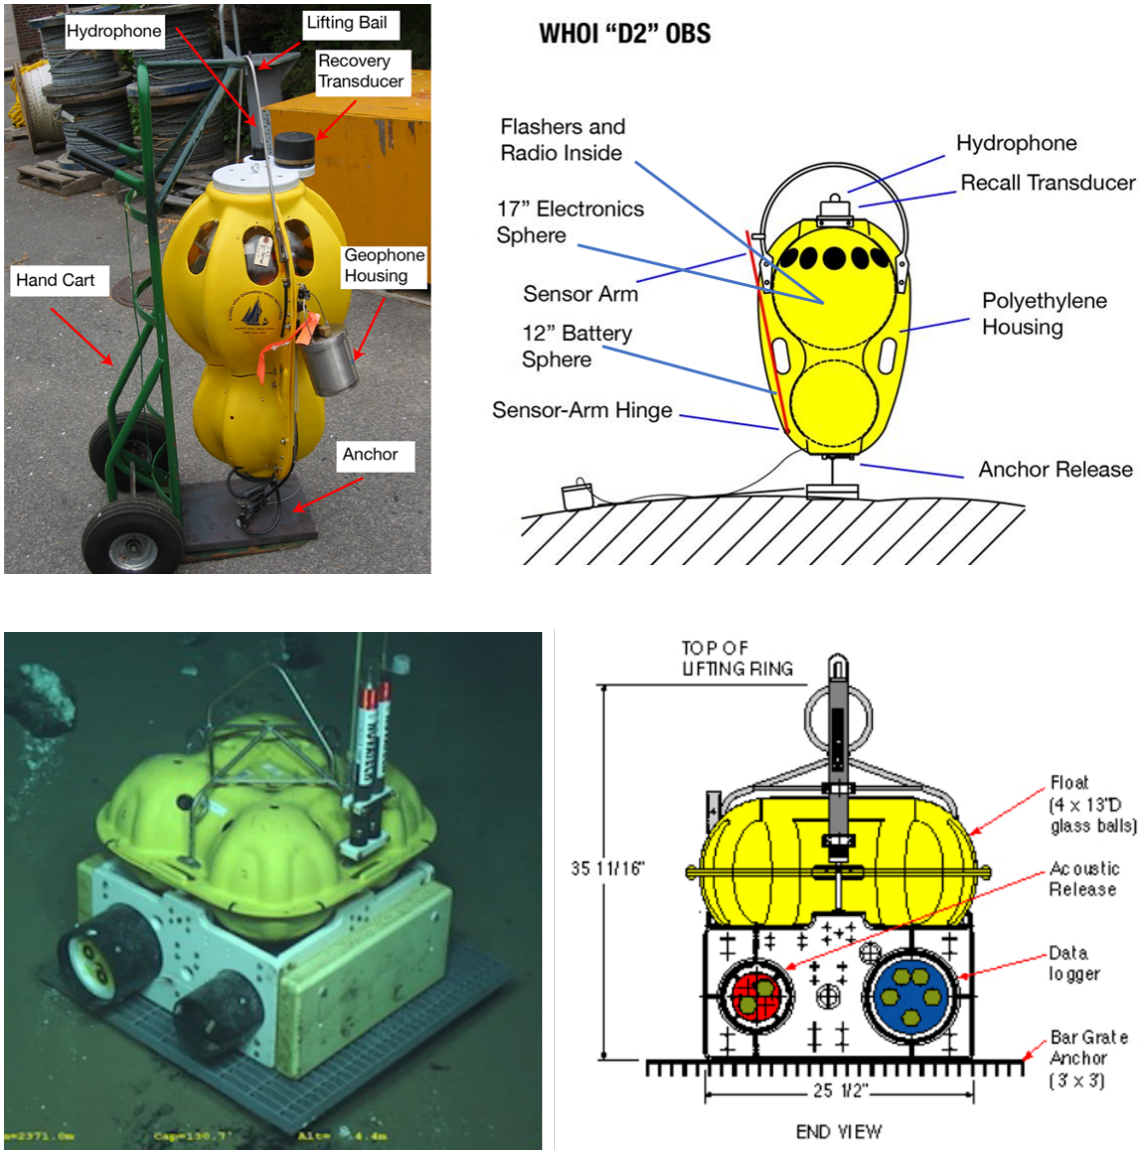


Figure 3.6: Two types of ocean-bottom seismometers deployed in the PROTEUS experiment, from (top) Woods Hole Oceanographic Institution and (bottom) Scripps Institute of Oceanography instrument-pool. Modified after WHOI (2021) and SIO (2021), respectively.

types of geophones were similar, with amplitude response to velocity flat above 4.5 Hz and dropping by about two-thirds at 3 Hz. All the OBSIP data are archived at the IRIS Data Management Center (http://www.iris.edu)).

Coverage The average OBS spacing was around 4 and 5 km to the W and E of the caldera, respectively (fig. 3.4). Even though more stations had unusable data either on pressure or velocity component on the western side of the caldera, the eastern side had overall sparser station coverage.

3.1.4 Land seismometers

Over 70 three-component seismometers provided by the Geophysical Instrument Pool at the German Research Center for Geosciences and the Aristotle University of Thessaloniki (AUTH). They were deployed on the islands of Thera (several teams coordinated by prof. M. Warner), Nea Kameni, Palea Kameni, Anafi (team coordinated by prof. J. Morgan), Anhydros, and Christiana (last two islets were accessed by a helicopter). Only 32 stations deployed on Thera and Kameni islands were analysed in this study. Other islands have even better-quality data (Anafi) but are further away from the target and will be used in the future. Twelve stations of the network were pre-existing, permanent stations of the Institute for the Study and Monitoring of the Santorini Volcano (ISMOSAV). Seven of them were older analog stations not suitable for FWI.

3.2 Observed data

3.2.1 Pre-processing

Pre-processing of the field data (such as correcting for a linear time drift of some of the instruments clocks) was done by science technicians of the Langseth vessel and other members of the PROTEUS team, primarily E. Hooft. Some steps that may be relevant for further processing of the input (ch. 7) are described below.

Data format The data traces aligned at shot origin time are stored in a binary, little-endian SEGY format. Each receiver gather is stored in a separate file of a name:

MGL1521_<subnetwork><station_no>_<channel>.sgy

where subnetwork may be S for SIO and W for WHOI, station number is a three-digit integer, and the channel number is one of the following: 1=vertical, 2&3=horizontals, 4=hydrophone. The naming station for land stations is similar. Important SEGY header words include: the station number (`tracf`), (shot-)line number (`ep`) and shot number (`fldr`).

Data size The raw data sampled at 200 Hz rate take about 0.5 GB of disk space per channel. The full OBS dataset (all 4 channels) then takes about 200 Gb of disk space. The land dataset counts fewer stations, only 3 channels and was sampled at 100 Hz. The land data subset analysed in this work (which excluded e.g. Anafi and Christiana islands) was about 25 GB in size.

Coordinates The Universal Transverse Mercator[3] coordinate system was used, centred at Nea Kameni island (about $36.4°\mathrm{N}, 25.4°\mathrm{E}$) rotated about 25.5° clockwise from north. The coordinates are

[3] It is *not* a single map projection. Instead, it contains 60 longitudal zones, six-degree-wide each. It allows for unambiguous positioning of points between 80S-84N range of latitude.

stored in SEGY headers with accuracy up to one 100th of a second of arc, giving a numerical accuracy of ~30 cm (M. Paulatto, personal communication). The initial positions of OBSs from GPS were improved by Heath et al. (2019) through minimising water-wave travel-time residuals within 2 km offset or larger, if necessary (Creager and Dorman, 1982).

Station response The correction for instrument responses was *not* performed. Instead, the instrument responses provided by the manufacturers were assumed to be identical (sect. 3.1.3) and effectively flat in the range of interest. Even if the frequency is in the non-flat, rolling-off part of the spectrum, the amplitude-drop tends to affect both signal and noise equally, keeping S/N ratio more or less constant. Relative amplitude differences at different frequencies were expected to play a minor role when inverting frequencies in a progression from lowest to highest. This is because at each step of such progression the band-width of the data are increased only by a small amount (0.5 Hz in this study) and the inversion iterates until convergence (after 20 iterations the fit changes very slowly) before the bandwidth is increased again. In this way, bandwidth of data inverted at each step is narrow and thus less affected by the amplitude roll-off.

3.2.2 Initial QC

Initial quality control of the whole dataset was performed by several members of the PROTEUS team (primarily E. Hooft, C. Papazachos and M. Paulatto). The conclusions of this analysis are summarised below. A more FWI-oriented quality control was carried out by the author for the data subsets presented e.g. in ch. 8.

OBS data Of the 90 recovered OBSs[4] eight have bad or no data in the hydrophone channel (107, 126, 129, 152, 173, 180, 182, 187). Ten OBSs had low SNR for both hydrophone and vertical channels. Hydrophone/vertical amplitude ratio is around 1 and 10 for all WHOI and SIO stations, respectively. Most energy is concentrated in 5-70 Hz range which consistent with the theoretical spectrum of the seismic source (fig. 3.3). There is a lot of high-frequency noise above 18 Hz. Many SIO stations suffer from a 6 Hz ringing noise on all four channels, not correlated to wind or marine traffic. Some WHOI hydrophones, especially in shallow water, exhibit a non-linear response – very severe for offsets $r < 0.5$ km, only slightly better at $r < 2$ km. Early arrivals across caldera are strongly attenuated, see fig. 2 of McVey et al. (2020). For most OBSs data quality rapidly deteriorates for offsets longer than 20 km giving maximum interpretable depth of about 5 km.

Land data As the land data is not the main part of the inversion results presented in this thesis, it will be only mentioned that their quality is varied, with a clear distinction between stations placed on a hard bed rock versus soft volcanic sediments. E.g. Anafi stations have good quality as they were located on a metamorphic basement.

3.3 Travel-time inversion

The travel-time tomography, which provided a starting model for this study, was performed by part of the PROTEUS team affiliated with the University of Oregon. The tomography resulted in three peer-reviewed papers (Hooft et al., 2019; Heath et al., 2019; McVey et al., 2020) and two theses (Heath, 2019; McVey, 2019) published so far. They all utilised the bathymetric map of Hooft et al. (2017).

[4] OBS 116 was not recovered.

3.3.1 Data

Hooft et al. (2019) and Heath et al. (2019) used approximately the same dataset of 200,000 travel-times of crustal refractions (Pg) in 6-65 km offset range picked on either p, v_z or pv_z-stacked traces filtered with a 4th-order causal Butterworth filter of 5-25 Hz bandwidth. Initially, only data with high S/N ratio in the 6-15 km offset-range were inverted (Heath et al., 2019). The predicted arrivals enabled picking noisier arrivals at offsets up to 65 km (Hooft et al., 2019) which were used in the final inversion. Offsets shorter than 6 km were excluded due to data sparsity and large grid spacing (200 m). Picking errors were assessed visually and ranged from 5 msec to 30 msec in agreement with the reciprocity analysis presented in appendix 5.D.

McVey et al. (2020) extended the dataset of Hooft et al. (2019) and Heath et al. (2019) by adding about 30,000 longer-offset travel-times of caldera-crossing crustal (Pg) refractions. Due to high attenuation beneath the caldera (their fig. 2), arrivals were picked on pv_z-stacked data, in a lower-frequency band (3-12 Hz) compared to previous studies, with a visually-estimated median error of 20 ms.

3.3.2 Results

Hooft et al. (2019) presented a 25x25x3 km model containing a cylindrical low-velocity anomaly, 3 ± 0.5 km in diameter, extending from 0.5 to 3 km depth between Kolumbo and Kameni lineaments in the caldera northern basin. It was interpreted as a volume of a collapsed caldera floor with excessive porosity of 4 to 28% filled with hot seawater. The fact that it directly overlies the 2011-12 inflation source (Parks et al., 2015) was concluded to support a link between magma recharge events and subsurface density anomalies.

Heath et al. (2019) further investigated the relationship between tectonic and magmatic activity at Santorini on a larger scale. Their 120x45x4 km perturbation model (relative to the 1d model) shows excellent agreement with reflection seismics (Tsampouraki-Kraounaki and Sakellariou, 2018) in terms fault locations. They re-defined Kolumbo lineament to be a border between seismically fast and slow regions at 3 km depth, subparallel to NE-SW trending faults. They associate it with faulting observed in multichannel profiles north of Kolumbo (Nomikou et al., 2016).

McVey et al. (2020) focused on the same area as Hooft et al. (2019) but extending the dataset allowed them to image the structure of the caldera down to 6 km depth. They detected a caldera low-velocity volume (CLV) directly below the caldera collapse column reported by Hooft et al. (2019) with a branch of linear low-velocity volume (LLV) extending about 15 km to NE terminating close to Kolumbo. They interpreted CLV as a mush zone of at least 4%–13% melt in 3-5 km depth-range with a maximum of 6%–17% (–21% v_P anomaly) at 3.4 km depth (11%–30% based on synthetic tests).

3.4 Imaging targets

The primary objective of the experiment was to image, in an unprecented detail, the magmatic plumbing system of an arc volcano, by means of the travel-time tomography and full-waveform inversion (ch. 4). In particular, to detect reservoirs of melt distributed throughout the crust and contrain the geometry and connections between them.

The roots of a volcano may reach down to the Earth's mantle (Sparks et al., 2019) which is at about 25 km below Santorini. The maximum array aperture (offset between source and receiver) of ~100 km

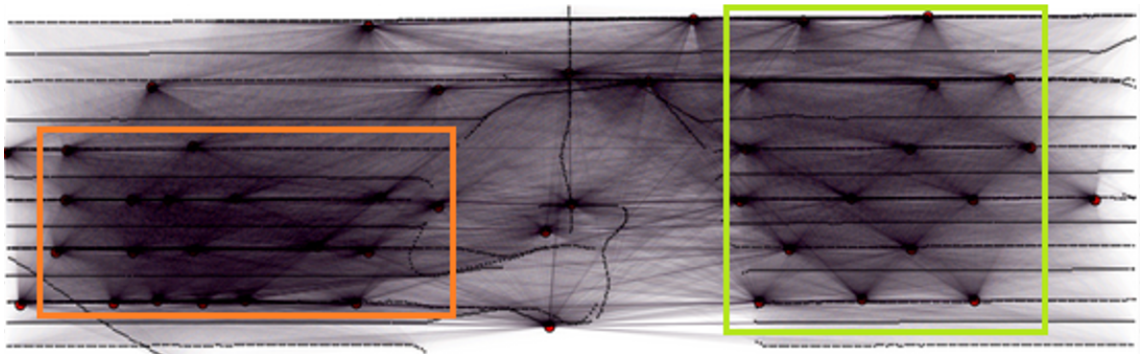


Figure 3.7: OBS coverage highlighted by straight lines connecting sources and receivers for offsets shorter than 20 km. Rectangles denote regions of the densest coverage.

allows for a meaningful interpretation of models not deeper than about 20-30 km which is enough for the local Moho depth, assuming good SNR for all offsets.

However, the data quality rapidly deteriorates for offsets longer than 20 km giving maximum interpretable depth of about 5 km. Other challenges involve data sparsity[5] (fig. 3.7), strong attenuation of waves passing through the target hampering the first-arrival picking, and strong effects of the rough topography that are notoriously difficult to simulate. These challenges constrain the target of high resolution imaging to the upper few kilometers of the subsurface.

[5]Compared to the exploration industry standards.

4

Full-waveform inversion

FULL-waveform inversion (FWI) is a high-resolution technique of imaging Earth's interior (Tarantola, 1984; Woodhouse and Dziewoński, 1984) used to refine the results of the travel-time tomography described in the previous chapter.

In a general sense, FWI is a method of deriving a physical model of an object's interior using finite-length recordings of waves that have propagated through that object. Using a medical analogy, it may be thought of as a powerful, all-in-one, combination of computer-aided tomography and ultrasonography. In some aspects, seismic imaging is much more difficult than a medical problem because of the curvature of wave trajectories and the sparsity of their recordings. The differences between the two applications have not prevented medical and seismic imaging from interacting and exchanging ideas over the years, with the recent introduction of FWI in the medical context (e.g. Guasch et al., 2020).

Mathematically, FWI is a highly non-linear inverse problem which requires a large amount of computational effort and quality control on the way. Currently, the only feasible method of solving it is by linearised optimisation that requires efficient numerical differentiation and hierarchical data selection.

In this chapter, after a brief historical overview, the theory of FWI's *inverse* step is presented from the mathematical, physical and computational standpoint (sect. 4.2, 4.3 and 4.4, respectively). The *forward* step is discussed in part II of the thesis devoted to the full-waveform *modelling*.

4.1 Introduction

Full-waveform inversion (FWI) compares observed and synthetic seismograms recordings to derive a model of elastic parameters of the subsurface. Comparing entire waveforms may seem like the most 'natural' way of recovering subsurface information. However, it is also the most challenging one. It was initially insurmountable computationally and two broad classes of alternative methods were developed – tomography (e.g. Aki et al., 1977) and migration (Claerbout, 1971). In the exploration industry, tomography has been used to prepare a background velocity model for migration. It was never

1984	Theoretical framework for time-domain FWI by Tarantola.
1989	Lab experiments and theoretical framework for frequency-domain FWI by Pratt.
1995	Multi-scale approach proposed by Bunks et al.
2002	2d wide-angle marine field-data study by Shipp and Singh.
2004	2d wide-angle land field-data study by Ravaut et al.
2004	Sparse frequency sampling proposed by Sirgue and Pratt.
2008	First 3d field-data study by Warner et al.
2009	First spectacular 3d field-data result by Sirgue et al.
2009	First 2d elastic field-data study by Brossier et al.
2009	First 3d anisotropic field-data study by Plessix and Perkins.
2011	First 2d visco-acoustic field-data study by Malinowski et al.
2013b	First spectacular 3d anisotropic field-data result by Warner et al.
2013	First 2d visco-elastic field-data study by Prieux et al.

Table 4.1: Time-line of the FWI development in the context of the exploration industry until 2013. Later years have seen a flurry of incremental improvements and new applications of the method, see Wang et al. (2021) for a recent example of high-frequency elastic inversion. Great overview of the history of the method as well as the underlying theory is provided by Virieux and Operto (2009), Bleibinhaus (2014) and Virieux et al. (2014). The word 'spectacular' refers here to improvements in model resolution and fidelity relative to the travel-time tomography.

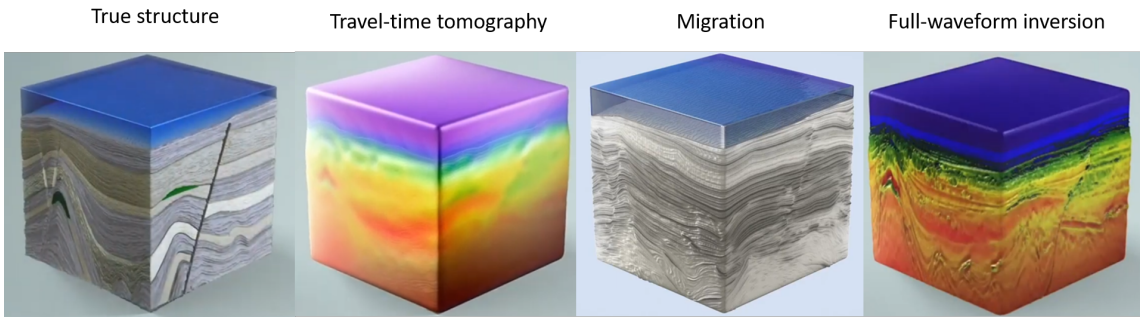


Figure 4.1: Subsurface, travel-time tomography, migration and FWI. Modified after ExxonMobil (2021).

a self-sufficient tool due to its limited resolution. In the academic world in turn, it proved to be an invaluable source of information about the Earth on a scale larger than the oil field (e.g. French et al., 2013).

Eventually, due to the ever-increasing computer power and algorithmic improvements, the two methods could be combined and applied to field-data. FWI was first proposed by Lailly (1983), Berkhout (1984) and Tarantola (1984), who merged the principles of reverse-time migration with linearised iterative optimisation techniques, especially the adjoint method of Lions (1972). It aimed to augment the models of hydrocarbon reservoirs already detected by other methods. Initially, FWI was not regarded as a huge advancement because the short-offset reflection data it was applied to, is fundamentally insensitive to the long-wavelength background features.

The importance of FWI increased when it was understood that FWI can reconstruct the medium at all spatial scales provided the wide-angle, forward-scattered arrivals are present in the data (Mora, 1987, 1988). In this way, FWI combined tomography and migration into one scheme (Mora, 1989). These concepts were first supported experimentally for laboratory-scale cross-hole data by Pratt (1989) and for the field data by Pratt and Worthington (1990) and Pratt and Shipp (1999).

Early reports for wide-angle 2d field data acquired at the Earth's surface were presented by Shipp and

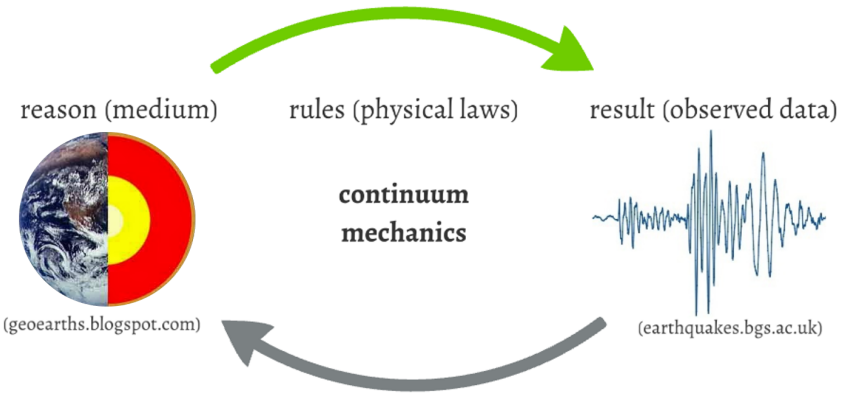


Figure 4.2: Idea of a forward and inverse seismic problem. The forward step (green arrow) is realised by nature when the waves propagate through the Earth. Scientists simulate it numerically to solve the inverse problem (gray arrow).

Singh (2002) and Ravaut et al. (2004), offshore and onshore respectively. The former studied relatively simple geology of flat-lying sedimentary layers in the marine Faeroe–Shetland basin, whereas the latter the geologically complex Southern Apennine thrust belt.

FWI was eventually fully acknowledged by the industry as a game-changing technology after significant improvements of 3d reservoir images reported by Sirgue et al. (2010) and Warner et al. (2013b) among others. Since then it has been primarily used to create accurate background velocity models for migration (Warner et al., 2012) but high-fidelity directly-interpretable models of the physical properties at reservoir depths are becoming feasible with high-frequency (Kalinicheva, 2019) and multi-parameter FWI (Gao et al., 2015a; Owusu et al., 2016).

Meanwhile, it was adopted in the academic community and further developed to face the challenges of regional and global seismology (Tromp et al., 2005; Fichtner et al., 2008; Morgan et al., 2013). Since then it has been applied to image the Earth's mantle (French et al., 2013; French and Romanowicz, 2015), impact craters (Morgan et al., 2011), deep crustal faults (Fichtner et al., 2013), mid-ocean ridges (Morgan et al., 2016), and subduction zones (Kamei et al., 2012; Górszczyk et al., 2017; Gray et al., 2019). There has not yet been an FWI study of an active volcano to date except for the ambient-noise tomography targetting Mount St. Helens (Flinders and Shen, 2017).

4.2 Inverse theory

A major part of the theory underpinning FWI is that of a generic *inverse problem*. An inverse problem is the one of deriving *model* parameters[1] from the experimental *data*. A procedure of solving an inverse problem is called *inversion*[2]. Usually a distinction between the model and data is such that the model is of actual interest but impossible to measure, and the data is of no interest itself, but depends on the model and can be measured directly.

The idiosyncrasy of FWI due to the sparsity and oscillatory character of the seismic data as well as the computational challenges related to typical sizes of Earth's models will be discussed in sect. 4.3 and 4.4, respectively.

[1] Physical quantities that capture the essential properties of the system under study but cannot be measured directly.

[2] In this general sense, inversion is equivalent to *inference*. Many everyday phenomena may be thought of as examples of inversion. Deriving a model of the reality by a human brain based on electromagnetic data (light) is just one of them.

4.2.1 General formulation

We will here be concerned with a model $m(\mathbf{x}, t)$ and data $d(\mathbf{x}, t)$. So far we do not specify anything about m and d other than they are some functions of space-time coordinates. The most general, implicit relation between model and data can be written as (Menke, 2012):

$$F(m, d) = 0. \tag{4.1}$$

This may, for instance, describe a system of partial differential equations. To emphasise – in general, we do not know F explicitly. Such an explicit form of F may not exist in the first place. Instead, F is usually a complicated computer code that takes m as an input and returns d as an output.

The theory presented will be general but one can think of F as a solver of the constant-density acoustic wave equation:

$$\partial_t^2 d = m^2 \partial_i^2 d + s, \tag{4.2}$$

because that is what this theory will be applied to in the subsequent chapters. Here the model is a P-wave velocity and the data is a pressure wavefield. We will not distinguish between the data and wavefield, as the difference is not relevant, here it is effectively assumed that there are wavefield recorders at every model point. This is a convenience allowing to forget about a measurement operator which can be added at any time if needed.

Note that in this case, F is a non-linear function of m:

$$F(m, d) = (m^2 \partial_x^2 - \partial_t^2) d + s. \tag{4.3}$$

To make the equation 4.1 explicit, we would have to be able to invert the differential operator A:

$$A(m) \coloneqq (m^2 \partial_x^2 - \partial_t^2), \tag{4.4}$$

such that $F(m, d) = A(m)d + s = 0$, in order to arrive at:

$$d(m) = -A^{-1}(m)s. \tag{4.5}$$

This method of solving for d given m is called a direct inversion (Virieux and Operto, 2009) and is often performed in the frequency domain. Another, iterative, approach is more feasible for the time-domain wave equation. To expose better the theory, in this section we will forget about specific equations and assume that the relation between the model and the data can be expressed as an explicit dependence of the data on the model:

$$G(m) = d, \tag{4.6}$$

where G will be called a forward function that encodes the physical laws governing the problem. This notation is chosen to be consistent with Menke 2012 which is probably the most popular textbook on this matter. Sometimes, however, using G is superfluous and it is handier to write it as $d(m)$ like in eq. (4.24).

In general, three problems can be associated with eq. (4.6):

- *forward* problem: find d given G and m
- *inverse* problem: find m given G and d
- *unknown-physics* problem: find G given m and d[3].

[3] In fundamental physics, G is often found by pure reasoning ('illumination') rather than analysing experimental data (Einstein was a prominent example). On the contrary, d is often used only as to check if G is indeed correct.

Here we will be concerned with the inverse problem. The two major ways in which one can divide the inverse problems are:

- linear vs. non-linear problems
- continuous vs. discrete problems.

Many important geophysical problems, FWI included, are continuous and non-linear. And most solution methods, like the one used in this study, proceed by converting them into discrete and linear ones.

In general, we cannot write:

$$m = G^{-1}d, \tag{4.7}$$

as G^{-1} simply may not exist, even if the explicit form (4.6) is available for a problem at hand. This is because G needs to be bijective to have an inverse, and this is not always the case. In essence, if two or more models produce the same data, then G is not bijective because the inverse would have to take more than one value, which contradicts the definition of a function[4]. In other words, a general inverse problem is *non-unique*. This is almost always the case for realistic problems where data are noisy. The data from different models are ambiguous up to the noise level. Actually, because of the noise, none of these indistinghuisable solutions may be exact solutions to the problem.

These two issues – lack of an exact solution and non-uniqueness of approximate solutions are typical for real-world problems and contribute to their *ill-posedness* or *ill-conditioning* (in the discrete case). There is another feature of ill-posed problems, namely *instability*. This happens when small changes to models result in large changes in data. This is characteristic of non-linear problems such as (4.2).

Attempts to tackle ill-posed inverse problems gave rise to the vast subject of inverse theory and inverse methods (Tarantola, 2005; Menke, 2012). They all hinge on the fact that the forward problem (unlike the inverse one) has a unique solution, at least in the deterministic, non-quantum world (Tarantola, 2006). One can proceed by comparing the calculated (synthetic) data to the observed data and choose the model that produces data most resembling[5] the observed. Such a process of comparing data calculated from different models is called forward modelling. Unfortunately, this approach is prohibitively computationally expensive for real problems. For this reason, most inverse methods consist of two iteratively repeated steps:

1. forward step – calculating synthetic data from the best-guess model. This is usually the computationally expensive part[6].
2. inverse step – take advantage of the previous step(s) to choose the new best-guess model.

The key difference between specific methods lies in the inverse step. Some give a single solution (e.g. best-fit solution), typically by means of local optimisation. In this case, we talk about parameter estimation (Gubbins, 2004). In some sense, this consists in extracting the model from the data. This approach is criticised by Tarantola (2006) who claims that 'observations cannot produce models, they can only falsify models'. This philosophy characterises the travel-time inversion method from sect. 3.3 which is referred to by its authors as 'adaptive inverse modelling' (Toomey et al., 1994).

Ideally, one will obtain not one but all plausible solutions within the data uncertainty. This is referred to as a (full) solution to an inverse problem by Gubbins (2004) and is contrasted against parameter estimation. All plausible solutions can only be achieved using Monte Carlo methods that are a natural

[4] For linear problems, it corresponds to a non-zero null space.
[5] This resemblance may be quantified by misfit functions and will be discussed for specific problems later.
[6] We will see that the inverse step consists of a forward calculation (typically more than one) as well.

choice for implementing Bayesian inference (Sambridge and Mosegaard, 2002). Unfortunately, Monte Carlo methods have not yet become feasible computationally for FWI in real-world problems although there have been constant progress in computing and the inverse methodology (Fichtner et al., 2019).

For the above reason, we will focus on the first class of methods i.e. the parameter estimation. It will from now on be used interchangeably used with the terms 'inversion' and 'solution to an inverse problem'. We can pose it as an optimisation problem if we introduce a scalar objective function

$$f \coloneqq |d - G(m)|^2 , \tag{4.8}$$

where $|\cdot| : V \mapsto \mathbb{K}$ is a general *norm* that maps elements of a linear space V (of discrete vectors or functions) into a field $\mathbb{K}$ of scalars for which the relation $<$ makes sense and thus enables one to talk about extremum values.

The specific optimisation scheme will strongly depend on the problem at hand. One class of such schemes is discussed in 4.2.3.

4.2.2 Discrete linear problem

Both data and models can be either continuous or discrete. While most data is collected as discrete, e.g. in form of digitised signals, most models of interest are continuous at a scale of interest (continuum approximation of matter), e.g. wave velocity distribution in the subsurface. Hence, most of the geophysical inverse problems consist in recovering continuous models from discrete data.

A linear[7] continuous problem can be defined as:

$$d_i = \int_V G_i(\mathbf{x}) m(\mathbf{x}) \mathrm{d}\mathbf{x}, \tag{4.9}$$

where $G_i = \frac{\partial d_i}{\partial m}$ is a (data) *sensitivity kernel*[8]. In this expression, each data sample is a weighted average of the entire model. The sensitivity serves as a weighting function in this average. A fully continuous problem:

$$d(t) = \int_V G_i(\mathbf{x}, t) m(\mathbf{x}) \mathrm{d}\mathbf{x}, \tag{4.10}$$

is used mostly for theoretical developments treating both model and data spaces as probability distributions in the Bayesian framework (Tarantola, 2005).

One can discretise eq. (4.9) by expressing $m(\mathbf{x})$ as a linear combination $m_j f_j(\mathbf{x})$ of some basis functions $f_j(\mathbf{x})$[9]:

$$d_i = \int_V G_i(\mathbf{x}) m_j f_j(\mathbf{x}) \mathrm{d}\mathbf{x} = \int_V G_i(\mathbf{x}) f_j(\mathbf{x}) \mathrm{d}\mathbf{x} \cdot m_j = G_{ij} m_j, \tag{4.11}$$

where

$$G_{ij} \coloneqq \int_V G_i(\mathbf{x}) f_j(\mathbf{x}) \mathrm{d}\mathbf{x}, \tag{4.12}$$

is a matrix form of a linear transformation $\mathbf{G} : M \to D$ which maps a vector space M of dimension $p \coloneqq \dim M$ into another vector space D of dimension $q \coloneqq \dim D$ according to the equation:

$$G_{ij} m_j = d_i, \tag{4.13}$$

[7] Non-linear problems can be linearised which is discussed in sect. 4.2.3.
[8] Often denoted by $K_i(x)$.
[9] Dirac delta functions may be used to obtain point-wise discretisation at grid nodes and Heaviside step functions to get grid cells.

or equivalently in the vector notation:

$$\mathbf{Gm} = \mathbf{d}. \tag{4.14}$$

In general, there are three possibilities: $p > q$, $p = q$ and $p < q$ called an under-, equi- and over-determined case, respectively. The number of linearly-independent columns equals the number of linearly-independent rows and is called a *rank* of a matrix. In a general case, it follows that:

$$\mathrm{rank}(\mathbf{G}) \leq \min(p, q). \tag{4.15}$$

The matrix G is called *full-rank* if $\mathrm{rank}(G) = \min(p, q)$. Otherwise, it is called *rank-deficient*.

For the over-determined case, there is no solution and some assumptions need to be made about the data errors. The canonical problem is a linear problem with Gaussian statistics (Menke, 2012). This leads to the least-squares inverse obtained by differentiating the squared L^2 norm of the residual $\delta d \coloneqq d_i - G_{ij}m_j$ with respect to the model vector:

$$\partial_j (d_i - G_{ij}m_j)^2 = 0, \tag{4.16}$$

$$2G_{ik}(d_i - G_{ij}m_j) = 0, \tag{4.17}$$

$$G_{ik}G_{ij}m_j = G_{ik}d_i, \tag{4.18}$$

which in the vector notation can be expressed as:

$$\mathbf{G}^T\mathbf{Gm} = \mathbf{G}^T\mathbf{d}. \tag{4.19}$$

For over-determined problems $\mathbf{F} \coloneqq \mathbf{G}^T\mathbf{G}$ is symmetric and positive semi-definite, i.e. all eigenvalues are non-negative. The symmetry is obvious if one writes it down as $F_{kj} \coloneqq G_{ik}G_{ij}$ because its transpose $F_{jk} = G_{ij}G_{ik} = G_{ik}G_{ij} = F_{kj}$. If $\mathbf{F}$ happens to be positive definite (all eigenvalues positive), it has an inverse and one can write:

$$\mathbf{m} = (\mathbf{G}^T\mathbf{G})^{-1}\mathbf{G}^T\mathbf{d}. \tag{4.20}$$

This is an important, least-squares solution to an over-determined linear inverse problem.

If $p < q$ the problem is under-determined. Its defining property is a non-trivial null-space. Note that the null-space is a property of the operator G, not m or d. There infinitely many solutions because adding a model from a null-space is also a solution. The preferred solution is a subjective choice. Usually, it is required to be optimal in terms of smoothness (e.g. having minimum norm) which is implemented by adding (Tikhonov) regularisation/damping terms in the objective function:

$$f = |\delta\mathbf{d}|^2 + \mu_1|\mathbf{L}_1\mathbf{m}|^2 + \mu_2|\mathbf{L}_2\mathbf{m}|^2 + ..., \tag{4.21}$$

where μ_i are weights $\mathbf{L}_i$ are finite-difference operators, for example:

$$\mathbf{L}_1 = \begin{pmatrix} 1 & -1 & 0 & \cdots & 0 \\ 0 & 1 & -1 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}. \tag{4.22}$$

Alternatively, the first term can be treated as a constraint and the latter term(s) as the actual objective of minimisation. This leads to *constrained optimisation* schemes and the method of Lagrange multipliers. If $\mathbf{G}$ in $\delta\mathbf{d} = \mathbf{d} - \mathbf{G}(\mathbf{m})$ is a PDE, e.g. a wave equation in the case of FWI, it is called a PDE-constrained optimisation.

Real problems like FWI are unlikely to be purely over- or under-determined, not to mention equi-determined. The matrix $\mathbf{G}$ will not be full rank and some model parameters will be over-, some equi- and some other under-determined. The easiest example of a mixed determined system is:

$$\begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \tag{4.23}$$

The parameter x is over-determined ($x = 0$ and $x = 1$) and has no solution whereas y is under-determined – there are ∞-many solutions for it. In theory, the most elegant solution to such problems are offered by the Moore-Penrose generalised inverse that can be calculated using singular value decomposition (Menke, 2012). However generalised inverse can only be afforded for small problems. For medium-size problems numerical-algebraic techniques taking advantage of, hopefully, the special structure of $\mathbf{G}$ matrix (sparsity, banding etc.) must be employed (Press et al., 1990). For large problems such as FWI even this is not enough and the approximation is the only way to go. This latter approach is described in the next section. In sect. 4.4.4 it will be shown that explicitly calculating, not to mention inverting, the matrix $\mathbf{G}$ (which for non-linear problems will turn out to be the Frechet derivative of the data with respect to the model) can actually be entirely avoided.

4.2.3 Non-linear problem

If $\mathbf{d}$ is a non-linear function of $\mathbf{m}$, the linear methods presented in the previous sections do not immediately work. There are two ways to proceed – Monte Carlo and linearisation. As mentioned before, currently the former approach is not feasible for real-world FWI problems. Linearisation of the data can be performed based on the Taylor expansion:

$$\mathbf{d}(\mathbf{m}_0 + \delta\mathbf{m}) = \mathbf{d}(\mathbf{m}_0) + \left[\frac{\partial\mathbf{d}}{\partial\mathbf{m}}\right]_{\mathbf{m}_0} \delta\mathbf{m} + \mathcal{O}(\delta\mathbf{m}^2), \tag{4.24}$$

where $\mathbf{d}$ was expanded around the current model $\mathbf{m}_0$ taking as a variable the true model $\boldsymbol{m}_{\text{true}} = \mathbf{m}_0 + \delta\mathbf{m}$ to be able to evaluate the Jacobian matrix $\mathbf{J} = \frac{\partial\mathbf{d}}{\partial\mathbf{m}}$ at the known $\mathbf{m}_0$. The linearisation truncates the higher-order terms to give:

$$\mathbf{d}(\mathbf{m}_0 + \delta\mathbf{m}) \approx \mathbf{d}(\mathbf{m}_0) + \left[\frac{\partial\mathbf{d}}{\partial\mathbf{m}}\right]_{\mathbf{m}_0} \delta\mathbf{m}, \tag{4.25}$$

and from now on we will replace $\approx$ with $=$ sign, denote the Jacobian matrix $\left[\frac{\partial\mathbf{d}}{\partial\mathbf{m}}\right]_{\mathbf{m}_0}$ by $\mathbf{J}$ and the synthetic data generated from the current model as $\boldsymbol{d}_{\text{syn}} \equiv \mathbf{d}(\mathbf{m}_0)$. Noticing that the data generated from the true model are (up to the errors) the data observed in the experiment $\mathbf{d}_{\text{obs}} \equiv \mathbf{d}(\mathbf{m}_{\text{true}})$ and defining a data residual:

$$\delta\mathbf{d} = \mathbf{d}_{\text{syn}} - \mathbf{d}_{\text{obs}}, \tag{4.26}$$

one can eventually write eq. (4.25) as:

$$\delta\mathbf{d} = \mathbf{J}\delta\mathbf{m}. \tag{4.27}$$

The idea of a linearised optimisation of the non-linear problem is to solve eq. (4.27) for $\delta\mathbf{m}$ to update the model:

$$\mathbf{m}_1 = \mathbf{m}_0 + \delta\mathbf{m}, \tag{4.28}$$

and repeat the whole procedure until $||\delta\mathbf{d}||$ is of the order of data uncertainty. The eq. (4.27) describes a linear system and the same methods of solving it can be used as discussed for the eq. (4.14). In particular, one may find the least-squares inverse like in eq. (4.20):

$$\delta\mathbf{m} = (\mathbf{J}^T\mathbf{J})^{-1}\mathbf{J}^T\delta\mathbf{d}, \tag{4.29}$$

provided $\mathbf{J}^T\mathbf{J}$ is not singular. This simple Taylor-series-based derivation has lead to an important equation (4.29) that is a foundation of FWI and least-squares RTM (Nemeth et al., 1999). A more in-depth interpretation of the terms in eq. (4.29) will become clear if eq. (4.29) is derived in a slightly different way. To this end, instead of expanding the data, let us start with a generic objective function:

$$f(\mathbf{m}_0 + \delta\mathbf{m}) \approx f(\mathbf{m}_0) + \left[\frac{\partial f}{\partial \mathbf{m}}\right]_{\mathbf{m}_0} \delta\mathbf{m} + \frac{1}{2}\left[\frac{\partial^2 f}{\partial \mathbf{m}^2}\right]_{\mathbf{m}_0} \delta\mathbf{m}^2 + \mathcal{O}(\delta\mathbf{m}^3). \tag{4.30}$$

So far we do not assume anything about f other than that it is of the form (4.8) and so the true model $\mathbf{m}_{\text{true}} = \mathbf{m}_0 + \delta\mathbf{m}$ minimises f. The gradient of f is hence zero for the true model:

$$\frac{\partial f}{\partial \mathbf{m}}(\mathbf{m}_{\text{true}}) = 0. \tag{4.31}$$

To find $\mathbf{m}_{\text{true}}$, let us differentiate eq. (4.30) and compare it to zero:

$$0 = \frac{\partial f(\mathbf{m}_0 + \delta\mathbf{m})}{\partial m} = 0 + \left[\frac{\partial f}{\partial \mathbf{m}}\right]_{\mathbf{m}_0} + \left[\frac{\partial^2 f}{\partial \mathbf{m}^2}\right]_{\mathbf{m}_0} \delta\mathbf{m} + \mathcal{O}(\delta\mathbf{m}^2), \tag{4.32}$$

where we exploited the fact that all terms evaluated at $\mathbf{m}_0$ are constant and do not survive the differentiation with respect to $\mathbf{m}$. Denoting $\nabla f \equiv \frac{\partial f}{\partial \mathbf{m}}$ and $\mathbf{H} = \frac{\partial^2 f}{\partial \mathbf{m}^2}$:

$$[\mathbf{H}]_{\mathbf{m}_0}\delta\mathbf{m} = -[\nabla f]_{\mathbf{m}_0}. \tag{4.33}$$

From now on evaluation at $\mathbf{m}_0$ will be implicit, for the sake of brevity. The eq. (4.33) may be solved for $\delta\mathbf{m}$ iteratively leading to so-called *truncated-Netwon* method (Nash, 1985; Métivier et al., 2013). Here a direct solution will be analysed. It leads to:

$$\delta\mathbf{m} = -\mathbf{H}^{-1}\nabla f. \tag{4.34}$$

This is a general expression independent of any particular definition of the objective function. It seems different from eq. (4.29) but it is not, as will turn out. To see this, let us take the same objective function as we implicitly used in the eq. (4.16) when calculating (4.29), that is:

$$f = \frac{1}{2}||\delta\mathbf{d}||_2^2 \equiv \frac{1}{2}\delta\mathbf{d}^T\delta\mathbf{d} \equiv \frac{1}{2}\delta\mathbf{d}\cdot\delta\mathbf{d} \equiv \frac{1}{2}d_i d_i.$$

The presence of the squared norm justifies truncating the series (4.30) after second-order unlike the (4.25). The gradient of this L^2 norm functional gives:

$$\frac{\partial f}{\partial m_j} = \frac{\partial}{\partial m_j}\frac{1}{2}(\delta d_i \delta d_i) = (\delta d_{i,j})\delta d_i = d_{i,j}\delta d_i, \tag{4.35}$$

where the notation $a_{i,j} \equiv \partial a_i/\partial m_j$ was used. In the second-to-last equality:

$$\frac{\partial(\delta\mathbf{d})}{\partial\mathbf{m}} = \frac{\partial\mathbf{d}}{\partial\mathbf{m}}, \tag{4.36}$$

we implicitly treated $\mathbf{d}_{\text{obs}}$ as a constant. Rewriting the left- and right-hand side of (4.35) in the vector notation gives:

$$\nabla f = \mathbf{J}^T\delta\mathbf{d}. \tag{4.37}$$

This equation reveals that the last term of the eq. (4.29) is a gradient of a particular, L^2 norm objective function. Remaining to interpret in the eq. (4.29) is the $(\mathbf{J}^T\mathbf{J})^{-1}$ term. Let us differentiate the gradient to get the Hessian:

$$H_{jk} = \frac{\partial}{\partial m_k}\left(\frac{\partial f}{\partial m_j}\right) = \frac{\partial}{\partial m_k}(J_{ij}\delta d_i) = J_{ij}\delta d_{i,k} + J_{ij,k}\delta d_i, \tag{4.38}$$

as $\delta d_{i,k} = J_{ik}$, see the remark under the eq. (4.36), we can rewrite it as:

$$\mathbf{H} = \mathbf{J}^T\mathbf{J} + \frac{\partial \mathbf{J}}{\partial \mathbf{m}}\delta d, \tag{4.39}$$

because $A_{ij}B_{ik}$ corresponds to A^TB in the vector notation. The third-rank tensor $J_{ij,k}$ expresses the second derivative of $\mathbf{d}$ with respect to $\mathbf{m}$:

$$J_{ij,k} \equiv \frac{\partial \mathbf{J}}{\partial \mathbf{m}} = \frac{\partial^2 \mathbf{d}}{\partial \mathbf{m}^2} \tag{4.40}$$

By definition, this is zero for linear problems. For non-linear ones, it is often neglected leading to the Gauss-Newton (GN) approximation of the Hessian (Nocedal and Wright, 2006):

$$\mathbf{H} \approx \mathbf{J}^T\mathbf{J} \equiv \mathbf{H}_{\mathrm{GN}}. \tag{4.41}$$

It now should be clear that the least-squares inverse of the Jacobian (4.29) is equivalent to the Gauss-Newton approximation:

$$\delta\mathbf{m} = (\mathbf{H}_{\mathrm{GN}})^{-1}\mathbf{J}^T\delta\mathbf{d}. \tag{4.42}$$

If the exact form of Hessian $\mathbf{H}$ is used, the scheme is called the Newton method (Nocedal and Wright, 2006) and has a faster convergence rate than GN. Other sensible approximations are discussed in sect. 4.4.5

A more basic approximation is the steepest descent method which approximates $\mathbf{H}^{-1}$ with a scalar α called *step-length* to get:

$$\delta\mathbf{m} = -\alpha\nabla f \tag{4.43}$$

The step-length scales the update in the gradient direction so that the data residual is minimised. There are various ways to evaluate it. The method used in this work requires only one extra forward modelling. Its starts by perturbing the model to get a new residual:

$$\delta\mathbf{d}_1 \coloneqq \delta\mathbf{d}(\mathbf{m}_1) = \delta\mathbf{d}(\mathbf{m}_0 + \delta\mathbf{m}). \tag{4.44}$$

Assuming the linearity between data and models, scaling this perturbation gives:

$$\delta\mathbf{d}_\alpha \coloneqq \delta\mathbf{d}(\mathbf{m}_0 + \alpha\delta\mathbf{m}) = \delta\mathbf{d}(\mathbf{m}_0) + \alpha\delta\mathbf{d}(\delta\mathbf{m}) \tag{4.45}$$

The $\delta\mathbf{d}(\delta\mathbf{m})$ can be evaluated using the linearity of $\mathbf{d}(\mathbf{m})$ once more:

$$\delta\mathbf{d}(\delta\mathbf{m}) = \delta\mathbf{d}(\mathbf{m}_1 - \mathbf{m}_0) = \delta\mathbf{d}(\mathbf{m}_1) - \delta\mathbf{d}(\mathbf{m}_0) \equiv \delta\mathbf{d}_1 - \delta\mathbf{d}_0 \tag{4.46}$$

Combining above two gives:

$$\delta\mathbf{d}_\alpha = \delta\mathbf{d}_0 + \alpha(\delta\mathbf{d}_1 - \delta\mathbf{d}_0) \tag{4.47}$$

The data residual from the model perturbed in the (anti-)gradient direction, scaled with α can be minimised by setting:

$$\frac{\partial(|\delta\mathbf{d}_\alpha|^2)}{\partial\alpha} = 0 \tag{4.48}$$

	time	**frequency**
space	Tarantola (1984); Miller et al. (1987)	Woodward (1992)
wavenumber	—	Wu and Toksöz (1987)

Table 4.2: Formulations of seismic tomography in different Fourier domains of space and time variables.

Method	Imaging operator	Parameter	Iterative	Non-linear	Reference
RTM	$\mathbf{J}^T\delta\mathbf{d}$	reflectivity	no	no	Baysal et al. (1983)
LSQ-RTM	$(\mathbf{J}^T\mathbf{J})^{-1}\mathbf{J}^T\delta\mathbf{d}$	reflectivity	yes	no	Nemeth et al. (1999)
FWI	$\mathbf{J}^T\delta\mathbf{d}$	velocity	yes	yes	Warner et al. (2013b)
LSQ-FWI	$(\mathbf{J}^T\mathbf{J})^{-1}\mathbf{J}^T\delta\mathbf{d}$	velocity	yes	yes	Pratt et al. (1998)

Table 4.3: Comparison of different inverse-scattering (imaging) methods. In linear methods the background velocity model, and thus the wave paths, remain constant.

which eventually gives the optimal step-length:

$$\alpha = \frac{\delta\mathbf{d}_0^T(\delta\mathbf{d}_0 - \delta\mathbf{d}_1)}{|\delta\mathbf{d}_0 - \delta\mathbf{d}_1|^2}. \tag{4.49}$$

By definition, the local optimisation scheme described above will always converge to the nearest minimum of the objective function. It is a separate task to design a starting model that is closer to the global rather than a local minimum.

4.3 Scattering theory

As seen in the previous section, FWI can be posed mathematically as a linearised local optimisation problem. From a physical point of view, FWI may be seen as an inverse-scattering problem. This perspective may be useful in establishing links between the optimisation theory and wave phenomena.

4.3.1 Born approximation

In the scattering theory (Hudson and Heritage, 1981) the Taylor expansion (4.24) is referred to as the *Born series*. The analysis of this series was originally performed in the field of quantum mechanics (see e.g. Shankar, 1994).

Truncating the Born series after the first-order term is called the *first Born approximation*, after the second-order – second Born approximation and so on. Here only the first approximation is analysed and from now on will be referred to simply as the Born approximation. This approximation was used in the previous section to derive equations 4.27, 4.33 and 4.49. The Born approximation assumes *single* or *weak* scattering (Hudson and Heritage, 1981).

4.3.2 Inverse scattering

Inverse scattering is a task of deriving a model from the data and can be formulated in various domains (tab. 4.2). A great overview of fundamental theorems of the inverse scattering theory (e.g. Fourier diffraction theorem) is given by Kak and Slaney (2001).

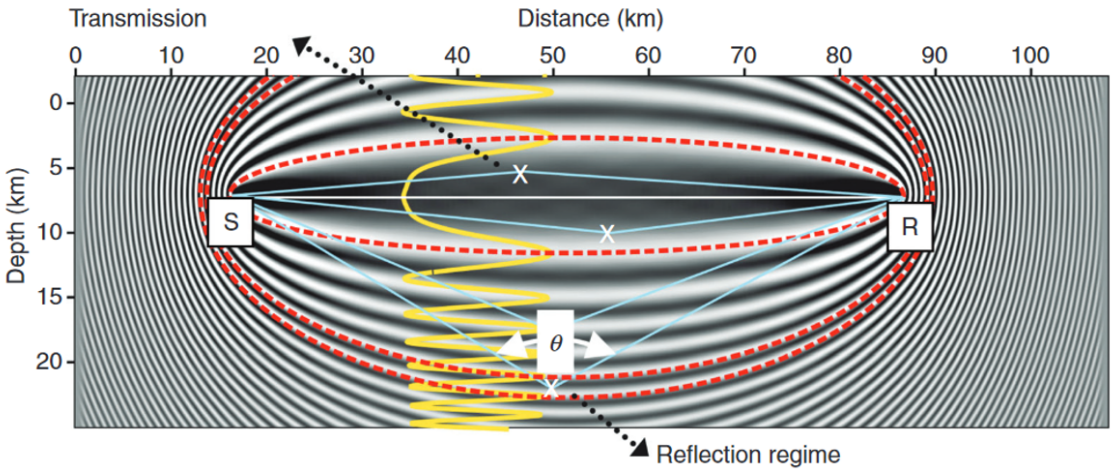


Figure 4.3: Monochromatic wavepath. From Virieux and Operto (2009).

Sensitivity kernel The monochromatic sensitivity kernel of a *finite* frequency, for a single source-receiver pair, is an oscillatory pattern that consists of Fresnel zones (fig. 4.3). It is often referred to as a monochromatic *wavepath* (Woodward, 1992). A *ray(path)* is simply a monochromatic wavepath of *infinitely high* frequency. The gradient is simply a *multi*-frequency wavepath (unless a monochromatic source is used). In FWI, usually multiple source-receiver pairs, resulting in the pattern focused on heterogeneities missing from the model. There are different waves of calculating sensitivity-like patterns (Tromp et al., 2005). For instance, one may backpropagate the observed data or the data residual, the latter case is shown in sect. 5.4.4. In general, if the duration of the source is long enough, sensitivity across the entire model domain is obtained.

Imaging operator The function that converts the data into model is called the *imaging operator*. Conventional migration uses the gradient, and does not iterate. FWI as used here, also uses the gradient (preconditioned with step-length) but is iterative (tab. 4.3). So called least-squares migration (Nemeth et al., 1999) uses Gauss-Newton Hessian to precondition the gradient. The difference between FWI is that it does not update the background model. In travel-time tomography, the data residual comes from any point along the ray connecting a given source and receiver – one shot samples the medium only along the ray. In FWI the data residual comes from any point in the whole model (one shot samples the whole model). Hence, the data are allowed to be spatially aliased (to some extent), and the reconstructed model may be meaningful outside source-receiver planes. The gradient is not a perfect imaging operator, a more accurate one involves the Hessian (Pratt et al., 1998), giving less acquisition footprint and fewer aliasing artifacts related to data sparsity (Kak and Slaney, 2001). Still when using gradient only, slightly wrong model updates are hoped to average out over the course of iterations.

4.3.3 Resolving power

Maximum theoretical resolution of FWI reaches half a wavelength provided normal-incidence reflections are present in the dataset (Virieux and Operto, 2009). Relationship between the experimental setup (fig. 4.4) and the spatial resolution of the reconstruction (scattering wavenumber $k = |\mathbf{k}|$) is

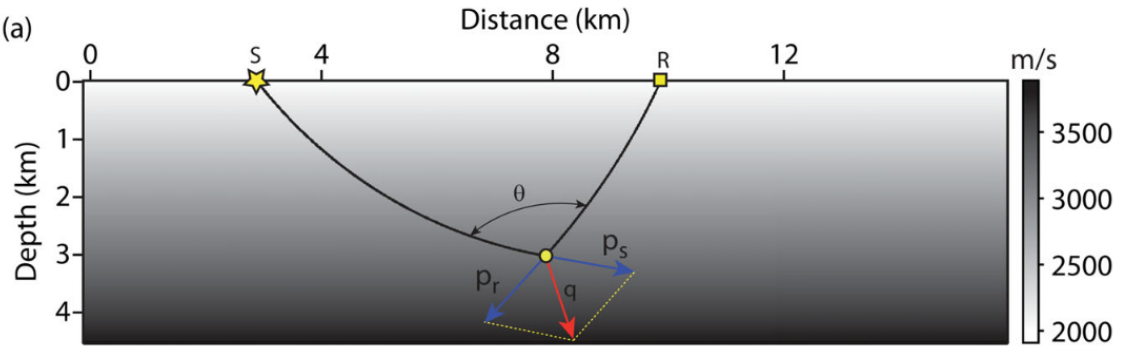


Figure 4.4: Scattering experiment. From Operto et al. (2015).

given by:

$$\mathbf{k} = \frac{2f}{c} \cos(\theta/2)\, \hat{\mathbf{q}}, \tag{4.50}$$

where c is the local phase velocity, f is the frequency of the harmonic wave, θ is the scattering angle and $\hat{\mathbf{q}}$ is the unit vector in the direction shown in fig. 4.4. The resolving power of FWI is thus limited primarily by two factors: (1) the source bandwidth and (2) scattering angle and thus aperture of the survey (Woodward, 1992).

4.4 FWI in practice

4.4.1 Challenges

Perhaps the most important aspects of FWI tha make it a particularly challenging inverse problem are:

- high *non-linearity* of $\mathbf{d}(\mathbf{m})$ – the oscillatory character of the data; its effects are amplified by significant heterogeneity of real models;
- *ill-conditioning* – often many model parameters are severely under-determined – even if $\mathbf{d}$ is longer than $\mathbf{m}$, the linearised system is called *rank-deficient*;
- very high *computational cost* of the forward problem, both in terms of memory and processor time due to the size of d and m vectors and number of operations per grid point to solve the wave equation;
- *accuracy* of the forward modelling – especially in the presence of irregular free surface boundary and strong viscoelastic effects.

The forward-modelling accuracy is the subject of part II. The remainder of the challenges are discussed below.

Non-linearity This is probably the most challenging of the above and numerous approaches have been proposed to tackle it and prevent a solution from being trapped in a local minimum (Virieux and Operto, 2009). The following overview touches only on a few arbitrarily selected avenues of research. Deriving sufficiently accurate starting models is an obvious prerequisite. Indispensible proved to be also a multi-scale strategy (Bunks et al., 1995), at the very least iterating from low to high frequencies of the data. Some version of this approach is discussed in ch. 8). An important area of active research is concerned with quantifying the waveform misfit so that the objective function has fewer (ideally one) minima (see sect. 4.4.3). Yet another approach is augmenting the objective function with an auxiliary

term that pulls the solution out of the local minimum (Bharadwaj et al., 2016). The ultimate solution would be to abandon the linearisation completely and switch to the Monte Carlo methods (Sambridge and Mosegaard, 2002). The full Monte Carlo approach is not yet achievable despite the ongoing efforts (Fichtner et al., 2019). A partial solution to this problem applies a global optimisation scheme only to parameters that effectively can be treated as smoothly-varying (such as quality factor Q) requiring much smaller model space to explore, while others (such as v_P) are still inverted in a local fashion (da Silva and Warner, 2017). To this end, various Monte Carlo methods can be used such as genetic (Debens, 2015) or quantum particle swarm (da Silva et al., 2019) optimisation.

Ill-conditioning It results from broadly defined data sparsity and is usually addressed by standard regularisation techniques discussed in sect. 4.2.2 or simply smoothing the gradient to remove short-wavelength features below a resolution limit. Sometimes more sophisticated methods (rather than just effectively smoothing the model) such as modifications of the total-variation method (Menke, 2012) are employed to preserve sharp features such as salt-dome boundaries (Esser et al., 2016).

Computational cost Improving computational efficiency is a vast topic on its own and reviewing it properly is beyond the scope of this thesis. The most obvious method is to reduce the number of forward simulations required to solve the inverse problem. The most important algorithmic advance that made FWI conceivable was the discovery of the adjoint method that allows to reduce the number of forward simulations per iteration to fewer than 5 for the steepest descent method compared to billions (size of typical models) otherwise required to calculate the Jacobian. A simple derivation of this method is presented in sect. 4.4.4.

4.4.2 Data requirements

The performance of transmission FWI is likely to improve when the input field data contain: (1) low frequency, (2) long offsets, and (3) good azimuthal coverage (Morgan et al., 2013). The first depends on a type of sources and receivers used, the latter two on acquisition geometry. Low-frequency content can be thus thought of as a *global* feature of a given experiment, given source and receiver types are the same across the whole experimental domain. The latter two are *local* features that depend on data coverage and size of the subdomain selected for the inversion.

Low-frequency Presence of signal in the low-frequency part of the spectrum is helpful for at least two reasons. First, it allows for a starting model to be further away from the true one, as the longer periods are less prone to cycle-skipping (Virieux and Operto, 2009). Second, inverting low-frequency data updates low-wavenumber features of the model (Sirgue and Pratt, 2004). Both of these imply that FWI should be started at the lowest possible frequency in order to be robust against the inadequacy of the starting model and thus to avoide cycle skipping. Low-frequency content depends on the acquisition method, in particular the effective frequency spectrum of the source. The lowest usable (i.e. with a S/N ratio above some threshold) frequency can be determined via phase plots of the observed data (fig. 8.8) which also allows for a quick detection of noise at different frequencies and offsets in the dataset (sect. 8.3.2).

Long-offset The importance of having long-offset data are also two-fold. First, similar to low frequencies, they provide additional control on long-wavenumber model updates needed at early stages of the inversion (Sirgue and Pratt, 2004). Second, the refracted waves recorded at long offsets sample

larger depths of subsurface. Increasing the penetration of short-offset data are only possible with waves reflected off the deep interfaces. However, due to relatively low reflectivity of these interfaces and long propagation distances, picking them out from noise, requires high-fold surveys relying on stacking many traces and thus suppressing the random noise at ca. $\sqrt{n}$ rate, where n is a number of traces. The exact penetration depth for a given offset is specific to model and acquisition geometry and can only be precisely determined by calculating a sensitivity kernel, or by ray tracing. However, as a first-guess estimate, one can assume that the offset required for refractions to illuminate a target is 3-6 times of its depth (Warner et al., 2010). This value should be treated as a lower bound, and increased accordingly for larger water and target depths.

Multi-azimuth Obtaining a non-distorted image of a three-dimensional structure requires illuminating it from different angles (azimuths). The only exception is when the structure is pseudo-3d, i.e. approximately constant in at least one dimension. This might be the case for a stack of planar layers or a fold with source-receiver lines perpendicular to its strike. However a 3d target will always benefit from a multi-angle illumination. A range of azimuths helps to make the inversion better posed and improve robustness against incoherent noise (Morgan et al., 2013).

4.4.3 Objective function

One of the areas of active research within the FWI community is concerned with designing an objective function such that the inversion converges to the global minimum as efficiently and effectively as possible. Inversion convergence depends jointly on the starting model and the objective function. Seismic data is oscillatory hence a lot of minima in an objective function based on the L^2 norm. The idea is to design a function that has as few minima as possible.

Metric One way is to treat data vectors as elements of metric spaces different from standard Euclidean space characterised by L^2 norm. Metrics such as L^1, Huber or Student-t norm are some of the possible choices (Virieux and Operto, 2009). Other measures include waveforms cross-correlation time-shift or difference between Dirac delta function and Wiener filter designed to map synthetic into observed data (Warner and Guasch, 2014; Guasch et al., 2019). It is also possible to separate phase and envelope misfit functions, which proved to be useful in FWI on a global scale (Fichtner et al., 2008). Envelopes are extracted by means of a Hilbert transform. Minimising a purely phase-driven functional based on the Rytov approximation is also possible (Shah et al., 2012b). Recently, the optimal-transport[10] misfit was proposed by Métivier et al. (2016) .

Search space Yet another approach is to optimise over both a wavefield and a model, thus increasing the search space and reducing the non-linearity of the misfit as a function of the model (van Leeuwen and Herrmann, 2013). Extending the data space may also be done by adding explicitly travel-time (Zhou et al., 1995) or already-mentioned phase misfits (Shah et al., 2012b).

4.4.4 Gradient calculation

Below, the adjoint method (e.g. Plessix, 2006) of calculating the least-squares gradient (eq. 4.37) is presented. In the machine learning community, it is known as the reverse-mode differentiation (No-

[10] Research on optimal transport has resulted in a few Fields medals in recent years.

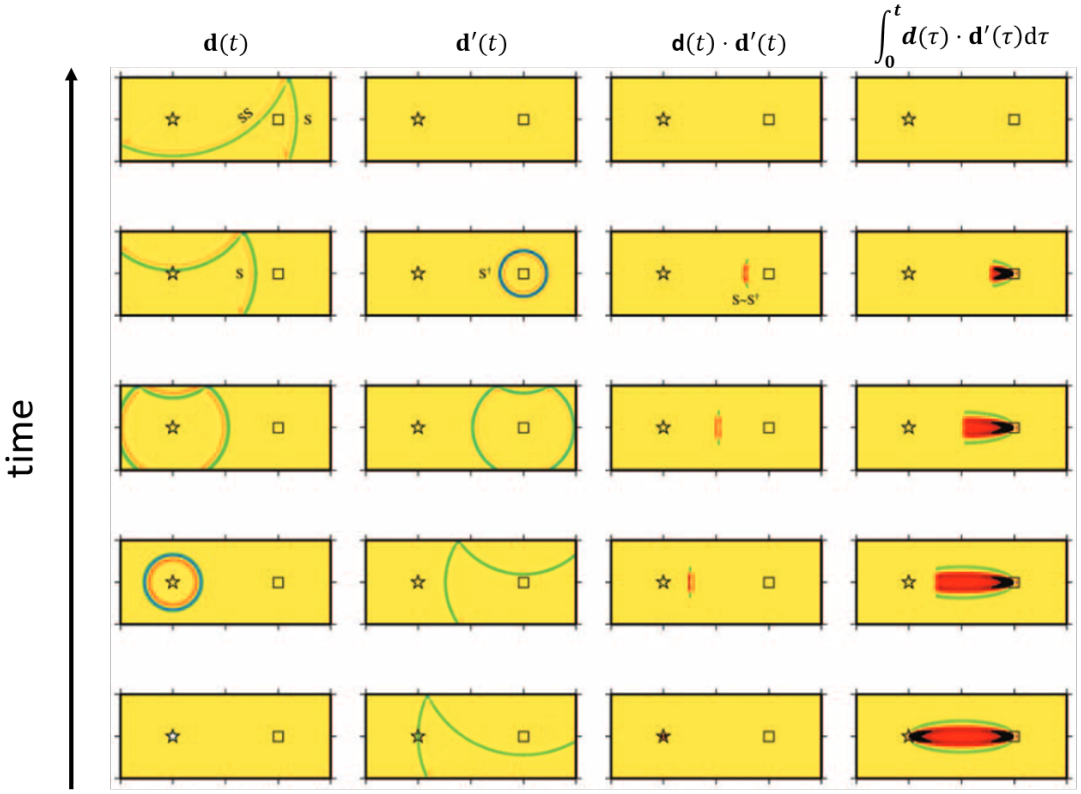


Figure 4.5: Time-lapse snapshots of the forward and adjoint wavefields, as well as their product, both instantanuous and time-integrated (columns from left to right, respectively). The adjoint wavefield was scaled with $\frac{\partial \mathbf{A}}{\partial \mathbf{m}}$. Modified after Tromp et al. (2005).

cedal and Wright, 2006), one of the flavours of *autodifferentiation* (Baydin et al., 2017). The insightful link between the two research fields is provided by Sambridge et al. 2007 and Ren et al. 2020.

In the derivation in the previous section, the starting point was the explicit non-linear dependence – eq. (4.24) which could be expanded as eq. (4.5). Here let us commence with the latter. Assuming that $\mathbf{A}$ is a linear function of $\mathbf{m}$ (recall the Born approximation) we can calculate the Jacobian by differentiating both sides of the eq. (4.5):

$$\partial_k (A_{ij} d_j) = \partial_k s_i \tag{4.51}$$

$$A_{ij,k} d_j + A_{il} d_{l,k} = 0 \tag{4.52}$$

$$d_{l,k} = -A_{il}^{-1} A_{ij,k} d_j \tag{4.53}$$

1540 The last line in the vector notation is:

$$\mathbf{J} = -\mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{d} \tag{4.54}$$

The tranpose needed in (4.37) gives:

$$\mathbf{J}^T = -\mathbf{d}^T \left(\frac{\partial \mathbf{A}}{\partial \mathbf{m}} \right)^T (\mathbf{A}^T)^{-1} \tag{4.55}$$

because $\left(\mathbf{A}^{-1}\right)^T = (\mathbf{A}^T)^{-1}$. Eventually the gradient can be written as:

$$\nabla f = -\mathbf{d}^T \left(\frac{\partial \mathbf{A}}{\partial \mathbf{m}} \right)^T (\mathbf{A}^T)^{-1} \delta \mathbf{d}. \tag{4.56}$$

The right term $(\mathbf{A}^T)^{-1}\delta\mathbf{d}$ can be seen as a wavefield $\mathbf{d}'$ generated by the source $\delta\mathbf{d}$ and propagated with the operator $\mathbf{A}^T$:

$$\mathbf{A}^T\mathbf{d}' = \delta\mathbf{d} \tag{4.57}$$

The operator $\mathbf{A}$ is often designed to be symmetric (self-adjoint) in space and anti-symmetric in time, so that acting on the residuals at the receivers (remember that the data is a function of space and time) with $(\mathbf{A}^T)^{-1}$ corresponds to propagating them from receivers back in time, or *backpropagating*. The backpropagation can be achieved perhaps most easily by reversing the data vector. That allows the use of the same operator for both forward and backpropagation. The gradient in (4.56) expresses a temporal zero-lag cross-correlation (or the inner product) of $\mathbf{d}$ and $\mathbf{d}'$, the latter being scaled with the operator $\frac{\partial\mathbf{A}}{\partial\mathbf{m}}$. This operator can be derived analytically and hard-coded.

The adjoint method requires only two forward simulations to calculate the gradient and it does not invert, compute or store the much larger Jacobian matrix. The finite-difference derivation of the latter would require perturbing and solving the forward problem for each of the model parameters in turn, which is prohibitively expensive for any real problem.

Even though inversion is driven solely by data, i.e. a wavefield sampled sparsely[11] at receiver locations, the entire wavefield is needed for the backpropation step (fig. 4.5).

4.4.5 Gradient preconditioning

The adjoint method discussed in the previous sect. provides an extremely efficient way to calculate the steepest-descent (anti-gradient) direction. However, despite its efficiency per iteration, the steepest descent method is known for its slow convergence. Although it is relatively feasible to calculate the Gauss-Newton Hessian, its inverse is usually too costly to evaluate. In some cases, additional iterations of the steepest descent algorithm will yield a better result than fewer but more time-consuming iterations with the quasi-Newton method (Tarantola, 2005). Moreover, the convergence rate of the steepest descent can be improved at little cost by tweaking the direction of the model update using some physical intuition on what the better direction should be (Tarantola, 2005). This correction is called *preconditioning* and has to be done before calculating the step length.

One way to spatially precondition the gradient is to neglect the off-diagonal elements of the approximate Hessian (Shin et al., 2002), since inverting a diagonal matrix is simply inverting each element of the diagonal. For the Gauss-Newton Hessian, this still needs one forward modelling per receiver (Shin et al., 2002) which is a considerable effort. Instead, a more heuristic method is to define the model update as (Warner et al., 2013b):

$$\delta m_i = -\frac{\alpha}{H_{ii}}\frac{\partial f}{\partial m_i}. \tag{4.58}$$

Note that the summation convention does not apply here. The scalar H_{ii} is approximately proportional to the wavefield energy so that the gradient is less dependent on the differences in wavefield illumination. In particular, it prevents anomalously large updates in model parts close to the sources and receivers but enhances contributions from poorly illuminated regions (Warner et al., 2013b).

[11] even for very dense acquistion geometries, number of receivers is usually an order of magnitude smaller than number of grid points.

Part II

Full-waveform modelling

5

Modelling multi-component data

BENEFITS of using multi-component data are well known in the context of reflection seismics. Combining pressure and vertical velocity components is routinely used to remove free-surface reflections from the data. Horizontal velocity components, dominated by shear waves, are crucial for inferring elastic properties of the subsurface.

Potential advantage of multi-component data in the full-waveform inversion is less understood, especially in the acoustic approximation. Furthermore, adding a capability of simulating and inverting multi-component data to the existing, highly-optimised acoustic FWI codes may be challenging. It is not always straightforward to accurately simulate both wavefields in a single numerical scheme, establish a reciprocal relationship between them, or to weight them properly in the inversion. As a result, despite multi-component acquisition being common nowadays, no case studies of the FWI applied to such field data have been reported to date.

This work demonstrates a method of acoustic full-waveform modelling and inversion of multi-component data. A reciprocal relation between pressure and velocity necessary for efficient inversion and forward simulation of sparse-receiver datasets is derived. Details of accurate implementation of dipole sources and receivers required to maintain reciprocity in the second-order acoustic wave equation, are presented. The method was implemented as a small, self-contained `Python` library `HicksPy` facilitating a rapid extension of existing FWI codes. The implementation was tested using a 3d field dataset from PROTEUS experiment, showing a good match between 2nd-order acoustic synthetics and observed vertical velocity data.

Key words: geophones · particle velocity · second-order acoustic wave equation · reciprocity · dipole sources · directional radiation patterns · joint inversion · vector-acoustic FWI

5.1 Introduction

5.1.1 Seismic sensors

There are two primary types of sensors used routinely in seismology – hydrophones and geophones which measure pressure and particle motion (displacement, velocity or acceleration), respectively[1]. Geophones are usually designed in a form of electro-magnetic coils, and hydrophones as piezoelectric discs stacked for signal amplification (Sheriff and Geldart, 1995).

Basic difference Hydrophones, as the name suggests, can only work immersed in water (on land surface, pressure is effectively zero). Geophones measure particle motion and can give meaningful measurements both in water and on land. Deploying geophones in water is more difficult (and thus expensive) than hydrophones, as they need coupling to the seabed. Poor coupling of a geophone may result in a non-linear instrument response.

Polarity reversal A buried vertical geophone records double polarity reversal of the ghost – one factual and one due to up/down change of the propagation direction (Sheriff and Geldart, 1995). In effect, no reversal is observed. A hydrophone records only the factual reversal.

AVO behaviour Each component v_i is a 1d projection of a 3d vector $\mathbf{v}$. E.g. amplitude of v_z relative to $|\mathbf{v}|$ varies with the angle of incidence θ like:

$$v_z = |\mathbf{v}| \cos(\theta) \tag{5.1}$$

The angle of incidence at the receiver increases with the source-receiver offset. Thus, v_z dominates the long-offset data in terms of amplitude. Pressure does not exhibit similar AVO behaviour.

5.1.2 Multi-component data

Dual-sensor acquisition provides multi[2]-component data. Marine geophones have been routinely deployed to improve recording of S-waves and inferring elastic properties since the last decade of the previous century (e.g. Operto et al., 2013). Some of the potential advantages of this type of acquisition, relevant for *acoustic* applications, are listed below:

- ***deghosting*** – this is a common preprocessing step in reflection seismic and is performed by PZ summation[3]; it relies on the opposite polarities in hydrophone and geophone response (see sect. 5.1.1);
- ***better AVO*** – vertical component dominates the particle motion due to P waves arriving at long offsets (see eq. 5.1);
- ***increased fold*** – more components increase the chance of high quality data for a given receiver; the components may be stacked or inverted separately, e.g. for different offsets (see the previous point);
- ***vector-acoustic FWI*** – joint inversion of pressure and velocity data may result in less free-surface-related gradient artifacts, as shown in 2d Marmousi synthetic tests by (Zheglova and Malcolm, 2020);The fidelity gain is more noticeable for sparse data but overall not overwhelming (fig. 5.2).

[1] Optical sensors such as fibres and laser rings (e.g. Nigbor, 1994) developed more recently are not discussed.
[2] Typically four-component, or 4C for short.
[3] Stacking carefully weighted pressure and vertical velocity data.

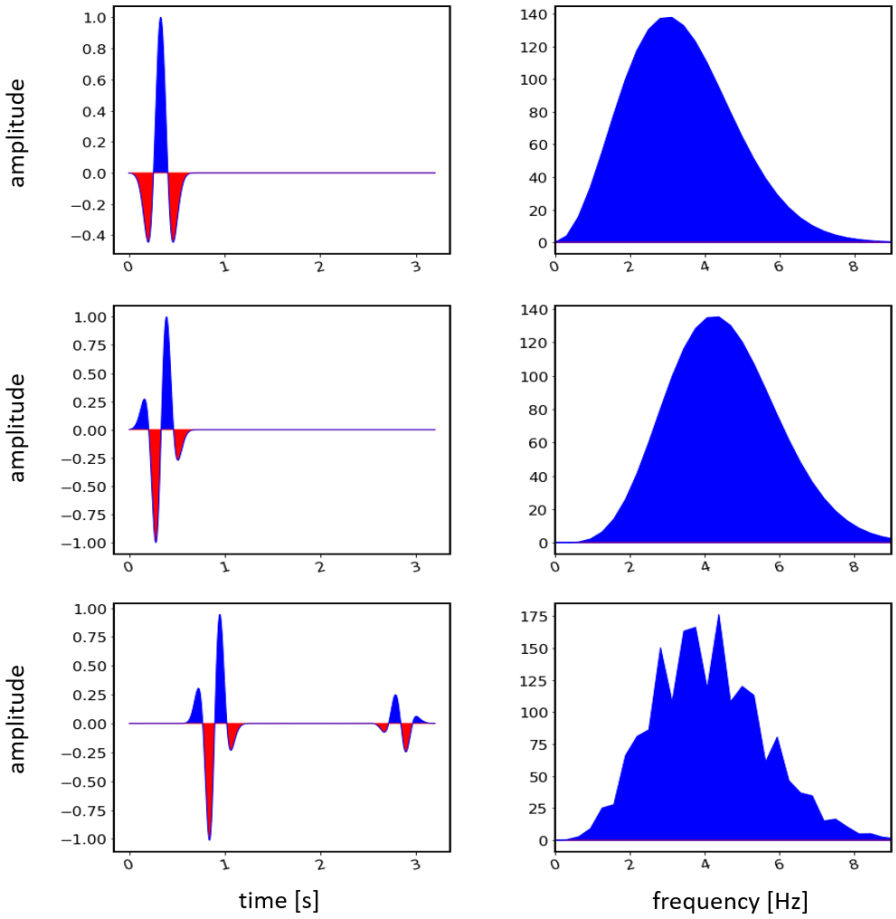


Figure 5.1: Amplitude spectrum of a direct wave (top) its time-derivative and (middle) direct wave (time-differentiated) with a ghost (bottom). In the last case, harmonic components of the two wave packets interfere leading to the notches in the spectrum.

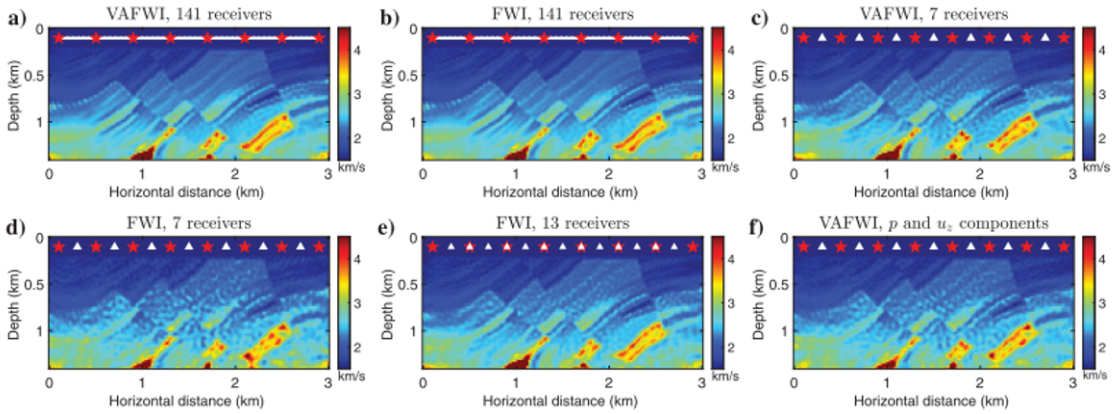


Figure 5.2: Final models of Marmousi 2d recovered by FWI using different station coverage and sensor types. Benefits of inverting more than one component are subtle but noticable, especially for sparse acquisition. From Zheglova and Malcolm (2020).

2004	Acoustic 2d inversion of land v_z data processed so as to mimic the AVO behaviour of the pressure wavefield by Ravaut et al.
2006	Acoustic 2d inversion of OBS v_z data using an acoustic solver with reciprocity (implementation details not provided) by Operto et al.
2008	Elastic 2d inversion of synthetic OBS data (p vs. v_z vs. v_z& v_x) by Sears et al. Different reverberation patterns for p vs. v_z data.
2010	Elastic 2d inversion of OBS data (v_z followed by v_x data) by Sears et al.
2011	Acoustic 2d inversion of land v_z data without reciprocity by Malinowski et al. Dipole receivers implemented via modification of the Hicks scheme (Hicks, 2002).
2012	Acoustic 2d inversion of land vibroseis v_z data by Plessix et al..
2018	Acoustic 2d inversion of land data by Vigh et al.
2019	Elastic 2d inversion of OBS v_z vs. p data (two 1C inversions) by Robin et al. Difference in results driven by data noise.
2019	Acoustic 3d inversion of land v_z data by Sedova et al.

Table 5.1: Time-line of previous approaches to modelling and inversion of velocity/multi-component data.

5.1.3 Modelling approaches

Different approaches to modelling multi-component data (or velocity data with 2nd-order acoustic wave equation) have been published to date. A short overview of major ones is given in tab. 5.1. Worth noting is the method of Plessix et al. (2012) who simulate v_z data at the land surface, $v_z(0)$, by shifting the pressure sources and receivers by a small amount (5 m) below the free surface. This is based on the observation that:

$$\rho \dot{v}_z(0) = -\partial_z p(0) = \frac{p(0+h) - p(0)}{h} = p(h)/h \tag{5.2}$$

because $p(0) = 0$ by definition of a free surface. A corollary of eq. (5.2) is that a shallow pressure receiver should record a wavefield that is proportional to particle acceleration at the free surface:

$$\dot{v}(0) \sim p(h). \tag{5.3}$$

Sedova et al. (2017) follow this approach and bury the geophones in their forward simulations too.

This study Here, a method of modelling multi-component data using a 2nd-order acoustic wave equation is presented. It uses the same differentiation method as Malinowski et al. (2011) but it extends it to experiments utilising reciprocity principle. The below description comprises both the theoretical background and implementation details.

5.2 Theory

5.2.1 Pressure & velocity

Below, conversion formulae between two related acoustic wavefields – pressure and particle velocity[4,5] (measured by hydrophones and geophones, respectively), are provided. The relationship between pressure and velocity can be derived from two basic equations governing the mechanics of continuum – the *Newton's* second law and *Hooke's* constitutive relation. The former states that the acceleration of

[4] It will be referred to as velocity for brevity, unless it can be mistaken for e.g. wave velocity.

[5] If needed, displacement and acceleration can be obtained from velocity by time-integration and differentiation, respectively.

a unit volume of a fluid is equal to the negative gradient of pressure:

$$\rho \dot{v}_i = -\partial_i p, \tag{5.4}$$

The latter expresses temporal variations of pressure in terms of the divergence of a velocity field:

$$\frac{1}{K}\dot{p} = -\partial_i v_i, \tag{5.5}$$

where $K = \rho c^2$ is the bulk modulus. Combined, eq. (5.4) and (5.5) form a first-order system of linear partial differential equatioms:

$$\begin{cases} \rho \dot{v}_i = -\partial_i p \\ \dfrac{1}{K}\dot{p} = -\partial_i v_i \end{cases} \tag{5.6}$$

equivalent to a second-order equation for pressure:

$$\frac{1}{c^2}\ddot{p} = \rho \partial_i \left(\frac{1}{\rho} \partial_i p \right), \tag{5.7}$$

which can be obtained by inserting (5.4) into a time derivative of (5.5):

$$\frac{1}{c^2}\ddot{p} = -\rho \partial_i \dot{v}_i = \rho \partial_i \left(\frac{1}{\rho} \partial_i p \right), \tag{5.8}$$

where we implicitly assumed a stationary density model ($\dot{\rho} = 0$). From (5.4), we have:

$$v_i(\mathbf{x}) = -\frac{1}{\rho(\mathbf{x})} \int \mathrm{d}t \partial_i p(\mathbf{x}), \tag{5.9}$$

in which the gradient operator introduces sensitivity to the direction of propagation whereas the time-integration cancels the influence of the gradient on the amplitude spectrum. We can express it as:

$$v_i(\mathbf{x}) = \frac{1}{\rho(\mathbf{x})} \diamondsuit p(\mathbf{x}) \tag{5.10}$$

where:

$$\diamondsuit \coloneqq -\int \mathrm{d}t \partial_i \tag{5.11}$$

is the operator that transforms pressure into velocity (up to the density scaling-factor). For completeness, the pressure analog of (5.9) is given by:

$$p(\mathbf{x}) = -\rho(\mathbf{x}) c^2(\mathbf{x}) \int \mathrm{d}t \partial_i v_i(\mathbf{x}). \tag{5.12}$$

Let us now consider how $\diamondsuit$ affects waves propagating in opposite directions. The amplitudes ψ_u and ψ_d of an upgoing and dowgoing wave propagating from a source at $(z, t) = (0, 0)$ in $-z$ and $+z$ direction, respectively, can be expressed as[6]:

$$\psi_u(z, t) = f(\omega t + kz) \tag{5.13}$$

$$\psi_d(z, t) = f(\omega t - kz) \tag{5.14}$$

The negative-gradient transforms them into:

$$-\partial_z \psi_u = -k f'(\xi)\Big|_{\xi=\omega t+kz} \tag{5.15}$$

$$-\partial_z \psi_d = +k f'(\xi)\Big|_{\xi=\omega t-kz} \tag{5.16}$$

[6] These are solutions to the two-way wave equation which can be easily checked or compared to eq. (3.39) of Shearer (2009).

where $f'(\xi) \equiv \frac{\partial f(\xi)}{\partial \xi}$. Integrating these with respect to time gives:

$$-\int \mathrm{d}t \partial_z \psi_u = -\frac{k}{\omega} f(\omega t + kz) + C(z) \tag{5.17}$$

$$-\int \mathrm{d}t \partial_z \psi_d = +\frac{k}{\omega} f(\omega t - kz) + C(z). \tag{5.18}$$

Finally, if we put $C(z) = 0$, then recall the definition of $\diamondsuit$ and notice that $k/\omega = 1/c$, we obtain:

$$\begin{aligned} \diamondsuit \psi_u &= -\frac{1}{c}\psi_u \\ \diamondsuit \psi_d &= +\frac{1}{c}\psi_d. \end{aligned} \tag{5.19}$$

This means that $\diamondsuit$ operator flips the polarity of upgoing waves and leaves the polarity of downgoing waves unchanged. In other words, a waveform recorded on the pressure and vertical-velocity channels will have the opposite polarity for upgoing waves and the same polarity for downgoing waves. Physically, a compressive[7] ($p > 0$) impulse generates outward first motions[8] so, at the receiver, a downgoing wave will move particles downwards ($v_z > 0$), and an upgoing wave – upwards ($v_z < 0$ which is the opposite polarity to p).

Lastly, let us note that pressure and velocity wavefields are not independent. Each component of velocity can be calculated from the pressure scalar field (eq. 5.9), and pressure can be calculated from the velocity vector field (eq. 5.12), provided both fields are differentiable in space.

5.2.2 Seismic sources

Below, it is shown that a single force source is equivalent to a dipole source of volume injection. So far sources of waves have not been discussed. Acoustic waves may be generated in two basic ways – by injecting a volume (explosion/implosion depending on a sign) or by exerting a force. Mathematically, this can be written down as a combination of two terms

$$s = \rho \partial_t I_V - \rho \partial_i \left(\frac{1}{\rho} f_i \right), \tag{5.20}$$

that can be added to the wave equation:

$$\frac{1}{c^2} \partial_{tt} p = \rho \partial_i \left(\frac{1}{\rho} \partial_i p \right) + s. \tag{5.21}$$

Here, I_V and f_i denote the volume injection and ith component of the force respectively. The first term will be referred to as a *volume source* s_V and the second as a *force source* s_f. Note that the latter term is a divergence ('source-ness') of a vector field. A volume point source at (x_s, y_s, z_s) can be described as a multi-dimensional Dirac delta:

$$s_V(x, y, z) = A\delta(x - x_s)\delta(y - y_s)\delta(z - z_s) \tag{5.22}$$

Let us know define a point vertical-force source as:

$$\mathbf{f} = \begin{bmatrix} 0 \\ 0 \\ \delta(\mathbf{x})S(t) \end{bmatrix} \equiv \begin{bmatrix} 0 & 0 & \delta(\mathbf{x})S(t) \end{bmatrix}^T. \tag{5.23}$$

[7] A positive perturbation of ambient pressure.
[8] See fig. 9.8 of Shearer (2009).

It has to be added to the scalar pressure wavefield as a divergence, hence:

$$\begin{aligned}\nabla \cdot \mathbf{f} &= \partial_x \cdot 0 + \partial_y \cdot 0 + \partial_z[\delta(\mathbf{x})S(t)] \\ &= S(t) \lim_{h\to 0} \delta(x)\delta(y)\frac{\delta(z+h) - \delta(z-h)}{2h}\end{aligned} \tag{5.24}$$

where a central finite difference was used in the definition of the derivative. It should now be clear that eq. (5.23) represents a vertical dipole. A volume source has an isotropic radiation pattern and will interchangeably be referred to as a *monopole source*. In contrast, a force source has a radiation pattern of a *dipole* aligned with this force. Note that sources of each type generate both types of wavefields, as there is no independent pressure or velocity wavefield. The *radiation pattern* in both cases refers to a shape of the pressure wavefront in the homogeneous, isotropic space. In the case of a receiver, a radiation pattern corresponds to sensitivity of the measuring apparatus to the wavefield incoming from different directions. For instance, a vertical geophone is insensitive to horizontal component of the particle motion. A source counterpart of this, is a single force vector.

5.2.3 Wavefield response

Below, analytical expressions for pressure and velocity responses to both types of sources discussed above are derived directly from the definition of the elastodynamic Green's function.

Elastodynamic Green's function A single body force is not a realistic seismic source, as it violates the law of conservation of momentum (Shearer, 2009). Nonetheless it is a useful concept as other, more realistic sources can be decomposed into ingredients of this type. The elastodynamic Green's function G_{ij} is a tensor encoding a response to such a body force. It is defined (Aki and Richards, 2009) by:

$$u_i(\mathbf{x}_A, t) = G_{ij}(\mathbf{x}_A, t; \mathbf{x}_B, 0) \cdot f_j(\mathbf{x}_B, 0). \tag{5.25}$$

To be precise, it is an impulse response of the displacement field. Responses of other types of wavefields to other types of sources can be constructed from the Green's function by spatial superposition and temporal convolution by the virtue of linearity and time-invariant of the wave equation (it is an LTI system).

Velocity response to explosion An explosion can be represented by a diagonal moment tensor, i.e. three perpendicular force couples without moment (Aki and Richards, 2009). A force couple without moment consists of two vectors of magnitude f pointing in opposite directions separated by a distance d along the direction parallel to these forces (fig. 5.3). Displacement field generated by such a force couple aligned with kth axis is given (Shearer, 2009) by:

$$u_i(\mathbf{x}) = G_{ij}(\mathbf{x}', t'; \mathbf{x}, t)f_j - G_{ij}(\mathbf{x}', t'; \mathbf{x} - d\mathbf{e}_k, t)f_j \tag{5.26}$$

where $\mathbf{e}_k$ is the Cartesian versor associated with the kth axis. This can be reduced to:

$$u_i(\mathbf{x}) = G_{ij,k}f_j d \tag{5.27}$$

For an explosion (three perpendicular force couples), summing over each axis gives:

$$u_i(\mathbf{x}) = A \cdot G_{ij,j}, \tag{5.28}$$

where $A \coloneqq fd$.

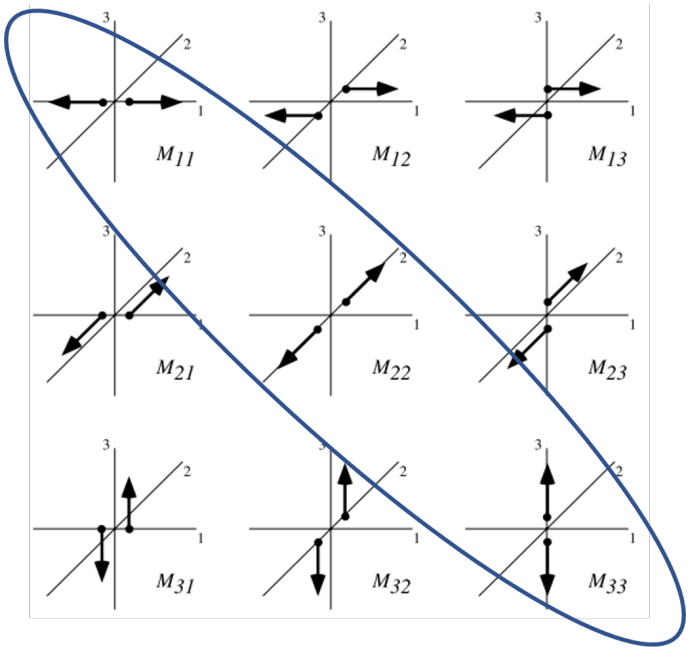


Figure 5.3: Graphical depiction of the components of a moment tensor. Diagonal elements in the ellipse describe an explosion. Modified after (Shearer, 2009).

5.2.4 Reciprocity principle

Below, it is shown that the pressure response at the point A to a single-force source at the point B is equal to the velocity response at the point B to a volumetric source at the point A.

Motivation For some experiments, a number of sources is orders of magnitude larger than the number of receivers. This is often caused by the higher cost of receiver acquisition. In time-domain, unless source-encoding is used, the computational cost is proportional to the number of sources. Hence, if there was a way to simulate seismic sources as receivers and receivers as sources, that would be a great computational gain. Luckily, there exists a fundamental symmetry between sources and receivers that allows to do that.

Pressure-velocity A basic principle of acoustic reciprocity states that the same waveform will be observed if the positions of the source and receiver are interchanged, no matter how far apart they are or how heterogeneous the medium is (Rayleigh, 1878). Stated like this, it implicitly assumes that both source and receiver are of certain types, namely that they are isotropic. A question arises if the symmetry holds for other types of source-receiver configurations. The answer is yes[9], but it brings in extra conditions which will be discussed below. A spatial reciprocity of the Green's function is expressed by:

$$G_{ij}(\mathbf{x}_A, t; \mathbf{x}_B, 0) = G_{ji}(\mathbf{x}_B, t; \mathbf{x}_A, 0), \tag{5.29}$$

where the LHS stands for an ith component of the displacement wavefield at point $\mathbf{x}_A$ and time t generated by a jth component of a point force at $\mathbf{x}_B$ and time 0, and the RHS denotes a jth component

[9] In fact, reciprocity is much more general and appears in flow and diffusion problems too (Wapenaar and Fokkema, 2004). Reciprocity principle can be integral (global) and differential (local). The former is of interest of interferometry and Marchenko imaging (Wapenaar et al., 2019) and migration (Vasconcelos et al., 2017). Reciprocity holds for anisotropic and dissipative media (Wapenaar and Berkhout, 1989).

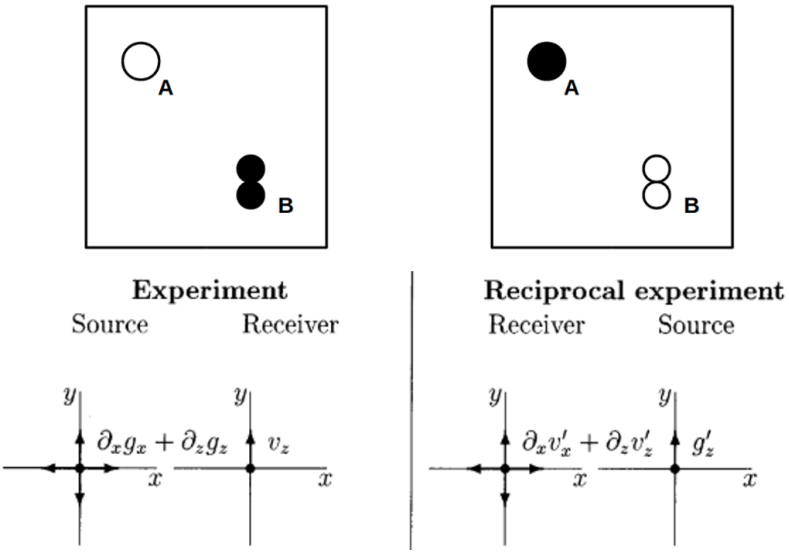


Figure 5.4: Schematic of reciprocal experiments involving monopole and dipole radiation patterns. Wavefield from a monopole source at location A recorded at location B by a dipole receiver (left) is equal to the wavefield from a dipole source at location B recorded at location A by a monopole receiver (right). The radiation pattern is denoted by the empty ('radiating out') and filled ('collecting in') shape (single/double circle for monopole/dipole) for the source and receiver, respectively. Bottom-row figures from Arntsen and Carcione (2000).

of the displacement wavefield at point $\mathbf{x}_B$ and time t generated by a ith component of a point force at $\mathbf{x}_A$ and time 0. The proof assumes homogeneous[10] boundary conditions and is given by Aki and Richards (2009). Differentiating both sides of eq. (5.29) with respect to time gives:

$$v_{ij}(\mathbf{x}_A, t; \mathbf{x}_B, 0) = v_{ji}(\mathbf{x}_B, t; \mathbf{x}_A, 0), \tag{5.30}$$

which is the reciprocal relation for velocity. Using reciprocity we can write eq. (5.28) as:

$$u_i = A \cdot G_{ij,j}(\mathbf{x}_A, t; \mathbf{x}_B, 0) \tag{5.31}$$

$$= A \cdot G_{ji,j}(\mathbf{x}_B, t; \mathbf{x}_A, 0). \tag{5.32}$$

The quantity on the RHS can be interpreted as $\dot{p}$ from a ith component source using the relationship between pressure and velocity derived earlier. Therefore one can write:

$$v_i = A \cdot \dot{G}_{ij,j}(\mathbf{x}_A, \mathbf{x}_B) \tag{5.33}$$

$$= A \cdot \dot{G}_{ji,j}(\mathbf{x}_B, \mathbf{x}_A) = A \cdot \partial_j v_j^i = -\frac{A}{c^2 \rho} \cdot \dot{p}^i. \tag{5.34}$$

Rewriting LHS and RHS gives:

$$v_i^p(\mathbf{x}_A, \mathbf{x}_B) = A \cdot \partial_j v_j^i = -\frac{A}{c^2 \rho} \cdot \dot{p}^i(\mathbf{x}_B, \mathbf{x}_A). \tag{5.35}$$

which is a reciprocal relationship between pressure and velocity.

[10] Homogeneity means that if the wavefield satisfies the boundary condition, its scaled version also does. It is the case for the pressure wavefield with the free surface and absorbing boundaries.

5.2.5 Numerical reciprocity

Below, it is shown that in order for the theoretical reciprocity to hold, the numerical scheme has to be self-adjoint in space. If the latter is not true, the source and/or receiver radiation patterns have to be modified accordingly.

The theoretical reciprocity of the acoustic wave equation shown in sect. 5.2.4 does not guarantee the reciprocity of a numerical simulation. The latter depends not only on the underlying physics but also on the numerical scheme chosen to simulate it (Pratt et al., 1998). Many of the definitions and the notation in the following analysis is similar, though not identical, to the one used by Ren et al. (2020).

Let us consider a general linear differential equation in a discrete form:

$$\mathbf{A}\mathbf{p} = \mathbf{s}, \tag{5.36}$$

where $\mathbf{A}$ is a differential operator which defines a field response $\mathbf{p}$ to the source $\mathbf{s}$. For convenience of the theoretical discussion, we may write eq. (5.36) as:

$$\mathbf{p} = \mathbf{B}\mathbf{s}, \tag{5.37}$$

where $\mathbf{B} = \mathbf{A}^{-1}$. In reality $\mathbf{A}$ may not be invertible, but the conclusions will hold as we will only be interested in the symmetry of $\mathbf{B}$. If $\mathbf{B}$ is symmetric, so is $\mathbf{A}$ (Riley et al., 2006). Given the domain consisting of N_z grid-cells and N_t time-steps and defining $N \coloneqq N_z \cdot N_t$, the operators $\mathbf{A}$ and $\mathbf{B}$ are $N \times N$ square matrices and $\mathbf{p}$ and $\mathbf{s}$ are vectors of length N. They are arranged so that the temporal and spatial coordinate are the slow and fast index, respectively. A data vector $\mathbf{d}$ of length N_t observed *at a single receiver* can be extracted from the wavefield by a $N_t \times N$ measurement matrix $\mathbf{D}$:

$$\mathbf{d} = \mathbf{D}\mathbf{p}. \tag{5.38}$$

Assuming the properties of the receiver do not change over time, the structure of $\mathbf{D}$ is as follows:

$$\mathbf{D} = \left[\mathbf{D}^1, \ldots, \mathbf{D}^{N_t}\right], \tag{5.39}$$

where each of $\mathbf{D}^i$ for $i = 1, \ldots, N_t$ is a $N_t \times N_z$ matrix of the form:

$$\mathbf{D}^i = \left[\underbrace{\mathbf{0}}_{1}, \ldots, \underbrace{\mathbf{r}^T}_{i}, \ldots, \underbrace{\mathbf{0}}_{N_t}\right]^T \equiv \begin{bmatrix} 0 & \cdots & 0 \\ r_1 & \cdots & r_{N_z} \\ 0 & \cdots & 0 \end{bmatrix}. \tag{5.40}$$

The vector:

$$\mathbf{r} = \begin{bmatrix} r^1 \\ \vdots \\ r^{N_t} \end{bmatrix}. \tag{5.41}$$

encodes the *radiation pattern* of the receiver. Its components r_i will be referred to as receiver weights/coefficients. For example, the radiation pattern of a point receiver located at the jth grid node is expressed by:

$$\mathbf{r} = \left[0, \ldots, \underbrace{1}_{j}, \ldots, 0\right]^T \tag{5.42}$$

or in the more concise index notation:

$$r_i = \delta_{ij}. \tag{5.43}$$

A source *wavelet* is a time-series of length N_t:

$$\mathbf{w} = \begin{bmatrix} w^1 \\ \vdots \\ w^{N_t} \end{bmatrix}, \tag{5.44}$$

1765 which can be converted into a source vector $\mathbf{s}$ of length N required by eq. (5.36) through a $N \times N_t$ matrix $\mathbf{S}$ such that:

$$\mathbf{s} = \mathbf{S}\mathbf{w}, \tag{5.45}$$

which has a block structure:

$$\mathbf{S} = \begin{bmatrix} \mathbf{S}^1 \\ \vdots \\ \mathbf{S}^{N_t} \end{bmatrix}, \tag{5.46}$$

where each of $\mathbf{S}^i$ for $i = 1, \ldots, N_t$ is a $N_z \times N_t$ matrix of the form:

$$\mathbf{S}^i = \begin{bmatrix} \underbrace{\mathbf{0}}_{1}, \ldots, \underbrace{\mathbf{e}}_{i}, \ldots, \underbrace{\mathbf{0}}_{N_t} \end{bmatrix} \equiv \begin{bmatrix} 0 & e_1 & 0 \\ \vdots & \vdots & \vdots \\ 0 & e_{N_z} & 0 \end{bmatrix}. \tag{5.47}$$

The vector $\mathbf{e}$ is a stationary source radiation pattern[11], similar to (5.41):

$$\mathbf{e} = \begin{bmatrix} e_1 \\ \vdots \\ e_{N_z} \end{bmatrix}, \tag{5.48}$$

1770 For instance, if N_t and N_z are equal to 2, and a point source and receiver are both located at the first grid node, eq. (5.38) gives:

$$\begin{bmatrix} d^1 \\ d^{N_t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} p_1^1 \\ p_{N_z}^1 \\ p_1^{N_t} \\ p_{N_z}^{N_t} \end{bmatrix} = \begin{bmatrix} p_1^1 \\ p_1^{N_t} \end{bmatrix}, \tag{5.49}$$

and eq. (5.45) reads:

$$\begin{bmatrix} s_1^1 \\ s_{N_z}^1 \\ s_1^{N_t} \\ s_{N_z}^{N_t} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} w^1 \\ w^{N_t} \end{bmatrix} = \begin{bmatrix} w^1 \\ 0 \\ w^{N_t} \\ 0 \end{bmatrix}. \tag{5.50}$$

Note that in this case $\mathbf{S} = \mathbf{D}^T$. This is true whenever the source and receiver are located at the same grid node. Let us now consider an experiment involving a single source of weights $\mathbf{e}$ and a single re-

1775 ceiver of weights $\mathbf{r}$. The observed data vector is given by:

$$\mathbf{d} = \mathbf{D}\mathbf{p} = \mathbf{D}\mathbf{B}\mathbf{s} = \mathbf{D}\mathbf{B}\mathbf{S}\mathbf{w}. \tag{5.51}$$

[11] As the symbol $\mathbf{s}$ is already used, the source radiation pattern will be denoted by $\mathbf{e}$ for *excitation* of the medium, after Claerbout (1985).

We can consider a scalar amplitude d observed as the lth sample at the receiver due to the jth sample of the source wavelet w^j. For simplicity, let us put $w^j \equiv 1$. It can be shown that:

$$d = \mathbf{r}^T \mathbf{C} \mathbf{e}, \tag{5.52}$$

where $\mathbf{C}$ is a block matrix located on the main diagonal of $\mathbf{B}$, its exact position depending on j and l. If we now design another experiment, in which the receiver has weights $\mathbf{e}$ and the source weights $\mathbf{r}$ (we effectively relabelled the source as a receiver, and the receiver as a source), the observed amplitude will be equal to:

$$d' = \mathbf{e}^T \mathbf{C} \mathbf{r}. \tag{5.53}$$

One can ask a question, what conditions are required for d' to be equal to d. As d is a scalar, it is equal to its transpose, $d = d^T$, and we can write:

$$d = \mathbf{r}^T \mathbf{C} \mathbf{e} = (\mathbf{r}^T \mathbf{C} \mathbf{e})^T = \mathbf{e}^T \mathbf{C}^T \mathbf{r}, \tag{5.54}$$

by the virtue of $(\mathbf{ABC})^T = \mathbf{C}^T \mathbf{B}^T \mathbf{A}^T$. Clearly, (5.54) is equal to (5.53), as required, provided $\mathbf{C}^T = \mathbf{C}$. This condition is met if $\mathbf{A}^T = \mathbf{A}$ which, in turn, implies $\mathbf{B}^T = \mathbf{B}$.

As the choice of j and l was entirely arbitrary, the conclusions hold for the full data vector due to the full source wavelet. We see that swapping source and receiver weights leads to a reciprocal experiment provided the matrix $\mathbf{A}$ defining the numerical scheme is self-adjoint. It was never assumed that $\mathbf{A}$ encodes the discretised wave equation, therefore the reciprocity principle holds for any equation which can be represented by a matrix with the above property.

5.2.6 Adjoint modelling

Below, it is shown that radiation patterns of the adjoint sources used in the back-propagation should be the same as the radiation patterns of the corresponding receivers, whether mono- or multi-component data are used. Acquisition configurations for which using the reciprocity in the forward step is beneficial are also given.

The adjoint method of calculating gradient for FWI, discussed in sect. 4.4.4 of ch. 4, is based on the reciprocity principle. Receivers become sources from which data residuals are propagated backwards in time and the resulting 'adjoint' wavefield is multiplied (cross-correlated) at each time-step by the wavefield propagated forward in time. The premise of the method is that these two wavefields interfere constructively only[12] at these portions of the model which produced the data residuals – the missing heterogeneities (diffractors). The reciprocity principle ensures that propagating the residuals from the receivers will focus at the diffractors after $t_{\max} - t$ time where t is the time taken by the forward wavefield to reach the diffractor and $t_{\max}$ is the duration of the simulation. For this reason, the cross-correlation gives the desired model update as an outcome. In reality, the constructive interference occurs elsewhere too, resulting in artifacts in the gradient. The gradient is hence an imperfect imaging operator but its errors are expected to substantially cancel out across iterations.

The above remarks lead to a conclusion that the adjoint modelling should be reciprocal to the forward one. As shown in sect. 5.2.5, this means that the adjoint sources should have the same weights (radiation patterns) as the receivers used previously in the forward step. A dipole receiver becomes a dipole adjoint source, and so on. Adjoint receivers are not needed for the cross-correlation. This

[12] This would be ideal. See the later remark.

principle works automatically for multi-component data (potential benefits of which are covered in sect. 5.4.4). For example, when inverting p and v_z data using reciprocity, each monopole receiver records data generated from two sources, a dipole and a monopole. Subsequently, two data-residuals are calculated which are then both backpropagated from a monopole adjoint souce. If we were *not* using reciprocity in the forward simulation of the same experiment, there would be two receivers per monopole source, a dipole and a monopole, and two adjoint sources, again – a dipole and a monopole.

The above example suggests that if the number of sources is at most twice/four times the number of two/four-component receivers, it is more computationally efficient *not* to use reciprocity in the forward simulation. However, in many real experiments, like the one described in sect. 5.5, the number of sources is 1-2 orders of magnitude larger than that of receivers. In such cases, using reciprocity is far more efficient.

5.3 Method

In the previous section it was shown that in order to simulate velocity data using a 2nd-order wave equation, dipole sources and receivers have to be implemented (the latter needed only if reciprocity is *not* applied). For this purpose, a highly accurate method based on the work by Hicks (2002) was implemented. It solves two problems at once – first, it allows for accurate representation of stations located between grid nodes. Second, it provides a means to endow them with directional radiation patterns. Implementation of the method is attached as a self-contained `Python` 3.x library `HicksPy` (see sect. 5.6).

5.3.1 Formulation

The goal here is to implement a method that:

1. simulates v_z data generated by an explosion, i.e. to implement the $\diamondsuit$ operator (eq. 5.11)
2. simulates a source that has a correct Fourier spectrum, and such that the experiment is reciprocal to the point 1.

Different approaches to resolving this problem are presented below.

Virtual sources One way to simulate an ideal dipole would be to use an additional station of the opposite amplitude. There is a severe problem with this approach though. To illustrate this, let us assume that the true station position is at i depth node. Then the extra station would need to be at $(i+1)$th or $(i-1)$th node. This effectively puts the net point station between grid nodes which is not its true position. One could put the true station at $i+1$ and the extra source at $i-1$ so that the net source is in the correct position. This, however, creates a dipole of an applied arm $d=2$ nodes, i.e. even worse an approximation of an ideal dipole ($d \to 0$) which will result in an even less accurate solution in the near field. Instead, one could use smaller dipoles which are located between grid nodes. This requires interpolation (see appendix 5.A for the error analysis). In 1d, general interpolation can be expressed as:

$$\psi(z) = \sum w_i(z)\psi_i. \tag{5.55}$$

Interpolation Note that as soon as the interpolation is required, the virtual receivers are no longer needed, as the differentiation (i.e. transformation of a monopole into a dipole) can be performed via

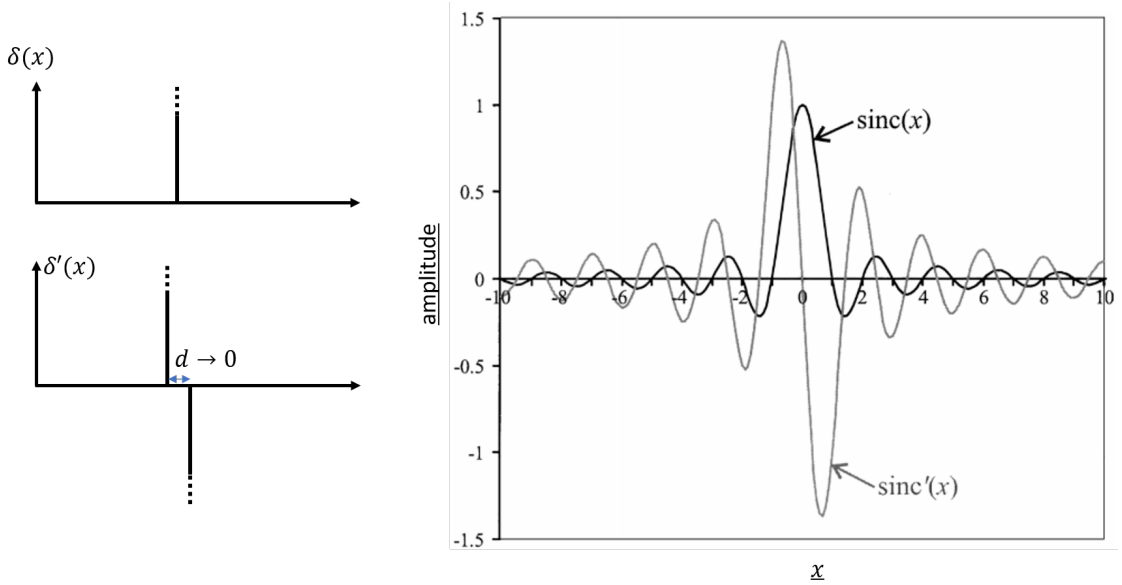


Figure 5.5: Schematic of the Dirac delta and its derivative (left, top and bottom) and their discrete versions – sinc function and its derivative (right, from Hicks, 2002).

interpolation:

$$\psi'(z) = \sum w_i'(z)\psi_i. \tag{5.56}$$

This method solves two problems at once – it allows to locate stations between grid nodes and allows for dipole radiation patterns. The question remains, what weighting function $w_i(z)$ to choose. If we were to simulate receivers only, the choice could be arbitrary. But the same weights need to describe sources, to honour the reciprocity and this imposes additional constraints on the wavenumber spectrum. A monopole source should have a spectrum of a *discrete Dirac delta*, i.e. be equal to one below the Nyquist wavenumber and zero above it. Implementation of such sources and receivers is described below.

5.3.2 Discrete delta

An ideal discrete delta is given by the sinc function[13]:

$$\mathrm{sinc}(x) = \begin{cases} \frac{\sin(\pi x)}{\pi x}, & x \neq 0 \\ 1, & x = 0 \end{cases}. \tag{5.57}$$

It is symmetric and takes zero values for all non-zero integers. Neither of these properties holds for its derivative (fig. 6.15). When implementing (5.57), one needs to guard against dividing by small numbers.

5.3.3 Kaiser window

The sinc function simulates an ideal delta only when it is defined on an infinite domain. In practice, it depends on a finite number of nodes and thus has to be windowed. There are many windows to choose from to achieve this. Here band-limiting with the Kaiser window (Kaiser and Schafer, 1980) is

[13] The N-dimensional delta function is a product of N one-dimensional ones.

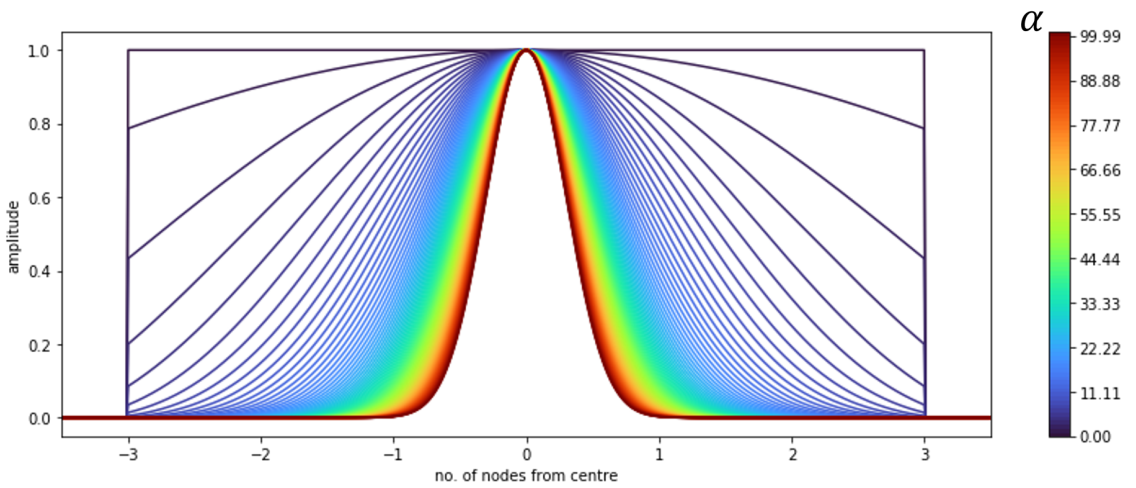


Figure 5.6: Kaiser windows of the radius r equal to 3 grid nodes and the shape parameter α ranging from 0 (blue) to 100 (red).

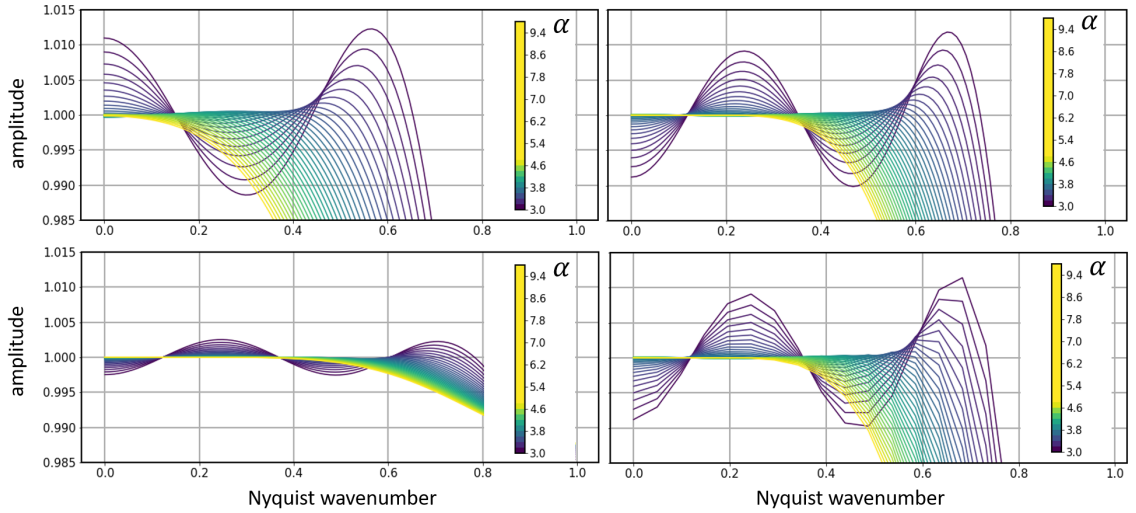


Figure 5.7: Wavenumber spectrum of the band-limited discrete Dirac delta defined as a Kaiser-windowed sinc function (a). Case (b) same as (a) but with a larger radius; (c) same as (a) but with smaller offset from grid η (d) same as (b) but for smaller number of samples. The colormap of the lines encodes a window shape parameter α. The units of the horizontal axis (wavenumber) are scaled so that the value 1 corresponds to the Nyquist wavenumber.

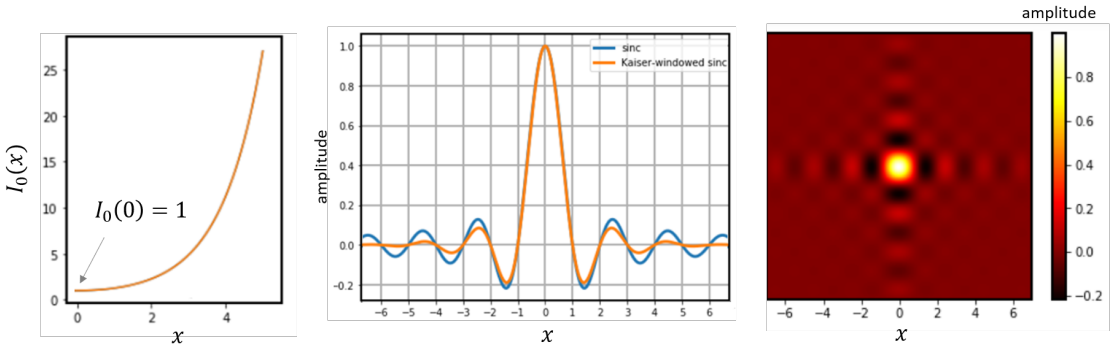


Figure 5.8: Sinc function tapered with the Kaiser window defined with a 0th-order modified Bessel function of the first kind (left) in 1d (middle) and 2d (right).

used after Hicks (2002) with properties discussed therein. It is defined as:

$$w(x) = \begin{cases} \frac{I_0\left(\alpha\sqrt{1-(x/r)^2}\right)}{I_0(\alpha)} & \text{if } |x| \leq r \\ 0 & \text{if } |\mathrm{x}| > \mathrm{r} \end{cases} \tag{5.58}$$

where I_0 is a 0th-order modified Bessel function of the first kind[14] (fig. 5.8). Its Fourier transform is given by (Kaiser and Schafer, 1980):

$$W(x) = \begin{cases} 2r\sinh(\alpha\sqrt{1-(k/k_\alpha)^2})\left[-I_0(\alpha)\cdot\alpha\sqrt{1-(k/k_\alpha)^2}\right]^{-1} & \text{if } |k| \leq k_\alpha \\ 2r\sin(\alpha\sqrt{(k/k_\alpha)^2-1})\left[I(\alpha)\cdot\alpha\sqrt{(k/k_\alpha)^2-1}\right]^{-1} & \text{if } |\mathrm{k}| > \mathrm{k}_\alpha \end{cases}, \tag{5.59}$$

where $k_\alpha = \alpha/r$. Parameter α (or b as denoted by Hicks (2002)) measures a width of the main lobe of the window in the time/wavenumber domain. The half width of the main lobe is given by:

$$R_1 = \frac{1}{2\pi r}\sqrt{\pi^2+\alpha^2}. \tag{5.60}$$

The larger the α, the narrower the window in the time/space domain. The Kaiser window is a simple approximation to the Slepian window[15] which has main lobe has a maximum possible energy. Hence, it is referred to as near-optimal. For $\alpha \to 0$ and $\alpha \to \infty$ the Kaiser window converges to the rectangular function and Dirac delta (like a increasingly narrow Gaussian function) respectively (fig. 5.6) Note the sharp edges for small α. Fig. 5.7 is an extension of fig. 3 by Hicks (2002). Its analysis allows to conclude that:

- a larger number of samples $2r$ smooths the spectrum $\Delta f \sim N$ (because $f_{\max}$ depends on Δt, and the number of samples is the same in both domains) but has no other effect;
- the offset from grid η affects mainly amplitude variations: the smaller the offset, the smaller the variations; the wavenumber content is barely influenced. In fact, it does tend to maximum content for $\eta \to 0$ but it does not drop for the maximum value $\eta = 0.5$, as much as for $r = 1$;
- the window size r affects mainly the wavenumber content, the larger the window, the more wavenumbers are present (up to Nyquist for the limit case); it makes sense, in particular infinite window preserves all frequencies; in general windowing in time domain convolves the spectrum with $\mathrm{sinc}(rk)$, so the larger r the more delta-like the sinc is. The amplitude variations are barely influenced;
- window shape parameter α affects both: the smaller value of η, the larger the wavenumber content (desired) and variations (undesired); and conversely – deviations of 0.5 from the optimal value seem to make no much difference.

5.3.4 Dipole radiation

Both, the continuous and the discrete delta function have the property:

$$\int_{-\infty}^{+\infty} f(z)\delta'(z)\mathrm{d}z = -f'(0), \tag{5.61}$$

i.e. a zero-lag cross-correlation with the derivative of the delta gives a *minus* derivative of the signal. This means that the v_z receiver can be simulated by a *positive* delta derivative. The source weights need

[14] Modified Bessel functions of the first and second kind are two linearly independent solutions to the modified Bessel's equation $x^2y'' + xy' - (x^2+\nu^2)y = 0$; they are exponentially growing and decaying functions, respectively (Riley et al., 2006).

[15] Also known as a *discrete prolate spheroidal sequence*, or DSSP for short.

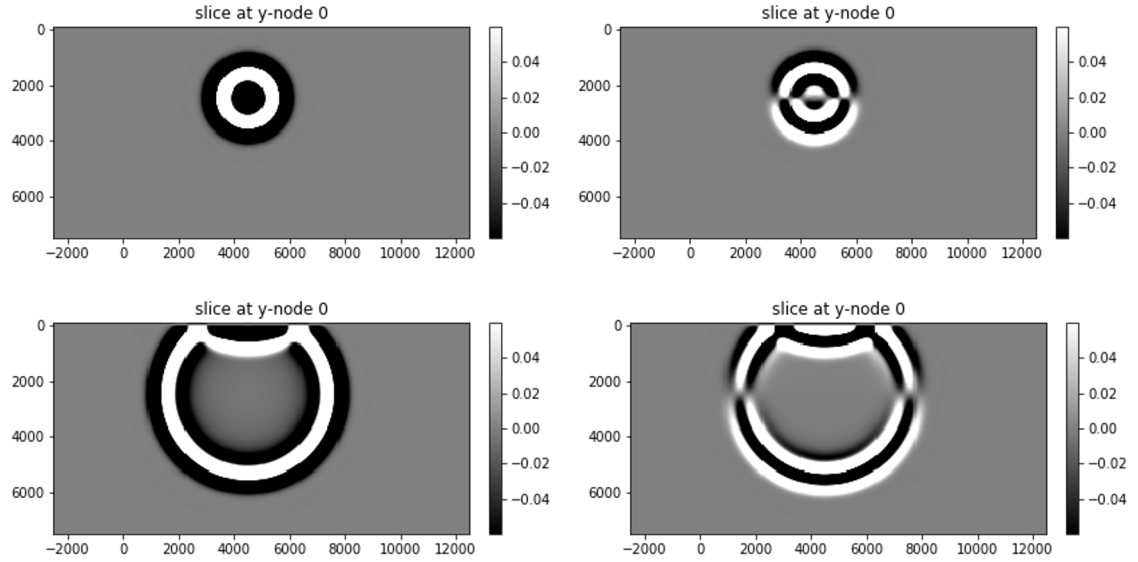


Figure 5.9: Radiation from a monopole (left) and a dipole source (right) shown as 2d wavefield propagated in homogeneous model with a free surface at the top. Rows show different time-snapshots. Axes are scaled in metres.

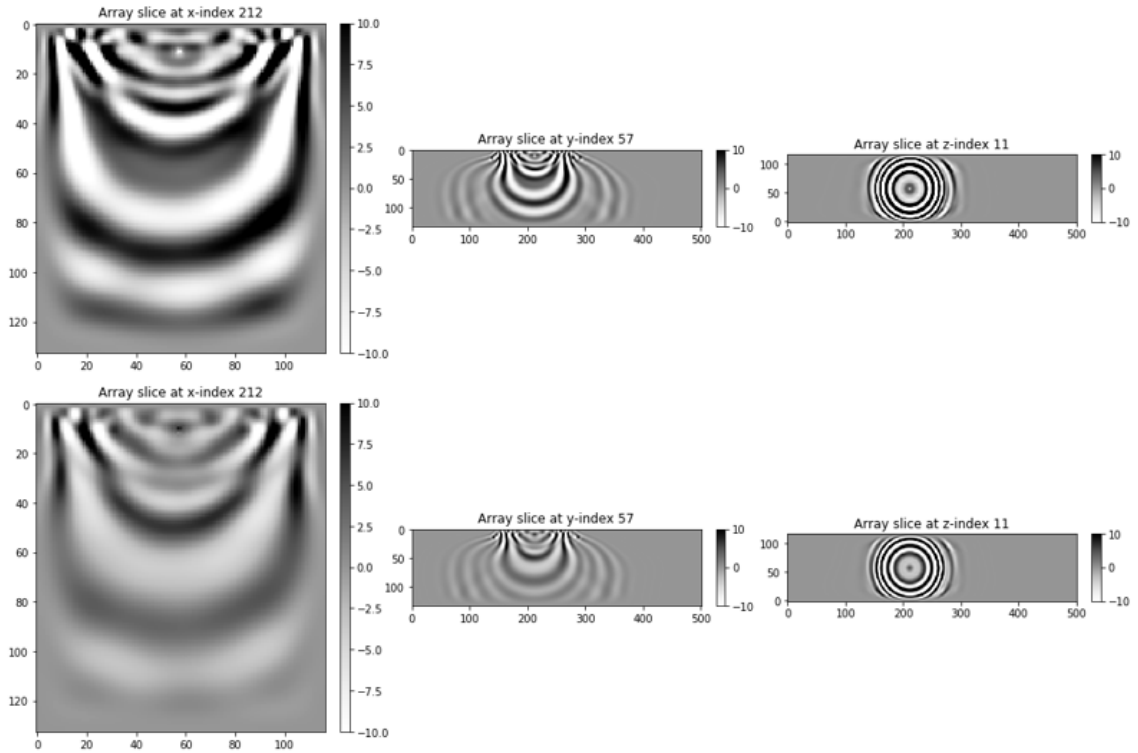


Figure 5.10: Radiation from a monopole (top) and a dipole source (bottom) shown as x/y/z slices (from left to right) of 3d wavefield propagated from the dual-sensor OBS 166 in a realistic model of Christiana basin near Santorini. Axes scaled in grid nodes ($h = 50$ m).

Algorithm 5.1 Quality control of the implemented particle-velocity simulator.

1. Implement geophone (particle-velocity) receivers for the second-order acoustic wave equation.
2. Check correctness of synthetic particle-velocity data generated by an airgun source against (a) an exact solution and/or (b) an output of a numerical solver of the first-order wave equation (here `SOFI3D` was used).
3. Design a reciprocal experiment with a single-force (dipole) source at the location of the previous receiver, and a hydrophone receiver at the location of the previous source. The data should be identical.
4. Use the configuration from the previous step to simulate real, field geophone data generated with an airgun source, taking advantage of the reciprocity.

to be oriented in the exaactly same way, i.e. the dipole source will radiate positive polarity upwards (fig. 5.9). The derivative of sinc is given by:

$$\frac{\partial \mathrm{sinc}(x)}{\partial x} = \frac{\cos(\pi x) - \mathrm{sinc}(x)}{x} \tag{5.62}$$

Note that for a dipole an optimal windowing parameter α is different from the one for a monopole. E.g. 4.40 versus 4.14 for a window half-width $r = 4$ as calculated by Hicks (2002). Fig. 5.9 and 5.10 show the difference in wavefield between a monopole and a dipole source for a homogenous and realistic model, respectively.

5.3.5 Time-integration

The time-integration needed in $\diamondsuit$ can be performed at any stage of the simulation since it is not sensitive to the direction of propagation. Here it is chosen to time-integrate the source wavelet.

5.3.6 Normalisation

In this study, the data are normalised so that scaling stations with multiplicative constants such as ρ^{-1} is not needed. If the true amplitudes were simulated, the density-factor would have to be included.

5.4 Numerical tests

First a full analysis is presented for a 1d case which is easy to implement. In 2d and 3d, a third-party code `SOFI3D` (Bohlen et al., 2015) is used as a benchmark for pressure-velocity conversion. The reciprocity is checked in a self-consistent manner using `FULLWAVE3D`.

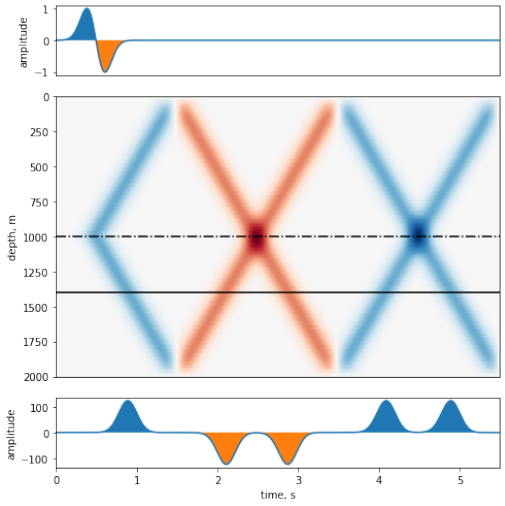
5.4.1 Simple 1d test

In this section, intuition behind reciprocity principle is presented with a one-dimensional example for which the entire space-time-dependent wavefield can be visualised in a single plot, instead of just temporal or spatial slices. Such exposition of reciprocity is novel[16] and is hoped to demistify this otherwise difficult to grasp concept. Additional benefit is that a 1d solver is straighforward to implement from scratch. The code for a constant-density finite-difference wave equation ($\mathcal{O}(h^2)$-accurate in both t and h):

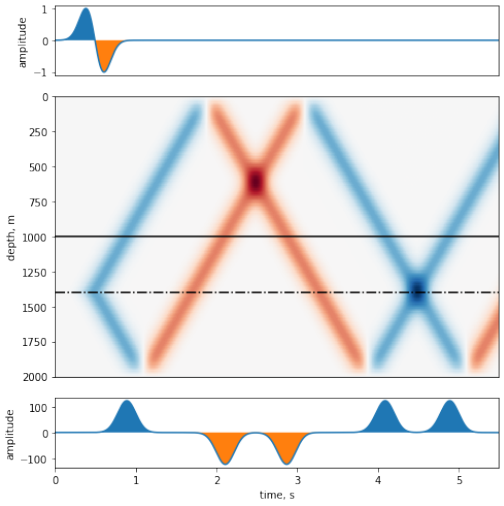
$$u_\xi^{\tau+1} \approx \delta t^2 c_\xi^2 \left(\frac{u_{\xi+1}^\tau - 2u_\xi^\tau + u_{\xi-1}^\tau}{\delta x^2} \right) + 2u_\xi^\tau - u_\xi^{\tau-1}. \tag{5.63}$$

is provided in list. 5.1.

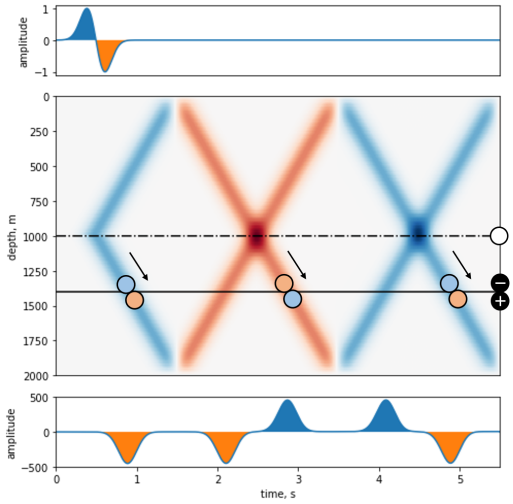
[16] It was inspired by plots of Slob et al. (2020).



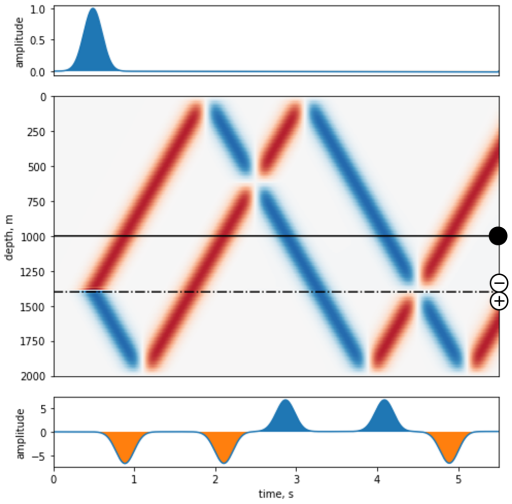
(a) Monopole source, monopole receiver.



(b) Same as 5.11a, but with source and receiver positions interchanged. Note that the two wavefields differ everywhere except at the receivers.



(c) Monopole source, dipole receiver. Note the wavefield is the same as in 5.11a except the recorded trace.



(d) Same as 5.11c, but with source and receiver positions interchanged. Note the time-integrated source wavelet.

Figure 5.11: Analysis of pressure-velocity reciprocity in 1d. Cases on the right and left are reciprocal. Each fig. shows source wavelet and recorded trace above/below the wavefield, respectively. Positive amplitude of traces/wavefield in blue. Dashed/solid black line – source/receiver location. Radiation-pattern symbols are explained in fig. 5.4. Constant v_P of 1 km/s with free surface at 0 and 2 km depth was used. Solver-code is supplied in list. 5.1.

Listing 5.1: Simple Python solver of the second-order 1d acoustic wave equation with first-order free-surface boundaries and first-order quiescent initial conditions. It uses the second-order approximation of the second derivative (both in time and space). Definitions of the source and receiver operators are provided in the appendix. Note how free-surface boundary conditions are imposed implicitly by preventing the first and the last grid node from being updated.

```
nx, dx, nt, dt = 101, 50, 1001, 0.001   # domain
c = 3000 * np.ones(nx)                  # velocity model
A = (dt/dx)**2
p_prv, p_now, p_nxt = np.zeros((3,nx)) # zero init. conditions
for it in range(nt):                    # time-stepping
  src.inject(u_nxt, ti)
  for ix in range(1,nx-1):              # free-surf. bnd. conditions
    laplacian = p_now[ix-1] - 2 * p_now[ix] + p_now[ix+1]
    p_nxt[ix] = 2 * p_now[ix] - p_prv[ix] + A*laplacian*(c[ix]**2)
  p_prv[:], p_now[:] = p_now[:], p_nxt[:]
  rec.detect(p_now, ti)
```

The analytical solution $\psi(z,t)$ in a one-dimensional infinite homogenous model of wave velocity c is given by a temporal convolution:

$$\psi(z,t) = G(z,t) * s(t) \tag{5.64}$$

of the source wavelet $s(t)$ and the Green's function (Rienstra and Hirschberg, 2019):

$$G(z,t) = \frac{1}{2c} H\left(t - \frac{|z|}{c}\right) \tag{5.65}$$

where H is a Heaviside step function and z and t are, as usual, space and time coordinates, respectively. Convolution with a step function is equivalent to a time-integration (see fig. 5.23), so in order to record a certain waveform, the source wavelet needs to be its time-derivative. Here a source wavelet of a first derivative of a *positive* Gaussian is used (top-left trace of fig. 5.11).

The spatial domain corresponds to a depth, with the axis pointing down. It is defined on a strip:

$$0 \leq z \leq 2\,\text{km}$$

split into 201 grid nodes and bounded by a free surface at 0 and 2 km. Upon each reflection off the free-surface boundaries, the wavefield is expected to experience a 180° phase-change (polarity reversal). The polarity of the primary impulse (direct wave) will be referred to as *positive*, and the reversed as *negative*. The experiments below were designed so that the only change to wave amplitudes is through polarity reversals. It is apparent from eq. (5.65) that amplitude of a 1d-wave does not experience any decay due to the geometric spreading (Aki and Richards, 2009). There is no physical dispersion or mode conversions either. The numerical dispersion is neglible as the waves are well sampled.

Case 1: monopole source at A, monopole receiver at B Let us first consider a monopole source at depth 1 km (the center of the domain) and a monopole receiver at 1.4 km (fig. 5.11a). Based on the above remarks, we expect to observe two Gaussian pulses propagating symmetrically from the source, without any decay of amplitude. Waves reflected an even number of times has the same polarity as the primary impulse, and waves reflected an odd number of times, a reversed one. These two cases will be referred to as *positive* and *negative polarity*, respectively.

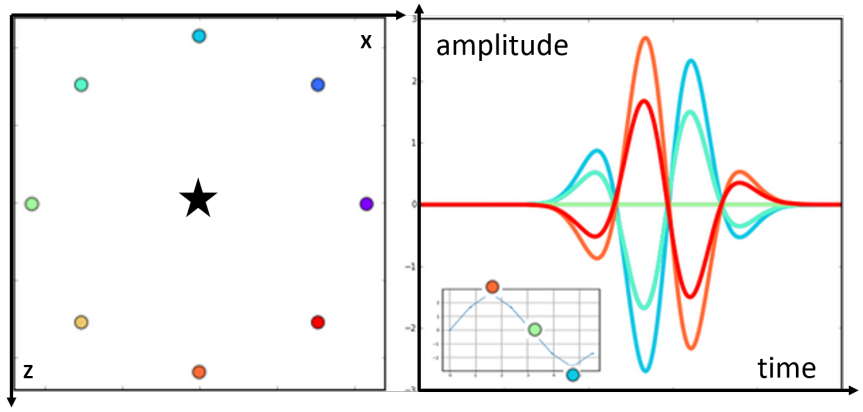


Figure 5.12: Sinusoidal amplitude versus angle dependence (inset) of the data generated with a monopole source and recorded with the circular array of v_z receivers shown on the left (black star and coloured points, respectively). Each observed trace on the right corresponds to the receiver of the same colour shown on the left. The source function was a Ricker wavelet and the observed data consist of its derivative (fig. 5.22) convolved with a 2d Green's function (fig. 5.23). The lightblue trace recording the upgoing wave has a negative polarity and the orange one recording the downgoing wave – positive. The amplitudes of recorded the traces match the analytical solutions.

Case 2: monopole receiver at A, monopole source at B This case is the same as the previous one, except the source weights are swapped with the receiver weights (fig. 5.11b). Both stations remain monopoles, so this corresponds to swapping their positions only. As the observed traces are identical for these two cases, the experiments are, by definition, reciprocal to each other. Note that the wavefields are different everywhere but at the receivers.

Case 3: monopole source at A, dipole receiver at B The next experiment (fig. 5.11c) is similar to the first one, except the receiver is now a dipole, not monopole and the recorded seismogram has been time-integrated to compensate for the effect of spatial differentiation on the amplitude spectrum (fig. 5.1). As the only change relative to the case 1 occurs upon detection at the receiver, the wavefield is exactly the same in both cases. The only difference occurs in the seismograms. Clearly, this experiment is neither reciprocal to case 1 nor 2. The dipole receiver flips the polarity of downgoing waves (the $-\int \mathrm{d}t\partial_z$ operator), e.g. the direct wave in this particular source-receiver configuration.

Case 4: monopole receiver at A, dipole source at B In this experiment (fig. 5.11d), the source and receiver have been interchanged relative to the previous case. Note that the source wavelet has been time-integrated in order for the wavefield to have the same wavenumber spectrum as in the other cases. If we only compared the observed traces, the time-integration might just as well be performed at the receiver. The source weights are distributed along the Z axis in the same order as the receiver's in the previous case, so the negative and positive pulse are radiated upwards and downwards, respectively. This experiment is reciprocal to the case 3, just like the case 2 was reciprocal to the case 1.

5.4.2 Particle-velocity test

The conclusions of the previous section were checked with the production-quality `FULLWAVE3D` code augmented with the `HicksPy` API intended for the inversion of realistic field datasets. Following alg. 5.1, the implementation was first tested in terms of the fidelity of the particle-velocity calculation (fig.

5.12). A monopole source placed in the middle of a 2d domain was surrounded by a circular array of vertical-dipole receivers. The result was compared to the one generated by the SOFI3D code that uses a first-order formulation of the acoustic wave equation (5.6) in which the particle-velocity wavefield is readily available. The SOFI3D code introduces an additional derivative as it adds the source wavelet to the stress tensor components which are then differentiated with respect to time. For FWI purposes, it is left to the inversion for the source-wavelet to correct for this effect (T. Bohlen, email communication 2020). To directly compare the two codes, no time-integration is involved in FULLWAVE3D in this experiment. The reader, however, should bear in mind that the real particle-velocity seismograms are the time-integrated versions of the ones presented in fig. 5.12.

The two traces for the same source-receiver configuration are identical between the two codes. The polarity of the downgoing wave was reversed with respect to the one that would be recorded with a monopole receiver and is consistent with fig. 5.11. As expected, the amplitude displayed an azimuthal dependence of the vertical component of partical-velocity ($\sin\alpha$, where α is the angle between the X axis $\mathbf{e}_x$ and the direction of propagation $\mathbf{k}$), with maximum sensitivity to the waves travelling along the dipole axis and no sensitivity in the orthogonal direction (fig. 5.12, right and inset).

The results of this numerical test show that the implementation correctly simulates particle-velocity data generated by an explosion. In the next section, a reciprocal experiment will be presented.

5.4.3 Reciprocity test

In many solvers like FULLWAVE3D, the computational cost of a simulation is directly proportional to the number of sources whereas the cost of extra receivers is neglible (even in the adjoint step, as data-residuals are backpropagated simultaneously from all the receivers). In experiments involving more sources than receivers, it is then desired to take advantage of the reciprocity of the wave equation and relabel sources as receivers and receivers as sources to speed up the simulation. In the previous section, correctness of simulating particle-velocity data was demonstrated, corrresponding to the case 3 of sect. 5.4.1. Here, the case 4, i.e. a reciprocal experiment is demonstrated in which pressure data are generated by a single-force (dipole) source.

Fig. 5.13 proves that these two source-receiver configurations give identical synthetic data. This confirms the conclusions of sect. 5.2.4, 5.2.5 and 5.4.1. It also implicitly shows that the FULLWAVE3D numerical solver has the self-adjoint property (sect. 5.2.5).

5.4.4 Inversion test

Once the method of simulating particle-velocity with or without reciprocity has been correctly implemented, one can investigate the potential differences when inverting either pressure or velocity data (or both data types simultaneously).

In fig. 5.14 the raw gradient prior to pre-conditioning, smoothing etc. (this may serve as a measure of sensitivity) was calculated for two source-receiver configurations, by backpropagating data residuals generated from a perturbed homogeneous model. Clearly, sensitivity of pressure and velocity data are different.

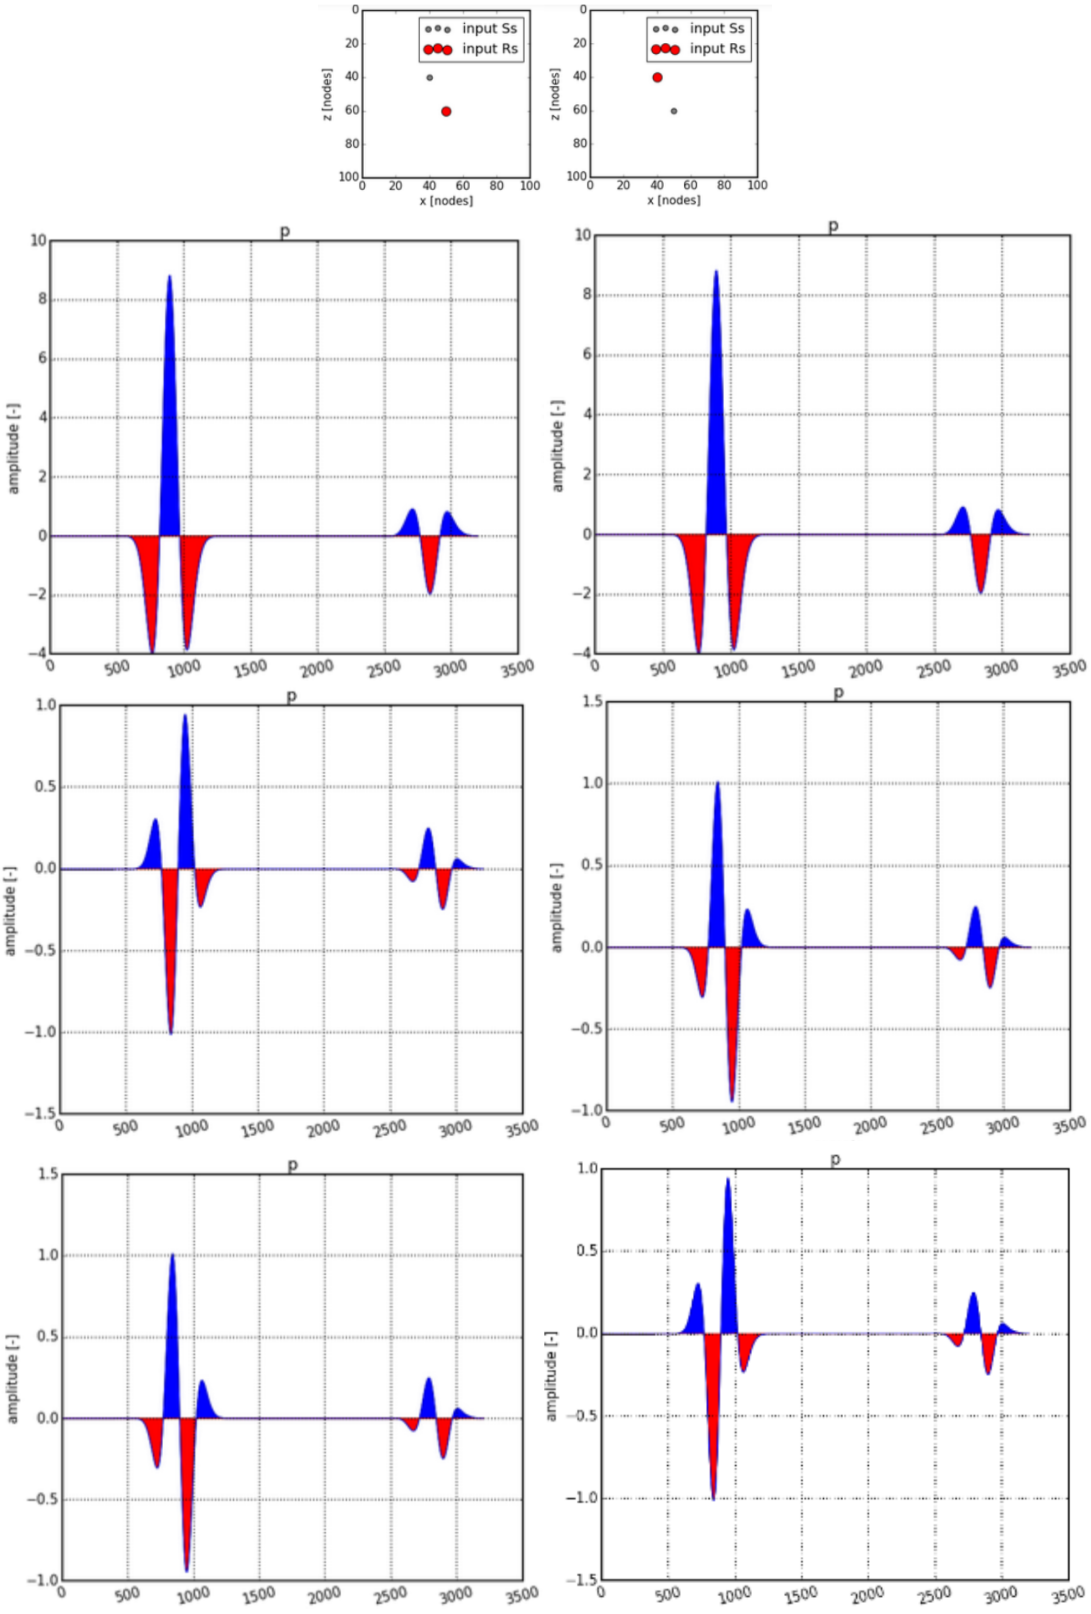


Figure 5.13: Reciprocity test in 2d: geometry (topmost), pressure source and pressure receiver (top), velocity source and pressure receiver (middle), pressure source and velocity receiver (bottom). X axis scale in time samples ($dt = 1$ ms).

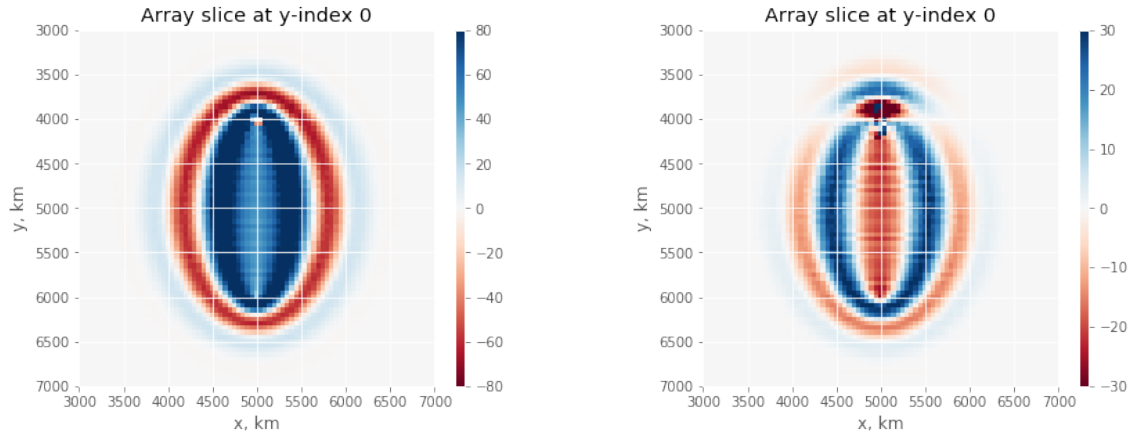


Figure 5.14: Raw gradient calculated for a monopole (left) and dipole (right) source and a monopole receiver in a homogeneous of 2 km/s and absorbing boundaries. Starting model was 10% faster than the true one and the data residuals were used as adjoint sources. Source function – a Ricker wavelet of 4 Hz peak frequency. Domain size was 10×10 km sampled at 50 m; simulated time was 2 s sampled with 2000 time-steps.

5.5 Field-data example

In this section it is shown that the conclusions from the previous sections hold for sparse, noisy field data from the PROTEUS experiment.

5.5.1 Field data

As described in ch. 3, in the PROTEUS experiment two physical quantities (data types / observables) were measured at the sea bottom – perturbation of the hydrostatic pressure and the velocity of the ground, both induced by the seismic energy generated by the airgun explosions. The particle-velocity data recorded by geophone sensors accounted for 75% offshore and 100% onshore data acquired in this experiment and should not be discarded.

The data quality is very variable as can be seen by comparing fig. 5.15 and 5.16. For some stations and shot lines it was the hydrophone component that displayed a better quality. This shows that multi-sensor data may be beneficial even for mono-sensor inversion allowing to choose the component with a higher S/N ratio.

5.5.2 Synthetic vs. field data

Let us compare synthetic data calculated from the starting model in fig. 5.17 with the field, pressure and vertical velocity data, applying the reciprocity principle in both cases.

Very similar pattern is obtained when comparing the observed vertical velocity data with synthetics calculated using time-integrated (fig. 5.18) dipole sources (fig. 5.19b). Observed data has similar S/N ratio to the previous case and the phase residuals are distributed almost identically, with similar RMS. This is what we expect, as both datasets constrain the same subsurface. Note that the phase of early arrivals at a hydrophone and a geophone is shifted about π radians relative to each other. In contrast, when using a monopole source to match the velocity data, one obtains a large modelling error and consequently severe cycle-skipping, as indicated qualitatively by the sharp red-to-blue transitions and quantitatively by a much larger RMS value (fig. 5.19c).

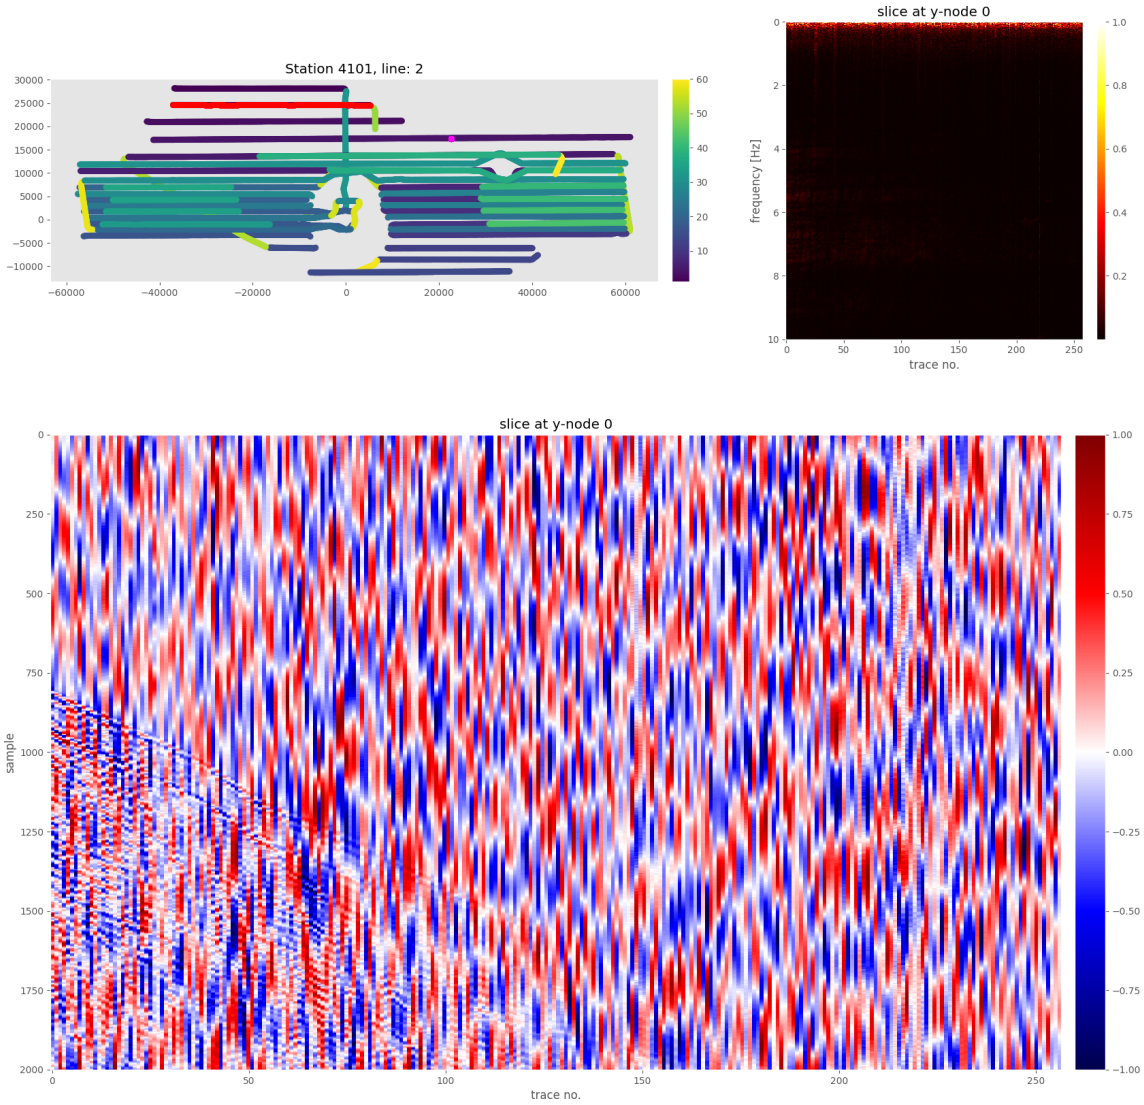


Figure 5.15: Hydrophone data recorded by OBS 101 (pink dot in the top-left figure) for the shot line highlighted in red (top-left figure, other colours denote the shot-line numbers). The low-frequency noise dominates the data as shown in both time and frequency domain (top-right and bottom figure, respectively).

5.5.3 Synthetic-data inversion

In order to get some insight into resolving power of vertical-component velocity compared to pressure data a synthetic-anomaly test was performed with the same set-up as in sect. 8.3.3. For both pressure and velocity, the data sparsity gives rise to spurious perturbations in the background model. However, the recovery is clearly different for pressure and velocity data. Pressure-derived results are characterised by the overshot anomaly and larger artifacts introduced in the anomaly-free background model, spurious low-velocity regions of amplitude up to 50% of the actual anomaly. Note that the amplitude of the anomaly is close to the true value for the velocity-data case. On the whole, the velocity-derived recovery is better, in terms of recovering both the geometry and amplitude of the anomaly. The differences in recovery suggest that the differeneces in models when inverting field pressure vs. velocity data, should not only be assigned to different S/N ratio of the two datasets.

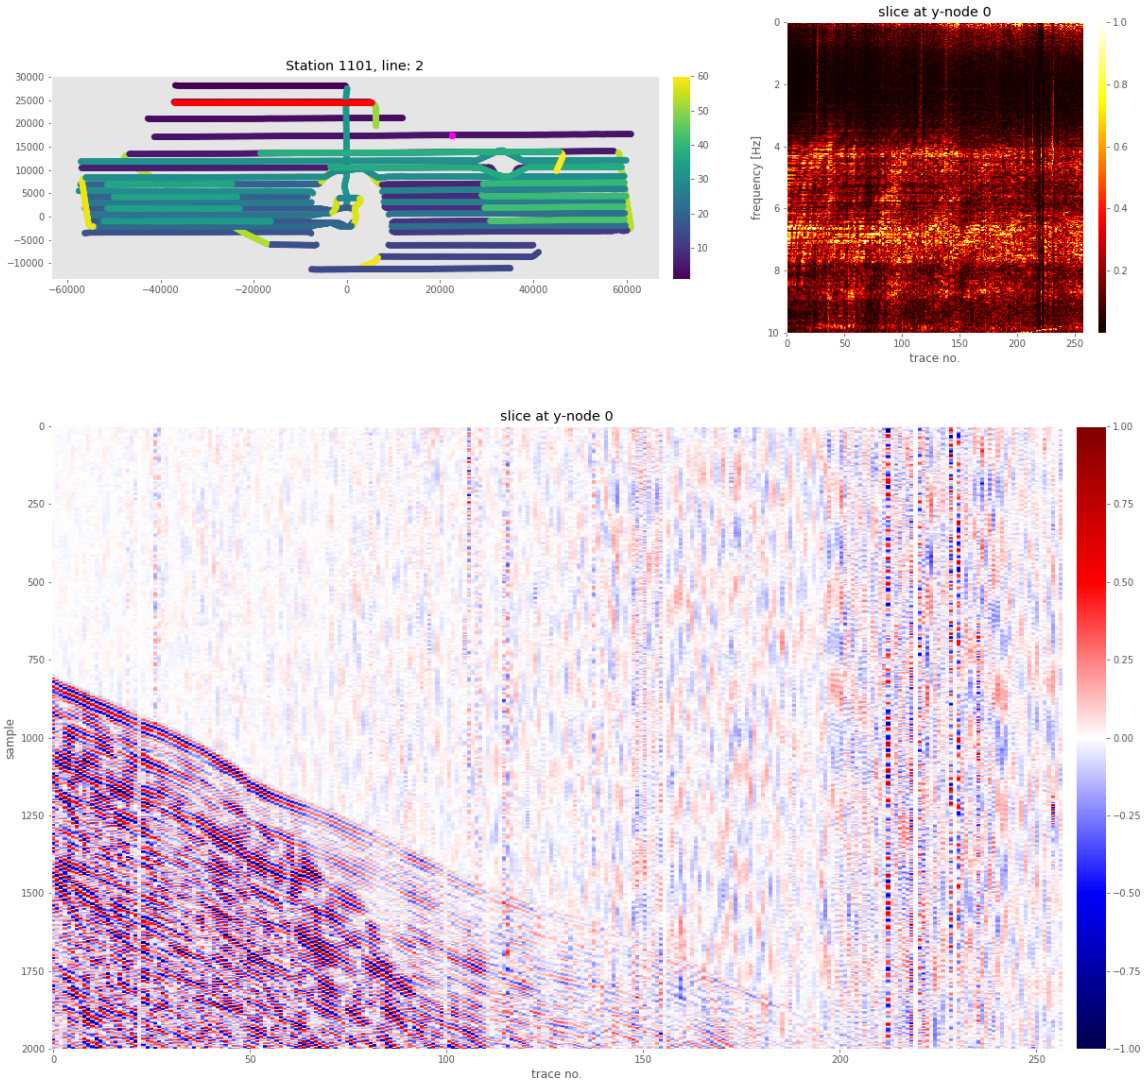


Figure 5.16: Geophone data recorded by the same OBS and shot line as in fig. 5.15, but with a much higher S/N ratio.

5.5.4 Field-data inversion

In particular, that for the velocity dataset, synthetic traces from the starting model match the observed ones with similar accuracy (up to the differences in the noise levels) as for the pressure dataset, wherever S/N ratio for these two datasets is comparable. Also, that the inversion of pressure data gives similar structural features in the final models as the velocity data, up to the different quality and the inherent difference due to different gradients (fig. 5.14).

The workflow from the previous section was used to invert the field pressure and velocity data – one data-type at the time. The quality-control of the data revealed that, while all hydrophones had acceptable S/N ratio, one geophone (station 137) had unusable data. Although it was included in the inversion, it did not contribute to the model update, effectively creating a gap in the acquisition coverage. This gap is visible in the results. Most structural features are similar (fig. 5.21), up to the differences explained above. The largest differences are in the bottom-right quadrant, where the inversion was adversely affected by the noisy station 137. This suggests, that the inversion of velocity data is correctly implemented and should be beneficial in cancelling the artifacts introduced by the inversion of the

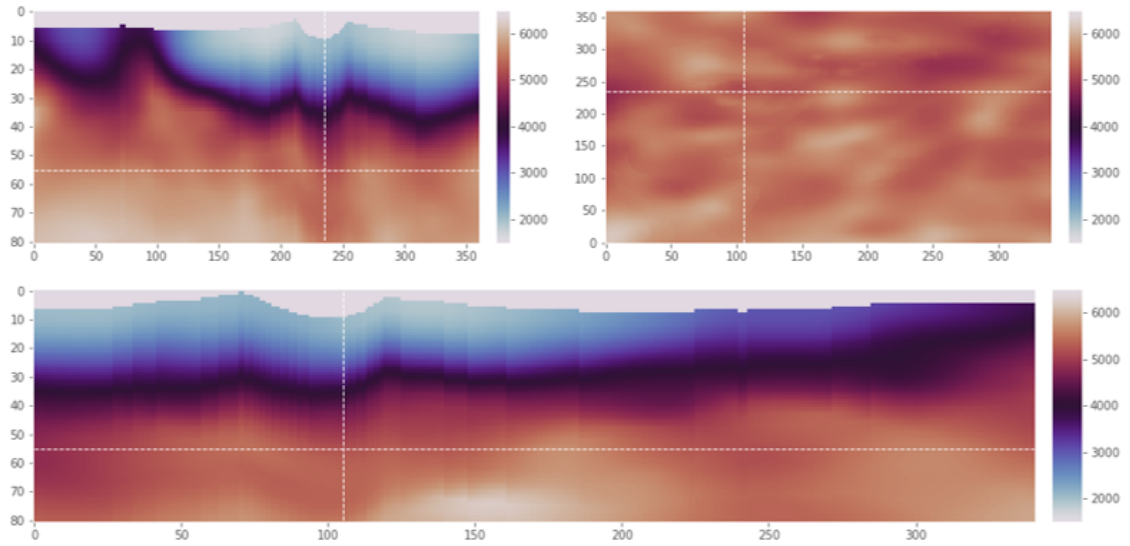


Figure 5.17: Starting v_P model (m/s) used in the field-data example. Top left – x slice, top right – z slice, bottom – y slice. Dashed white lines denote location of slices. Axes scaled in grid nodes ($dx = 50$ m), with variable aspect ratio.

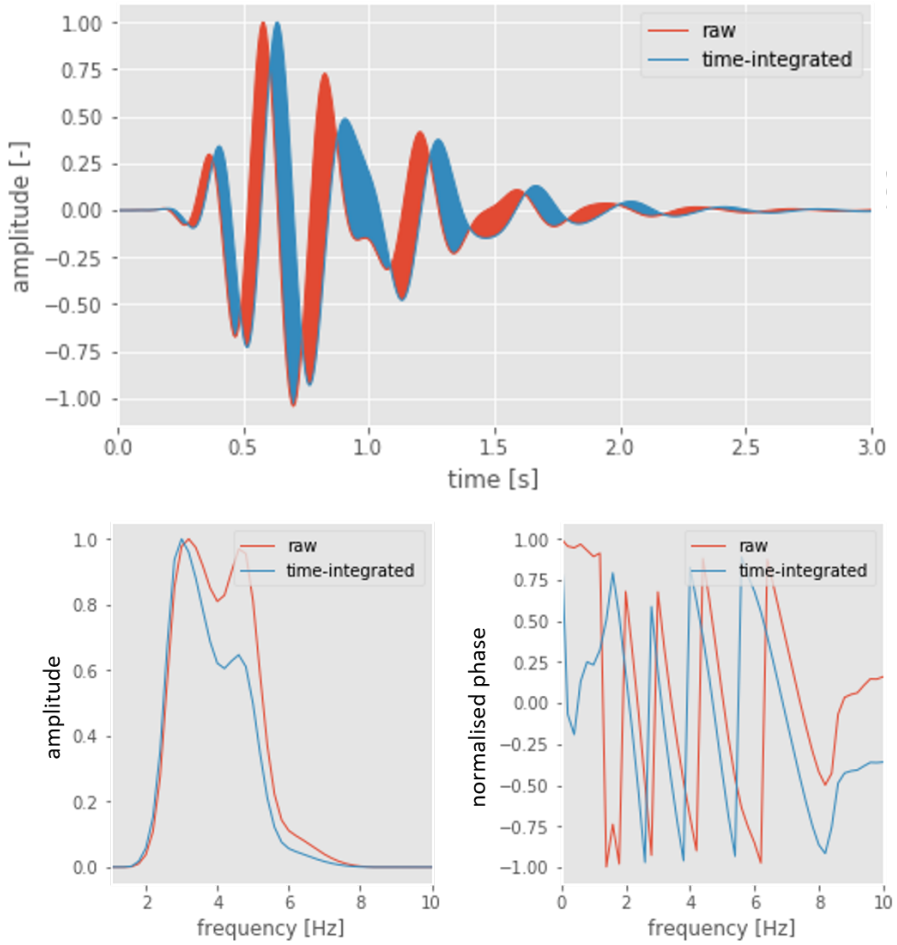
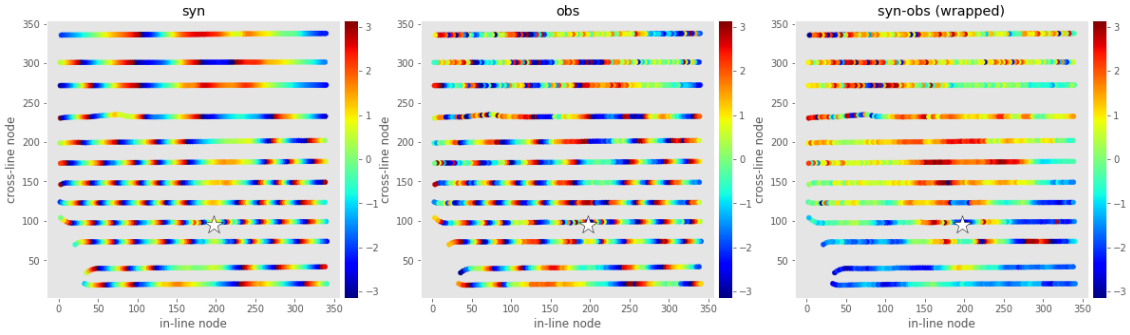
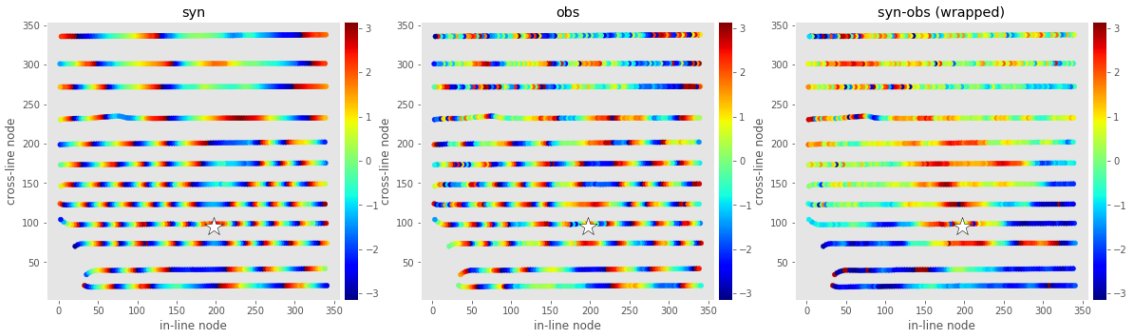


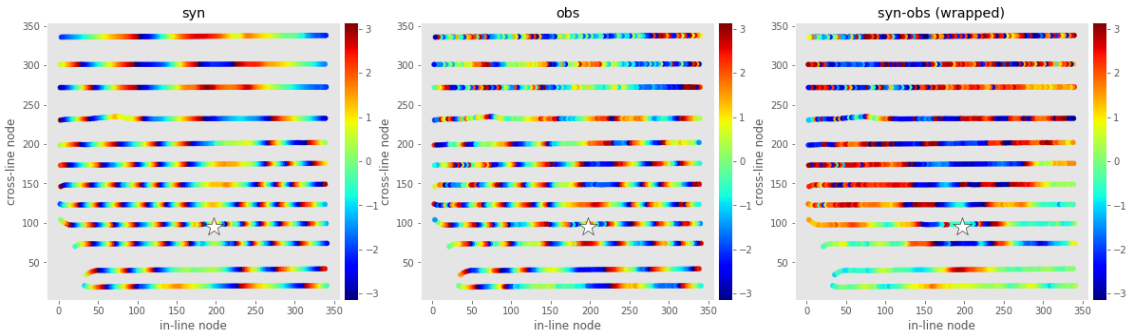
Figure 5.18: Source wavelet used in the field-data example, here before and after the time-integration required to simulate the particle-velocity data.



(a) *Pressure* synthetics (left) versus *pressure* observed data (middle); RMS of phase residuals (right) equal to 1.52 rad.

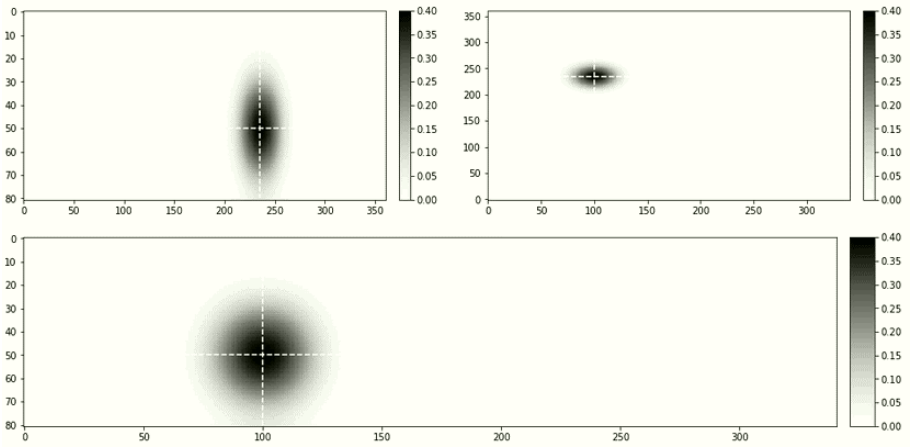


(b) *Velocity* synthetics (left) versus *velocity* observed data (middle); RMS of phase residuals (right) equal to 1.57 rad.

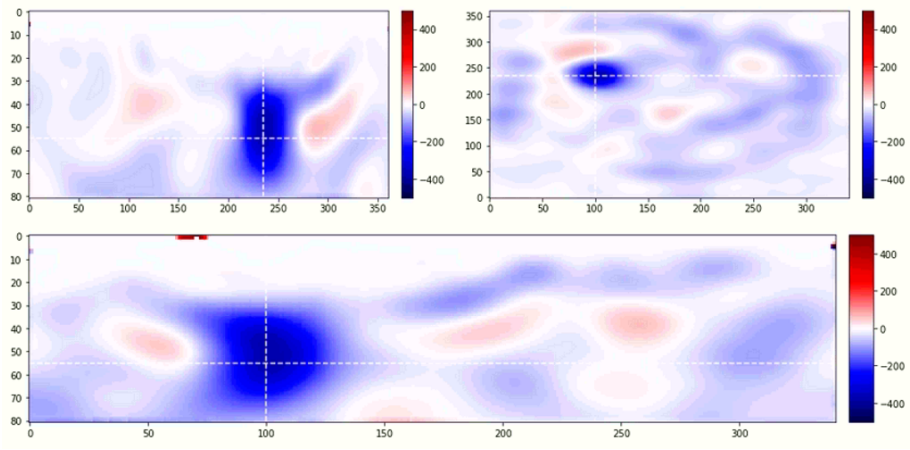


(c) *Pressure* synthetics (left) versus *velocity* observed data (middle); RMS of phase residuals (right) equal to 2.06 rad. Almost all the traces are cycle-skipped.

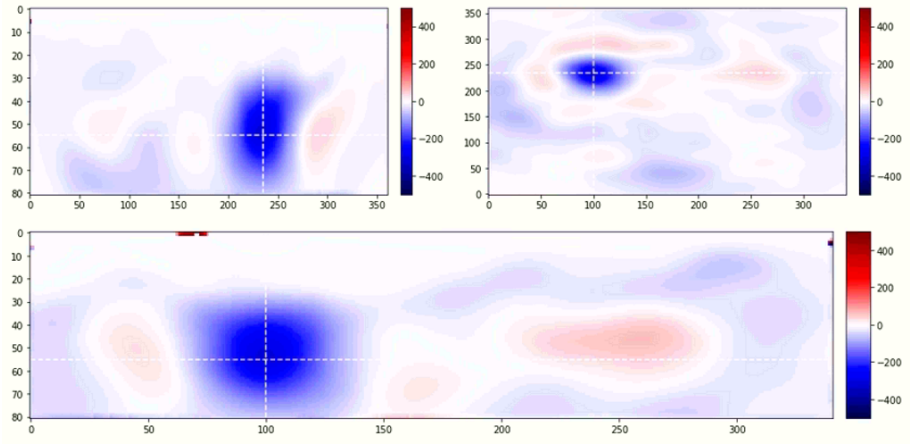
Figure 5.19: Phase analysis of pressure and vertical component of particle-velocity data for the OBS station no. 157. The white star and coloured points denote horizontal locations of the reciprocal source and receivers, respectively. The colourmap the phase of data windowed in frequency at 3 Hz and time (1 s around the first arrivals), wrapped into $[-\pi, \pi)$ range. Left, middle and right panels show synthetic ϕ_{syn}, observed ϕ_{obs} and the residual phase, respectively.



(a) Gaussian anomaly which was added to the starting model from fig. 5.17 according to eq. (8.2).



(b) Model recovered (final minus starting model) with pressure data.



(c) Recovery with velocity data.

Figure 5.20: Anomaly-recovery test using pressure vs. velocity data. The models recovered after 10 iterations at 3 Hz cut-off were calculated as the difference between the final and starting model (fig. 5.17, sliced in the same way).

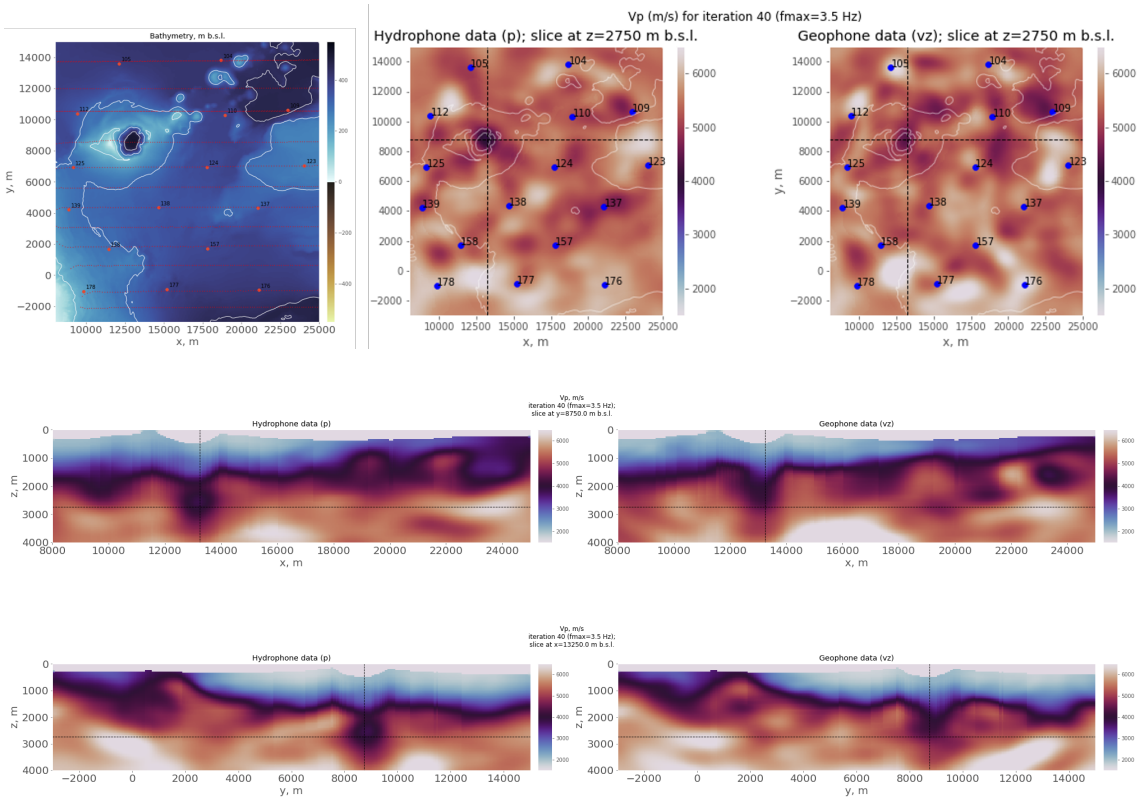


Figure 5.21: Final P-wave velocity model (m/s) at Kolumbo (slices along the dashed black lines at z, y and x, respectively from top to bottom) recovered after 40 iterations using pressure (left) and velocity (right) data. Top-left – acquisition geometry overlain on bathymetry.

pressure data alone.

5.6 Conclusion

Full-waveform inversion (FWI) of airgun-generated waves recorded by a sparse array of four-component (4C) ocean-bottom seismometers (OBS) far from land, using second-order acoustic isotropic wave equation and exploiting the reciprocity principle was presented. This work was aimed at inverting data from the PROTEUS experiment (Hooft et al., 2019) but its conclusions are applicable to many other case studies, especially sparse, longer-offset OBS datasets typically acquired in academic projects. The applications include (1) marine experiments with hydrophones and geophones, (2) land experiments with geophones; (3) marine-land settings such as Santorini (Hooft et al., 2017) or Hikurangi (Gray et al., 2019).

In this chapter, sources were implicitly assumed to be far from boundaries, In the next chapter, accurate simulation of reflections from irregular boundaries, such as the surface of the Santorini caldera will be discussed.

Supplementary information

A supplementary `Jupyter` notebook is available at: `http://github.com/kmch/PhD-Santorini-FWI/notebooks/`. See sect. 8.6 for more information.

Acknowledgements

The author wants to thank Michele Paulatto, Mike Warner and Adrian Umpleby for the idea of the method and important insights into the reciprocity principle. This works also benefited from discussions with Lluís Guasch, and Oscar Calderón Agudo.

Code availability

The code of the `HicksPy` library will be made available at *https://github.com/kmch/HicksPy* as soon as the FULLWAVE consortium gives a permission to do so. All contributions are looked forward to.

Appendix

5.A Locating stations

Before considering a dipole radiation pattern, let us analyse a problem of locating a monopole station at an arbitrary position relative to the finite-difference grid, in particular between grid nodes. If the position of the source coincides with a grid node, its injection is trivial and simply consists in adding a source wavelet to the wavefield at this node at each time step. If the position falls between nodes, it must be somehow *extrapolated* (spread) onto the grid. Similarly, a receiver needs to perform *interpolation*.

Nearest node The simplest solution to this problem is the nearest-neighbour extrapolation. However, this may not be accurate enough, especially for coarse grids. The mislocation error δx is proportional to the grid sampling h:

$$\delta x \propto \mathcal{O}(h) \tag{5.66}$$

with the maximum error $\delta x_{\max}$ equal to $\frac{1}{2}h$ or $\frac{\sqrt{2}}{2}h$ or $\frac{\sqrt{3}}{2}h$ in 1d, 2d and 3d, respectively. The mislocation δx corresponds to a travel-time error:

$$\delta t = \frac{1}{c(x)}\delta x \tag{5.67}$$

where $c(x)$ is the average velocity between the true and the effective location. For example, using 200-m grid gives approximately a 67 ms travel-time error for a station immersed in sea water of 1.5 km/s P-wave velocity.

Finite extent Apart from the travel-time error (eq. 7.3), mispositioning may result in a strikingly different interference pattern of the ghost and primary impulse, especially for stations located close to the sea surface, which is often the case. The improvement of the nearest-neigbhour method can be achieved by representing a point source with a band-limited, discrete Dirac delta (Hicks, 2002). See the main text of sect. 5.3 for the details.

5.B Analytical solutions

Below Green's functions G of a scalar 2nd order wave equation in 1d, 2d and 3d are provided. For the derivation, see e.g. (Rienstra and Hirschberg, 2019). G encodes the information about the differential operator and boundary conditions of a linear differential equation. By definition, response to any source function is a convolution with such Green's function. Here the first derivative of a Gaussian is used. The Green's function in 1d is a step function, which is an integral of the Dirac delta:

$$G_{1\mathrm{d}} = \frac{1}{2c} H\left(t - \frac{|r|}{c}\right). \tag{5.68}$$

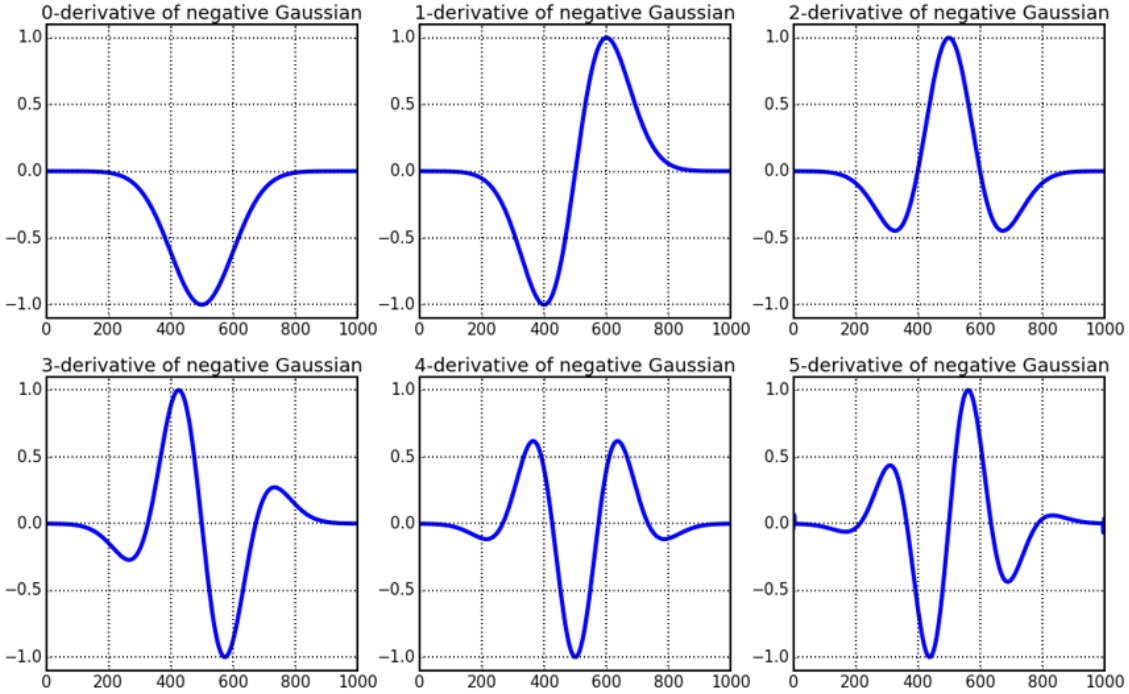


Figure 5.22: Gaussian function (negative) and its derivatives used to benchmark results of numerical tests (sect. 5.4). The second derivative is called the Ricker wavelet (e.g. Sheriff and Geldart, 1995).

Hence, the seismogram will be an integral of the source function. The situation in 2d is similar:

$$G_{\mathrm{2d}} = \frac{1}{2\pi c^2 \sqrt{t^2 - \frac{r^2}{c^2}}} H\left(t - \frac{|r|}{c}\right), \tag{5.69}$$

but not quite identical – the amplitude decays. This is because in 2d it is actually an infinite line source (Strang, 1986). In 3d, the Green's function reads:

$$G_{\mathrm{3d}} = \frac{1}{4\pi c^2 r} \delta\left(t - \frac{|r|}{c}\right). \tag{5.70}$$

which means the observed trace will have the same waveform as the source function up to the geometric-spreading factor $\frac{1}{4\pi c^2 r}$.

5.C Caldera phase-plots

A similar test to the one presented in fig. 5.19 was performed for a larger domain comprising the Santorini caldera (fig. 5.24). The same behaviour has been observed and the conclusions of the main text hold.

5.D Travel-time picking

The FWI posed as a local optimisation of a linearised l_2 data-misfit depends strongly on a quality of the starting model. This, in turn, relies on a correct identification ('picking') of certain seismic phases, if the starting model was obtained through first-arrival travel-time tomography. The picking can be challenging especially in the presence of shadow zones related to low-velocity anomalies, such as melt or volatile-rich zones.

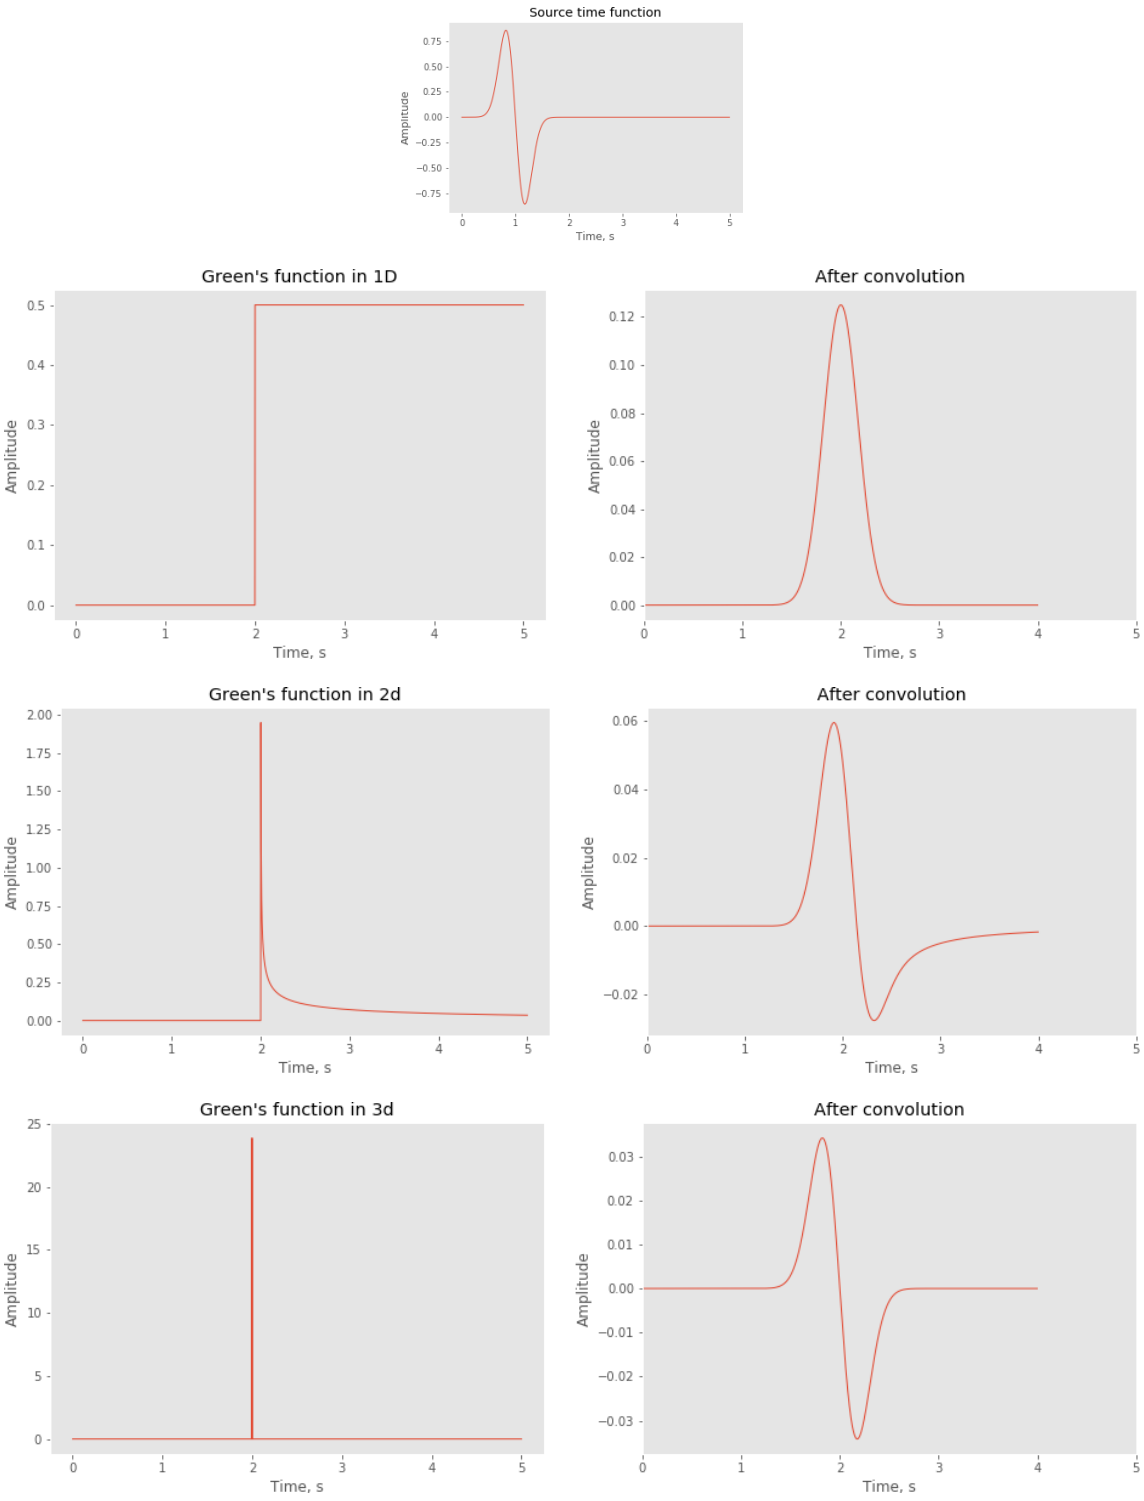
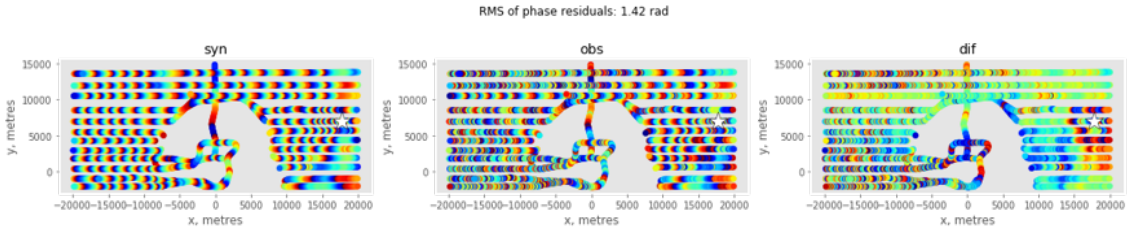
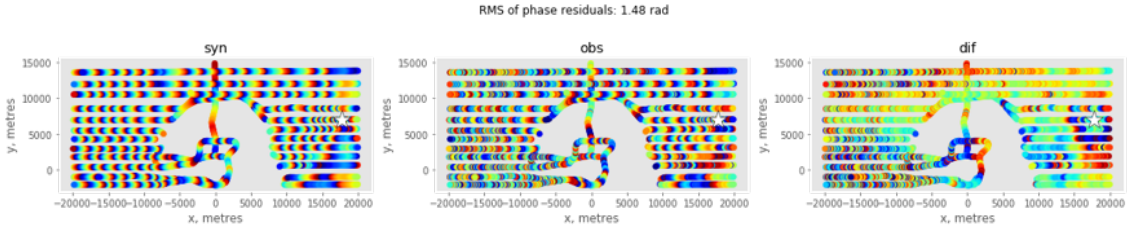


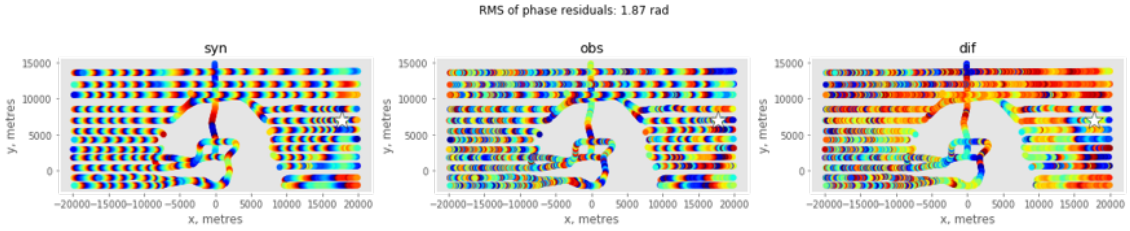
Figure 5.23: Analytical solutions of acoustic wave equation in homogeneous space. Green's functions (left panel) convolved with the source wavelet (inset) give the observed traces (right panel). The Green's function depends on the dimensionality of the problem. Only in 3d the observed trace has the same shape as the source wavelet. In 1d it is its time integral and in 2d it has a non-trivial tail. Calculated based on Igel (2016) after derivation provided by Rienstra and Hirschberg (2019).



(a) *Pressure* synthetics (left) versus *pressure* observed data (middle).



(b) *Velocity* synthetics (left) versus *velocity* observed data (middle).



(c) *Pressure* synthetics (left) versus *velocity* observed data (middle). Almost all the traces are cycle-skipped.

Figure 5.24: Phase analysis of pressure and vertical component of particle-velocity data for the OBS station no. 157. The white star and coloured points denote horizontal locations of the reciprocal source and receivers, respectively. The colourmap encodes the phase of data windowed in frequency at 3 Hz and time (1 s around the first arrivals), wrapped into $[-\pi, \pi)$ range. Left, middle and right panels show synthetic ϕ_{syn}, observed ϕ_{obs} and the residual phase, respectively.

Algorithm 5.2 Estimation of travel-time picking error using the reciprocity principle.

1. In the receiver gather A, check travel-time of picked phase for a shot localized at the receiver B; this is unlikely to be the same location so use two shots closest to receiver B instead.
2. Use coordinates of receiver B and shots to calculate distance r_1 and r_2 (fig. 5.27) and correct travel-times of both shots: $t_i = t_i - \Delta t_i$, where $\Delta t_i = \frac{r_i}{v_{\mathrm{water}}}$ and $v_{\mathrm{water}} = 1500$ m/s.
3. Take an arithmetic mean of both shots as a final estimate.
4. Repeat steps 1-3 with receivers A and B interchanged.
5. Calculate a difference of the average travel-times from step 3 in both cases (A and B, B and A). It gives an idea of the picking error and can serve as its lower-bound estimation.
6. Repeat for all receivers along the profile.

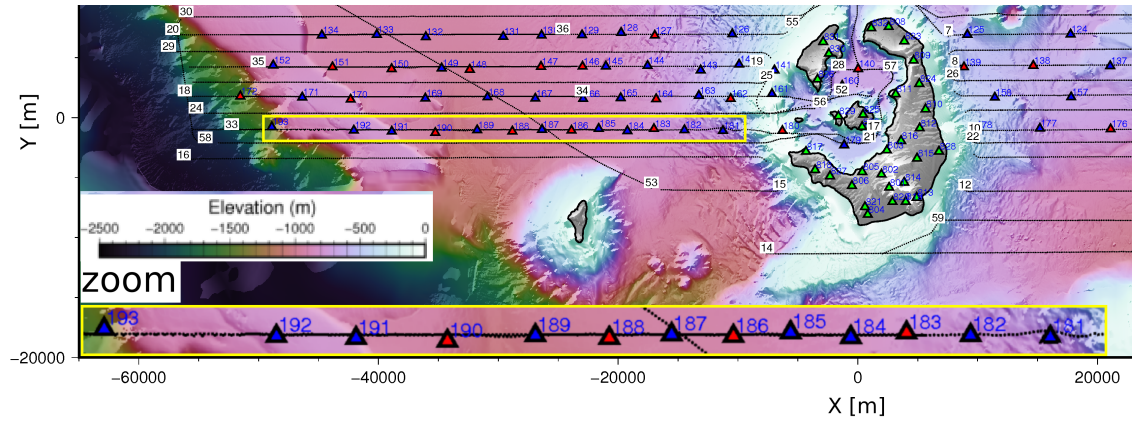


Figure 5.25: Acquisition geometry of the PROTEUS experiment with the profile in the Christiana sedimentary basin (zoomed in the yellow box) used in the picking-error estimation. Only stations 181, 183, 184, 185, 186, 188, 189, 190, 192, 193 were used.

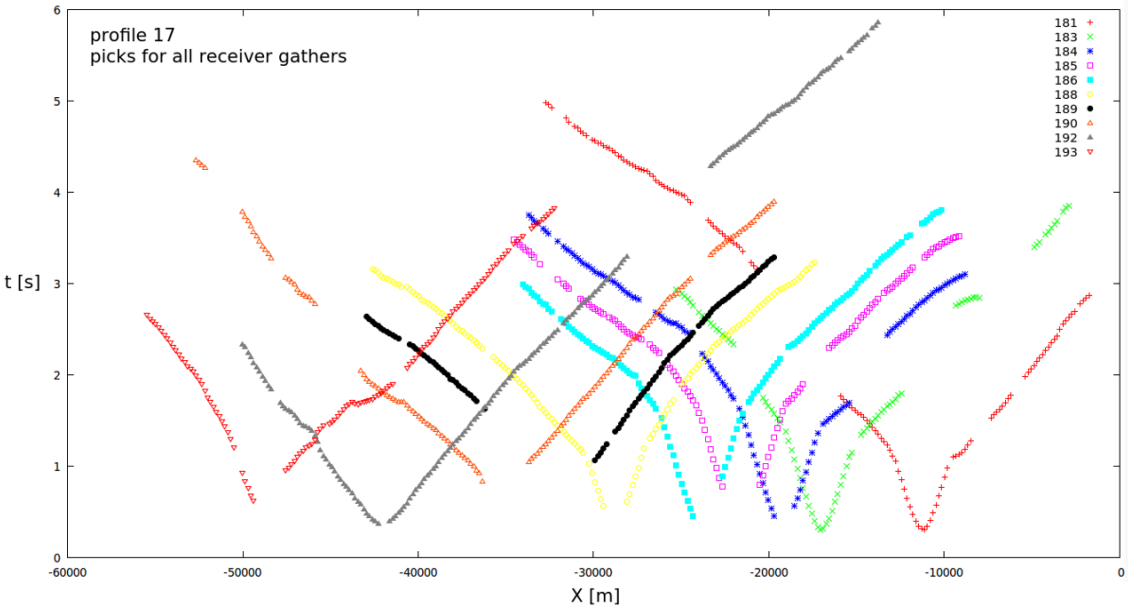


Figure 5.26: First arrivals picked for the stations in fig. 5.25. Travel-time of the arrival at station A generated by the source at station B should remain the same after interchanging A and B.

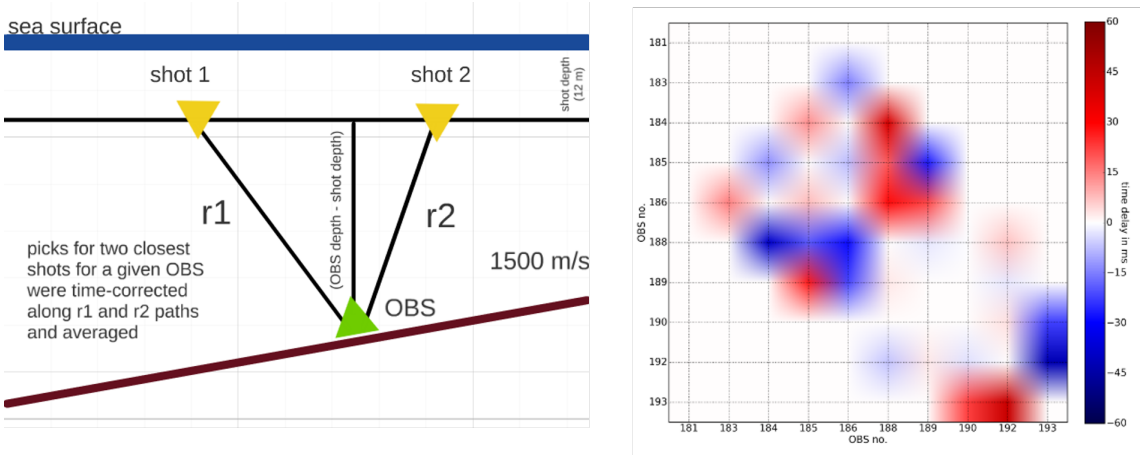


Figure 5.27: Left – schematic plot of shot and OBS location used for a geometrical correction of the travel-times in the reciprocity check. Right – reciprocity matrix with rows and columns representing respective receivers. Value at row i and column j stands for a delay of receiver i with respect to receiver j. Note that the matrix is by definition skew-symmetric.

The reciprocity principle can be used as a simple means to estimate uncertainty of picking travel-times used in a tomographic inversion for the starting P-wave velocity model. Although picking errors do not map into model artefacts in a simple way, arrival picks need to be re-assessed wherever there is a significant discrepancy. An example of an iterative workflow of alternating picking, travel-time tomography and FWI is presented by Górszczyk et al. (2017). Below, a reciprocity method allowing verification of the picking-correctness is presented. A general workflow of a reciprocity check along a seismic profile is summarised in alg. 5.2. Note that in many experiments shots are not fired exactly at receivers locations and some corrections must be made to take it into account (fig. 5.27, left). In principle only the nearest shot could be used, nonetheless a pair of surrounding shots was analysed here to possibly reduce uncertainty of the estimation (advantage of this approach remains to be proved). Note also that real water velocity might have differed slightly from the one used here. This should not affect the accuracy of the result though, since it was used consistently for both receivers of every pair.

Analysis described above has been performed for one of the profiles within the Christiana sedimentary basin containing OBS stations of numbers: 181, 183, 184, 185, 186, 188, 189, 190, 192, 193 (fig. 5.25). This fully marine profile characterised by the most dense array of stations was a natural candidate to be included in first FWI runs. The travel-times are shown in fig. 5.26. Results are presented in a matrix form (fig. 5.27, right) and indicate discrepancy between picks of at most 30 ms for some receiver pairs. This is in agreement with visually-assigned one-sigma errors of Hooft et al. (2019). The question remains how this inconsistency maps into a model error. Even without quantitative estimation of model uncertainty, these results highlight a subset of picked phases that may potentially need to be reassesed in order to improve performance of the tomography and subsequently the starting model for the FWI. Planned extensions of above analysis include better justification of used time-corrections as well as application to other profiles.

6

Modelling irregular boundaries

THE PROTEUS experiment around Santorini included marine shots recorded on ocean bottom and land seismometers. In this chapter, the immersed boundary method is used to allow land topography to be incorporated in velocity models. The method is formulated by imposing an irregular boundary condition of a PDE in the finite-difference (FD) scheme.

Important novel features of the presented method include an efficient 3d orthogonal point-projection and accurate injection of seismic energy at the free surface. Additionally, a variation of the vacuum method which does not require modifying the model is presented – it may replace IB whenever accuracy is not a critical issue. Numerical tests and a proof-of-concept example using 3d field-data from the PROTEUS experiment confirms the superior performance of IB over the air-layer method.

The method was implemented as a `Fortran90` package `IMMERSE` consisting of a stand-alone precalculator of wavefield-extrapolation weights and a thin interface for third-party FWI codes.

Key words: land acquisition · irregular topography · free surface · Dirichlet boundary conditions · acoustic wave equation · finite differences · staircase error · immersed-boundary method

6.1 Introduction

6.1.1 Forward problem

An accurate solution of the forward problem is a prerequisite of FWI. As FWI relies on comparing synthetic and observed data wiggle-for-wiggle, the former needs to simulate not less nor more than all physical phenomena recorded in the latter (no false positives or negatives). The first requirement, i.e. modelling the full physics of the problem in a computationally feasible manner is a subject of extensive research (e.g. Guasch et al., 2010; Alkhalifah and Plessix, 2014; da Silva and Warner, 2017) and remains one of the biggest challenges of the method. The second has to do with spurious effects introduced due to discretisation of a continuous problem. These errors can be grouped into two categories discussed below.

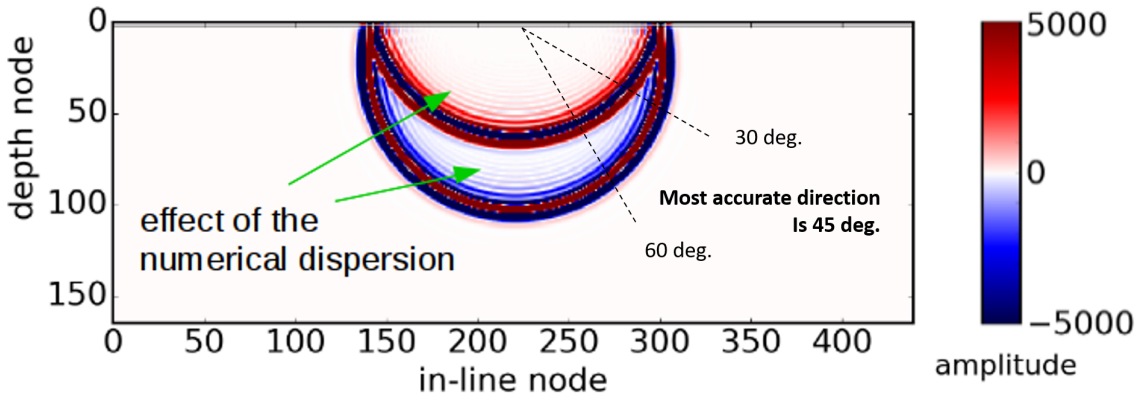


Figure 6.1: Example of a pressure wavefield affected by numerical dispersion. Note that the error amplitude varies like $|\cos(2\theta)|$ where θ is the inclination from the grid. The vertical and horizontal directions suffer from the largest error, while the diagonal directions ($\theta = 45°$) are least affected. Parameters used in the simulation: $v_P = 1500$ m/s, $\Delta x = 50$ m, Ricker source with $f_{\text{peak}} = 3.5$ Hz containing some energy above 6 Hz; the spatial-sampling criterion is violated since the scheme requires at least 5 nodes per wavelength.

Numerical dispersion The first category comprises a *numerical dispersion*. A phase velocity of the wave is incorrect if spatially under-sampled, i.e. if there are too few nodes per wavelength. This leads to nonphysical dispersion and thus undesired spreading of the wavepackage (fig. 6.1) – the discrepancy that increases proportionally to the propagation distance. The error is anisotropic with most accurate direction being 45° to the grid. To mitigiate this error a spatial *sampling criterion* specific to each numerical scheme must be met and/or the propagation distance should be sufficiently small. This class of errors can be successfully diminished by meeting the spatial-sampling criteria imposing a lower limit for a number of grid nodes per wavelength. This number is specific for a scheme used. It can be lowered by a design of the stencil so that it minimises the dispersion in a least-squares sense.

Numerical scattering The second category is concerned with misalignment of a grid and interface/boundary of a model. Such situations lead to a staircase approximation of smooth surfaces and subsequently first-order $\mathcal{O}(\Delta x)$ errors related to artificial diffraction and wrong travel-times (Yao et al., 2017), insensitive to accuracy of the scheme used (van Vossen et al., 2002; Symes et al., 2008). The second group of errors still poses a challenge and will be discussed in detail below. This inherent difficulty of a regular-grid method such as FD remains one of its major challenges and can lead to substantial errors if not treated properly (Mulder and Huiskes, 2017).

6.1.2 Boundary conditions

Stress $\tau_{ij}(\mathbf{x})$ and (particle) velocity $v_i(\mathbf{x})$ are continuous functions of the position $\mathbf{x}$. If stress was not continuous, one would have to deal with infinite forces. If velocity was not continuous, neither would be displacement resulting in gaps between rock layers, or rock layers permeating each other (Sheriff and Geldart, 1995). Continuity enforces stress and velocity to tend to the same limiting value at any interface[1], regardless from which side the interface is approached. One can write down these continuity

[1] real, i.e. a constrast of material properties, or imaginary, if needed for a thought experiment

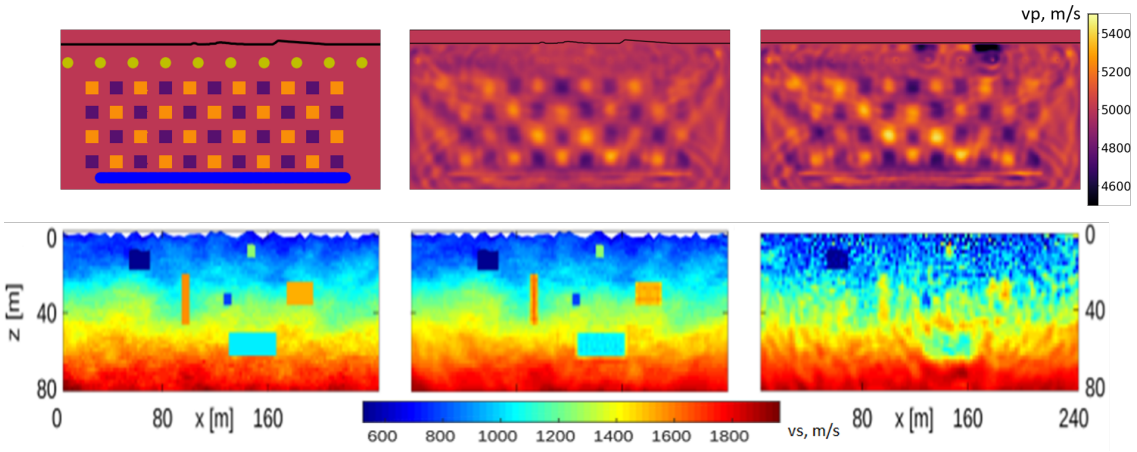


Figure 6.2: Effect of land topography on the full-waveform inversion of acoustic (top) and elastic (bottom) data. From left to right – true models, results of inversion using the actual and the flat free surface. Yellow and blue circles in the top-left figure denote sources and receivers respectively. Bottom row adapted from Nuber et al. (2016).

conditions as:

$$v_i^{(1)} = v_i^{(2)} \tag{6.1}$$

$$\tau_{ij} n_j^{(1)} = \tau_{ij} n_j^{(2)} \tag{6.2}$$

where the superscripts denote the medium on each side of the interface. The stress vector $\tau_{ij} n_j$ is sometimes called the *traction* and denoted as t_i^n. Equations (6.2) are general and hold for any medium. Two important special cases to consider are a liquid and a vacuum. In a liquid, shear stress is zero. Boundary of a vacuum is called a *free surface*, and is discussed below.

Free surface The free surface is an idealisation used in seismology to simulate a rock-air or water-air interface[2]. In this way, a model-parameter discontinuity is simulated implicitly through a condition imposed on the wavefield. The only kinematic boundary condition we can apply for the free surface is continuity of normal displacement, because cavitation does not occur (apparently changes of pressure are not that rapid) and diffusion of air into medium would require longer time scale. We cannot apply continuity of tangential displacement as the Earth's atmosphere can be treated as completely inviscid (Aki and Richards, 2009), enabling tangential components of the displacement to be discontinuous (slip parallel to the free surface). For these reasons, the free surface boundary condition is expressed by a dynamic boundary condition for a normal traction:

$$t_i := \tau_{ij} n_j = 0, \tag{6.3}$$

where τ_{ij} is the stress tensor and n_j is the jth component of the surface normal vector.

6.1.3 Free-surface effects

The shape of a free-surface has a significant effect on results of seismic wave modelling (e.g. Lombard and Piraux, 2004) and inversion (fig. 6.2). Bleibinhaus and Rondenay (2009) investigate the effects of

[2]This is justified as the air density is neglible compared to both solid rocks and sea water, which makes a perfect-reflector boundary condition appropriate.

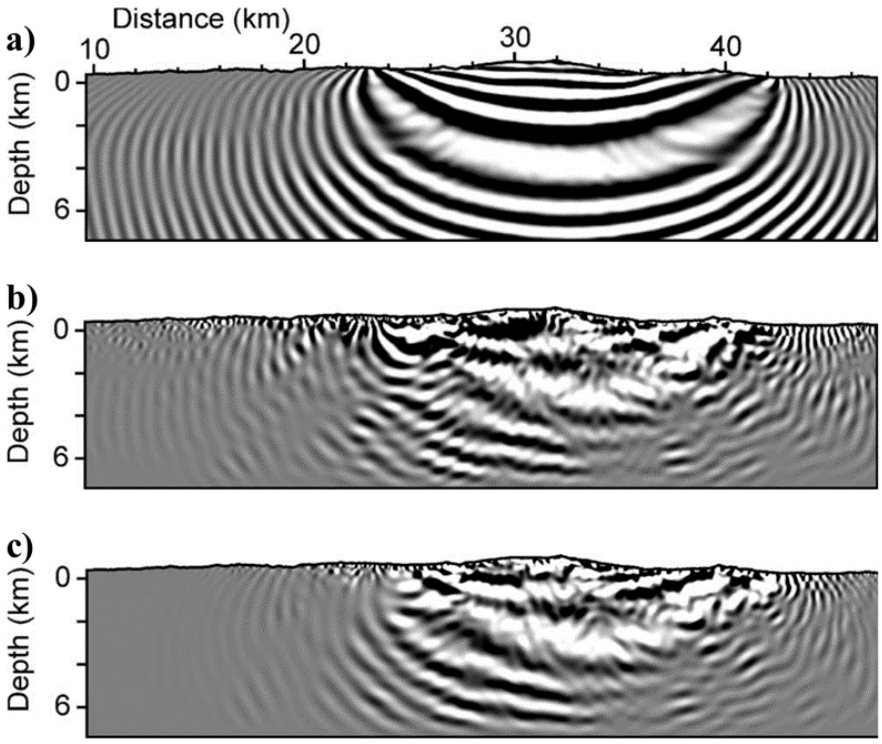


Figure 6.3: Effect of land topography on a seismic wavefield shown by 4-Hz pressure wavepaths calculated using (a) absorbing, and (b,c) free-surface boundary with (c) and without (b) time-windowing. Different code was used in (a) vs. (b,c) cases. Adapted from Bleibinhaus and Rondenay (2009).

neglecting surface-related effects either by putting an absorbing layer above topography or by approximating it with a flat surface (fig. 6.3). A heuristic conclusion of Huiskes et al. (2017) is that elevation variation of $\lambda/4$ (where λ is a local wavelength) may produce significant inversion artifacts. According to this rule of thumb, for a wavelength of 500 m (e.g. a 4-Hz wave travelling in shallow sediments at 2 km/s), elevation variations of 100 m may already cause a problem.

6.1.4 Modelling methods

Many, if not most, numerical solvers used in seismology are based on the finite-difference (FD) method (Warner et al., 2013b; Bohlen et al., 2015; Métivier and Brossier, 2016; Louboutin et al., 2019). In FD, to model a flat free surface, it is easiest to define a grid so that the free surface is one of its planes. For a staggered grid even this might pose a problem. Various solutions have been proposed to avoid Rayleigh wave numerical dispersion, including the image method (Levander, 1988) and its improvements such as AFDA (Kristek et al., 2002). Modelling an irregular free surface is much more difficult problem on any finite-difference grid and no fully satisfactory solution has been found so far (Moczo et al., 2014). Below, different takes on handling an irregular free surface in the FD method are presented.

Absorbing layer One idea is to replace the free surface with an absorbing boundary. One would have to suppress all surface-related reflections in the field data which in most cases is impossible.

Flat approximation Another approach is to approximate an irregular topography with a flat (planar) free surface. This may be enough for marine receivers located far from land. Land stations would

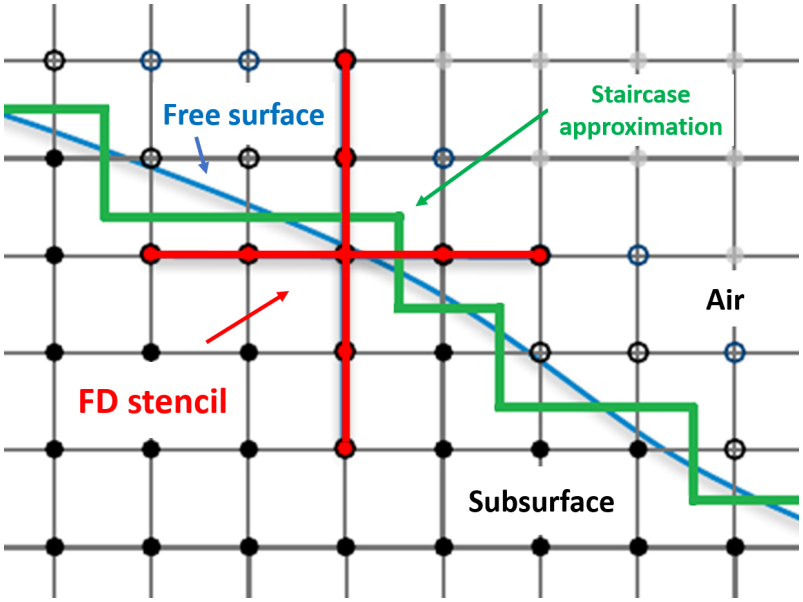


Figure 6.4: Staircase approximation of the irregular free surface arising in vacuum/air-layer finite-difference methods. Modified after M. Paulatto (unpublished).

require static corrections. However, these corrections assume sub-vertical propagation and thus work only for short-offset data (Sheriff and Geldart, 1995).

Model adaptation In this method referred to as the *vacuum* or *air-layer* method (e.g. Opršal and Zahradník, 1999), model parameters above the free surface are assigned a very low acoustic impedance ρv. This is typically achieved by setting the density close to 0 as the velocity needs to be rock-like to retain the numerical stability, unless small Δx or large Δt are used. This approach, despite being perhaps the most intuitive one, leads to a staircasing error of $\mathcal{O}(\Delta x)$ order (fig. 6.4). A flavour of the vacuum method which does not require model adaptation is proposed in the appendix 6.A.

Stencil adaptation The idea of this approach is to change the FD stencil so that it has non-zero coefficients only for the grid nodes below the free surface. This change depends on the position relative to the free surface and thus is difficult to implement in a generic way, except for very simple cases (Moczo et al., 2014).

Grid adaptation In this approach, the finite-difference grid (mesh) is transformed in one of the two major ways:

- *refinement* – using a smaller grid spacing better samples the boundary, decreasing the roughness of the staircase approximation. Unfortunately this method is prohibitive for large 3d models where performance is the biggest issue. Computational overhead increases with the number n of the grid nodes in one direction as n^4 (the time step must be reduced too due to the stability criterion). Refining grid only locally (Robertsson and Holliger, 1997) is an option, however the interface between coarser and finer parts of the grid can also generate artificial scattering;
- *coordinate transformation* – in some cases a transformation between Cartesian and curvilinear coordinates can be found such that it maps an irregular surface into a 'flat' surface in the new coordinates (Zhang and Chen, 2006); one of the downsides of this method is due to the fact that

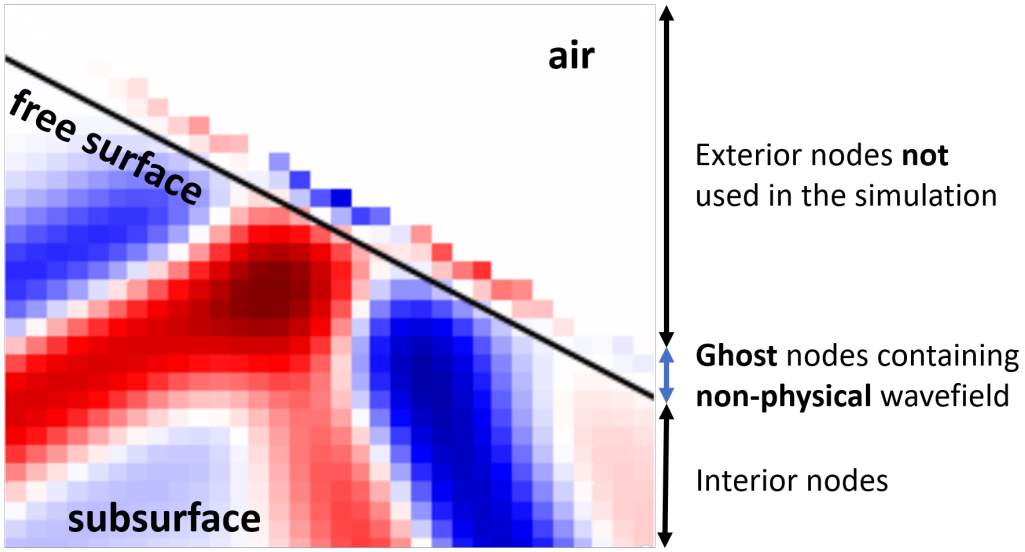


Figure 6.5: Idea of the immersed-boundary method. The physical wavefield propagating in the subsurface (interior grid nodes) is anti-symmetrically extrapolated above the free surface so that the non-physical wavefield at the ghost nodes imposes the zero-pressure boundary condition at the free surface.

automated generation of grid conformal to the highly irregular free surface often creates very small cells enforcing a smaller time step (Hu, 2016). This problem is also present in the finite-element method, and meshing in general is a challenging task (Igel, 2016).

Wavefield adaptation In this approach, the wavefield is assigned some (not necessarily physical) values above the free surface so that the BC is imposed implicitly. Its major variants include:

- *image* method – proposed for staggered grids and planar surfaces by Virieux (1986) and Levander (1988). In this method, stress or velocity above the surface is assigned values opposite to those below the surface. This method is accurate only for horizontal, planar free surfaces. A straightforward extension to irregular surfaces leads to stair-casing errors (Robertsson, 1996);
- *immersed*[3] *-boundary* (IB) method – can be thought of as an extension of the image method that preserves the accuracy of the FD scheme in the presence of irregular surfaces.

As no stencil manipulation is needed in the IB method, its implementation can be decoupled from the FD scheme. This is highly beneficial when augmenting the existing solvers. Furthermore, the IB method does not require special data processing or meshing. Because of these advantages, IB was the method of choice in this study. Its further discussion is provided below.

6.1.5 Immersed boundary (IB)

The bottom line of the immersed-boundary method is very simple and in fact it was already stated in the description of the image method in the previous section. The idea is to assign some non-physical ('virtual' or 'ghost') values to the wavefield above the free surface (at so called *ghost nodes*) based on the values of the physical wavefield in the subsurface (fig. 6.5). Although the basic idea is simple, a

[3]The name of the method comes from the fact that the irregular surface is 'immersed' in a regular grid, no modification of the grid is required (Mittal et al., 2008).

1976	Introduction of a 2nd-order accurate staggered-grid scheme for elastic problems in velocity-stress formulation by Madariaga.
1984	Applying Madariaga's grid to SH elastic waves by Virieux.
1986	Applying Madariaga's grid to P-SV elastic waves by Virieux.
1988	Improving the accuracy of Madariaga-Virieux grid up to the 4th order in space and introducing the *image method* by Levander.
1996	Extension of the image method to irregular surfaces by Robertsson.
2001	Immersed interface method for 1d acoustic waves by Piraux and Lombard.
2004	Immersed interface method for 2d acoustic/elastic waves by Lombard and Piraux.
2008	Immersed boundary method for 2d elastic waves by Lombard et al..
2010	Immersed boundary method for 2d E-M waves by Zhao.
2015b	Immersed boundary method for 3d elastic waves by Gao et al. (in progress?).
2016	Immersed boundary method for 2d acoustic waves for the system of 1st-order PDEs by Hu.
2017	Simplified immersed boundary method for 2d acoustic waves by Mulder.
2021	Extension of Mulder's method to 3d acoustic waves by Caunt et al. (in progress).
2021	Immersed boundary method for 3d acoustic waves for the 2nd-order PDE (this study).

Table 6.1: Time-line of major developments of the immersed-boundary method for modelling seismic waves.

rigorous analysis of the stability and accuracy of the method can be challenging even for homogeneous 1d models (Beyer and LeVeque, 1992), not to mention more realistic geometries (Kreiss and Petersson, 2006).

Wavefield extrapolation The above assignment can be thought of as wavefield *extrapolation*. The values extrapolated across ghost nodes should be such that the desired boundary condition is imposed implicitly. Due to the linearity of the FD stencil, the pressure wavefield is an odd (antisymmetric) function *across* the free surface. Hence, for an irregular surface, the extrapolation should be performed along *normal vectors* of the surface in order to retain the accuracy of the scheme. Other approaches, which trade in some accuracy for the robustness, have also been reported, e.g. by Mulder (2017) – see the summary below.

6.1.6 Previous IB studies

Although the main idea of the IB method is straightforward, there are a lot special cases to consider which make the implementation challenging. Major approaches published to date are summarised below. The overview of methodological developments related to wave phenomena is given in tab. 6.1.

Peskin (1972) This is the first[4] immersed-boundary study – a Navier-Stokes simulation of the flow around heart valve immersed in blood. Although both the equation and the boundary condition were different to the ones used in seismology (e.g. boundary shape changed at each time step), the method proved to be applicable to seismic problems. The author used a simple linear extrapolation scheme along the normal to the boundary. In most modern approaches, such as by Kreiss and Petersson (2006)[5], a higher-order polynomial extrapolation is applied.

[4] Or at least the one which gave rise to the method's name.

[5] The method of Kreiss and Petersson (2006) seems to be a prototype of many subsequent studies using the 'ray-casting' method, e.g. by Zhao (2010), although it is not acknowledged there.

Lombard et al. (2008) The authors perform a 2d extrapolation by kth order Taylor expansion using interior points contained within a circle of a certain radius centred on the exterior node to be evaluated. A value of k larger than a scheme order r is claimed to ensure stability, although numerical experiments allow values as low as $k = r - 1$. Finding of the extrapolation polynomial is an overdetermined problem solved in least-squares sense. The free surface is treated locally as a straight line (that is why it is assumed to be sufficiently smooth), perpendicular to the line joining the exterior node and the intersection of the free surface and its normal vector. The radius of the circle is specific for each k ensuring an over-determination of the problem. Sharp corners are revealed to invalidate the accuracy of the method and are replaced by an arc of the circle. They use a quite terse matrix notation starting from the 2d wave equation in the stress-velocity formulation:

$$\partial_t U = A\partial_x U + B\partial_z U,$$

where $U = (v_x, v_z, \tau_{xx}, \tau_{xz}, \tau_{zz})$ is the wavefield vector and A, B are 5×5 matrices. This can be expanded as:

$$\begin{aligned}
\dot{\tau}_{xx} &= \lambda(\partial_x v_x + \partial_z v_z) + 2\mu\partial_x v_x = (\lambda + 2\mu)\partial_x v_x + \lambda\partial_z v_z \\
\dot{\tau}_{zz} &= \lambda(\partial_x v_x + \partial_z v_z) + 2\mu\partial_z v_z = (\lambda + 2\mu)\partial_x v_x + \lambda\partial_z v_z \\
\dot{\tau}_{xz} &= \mu(\partial_x v_z + \partial_z v_x).
\end{aligned}$$

The method is reported to lack stability in the case of a staggered grid, including the more recent attempt by Gao et al. (2015b).

Hu (2016) The author follows the approach of Zhao (2010) and proposes a two-stage interpolation scheme along with an iterative procedure of populating ghost values above the highly irregular free surface. He uses the 1st order formulation of the acoustic wave equation:

$$\begin{cases} \rho i\omega v_j = -\partial_j p + f_j \\ \dfrac{1}{K} i\omega p = -\partial_j v_j + q \end{cases}. \tag{6.4}$$

In this formulation, the 0th order (Dirichlet) boundary condition $p = 0$ is less accurate. This is because v has to be updated based on p and the boundary condition is not imposed for v directly. The 1st order (Neumann) boundary condition, which is more accurate for this formulation, reads:

$$\begin{cases} v_{x'}|_\Gamma = 0 \\ \partial_{z'} v_{z'}|_\Gamma = 0 \end{cases},$$

where x', z' are coodinates in a rotated Cartesian frame, so that x' and z' axes are tangential and normal to the surface, respectively. At the boundary, the tangential velocity is zero, and the normal velocity is 'flat' (its slope is zero).

Mulder (2017) The author implements a simpler variant of IB, after facing stability issues with the method of Lombard et al. (2008). He extrapolates the wavefield along the grid lines, instead of along the normal vectors. This approach may be regarded as a stencil-modification scheme (see sect. 6.1.4), the difference between the two being subtle. He extrapolates the wavefield across several exterior notes at once using polynomials of degree as high as 7, not reporting any issues due to the Runge's phenomenon[6]. Mulder and Huiskes (2017) extend this approach to the first-order formulation of the acoustic wave equation (eq. 6.4).

[6] This phenomenon consists in spurious oscillations of the polynomial near the ends of the interpolating interval. The effect increases with the degree of the polynomial and can be alleviated by choosing the Chebyshev sampling (Strang, 1986).

Algorithm 6.1 Simulating an irregular free surface.

Pre-calculation *(one per FWI case study)*

1. Read in and check surfaces & acquisition geometry.
2. For each surface calculate wavefield-extrapolation weights.
3. For each surface calculate source and receiver weights.

Wavefield extrapolation *(at each time step of forward/adjoint modelling)*

1. Initialise source-injection and wavefield-extrapolation weights.
2. Inject the forward/adjoint source using pre-calculated weights.
3. Set ghost wavefield to zero (vacuum step).
4. Iteratively update the ghost wavefield using pre-calculated weights.
5. Update the wavefield inside the model.
6. Go back to point 2 to start a new time step (time-domain codes).

6.2 Basic IB method

The immersed-boundary method introduced in the previous section will now be described in detail, focusing on practical, implementation-related aspects. The method was implemented as a stand-alone `Fortran` package `IMMERSE` (see appendix 6.B), and the following description, to a large extent, applies to this implementation. This work is to some extent based on the method of Hu (2016). However, it differs from Hu (2016) in several important aspects, including: (1) the second order formulation, (2) three-dimensional implementation, (3) surface sources (4) and including a field-data example.

6.2.1 Workflow summary

One of the benefits of the immersed-boundary method is that its workflow can be decoupled almost entirely from the original numerical scheme, so can be added to existing codes easily, at least in theory (Virieux et al., 2012). In practice, one has to take care of this decoupling by an appropriate design of the workflow. This workflow can be naturally divided into two parts, the pre-calculation and the wavefield extrapolation (alg. 6.1). The former is much more challenging and time-consuming, and is where the majority of the theory is implemented. Any potential accuracy and stability issues will most likely have their source there. The latter requires a minor modification to the time-stepping scheme in the existing FWI code. Its computational cost per time step[7] is much smaller than for the pre-calculation.

6.2.2 Wavefield grid

Firstly, the wavefield and model are assumed to be sampled on the same uniform rectilinear grid padded with absorbing boundaries if needed.

Grid spacing In the finite-difference method, a solution of the wave equation (i.e. the wavefield) is sampled in space and time on a *grid* (*mesh*). In this study, the spatial grid is assumed to be uniform rectilinear (i.e. cuboidal), with the same spacing h, constant everywhere in all dimensions:

$$h \coloneqq \Delta x = \Delta y = \Delta z. \tag{6.5}$$

The model is assumed to be defined on the same grid, otherwise some interpolation would be required to update the wavefield, because the model coefficients appear in the finite-difference stencil. The temporal grid is usually determined from the spatial grid so as to honour the CFL condition (Strang, 1986) and retain the stability of the scheme.

[7] Although the time-stepping is being constantly referred to, the method works just as well for modelling in the frequency or other domains.

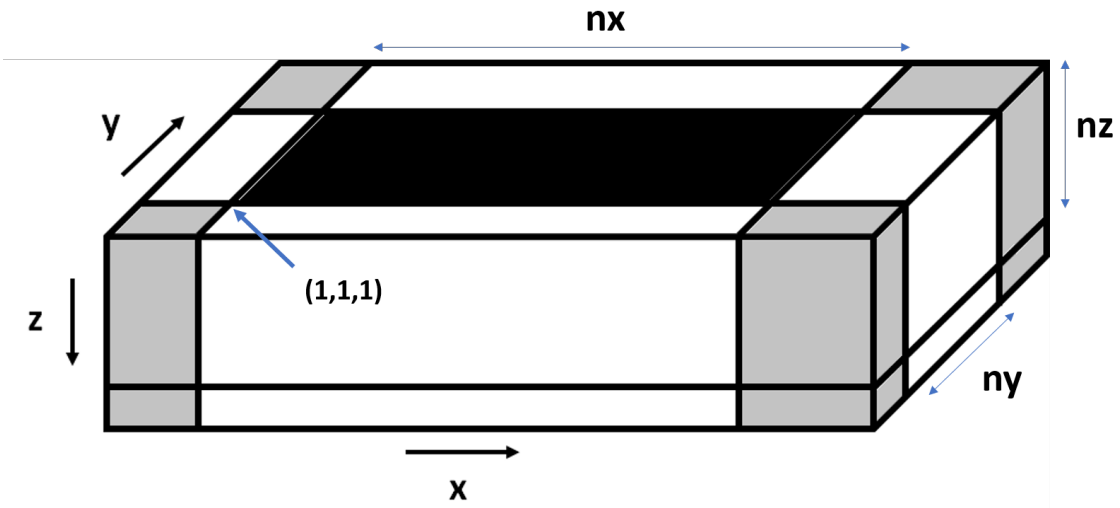


Figure 6.6: Wavefield grid – inner grid (black) with the origin at node (1,1,1) and nx, ny,nz in size is extended by the left, right, front, back and bottom paddings (white with grey corners).

Extra padding In this study, unless stated otherwise, all model boundaries except for the top one are effectively moved to infinity by imposing absorbing boundary conditions (Cerjan et al., 1985). These boundaries have finite width and require padding the grid with layers of extra nodes (fig. 6.6). Consequently, we will define (1) an *inner grid*, comprising the actual model and all the shots and receivers and (2) an *extended grid*, i.e. the inner grid extended over the extra nodes. The number of extra nodes depends on the desired effectiveness of absorption. The method of Cerjan et al. (1985) requires the layers to be at least twice as wide as the longest wavelength simulated.

6.2.3 Surface representation

In this section, a continuous representation of the free surface $z(x, y)$ sampled horizontally on a uniform rectilinear model grid is given by interpolation between the samples using cubic B-splines (in 2d/3d) or bilinear patches (in 3d).

Assumptions The surface is assumed to be an *explicit function* $z(x, y)$ of horizontal coordinates (x, y), i.e. a graph of continuous function of two variables[8] defined over a connected open subset of the Euclidan plane $\mathbb{R}^2$. In other words, for any pair there can be only one surface point. The explicit form cannot represent overhanging or closed surfaces but it greatly simplifies the analysis as meshing of general surfaces/volumes is a challenging task[9] (Igel, 2016).

Point cloud At the beginning, the surface is almost always provided in the form of a *point cloud*, i.e. as a set of (x, y, z) tuples. In crustal-scale seismic experiments, such as the one described in sect. 6.6, such a point cloud may be acquired using a LIDAR or a multibeam echo-sounder, in the case of land and seabed topography, respectively. Before being fed into the IB algorithm, the surface is expected to be defined on the wavefield grid (sect. 6.2.2). Such representation can be easily obtained from a point cloud by interpolation.

[8] Which can be thought of as elevation.
[9] This is one of the reasons for popularity of the finite-difference method and for the substantial entry barrier of codes which require meshing, such as SpecFEM (Komatitsch and Tromp, 1999).

Extra padding Absorbing boundaries require extra padding of the wavefield grid (sect. 6.2.2). The surface is extrapolated over the paddings using the nearest neighbour method. Hence, in the padding it follows the topography along the corresponding edge of the inner grid and is flat in the corners (fig. 6.6).

Continuous form Transforming a set of discrete points into a continuous function is a process of interpolation. It is unlikely that a single, reasonably behaving, continuous function can be found to represent a real-world surface. Two basic methods of solving this problem are available: *tesselation* and *piece-wise interpolation*, where the former is a special case of the latter. Here, *piecewise interpolation* between a uniform grid of data points is performed resulting in a *mesh* of continuous surface *patches*. Two interpolation methods, cubic B-splines (in 2d/3d) and bilinear patches (in 3d), were implemented and are discussed below.

Cubic B-splines In this study, B-spline patches are used both in 2d and 3d. This is a special, non-rational case of NURBS (e.g. Bingol and Krishnamurthy, 2019), but the non-rationality[10] makes little difference for seismic applications. A uniform grid of control points is assumed. Such a grid can be easily obtained from a point cloud by any interpolation algorithm. The grid spacing is by default expected to be the same as the wavefield's grid[11] (sect. 6.2.2). The surface is interpolated by cubic B-splines centred at x' defined by Hobro (1999) as:

$$A^{1\mathrm{d}}(x,x') = \begin{cases} \frac{1}{2}\left(\frac{|x-x'|}{h}\right)^3 - \left(\frac{|x-x'|}{h}\right)^2 + \frac{2}{3}, & \text{if } |x-x'| \le h, \\ \frac{1}{6}\left(2 - \frac{|x-x'|}{h}\right)^3, & \text{if } h < |x-x'| \le 2h \\ 0, & \text{if } 2h < |x-x'| \end{cases} \tag{6.6}$$

The coefficients are obtained by solving a linear system demanding continuity of the function up to the second derivative. For a 2d surface:

$$A^{2\mathrm{d}}(x,y;x',y') = A^{1\mathrm{d}}(x,x')A^{1\mathrm{d}}(y,y') \tag{6.7}$$

Interpolation is performed using the formula for 2d and 1d:

$$z(x,y) = \sum_{i=1}^{16} A^{2\mathrm{d}}(x,y;x_i,y_i)\cdot z_i \tag{6.8}$$

$$z(x) = \sum_{i=1}^{4} A^{1\mathrm{d}}(x;x_i)\cdot z_i. \tag{6.9}$$

Bilinear patches Data points $Q_{ij} = (x_i, y_j)$ are assumed to be on a regular grid, e.g. $Q_{12} = (x_1, y_2)$. Extrapolant is given after performing a linear interpolation in each direction:

$$z(x,y) = \frac{y_2 - y}{y_2 - y_1} z\left(R_1\right) + \frac{y - y_1}{y_2 - y_1} z\left(R_2\right), \tag{6.10}$$

where:

$$z(R_1) = \frac{x_2 - x}{x_2 - x_1} z\left(Q_{11}\right) + \frac{x - x_1}{x_2 - x_1} z\left(Q_{21}\right) \tag{6.11}$$

$$z(R_2) = \frac{x_2 - x}{x_2 - x_1} z\left(Q_{12}\right) + \frac{x - x_1}{x_2 - x_1} z\left(Q_{22}\right). \tag{6.12}$$

Note that the interpolant is a quadratic surface, even though each step is linear.When implementing this interpolation, it is important to remember that the order of Q_{ij} corners matters (fig. 6.13, right).

[10] One of its consequences is inability to represent exactly circles/spheres.
[11] This is because topography and wavefield grid are the same in the field-data experiment (sect. 6.6).

6.2.4 Nodes classification

In this section, the grid nodes are classified according to their positions relative to the surface. Once the input surface has been loaded and checked for errors and its continuous representation has been chosen, the next step is to classify grid nodes according to their positions relative to the surface. Perhaps the simplest option is to assign two kinds of flags. The first kind of flags divides all grid nodes into exterior/surface/interior nodes. The second, divides exterior grid nodes into non-ghosts and ghosts of different levels. These two kinds of flagging are discussed below.

Exterior/interior Firstly, nodes are classified according to their position relative to the free surface. They can be located *above*, *below* or *at* the surface. In a general case of overhanging/closed surface, the terms 'below' and 'above' change to 'inside' and 'outside'. For closed surface, normal vectors must be defined to discriminate between the exterior and interior. It may or may not happen that there are nodes exactly at the surface. It may be useful to treat as such the nodes that are close enough to the surface. In the `IMMERSE` code the threshold is set as a runfile parameter `minGhostFS` with a default value of 0.001 grid nodes. . Note that the algorithm of classifying the nodes strongly depends on the chosen surface parametrisation.

Ghost nodes Secondly, exterior grid nodes are divided into non-ghosts and ghosts of different levels, based on their proximity to the surface. Every ghost node encodes the actual, fractional position of the surface. Ghost nodes, coating the free surface with no gaps and at multiple levels above it, are able to capture the surface curvature with a subcell resolution (Hu, 2016). Independent ghost levels reduce coupling between ghosts and decrease the number of iterations in the iterative update. Coupling occurs when the wavefield at a ghost node depends on the wavefield value at another ghost node instead of solely on the interior nodes. Let us define first-level ghosts as exterior nodes having at least one interior-node as a neighbour, and ith level ghosts as exterior nodes having at least one $(i-1)$th-level ghost as a neighbour. In such case, number of ghost levels equal to half-size of the FD stencil (defined here as a number of nodes from its centre to its most distant point) prevents the stencil from using the nodes which are neither interior nor ghost nodes, i.e. which have undefined values. Following this rule, 4th-order accurate scheme needs $4/2 = 2$ ghost levels, 10th-order accurate scheme needs $10/2 = 5$ ghost levels, and so on. Unlike Mittal et al. (2008) who defines neighbours along grid lines, here diagonal neighbours are also considered, because of the off-diagonal elements of the finite-difference stencils used. Ghost nodes are further split into three groups: (1) *ordinary* (2) *surface* and (3) *margin* ghost nodes. The surface ghost nodes are defined as located less than half a grid node from the discretised surface. The margin ghost nodes are the ones for which no other points have to be calculated.

6.2.5 Point-projection

The premise of the IB method is based on extrapolating the wavefield across ghost nodes along the surface normal vectors. Finding a normal for each ghost node (fig. 6.7) is arguably the most important step of the IB algorithm, essential for its robustness and efficiency. This is because:

- the extrapolation is very sensitive to the position of the intersection of the surface and the normal vector – inaccurate positions of intersections can destabilise the extrapolation (Mittal et al., 2008),
- as a result of the required accuracy, this is also the most computationally intensive part of the IB workflow.

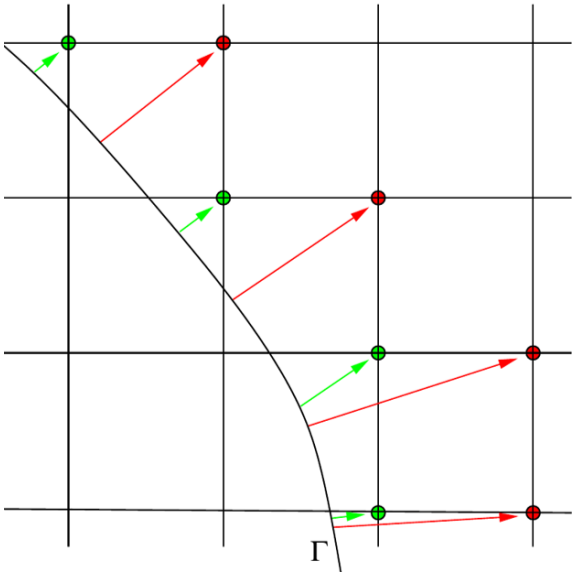


Figure 6.7: Orthogonal point-projection in 2d. Circles of different colors denote ghost nodes of different levels projected onto the surface Γ. From Zhao (2010).

Algorithm 6.2 Orthogonal point-projection posed as a global optimisation problem.

1. Find coordinates of the point G which you want to project onto the surface Γ.
2. Decompose the surface Γ into continuous patches γ_i from which you can generate points.
3. Exclude patches that contain no solution using the spherical clipping method.
4. Find a (local) minimum distance to each of the patches γ_i using a grid search (not feasible for most realistic 3d problems) or a local optimisation algorithm such as the Newton method (sect. 6.3).
5. Choose the shortest minimum distance as a (global) minimum distance to the surface Γ. The point corresponding to this global minimum is considered to be the orthogonal projection of G onto Γ.

Yet, surprisingly little attention is paid to this issue in the existing literature, at least in seismological applications. Here, after formulating the problem in the general case, a brute-force method suitable for 2d geometries is presented. The 3d case, which is far more challenging, is analysed in sect. 6.3.

Formulation For every ordinary ghost node (see sect. 6.2.4) let us define a surface normal that passes through this ghost node (fig. 6.7). The origin of this normal vector is an intersection I such that $\mathrm{I} \in \Gamma$ and:

$$\overrightarrow{\mathrm{GI}} \cdot \mathbf{t} = 0, \tag{6.13}$$

where $\mathbf{t}(I)$ is tangent to Γ at point I and $\overrightarrow{\mathrm{GI}}$ is a vector pointing from G to I. Finding I for a given G is called the (*orthogonal*) *point-projection* problem (e.g. Mittal et al., 2008). The point-projection can be posed as a problem of finding the shortest L^2 distance between a given point $G = (g_1, g_2, g_3)$ and a point $X = (x_1, x_2, x_3)$ on the surface:

$$|d|^2 := \sum_{i=1}^{3} (g_i - x_i)^2. \tag{6.14}$$

If there exists a unique global minimum X, it can be regarded the desired intersection I. The other, special, cases are discussed below.

Degenerate case In general, more than one surface normal may be passing through a given ghost node. Let us denote the intersections associated with these normals by I_i where $i = 1, 2, ...$ etc. If the distances $|GI_i|$ are different, the optimisation approach will select the shortest one. This is desirable, as the normal vector is expected to capture the *local* slope of the surface.

Grid search In 2d, probably the easiest to implement method of optimising (6.14) is a *grid search* within a certain radius. The search radius r is defined along grid lines so it corresponds to a L^1 (taxicab/Manhattan) metric $d(a, b) = \sum |a_i - b_i|$. Testing a range of values showed that the radius r should not be shorter than the number of ghost levels (scheme order):

$$r \geq k. \tag{6.15}$$

A slighly larger number may be a safer option ensuring the global minimum is not missed. This process of excluding points that are further away from the point to be projected is referred to as *spherical clipping* (Chen et al., 2008, 2009). Within the search radius, the grid should be *locally refined*. The new grid spacing directly translates into the error of the position-determination of the closest surface point that can be found. This has to be as small as possible, as it is has the largest effect on the accuracy of the immersed boundary method. To determine the acceptable accuracy value, normal vectors were plotted for a case with a complex topography (fig. 6.9). A visual inspection of the normal vectors suggested that the `accuracy` parameter a defined as:

$$a \coloneqq h_{\text{original}}/h_{\text{refined}}, \tag{6.16}$$

where h is the grid size, should be at least 10^4. Smaller values resulted in a distinct mispositioning of the surface points. This effect increases with the scheme order (where more fictitious nodes are used) because the intersection-mislocation error increases with the distance from the surface.

Limitations The grid search was implemented in `IMMERSE` for both 2d and 3d. The latter case, however is too expensive computationally for surfaces typically encountered in real-world problems as the one discussed in sect. 6.6. In 3d, for every cell of the original grid within the search radius one needs to traverse through $(10^4)^3 = 10^{12}$ points of the refined grid assuming $a = 10^4$. The number of cells to search through is proportional to r^2, i.e. the search radius squared. Consequently, the total number of points N_{search} to traverse is:

$$N_{\text{search}} \sim N_G r^2 a^3, \tag{6.17}$$

where N_G is the number of ordinary ghost nodes. For each of these N_{search} points, one needs to calculate the surface elevation and the distance from the surface point to the ghost node. On the tested machine (Intel(R) Core(TM) i7-6700 CPU @ 3.40GHz with 32 GB RAM), the value $a = 10^4$ cannot be handled in 3d due to a memory overflow. For this reason, an alternative approach has been developed. It is discussed in sect. 6.3.

6.2.6 Extrapolation data

In this section, data points and associated weights required for the wavefield extrapolation are defined. These data can be pre-calculated and re-used in multiple FWI runs provided the free surface and stencil order remain the same.

Extrapolating the physical wavefield along the free surface normal vectors into ghost nodes is a key part of the immersed-boundary method. This is where the boundary condition is actually being imposed.

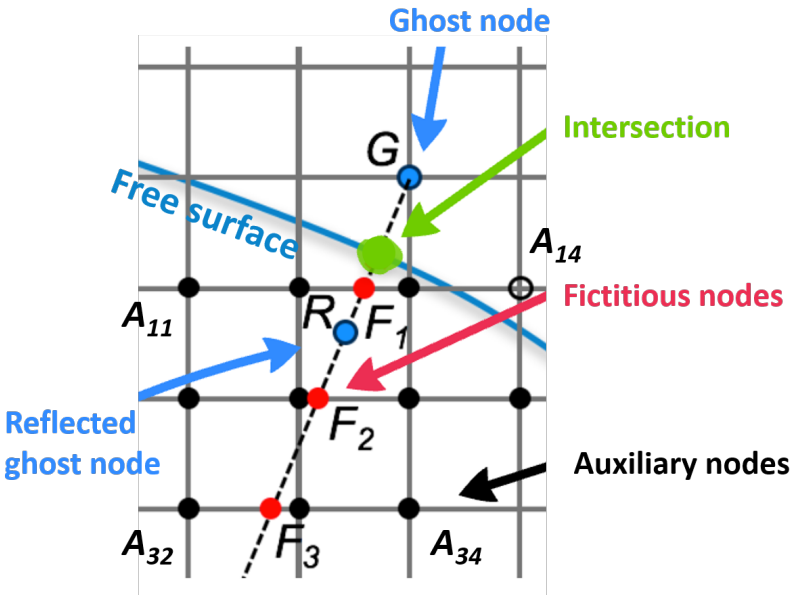


Figure 6.8: Points used in the immersed-boundary method to extrapolate the wavefield along the normal vector (dashed line) across the free surface into the ghost node G. Here, the extrapolation is done in three steps: 1. values at fictitious nodes F_i are found by interpolation between auxiliary nodes A_{ij}; 2. value in R is found by interpolation between fictitious nodes and intersection I. 3. value in G is set equal to the value in R with the opposite sign. Modified after M. Paulatto (unpublished).

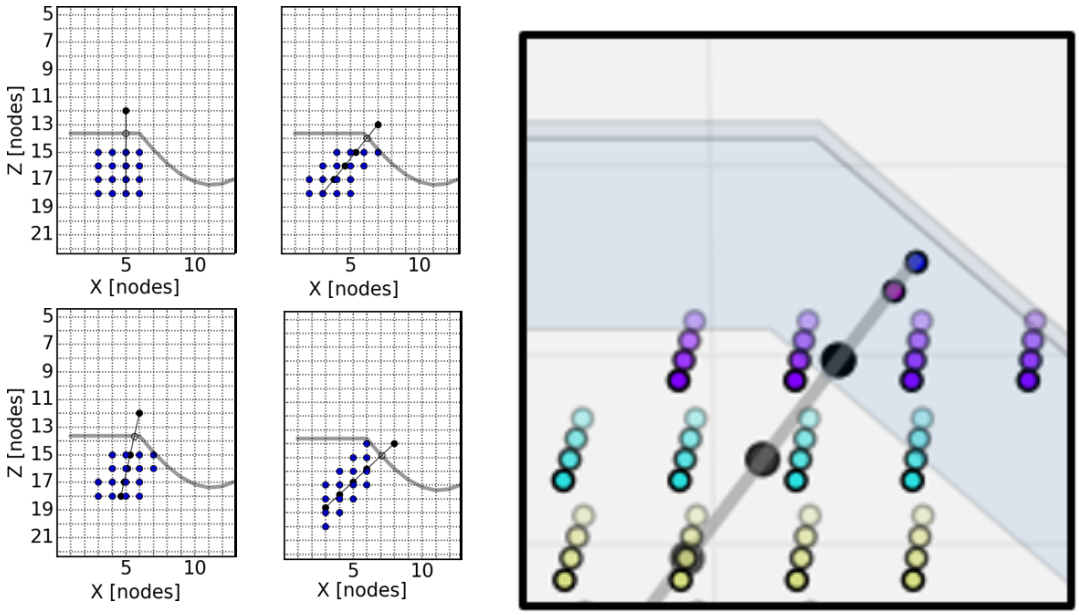


Figure 6.9: Variable orientation of the extrapolation stencil (fictitious and auxiliary nodes) for the free surface of a variable slope locally exceeding 45°. The same behaviour was implemented in 3d (right).

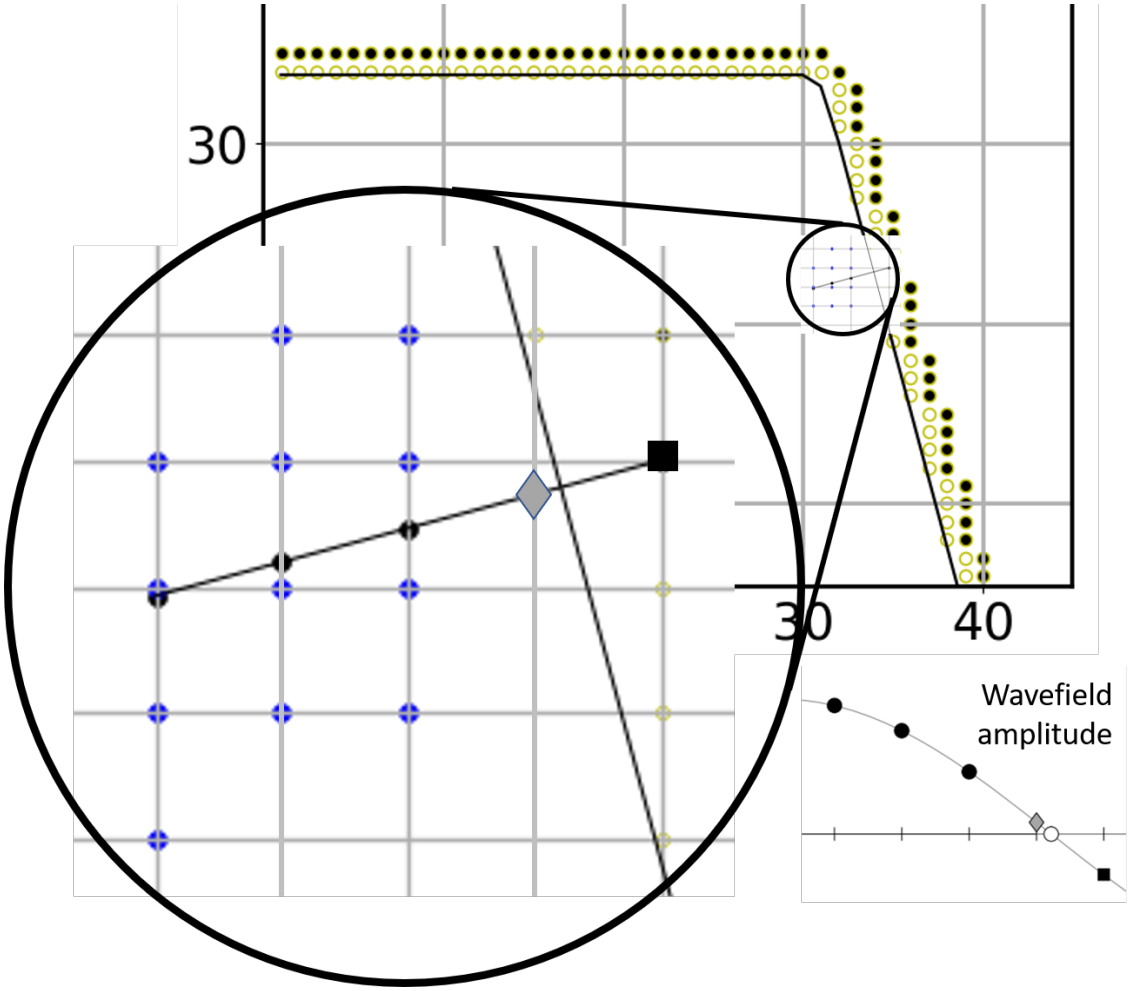


Figure 6.10: Exclusion of a fictitious node (grey rhombus) located less than half grid spacing from the free surface to avoid instability of some interpolators. The bottom-right inset adapted from Mulder (2017).

Experience shows that the extrapolation is also a potential source of instabilites. Overshot wavefield values may cause a kind of *resonance* near a reflecting boundary (Trefethen, 1982) leading to a numerical error that grows to infinity exponentially. For this reason, designing an extrapolation scheme requires a special care. Here, the approach of Hu (2016) is followed.

Formulation Extrapolation can be formulated as an interpolation problem (Press et al., 1990), treating the point of interest (here, the ghost node G) as one of the data points, and one of the data points (here, the intersection I) as the point of interest. Interpolation gives a value of the wavefield in I according to the formula:

$$p(I) = w_0 p(G) + \sum_{i=1}^{N_f} w_i p(F_i) \tag{6.18}$$

where w_j, with $j \in \{0, \ldots, N_f\}$, are the interpolation weights associated with data points located along the normal vector: the ghost node G and N_f *fictitious nodes*[12] F_i (fig. 6.8). The accuracy of the

[12]This name is used by Zhao (2010) probably because F_i are not necessarily grid nodes. It is rather unfortunate as they store interpolated values of the physical, not fictitious, wavefield. Hu (2016) calls them 'fictitious points', whereas 'fictitious values' are used by Lombard et al. (2008) for what is called here a ghost wavefield. Despite the potential confusion, the names 'fictitious nodes/points' will be used here for the sake of consistency with the similar approaches.

interpolation depends on N_f and is discussed below. By definition, the pressure wavefield is exactly zero on the free surface, therefore:

$$p(I) = 0,$$

and solving eq. (6.18) for $p(G)$ gives:

$$p(G) = \sum_{i=1}^{N_f} f_i p(F_i) \tag{6.19}$$

where the weights f_i are

$$f_i := -\frac{w_i}{w_0}. \tag{6.20}$$

The specific values of weights w_j depend on the interpolation method chosen (discussed below). In general, a normal vector will not coincide with the grid lines, and, consequently, the values at F_i have to be found by interpolation between *auxiliary* (grid) *nodes* A^i_{jk}:

$$p(F_i) = \sum_{j=1}^{N_a} \sum_{k=1}^{N_a} a^i_{jk} p(A^i_{jk}). \tag{6.21}$$

where a^i_{jk} are the interpolation weights of the 2d array of auxiliary nodes A^i_{jk} associated with the ith fictitious node F_i of the ghost G. Note that N_a is a number of auxiliary nodes per fictitious node per dimension. The total number of auxiliary nodes per F_i in 3d is then N_a^2. Similar to N_f, N_a affects the accuracy of the interpolation. The nodes A^i_{jk} surround F_i with the same number of nodes at each side (fig. 6.8), so that F_i is close the center of the interpolation stencil, where the accuracy of interpolation is maximum (Press et al., 1990). The 2d arrays A^i_{jk} have a variable orientation parallel to either XY, XZ or YZ grid planes, depending on the direction of the surface normal vector in such a way the arrays remain sub-parallel to the local tangent plane of the surface (fig. 6.9). Combining eq. (6.21) and (6.19), one obtains an extrapolation formula for a ghost wavefield above the free surface as a function of the wavefield inside the model:

$$p(G) = \sum_{i=1}^{N_f} f_i \sum_{j=1}^{N_a} \sum_{k=1}^{N_a} a^i_{jk} p(A^i_{jk}). \tag{6.22}$$

In 3d, the above formulation consists of two 1d interpolations between auxiliary nodes, followed by the extrapolation into the ghost node. This multi-step scheme may be replaced by a single-step scheme as in Lombard et al. (2008), with the caveat of potential stability issues having their cause obscured by more complex mathematics.

Interpolator For stability reasons (Trefethen, 1982), specific functions (interpolators) chosen to calculate interpolation weights f_i and a^i_{jk} should not overshoot the values between data points, e.g. due to the Runge phenomenon (Press et al., 1990). Here, interpolation using a simple type of *radial basis function* (RBF) referred to as *inverse distance weighting* or IDW (Shepard, 1968) was used:

$$I(x) = \begin{cases} \sum_{i=1}^{N} w_i(x) f(x_i), & \text{if } x \neq x_i \\ f(x_i), & \text{if } x = x_i \end{cases}$$

where $I(x)$ and $f(x)$ denote the interpolating and the original (unknown) function, respectively, x_i are the data points for which the values of f are known, $w_i(x)$ are the interpolation coefficients/weights and N is the number of data points. The weights $w_i(x)$ are given by:

$$w_i(x) = \frac{1}{\mathcal{N}(x)} \left[d(x, x_i)\right]^{-p},$$

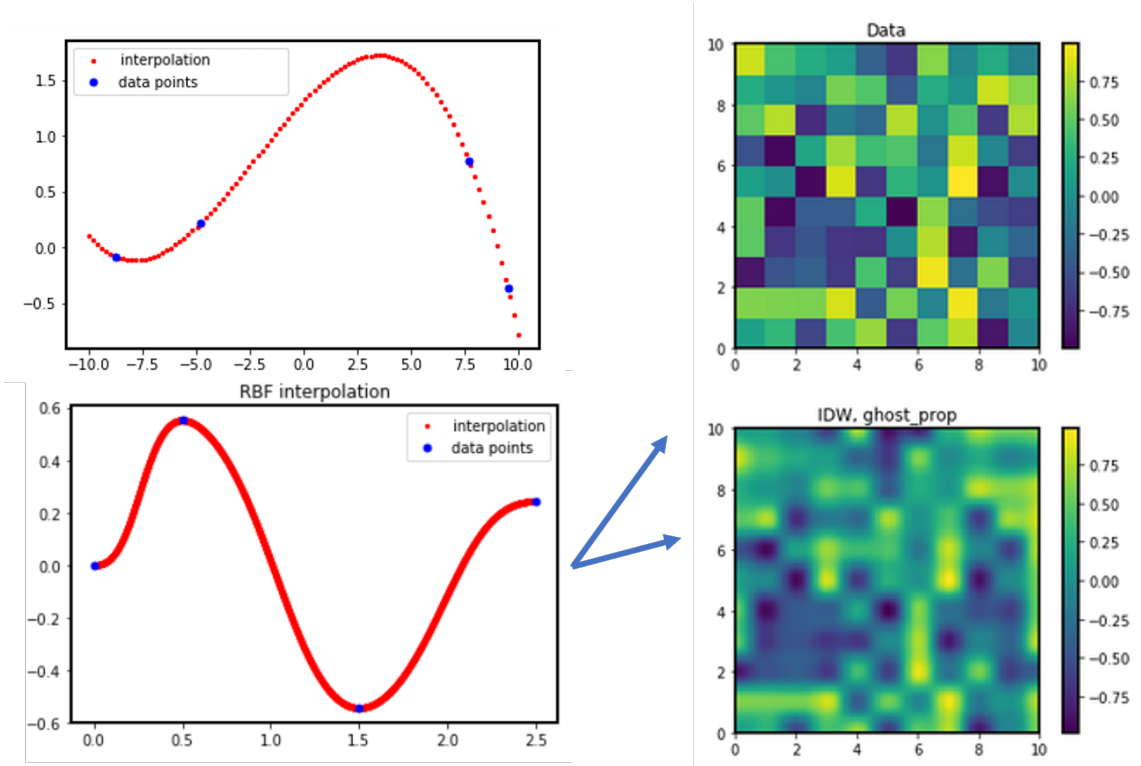


Figure 6.11: Polynomial (top-left) vs. radial-basis-function (IDW with $p = 2$) interpolation between four randomly generated data-points. The two-dimensional version of IDW is presented on the right.

where $d(x, x_i)$ is a distance between x and x_i, the power $p \in \mathbb{R}^+$ is here set to $p = 2$ (other values were also tested), and $\mathcal{N}(x)$ is a normalisation factor defined as:

$$\mathcal{N}(x) = \sum_{i=1}^{N} [d(x, x_i)]^{-p}, \tag{6.23}$$

which makes the interpolation stable as $x \to x_i$. In 1d, the distance is defined simply as:

$$d(x, x_i) = |x - x_i|.$$

In higher dimensions, L^1 and L^2 metrics seem to be reasonable choices. The interpolator defined above is discussed in sect. 3.7.3 of Press et al. (1990). The RBFs have several desirable properties, e.g. the calculation of weights is fast with no solution of linear equations required. The $\mathcal{O}(N)$ complexity is acceptable even for large N. More importantly for this study, IDW does not overshoot function values between the evenly-spaced data points, unlike high-order polynomials (fig. 6.11). The anti-symmetry of the wavefield across the free surface in the extrapolation is imposed, at least approximately, because of the minus sign in eq. (6.20). As all IDW weights are non-negative ($w_j \geq 0$), wavefield $p(G)$ will have the sign opposite to the sum $\sum_i f_i p(F_i)$, which for IDW is dominated by the value in the nearest point $p(F_1)$. Several other interpolators were implemented and tested, including Lagrange polynomials and a windowed sinc function. They are review in appendix .

Stability Note that extrapolation weights f_i (eq. 6.20):

$$f_i = \frac{w_i}{w_0} = \frac{1}{\mathcal{N}(x)} [d(x, x_i)]^{-p} \mathcal{N}(x) [d(x, x_0)]^p = \left[\frac{d(x, x_0)}{d(x, x_i)}\right]^p \tag{6.24}$$

may effectively go to infinity for large $d(x, x_0)$ or small $d(x, x_{i>0})$, making the interpolation unstable. The former case can be ruled out, because the distance in grid nodes between a ghost node of the kth level and the associated intersection is bounded from above:

$$d(x, x_0) < \frac{k}{2}. \tag{6.25}$$

The latter case, however, is possible and should be guarded against. To this end, the minimum distance of the first fictitious node F_1 from the free surface is set to be half a grid spacing h after Hu (2016) and Mulder (2017). If $d(x, x_1) < h/2$, the F_1 is skipped and an extra fictitious node is appended at the other end of the interpolation stencil (fig. 6.10). To further stabilise the scheme, the extrapolation may be posed as interpolation using a mirror-reflection of G. As the free surface effectively acts as a source of wavefield equal to the impinging wave but with inverted polarity, close to the free surface the pressure wavefield p is anti-symmetric along the normal direction, giving:

$$p(R) = -p(G), \tag{6.26}$$

where R is defined as a reflection of a given ghost node across the free surface, such that $|GI| = |RI|$ (fig. 6.8). In this case, instead of eq. (6.19), the following is used:

$$p(R) = w_I p(I) + \sum_{i=1}^{N_f} w_i' p(F_i) = \sum_{i=1}^{N_f} w_i' p(F_i), \tag{6.27}$$

where the apostrophes signify the difference between these weights and the ones in (6.18). The weight w_I does not matter, this term vanishes anyway, as $p(I) = 0$. From (6.27) the value in the ghost node is given by:

$$p(G) = -p(R) = -\sum_{i=1}^{N_f} w_i' p(F_i).$$

Note that this time, unlike in eq. (6.19), the normalisation factor (6.23) does not cancel and may stabilise the interpolation. Note that the distance between R and the intersection:

$$|IR| < \frac{k}{2} \tag{6.28}$$

just like $|GI|$ in eq. (6.25). The analogous distance for the furthest fictitious node is:

$$|IF_{N_f}| > N_f - 1 = k - 2, \tag{6.29}$$

with the last equality implied by eq. (6.30). Combining (6.28) and (6.29) gives:

$$|IR| < |IF_{N_f}|$$

for $k \geq 4$. As only schemes of accuracy order of 4 or more are considered here, the point R will always fall between fictitious nodes. Consequently, eq. (6.27) indeed describes an interpolation problem.

Accuracy Accuracy of interpolation depends on several factors, including number of data points. Mittal et al. (2008) show (their eq. A.5) that considering only the nearest grid nodes results in a truncation error of $\mathcal{O}(h^2)$. In general, using $k/2$ data points on each side of the point of interests gives a kth-order accuracy. Hence, to retain the accuracy of the FD scheme, the numbers N_a and N_f are set to:

$$\begin{aligned} N_a &= k \\ N_f &= k - 1, \end{aligned} \tag{6.30}$$

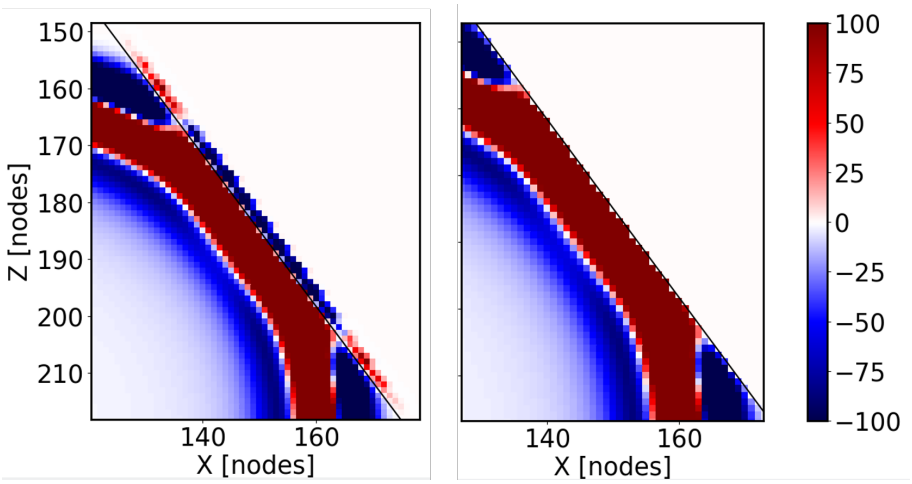


Figure 6.12: Ghost wavefield before (left) and after (right) setting to zero during the vacuum step of the wavefield-extrapolation.

the difference between the two resulting from the extra point (I) included in the extrapolation stencil.
2550 For example, there will be 3 fictitious nodes per ghost node and 4x4 auxiliary nodes per fictitious node (4 in 2d) for a 4th-order accurate FD stencil of 5x5x5 size.

Storage The only data needed to extrapolate the wavefield into a given ghost node are:

- positions and interpolation weights of auxiliary nodes,
- interpolation weights of fictitious nodes.

2555 Note that one does not need coordinates of fictititious nodes once the auxiliary nodes have been found. The extrapolation data can be saved on disk and re-used in multiple FWI runs saving the computational time otherwise spent on projection of ghosts calculation of extrapolation weights which are the most computationally expensive parts of the algorithm.

6.2.7 Extrapolation

2560 In this section, a procedure of setting the wavefield values at the ghost nodes is discussed. In the time-domain formulation, the extrapolation has to be performed at every time step. An iterative solution to the problem of *coupled* ghost nodes is presented. The method requires initial conditions – here the ghost wavefield is set to zero before iterating until convergence. If iterating is skipped, the scheme is equivalent to the *vacuum method* (VM) which may serve as a more affordable alternative to IB
2565 whenever the accuracy is not a critical requirement.

Non-degenerate case The wavefield at ghost nodes is updated by extrapolating the physical wavefield using the weights pre-computed weights of fictitious points and auxiliary nodes f_i and a^i_{jk} discussed in the previous section. This is a simple matrix-vector multiplication. The update proceeds according to eq. (6.22), one ghost-level at a time, starting from the first level. Here, only ordinary
2570 ghosts are updated. The type of ghosts that does not need update are surface-ghosts which store zero pressure wavefield at all time-steps.

Coupled ghost nodes The basic scheme described above corresponds to the situation where ghost nodes do not depend on each other. For some shapes of the free surface, it may happen that some

auxiliary nodes associated with a given ghost node are ghost nodes themselves. An example of such an auxiliary node is denoted by the empty circle in fig. 6.8. In this case, some ghost nodes are coupled and the corresponding system of linear equations for ghost wavefield is non-diagonal. This problem can be solved directly (Lombard et al., 2008) or iteratively (Hu, 2016)[13]. The former is likely to have a smaller operation count but is less memory-efficient and may lead to instabilities (Hu, 2016). For this reason, here, the latter was chosen. As the update is iterative, it needs starting values. The starting values at ghost nodes may be arbitrary, probably without making much difference – here 0 was used because even if the iterating is skipped this is equivalent to the *vacuum method* (VM) and the wavefield will be reflected off the staircased free surface (see appendix 6.A). This is not as accurate as the IB but may suffice in some cases.. Henceforth, setting the wavefield to zero will be referred to as the vacuum step (fig. 6.12). Following Hu (2016), the ghost wavefield update is performed iteratively until the ghost wavefield becomes stationary, i.e. does not change between subsequent iterations:

$$p^{n+1}(G) = C^A P(A) + P^n(G). \tag{6.31}$$

Only one iteration means that only independent ghost nodes needed to be updated at this time step. For realistic topographies studied here less than 5 iterations were required for convergence.

6.3 Point-projection in 3d

In sect. 6.2.5 the orthogonal point-projection was introduced and posed as an optimisation problem and a brute-force method of solution was presented. In 3d, this method is prohibitively expensive in terms of memory and CPU usage. For this reason, a new method based on the *Newton optimisation* scheme was developed. The idea is to apply the Newton method to multiple surface patches within a region known to contain a global minimum. The method was conceived from scratch as no similar approach was found in the existing seismological literature. It was later discovered that the problem is tackled in computer-aided design, robotics, computer graphics etc. (Johnson and Cohen, 1998; Ko and Sakkalis, 2014). The existing methods developed in those fields were found to take a very similar route to the one presented here. The procedure is summarised in alg. 6.2. The details of each step are discussed below.

6.3.1 Surface patches

The first of the method consists in tesselation of the surface with bilinear patches. The choice of bilinear interpolation method (see sect. 6.2.3) is motivated as follows:

- bilinear is the simplest[14] 2d interpolation that allows for curved surfaces. This may seem counterintuitive, but even though the surface is linear along grid lines, in general it is curved (quadratic) in any other direction;
- a bilinear surface has no saddle/inflection points (i.e. it does not change its curvature) in any cross-section. This simplifies a local optimisation problem as one does not need to performed any additional checks of curvature sign to determine if an extreme or inflection point has been reached;
- a bilinear surface never suffers from overshooting which may be the case for higher-order methods (Press et al., 1990).

[13] Other solutions to this problem include shifting the auxiliary nodes below the free surface (Zhao, 2010) nonetheless it does not seem to be as robust as the other two mentioned.

[14] Simplicity of computation is important for clarity and efficiency of the algorithm.

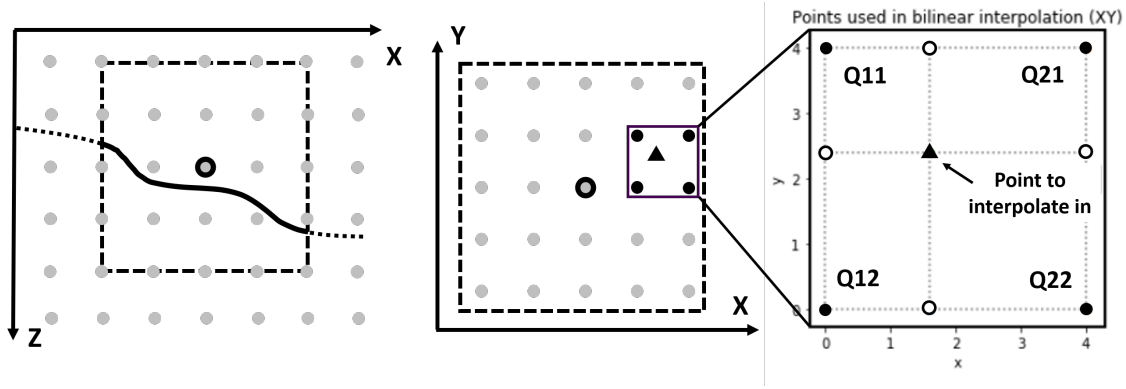


Figure 6.13: Geometry of the 3d orthogonal projection of a point (grey dot with black outline) onto a $z(x, y)$ surface (black curve). The part of the surface enclosed in the search region (dashed square) is denoted by the thick solid line (left). The XY projection of the search region is split into squares with corners being the grid nodes (middle). Each square corresponds to a surface patch generated by a bilinear interpolation between its corners (right). The shortest distance to each patch is found by the damped Newton optimisation. The scheme is readily generalisable to arbitrary surfaces.

The surface is assumed to be sampled on a regular grid and to contain no overhangs (see sect. 6.2.3). This is because a bilinear interpolation is defined on a rectilinear 2d grid (Press et al., 1990). If it is not the case, and it is represented e.g. as an arbitrary point cloud, it needs to be gridded which should not pose any problem. If the surface contains overhangs, one could split it into fragments each of which is an explicit function of two independent variables (not necessarily x and y) of the local coordinate system.

Higher-order patches Options other than bilinear interpolation are e.g. bicubic or any flavour of spline interpolation. A general advantage of bicubic and higher-order interpolation is that the patches may join seamlessly, i.e. one can solve the interpolation problem with extra constrain about the slopes at the edges. Intuitively, the boundaries between patches can be made invisible. If the slopes are chosen to be the same for neighbouring patches, the surface is C^1 smooth globally, i.e. has a continuous first derivative. This is useful if the interpolation is performed for visualisation, i.e. in computer graphics or film industry. For our purposes however, it is not necessary. Moreover, higher-order interpolation may suffer from overshooting as mentioned above.

6.3.2 Spherical clipping

A key step of the shortest-distance search is to exclude, at the very beginning of the procedure, the parts of the surface that are obviously too far to be of interest. This makes the global optimisation problem more localised. Constraining the local search to a limited region is a key step and it gets more challenging with increasing surface complexity. The algorithms dealing with NURBS (e.g. Chen et al., 2009) and similar types of surfaces that allow for overhangs need to ensure that the exclusion criteria are mathematically proved to eliminate only those patches that do not contain a global minimum. Furthermore, in these application, often a computational performance is a major issue as they are used in real-time rendering (Ko and Sakkalis, 2014).

Search region As this study is constrained to rather simple surfaces (see sect. 6.2.3) and time-efficiency is not a great concern, a simple exclusion criterion was derived. As in the 2d case, a search

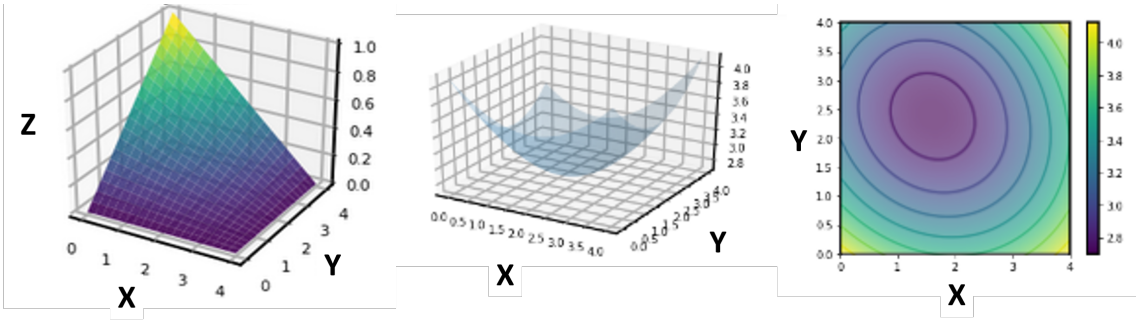


Figure 6.14: Finely sampled surface patch (left), the corresponding objective function (middle), and the projection of the objective function onto the XY plane (right). The objective function expresses the distance between a point on the surface and the reference point (not shown).

Algorithm 6.3 Damped Newton method of finding the shortest distance from a surface patch.

1. Choose a starting model in between the patch corners.
2. Calculate the gradient ∇f. In a case of a bi-linear patch there is an analytical formula but approximate calculation is affordable and more generic.
3. Calculate the Hessian $\mathbf{H}$. The same remark as above.
4. Invert the Hessian. Stabilise if needed.
5. Calculate the model update $\mathbf{H}^{-1}\nabla f$.
6. Check if the updated model stays inside a patch.
7. If not, multiply the update by $1/2$ and go to step 5.
8. Update the model and go to step 1 iterating until convergence (the update falls below a certain, heuristically-derived, threshold) or a maximum no. of iterations is reached. In the latter case, return with an error.

radius r is defined to be equal to the level of a ghost which is being projected onto the surface (eq. 6.15). E.g. for a ghost of the 2nd level, the search is performed up to 2 grid nodes away from the ghost. As non-overhanging surfaces are only considered, this is equivalent to considering grid nodes in the region A_{search}:

$$A_{\text{search}} = \{x_{i-r}, \ldots, x_{i+r}\} \times \{y_{j-r}, \ldots, y_{j+r}\} \tag{6.32}$$

where the ghost node is located at a node (x_i, y_j) horizontally and $\times$ is the Cartesian product. The geometric configuration of the search is shown in fig. 6.13.

6.3.3 Local optimisation

Once the surface patches have been selected, each of them is subject to a local-optimisation procedure that aims to find on it a point of the shortest distance to the ghost node. The Newton method is used because the dimension of the model space is small and thus the Hessian operator is cheap to compute and store. The workflow is summarised in alg. 6.3.

Objective function Let us denote a reference point from which we measure distance to the surface as G. Then objective function f can be defined as:

$$f(\mathbf{m}) = (x - x_G)^2 + (y - y_G)^2 + (z(x, y) - z_G)^2, \tag{6.33}$$

where z is a known function returning a z coordinate of the surface at (x, y), and $\mathbf{m} = (x, y)$ will be referred to as a *model* to conform to the optimisation nomenclature. The goal is to find a model $\mathbf{m}$

such that $f(\mathbf{m})$ is minimal. In a linear approximation, the update of the model is given by the Newton direction:

$$\delta\mathbf{m} = \left[\mathbf{H}^{-1}\right]_{\mathbf{m}_0} \left[\nabla f\right]_{\mathbf{m}_0}, \tag{6.34}$$

where $\mathbf{H}$ stands for the Hessian matrix. See sect. 4.2.3 for the derivation. For the bilinear interpolation, the analytical formulae for the gradient and Hessian of (6.33) are trivial to derive. However, a finite-
2655 difference approximation of the gradient and Hessian was found even more convenient to implement. In the following description, the finite-difference intervals Δx and Δy are assumed to be equal to $h = 10^{-5}$ grid nodes.

Gradient As the model has only two parameters (x, y), the gradient of the objective function is given simply by:

$$\nabla f(\mathbf{m}) = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right) \approx \frac{1}{2h} \begin{pmatrix} f(x+h) - f(x-h) \\ f(y+h) - f(y-h) \end{pmatrix}^T. \tag{6.35}$$

2660 For the sake of completeness, the analytical expressions read:

$$\frac{\partial f}{\partial \xi} = 2(x - x_G) + 2\left(z(x,y) - z_G\right) \frac{\partial z}{\partial \xi} \tag{6.36}$$

where ξ denotes x or y, and:

$$\frac{\partial z}{\partial \xi} = \frac{(\xi_2 - \xi) \cdot [z(Q_{21}) - z(Q_{11})] + (\xi - \xi_1) \cdot [z(Q_{22}) - z(Q_{12})]}{h^2} \tag{6.37}$$

using the same variables as in sect. 6.2.3.

Hessian The Hessian captures a curvature of the objective function f. If f is flat, the Hessian will not change throughout the iterations. For the problem at hand, the Hessian is a 2×2 matrix:

$$\mathbf{H} = \begin{pmatrix} \frac{\partial^2 f}{\partial x^2} & \frac{\partial^2 f}{\partial x \partial y} \\ \frac{\partial^2 f}{\partial y \partial x} & \frac{\partial^2 f}{\partial y^2} \end{pmatrix}, \tag{6.38}$$

2665 where, denoting x and y jointly by ξ:

$$\frac{\partial^2 f}{\partial \xi^2} = \frac{f(\xi + h) - 2f(\xi) + f(\xi - h)}{h^2}. \tag{6.39}$$

Taking a centred finite-difference of both first derivatives (eq. 6.36), one obtains a second-order accurate approximation of the mixed derivatives:

$$\begin{aligned} \partial_{yx} f = \frac{1}{4h^2} [f(x+h, y+h) + f(x-h, y-h) - \\ - f(x-h, y+h) - f(x+h, y-h)] = \partial_{xy} f. \end{aligned} \tag{6.40}$$

The equality of the mixed partial derivatives is implied by the continuity of the second partial derivatives (Riley et al., 2006). The inverse of the Hessian follows from the inverse of a general 2×2 matrix:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1} = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}. \tag{6.41}$$

If the mixed partial derivatives are small compared to the proper ones (i.e. if the Hessian is *diagonally dominant*), ad is much larger than bc yielding $ad - bc \approx ad$ and consequently:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1} \approx \frac{1}{ad} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix} \approx \begin{pmatrix} 1/a & 0 \\ 0 & 1/d \end{pmatrix} = \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix}^{-1}. \tag{6.42}$$

In this case, a diagonal Hessian is a reasonable approximation. Even though the mixed derivatives are easy to calculate and store in this simple case, a diagonal approximation of Hessian that simplifies the calculation even more, was found to converge quickly too. If either a and d are small and could destabilise the inversion, they are set to 1 so that the model update is equal to the minus gradient for this row.

Damping An important caveat of the method is that one needs to make sure that the local optimisation scheme is constrained to a given patch, i.e. the model (x, y) stays between the patch corners, wich is not guaranteed *a priori* by the Newton method. This problem can be seen as a constrained optimisation problem where the constraints are the boundaries of a patch. Here, at every iteration it is checked if the updated model remains inside the patch. If not – the update $\delta\mathbf{m}$ is multiplied by 0.5 as many times as needed to bring it back between the patch corners. Damping the model update $\delta\mathbf{m}$ by a multplicative scalar factor leads to the *damped Newton* method (Nocedal and Wright, 2006).

Convergence Here, the starting model is chosen in the middle of the patch and the optimisation proceeds until the convergence or total number of 50 iterations is reached (the latter case is discussed below). The convergence criterion is defined as:

$$|\mathbf{m}_{i+1} - \mathbf{m}_i| < \epsilon, \tag{6.43}$$

where $|\cdot|$ is the L^2 norm, ϵ is a small number (here heuristically chosen to be 10^{-5}) and $\mathbf{m}_i$ is a model vector at ith iteration. If the minimisation does not converge, the starting model is taken as the final result and the search is performed for the next patch. This seems reasonable, as the lack of convergence indicates that the minimum lies outside the currently investigated patch.

6.3.4 QC of the algorithm

Correctness and efficiency of the above optimisation scheme was extensively tested. For the assessment of the local optimisation, all possible shapes of bilinear patches were utilised. These possibilities correspond to different relative positions of the corners of a patch as well as different positions of the ghost node (the point being projected onto the surface). A benchmark solution in each of the experiments was the minimum distance to a given patch found by a brute-force grid-search.

There are only a few possible configurations of a four-corner surface patch, namely: all corners at the same level (horizontal plane), two adjacent corners lifted by the same amount (tilted plane), and several remaining cases resulting in a curved surface, e.g. two non-adjacent (diagonal) corners lifted (fig. 6.14). As expected, the local scheme failed to converged to a minimum if that was located outside the patch. In such a case, the next patch was considered. Finally, positions of corners and the ghost node were randomised (within a range of plausible values) and the procedure was repeated hundreds of times to check for any unexpected behaviour but none were found.

6.4 Source injection

In this section a problem of injecting sources at/near to a free surface is discussed. Examples of such sources include shallow airgun explosions or waves numerically back-propagated from land geophones by virtue of the *reciprocity principle* (see sect. 5.2.4).

6.4.1 Problem statement

In active-source seismology, sources are usually treated as points (i.e. as infinitisemally small) because their size is typically much smaller than the dominant wavelengths of the energy they produce. Such a source can be added at any point of the model, including at or just below the free surface boundary. The former case is limited to force sources only – even if volume injection at an infinitely thin surface was possible, no waves should be produced anyway because pressure is zero there by definition. Therefore, only shallow volume sources and shallow/surface force sources potentially need to be implemented. In theory, it is straightforward to add a point source arbitrarily close to the free surface, or at a surface remembering that a volume source should be zero there. In practice, however, one always deals with grids of finite spacing, and sources rarely happen to coincide with their nodes. In such cases, a point source has to be extrapolated over neighbouring grid nodes. The nearest-node approximation is usually not accurate enough for full-waveform inversion (Warner et al., 2013b) so the extrapolation is performed onto more than one node. Consequently, the source is no longer a point. It has a finite size, or *spatial support* (if seen as a function). Here, it will also be referred to as a *source region* or a *source cube*. Homogeneity of the source region is required for the far-field radiation pattern to be effectively the same as the one of the point source (Hicks, 2002).

Flat surface If the source region intersects a model interface, the assumption of homogeneity no longer holds. In particular, if it is a free surface, the most extreme contrast possible, the resulting radiation pattern will be incorrect. To correct for this, the part of the source region located above the free surface has to be reflected into the model with a reversed sign (Virieux et al., 2012, p.271). Hicks (2002) discusses this procedure in the case of a flat free surface.

Irregular surface In a general case, the reflection needs to be performed off an irregular free surface, and this poses some challenges described below. This additional effort is necessary as radiation patterns of shallow/surface sources strongly depend on the proximity and shape of the free surface, significantly deviating from the radiation patterns of deeply buried sources. Such radiation patterns have been derived both analytically (Burridge et al., 1964) and experimentally (Gupta and Kisslinger, 1966). Injection of such sources does not work automatically even with the implemented immersed boundary method. This is because there is no built-in mechanism that would put the outlying part of the finite source region back inside the model. The part already in the subsurface will be reflected correctly. However, the part above the free surface, just like any other ghost node, will be set to zero in the vacuum step that follows the source injection. This spatial clipping of the source region will result in an incorrect radiation pattern.. The only special case when it works 'out-of-the-box', is if the source happens to be located at a grid node. Note that this is not an issue for receivers which can naturally make use of the ghost wavefield. This is because detection at a receiver consists in interpolation, and this case is accounted for by the IB automatically. However, in the full-wavefom inversion as it is usually implemented, during the gradient calculation receivers become adjoint sources. So their position relative to the free surface matters and the special injection scheme has to be applie to them as well.

Possible approaches Huiskes et al. (2017) seem to take this into account with their approximate scheme by modifying the source and receiver operators in the same way as the finite-difference stencil.

An alternative method worth considering consists in putting an image point source on the other side of the surface. Such a source has a reverse polarity and is located along the surface normal passing

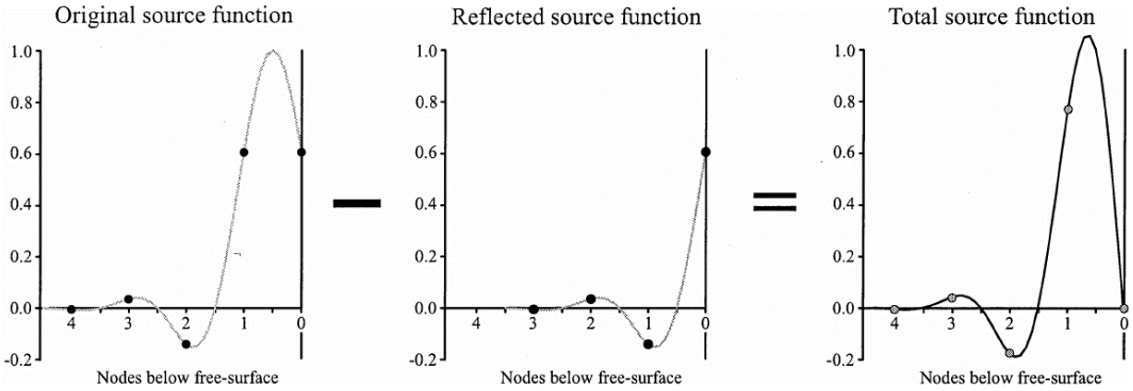


Figure 6.15: Schematic of injecting a source close to a 1d free surface. Here, the method was generalised to 3d, requiring an iterative solution. From Hicks (2002).

Algorithm 6.4 Source injection at/near an irregular free surface.

1. If a composite source is to be injected, decompose it into point sources and proceed one point source at a time.
2. Check if the point source is located at a grid node or not. An offset from a grid node of a few metres may be neglible. If it is a buried source, located at a grid node it will not require any special treatment – start this step again using the next source.
3. Create a cube for the final weigths (*final source region*, see point 5). Initially it may be larger than the one actually used by the FWI code, see the last step.
4. Extrapolate the point source over neighbouring nodes (*source region/cube*) following the method of Hicks (2002) or similar.
5. Reflect (with a reversed sign) a part of the source region that is above the free surface back to the model along a normal vector determined in the immersed-boundary pre-calculation step (sect. 6.2).
6. If the reflection lands at a grid node, add its amplitude to the current value of the final source region at the same node. If not, append it to a list of secondary sources.
7. Use all the secondary source as an input of the step 2 and repeat until amplitudes of all new sources fall below a certain threshold.
8. Clip the size of the final source region excluding nodes too far apart from the original point source, so that it is a good approximation of a point source.

through the original source, at the same distance from the surface. A cube of the image source extending inside the model is added to the original source cube, wherever the two overlap. This avoids iterating at the cost of accuracy loss (surface is implicitly treated as planar within the source region) and a need for solving extra point-projection problem, instead of reusing projection vectors of the ghost nodes.

6.4.2 Iterative solution

An iterative method has been proposed to accurately simulate radiation pattern of a general composite source located near/at an irregular free surface. Its main points are summarised in alg. 6.4 and discussed in more detail below.

The method works for both point and composite sources, as the latter can always be decomposed into a larger number of point source by virtue of linearity of the wave equation. This is the premise of the Green's function method of solving linear differential equations analytically.

The first thing to do is to check if the point source is located at a grid node or not. If the offset is

smaller than the uncertainty of the source position determined in the experiment, such a source can be snapped to a grid and injected in a standard way if it is a buried source. For a source exactly at the free surface, one needs to follow the same steps as for a source offset from grid. Here a relative position-determination precision of 0.05 grid node is used which corresponds to 2.5 metres for 50-m grid spacing.

Next, the point is extrapolated ('spread') over neighbouring nodes. The Shannon's sampling theorem provides an optimal way of doing it – by using a band-limited, discrete Dirac delta. A convenient choice is a sinc function tapered with the Kaiser window. This is because a shape of this window is defined by a single, easy to optimise parameter b (or β). Optimal values of b for a given window width r are given by (Hicks, 2002). By differentiating the source region, one can obtain anisotropic sources, such as dipoles (ch. 5), with a different optimal window shape. Regardless of the source type, the only free parameter to decide upon is the window width. In this study a radius (or half-width) of $r = 3$ is used as compromise between accuracy and compactness. This means we need at least three ghost levels for which we know intersection points. In general, in order to reuse normal vectors found in the immersed boundary method, r must be no larger than the number of ghost levels. The same r for neighbours and the window is used, as outside the window values are zero by definition.

The source cube is then split into parts above, at, and below the surface. One way of doing it is to check for each source node if it is a ghost node. For a non-overhanging surface like here, this is slower than just comparing the depth of each node with surface elevation. But the former method works automatically for arbitrary surfaces. Furthermore, it allows to immediately find the intersection I associated with a given ghost node which is used in the next steps.

Each source-node above the free surface is orthogonally projected onto the surface. Or, using the naming convention from the previous section, an intersection I is found. These were already determined in the immersed-boundary pre-calculation and do not need to be recalculated, provided that the source region does not stick above the ghost-node coating. Given the intersection I, a reflection R of the source node G across the point I is easily obtainable according to a vector equation:

$$R = 2I - G \tag{6.44}$$

The reflection is assigned the amplitude of the source node but with a reversed sign. This could be the end of the procedure for this node, but there is an important complication. If the reflection R does not land exactly at a grid node, which is most likely the case, it needs to be extrapolated over the neigbhouring nodes as well. The finite region of this reflection will almost inevitably intersect the free surface and the whole procedure has to be applied in the same way as it was for the primary source. All reflected sources that do not fall at grid nodes are treated as such secondary sources. For high accuracy, the procedure has to be repeated iteratively until amplitudes of all new sources fall below a certain threshold. To a good approximation all energy now have been correctly injected in the model and is ready for the next time-marching step.

The method has been benchmarked against known radiation patterns of sources located near/at planar free surface rotated by different angles. These tests are discussed in sect. 6.5. In particular, the method correctly reflects the upper half of a dipole associated with a vertical-force source located at a tilted planar free surface. In the same setting, but for a volumetric source, it gives zero radiation, as expected.

The iterative process of extrapolating point sources onto neighbouring nodes and reflecting the outlying nodes along the surface normals will inevitably expand the final source region. The source, however

should stay as compact as possible. It still should be a good approximation of a point source, especially if there are receivers nearby since in reality they almost always record a far-field radiation. What is more, the larger the source, the worse the assumption of the model homogeneity. For these reasons, the source cube should be spatially limited. This clipping is justified also by substantially smaller amplitude further away from the original point source. Another option could be decreasing the sinc window-width when extrapolating secondary sources. This could facilitate a faster convergence and smaller source region at the end.

6.4.3 Implementation

In this study a cube of a taxicab radius equal to 5 is used which is only 2 nodes more than what would be used for source far away from the free surface. It was determined heuristically, looking at a typical spatial distribution of amplitudes away from the original point source. Such a $10 \times 10 \times 10$ cube is found for all sources and receivers. It contains the source extrapolation weights and associated grid coordinates. Some of the weights are always zero because it is much more convenient to store a cube, even if it contains redundant nodes lying above the free surface.

In the `FULLWAVE3D` acoustic code the models and wavefields are defined on the same grid and stored as one-dimensional arrays, so only two numbers are needed to fully characterise every source node – a weight and a grid index. The injection proceeds rapidly by simply adding the amplitudes (source weight times source wavelet at the current timestep) to the current wavefield at all source node positions.

The method has been implemented as a standalone `Python` code that reads in a list of sources and receivers and immersed-boundary ghost data. It outputs two ASCII files – `SourceData.txt` and `ReceiverData.txt` containing a list of injection cubes for each source and receiver, respectively, with the IDs or in the order meaningful for the FWI code that will read it in. Such implementation is decoupled from any specific FWI scheme and only requires modification of a thin API to port it to a new code.

6.5 Numerical tests

6.5.1 Effect of slope

In this section, the immersed-boundary method is shown to give more accurate results than the air-layer method when simulating reflections from a tilted planar free surface, regardless of the tilt-angle.

A 2d planar free surface inclined at different angles was analysed (fig. 6.16). The analytical solution in this case is trivial, consisting of a primary impulse of a waveform discussed in sect. 5.B) and a reflection with a reversed polarity. No physical diffraction is involved. If the relative position of the free surface and sources and receivers is kept fixed, the solution should be independent of the tilt angle. This allows to benchmark the method of simulating a tilted free surface against the zero-slope case, with most other sources of errors the same in both cases. There are, however, a few caveats to take into consideration: (1) rotation of acquisition geometry needs to be accurate and preserve the distance from the free surface (2) the model needs to be deep enough to contain steeply dipping surfaces (3) the distance from absorbing boundaries needs to be the same or large enough in each case (4) in general, rotated stations will land between grid nodes[15] so it is good to bear in mind the amplitude and

[15] One could instead consider a few special angles (characteristic of Pythagorean triangles) where grid nodes remain grid nodes after rotation.

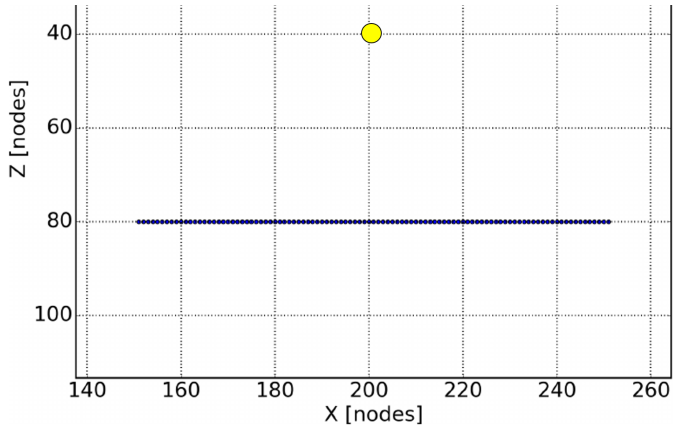


Figure 6.16: Acquisition geometry of the experiments using the flat/rotated free surface: a single source and 101 receivers (yellow and blue dots, respectively); free surface at node $z = 0$. Relative positions and distances between points do not change under rotation of the frame of coordinates around its origin – this property was used to compare results for the flat and rotated free surface.

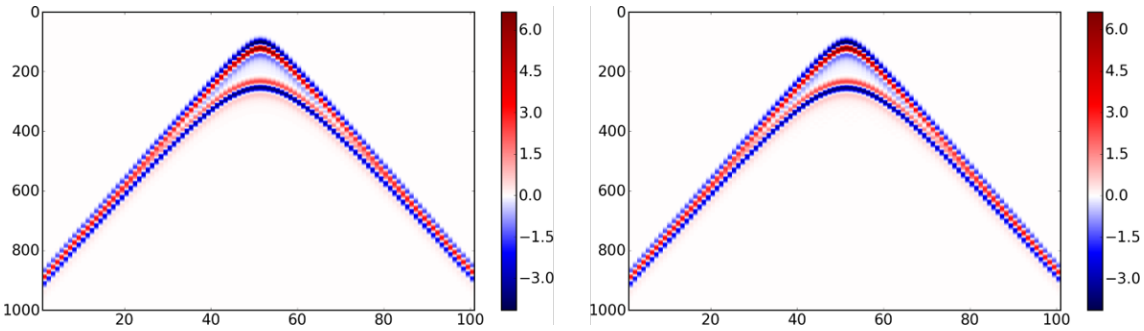


Figure 6.17: Shot gathers from the numerical test involving a horizontal planar free-surface using (left) the default `FULLWAVE3D` method and (right) the immersed-boundary method. The relative positions of sources, receivers and the free surface were the same in both cases.

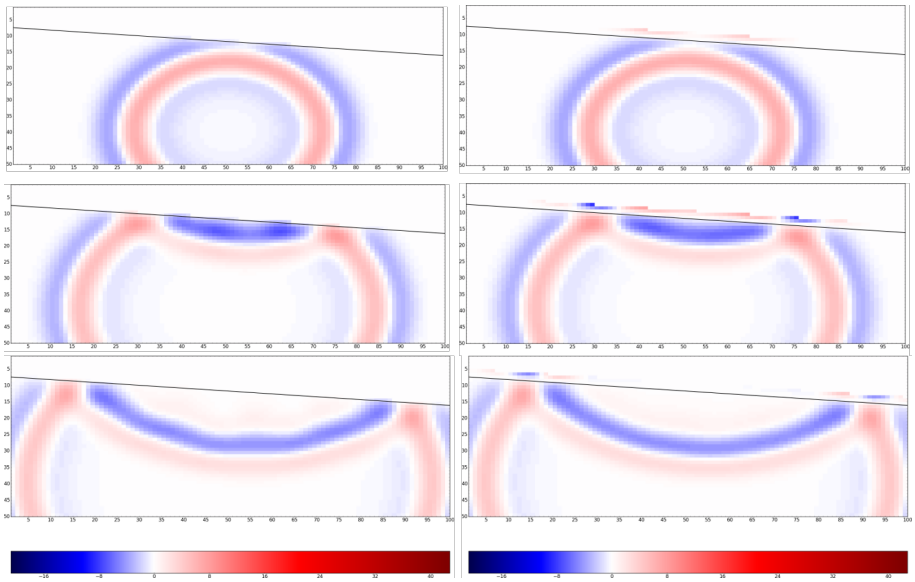


Figure 6.18: Wavefield snaphots from the numerical test involving a planar free-surface tilted by 5° from the horizontal: (left) air-layer method; (right) immerse boundary method.

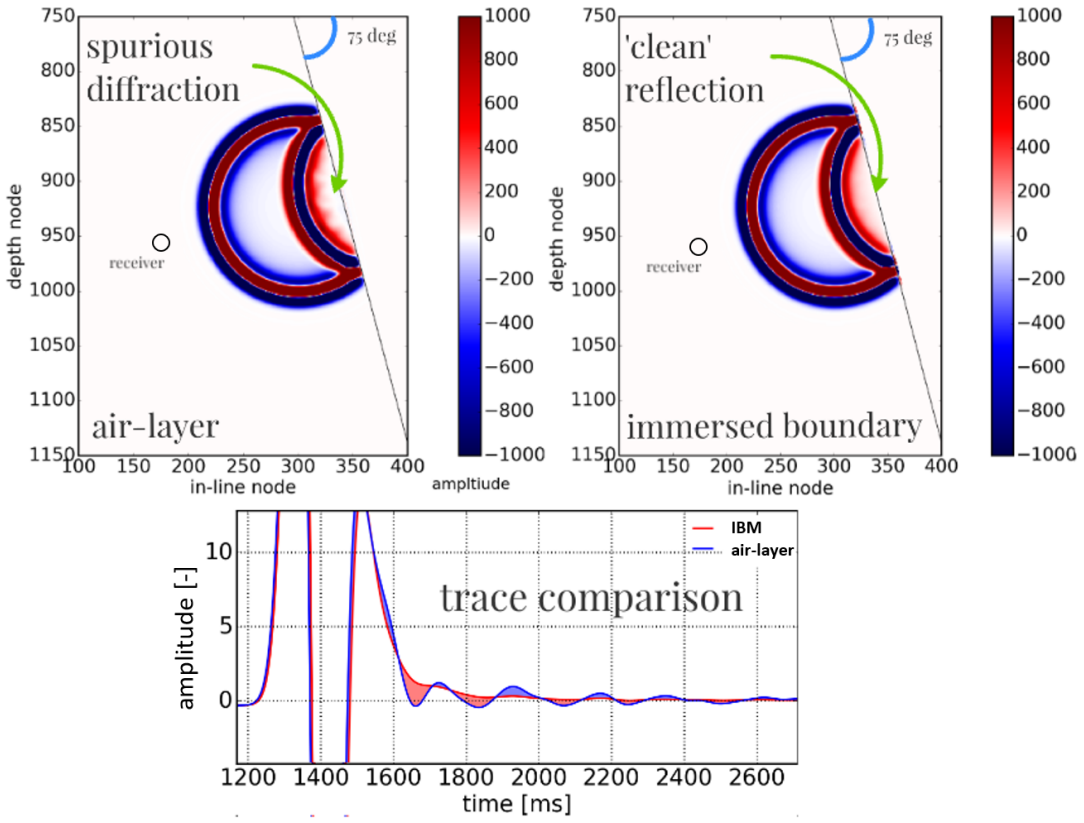


Figure 6.19: Wavefield snapshots from the numerical test involving a planar free-surface tilted by 75° from the horizontal: (top-left) air-layer method; (top-right) immerse boundary method. Note the spurious diffraction pattern resulting from the staircase-approximation for the air-layer method. Bottom – wave reflected from the free surface recorded at the central (51th) receiver for: (blue) air-layer method; (red) immersed-boundary method.

phase errors associated with the Hicks scheme.

First, results for the flat free surface were compared to the default `FULLWAVE3D` code implementation and proved to be nearly identical (fig. 6.17). Then, the whole geometry was rotated by the angle $\alpha \in \{5, 15, 30, 45, 60, 75\}$ deg.[16] and juxtaposed against the air-layer method where the boundary condition was imposed through assigning $v_P = 343$ m/s above the free surface.

In the air-layer method, the amplitude of artificial scattering is prominent already for angles as small as 5 deg. (fig. 6.18). In contrast, the immersed-boundary method displayed distinctively higher accuracy which was almost insensitive to the angle of rotation (fig. 6.19).

6.5.2 Effect of curvature

After testing performance of the method with planes, a natural next step is to test non-zero-curvature surfaces. The immersed-boundary method is shown to be long-term stable when applied to free surfaces with large curvatures.

A sinusoidal shape of the surface was chosen for several reasons. First, sine is a smooth function, so no scattering off sharp edges obscure detection of artificial scattering due to surface discretisation error.

[16] Note that the topographic effect repeats with a period of 45°. For example the 75° case should be the same as the 15° case.

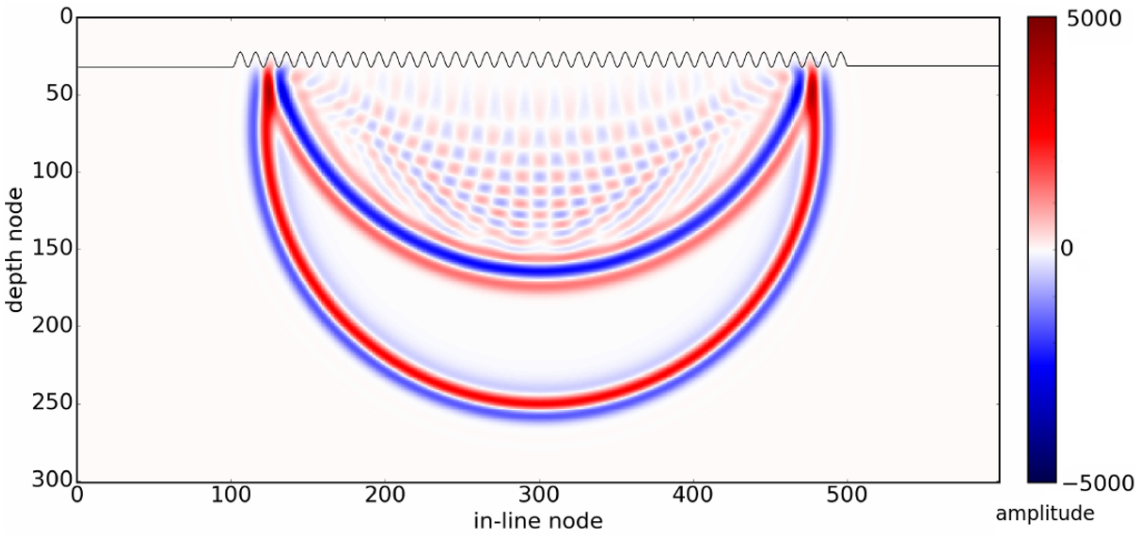


Figure 6.20: Wavefield snapshots from the numerical test involving a sinusoidal free surface with the wavelength of 10 grid nodes.

E.g. an arc of circle tested by Hu (2016) does not meet this criterion. Second, its curvature can be easily calculated and modulated by tweaking either the amplitude A or the wavenumber k. If the intent is to keep the curvature approximately the same for different surfaces, the A/k^2 ratio should be constant.

A large number of combinations of f_{peak} of the source wavelet, k, and A and relative positions of source and surface was analysed. In particular, highly irregular sine-shape topography were tested for a fixed amplitude and various wavelengths of the sine function down to 10 nodes per sine-wavelength, significantly fewer than used by Lombard et al. (2008). Simulation times up to 10 s (10000 time steps) were investigated. No change of accuracy or stability was observed (fig. 6.20). These observations were further quantified by a suite of convergence tests described in appendix 6.C.

6.5.3 Realistic topography

In this section, the immersed-boundary method is shown to give a physically plausible wavefield scattered off the real topographic surface, both in 2d and 3d. The tests were performed using the topography of the Santorini caldera (Hooft et al., 2017), a Ricker source wavelet of $f_{\text{peak}} = 3$ Hz ($f_{\max} \approx 5$ Hz), and a homogeneous model with $v_P = 1500$ m/s. The results for 2d and two 3d settings are shown in fig. 6.21, 6.22 and 6.23, respectively.

Qualitative analysis of the results showed that the method is stable for simulations including the entire Santorini caldera (fig. 6.23) and that it predicts real diffracted waves related to the presence of steep caldera cliffs (fig. 6.21).

6.6 Field-data example

No field-data application of the immersed-boundary has been published to date in the seismological literature. To fill this gap, a marine-land dataset at Santorini volcano acquired in the PROTEUS experiment (Hooft et al., 2017) was used.

An explicit, time-domain solver of the acoustic, isotropic, 3d wave equation of the `FULLWAVE3D` package (Warner et al., 2013b) was used for modelling. A compact ($5 \times 5 \times 5$) operator was used with a

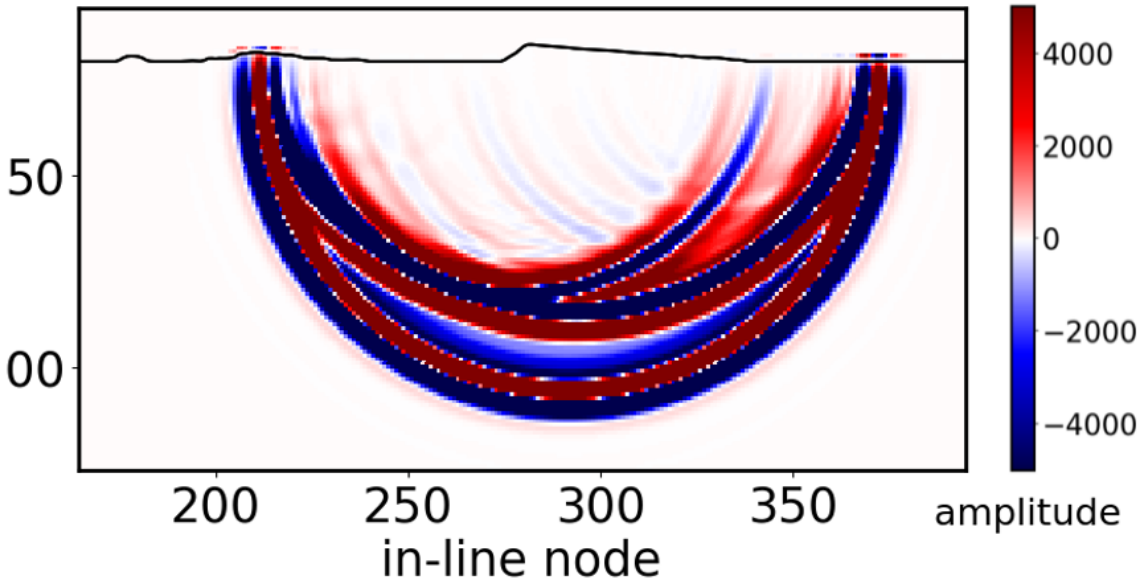


Figure 6.21: Wavefield snapshot after 3 s of the simulation using a 2d topographic surface cut horizontally through the center of the Santorini caldera, homogeneous model with $v_p = 1500$ m/s and Ricker source wavelet with $f_{\text{peak}} = 3$ Hz.

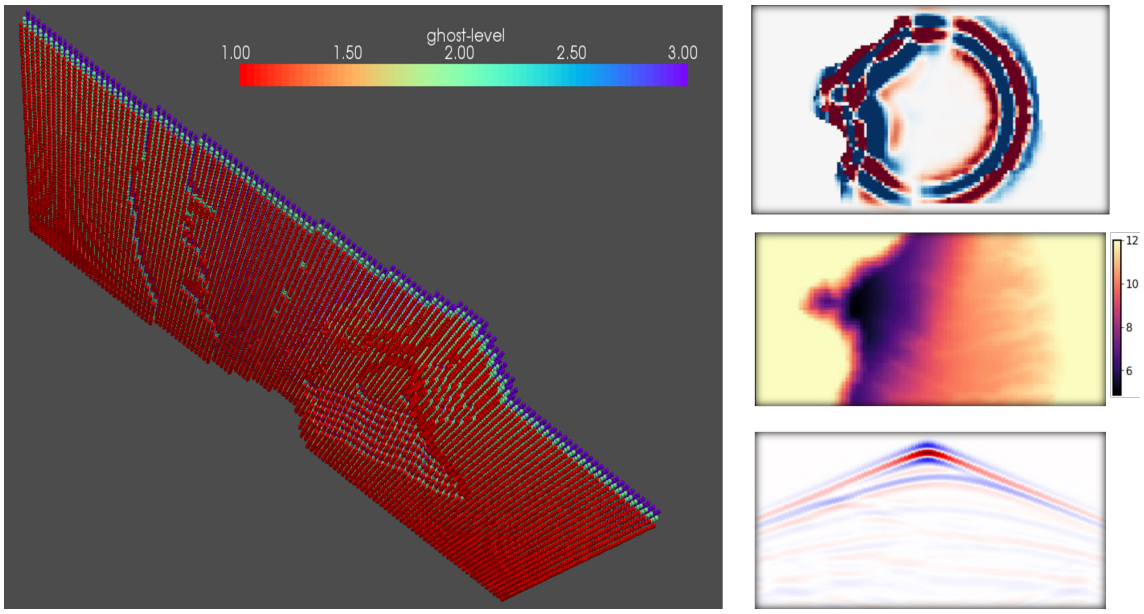


Figure 6.22: Numerical experiment involving a 3d fragment of the Santorini caldera: (left) three levels of ghost nodes coating the free surface; (top-right) wavefield snapshot sliced above the sea-level – the righthand part of the wavefield with the inverted polarity is the ghost wavefield; (middle-right) free surface topography measured in grid nodes below the top of the model; (bottom-right) data recorded by receivers located in the in-line direction across the center of the domain, below the sea level.

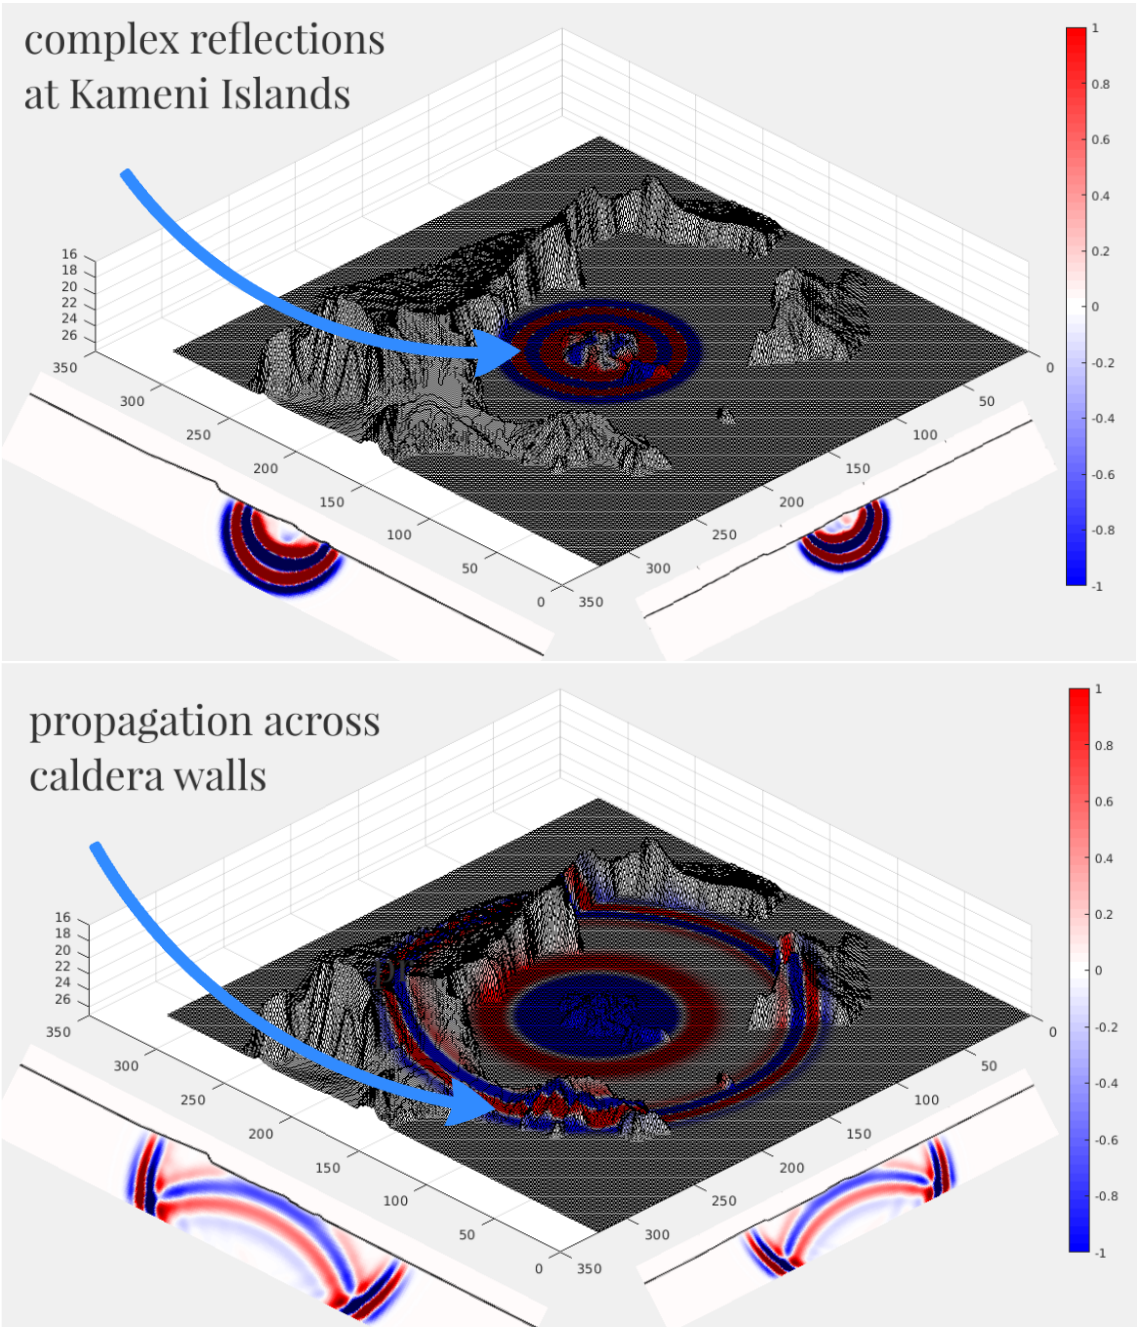


Figure 6.23: Wavefield snapshots of the simulation using a 3d topographic surface of the Santorini caldera. The 2d cross-sections were cut through the centre of the Nea Kameni island in the middle of the caldera. All boundaries are reflecting to test the stability of the scheme in the presence of strong multiples.

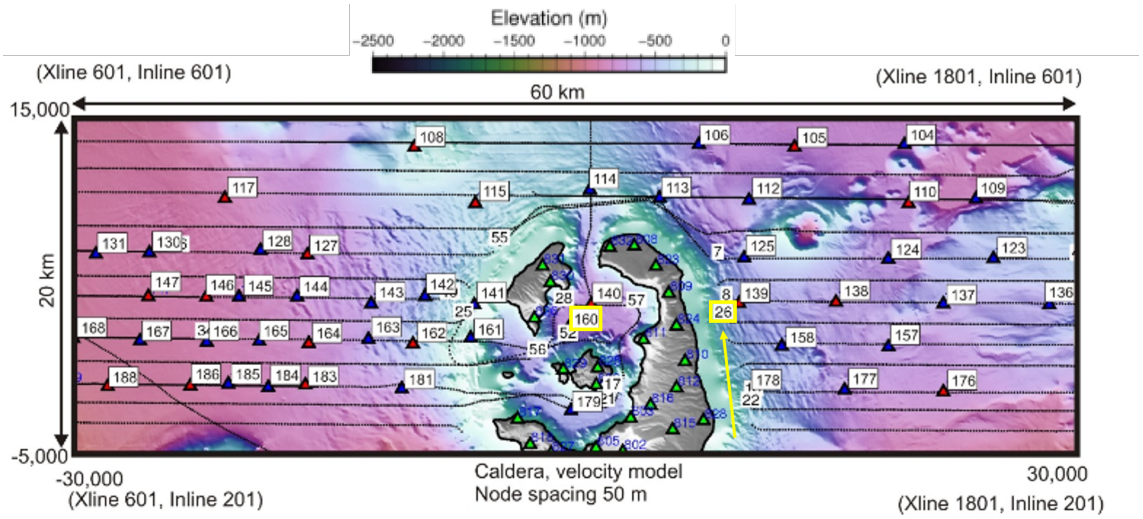


Figure 6.24: Acquisition geometry of the first field-data example. IDs of the OBS and the shot-line used in fig. 6.25 and 6.26 are highlighted in yellow. Modified after J. Morgan and M. Paulatto (unpublished).

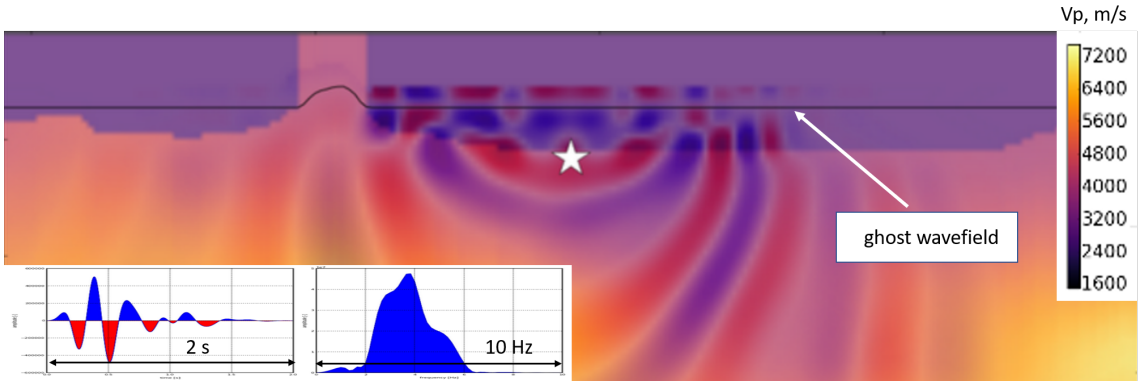


Figure 6.25: Wavefield snapshot from the first field-data example, superimposed over v_P model, sliced in the in-line direction through the source (OBS 160). Inset – data-derived source wavelet in the time (left) and frequency (right) domain.

nominal accuracy of 4th order in space, 2th order in time, least-squares-optimised to give the effective accuracy of about 6th order in space, 4th order in time. This sets an approximate criterion for spatial discretisation to about five (down to four) grid nodes per wavelength. For the chosen, low-order stencil, stability considerations constrain the ratio $\mathrm{d}x/\mathrm{d}t$ to be at most $1/2$. Values of $\mathrm{d}x = 50$ m and $\mathrm{d}t = 0.0025$ s (2.5 ms) were used. The former allows for accurate modelling in sea water with source frequency up to 6.5 Hz.

6.6.1 Long-offset OBS data

In the first case study, the effect of irregular land topography on long-offset OBS data was analysed. A large domain of $60 \times 20 \times 6$ km comprising most of the Santorini caldera and 49 OBS stations (reciprocal sources) was considered (fig. 8.5). Using $\mathrm{d}x = 50$ m, the model reached $1201 \times 401 \times 121$ grid nodes counted without absorbing boundaries which were 20-node thick each. Simulated time was 5 s corresponding to 2000 time steps.

To determine the effect of the land topography on the OBS data, two different v_P models were used

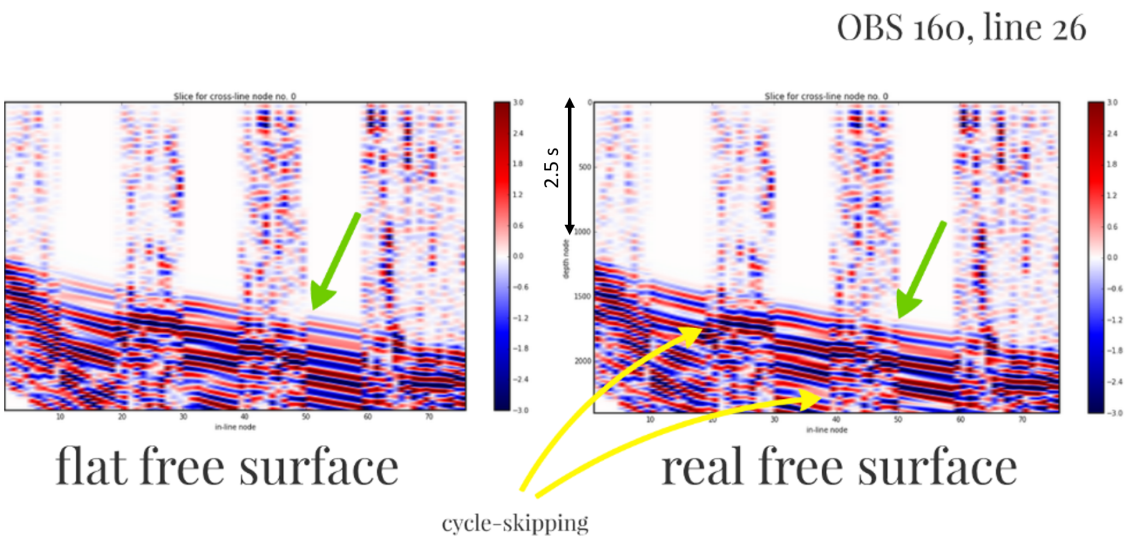


Figure 6.26: Interleaved observed and synthetic data in the first field-data example shown for OBS 160 and shot line no. 26. Left – model with the flat free surface. Right – model with actual topography simulated using the immersed-boundary method. Vertical axis – no. of time-samples ($\mathrm{d}t$ was 0.0025 s).

in the forward modelling. In the first, the topography was fully included, with the sea level at the 18th node from the top of model. The second model had the topography chopped at the sea level and comprised 113 nodes in the vertical direction.

The IB simulation proved to be stable in the presence shallow model heterogeneities reaching the steep cliffs of the Santorini caldera (fig. 6.25). No clear spurious diffractions related to the free surface can be identified. The synthetic data calculated using full topography simulated with IB match the field data better than in the flat-topography case (fig. 6.26). Unlike in the latter, the early arrivals are not cycle-skipped. This shows that approximating the land topography with a flat surface is not sufficient for FWI, even for long-offset OBS data.

6.6.2 Short-offset land data

In the second case study, simulation of short-offset land data recorded around the Kameni islands inside the Santorini caldera (fig. 6.27) was performed. The main objective was to test the injection of surface sources (sect. 6.4).

A small domain, about $4.5 \times 6 \times 5$ km in size, comprising 3 land and 3 OBS stations, was used. Here, only the land stations are analysed. In the simulation they were used as reciprocal vertical-force sources (see ch. 5) injected with the scheme described in sect. 6.4. Fig. 6.28 shows the wavefield snapshots zoomed at all three reciprocal sources. The radiation pattern is characteristic of the land injection with ghost wavefield above the free surface, following the topographic relief, and the ordinary wavefield reflected into the subsurface.

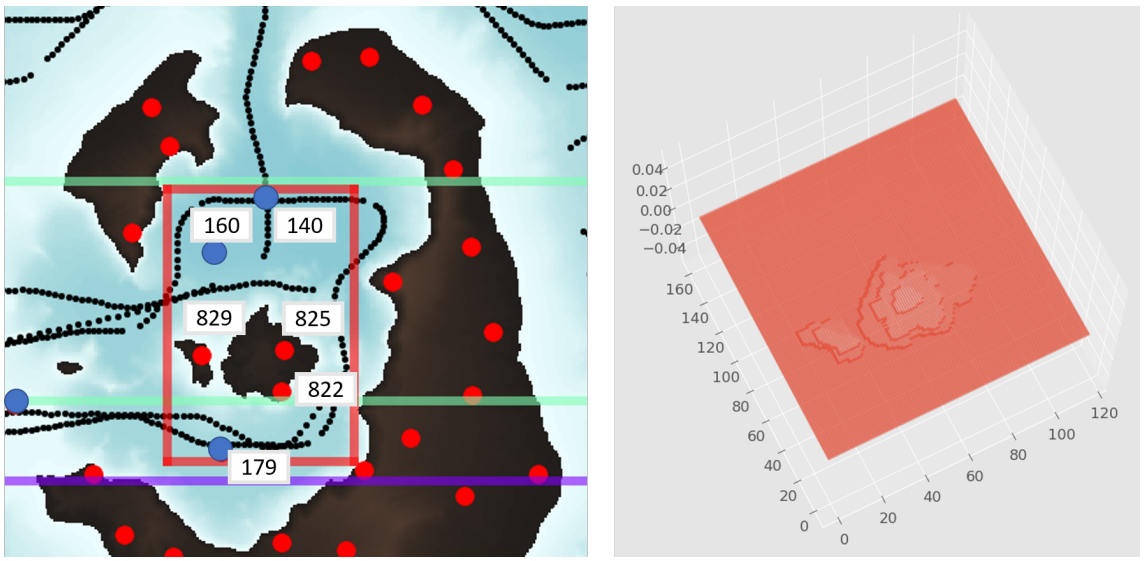


Figure 6.27: Acquisition geometry (left) of the second field-data example (red rectangle), on land stations were analysed (red dots). Right – position of ghost nodes of teh first level.

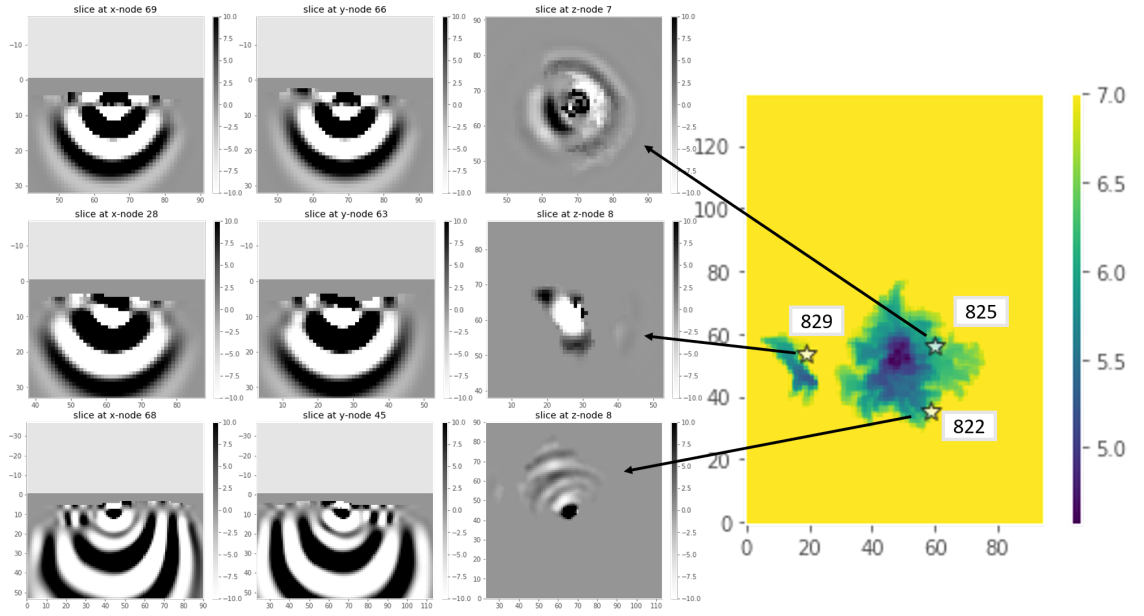


Figure 6.28: Wavefield snapshots from the second field-data example sliced through stations 825, 829, 822 (from top to bottom) with YZ, XZ and XY planes (from left to right).

6.7 Conclusion

6.7.1 Computational cost

The computational cost of the implemented immersed-boundary method is split between the pre-calculation of the ghost data and source weights, and the time-stepping. The former is done once per case study, but the latter is repeated at every FWI iteration and is a possible source of extra computational cost. Both are analysed below in more detail.

Extrapolation data The calculation of ghost data took less than hour on an Intel(R) Core(TM) i7-6700 CPU @ 3.40GHz machine with 32 GB RAM, for the largest Santorini models discussed in sect. 6.6.1 using three ghost levels. In these experiments ghost nodes were found in the whole extended domain, including the flat sea surface and absorbing boundaries which is much more than is actually needed normally. This provides room for reducing this time easily, by contraining the ghost calculation to the land area (see sect. 6.2.4). The absorbing layers should not be depended upon anyway, as they are runfile parameters that should be flexible to change without a need of re-calculating extrapolation weights. Note that these tests were run as a standalone `immerse` code on a local machine. It can, instead be run on a cluster from a FWI code, as such a routine is provided by the interface module of the `immerse` package. Once the pre-calculation has been done, a binary file with necessary data is dumped and the next time the FWI code is started, it will read the file instead of redoing the precalculation.

Sources and receivers The calculation of source weights is even faster. For Santorini this is because the prevailing number of stations is the airgun shots which require only one iteration of the source-cube calculation (see sect. 6.4) as the sea surface is assumed to be flat. The OBSs do not require iterating at all, as they are deep enough so that their source region do not intersect the free surface. The only stations that need more than one iteration are the land stations, but their number is usually much smaller than the OBSs.

Parallelisation If the computational performance of the pre-calculation is still of concern, e.g. in a case of very large land model, or multiple, non-stationary irregular sea surfaces (see sect. 6.7.3), it can be optimised as both, ghost-data and source-weights calculation are easy to parallelise. In the case of the ghost-data, weights of different ghosts are effectively uncoupled[17] and can be parallelised using `OpenMP`. One needs to split one of the bottleneck-loops, e.g. point-projection loop, into `nthreads` chunks. For source-weights, the code could be either ported to `Fortran` or optimised using `Numba` which seems like the most straightforward way to do.

Memory In terms of memory, the size of the ghost data is proportional to:

$$\frac{k}{2}((N^{1/3})^2)^2 = \frac{k}{2}N^{4/3}, \tag{6.45}$$

where N is the number of grid points and k is the scheme order. The number of ghost levels is unlikely to exceed $k = 5$ for most accurate schemes used today. The size of source weights is proportional to

$$((N^{1/3})^2)^3 = N^2, \tag{6.46}$$

since usually a fixed, scheme-independent width of the source region is used.

[17] The coupling is addressed implicitly by the iterative update.

Extrapolation During time-stepping, a possible bottleneck introduced by the IB method is the iterative update. A computational complexity of a typical d-dimensional time-domain FD code is:

$$\mathcal{O}\left(N^{\frac{d+1}{d}}\right), \tag{6.47}$$

where N is number of grid points. In 3d, the complexity of the time-stepping is then of the order:

$$\mathcal{O}\left(N^{\frac{4}{3}}\right). \tag{6.48}$$

A number of ghost nodes scales as a surface area $N^{d-1/d}$ and is proportional to the scheme order k which one writes jointly as $\sim \frac{k}{2}N^{d-1/d}$. Number of times steps scales as $N^{1/d}$ because of the CFL criterion. Denoting a number of iterations for the ghost-update to converge as n, one finally obtains the computational complexity of the immersed-boundary method as $\mathcal{O}\left(n\frac{k}{2}N^{d-1/d}N^{1/d}\right)$ or in a simplified form:

$$\mathcal{O}\left(\frac{k}{2}nN\right), \tag{6.49}$$

where n is the number of iterations in the recusrive update, and k is the order of the FD scheme. Comparing eq. (6.48) and (6.49), and remembering that N is a total number nodes (i.e. $N \sim a^3$, where a is an average number of nodes in a single dimension), it becomes clear that the cost of the immersed boundary is a small fraction of the basic time-stepping algorithm, even for $n > 10$, which is unlikely to be needed even for most complex examples of real surfaces.

6.7.2 Summary

The improved immersed-boundary method of Hu (2016) was extended to handle realistic 3d geometries and seismic sources located arbitrarily close to the free surface. It allows for accurate simulation of the irregular free surface using finite-differences addressing two important problems: 3d orthogonal point-projection and injection of source energy near/at the free surface. Experience shows that the key steps of the method are – representation of the surface consistent across the whole workflow, accurate point projection and stable non-overshooting interpolation method used in the wavefield extapolation. The method is both *usable* and *useful*, as summarised below.

Usability It can be added to existing FWI codes with minor modifications thanks to modular design of the code. The method was implemented in `Fortran90` as a standalone program performing the precalculation of extrapolation data and a thin API to interface it with third-party finite-difference solvers wave-equation solvers. The implementation was tested on a highly parallel, industry-standard `FULLWAVE3D` code by Warner et al. (2013b) and a 3d marine-land field dataset from the Santorini volcano (Hooft et al., 2017).

Usefulness Numerical experiments proved the method to be stable, accurate and have a low computational cost compared to the wavefield propagation far from the boundary. Its superior accuracy relative to flat approximation has been shown for cross-borehole synthetic study and surface-aquisition field-data example. Its superior accuracy relative to vacuum method has been demonstrated in the tilted-plane numerical test (fig. 6.19) and its stability for the sinusoidal surface and all boundaries reflecting. With appropriate data processing, acoustic FWI augmented with immersed boundary and accurate source injection should allow for improving models when inverting challenging land or mixed land-marine datasets.

6.7.3 Next steps

Arbitrary geometries Handling caverns, overhanging walls, mines, karst, bone surfaces (in medical imaging) i.e. allowing the free surface to be an arbitrary surface, not necessarily the explicit function $z = f(x, y)$, can be easily achieved by modifing the algorithm of classifying the nodes. The dot product with an outward normal vector of the nearest surface patch will discriminate between the interior/exterior (Mittal et al., 2008).

Time-varying surfaces In general, the movement of the sea-surface can be ignored when inverting marine seismic data. In high-frequency studies of the near-surface, for example when locating offshore windfarms, the movement of sea between shots affects the source and receiver ghosts and needs to be taken into account. This creates a large amount of information to pass around. Additionally, for sources and receivers close/at the free surface, one needs to carefully calculate spreading factors. In the case of multiple surfaces there has to be associated metadata providing a link that connects a surface with a given source/receiver.

Rock interfaces In crustal seismology usually no model interface is known *a priori* with accuracy better than grid sampling Δx except for the seabed Warner et al. (2013b). Typically it is dealt with in an approximate manner using some sort of homogenisation approach which usually involves interpolation (Symes and Vdovina, 2009; Hobro and Rickett, 2014; Yao et al., 2017). Handling model-parameter discontinuities was not analysed in this study. The *jump conditions* at the interface can be accurately simulated using a very similar, immersed-interface method (Lombard and Piraux, 2004).

Anisotropy The VTI anisotropy is straightforward to account for with the only difference from the isotropic case being additional wavefields to handle. The TTI anisotropy, however, involves calculating mixed partial derivatives which may pose some challenge for the method.

S waves Modelling S-waves is typically done on staggered grids. It is challenging because in some positions of stencil relative to the surface, 0th-order boundary conditions are never imposed on the velocities, leading to exponentially growing oscillations near the boundary (Lombard et al., 2008).

Acknowledgements

The author wants to thank Michele Paulatto for the invaluable introduction to the method and his 2d code used as a benchmark at the early stage of developing `IMMERSE`. The author is also grateful to Adrian Umpleby for helping to understand the inner workings of `FULLWAVE3D` software, and his encouragement to implement the source-injection scheme. This work also benefited from discussions with and talks by Mike Warner as well as Rhodri Nelson, Edward Caunt and Gerard Gorman.

Supplementary information

A supplementary `Jupyter` notebook is available at: `http://github.com/kmch/PhD-Santorini-FWI/notebooks/`. See sect. 8.6 for more information.

Data availability

See the analogous section of ch. 8.6.

Code availability

The code will be made available at http://github.com/kmch/IMMERSE as soon as the FULLWAVE consortium gives a permission to do so. All contributions are looked forward to.

Appendix

6.A Vacuum method

The standard *vacuum/air-layer method* (Moczo et al., 2014) requires: (1) a thin layer of 'vacuum' coating the model (2) a thick absorbing layer coating the vacuum layer. The vacuum layer needs to be at least k nodes thick, where k is the order of the FD scheme (half-width of the FD stencil) and have a very low value of the acoustic impedance ρv. This is typically achieved by setting the density close to 0 as the velocity needs to be rock-like to retain the numerical stability, unless small Δx or large Δt are used which is rarely a feasible option.

Instead of modifying the velocity and/or density model, an alternative approach is proposed here. It consists simply in setting the wavefield above the free surface to zero. This is similar to the modified image method of Robertsson (1996) who proposed the particle velocity to be set to zero above the free surface. Although the staircasing error is not avoided, this approach is fast and may be accurate enough e.g. when simulating OBS data or when finer grids can be afforded (see appendix 6.C). As a 'by-product' of the immersed-boundary method (see sect. 6.2.7), it does not require any additional coding. In the `IMMERSE` code, the vacuum method is activated by setting the `vacuum` runfile parameter to 1.

6.B `IMMERSE` software

The immersed-boundary method described in sect. 6.2 and 6.3 was implemented as a stand-alone package written in the `Fortran90` programming language.

For portability, the code was split into two basic units: a pre-calculator and an extrapolator (compare alg. 6.1). The latter has a form of a thin interface that uses the output from the pre-calculator and extrapolates the wavefield. It can be quickly rewritten to other languages like `Python`. This interface

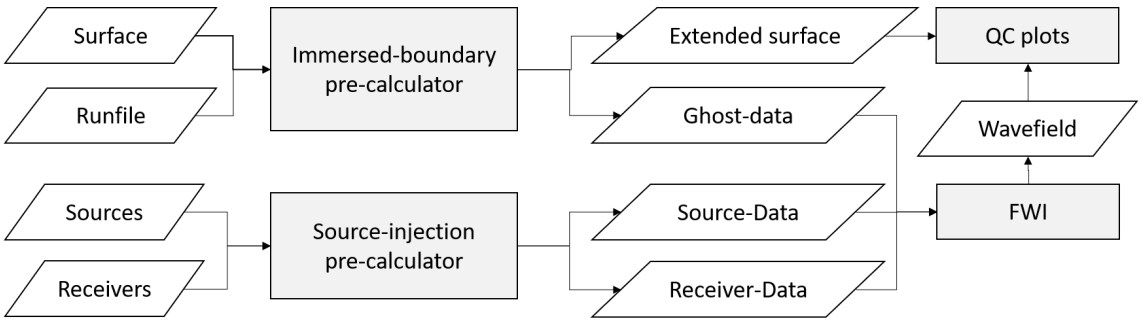


Figure 6.29: Flow of data in the `IMMERSE` software. The source-injection pre-calculator is implemented in the HicksPy code (sect. 5.6).

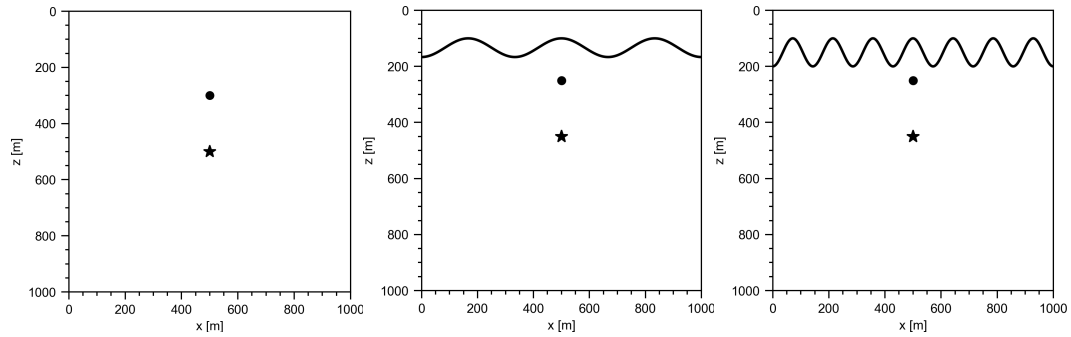


Figure 6.30: Acqusition geometry of the numerical experiments performed to study the convergence of the IB method. The cross-sections at $y = 500$ m are shown. The black star, circle and curve denote the source, receiver and irregular free surface, respectively.

was used to extend functionality of `WAVEFOR2D` (L. Guasch, unpublished) and `FULLWAVE3D` (Warner et al., 2013b) codes beyond purely-marine experiments. The latter package has been widely used for marine data by the exploration industry, and now it can also be applied to land, or marine-land datasets, such as the ones acquired in the PROTEUS (Hooft et al., 2017) or Hikurangi (e.g. Gray et al., 2019) experiments.

6.C Convergence tests

A numerical scheme is said to be convergent if, for an increasingly refined grid, the corresponding series of solutions approaches a certain limit. If this limit is the solution of the continuous equation being solved, the scheme is consistent, and the rate of convergence defines the accuracy of the scheme (Strang, 1986).

The convergence of the implemented immersed-boundary method was investigated by performing a suite of numerical experiments using increasingly refined grids. The same `FULLWAVE3D` kernel was used as the one for the inversion of the PROTEUS field data. A simulated time of 2 seconds was discretised using a 0.001 s time interval. A fixed domain of a homogeneous ($v_P = 1000$ m/s) model extended from 0 to 1000 m in all directions of a 3d regular grid. Three different grid intervals: 10 m, 5 m and 2.5 m were used, resulting in a $401 \times 401 \times 401$ grid for the last case. Two types of an irregular free surface were investigated, both constant along Y and sinusoidal in the XZ plane, with the peak-to-peak amplitude of 100 m and the wavelength along X of 333 m and 143 m, respectively. All other boundaries were also reflecting but flat. A single Ricker source (with a peak frequency at 4 Hz) located at $(500, 500, 450)$ m and a single receiver at $(500, 500, 250)$ m were used in all experiments (fig. 6.30).

The reference 'true' solution was defined as the solution obtained for the finest (2.5-m) grid. The error was defined as the maximum absolute difference between a given solution and the true one. To assess the effective accuracy of the scheme in the studied domain, a series of simulations with the default, flat free surface at the top of the model was also performed. Halving the grid resulted in the maximum error decreased ~10 times (fig. 6.31, top). The performance of the immersed boundary method was compared against the vacuum mode (appendix 6.A).

For both surface shapes, the immersed boundary turned out to have superior accuracy, with the larger difference in the case of the rougher topography. Apart from the faster convergence, the IB method was characterised by significantly smaller absolute errors than the vacuum case. The difference was more

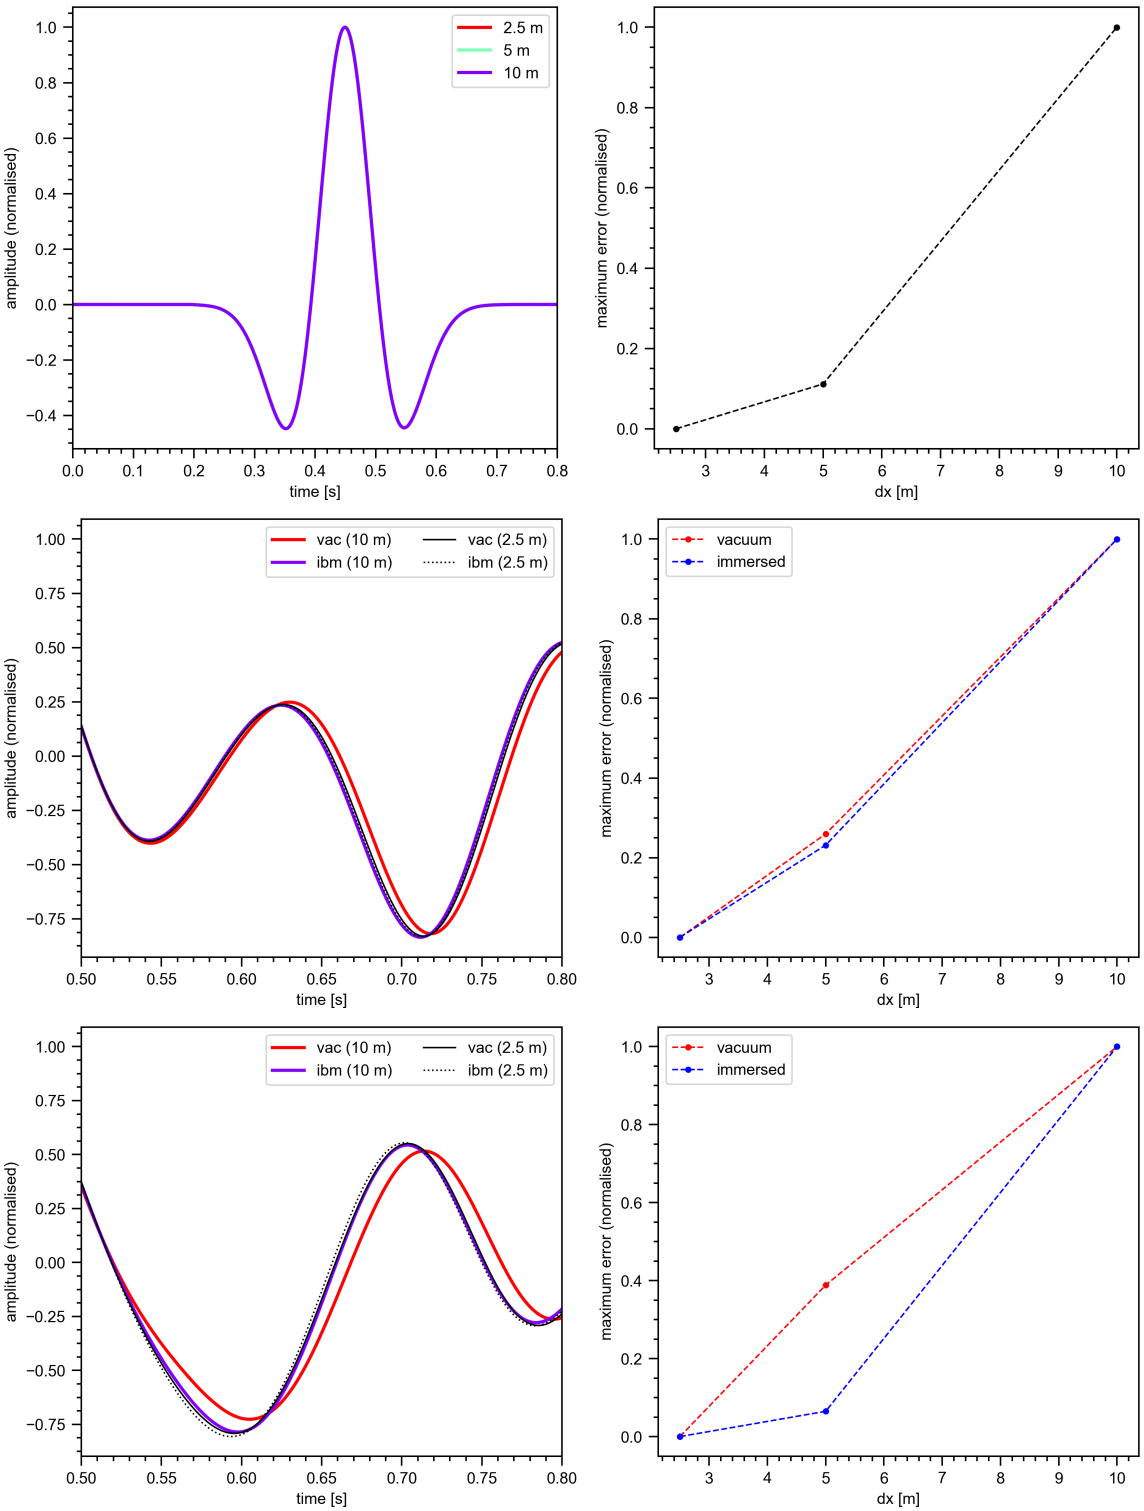


Figure 6.31: Convergence of the results of the experiments from fig. 6.30 to the reference 'true' solutions. Left: traces observed at the receiver. Right: maximum error normalised to the coarsest-grid case. The 'ibm' and 'vac' stand for the immersed-boundary and vacuum method, respectively, with the grid interval indicated in the parentheses.

pronounced for the rougher topography which shows that using the IB method is especially beneficial in realistic 3d simulations where grids to need to be as coarse as possible for computational efficiency.

Part III

Inversion of field data

7

Input preparation

Full-waveform inversion (FWI) requires three basic pieces of input: (1) high-quality observed[1] data, (2) an accurate estimate of the seismic-source wavelet, and (3) a good starting model[2]. All three have to be correctly located relative to each other in time and space. Also, the observed data and the source wavelet have to be processed in a consistent manner. The basic requirement for a successful transmission[3] FWI is that the majority of the synthetic data should match the early seismic arrivals observed in the field data to better than half of the longest inverted period. The data match depends on the velocity model, source wavelet, station positions and the numerical solver. Here, the solver is assumed to be accurate for the inverted data. Preparation of the input is discussed below.

7.1 Metadata

Metadata will here refer to the information about data acquisition geometry and timing. All of the input need to be located in space and time in a self-consistent manner. Data and source traces need to start at the same time, typically chosen as the time of firing a given shot. Sources and receivers have to be accurately positioned relative to each other and the model. Both the time origin and station locations are kept fixed during the inversion and their contribution to the misfit can only be reduced by modifying the velocity structure.

7.1.1 Station positions

The accuracy to which these positions are known is compromised mainly by GPS errors in case of the airgun shots and sea currents/tides in the case of OBSs (Creager and Dorman, 1982). The average position of the ship is typically determined within ±10 m accuracy (Nomikou et al., 2014b). The positions

[1] In seismology, 'observed' is synonymous with 'field'.

[2] There are other FWI variants that can do without the last two ingredients (with a varied degree of success), discussion of which is beyond the scope of this chapter.

[3] The transmission FWI focuses on early-arriving, forward-scattered (refracted) waves (Mora, 1988).

of the land stations are constrained only by the uncertainty of the GPS which is assumed to be negligible given the accuracy of modern GPS devices.

Relocation method Typical OBS position uncertainty was estimated by Heath et al. (2019) to be ± 10 m vertically and ± 4 m horizontally and will here be assumed to be 3σ bounds on these errors. As the source array was only 15-metre long (fig. 3.2), its deviation from the intended towing position behind the vessel can be neglected. Given the large number of shots, the problem is highly over-determined and OBS positions can be very well constrained. Station locations were determined by Heath et al. (2019) using a least-squares inversion, and placed on the seafloor using the bathymetric map from Hooft et al. (2017). Here the procedure is assumed to be accurate with the only source of uncertainty being the assumption that the water velocity is constant. The potential travel-time error related to this uncertainty is discussed below.

Water-velocity error The RMS of water-wave time-residuals was $\delta t \approx 5$ ms for *most* stations and was taken as a picking error for these phases (Heath et al., 2019). This value corresponds to the position uncertainty of $\delta x \approx 7.6$ m assuming $\delta x = v_{\text{water}} \delta t$ and $v_{\text{water}} = 1520$ m/s.

Sensitivity of travel-time residuals to OBS positioning average water velocity errors can be quantified using analysis of error-propagation (Taylor, 1997). Taking coordinates so that the source is at $(0, 0, 0)$ and Y axis is along source-receiver plane, we can express the total distance between the source and receiver as:

$$L = \sqrt{x^2 + z^2} \tag{7.1}$$

Neglecting the ray path curvature, the travel-time is given by:

$$t = \frac{L}{v_{\text{water}}} = \frac{\sqrt{x^2 + z^2}}{v_{\text{water}}} \tag{7.2}$$

From the theory of uncertainty-propagation (Taylor, 1997), the travel-time residual can be expressed as:

$$(\delta t)^2 = \left(\frac{\partial t}{\partial v}\delta v\right)^2 + \left(\frac{\partial t}{\partial x}\delta x\right)^2 + \left(\frac{\partial t}{\partial z}\delta z\right)^2, \tag{7.3}$$

with:

$$\frac{\partial t}{\partial x_i} = \frac{1}{v_{\text{water}}} \cdot \frac{x_i}{\sqrt{x^2 + z^2}}, \tag{7.4}$$

where x_i is any of the x, z and:

$$\frac{\partial t}{\partial v} = -\frac{\sqrt{x^2 + z^2}}{v_{\text{water}}^2} \tag{7.5}$$

For position uncertainties from Heath et al. 2019 and 10 m/s of uncertainty in the average water velocity, one gets about 0.01 s of the travel-time error at $z \approx 400$ m and $x \approx 2000$ m (water depth and maximum offset used in the relocation, respectively).

7.2 Observed data

Processing of the data was aimed at enhancing signal-to-noise ratio (SNR) and excluding data other than early arrivals. This was achieved by simple filtering and time-windowing discussed below. Other

Algorithm 7.1 Processing of the field data.

1. Identify the lowest usable frequency determined by incoherent noise.
2. Define the maximum frequency to be simulated on a given grid.
3. Band-pass filter.
4. Time-window around early arrivals.

Algorithm 7.2 Filtering workflow that minimises wrap-around effects. Applied to each trace.

1. Taper both ends.
2. Pad both ends with zeroes.
3. Filter (minimum-phase band-pass Butterworth)
4. Trim the paddings.

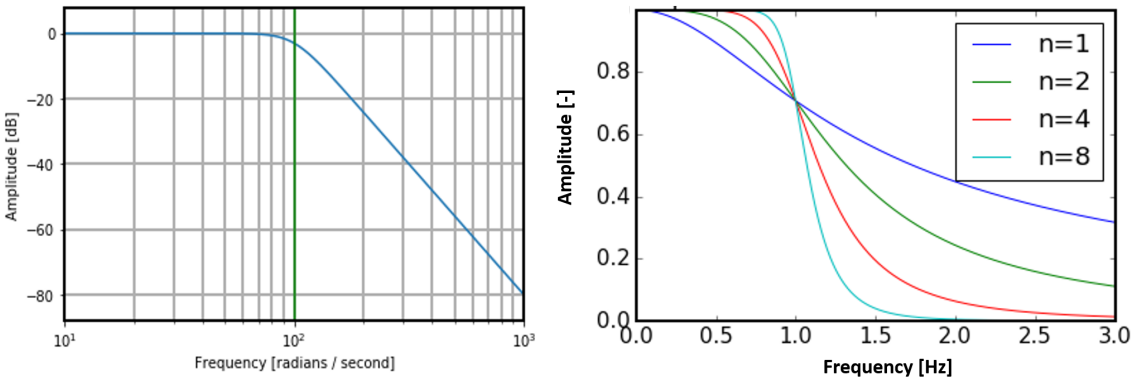


Figure 7.1: Low-pass Butterworth filter: (left) power spectrum of a filter with 4 poles and cut-off frequency of $\omega_c = 100$ rad/s denoted by the green vertical line. Note the semi-log scale and the flatness of the passband; (right) amplitude spectrum of Butterworth filters of 1 Hz cut-off frequency and different order n.

processing steps were not pursued partly because they can damage low-frequency part of the spectrum, and partly because some amplitude balancing is done by the `FULLWAVE3D` code (Warner et al., 2013b). The pre-processing performed by other members of the PROTEUS team is described in sect. 3.2.1.

To ensure that the synthetic data matches the input field data, the source wavelet was extracted from the processed field data.

7.2.1 Filtering

Two kinds of filters were applied to the field data. A high-pass filter to remove the swell noise dominant in the 1-3 Hz frequency range and a low-pass filter removing the frequencies above 6.5 Hz. The latter was aimed to avoid numerical dispersion of the wavefield in the water layer sampled at 50 m spacing (fig. 8.9) when using the low-order `FULLWAVE3D` solver (sect. 8.3.1). The low-pass filtering could be instead applied during the inversion, which would result in a greater flexibility of the workflow as exactly the same filtering should be applied to the source wavelet (sect. 7.3).

The filtering was applied as a minimum-phase band-pass filter of corner frequencies 2.0-3.0-4.5-6.5 Hz using `subfilt` program, a part of the `Seismic Unix` (SU) package (Stockwell and Cohen, 2008). A distinguishing feature of a Butterworth filter is a maximally flat frequency response of the passband

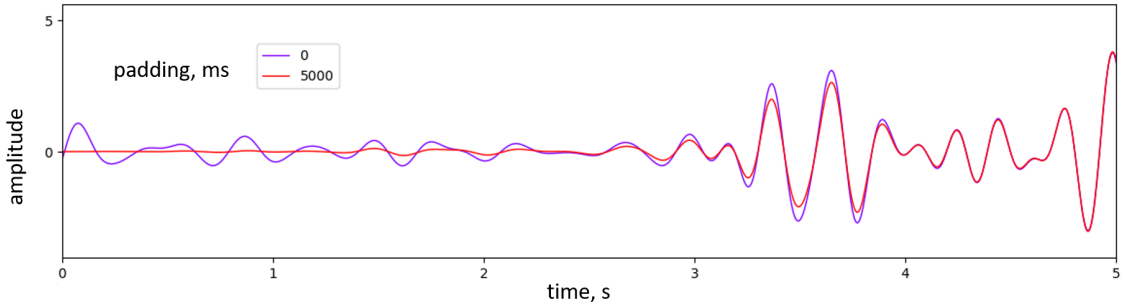


Figure 7.2: Effect of trace padding on filtering the field data (OBS 108, hydrophone channel, line 29, trace no. 250); purple – no padding, red – 5000 ms applied to both trace ends.

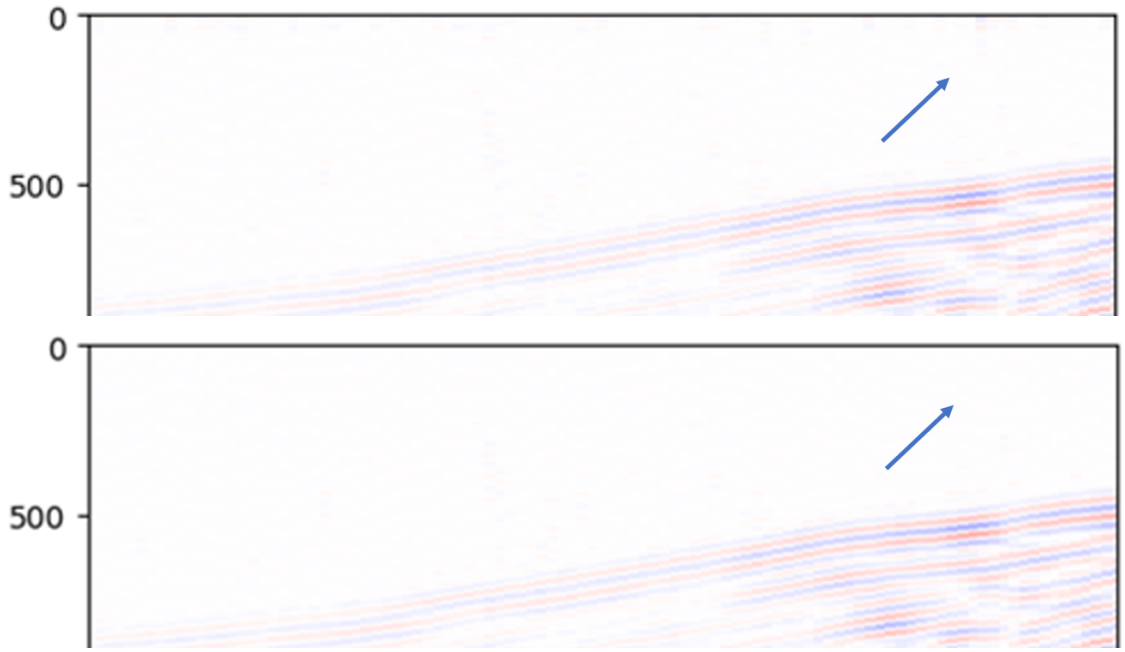


Figure 7.3: Data filtered (top) with and (bottom) without taper applied prior to the filtering. The difference indicated by the arrow is quite subtle. The Y axis in time-samples (0.0025 ms each).

(Gubbins, 2004). A power spectrum of a low-pass Butterworth filter is given by:

$$|A(\omega)|^2 = \frac{1}{1 + (\omega/\omega_c)^{2n}} \tag{7.6}$$

where n is the filter order (no. of poles) and ω_c is a *cut-off frequency* defined as a point where the power drops by half (and amplitude by a factor $\frac{1}{\sqrt{2}}$) relative to the maximum value in the passband (fig. 7.1). This corresponds to $10 \log_{10} \frac{1}{2} \approx -3$ dB hence the cut-off frequency is also known as the *half-power* or *-3 dB point*. The order of the filter controls the sharpness of the roll-off (fig. 7.1). A high-order Butterworth filter a resembles a trapesoidal filter such as the Ormsby filter (Sheriff and Geldart, 1995).

To prevent the wrap-around effects of frequency-domain filtering (Press et al., 1990), zero-padding was first applied to both ends of each trace (fig. 7.2). In SU this can be accomplished in two steps: padding with the trailing zeroes (twice as wide as intended) and shifting in time so that the resulting padding is the same at both ends. Before the padding, traces were tapered at both ends with a 500-sample

Algorithm 7.3 Muting workflow.

1. Calculate synthetic data.
2. Extract times of the first-arrivals from the synthetic data.
3. Define a top mute as the time of the first arrival and a bottom mute as the top mute plus a time-window
4. Apply both the top and bottom tapered mutes.

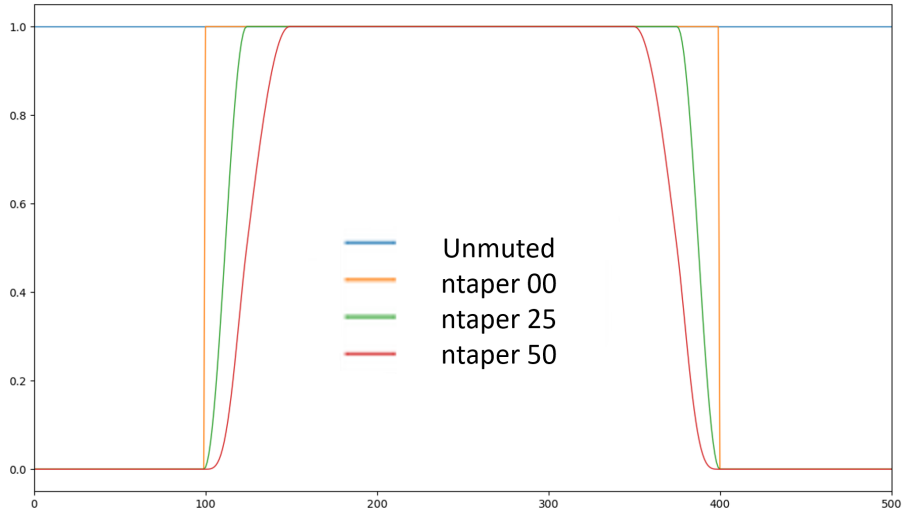


Figure 7.4: Top and bottom mutes with different widths of the sine-squared taper starting at the 100th and 400th sample, respectively.

squared-sine taper (fig. 7.3) to reduce spectral leakage (Gubbins, 2004). The traces were trimmed to the original length after the filtering. The full workflow is presented in the alg. 7.2 and the final effect compared to raw data can be seen in fig. 7.5.

7.2.2 Muting

After filtering, two kinds of mutes, top and bottom, were applied (fig. 7.4). The top mute was intended to remove the noise ahead of the first breaks. The bottom mute was meant to exclude the data other than early-arriving phases, mostly refractions and their ghosts. Such time-windowing is a key step for effective transmission FWI. The FWI code does not try to match any parts of the input data that have a zero amplitude..

The first breaks were obtained by calculating synthetic data through the starting model. They were defined as a first sample with at least 0.01 times the maximum amplitude of the trace. Note that this needs to be recalculated every time the data is processed in a different way. Current version of `FULLWAVE3D` has a capability to extract the first breaks but not yet to mute the data in a way intended here. If the FWI code had a capability to mute the data after the low-pass filtering at a given iteration, this would be more accurate and convenient. Ideally, the muting should be the last part of the pre-processing because any subsequent filtering introduces spurious high-frequency oscillations due to the finite extent and sharp edges of the time-windowed data. This effect is smaller for gently sloping edges which can be achieved by tapering the mutes.

Here a sine-squared taper was used as implemented in the `SU` package (program `sutaper`). Fig. 7.4 shows that in this implementation the tapering starts/stops at a given time sample (instead of being

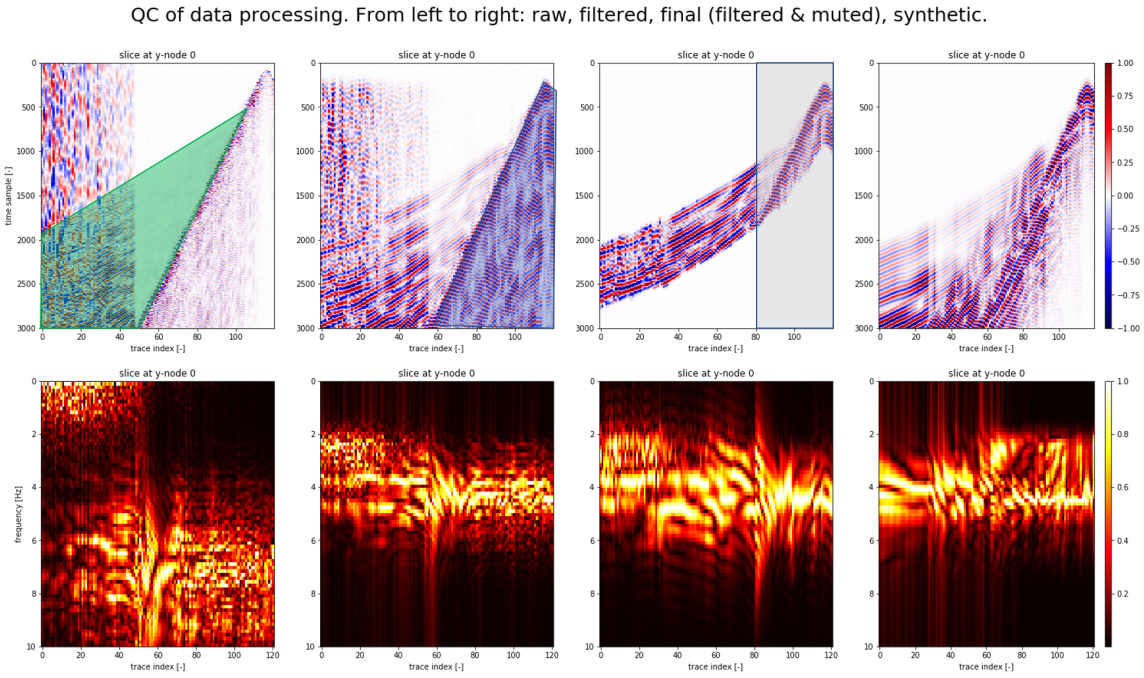


Figure 7.5: Different stages of processing the field data compared to the synthetic ones shown with a receiver gather (top panel with 0-7.5 s time on the Y axis) and its amplitude spectrum (bottom panel with 0-10 Hz frequency on the Y axis). Note the effect of muting on the amplitude spectrum. Arrivals before/after the direct wave highlighted in green/blue, respectively. Offsets shorter than 5 km excluded from the inversion highlighted in grey.

centered at it) for a top/bottom mute respectively, always tapering the useful signal inside the time window. In this way, the effective width of the time-window Δt_{win} is equal to:

$$\Delta t_{\text{win}} = (t_{\text{bot}} - t_{\text{top}}) - (\Delta t_{\text{top}} + \Delta t_{\text{bot}}) \tag{7.7}$$

where t_i are the window edges, Δt_i are the taper widths and i refers to the top or bottom mute respectively. Here t_{top} was defined as the time of the first break, $\Delta t_{\text{top}} = \Delta t_{\text{bot}} = 0.25$ s and t_{bot} was defined so that various time windows Δt_{win} were obtained as discussed in sect. 8.2.3. Figure 7.5 shows an example of a raw receiver gather and its amplitude spectrum, along with the same gather after filtering and muting, as well as the synthetic gather.

7.3 Source wavelet

A seismic source function is a time series that when convolved with the Earth's response and the receiver's response will produce a seismogram recorded in the field. This corresponds to a direct wave (without a ghost etc.) propagated through a perfectly elastic, homogeneous medium, shifted in time to start at zero and with amplitude corrected for geometric spreading.

In travel-time tomography the wavelet does not need to be known, only the onset time matters. In migration, It is usually debubbled, and filtered to provide a broadband zero-phase wavelet (Warner et al., 2013b). Its low-frequency content does not need to be known precisely. This may compromise the accuracy of the final result, but in a 'linear' way in a sense that small wavelet errors will translate into image errors but will not lead to a complete failure (Warner et al., 2013b).

The wavelet can be derived by different methods (Nangoo, 2013):

Algorithm 7.4 Source-wavelet derivation. See also the flowchart 7.6.

1. Select a number of relatively clean WHOI traces with linear response, within 100-300 m offset range.
2. Align all to one arbitrary trace, via cross-correlation.
3. Stack them.
4. Calculate the arrival of the first multiple using 1500 m/s velocity and coordinates of the trace used for alignment.
5. Bottom mute the stacked trace with 50-sample sine-squared taper at the time just before the arrival time from previous step.
6. Prepare an initial-guess wavelet as a unit spike at time zero, filtered with minimum-phase Butterworth filter of corner frequencies 2-3-4.5-6.5 Hz.
7. Generate synthetic data in water velocity 1500 m/s model with a free surface.
8. Design a Wiener filter transforming synthetic into observed data in an optimal way in the least-squares sense.
9. Obtain the new wavelet by convolving the Wiener filter with the initial guess.
10. Generate synthetics using the new wavelet.
11. Compare with the observed trace. Perfect match indicates good performance of the Wiener filter.

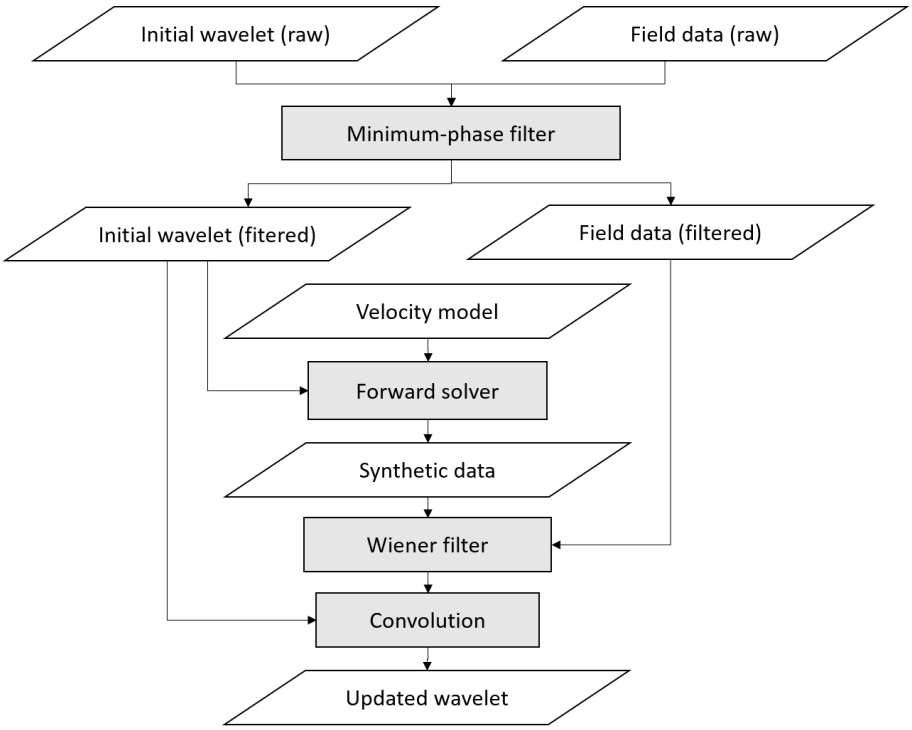


Figure 7.6: Flowchart of deriving the source wavelet from the field data using the Wiener filter.

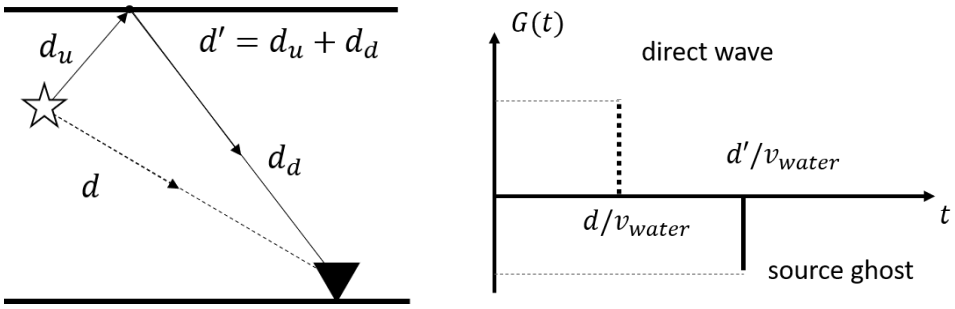


Figure 7.7: Seismic phases and corresponding impulse response (Green's function).

- joint inversion for wavelet and velocity model (Pratt, 1999; Virieux and Operto, 2009) – this may potentially compromise the final model quality;
- calculating airgun-array far-field signature – this is the theoretical Langseth source presented in fig. 3.3. It works well for the bandwidth commonly used in migration which is higher-frequency than FWI, and should not be trusted when low frequencies are inverted (Warner et al., 2013b);
- extracting of source from the field data.

The last approach was adopted here. The free surface was intended for FWI, so the wavelet had to be deghosted. In a real experiment recording of the direct waves is contaminated by other arrivals. One can keep their presence to a minimum by looking only at the shortest offsets and using narrow time windows. This approach eliminates refractions and most reflections, especially for deep water settings (surface multiples arrive outside the window). A ghost (reflection from the sea surface), however, remains in the data since the airgun is fired just below the sea surface.

The deghosting was achieved using the Wiener filter. The workflow is summarised in alg. 7.4 and visualised in fig. 7.6. There are some practicalities to take into consideration that will now be discussed in detail.

7.3.1 Theory

Let us assume that observed data d_{obs} at any of the receivers can be expressed, up to the data noise e, as a convolution of the source function s, the Green's function G and the receiver response r:

$$d_{\text{obs}}(t) = s(t) * G(t) * r(t) + e(t). \tag{7.8}$$

All of above are the functions of time only as the position is fixed at the receiver. From now on, the dependence will be implicit for the sake of brevity.

At first, we will not deal with the receiver response. One can incorporate it into s or just assume it to be a perfect Dirac delta $\delta(t)$ for the frequency range of interest. We will, at least initially, neglect e too, so that eq. (7.8) simplifies to:

$$d_{\text{obs}} = s * G. \tag{7.9}$$

This gives a formula for the unknown wavelet:

$$s = G^{-1} * d_{\text{obs}}, \tag{7.10}$$

where G^{-1} is a convolutional inverse G defined by the identity:

$$G^{-1} * G = \delta. \tag{7.11}$$

We do not know G (or G^{-1}) directly. However, it can be designed analytically, using the known source depth and assuming all phases arriving later then the source ghost have been muted in the field data, so that the seabed reflectivity does not need to be known. Such an impulse response will consist of two spikes of opposite polarites separated by the time difference between the direct and source ghost arrival (fig. 7.7). The inverse can be then calculated using deconvolution (e.g. by designing a Wiener filter).

Another method, chosen here, is very similar and involves synthetic data simulation. We can generate synthetic data using an initial-guess source s_0:

$$d_{\text{syn}} = s_0 * G_0, \tag{7.12}$$

where G_0 is the approximation to the true Green's function G. It captures the best-guess model of the subsurface as well as the discrepancy between the true physics and the simplified wave equation used by the numerical solver. Its inverse:

$$G_0^{-1} = s_0 * d_{\text{syn}}^{-1} \tag{7.13}$$

is an approximation $G^{-1} \approx G_0^{-1}$ that can be plugged into (7.10) to give:

$$s = s_0 * d_{\text{syn}}^{-1} * d_{\text{obs}}. \tag{7.14}$$

It now becomes clear that the true wavelet can be obtained by convolving the initial wavelet with a filter w defined as a deconvolution of d_{syn} from d_{obs}:

$$w \coloneqq (d_{\text{syn}})^{-1} * d_{\text{obs}}. \tag{7.15}$$

Equivalently, w is a filter that transforms d_{syn} into d_{obs}:

$$d_{\text{syn}} * w = d_{\text{obs}}. \tag{7.16}$$

In the absence of noise, w could be found directly from eq. (7.15). In reality, this may not be possible e.g. if d_{syn} is not minimum-delay or has notches in the spectrum (Gubbins, 2004). These effects can all be expressed as some error e:

$$\underbrace{d_{\text{syn}}}_{\text{input}} * \underbrace{w}_{\text{filter}} = \underbrace{\tilde{d}_{\text{obs}}}_{\text{filter output}} = \underbrace{d_{\text{obs}}}_{\text{desired output}} + \underbrace{e}_{\text{error}} \tag{7.17}$$

and the problem can be turned into finding a filter w that minimises e in a least-squares sense:

$$Aw = d_{\text{obs}} + e \tag{7.18}$$

where Aw is a matrix form of $d_{\text{syn}} * w$ and:

$$A = \begin{bmatrix} a_0 & & & \\ \vdots & a_0 & & \\ a_{N-1} & \vdots & \ddots & \\ & a_{N-1} & & a_0 \\ & & \ddots & \vdots \\ & & & a_{N-1} \end{bmatrix} \tag{7.19}$$

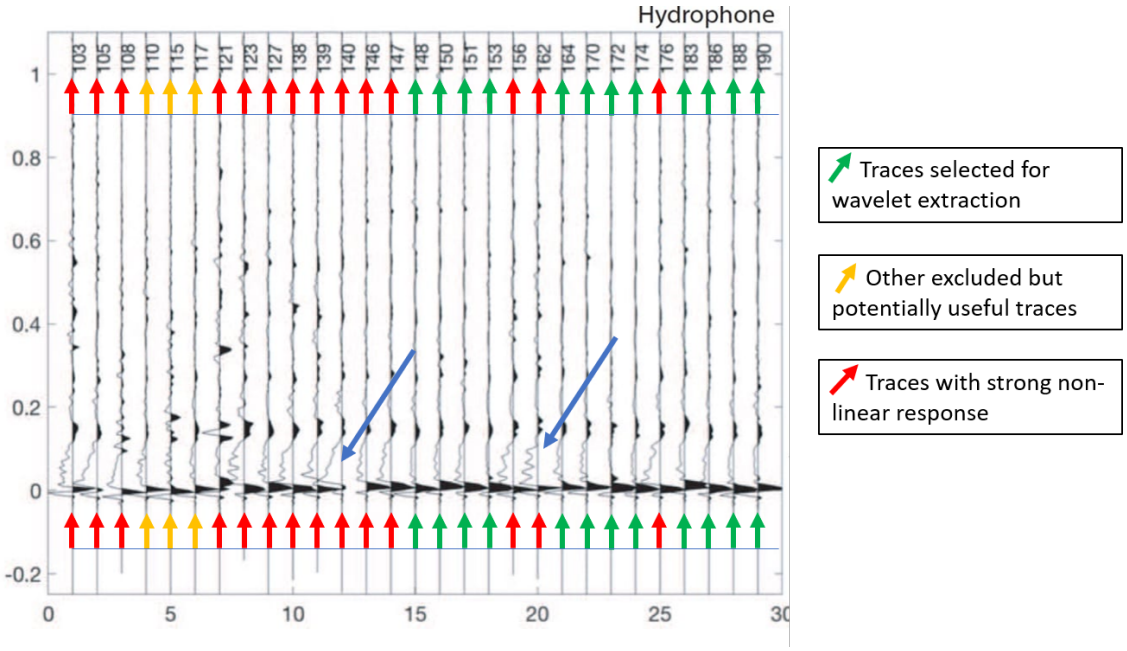


Figure 7.8: Nearest-shot traces for all WHOI stations, unfiltered and aligned on the predicted water-wave arrival time;. The hydrophone component displays a variable shape of the negative pulse arriving right after the strongest positive pulse. Station IDs are shown next to the traces. Y axis in seconds. Modified after Heath (2019).

in which $a_0, \ldots, a_{N-1}$ are elements of d_{syn}. Applying a generalised inverse (Gubbins, 2004) one gets:

$$w = (A^T A)^{-1} A^T d_{\text{obs}}, \tag{7.20}$$

where $A^T A$ is a autocorrelation of d_{syn} and $A^T d_{\text{obs}}$ is a matrix form of the cross-correlation $d_{\text{syn}} \star d_{\text{obs}}$. The autocorrelation is a Toeplitz matrix hence its inverse can be quickly calculated by the Levinson algorithm (Press et al., 1990). In practice, eq. (7.20) should be damped for stability:

$$w = (A^T A + \lambda I)^{-1} A^T d \tag{7.21}$$

which can be shown to be equivalent to the water-level method used in the frequency-domain deconvolution (Gubbins, 2004). The quality of deconvolution depends on the length of the filter w – the longer the better, sometimes only an infinitely-long filter can reproduce the desired output precisely. However, for a given length, w is always optimal in the least-squares sense.

7.3.2 Workflow

A basic requirement for the observed data is that it should contain a clean direct wave with no other phases except for the source ghost. This constrains the candidate traces to the shortest offsets for which the direct wave is the first arrival. At larger offsets, waves refracted in the solid subsurface start to arrive at a similar time as the direct wave, and eventually become the first arrival. On the other hand, intervals between arrival times of multiples:

$$t = \frac{2z}{v_{\text{water}}} \tag{7.22}$$

increase with the offset, so they are easier to separate for longer offsets.

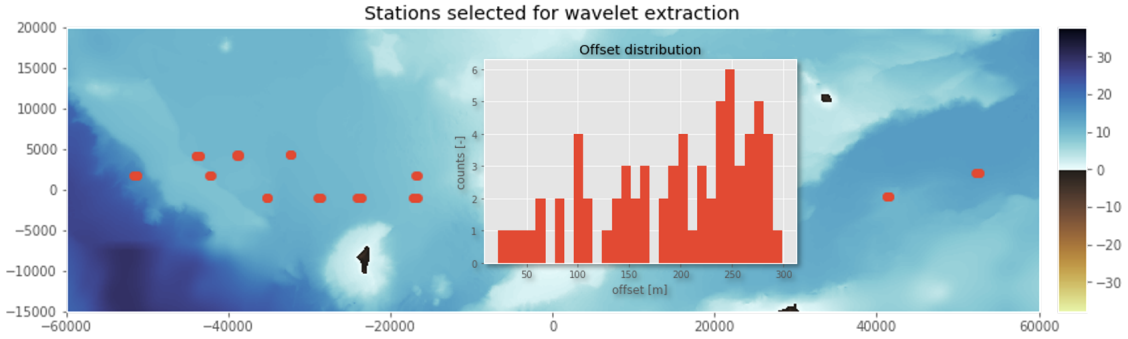


Figure 7.9: Map of the stations from which the traces were selected for stacking and wavelet extraction; axes scaled in local coordinates (in metres). For some stations more than one trace was used. Inset – offset distribution of the selected traces.

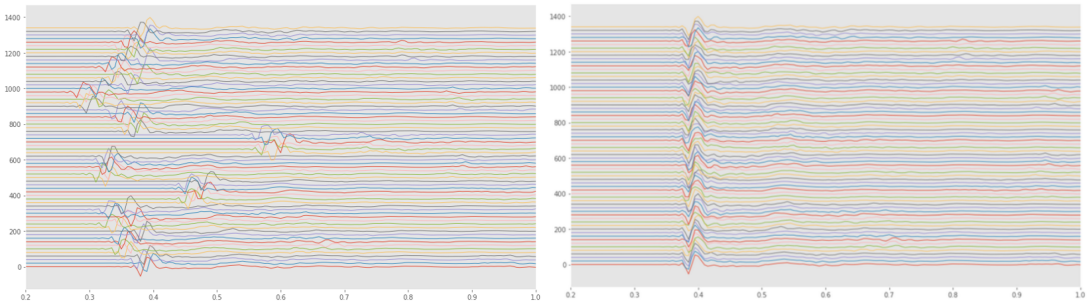


Figure 7.10: Selected traces (left) before and (right) after alignement using the cross-correlation lag. Different colors are used only to help discriminate between neigbhouring traces. X axis in seconds.

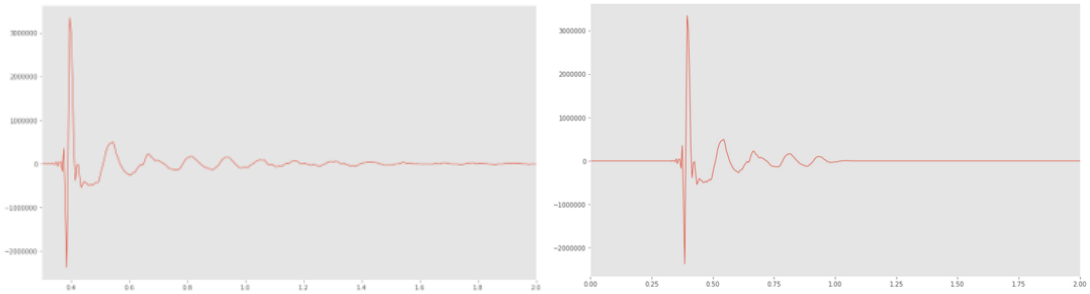


Figure 7.11: Stacked trace (using all the traces from 7.10) raw (left) and (right) bottom-muted with a 50-samples sine-squared taper at a time just before the arrival of the first multiple; the arrival time was calculated using 1500 m/s water column) and the coordinates of the receiver same as for the trace used for alignment. X axis in seconds.

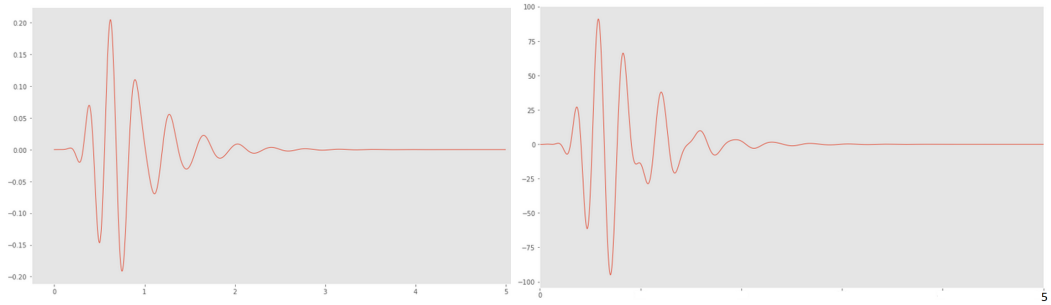
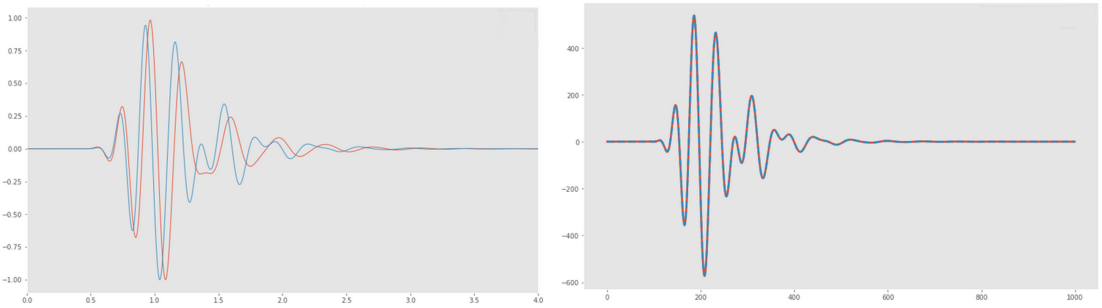


Figure 7.12: Initial-guess (left) and the final (right) wavelet. The former was defined as a Kronecker delta at zero time filtered with a minimum-phase filter. X axis in seconds.



(a) (b)

Figure 7.13: Observed (blue) vs. synthetic (red) data generated using (left) the initial-guess and (right) final wavelet. The traces on the right overlap almost entirely. X axis in seconds.

Following conclusions of sect. 3.2.2, many SIO stations had problems with amplitude clipping. That is why the selection was limited to the WHOI stations. These stations had their own problems though. Some of them suffered from a non-linear response (fig. 7.8). Eventually WHOI stations with IDs 148, 150, 151, 153, 164, 170, 172, 174, 183, 186, 188, 190 (fig. 7.9) as having similar waveforms of the direct wave. At least one trace within 100-300 m was picked for each of these stations for further analysis. The filter could be designed to match all the traces simultanously but instead a stacked traces was obtained (fig. 7.10 and 7.11).

As an initial-guess wavelet a minimum-phase filtered delta function was used (fig. 7.12). It is important to keep track of all processing steps applied during extraction of the source wavelet as we will need to apply them to the data prior to the inversion. Here the same filtering was applied to the initial-guess wavelet as to the field data (sect. 7.2). The corner frequencies were the same as for the field-data filtering. The trace was also muted to remove the surface multiples (fig. 7.11).

To generate the synthetic data (fig. 7.13) a homogeneous 3d model with 1.5 km/s velocity was used. It was about $10 \times 10 \times 10$ km in size so that both source and receiver were about 5 km from the nearest edge. Synthetic data was padded and shifted in time to capture the full length of the Wiener filter. The implementation of the Wiener filter (7.14) used was a part of the `SU` package. Numerous tests have been performed to get a feeling of sensitivity of the filter to the input.

The final wavelet (fig. 7.12b) has a sampling rate 0.005 s, same as the raw data, and 2000 samples which gives 10 s of total time. It perfectly matches the stacked observed trace (fig. 7.13b) and other first arrivals at slightly longer offsets (fig. 7.15).

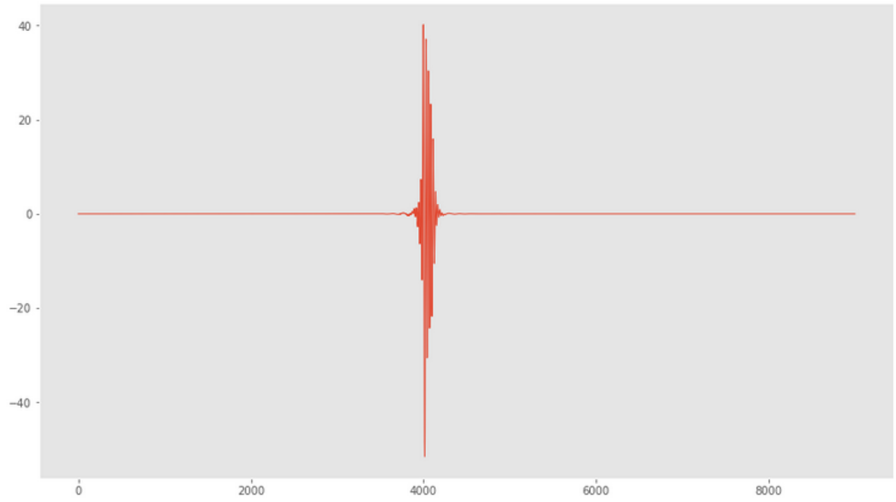


Figure 7.14: Wiener filter used to convert the synthetic into observed data in fig. 7.6b. X axis in time-samples ($dt = 5$ ms).

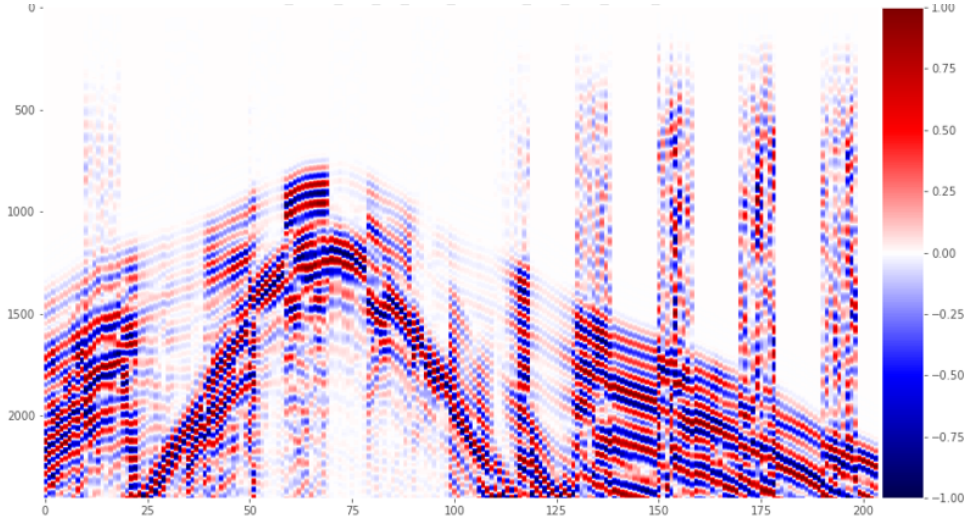


Figure 7.15: Observed vs. synthetic data calculated for OBS 104, line 37 using the final wavelet from fig. 7.12. Y axis in time-samples ($dt = 2.5$ ms).

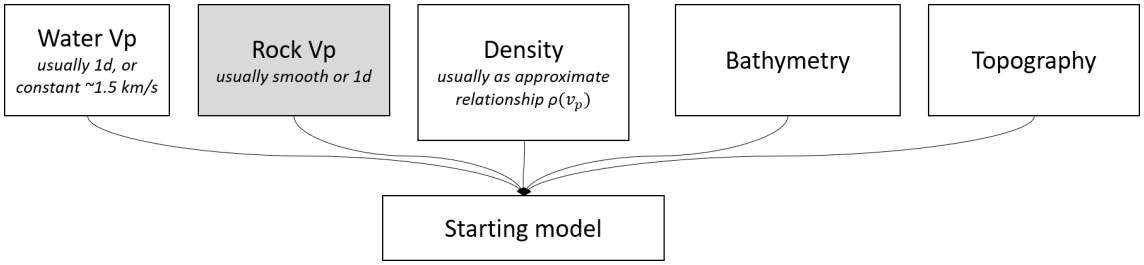


Figure 7.16: Ingredients needed to build an accurate starting model for FWI.

7.4 Starting model

The ingredients needed to build a starting model suitable for acoustic FWI are shown in fig. 7.16. Below, the process of building such a model is described in detail.

7.4.1 Parametrisation

Acoustic isotropic medium can be fully characterised by two parameters – P-wave velocity v_P and density ρ. Their product called the (acoustic) impedance:

$$I = \rho v_P \tag{7.23}$$

controls the amplitude of a reflection. Such amplitude is proportional to the *reflectivity*, i.e. the impedance *contrast* (Sheriff and Geldart, 1995). Note that doubling the velocity and halving the density in some part of medium will produce no reflections at the interface although it will affect the travel-times. On the other hand, a density contrast in a constant velocity model will generate a reflection without refracting or changing a travel-time of the wave. The model parameters may vary smoothly (on a scale of the seismic wavelength) in some parts of the model, and change abruptly at interfaces such as the seabed or land surface.

Model interfaces Assuming that the station positions and source wavelet are accurate, travel-time residuals are caused in general by inaccurate background velocity model and mispositioned impedance contrasts (reflectivity model). Transmission FWI can only correct the velocity model in a tomographic fashion, with resolution constrained by a local wavelength. Therefore, features finer than the local wavelength such as sharp impedance contrasts should not be present in the starting model unless their position and reflectivity are known with certainty and they can be accurately represented on a model grid. Usually the only interfaces meeting the above criteria are seabed and land topography. Here the focus is on the former, dealing with the latter is discussed in ch. 6.

Seabed reflectivity In the study area, seabed reflectivity varies laterally from hard metamorphic rocks to soft basin sediments (Heath et al., 2019). In sedimentary basins, the reflectivity is controlled mainly by the density difference because typical velocity of shallow seafloor sediments (about 2 km/s) is not much higher than the water velocity (Nangoo, 2013). On metamorphic ridges in turn, the reflectivity is expected to be higher, with contribution from both sharp velocity and density contrasts. In either case, the seabed is expected to be the strongest reflector in any vertical profile and mispositioning it can result in considerable phase errors and potential cycle-skipping. The mislocation especially affects seabed multiples but they were excluded from the inversion. As the data cannot be easily 'unskipped', it is better to first smooth the interface and let the inversion refine it gradually.

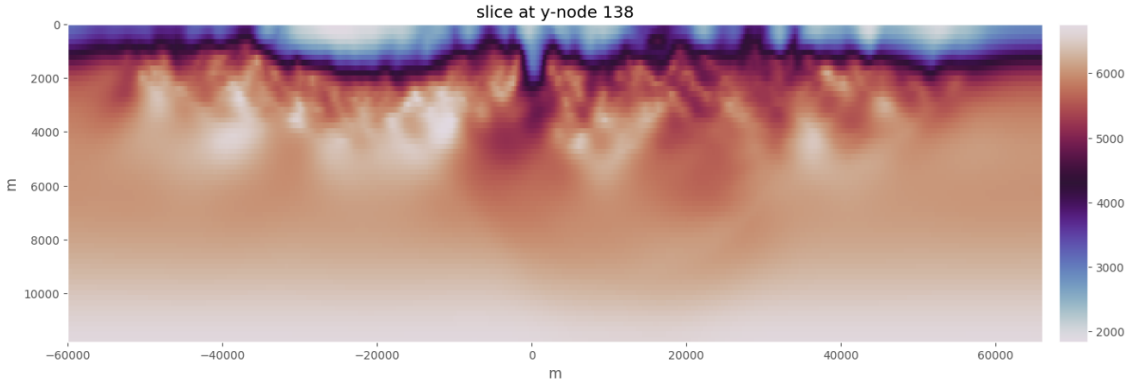


Figure 7.17: Travel-time tomography model of Heath et al. (2019) sampled on a vertically-sheared 200x200x200 m grid. Minimum and maximum value of v_P is 1.69 and 9.22 km/s, respectively.

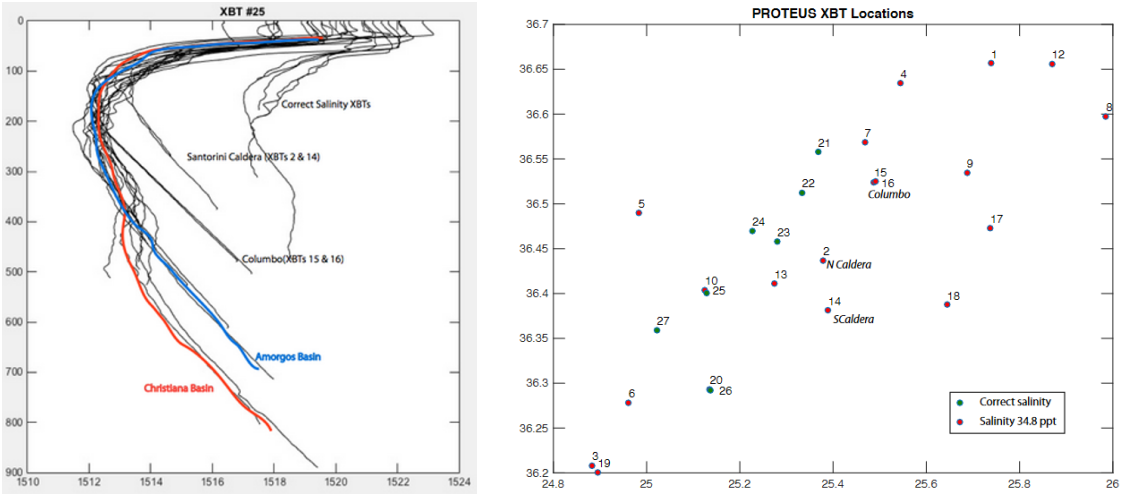


Figure 7.18: Water column velocity profiles (left, with v_P [m/s] and depth [m] on the horizontal and vertical axis respectively) determined from the bathythermographic (XBT) measurements at various locations in the study area (right, axes in geographical coordinates). Modified after Hooft (2015a).

7.4.2 P-wave velocity

The model of velocity was constructed from the isotropic slowness models obtained with travel-time tomography by Heath et al. (2019) and McVey et al. (2020). The difference between the two is that the latter used $\sim$ 30k more travel-time picks especially at long offsets across the caldera (McVey et al., 2020). That allowed to constrain the structure beneath Santorini down to 6 km b.s.l. as opposed to 3-4 km b.s.l. in the former case (Heath et al., 2019). Which one of the tomographic models was used will be clarified when discussing a specific case study. Here, the model of Heath et al. (2019) is shown (fig. 7.17). As the forward modelling was performed on a vertically-sheared, seafloor-conforming grid (Toomey et al., 1994), it had to be re-sheared. There are a few caveats related to this procedure.

First, to shear the grid Heath et al. (2019) used the bathymetry map from (Hooft et al., 2017). In the travel-time inversion 1.52 km/s constant water velocity was used. The bathymetry measured by multi-beam swath is a two-way travel-time section that needs to be converted to depth using the same water velocity model as in tomography. If this model was different than the one used in tomography, e.g. the one from fig. 7.18, the seabed would be vertically stretched relative to its true shape. Ideally, bathymetry should be reconverted to depth, or a different water velocity should be embedded in the

Algorithm 7.5 Embedding finely-sampled seabed on a coarse model grid.

1. Upsample the model vertically.
2. Put the seabed at a nearest grid node.
3. Smooth the seabed interface over an area approximately equivalent to the final grid spacing *doubled*.
4. Downsample the model to the desired rate.

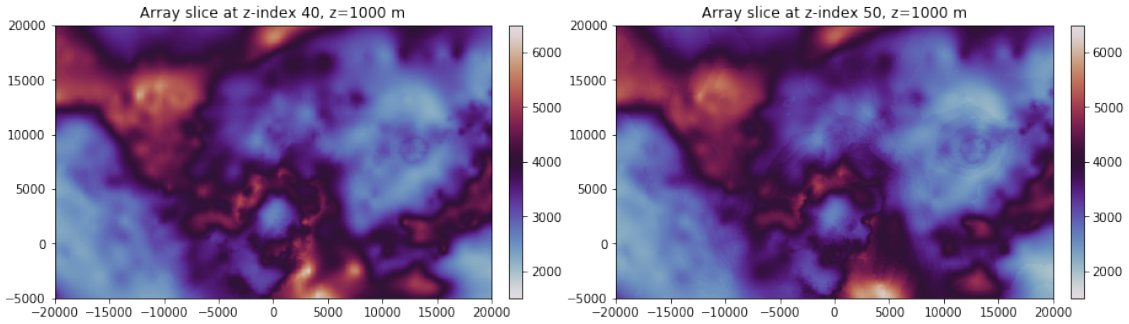


Figure 7.19: Starting model of v_P (m/s) prepared from fig. 7.17 following alg. 7.5 (left) is characterised by a smaller topography imprint compared to the original version from Heath (2019). Axes scaled in local coordinates, in metres.

starting model (fig. 7.18), to make the two consistent. The related error, however, is expected to be neglible given the shallow water depth in the study area.

Second, the seabed used by Heath et al. (2019) was sampled on a 200-m grid laterally[4]. In FWI a 50-m grid was used. For this reason, another bathymetric map was created by M. Paulatto with a 50-m lateral sampling (Paulatto, 2017, accessed 25 Apr 2021). Even though the two used the same PROTEUS bathymetric data (Hooft et al., 2017) in the area of interest, their shapes were slightly different, most likely due to different coordinate-transformation procedure used which needs further investigation (M. Paulatto, personal communication). For the time being, the 50-m was used for the re-shearing.

Third, even if the shape of the seabed is known accurately and it is correctly sampled laterally, the elevation cannot be represented on a model grid which coarse in the vertical dimension. This gives a first-order discretisation error (Symes et al., 2008) which in the case of 50-m grid reaches up to 25 m. For this reason the seabed was resampled as in alg. 7.5 resulting in a minor topography imprint (fig. 7.19).

The final model (fig. 7.20) derived from the one by Heath et al. (2019) extended from -60 to 64 km inline, from -14 to 29 km crossline and from -1.5 to 5 km in depth below the sea level (b.s.l.). Sampled at 50 m in all directions, this gives a $2481 \times 861 \times 131$ cubic grid. A water layer added on top had an approximately constant value of 1.52 km/s.

7.4.3 Density

Although a density model does not influence the travel-times, it can affect a shape of early-arriving waveforms and producing an apparent phase shift by the interference of different phases in the shallow subsurface (fig. 4.19 from Nangoo, 2013).

[4] The original bathymetry had a a finer spacing but it was downsampled.

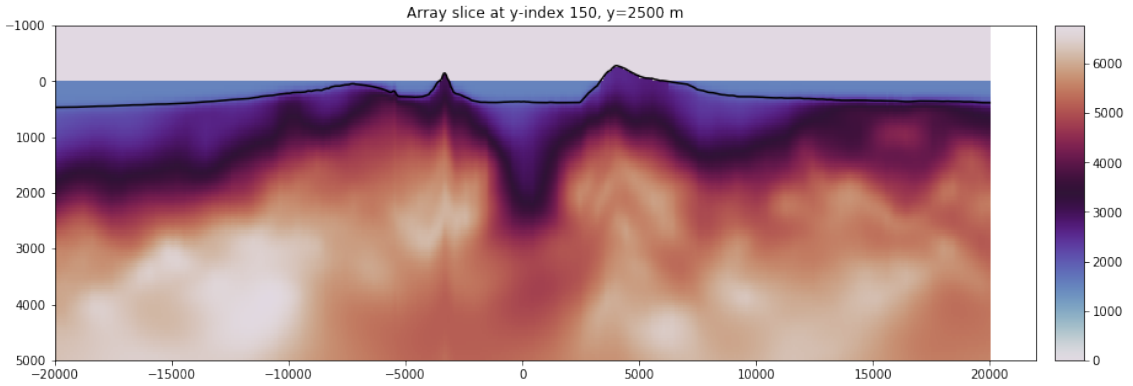


Figure 7.20: Starting model of v_P (m/s) overlain with the topography of the seabed and land surface (black line). Axes scaled in local coordinates, in metres.

In the FULLWAVE3D code, a density model does not need to be provided explicitly. Instead, it is calculated as a function of velocity at a current iteration:

$$\rho = 310 \cdot (v_P)^{1/4} \tag{7.24}$$

where ρ and v_P are in kg/m3 and m/s, respectively. This is a special case of a more general relation $\rho = a \cdot (v_P)^b$ derived empirically by Gardner et al. (1974) for shallow-water continental sedimentary rocks. This relation is enforced for $v_P > v_{\text{Gardner}} = 1600$ m/s. For the water velocity $v_{\text{water}} = 1550$ m/s the density is set to $\rho_{\text{water}} = 1000$ kg/m3. For $v_{\text{water}} < v < v_{\text{Gardner}}$ a linear interpolation between ρ_{water} and $\rho(v_{\text{Gardner}})$ is performed. For $v \leq v_{\text{cutoff}} = v_{\text{water}}$ the model was fixed throughout the inversion. A more detailed discussion of strategies related to density, including benefits of setting non-physical values in the water layer, is provided by Nangoo (2013); Silverton et al. (2016).

Appendix

7.A `Jupyter` notebooks

For the sake of reproducibility of the described workflow, each of the above sections is accompanied with an associated `Jupyter` notebook. Data and wavelet processing is delegated to `Seismic Unix`. Model manipulation to `NumPy` and `SciPy`, whereas metadata analytics to `Pandas`. They are all wrapped in high-level routines and objects implemented in the `FullwavePy` package (`http://github.com/kmch/FullwavePy`). Advice on how to use the notebooks is described in sect. 8.6.

8

Kolumbo volcano

SUBMERGED volcanoes of continental arcs are hazardous due to their high explosivity (driven by magma-water interaction), potential for tsunami generation and proximity of human population. To properly assess the volcanic hazard, tight constraints on geometry of underlying magmatic systems (volumes, depths and connections between melt reservoirs) are needed. Unfortunately, standard geophysical methods have too low resolution to fulfil this objective.

This study applies a state-of-the-art seismic imaging method, full-waveform inversion, to image the topmost 4 km of the subsurface beneath Kolumbo, an active submarine volcano of the Hellenic arc located a few kilometres away from Santorini, Greece. Kolumbo last erupted in 1650 AD in a violent, phreatomagmatic[1] event, producing several km^3 of hydrous, biotite-rich rhyolitic lava (Cantner et al., 2014; Konstantinou, 2020), gas which poisoned a few tens of people (Nomikou et al., 2014a), and generating tsunami (Ulvrova et al., 2016).

The final P-wave velocity model reveals previously undetected, strong low-velocity anomaly of ca. -1.5 km/s relative to the starting model, at around 3 km b.s.l. directly beneath the crater. The effective-medium calculation shows that the anomaly may be associated with a high-fraction reservoir (magma chamber) of granitic melt.

The existence of a shallow magma chamber sheds new light on the Kolumbo plumbing system and continental arc volcanoes in general. It demonstrates that small shallow magma chambers, places where magmas pond prior to reaching the Earth's surface, may be re-assembled less than 500 years after the previous eruption.

Key words: back-arc volcanism · Aegean sea · Santorini volcanic field · full-waveform inversion · magma reservoir · submarine eruption · tsunami hazard

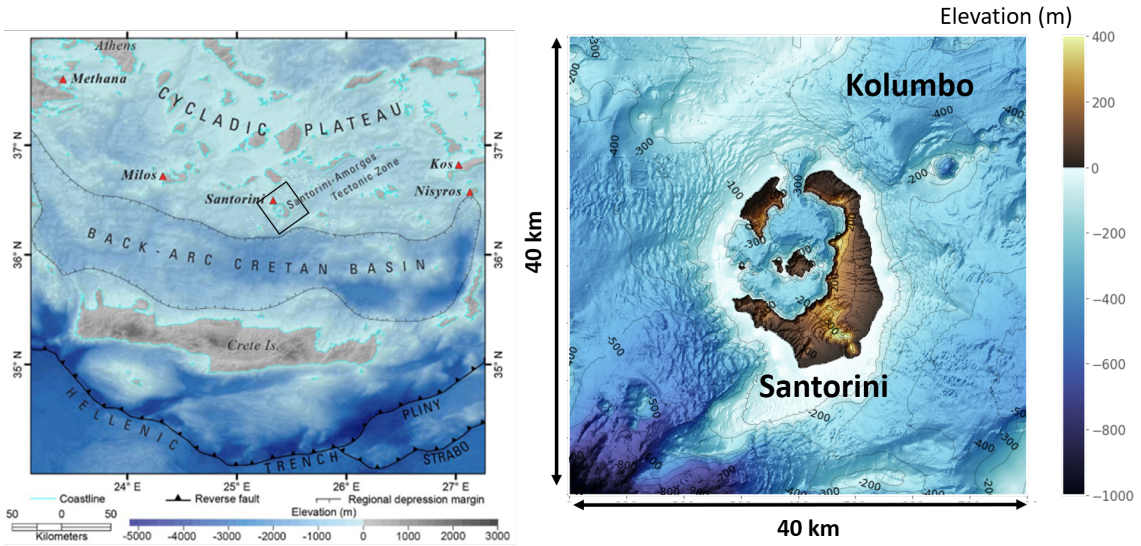


Figure 8.1: Topographic map of the Aegean Sea: (left) regional tectonics from Huff et al. (2021), (right) zoom into the study area.

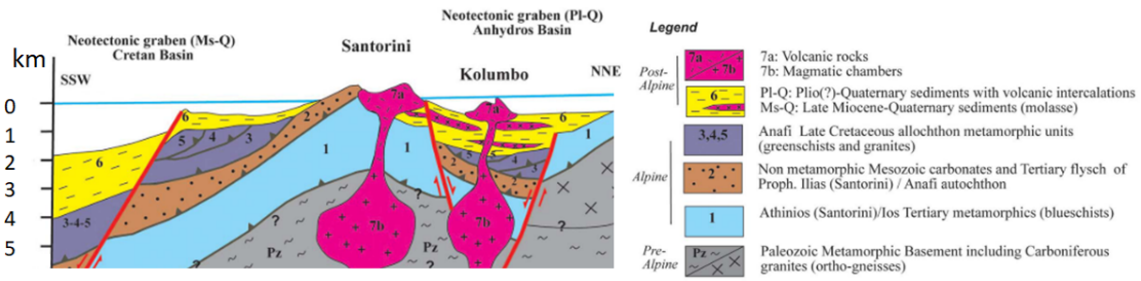


Figure 8.2: Hypothesised geological cross-section through the Hellenic Volcanic Arc, from the molassic back-arc Cretan Basin to the Cycladic Plateau near Ios. The country rock in the topmost 4 km beneath Kolumbo consists of interlaced metamorphic and non-metamorphic units. Modified after Kilias et al. (2013).

8.1 Introduction

Kolumbo (Coloumbo / Kolombo) is the largest of 23 submarine volcanic cones[2] stretching in the NE direction 7 km from Santorini (fig. 8.1). The crater is 1.7 km in diameter (7 km at the elongated base at around 504 m b.s.l.), up to 500 m deep, and reaches the height of 18 m b.s.l. The entire Kolumbo volcanic complex has a height of $\geq$ 1 km and a diameter of $\geq$ 11 km.

8.1.1 Tectonic setting

Kolumbo lies in the extensional sedimentary basin of Anhydros (fig. 2.4) bordered by marginal fault zones: Ios Fault to NE and Anhydros and Amorgos faults to SE. It is part of a longer lineament – the Christiana-Santorini-Kolumbo (Nomikou et al., 2019).

The edifice is built on top of the 10-15 km-thick continental crust of the Cyclades. The outcrops of the Palaeozoic metamorphic basement on Ios and Santorini suggest that the host rocks beneath Kolumbo

[1] i.e. involving interaction of magma and water.
[2] Including four discovered in the PROTEUS experiment (Hooft et al., 2017).

1879	First mapping of Santorini, including effects of the Kolumbo 1650 AD eruption by Fouqué.
2006	Discovery of a hydrothermal vent field by Sigurdsson et al.
2006	Discovery of a shallow bright spot beneath the crater in the 2d multi-channel[3] seismic data by Hübscher et al.
2006	Seismicity study by Bohnhoff et al.: most hypocenters found in 6-12 km depth range.
2009	Seismicity study by Dimitriadis et al.: most hypocenters found in 6-9 km depth range.
2010	Local-earthquake tomography by Dimitriadis et al.: low-velocity zone (-5% v_S) at 6-9 km b.s.l. offset from the edifice.
2013	ROV[4] study of the hydrothermal vent field by Kilias et al.: massive-sulphide deposits and iron microbial mats.
2013	ROV study of the hydrothermal vent field by Carey et al.: hazard related to accumulation of dense, CO_2 -rich water at the crater floor.
2014	Petrological study by Cantner et al.: calculation of pre-eruption p-T conditions and volatile content.
2016	Petrological study by Klaver et al.: isotope signature found to be distinct from Santorini.
2016	ROV study of the hydrothermal field by Rizzo et al.: high $^3He/^4He$ suggesting direct degassing of a MORB-like mantle through lithospheric faults.
2016	Tsunami modelling by Ulvrova et al.: caldera-collapse ruled out as a possible source mechanism.
2019	Active-source travel-time tomography by Heath et al.; Hooft et al.; McVey et al.
2020	Modelling of evolution of the magma chamber from Dimitriadis et al. (2010) by Konstantinou: not likely to erupt in 1 ka, unless its radius is smaller than 0.8 km.
2020	Acquisition of 3d multi-channel seismic data by Karstens et al.
2021	Seismicity study by Andinisari et al. refining the results of Dimitriadis et al.: most hypocenters found in 8–16 km depth range.
2021	Full-waveform inversion (this study).

Table 8.1: Time-line of major studies of the Kolumbo volcano.

consist of pre-Alpine (Palaeozoic) metamorphosed granites, and various Alpine metamorphic and non-metamorphic units including granites and flysch, covered with Cenozoic molasse interlaced with volcanic sediments (fig. 8.2).

8.1.2 Current activity

Kolumbo has been the most active source of microseismicity in the region, with earthquake hypocentres concentrated at 3-15 km depth (Andinisari et al., 2021). It contains an active hydrothermal field with polymetallic chimneys of sulphides and sulphates. Since its discovery by Sigurdsson et al. (2006), the interest in studying Kolumbo increased significantly (tab. 8.1). The source of the vigorous hydrothermal activity involving polymetallic chimneys (e.g. Kilias et al., 2013; Rizzo et al., 2016) was imaged to be at around 200 metres below the seafloor (Hübscher et al., 2015). The hydrothermal activity is thought to lead to the outgassing of potentially hazardous mantle-derived volatiles (Rizzo et al., 2016) and also provide nourishment to some marine life (Kilias et al., 2013).

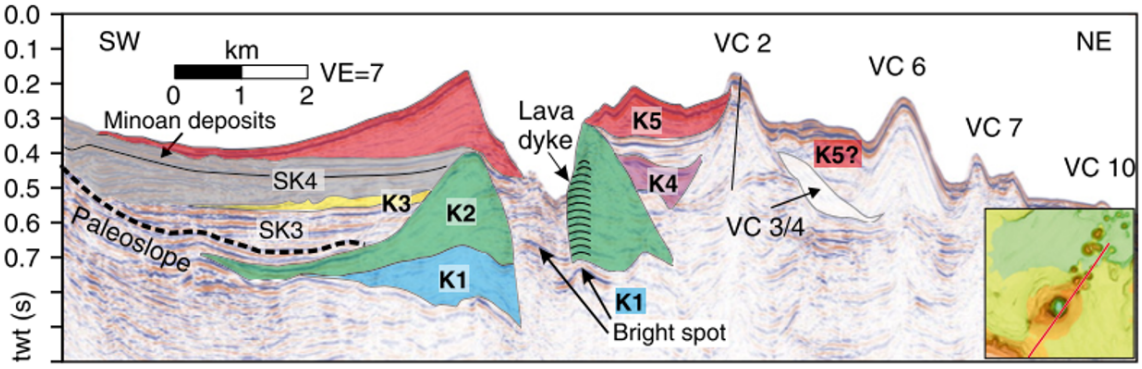


Figure 8.3: Seismic-reflection time-section through shallow sediments around Kolumbo. The K5 unit is associated with the most recent 1650 AD eruption. Note the bright spot at about 200 m below the crater floor (Hübscher et al., 2006). VC – volcanic cone. Modified from Hübscher et al. (2015).

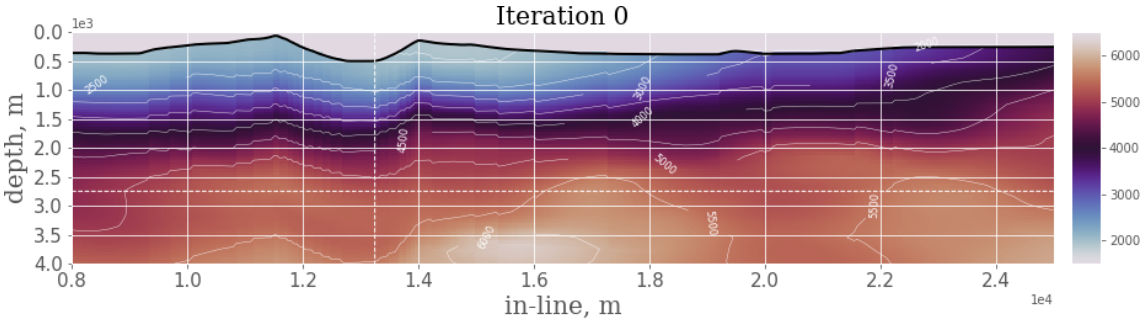


Figure 8.4: P-wave velocity model (in m/s) across Kolumbo obtained by Heath et al. (2019), sliced at $y = 8.75$ km. It served as a starting model for this study.

8.1.3 Eruption history

Reflection seismic images revealed cone-shaped units beneath the Kolumbo volcano that are interpreted to have evolved through explosive eruptions (Hübscher et al., 2015). A record of three major (K1, K2 and K5) and one or two (K3 and K4) smaller eruptive phases has been recognised (Nomikou et al., 2016), where K1-K5 are volcanic units produced during different eruptive phases as intepreted from seismic reflection profiles (fig. 8.3). The total volume of volcanic sediments was estimated to fall between 13–22 km^3 (Hübscher et al., 2015). On a reflection profile across Kolumbo there are phase-reversed bright spots a few hundred meters below the crater floor (fig. 8.3) which were interpreted as a direct indicator for a gas reservoir by Hübscher et al. (2015).

8.1.4 Magmatic system

Despite its eruptive history, ongoing activity and tsunami hazard (Nomikou et al., 2014a), very little is known about the magmatic system below Kolumbo, in particular in relation to Santorini. Whereas petrologic analyses suggest the existence of distinct crustal reservoirs feeding these two volcanoes (Klaver et al., 2016), local-earthquake and active-source tomography (Dimitriadis et al., 2009, 2010) indicate that they may be linked below the surface. Subsurface models of seismic velocity (Hooft et al., 2019; Heath et al., 2019; McVey et al., 2020) show an elongated low-velocity zone (fig. 3 of McVey et al. 2020) at

[3] In this method, multiple source-receiver pairs which image the common mid-point (CMP) in the subsurface are stacked to enhance the S/N ratio.

[4] Remotely-operated vehicle.

about 3-km depth below surface along the so-called Kolumbo line of volcanoes between Santorini and Kolumbo. This low-velocity zone correlates with reflection-imaged faults (Nomikou et al., 2016) and dikes in the walls of Santorini caldera (Heath et al., 2019). Whether the two volcanoes are connected, however, remains unclear due to resolution limits that are inherent in travel-time tomography. Apart from that, McVey et al. 2020 did not confirm the presence of crystal mush zone at 5-7 km depth reported by Dimitriadis et al. (2010), nor any other, more shallow, reservoir of melt. Such reservoirs would also be difficult to image using conventional seismic reflection data due to the structural complexity of the volcanic caldera close to surface.

The depth of the magma chamber beneath Kolumbo submarine volcano was estimated to be at around 5-7 km depth using local earthquake tomography (Dimitriadis et al., 2010) and geochemical analysis (Konstantinou, 2020). The model presented by Hooft et al. (2019) contains slightly reduced P-wave velocities (relative to the average 1d model of Anydros metamorphic block) at depths greater than 2 km b.s.l. directly underneath the Kolumbo edifice (fig. 8.4) but this result was not further discussed by the authors. McVey et al. (2020) who utilised data from the same experiment but with longer offests were able to probe larger depths down to 6 km b.s.l. They defined a linear low-velocity volume (LLV) by 5.3 km/s contour at 3.4 km depth which in their model does not extend to the Kolumbo where velocities are a bit higher (5.4 km/s). This velocity is consistent with very little to no melt (0-1% volume fraction) in contrast with the low-velocity volume under the Santorini caldera with at least a few percent of melt (McVey et al., 2020).

Klaver et al. (2016) argues that Kolumbo and Santorini have different magma differentiation history based on geochemical signature of amphibole fractionation and the assimilation of crystalline basement rocks. Hydrous phenocrysts and high Nb/Yb>3 and low 206Pb/204Pb<18.82 seem to indicate that the Kolumbo magmas are more enriched in mantle material. Whereas eruption as soon as four centuries after the last major event does not seem very probable (the previous eruption happened 22 ka), there is always a possibility of a magma recharge and remobilisation. Although the magnitude of any eruptions is likely to be small relative to Santorini, Kolumbo still poses a major threat due to a multitude of possible tsunami-generating mechanisms (Nomikou et al., 2014a).

8.1.5 This study

The goal of this study was to refine the travel-time tomography model of the upper 4 km beneath Kolumbo (Heath et al., 2019; McVey et al., 2020) by means of *full-waveform inversion* (Virieux and Operto, 2009) using data collected in the same experiment. The method of choice promised to recover structure 2-4 times finer than the travel-time tomography, even with the relatively sparse acquisition coverage (Morgan et al., 2016). The inversion was intended to focus on *transmission data*, i.e. forward-scattered (refracted) waves (Mora, 1988) which allow to image a wide wavenumber spectrum of the structure (Virieux and Operto, 2009). The unprecedented resolution was expected to reveal potential shallow magma reservoirs which could feed the eruptions in coming years.

8.2 Data

8.2.1 Acquisition

The data were collected in 2015 as a part of the PROTEUS experiment (Hooft, 2015a; Hooft et al., 2017). Details of the pre-processing, initial QC, station-positioning, source spectrum, receiver responses etc.

are given in ch. 3. A short summary relevant for this case study is provided below.

Airgun source Most of the recorded data consist of seismic waves generated by a 36-element, 6600 in.^{3} airgun array towed at 12 m below the sea surface behind the R/V Marcus G. Langseth, shot ca. every 150 m along many profiles over a distance of 2500 km in total.

OBS sensors The data were recorded by 61 ocean-bottom seismometers (OBSs) from Scripps Institution of Oceanography (SIO) and 30 OBSs from Woods Hole Oceanographic Institution (WHOI). Each OBS consisted of (1) a hydrophone (High Tech HTI-90-U) with amplitude response to hydrostatic-pressure variation flat between 6-80 Hz and dropping by two-thirds at 3 Hz; (2) a three-axis geophone (Sercel L-28-3D or Geospace GS-11D, in SIO and WHOI OBS respectively) with amplitude response to velocity flat above 4.5 Hz and dropping by about two-thirds at 3 Hz. This gives a total of four components (p, v_x, v_y, v_z) recording 60 s of signal sampled at 200 Hz.

Instrument response Even if the frequency is in the non-flat, rolling-off part of the spectrum, the amplitude-drop tends to affect both signal and noise equally, keeping S/N ratio more or less constant. Relative amplitude differences at different frequencies were expected to play a minor role when inverting frequencies in a progression from lowest to highest. This is because at each step of such progression the band-width of the data are increased only by a small amount (0.5 Hz in this study) and the inversion iterates until convergence (after 20 iterations the fit changes very slowly) before the bandwidth is increased again. In this way, bandwidth of data inverted at each step is narrow and thus less affected by the amplitude roll-off.

8.2.2 Data subset

The PROTEUS acquisition described in the previous section covered a large area and gathered a substantial amount of data. In principle, one could attempt to constrain the entire model domain by inverting all the data. In practice, this may not be feasible because FWI requires a careful input preparation and extensive quality control at every stage of the workflow (Morgan et al., 2013). Instead, the approach followed here was to use the smallest possible subdomain of the experiment (and the corresponding data subset) that had a potential to sufficiently image the target. This is because:

- small datasets can be visually inspected and hand-picked,
- small models allow for extensive parameter testing in a reasonable computational time.

The process of selecting such a data subset and model subdomain (box) is described below.

Constraints As this study was targeted at the topmost 4 km of the subsurface (sect. 8.1.5), this puts a constraint on a minimum offset. The rule of thumb mentioned in sect. 4.4.2 suggests that a minimum offset of 12 km, and more realistically about 15 km, is needed to illuminate the subsurface down to the intended depth of 4 km. Another constraint comes from the proximity of the Kolumbo to the Santorini caldera. The lateral extent of the model was limited by the topographic relief of the island. Otherwise, one would have to either invert also the land data which has proven to be quite challenging (see ch. 6), or include regions of very sparse OBS coverage[5]. Benefits of a few (if any) clean, long-offset traces

[5] In general extending the model into regions of sparse coverage makes the inverse problem more mixed-conditioned. Some parts of the model (hopefully the target) might become better constrained, whereas others, most likely the ones just appended, are going to be poorly constrained, if at all.

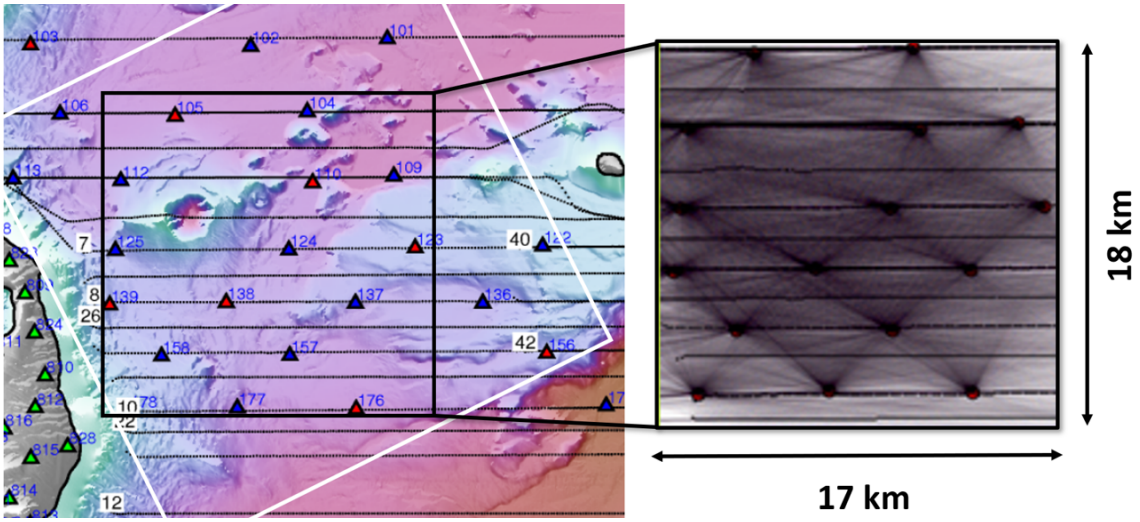


Figure 8.5: Acquisition geometry of the domain of study (black rectangle) with the coverage density represented by lines between sources and receivers with offsets up to 20 km. The local system of coordinates was based on the Universal Transverse Mercator coordinates rotated 25.5° counterclockwise from north with the centre $(0, 0)$ m at ca. $36.4°\mathrm{N}, 25.4°\mathrm{E}$. Red and blue triangles denote OBSs from WHOI and SIO pools, respectively. The leading digit '4' is not shown in the station numbers. Shots are marked by black dots with line numbers in white squares. The large white rectangle denotes a potential extension of the domain for future inversions.

recorded by intracaldera OBS 4140 and 4160 (fig. 3.4) do not seem to outweigh the computational cost of the larger model domain. Thus, it was decided to effectively move the target from the centre towards the top-left corner of the model box (fig. 8.5).

Model box Given the above constraints a few different caldera-bounded boxes were investigated by synthetic tests in terms of illumination of the target area. The best trade-off between a computational cost and target illumination was found to be provided by a $17 \times 18 \times 4$ km box (fig. 8.5) corresponding to the local coordinates of $x \in [8, 25]$ km in-line, $y \in [-3, 15]$ cross-line, and $z \in [0, 4]$ depth b.s.l. It contains 16 OBS stations (including six WHOI and ten SIO instruments) and 1505 shots which gives 24,080 traces per component. All of them have acceptable quality of the hydrophone component and all but one on the vertical velocity component which gives a total number of traces reaching 46,655 of varied, but overall relatively high data-quality. In this study, the focus is on the hydrophone component only (denoted by '4' in the leading digit of the station number). The inversion of geophone component is presented in ch. 5 and the joint inversion of hydrophone and geophone data is the subject of current work.

Water depth A maximum water depth in the selected model box was ~500 m with most shots and receivers located in areas of 250-450 m water depth (fig. 8.6). This relatively shallow water caused Pg refractions to become first arrivals at offsets shorter than 5 km (fig. 7.5) and the multiples to arrive quite early (around 0.67 s, assuming vertical ray paths and 500-m water depth) after the primary pulses.

Coverage The maximum offset contained the model box was 20 km but it was reached by few traces (fig. 8.7). Most traces fell in the 10-15 km offset range (fig. 8.1). There was a good azimuthal coverage for offsets up to about 15 km as a result of a square-like shape of the domain and relatively regular acquisition geometry. As in most other parts of the experimental domain, shots were about 150 m

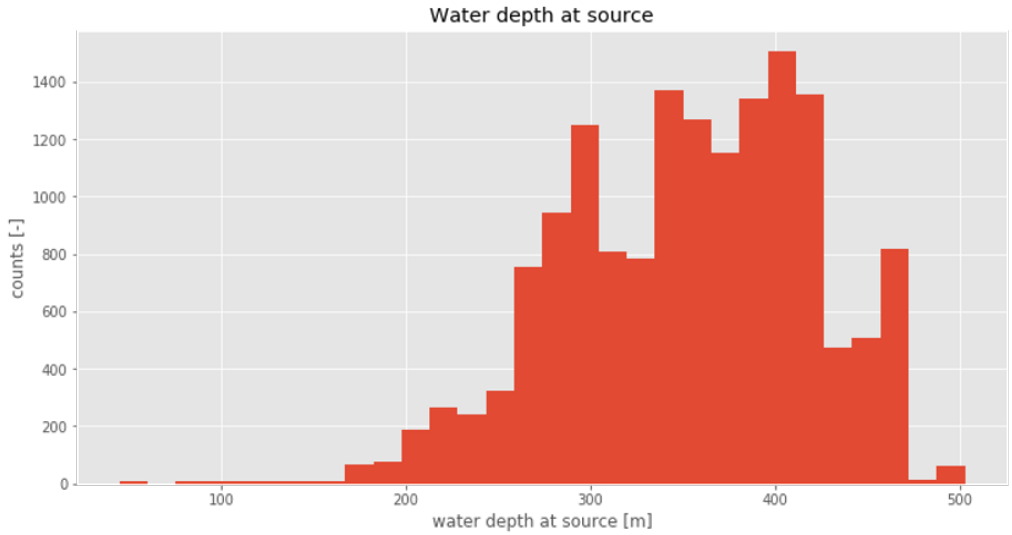


Figure 8.6: Histogram of water depths at shot positions in the study area. Most shots were fired in water of 250-450 m depth.

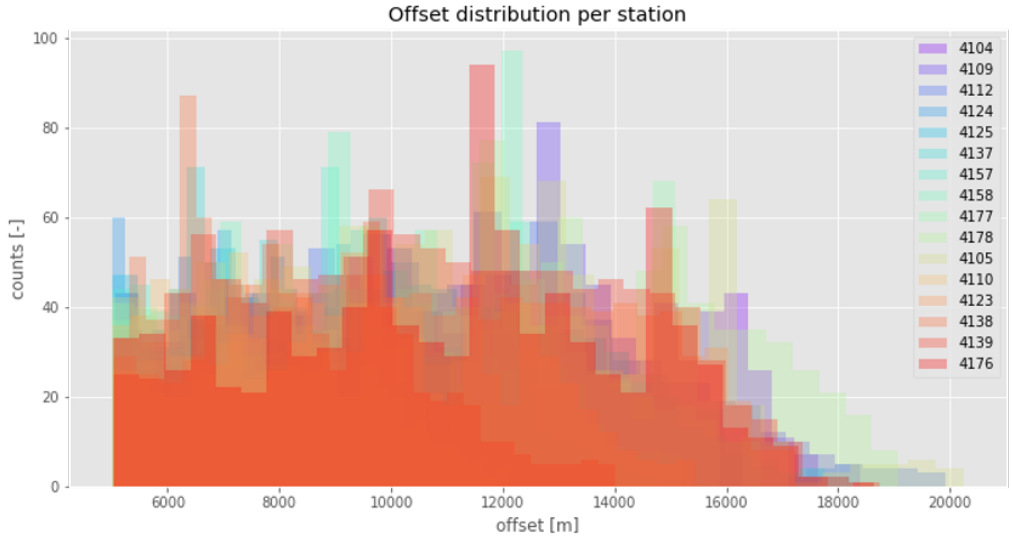


Figure 8.7: Histograms of source-receiver offsets for every OBS (semi-transparent colourscale) in the study area.

apart with shot-line spacing of about 1 km. OBSs were about 5 km apart and although it was as dense as possible in the PROTEUS experiment, data sparsity was expected to be the main factor controlling the performance of the inversion. A case study with similar acquisition features (though a slightly denser shot coverage) was published by Morgan et al. (2016).

Time duration Acquisition geometry and velocity of the subsurface both control duration $t_{\max}$ of seismograms required to record arrivals of interest. At minimum, this should be the time of the latest first arrival time of the latest first break plus the width of the applied time window. Choice of such a window is discussed in 8.2.3. To save computational resources, $t_{\max}$ was optimised to include just the entire time window for all the traces but no arrivals beyond that window. The optimum value was found to be 5 s.

8.2.3 Processing

The data were only minimally processed in order to preserve the low-frequency part of the spectrum (Warner et al., 2013b). Above-mentioned low and high-pass filtering was applied simultaneously as a minimum-phase frequency-domain band-pass filter with corner frequencies 2.0 – 3.0 – 4.5 – 6.5 Hz. Details of this procedure, involving tapering and zero-padding the traces, as well as the implementation of the filter are described in sect. 7.2.

High-pass filter Analysis of phase plots generated for the 60-km-long box (fig. 8.8) revealed that the S/N ratio was low up to frequencies lower than 3 Hz due to the presence of swell noise affecting shallow ocean-bottom seismometers. A higher S/N ratio characterised the 4 Hz cut-off, but it is desirable to start at the lowest possible frequency to first adjust the low-wavelength features of the model. As expected, longest offsets had overall lowest S/N ratio but the noisy traces were not pervasive enough to justify clipping the offset to some maximum distance. Clearly, the 2-Hz frequency was much more dominated by noise for almost all stations, and hence unusable. The 4-Hz frequency exhibited similar or better S/N ratio compared to 3 Hz, but eventually the latter was chosen to avoid cycle-skipping and enhance long-wavelength updates as explained in sect. 8.2.3.The amplitude rolled off at about 18 dB/octave, falling below instrument nominal cut-off frequency at 2 Hz.

Low-pass filter Initially, to avoid potential problems with higher-frequency noise, e.g. a 6-Hz ringing noise detected on all four channels and not correlating with atmospheric conditions or boat traffic (Heath et al., 2019), low-pass filtering cutting off frequencies above 6 Hz was performed. The intention was to relax this later to include higher frequencies in the inversion (see sect. 8.6.2).

Time window As mentioned before, only early-arriving refractions were planned to be inverted. Other arrivals, including direct (water) waves, primary reflections, free-surface multiples, P-S-P converted reflections, etc., increase non-linearity of the forward problem making the Born approximation utilised in the inverse step less valid. Excluding deeper reflections was done by applying a bottom mute. An additional, top mute was also applied to remove noise ahead of the first arrivals. The procedure requiring a knowledge of the first-break times is discussed in detail in 7.2. Two time windows were tested, with lengths of 1 and 2 seconds, and with 0.25 s tapers at both ends. The 1-sec window contains the higher-amplitude part of the early-arriving wavefield, whereas the 2-sec window includes its full coda (fig. 5.18). The results in sect. 8.4 are shown for the 1-sec window.

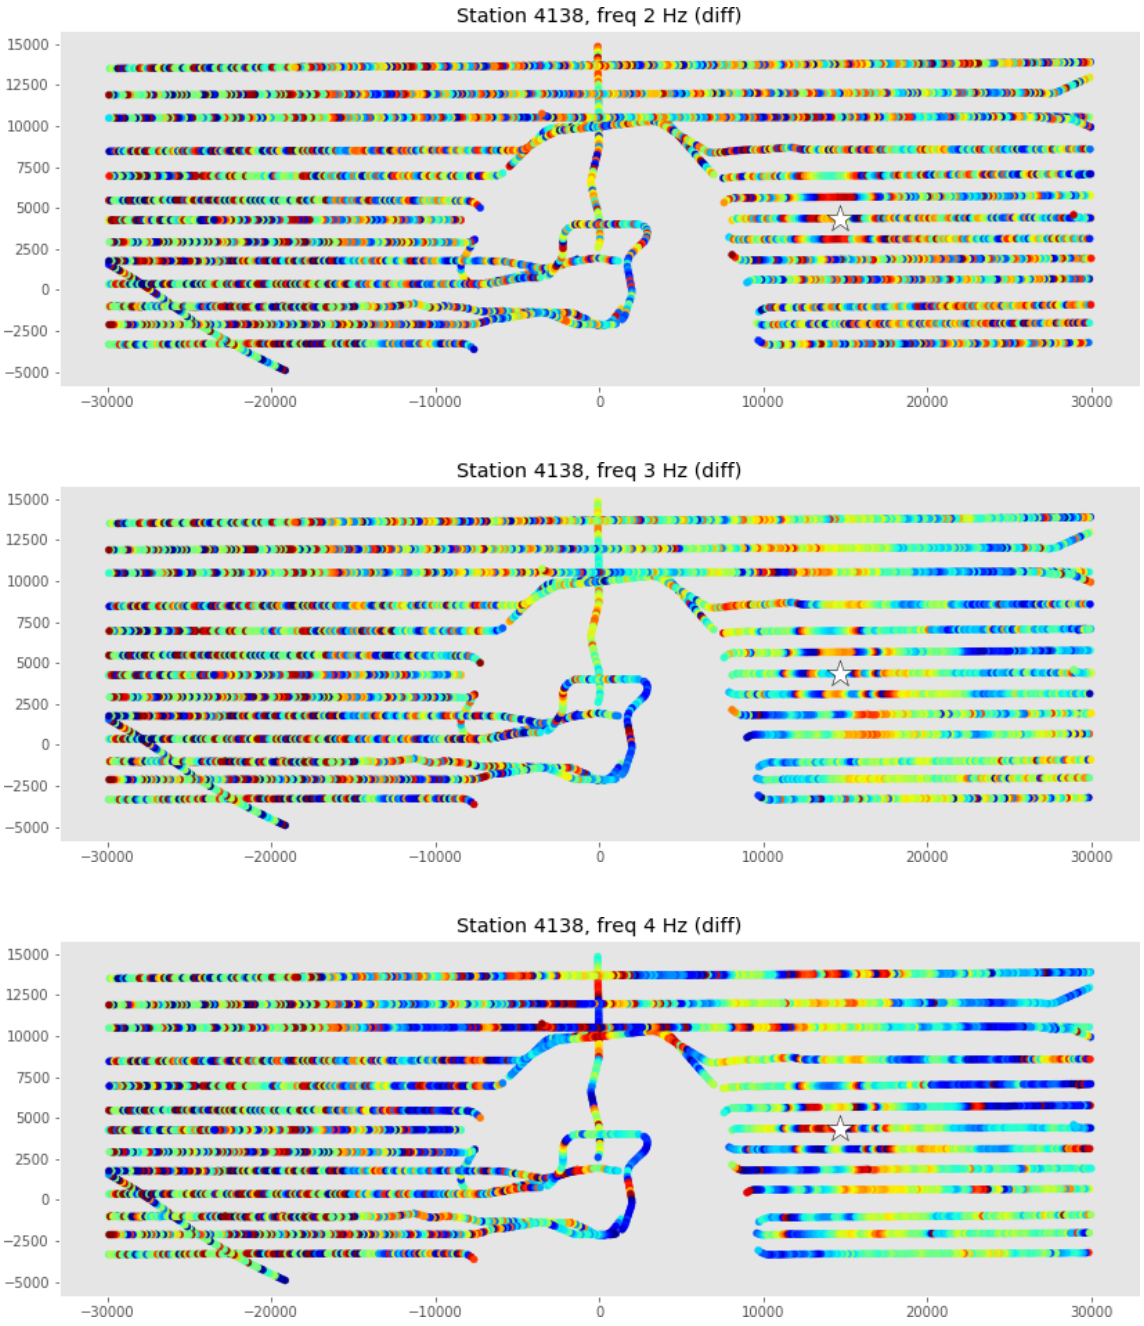


Figure 8.8: Phase plots (Shah et al., 2012a) used to determine the minimum invertible frequency. Each dot represents a phase residual (phase of synthetic data minus phase of observed data, wrapped into $[0, 2\pi)$ rad interval) of the signal extracted at 2, 3 or 4 Hz (from top to bottom, respectively). Colour scale from $-\pi$ (blue) to π rad (red), green denoting a zero residual, and thus good data fit. Incoherent colour variations indicate low S/N ratio. Note that the noise affects the long-offset data at each frequency. White star denotes the location of the OBS. Axes scaled in local coordinates in metres.

Offset window The traces with source-receiver offsets smaller than 5 km were excluded from the dataset. The first arrivals in this offset-range consist of water waves and subhorizontal head waves refracted in the seafloor sediments (fig. 7.5). They have large amplitudes that would dominate the inversion and they are sensitive to the seabed topography and shallow velocity structure which are not resolvable with this sparse dataset. The value of 5 km was determined as a minimum value valid for all the stations contained in the domain. For some stations this may be overestimated, but having a single, station-independent condition is easier to implement. For deeper basins around Santorini this value may be slightly underestimated. As the SNR was acceptable and the data were not cycle-skipped for the longest offsets (fig. 7.5 and 8.13), the upper limit for offsets was not imposed.

8.3 Method

8.3.1 Synthetic-data calculation

To calculate synthetic seismograms one needs a numerical solver of an appropriate wave equation, a model and a source wavelet, which constitute the space-dependent coefficients and time-dependent forcing term of that wave equation, respectively.

Numerical solver An explicit, time-domain finite-difference solver of the acoustic, isotropic, 3d wave equation was used. It is a part of the `FULLWAVE3D` package developed at Imperial College London and described by Warner et al. (2013b). Among the suite of implemented finite-difference stencils, the one selected for this study was a compact ($5 \times 5 \times 5$) operator with a nominal accuracy of 4th order in space, 2th order in time, least-squares-optimised to give the effective accuracy of about 6th order in space, 4th order in time. This sets an approximate criterion for spatial discretisation to be about five (down to four) grid nodes per wavelength. These numbers are not entirely meaningful without providing information on propagation distances (Igel, 2016). They were, however, numerically estimated by the code developers based on the realistic 3d acquisition geometries, with offsets larger than the ones considered here (`FULLWAVE3D` manual, unpublished). Therefore, it was assumed that unless the above accuracy criterion is violated, numerical dispersion will not change the waveform enough to affect the inversion. As the number of shots (> 1000) was much greater than of receivers (16), source-receiver reciprocity was utilised, labelling sources as receivers and receivers as sources. The sea surface was simulated as a flat, planar free surface located one node above the top of the model grid. Instead, an absorbing layer could have been used which would require suppressing all sea-surface-related ghosts and multiples in the field data and it is preferential in FWI to perform as little data pre-processing as possible (Warner et al., 2013b).

Discretisation Solving the wave equation numerically requires a discretisation of the domain defined in sect. 8.2.2 either by point-wise sampling or basis functions, the former being a method of choice in finite-difference schemes like here. The time interval, defined as the maximum time included after the source origin time (see 8.2.2), has to be discretised as well. The `FULLWAVE3D` code uses a regular sampling of both space and time with dx (same for all three spatial directions) and dt being their respective cell sizes. A value of $dx = 50$ m was chosen to allow for accurate modelling in water up to 6.5 Hz (fig 8.9). For computational efficiency, one could downsample the grid down to $dx = 75$ m for frequencies up to 4 Hz, but this was not pursued as the model size with a 50-m grid spacing was not too large. The velocity model has $341 \times 361 \times 81$ grid cells and $401 \times 421 \times 111$ grid cells when absorbing boundaries are included, which is a similar order of magnitude to the $321 \times 261 \times 81$ grid cells

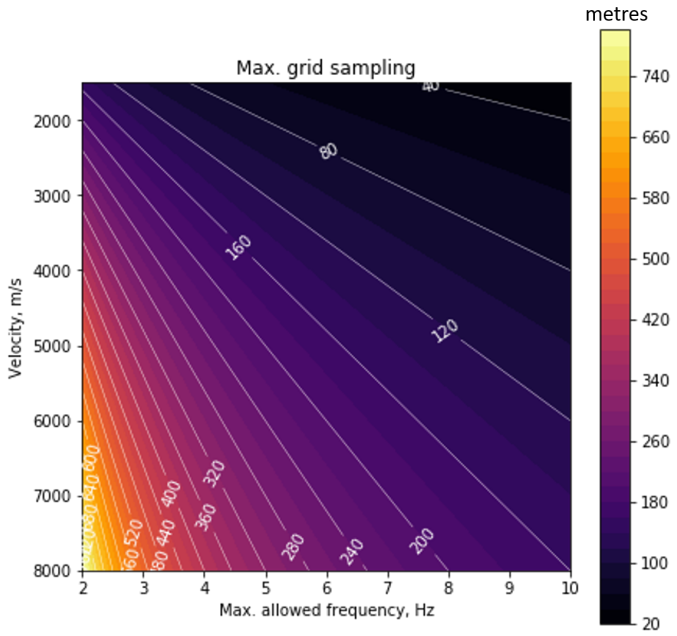


Figure 8.9: Maximum grid sampling (colour scale) for a numerical scheme accurate for at least 5 grid nodes per wavelength, as a function of the highest frequency of the source (X axis) and the lowest velocity present in the model (Y axis).

(counted with boundaries) used to invert the Tommeliten data (Warner et al., 2013b). The total computation time for the Kolumbo inversions is shorter, however, as many more shots are included in the Tommeliten inversions. For the chosen scheme, the stability considerations constrain the ratio dx/dt to be at most $1/2$. The constant quotient explains why computational cost grows as 4th power of the model size (actually as a 6th power, if one considers a growing number of sources, assuming surface acquistion). Halving dx enforces halving dt for the algorithm to remain stable. The chosen value was $dt = 0.0025$ s (2.5 ms) which was half the time sampling of the raw data.

Source wavelet A *source wavelet* will be defined as a theoretical, source-generated waveform shifted to zero time, that would be recorded in a far field of a homogeneous 3d space. Note that this definition implies that the wavelet contains no source ghost, but it does contain receiver response and source radiation pattern. In reality, it is a superposition of a direct wave, a source ghost and other phases sampling the medium around the station. Errors in the source wavelet may lead to significant artifacts in the final models (Warner et al., 2013b), so it is essential to either invert for it or estimate it *a priori* as accurately as possible. Here the latter was chosen, with the workflow described in detail in sect. 7.3.

Starting model One needs only two parameters to fully describe wave propagation in acoustic, isotropic medium (validity of these modelling assumptions are discussed in sect. 8.5.1). These are density and (P-wave) velocity. In this study, this number was further reduced to one, as the density was calculated from the P-wave velocity according to Gardner's relation (Gardner et al., 1974). This relation was estimated empirically for marine sediments and it is certainly less accurate for volcanic products and crystalline basement probed by this survey. This should not be a major issue though, as refracted P-waves are known to be much less sensitive to density than velocity (Płonka et al., 2016). The model of velocity was prepared from the slowness model obtained from travel-time tomography performed

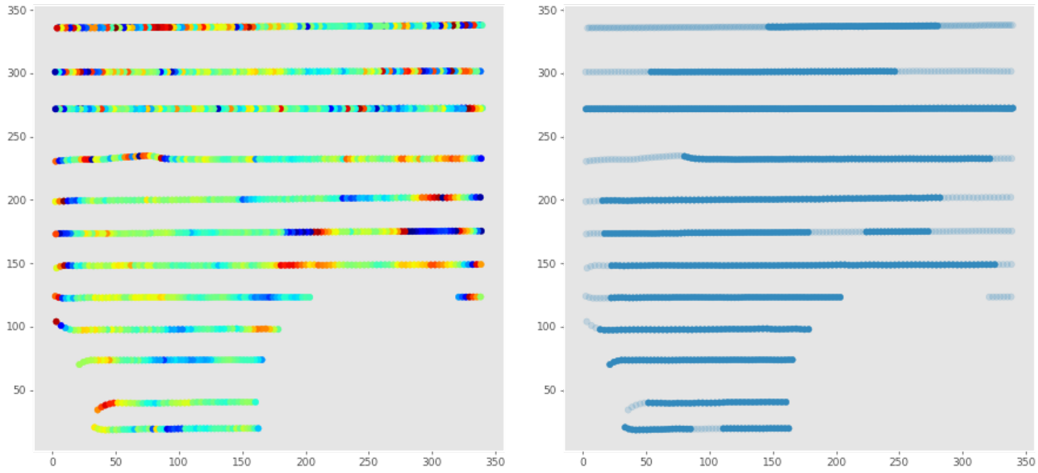


Figure 8.10: Phase residuals (left) at 3 Hz calculated for shots recorded by OBS 4176 and acquisition from fig. 8.5 with source-receiver offsets < 5 km excluded. The colour scale is the same as in fig. 8.13. The right figure shows the shots input to the inversion (dark blue) and shots excluded due to being noisy or cycle-skipped (grey). Axes scaled in grid nodes ($dx = 50$ m).

by Heath et al. (2019). The procedure is described in more detail in ch. 7.4.

8.3.2 Synthetic vs. field data

In order to estimate S/N ratio of the field data and adequacy of the synthetic data, these two datasets were compared in terms of a phase spectrum of the early arrivals wrapped into $[-\pi, \pi]$ rad range (Shah et al., 2010, 2012a). The method enables a rapid quality assessment of entire shot/receiver gathers at a set of discrete frequencies. This is much faster than line-by-line or trace-by-trace analysis in the time domain.

Quality of field data Phase plots of observed data were used to detect and further exclude noisy traces with S/N ratio below a given threshold. In particular, they were used to constrain the usable offset range and to guide the choice of the lowest possible starting frequency for FWI. The exclusion criterion was the lack of continuity (lack of coherency) mentioned above. Note that preparation of these plots requires knowledge of travel times of the first arrivals. They were estimated by full-waveform modeling through the starting model. Furthermore, a manual removal of bad-quality data via lasso-selector implemented in `FullwavePy` package was applied for all shot-gathers (fig. 8.10). This was aimed at contiguous regions of noisy or potentially cycle-skipped data. The rationale was to decrease amount of noise-related errors mapped into final models as artefacts. Single, isolated traces with extreme values close to $\pm\pi$ were omitted as expected to have little influence on the inversion.

Adequacy of synthetic data Phase plots of synthetic data (Shah et al., 2012a) can be used for stations with clean observed data to assess adequacy of the synthetic data. Both adequacy of the input and modelling method (approximation of wave physics, absorbing boundaries etc.) contribute to the phase residual and are impossible to separate without further assumptions. Here, it is assumed that

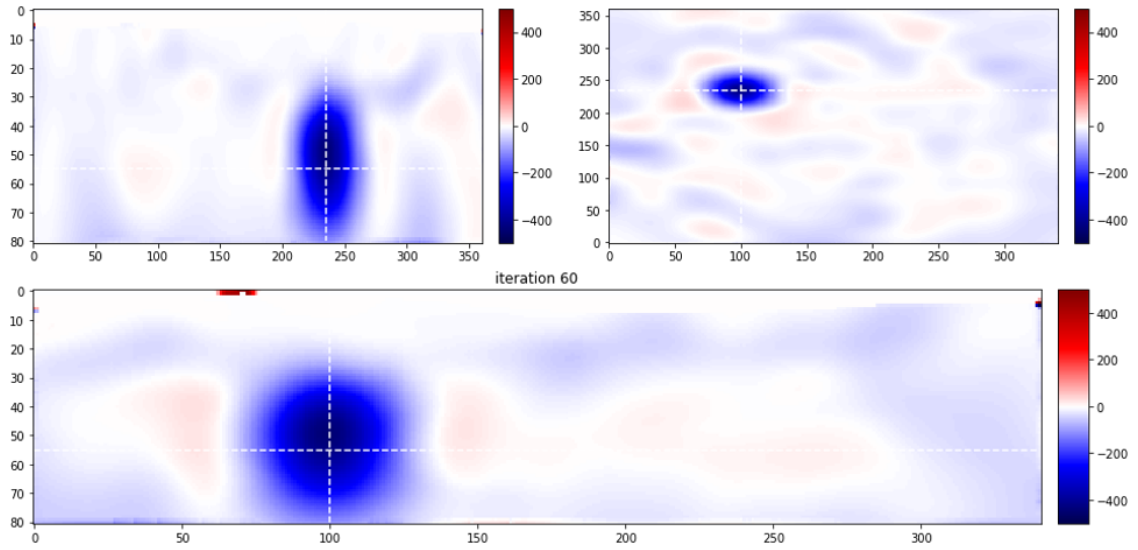


Figure 8.11: Anomaly recovered (final minus background model) after 60 iterations of inverting synthetic data generated from the same model as in fig. 5.20a. Top left – x slice, top right – z slice, bottom – y slice. Dashed white lines denote location of slices. Axes scaled in grid nodes ($dx = 50$ m).

the acoustic solver accurately models the early arrivals (see sect. 8.5.1) from which the phase was extracted, and that the source-wavelet errors can be neglected as shown by the good match of direct waves (see sect. 7.3). Under above assumptions phase-residual plots can be used to assess the quality of the starting model, in particular to detect problems with data cycle-skipping. Such problems are indicated by abrupt red-to-blue or blue-to-red[6] changes between neighbouring stations. Overall, there were no signs of cycle-skipping of the refracted waves, even for the longest offsets of the survey (fig. 8.13). This suggest not imposing any maximum-offset constraint in the inversion.

8.3.3 Synthetic-data inversion

Initial inversions led to the recovery of a low-velocity anomaly beneath the Kolombo caldera. A suite of synthetic tests were performed in order to explore how well constrained this feature was, and to identify an effective inversion strategy for its recovery.

True model The background model was defined to be the final model from the travel-time tomography (sect. 8.3.1) which will be later used as a starting model in the field-data inversion. The true model was created by perturbing the background model with a negative low-velocity anomaly of a spherical symmetry. Anomalies of this kind are a basis of the so called *spike test* (Rawlinson and Spakman, 2016). The recovered anomaly can be thought of as *point spread function* (PSF) in a sense of Fichtner and Trampert (2011). The shape of anomaly was defined as a not normalised, zero-correlation, 3d normal distribution:

$$P(x,y,z) = A \exp\left[-\frac{(x-\mu_x)^2 + (y-\mu_y)^2 + (z-\mu_z)^2}{2\sigma^2}\right], \tag{8.1}$$

where $\sigma = 1$ km and $\overrightarrow{\mu} \equiv (\mu_x, \mu_y, \mu_z)'$ was equal to (13.25;8.75;2.75) km or $(105, 235, 55)$ grid nodes with various peak amplitudes A relative to the background. The following analysis considers the case

[6] Or another pair of end-member colours for a different colourmap.

of $A = -0.10$. The true model was calculated as:

$$v_{\text{true}}(x, y, z) = v_{\text{bckgr}}(x, y, z) \cdot [1 + P(x, y, z)] \,. \tag{8.2}$$

These formulae are also the basis of the recovery tests described in appendix 8.B.

Inversion strategy The data were normalised in a way described by Warner et al. (2013b). This normalisation scheme reduces the effect of wave phenomena not modelled by the selected solver. Water velocity and seafloor topography were known a priori and thus kept fixed. For dense datasets the FWI gradient is self-regularising due to contribution of many stations (Warner et al., 2013b). For sparse datasets like here, regularisation is needed. The results show that the inversion introduces spurious perturbations to the background model outside of the anomaly (fig. 8.11). These artifacts however, have negligible amplitude compared to the recovered anomaly, and the recovery after 60 iterations is acceptable (fig. 8.11). This shows that, for the model under study, the true acquisition geometry, model grid and data processing discussed in sect. 8.2.3, combined with smoothing over 1-2 wavelengths result in a unambiguous recovery of the LVZ beneath Kolumbo. A range of parameter suites related to frequency, offset and regularisation was inverstigated using similar tests. The final strategy involved inverting 1 sec of pressure data for offsets > 5 km at subsequent low-pass frequencies of 3.0 to 6.5 Hz (with a step of 0.5 Hz), running 20 iterations at each frequency block. The data were low-cut filtered at 2-3 Hz which means there is still significant energy in this range, and it affected both noise and signal. The final strategy was applied to anomalies of different sizes and amplitudes (see appendix 8.B).

Parameter tuning The sensitivity to most free parameters of the used inversion code was also tested. These parameters included different objective functions (L^2 and L^1 norm of data residuals), model parametrisations (velocity, slowness), optimisation algorithms (steepest descent, conjugate gradient), finite-difference kernels (of different accuracy than the one used to generate synthetic observed data), absorbing boundaries. The inversions were stopped after the first block of iterations at the lowest frequency, to check the quality of the low-wavenumber update. It is well understood that inverting higher-frequency data usually leads to improvements in models that are already proceeding in the right direction, whereas models that appear worsen during the low-frequency inversions tend to get worse.

8.4 Results

In this section, results of the field-data inversion are presented. The strategy and parameters derived in the previous section were applied to the inversion of the field data. The optimal smoothing of the gradient, however, was again searched for, in the range between 0.5 and 3 local wavelengths. The rationale for investigating the smoothing again was to check if stronger regularisation might be required when inverting the field data, due to the presence of noise. Other parameters were kept at their optimal values as they were considered to be less likely to depend on data quality. A double smoothing (i.e. gradient smoothed over two local wavelengths) was found to be the best compromise between model roughness and data fit.

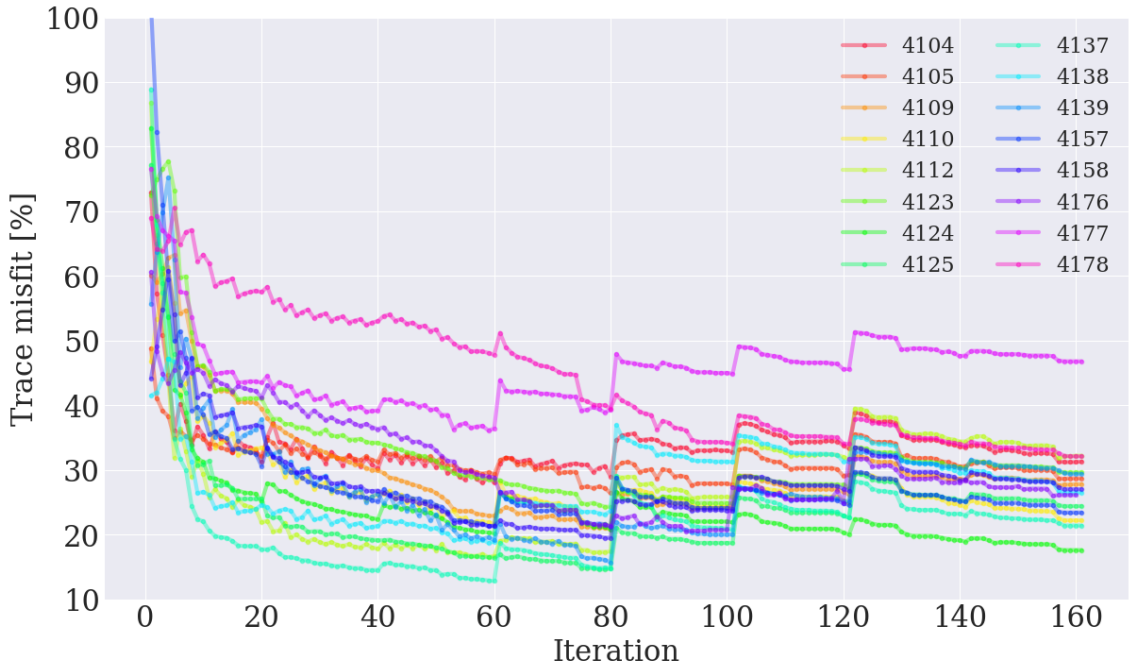


Figure 8.12: Evolution of the objective function across iterations for each OBS (colour-coded).

8.4.1 Objective function

As a first step of assessing the inversion performance, the objective function defined as a difference between normalised data:

$$\frac{100\%}{N} \cdot \left\| \frac{\mathbf{d}_{\text{obs}}}{||\mathbf{d}_{\text{obs}}||} - \frac{\mathbf{d}_{\text{syn}}}{||\mathbf{d}_{\text{syn}}||} \right\| \tag{8.3}$$

where $||\cdot||$ denotes the L^2 norm and N is a normalisation constant, was plotted for every OBS and iteration (fig. 8.12).

The reduction in misfit across all iterations was significant for all the shots (30-80%). This confirms that all of the OBS stations (reciprocal shots) had usable data of similar quality because noisy data would be cycle-skipped at the start and barely improve across iterations. Most of the misfit-reduction occured during the first iteration block with a cut-off frequency of 3 Hz. Rather mild changes of misfit characterised each of the three iterations blocks with high-cut frequencies in the range 4.5-5.5 Hz (iterations 60-120). Station 4177 stands out as showing the least improvement in data fit (fig. 8.12), with the second worst station being its immediate neighbour on the bottom-most shotline, 4178 (fig. 8.1). Station 4176 which is located on the same shot line, has comparable misfit reduction to the stations on all the other shot lines.

8.4.2 Data fit

As already indicated by the shot-specific objective function, the data fit, as it progressed across iterations, improved substantially for all the stations. The phase-residual plots introduced in sect. 8.3.2 were utilised to track the evolution of the data fit trace by trace. In fig. 8.13 there is no indication of prominent cycle-skipped regions. Clearly, inversion proceeded in the right direction, with only a few small, isolated clusters of poorly-fit, noisy data. The most problematic seem to be the traces in the south-eastern quarter of the domain which will be further discussed when analysing the final model.

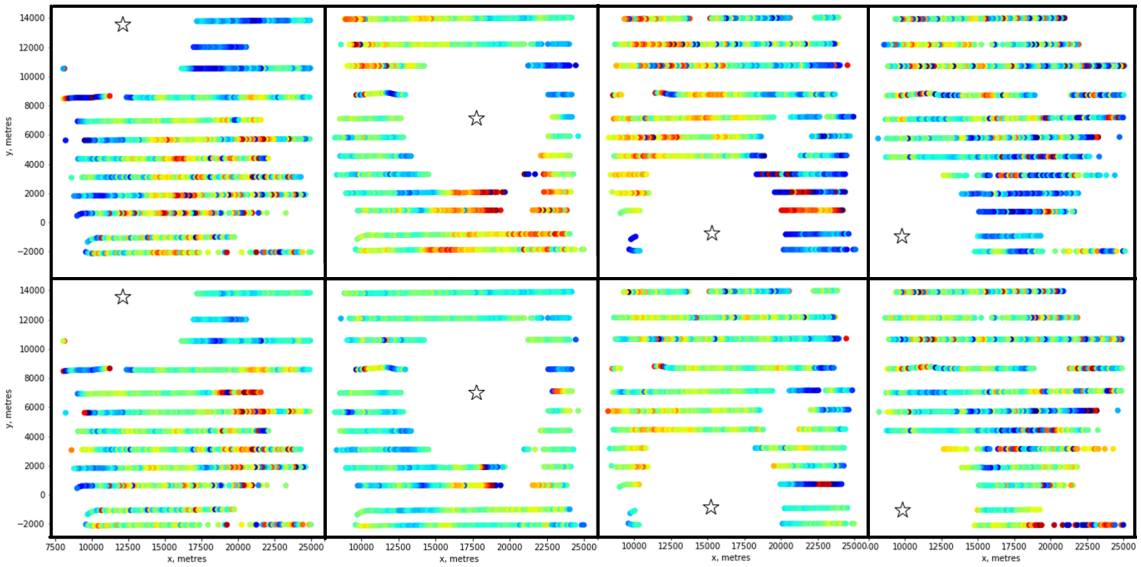


Figure 8.13: Phase residuals at 3 Hz for four stations (from left to right: 4105, 4124, 4177, 4178) at iterations 0 (top) and 80 (bottom). The white star shows the OBS location, the coloured dots show shot locations.

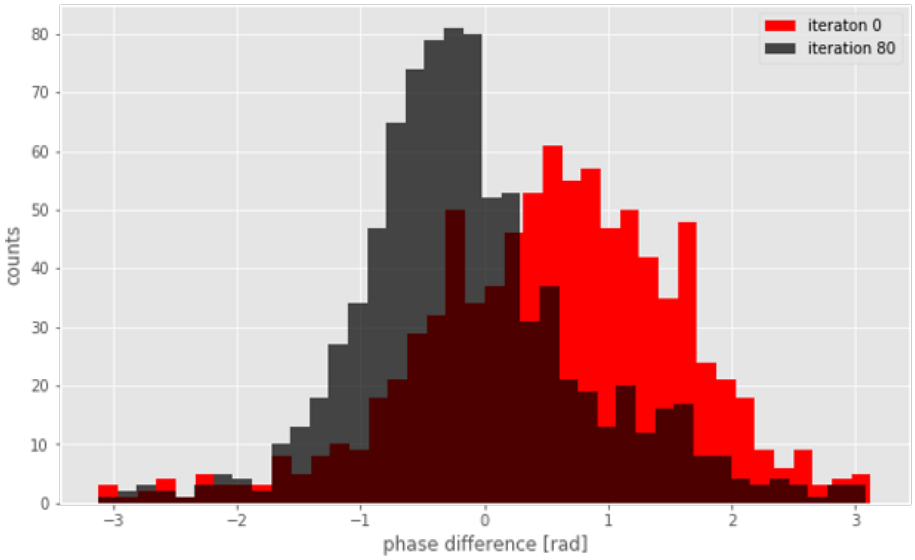


Figure 8.14: Histograms of phase residual at 3 Hz for OBS 4137 at iteration 0 (red) and 80 (black, semi-transparent).

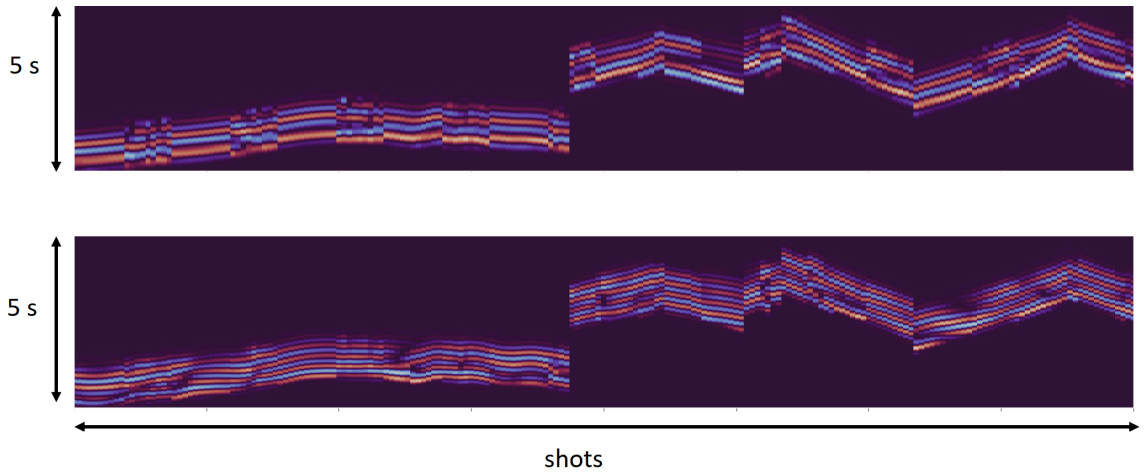


Figure 8.15: Data fit in the (left) time and (right) frequency domain for station 4137, and multiple shot-lines at the iteration 0, and 80 (top to bottom), filtered at high-cut frequency of 3 and 4.5 Hz, respectively. Each plot consists of interlaced blocks of synthetic and observed traces, 10 traces each, starting with synthetics. Data are normalised so that the maximum value of each trace is equal to 1. Each vertical arrow denotes 5 s time-window.

On average, a decrease of about 0.30 rad in RMS of phase residual at 3 Hz across 80 iterations was observed. The best improvement of 0.70 rad drop from 1.51 to 0.81 rad was for the station 4157, second best of 0.60 rad drop from 1.46 to 0.86 rad was for the station 139, and the worst, of 0.16 rad drop from 1.24 to 1.08 rad., was observed for the station 4104. For all stations, and for each iteration block, the residual decreased between the first and last iteration (fig. 8.13). The reduction in phase residual was smaller for the higher frequencies, which is consistent with the observed changes in the objective function (fig. 8.12).

The increased match to the field data between the first and last iteration is also visible in the histograms of phase residuals (fig. 8.14). It is indicated by narrowing and centering of the phase-residual distribution around zero. Additionally, the direction of the observed shift suggests that the average velocity of the model increased between the iterations. i.e. the starting model was overall too slow, at least as seen by this station.

To further investigate the data fit, juxtaposition of observed and synthetic traces, for both the starting and final velocity models, is shown in fig. 8.15. Such a way of presenting interleaved blocks of synthetic and observed traces allows to detect subtle differences of arrival times of particular phase between the two datasets. Fig. 8.15 shows that the inversion led to a significant improvement in the continuity of both first arrivals and later waveforms. A similar improvement was observed for all stations.

It is worth mentioning that both the phase plots (fig. 8.13) and interleaved traces (fig. 8.15) used the data in the same form in which contribute to the objective function. That is, after low-pass filtering, normalisation, spectrum-shaping and all the other steps our code takes to balance them optimally. This allows to determine the misfit of which traces really drives the inversion.

8.4.3 Recovered model

The most important difference between the recovered models and the starting one, already visible after the first 20 iterations at 3 Hz (fig. 8.16), is the appearance of a distinct low-velocity zone (LVZ) centred roughly at $(13.25; 8.75; 2.75)$ km. The velocity change relative to the starting model exceeds 20% in the

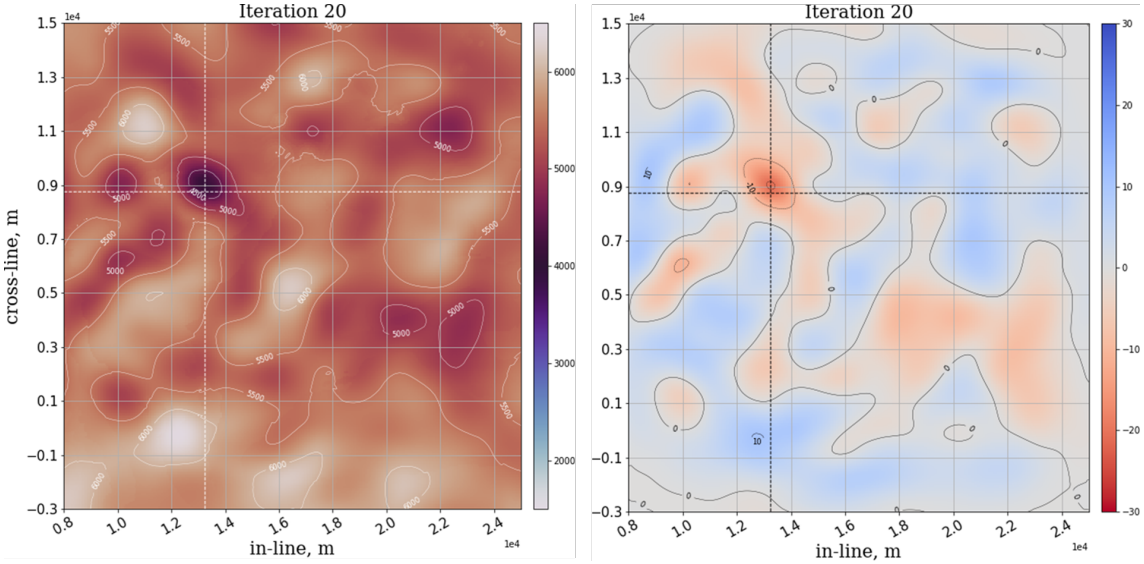


Figure 8.16: P-wave velocity model (left) and the anomaly relative to the starting model, in % (right), recovered after 20 iterations at the low-pass frequency of 3 Hz.

LVZ whereas it is less than 15% in all other regions (fig. 8.16). As the LVZ was already present (although much weaker) in the starting model (fig. 8.4), the contrast relative to a different reference model would be even more prominent, as shown in sect. 8.5.2. The fact that the velocity under Kolumbo changes quite dramatically in the first iterations suggests that this part of the starting model was the furthest from the true structure.

The final model recovered after 160 iterations (fig. 8.17) provides a substantial refinement relative to the starting one. The absolute value of anomaly beneath Kolumbo reaches 1.5 km/s which corresponds to almost 30% of velocity reduction relative to the starting model in which the medium at the same depth is characterised by velocities of around 5.0 km/s. Due to the smoothing effect of the inversion, this value is more likely to be under- than over-estimated. For the same reason, defining the spatial extent of the anomaly is arbitrary. Here the -10% isoline is taken as the approximate boundary of the LVZ under Kolumbo. It delineates a closed domain of size $2.0 \times 2.0 \times 1.5$ km, which gives a volume of 6 km^3.

Worth noting are two other, much smaller LVZs beneath Kolumbo directly above the main anomaly discussed in the previous paragraph. The spherical LVZ of 0.5 km radius centred at 1.5 km b.s.l. seems to be connected to the main one by a narrow conduit. The tiny LVZ right beneath the crater seafloor started to be picked up by the inversion in the last block of iterations at 6.5 Hz. Feasibility of detection such small features, along with their geological interpretation are discussed in sect. 8.5.

The above mentioned structures are not the only strong negative anomalies present in the final model. At later iterations another LVZ of roughly similar size appears at $(20.0; 4.00; 2.75)$ km. Unlike the anomalies beneath Kolumbo, it reaches the surface cutting through a high-velocity block present in the starting model (compare fig. 8.4 and 8.16). Its velocity is almost 1 km/s greater than the LVZ under Kolumbo (4.0-4.5 km/s vs. 3.5 km/s). The difference is even more pronounced at the earlier stages of the inversion (fig. 8.17), suggesting that these two features are of distinctly different nature. It is possible that this part of the model was cycle-skipped (fig. 8.13) and its recovered properties are not fully credible.

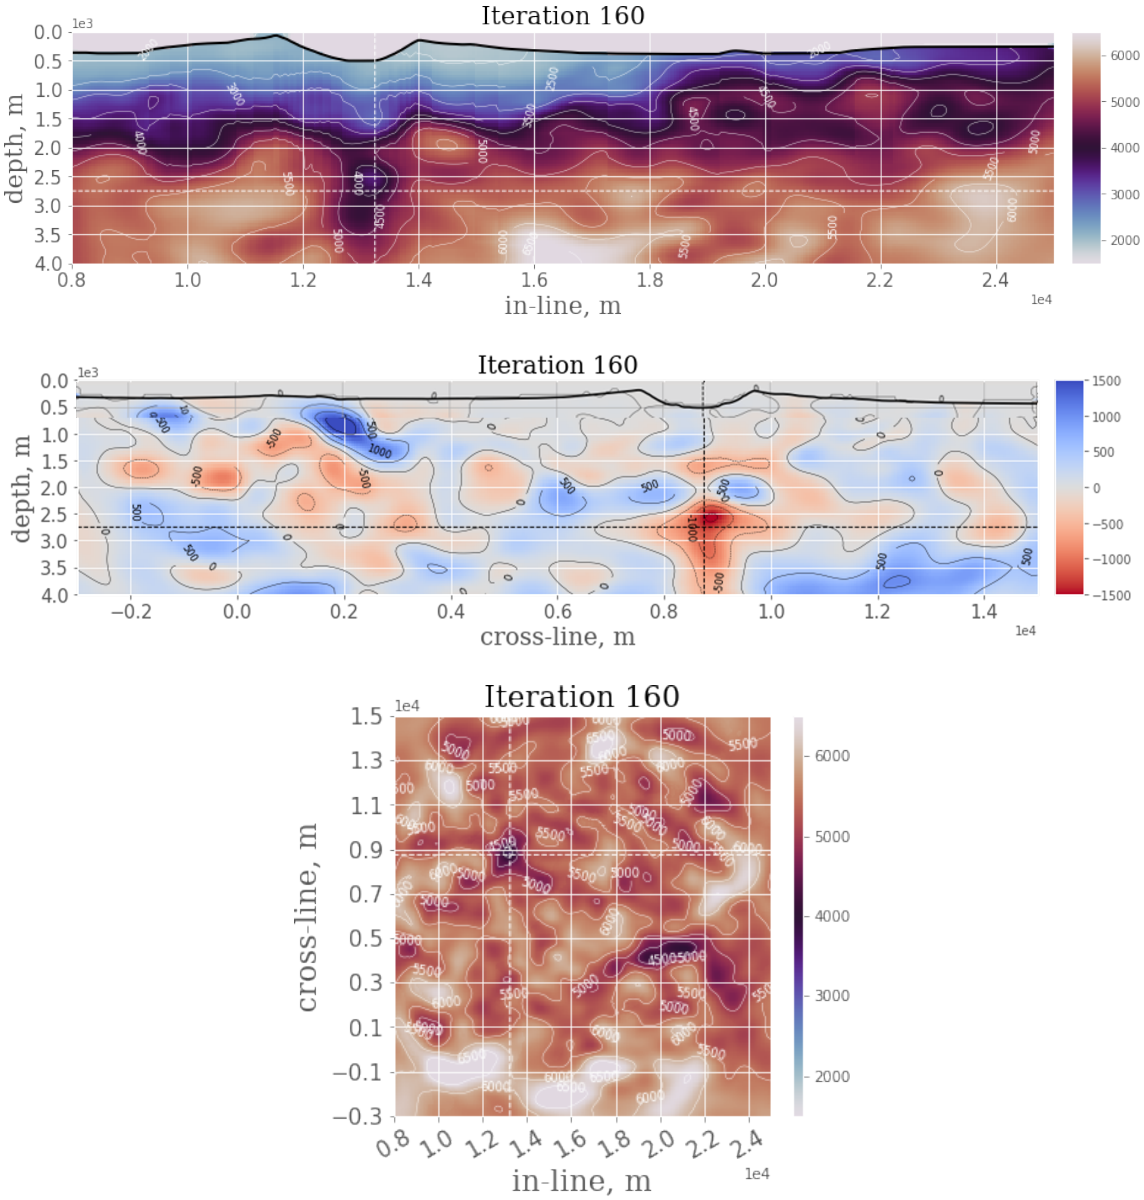


Figure 8.17: Slices through the final P-wave velocity model around Kolumbo: (top) NW-SE profile at $y = 8.75$ km; (middle) SW-NE profile at $x = 13.75$ km (shown as a perturbation of the starting model); (bottom) depth slice at $z = 2.75$ km. Dashed lines show locations of the slices.

8.5 Discussion

8.5.1 Uncertainty

In this section, the low-velocity zone beneath Kolumbo recovered in the final P-wave velocity model (fig. 8.17) is proved to be a well-resolved, robust feature based on the theoretical analysis and synthetic recovery tests.

In sect. 8.4 it was shown that the inversion substantially improved the match between the synthetic and observed data for all the stations and at all frequencies (fig. 8.12-8.15). This indicates that the inversion converged to some minimum of the objective function but it does not say much, even though the first model updates look geologically credible, about plausibility of the recovered structures. The latter is typically assessed using independent constraints – wireline logs (Warner et al., 2013a) and seismic reflection images (Davy et al., 2018). However, there are no boreholes or sufficiently deep reflection sections available in the area of study. The only reflection image of the target region published so far (Hübscher et al., 2015) reaches the depth of only 1 km (fig. 8.3). In this situation, the plausibility of the model was analysed in terms of its resolution, validity of the modelling assumptions and the sensitivity of the recovered structure to different inversion parameters. It was also noted that neither the lowest, nor the highest rock-velocity of the starting model (2 km/s and 6.5 km/s, respectively) were exceeded in the final model. As these values are geologically plausible, it is another indicator of the fidelity of the recovered model.

Modelling assumptions Here, a validity of the modelling assumptions used in the inversion (acoustic, isotropic, lossless medium) and their expected effect on the recovered structure is dicussed. The acoustic approximation seems reasonable as the shots were fired in water and most energy generated was compressional waves. The double-converted waves (P-S-P etc.) were excluded from the inversion by muting the data. Other elastic effects were suppressed, at least partially, by data normalisation (Warner et al., 2012, 2013b). Significant azimuthal (HTI) anisotropy of the P-wave velocity in the study area was reported by Heath (2019). Clearly, the assumption of isotropy is only a crude approximation, however, the effects of neglecting anisotropy are known to mainly distort the image by mispositioning anomalies, but not to severely affect their amplitudes (**?**). That is why from the point of view of this study (estimation of melt fraction), taking anisotropy into account is not critical. Accounting for attenuation proved to improve P-wave velocity models in terms of resolution and fidelity, especially in the deeper parts of the recovered models (Calderón Agudo et al., 2018). However, it is not a first-order factor which would drive a transmission FWI (Warner et al., 2012). Furthermore, with anomalies that small, their attenuative effect is not expected to be dominant. It is likely to be larger for the data crossing the northern basin of the Santorini caldera (McVey et al., 2020). The above analysis shows that the presence of anomalies is unlikely to be an artifact due to the modelling assumptions.

Resolution analysis As shown in sect. 8.3.3, the resolution of the recovered model is controlled mainly by the data acquisition geometry, length of the time-window and the local wavelength and falls somewhere between the travel-time tomography and the best theoretical resolution of FWI. As the latter is not reached, the model updates were smoothed over two local wavelengths determined from the best-trade-off criterion (Hansen, 1992). The average local velocity at the depth of the anomaly is around 5 km/s (fig. 8.17). This gives a local wavelength of $\lambda_{3.0\text{Hz}} = \frac{5.0}{3.0}\frac{\text{km/s}}{1/\text{s}} \approx 1.7$ km, $\lambda_{6.5\text{Hz}} = \frac{5.0}{6.5}\frac{\text{km/s}}{1/\text{s}} \approx 0.77$ km in the first and last iteration block with frequencies of 3.0 and 6.5 Hz, re-

spectively. The best theoretical resolution of the final model would be $\frac{1}{2}\lambda_{6.5\text{Hz}} \approx 0.38$ km (Virieux and Operto, 2009) but this is likely to be beyond the reach of this dataset, given its sparsity. The gradient was smoothed over a distance of 2λ (local wavelengths) horizontally and 1λ vertically. Consequently, features of the final model at 2.75 km b.s.l. of a size 1.54 km horizontally and 0.77 km vertically may be taken the smallest features recoverable with good fidelity. Smaller features would be smoothed – their spatial extent increased and amplitude reduced. The LVZ of the final model was also present in the starting model from travel-time tomography but it was much more subtle (Heath et al., 2019; McVey et al., 2020). The first breaks were picked from the data at 5-25 Hz frequency range. The dominant frequency was probably somewhere in the middle, but let us take 25 Hz to get an upper bound limit for the resolution. The smallest feature which can be reasonably accurately reconstructed using ray traveltime tomography, is given by a radius of the first Fresnel zone (Sheriff and Geldart, 1995):

$$r_{\text{min}} = \frac{1}{2}\sqrt{\pi L\lambda} \approx \sqrt{L\lambda},$$

where L is the propagation distance and λ is the dominant wavelength. Taking $f = 25$ Hz and the average velocity $\langle v \rangle = 5$ km/s, we get $\lambda_{\text{min}} = \langle v \rangle / f_{\text{max}} = \frac{5}{25} = 0.2$ km. Offset of at least 10 km is needed to penetrate down to 2.7-3.5 km depth where LVZ is. Here the distance along-the-ray is assumed to be equal to the offset, which is an underestimation. This gives $r_{\text{min}} = \sqrt{L\lambda} = \sqrt{10 \cdot 0.2} \approx 1.4$ km. Halving the dominant frequency or doubling the distance would give a value of 2 km which is still smaller than the anomaly size. This suggests that the LVZ in the starting model is relatively subtle due to smoothing applied during the travel-time tomographic inversion (not analysed here), and/or that the frequency was effectively lower than used in the simplistic resolution calculation above.

Robustness of recovery Another possibility to investigate was whether the recovered anomalies are not artifacts due to the data noise. To this end, different inversion strategies were tested using input data with different offset ranges, time windows, and starting frequency, all of which aimed to decrease the influence of noise in the inversion. The LVZ under Kolumbo proved to be a robust feature that appeared in all recovered velocity models, unlike the LVZ in the metamorphic block which thus will not be further interpreted. For this reason, the LVZ under Kolumbo seems to be a real geological feature.

Fidelity of recovery The recovered model is a smoothed version of the true v_{P} structure. To investigate the smoothing effect of the inversion, several spike tests (Rawlinson and Spakman, 2016) were performed using the same strategy and parameters as in the field-data inversion. The rationale was to find a synthetic model that leads to recovery of similar anomaly to the observed one. This problem is non-unique as there is a trade-off between the size and amplitude of the anomaly. For this, reason multiple 3d Gaussian anomalies of spherical symmetry and different maximum amplitudes (1%, 5%, 10%, 50% of the starting model) and sizes (full width at half-maximum between 0.2-2.0 km) and were tested. Fig. 8.18 shows that anomaly as small as 200 m (FWHM) can be recovered provided it is strong enough (here -50%). The recovered anomaly is smeared with the underestimated maximum amplitude. Results of similar tests are supplied in appendix 8.B.

8.5.2 Melt fraction

In this section, the effective-medium theory is used to show that the low-velocity (down to 3.5 km/s) anomaly observed around 3 km b.s.l. at Kolumbo is consistent with a presence of partially molten granitic intrusion of 20-45% porosity with melt-filled penny-shaped voids.

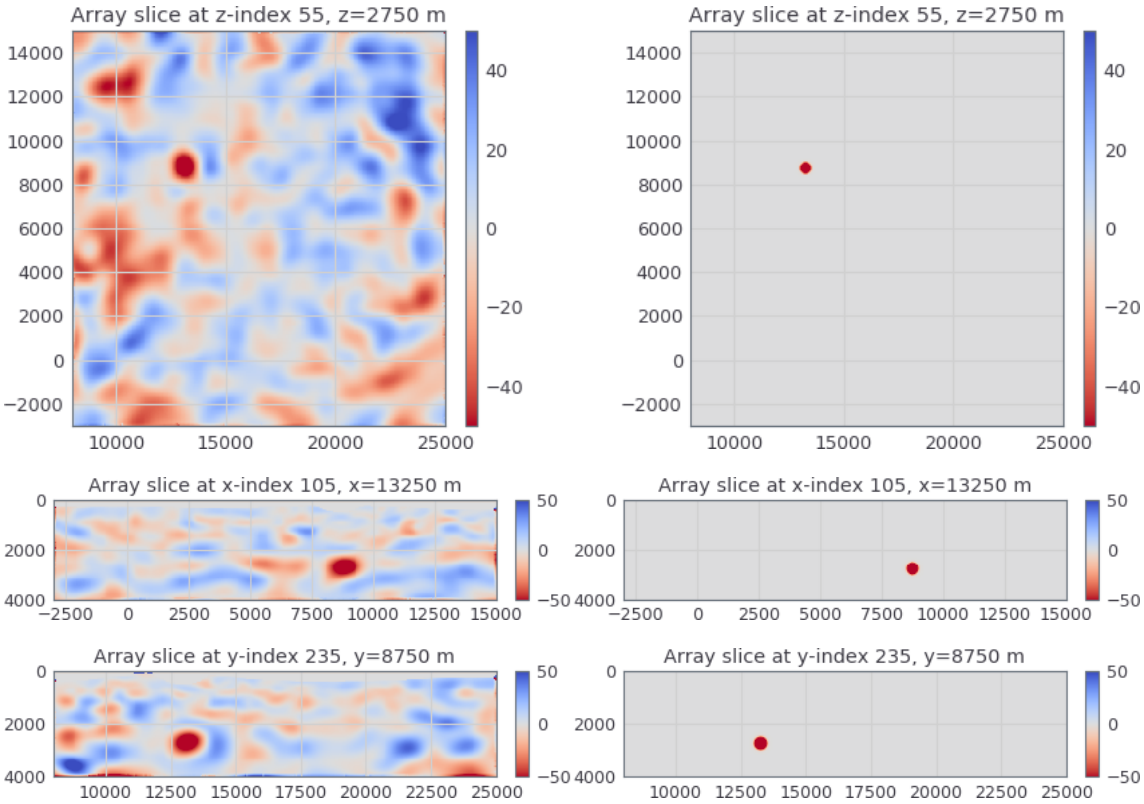


Figure 8.18: Model recovered after inverting synthetic data generated from the starting model perturbed by a spherical LVZ added underneath Kolumbo. An anomaly as small as 200 m (FWHM) can be recovered, provided it is strong enough (here -50%). The final anomaly is smeared with a slightly under-recovered amplitude.

To interpret a tomographic model of seismic velocity in terms of melt-content one needs to understand elastic properties of partially molten rocks. Such rocks are multi-phase (composite) media, consisting of solid matrix, and pores filled with fluids such as melt and volatiles (primarily supercritical H_2O and CO_2). Seismic velocity is a function of density and any two of the elastic moduli[7] (Landau and Lifshitz, 1970), for example:

$$v_P = \sqrt{\frac{K + \frac{4}{3}\mu}{\rho}} \tag{8.4}$$

$$v_S = \sqrt{\frac{\mu}{\rho}} \tag{8.5}$$

and conversely:

$$K = \rho\left(v_P^2 - \frac{4}{3}v_S^2\right) \tag{8.6}$$

$$\mu = \rho v_S^2 \tag{8.7}$$

where K and μ[8] are the bulk and shear modulus, respectively. The elastic moduli depend primarily on:

- pressure – usually positively correlated with v_P,

[7] They all have units of stress as they are ratios of stress and strain (the latter is dimensionless).
[8] Denoted by G by some authors.

SiO_2	Al_2O_3	Na_2O	K_2O	FeO[10]	CaO	MgO	MnO	TiO_2	P_2O_5	*sum*	H_2O
74.34	13.75	4.77	4.35	1.43	0.86	0.24	0.11	0.09	0.02	99.96	6.50

Table 8.2: Major-element composition of the K5 pumice (fig. 8.3), given on a volatile-free basis and normalised to 100 wt.%. The sum of major elements without water is 99.96 wt.%. The values are the average of NA007-058 NA007-071, NA007-073 NA007-075, NA007-085, NA007-090, NA014-019 samples reported in the supplement of Klaver et al. (2016).

- temperature – usually negatively correlated with v_P,
- porosity – usually negatively correlated with v_P,
- composition of solid matrix,
- composition of pore fluids,
- geometric arrangement (shapes and connectivity) of the pores.

Melting the rock changes its elastic moduli and thus the seismic velocity. In general, P-wave velocity increases with confined pressure and decreases with temperature – the *normal geotherm* will give no net change of v_P with depth (Christensen and Stanley, 2003).

Here, the low-velocity zone will be associated with a partially molten intrusion (alternatives are much less likely as discussed in sect. 8.5.3). Assuming absence of volatiles[9] and compositional homogeneity, the intrusion constitutes a two-phase system consisting of the solid-rock matrix and the melt entirely filling the pore spaces. Note that under these assumptions, porosity is equivalent to melt-fraction.

Composition The intrusion will be assumed to be a biotite-bearing silica-rich *granite*, based on the analysis of volcanic products from the 1650 AD eruption (Cantner et al., 2014). The average major-element whole-rock composition was calculated from seven *in situ* samples of 250-m-thick 1650 AD pumice layer (K5 in fig. 8.3) analysed by Klaver et al. (2016)[11]. The water content was taken as the middle value 6.5 wt.% of the 6-7 wt.% range determined by Cantner et al. (2014) for different possible pre-eruption pressures and temperatures. These numerical data, gathered in tab. 8.2, were later used to derive properties of melt (see below).

p-T conditions Elastic properties of the solid and liquid phases depend not only on their composition but also on pressure and temperature conditions. The lithosthatic pressure p_{lith} at the depth of 3 km (strongest v_P contrast) was calculated to be 66 MPa using:

$$p_{\mathrm{lith}}(z) = p_0 + g \int_0^z \rho(\xi) \mathrm{d}\xi \tag{8.8}$$

where $g = 10\,\mathrm{m/s^2}$ and $p_0 = p_{\mathrm{atm}} = 0.1$ MPa. The density ρ was calculated from a 1d depth v_P profile averaged over a 10x10 km box centred at Kolumbo using the empirical Nafe-Drake relation quantified by Brocher (2005) and a value of 1 g/cm3 in the water column. The temperature of solidus T_{solidus} of a hydrous granite at p_{lith} was assumed to be 700 °C[12] from Boettcher and Wyllie (1968).

Pore geometry The best bounds for an isotropic linear elastic composite, defined as giving the narrowest possible range without specifying anything about the geometries of the constituents, are the

[9] Effect of volatiles will be discussed qualitatively. Quantitative analysis is beyond the scope of this study.
[10] Total ferrous iron is given instead of separate FeO and Fe_2O_3 oxides.
[11] This dataset largely overlaps with samples from Cantner et al. (2014), see this paper for a thorough description of the samples.
[12] At low pressures like this, the melting curve is negative trending (Boettcher and Wyllie, 1968).

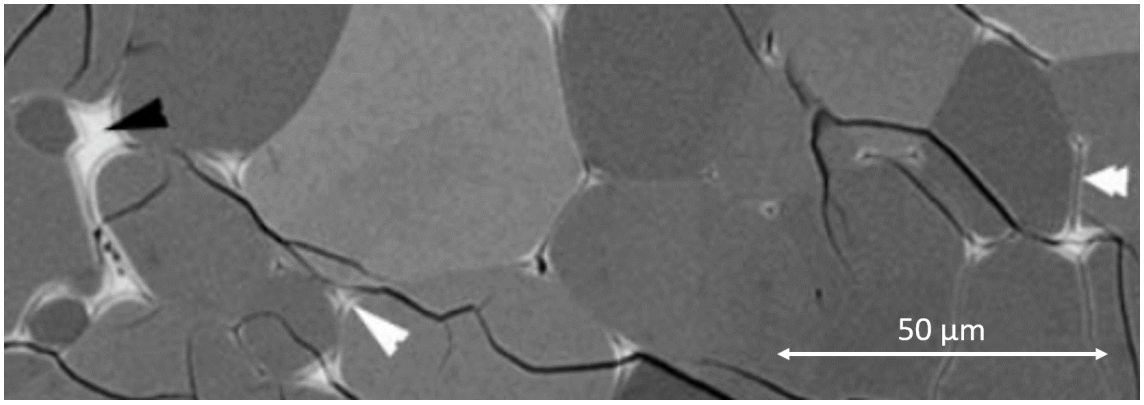


Figure 8.19: Partially molten olivine grain boundaries in an olivine-basalt aggregate. The light grey grains are basalt glass, the darker are olivines. Melt resides primarily in large pockets (black arrow), triple-junctions (white arrow) but also thin layers between grains (double white arrow), the latter two are poorly resolved in this back-scattered-electrons (BSE) image. Black cracks are a result of quenching. From Cmiral et al. (1998).

Hashin–Shtrikman bounds (Hashin and Shtrikman, 1963). These bounds are too wide (fig. 8.21) for a meaningful discussion and certain assumptions about pore geometry must be made. The method used to calculate the effective properties of the two-phase system (see below) assumes a *spheroidal* shape of pores. A spheroid is a special case of an ellipsoid with two equal semi-diameters. It can be constructed by rotating an ellipse about one of its principal axes. For spheroidal inclusions, one can define the *aspect ratio*:

$$\alpha = \frac{r_3}{r_1}. \tag{8.9}$$

Spheroids with $\alpha > 1$ are called *prolate* (rugby-ball shape) and with $\alpha < 1$ *oblate* (penny shape). The latter type is expected for melt pockets/films/layers at the grain boundaries (fig. 8.19). The exact aspect-ratio of oblate spheroids for a specific rock like granite is poorly constrained. Here, the values between 0.05 and 0.5 were considered plausible as discussed by Paulatto et al. (2019).

Properties of solid The properties of the solid matrix were taken as:

$$v_P^{\text{lab}} = 6.25 \pm 0.13\,\text{km/s}$$
$$v_S^{\text{lab}} = 3.67 \pm 0.12\,\text{km/s}$$
$$\rho_{\text{lab}} = 2.66 \pm 0.02\,\text{g/cm3}$$

from the laboratory measurements of more than 100 granite samples at $p_{\text{lab}} = 200$ MPa and $T_{\text{lab}} = 30$ °C (Christensen and Stanley, 2003)[13]. The P-wave velocity of the solid matrix v_P^{solid} was taken as a value v_P^{lab} corrected to correspond to pressure and temperature at the depth of anomaly:

$$v_P^{\text{solid}} = v_P^{\text{lab}} + \left(\frac{\partial v_P}{\partial p}\right)_T \Delta p + \left(\frac{\partial v_P}{\partial T}\right)_p \Delta T \tag{8.10}$$

where:

$$\Delta p = p_{\text{lith}} - p_{\text{atm}} \approx 66\,\text{MPa}$$

and:

$$\Delta T = T_{\text{solidus}} - T_{\text{lab}} \approx 670\,^\circ\text{C}$$

[13]The authors do not report the temperature of rocks explicitly but it can be concluded from the text.

The sensitivity of P-wave velocity to pressure, $(\partial v_P/\partial p)_T$ was assumed to be $2 \cdot 10^{-4}$ km/s/MPa after Rudnick and Fountain (1995) which gives a difference of about 0.01 km/s for the above value of Δp. Influence of pressure variations on v_P can thus be neglected as a second-order factor. The P-wave velocity sensitivity to temperature $(\partial v_P/\partial T)_p$ was taken to be -0.39×10^{-3} km/s/°C after Christensen (1979). This derivative is only valid for temperatures up to 300 °C and a more sophisticated calculation (Paulatto et al., 2019) based on the model of Karato (1993) could be used for higher T. However, fig. 8.31 shows that the potential difference may be expected to be negligible compared to the effect of other unknowns such as the pore aspect-ratio.

Properties of liquid The model of Ueki and Iwamori (2016) was used to calculate the bulk modulus and density of the silicate melt of the composition from tab. 8.2. This model, simpler than but consistent with the MELTS software (e.g. Asimow and Ghiorso, 1998) widely used in the petrological community, is based on the *Birch–Murnaghan* isothermal *equation* of state:

$$p = \frac{3}{2} K \left(\delta V^{-7/3} - \delta V^{-5/3}\right) \left[1 + \frac{1}{4}\left(3K' - 4\right)\left(\delta V^{-2/3} - 1\right)\right] \tag{8.11}$$

where p is the pressure, K is the isothermal bulk modulus at 1 bar and $K' \coloneqq \partial K/\partial p$ (both obtained from fits to experimental data) and:

$$\delta V \coloneqq \frac{V}{V_0}, \tag{8.12}$$

where V is a partial molar volume under the given p-T conditions and V_0 is a partial molar volume at the reference pressure p_{ref}:

$$V_0 \coloneqq V(p_{\mathrm{ref}}, T_{\mathrm{ref}}) + \left(\frac{\partial V}{\partial T}\right)_p (T - T_{\mathrm{ref}}). \tag{8.13}$$

The value of K is calculated from:

$$K = -\frac{V_0}{(\partial V/\partial p)_T + (\frac{\partial^2 V}{\partial T \partial p})_T (T - T_{\mathrm{ref}})}. \tag{8.14}$$

The unknown parameters are optimised based on the experimental data. The final bulk modulus is obtained from K by linear extrapolation. Eq. (8.11) was originally developed for cubic crystals (Birch, 1947) but since then successfully applied to various silicate melts (see Ueki and Iwamori (2016) for references).

For values from tab. 8.2, $p = 0.06$ GPa and $T = 700$ °C, the obtained values were $\rho = 2.114$ g/cm3 and $K = 8.080$ GPa. The subsequent analysis, with one parameter varied and other fixed at the above values, showed that the results are primarily sensitive to the water content, giving $\rho = 2.385$ g/cm3, and $K = 11.986$ GPa for 0 wt.% water, and $\rho = 2.060$ g/cm3, and $K = 7.574$ GPa for 8 wt.% water. Other factors seem to be of less importance, e.g. increasing pressure $p = 0.5$ GPa (order of magnitude larger) gives $\rho = 2.221$ g/cm3, and $K = 9.796$ GPa. Also, varying the temperature ± 100 °C has a negligible effect, yielding $\rho = 2.126$ g/cm3, and $K = 7.864$ GPa for $T = 600$ °C, and $\rho = 2.102$ g/cm3, and $K = 8.307$ GPa for $T = 800$ °C. Decreasing the silica content SiO_2 to 60 wt.% results in $\rho = 2.132$ g/cm3, and $K = 8.224$ GPa.

Properties of composite The 'self-consistent' scheme of Berryman (1980) was used to calculate (fig. 8.20) the effective properties of the two-phase medium. The P-wave velocity predicted for a given melt fraction strongly depends on the assumed pore aspect-ratio (fig. 8.21). If one takes 0.05-0.5 as a range of plausible values for the effective aspect ratio influencing the seismic velocity (Paulatto et al., 2019), this gives 20-45% of melt fraction at $v_P = 3.5$ km/s.

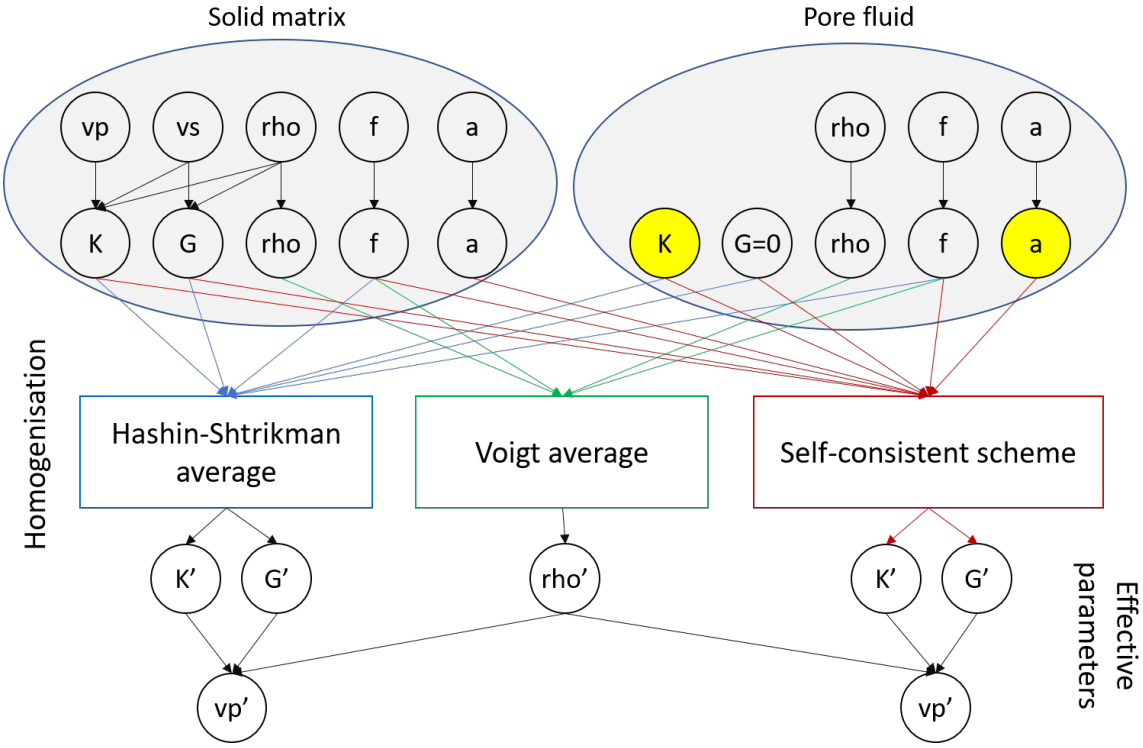


Figure 8.20: Flow of calculation of the effective P-wave velocity using the self-consistent scheme of Berryman (1980) and the Hashin-Shtrikman average (Hashin and Shtrikman, 1963). The input variables are in grey ellipses. The Voigt average is an arithmetic mean weighted by volume fractions (Mavko et al., 2009). The effective velocity is by far the most sensitive to pore-fluid bulk modulus and aspect-ratio highlighted in yellow.

Effect of volatiles Cantner et al. (2014) reported a volatile content to be 6.5 wt.% (median value[14]) based on the analysis of plagioclase-hosted melt inclusion. The presence of volatiles reduces the elastic moduli of melt and hence of the composite medium. For this reason, ignoring the volatiles leads to *overestimation* of the melt content (Paulatto et al., 2019). On the other hand, v_P observed in the LVZ of the final model was likely too high due to the smoothing effects of sparse-data FWI leading to the *underestimation* of the melt content. Quantifying the relative weight of these two factors will be a topic of future research. The presence of volatiles in tomographic images may be indicated by low v_P/v_S ratio of the effective medium because of the much lower bulk modulus (and hence v_P) of gas.

8.5.3 Interpretation

In this section, the v_P anomaly under Kolumbo is interpreted as evidence for a small but potentially eruptible melt reservoir stored as magma rather than mush.

Melt vs. water Similar calculation led Hooft et al. (2019) to a conclusion that the LVZ under the Santorini caldera is associated with an excessive porosity filled with seawater. In that case, presence of water seems more likely than melt because:

- the Santorini anomaly is much weaker and more shallow (stretching from the seafloor, with the maximum anomaly at 1.5 km),

[14] Authors remark that this value may be under-estimated.

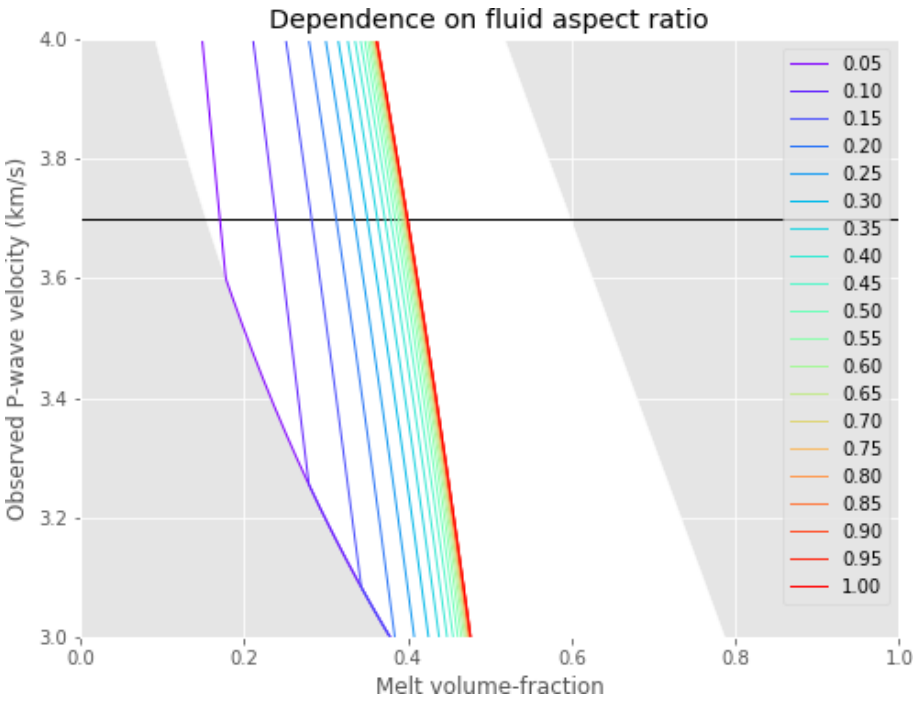


Figure 8.21: P-wave velocity of a two-phase composite calculated with the self-consistent approximation (Berryman, 1980) in the low-frequency, equilibriated, limit starting with air-filled pores and subsequent fluid substitution according to Gassman's equation (Mavko et al., 2009). White area denotes Hashin-Shtrikman bounds (Hashin and Shtrikman, 1963). coloured lines correspond to different values of the pore aspect-ratio. Sensitivity to other parameters can be found in appendix 8.C.

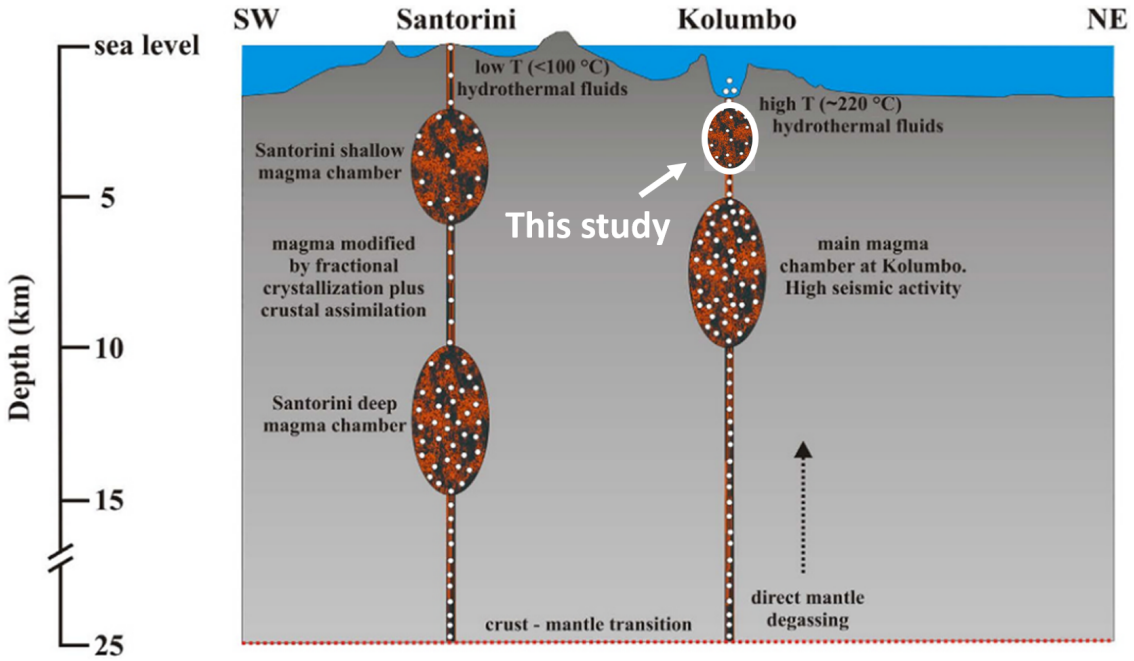


Figure 8.22: Schematic model of Santorini and Kolumbo magmatic system(s). Not to scale. Modified after Rizzo et al. (2016).

- the collapse of the Santorini caldera must have resulted in a large volume of subsided porous infill, likely to reach the depth of anomaly,
- the depth of the melt reservoir at 4-6 km beneath Santorini is relatively well constrained (see e.g. fig. 3b of Parks et al. 2012) and the presence of another, more shallow reservoir is not required by the satellite data to explain the 2011-12 inflation (Parks et al., 2015),
- the hydrothermal vents inside the Santorini caldera have low temperature.

The water-porosity seems to be much *less likely* at Kolumbo which hosts numerous high-temperature hydrothermal vents and where heat and gas fluxes are much higher (Sigurdsson et al., 2006; Rizzo et al., 2016). The caldera is much smaller, and the anomaly starts at 2.5 km b.s.l. stretching to larger depths. The velocity reduction is much larger. Furthermore, the deeper melt reservoir is not just as well constrained, and its location may be less certain. The local-earthquake tomography model (Dimitriadis et al., 2010), which forms the basis of all models of the Kolumbo magmatic systems published to date (e.g. Konstantinou, 2020), is not well resolved around Kolumbo (see fig. 8 of Dimitriadis et al. (2010), including the caption). In the cross-section 7 (fig. 13b and 15 of Dimitriadis et al., 2010), which is usually interpreted by petrologists (e.g. Cadoux et al., 2014; Cantner et al., 2014), much stronger anomalies are present between Santorini and Kolumbo, rather than beneath the latter. This is well visible in the depth sections and the cross-section 5, where the localised anomaly is offset from Kolumbo towards Santorini.

Magma vs. mush The critical porosity (Mavko et al., 2009), or the melt fraction at which the mush-magma transition occurs (resulting in a step-change increase in melt-mobility) can be expected to be in 30-50% range (Sparks et al., 2019). For example, it was estimated to be around 35% at Montserrat (Paulatto et al., 2019). Even the lower bound of the calculated melt fraction 20-45%, may be close to the critical value. For this reason, the reservoir is interpreted as a shallow *magma chamber* rather than a mush zone. The small depth enables it to supply the Kolumbo high-temperature hydrothermal system with sufficient heat (fig. 8.22). One of possible explanations of the presence of melt in the shallow reservoir only 400 yr after the last major eruption is that the shallow felsic mush zone has been heated by influxes of more mafic melt from the deeper reservoir reported by Dimitriadis et al. (2010). These recharging events may be partly responsible for the ongoing high microseismic activity around Kolumbo.

8.6 Conclusion

8.6.1 Summary

Acoustic isotropic 3d full-waveform inversion of early-arriving seismic waves recorded by 16 ocean-bottom hydrophones was used to constrain the topmost 4 km of the subsurface below Kolumbo volcano, the most seismically-active center of the Santorini volcanic field.

The inversion revealed a previously undetected, sharp low-velocity zone of several km^3 in volume, extending from 2 km down to the bottom of the model at 4 km depth. The minimum value in the LVZ was 3.5 km/s corresponding to -1.5 km/s difference relative to the starting model which already contained the smoothed version of this LVZ.

The effective-medium calculation showed that the anomalous velocity structure is consistent with a partially molten granitic intrusion of at least 20-45% melt fraction. This was interpreted as a shallow high-fraction melt reservoir that is ready to feed future eruptions.

8.6.2 Next steps

The results of this study will be compared with newly acquired high-quality multi-channel seismic reflection data, as soon as they become available (Karstens et al., 2020). Other steps recommended for future research are listed below.

More data To further constrain the structure under Kolumbo, more data should be inverted. Increasing the amount of data can be done in a few steps:

1. increasing the model size to include longer offsets – this should better constrain the low-wavenumber features of the model, although the amount of data with acceptable S/N ratio is limited,
2. adding a vertical component of particle velocity in order to either invert it jointly with pressure, or stack the final models – this approach may reduce the number of artifacts in the final model (ch. 5),
3. widening the time window – iterating over time-window-width as an innermost loop of the inversion workflow, so called *layer-stripping* strategy (e.g. Górszczyk et al., 2017),
4. inverting higher frequencies – it will need reprocessing of the source wavelet and field data (using higher low-pass frequency) as well as the model (upsampling to $\mathrm{d}x = 25$ m or less).

Less approximation After exploiting the acoustic isotropic approximation to the maximum, some of the modelling assumptions may be relaxed by:

1. anisotropic forward modelling[15] – using the HTI model of Heath (2019) which is soon to be published in a peer-reviewed journal,
2. (visco)elastic forward modelling – to investigate sensitivity of data to (un)elastic effects associated mainly with the presence of partial melt and volatiles,
3. inverting for attenuation – it may be done using an efficient, semi-global scheme (da Silva et al., 2019) based on a long-wavenumber approximation of Q which should be enough to address most of the viscoacoustic effects (Calderón Agudo et al., 2018),
4. inverting unnormalised seismograms – tighter constraints on location, shape and absolute velocity of the anomalies will require including true amplitudes of seismograms. To this end, starting models of anisotropy, attenuation and perhaps v_S will be necessary (Warner et al., 2012).

Acknowledgements

The author wants to thank M. Paulatto, J. Morgan and E. Hooft for numerous pieces of useful advice, enthusiasm and constant encouragement to complete this study.

Supplementary information

A supplementary `Jupyter` notebook is available at: `http://github.com/kmch/PhD-Santorini-FWI/notebooks/ch_kol` (fig. 8.23). The structure of the notebook follows the one of the chapter and contains commands run to produce many of its figures. It can be launched on `http://mybinder.org` by clicking the button at the top of the `README.md` file of the repository. `Binder` allows to interactively run the commands and reproduce the figures[16] without a need to have any software installed on a local machine. The `ch_kol` directory, apart from the notebook itself, contains small-size files of open-access

[15] Full-waveform inversion for anisotropy would be a much more challenging task
[16] Provided the libraries are correctly set up.

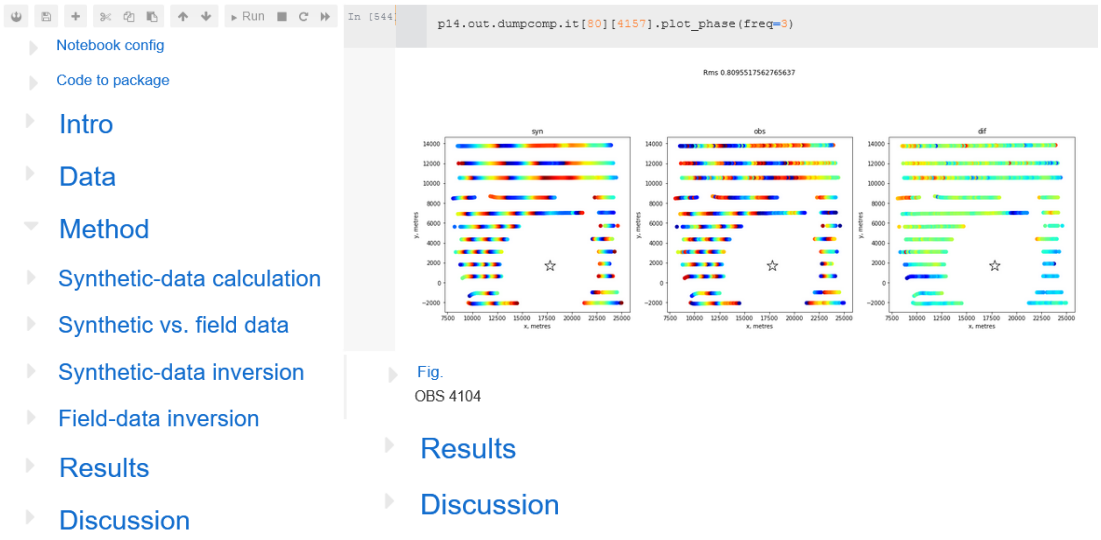


Figure 8.23: Screenshots of a supplementary `Jupyter` notebook `Kolumbo_volcano.ipynb`.

data needed to reproduce some of the figures. Medium-size files can be scraped from the data server (see below) e.g. by calling a wget command from the `postBuild` file created in the `GitHub` repository (see `http://github.com/binder-examples/getting-data` for detailed instructions).

Data availability

The entire OBS dataset analysed in this study has been archived at the IRIS Data Management Center (`http://www.iris.edu`). Search for 'PROTEUS experiment' and/or 'Emilie E. Hooft'.

Code availability

The FWI-diagnostics `FullwavePy` software used to QC the inversion and produce most of the figures (see the supplement) can be downloaded from `http://github.com/kmch/FullwavePy`. The `FULLWAVE3D` software used to invert the data can be accessed through academic collaboration with Imperial College London or a membership of the FULLWAVE consortium (Warner et al., 2013b). The effective-medium properties were calculated using the `ElasticC` package by M. Paulatto, available at `http://github.com/michpaulatto/ElasticC`.

Appendix

8.A Final inversion results

Additional plots showing the results of the final inversion are provided in fig. 8.24-8.26. Slices are taken along the dotted white lines visible in each plot.

8.B Anomaly-recovery tests

In this appendix, results of additional recoverability tests mentioned in sect. 8.5 are provided. The same inversion parameters were used as for the final model from the field data. Anomaly with a peak amplitude of -1% is not recovered (fig. 8.27), of -5% is quite well recovered both in terms of shape and amplitude (fig. 8.28), and of -50% leads to the cycle-skipping of the data (fig. 8.29).

8.C Melt sensitivity analysis

Here, sensitivity analysis for the melt-fraction estimation (see sect. 8.5.2) is provided. For a given v_P (black line), melt-content estimation depends primarily on the assumed aspect ratio of melt-filled spaces and bulk modulus of melt. The latter is well constrained given the known composition of the volcanic products from the last eruption. Even with completely wrong water content, the difference in K would not affect the melt-content. The former, however, is less well constrained, although values between 0.05-0.2 are recommended by Paulatto et al. (2019).

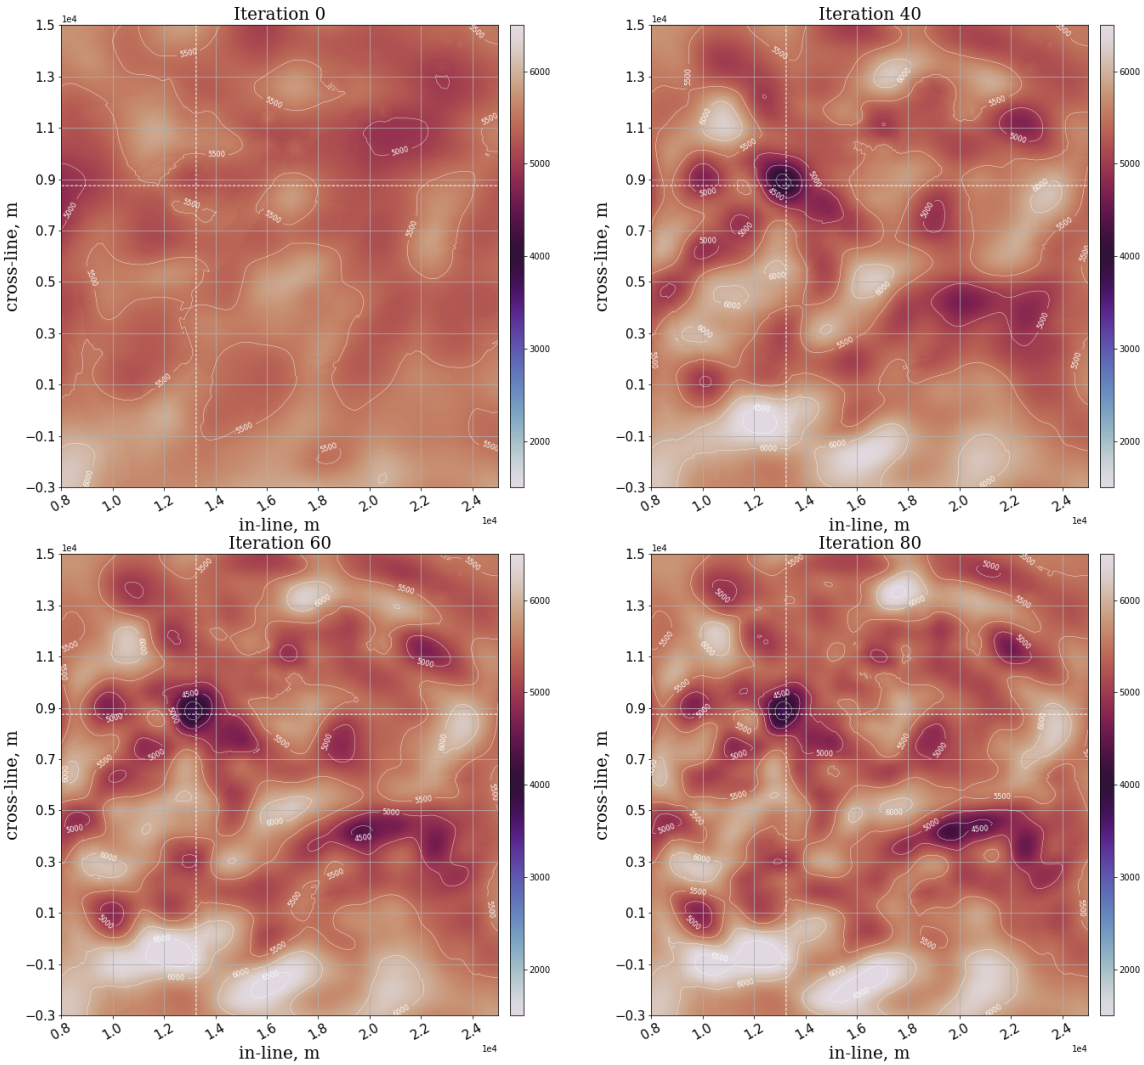


Figure 8.24: Slices through the final model at $z = 2.75$ km b.s.l. The slice after 20 iterations is shown in fig. 8.16.

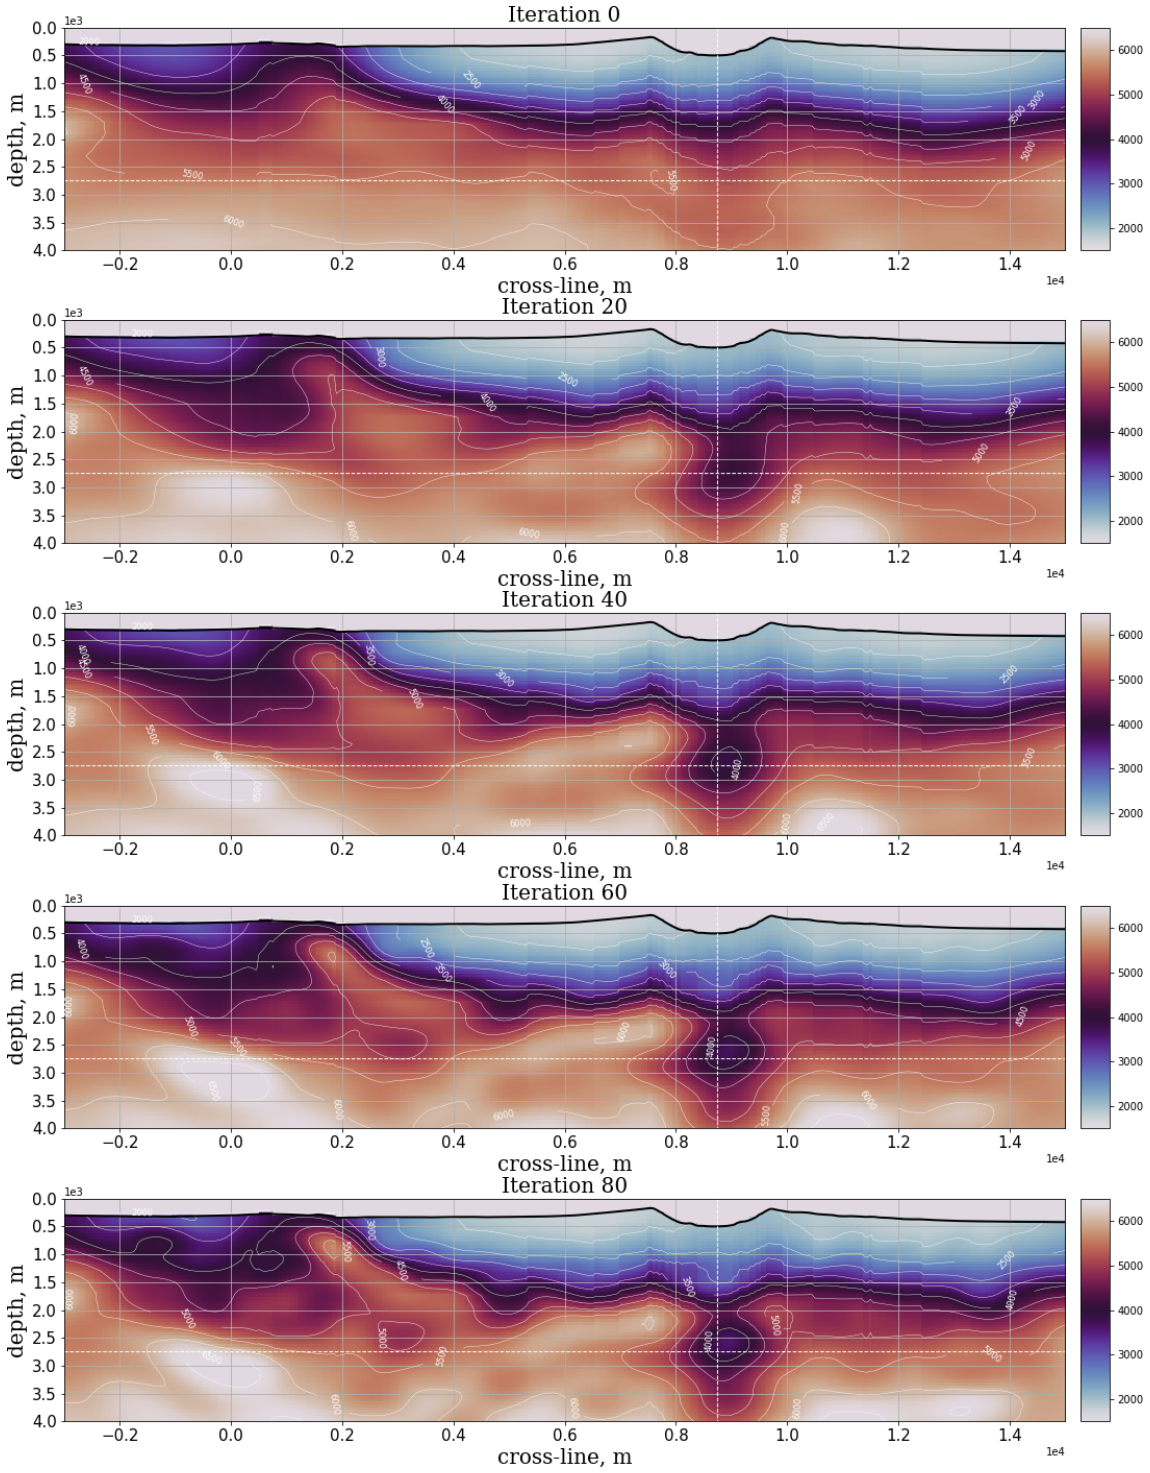


Figure 8.25: Slices through the final model at $x = 13.25$ km.

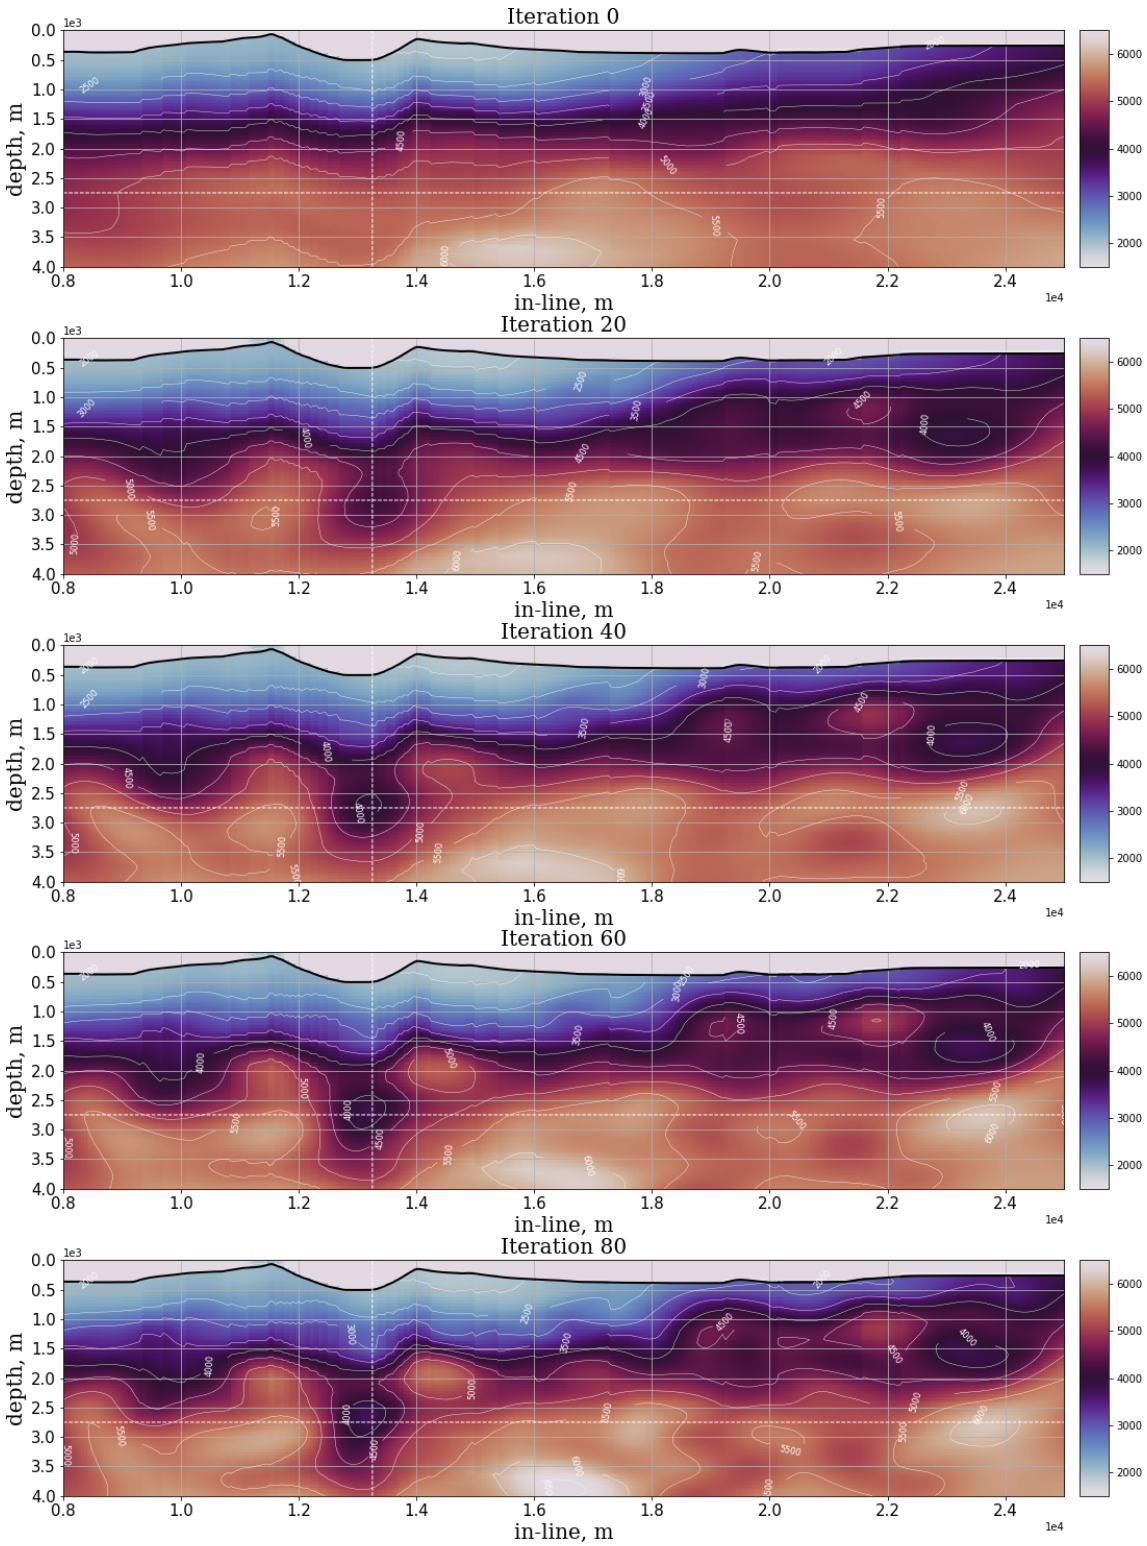


Figure 8.26: Slices through the final model at $y = 8.75$ km.

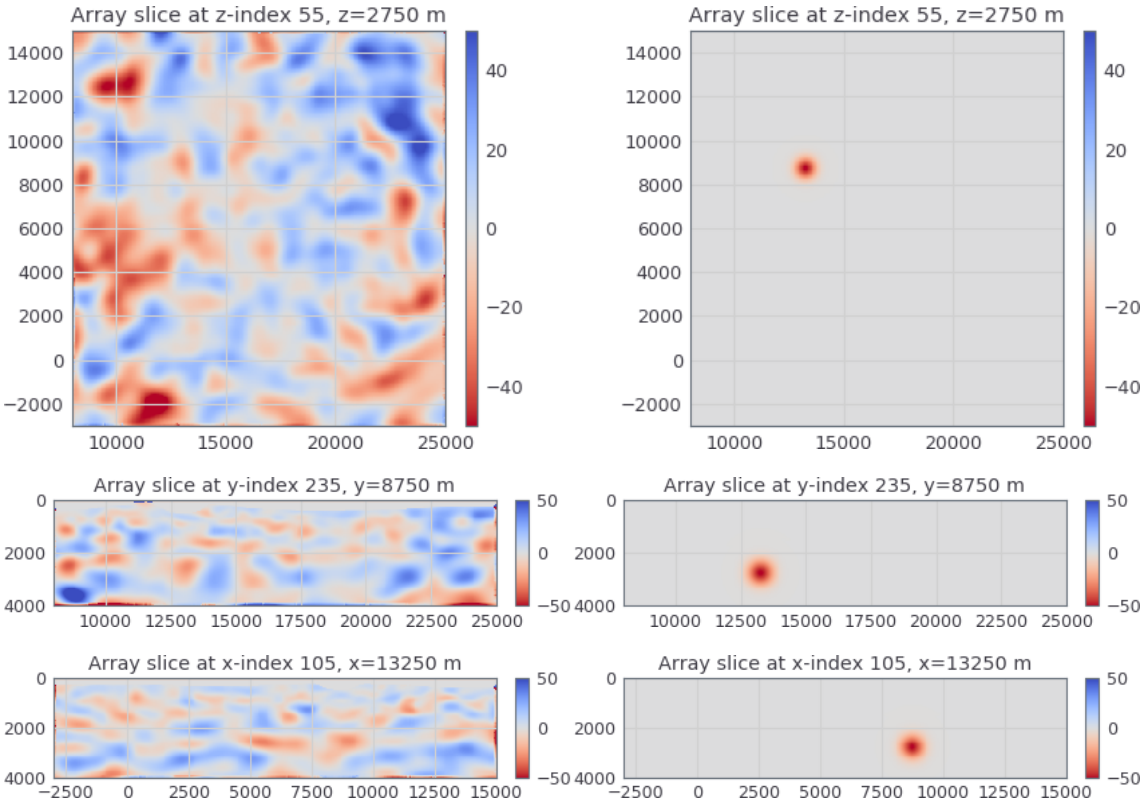


Figure 8.27: Final minus starting model (left) in the inversion of synthetic data generated from the 'true' model which was the starting/background model perturbed with a negative anomaly underneath Kolumbo. Right – true minus background model. The anomaly was a spherically-symmetric 3d Gaussian with a peak amplitude of -1% of the background model and a full-width at half-maximum (FWHM) of 800 m (right).

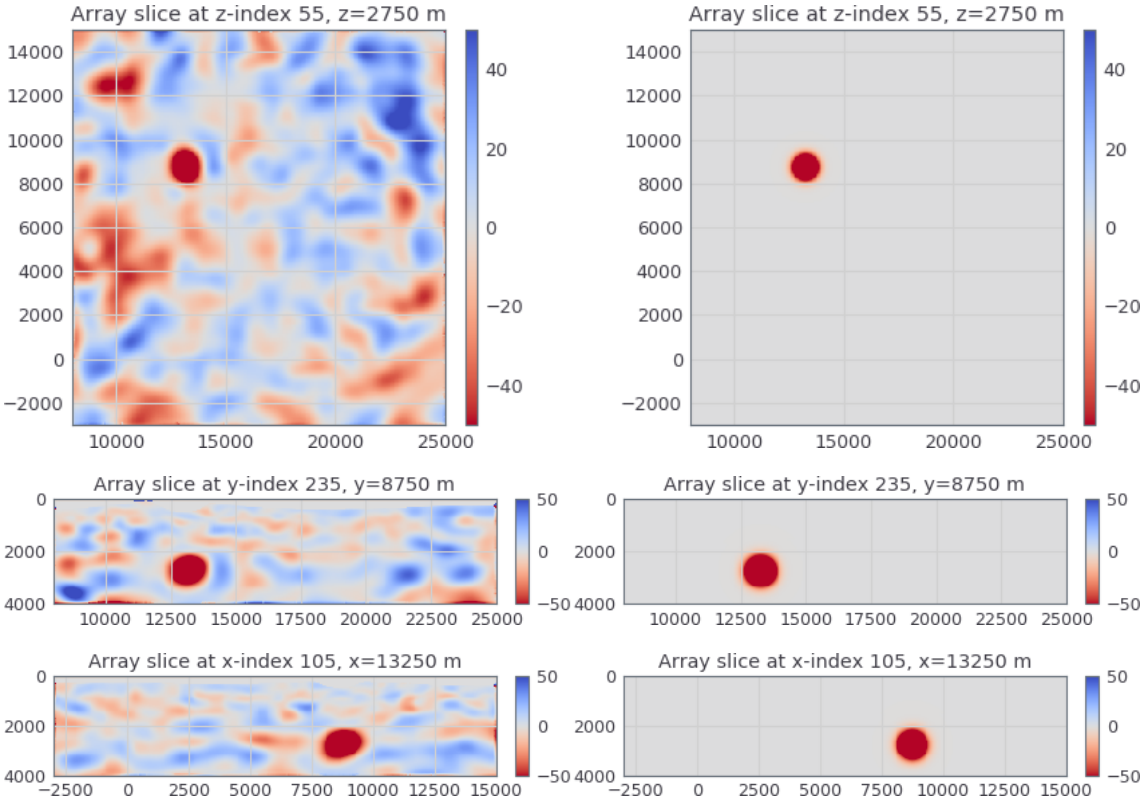


Figure 8.28: Same as fig. 8.27 but for anomaly with a peak amplitude of -5%.

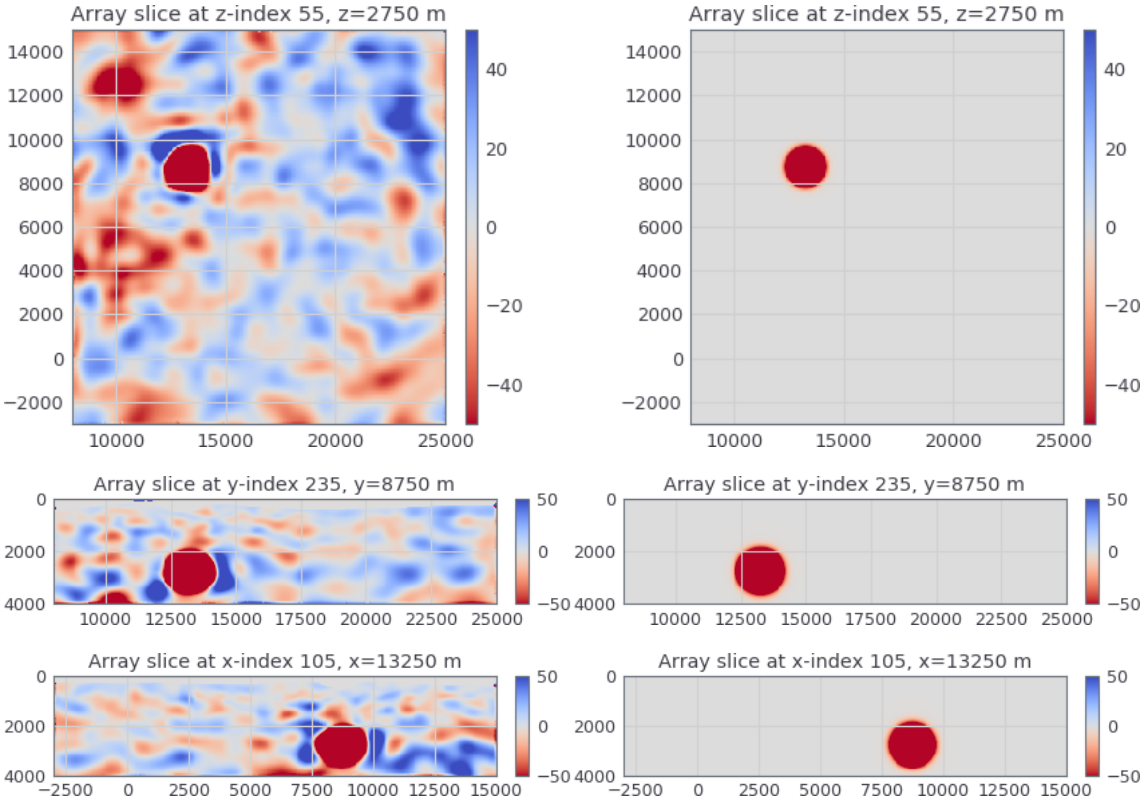


Figure 8.29: Same as fig. 8.27 but for anomaly with a peak amplitude of -50%.

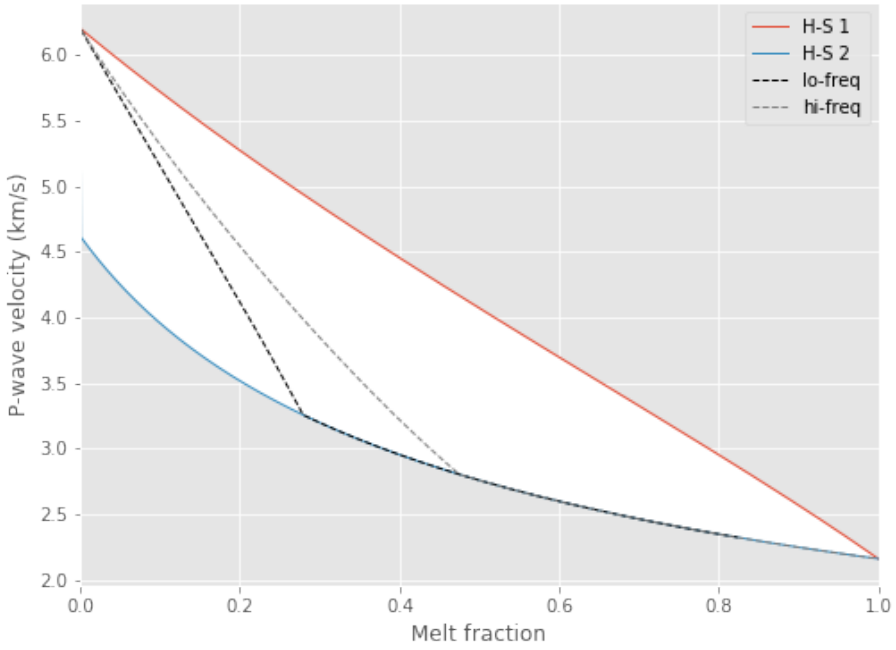
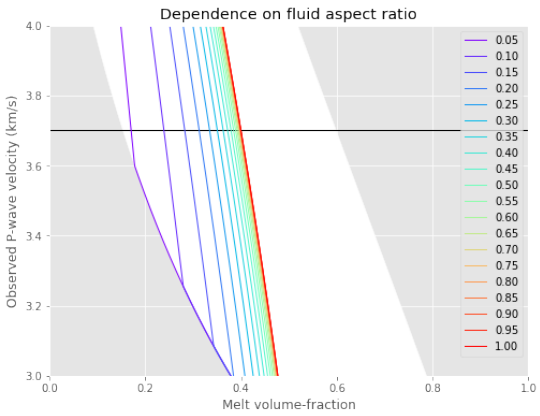
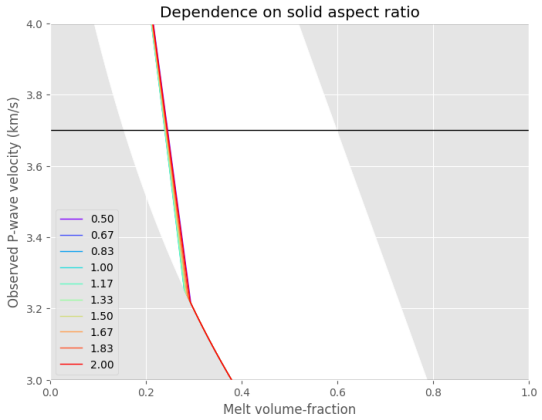


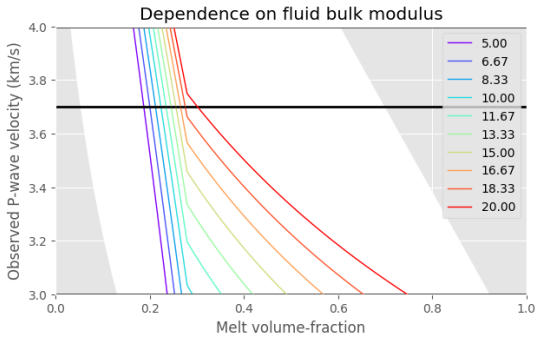
Figure 8.30: High versus low frequency limit (Mavko et al., 2009) used in the calculation of the effective parameters. All other figures of this appendix are for the low-frequency limit. Note that for a given v_P, melt fraction is lower in the low-frequency limit, and the difference increases with decreasing v_P. Blue and red curve denote the Hashin-Shtrikman lower and upper bound, respectively.



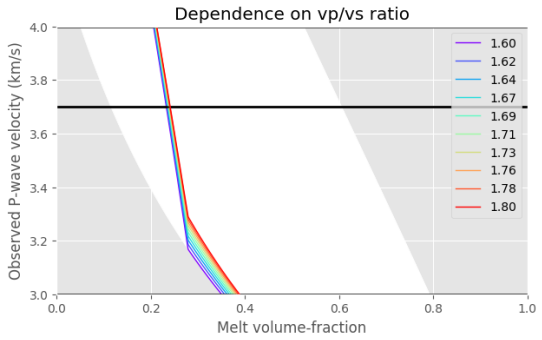
(a) Dimensionless.



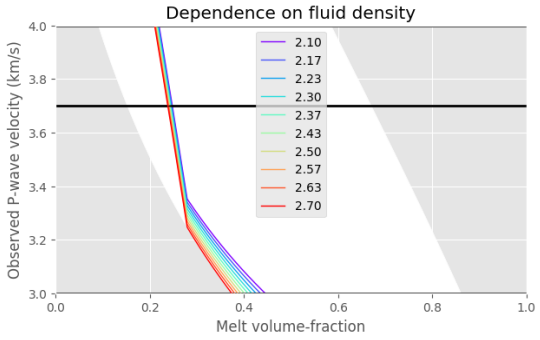
(b) Dimensionless.



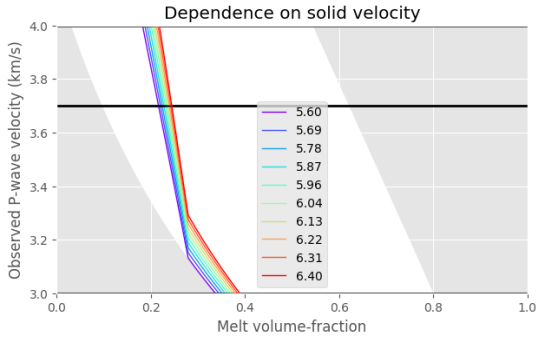
(c) In GPa.



(d) Dimensionless.



(e) In g/cm3.



(f) In km/s.

Figure 8.31: Analysis of the relationship between v_P and melt-fraction of a two-phase effective medium (Berryman, 1980) in terms of sensitivity to different properties of the solid and fluid phases. Each figure presents a single-parameter study (values in color, units in captions) with other parameters fixed. Area between Hashin-Shtrikman bounds denoted in white.

Concluding remarks

The general objective of this PhD project was to apply 3d full-waveform inversion (FWI) to the PROTEUS dataset and obtain a high-resolution model of P-wave velocity across the Santorini volcanic field which could be interpreted in terms of potential melt reservoirs distributed throughout the crust. The FWI was intended to built upon the travel-time tomography model obtained by the part of the PROTEUS team based at the University of Oregon.

The data were acquired in the PROTEUS experiment in which tens of ocean-bottom and land seismometers were deployed to record waves generated by thousands of airgun explosions. Although local-earthquake tomographers can only dream of sampling the study area with a similar density and regularity of sources, such coverage is extremely sparse compared to the industry standards such as the Valhall (Operto et al., 2015) or Tommeliten (Warner et al., 2013b) hydrocarbon fields. The amphibious acquisition centred around the steeply-dipping cliffs of the Santorini caldera recorded data adversely affected by surface-related coherent noise and thus challenging for acoustic modelling. Hence, the major part of the research was devoted to (1) developing the framework of quality-controlling the marine-data inversions and (2) implementing methodology to make inversion of the land data possible. These efforts are discussed below in the reversed order compared to their exposition in the main body of the thesis.

Part III

The preparation and inversion of the marine data were described in the respective chapters of part III. The main result of this part was a discovery of a previously undetected low-velocity zone directly under the Kolumbo edifice. It has been interpreted as a shallow lens of granitic melt, based on the calculation of the effective properties of a partially molten granitic intrusion. In this way, the second question of sect. 2.4 has been answered. On the other hand, the first question of sect. 2.4 was to some extent already addressed with the travel-time tomography performed prior to this project. The methodology developed in part II (discussed below) will help to further refine these models and also to answer questions 3-5 which require long-offset land data[17].

Importance The most hazardous volcanoes on Earth are underlain by complex systems of molten-rock reservoirs ranging from melt-poor mush zones to melt-rich magma chambers (Sparks et al., 2019). Despite being a cornerstone of volcanological models of melt storage, differentiation and transport towards the Earth's surface, no magma chamber has ever been imaged unambiguously. Recently the paradigm shifted towards the conclusion that magma chambers are too short-lived to capture, and melt must spend most of its life in mush zones. However some chambers may have gone undetected

[17] The long-offset OBS data are severely affected by noise and require more processing and certainly meticulous selection.

due to the inherent resolution limits of the seismic methods used to date. Here, the full-waveform inversion, a next-generation imaging method, was applied to study Kolumbo, the centre of seismic and hydrothermal activity of the Santorini volcanic field, Greece. Combining FWI and active-source data for the first time in volcanology led to the discovery of a small magma chamber extending from 2.1 to at least 4 km b.s.l. The image is the first-ever tomographic evidence of a high-melt-fraction reservoir. The chamber is a missing link between the high-temperature hydrothermal system venting at the crater floor of Kolumbo and the mantle source or a deeper reservoir yet to be imaged. Though too small to be detected by standard seismic tomography, the melt reservoir is large enough to threaten the nearby islands with tsunamigenic eruptions. Its location, size and melt fraction can be readily used in numerical models of magma differentiation and eruption forecasting. The results suggest that similar melt reservoirs are yet to be discovered by full-waveform inversion at other active volcanoes.

Outlook The above results suggest that despite of Santorini being much more famous, its smaller submarine neighbour may be worth more attention of the local authorities in the near future. The marine data subset at Kolumbo has a potential to become a great testing ground for FWI of more data and with more complete physics. Merging the obtained model with the recently acquired high-resolution 3d multi-channel data (Karstens et al., 2020) is an exciting prospect. Upcoming IODP drilling campaign across the Santorini volcano field will hopefully verify and complement the joint interpretation of these two projects. Experimental results on elastic properties of partially molten rocks that could better constrain the effective-medium calculations are badly missing. This is likely to change in the near future due to efforts of M. Paulatto and others.

Part II

Inverting land data with the `FULLWAVE3D` high-performance FWI code developed at Imperial was not possible before the start of this project. Two missing features were added to the code as described in the respective chapters of part II. First, dipole sources and receivers were implemented allowing to accurately model multi-component data using the 2nd-order acoustic wave equation with and without reciprocity. The reciprocity principle is key to understand the adjoint method which underpins most of the imaging methods used in seismology. Second, the implementation of the immersed-boundary method for simulating irregular boundaries such as the free surface was provided. The IB method is notoriously troublesome to implement. Writing it in a low-level programming language such as `Fortran` (as it is now considered), and having to port it to a large, highly-optimised and MPI-permeated codebase was a challenge that took a lot of time of this project. Hopefully, its description in ch. 6 highlights the crucial practicalities that other published reports seem to skip over.

Importance In order to image fine-scale features of the subsurface such as faults, melt lenses, gas clouds etc. one has to use a high-resolution method called full-waveform inversion. The success of inversion depends strongly on the accuracy of the forward step which seeks to find a numerical solution to the wave equation. This solution is most commonly obtained by the finite-difference (FD) method which uses regular grids and thus avoids a challenging meshing step. However, full-waveform modelling with the FD method is not accurate enough if the model contains an irregular free surface such as the topography of mountainous terrain. The issue may be addressed with the immersed boundary (IB) method (Mittal and Iaccarino, 2005) which can complement existing FD schemes without a need for modifying their kernels. Unfortunately, no immersed boundary method that could simulate realistic 3d surfaces and land seismic data has been reported to date, limiting the application of

finite-difference FWI to purely marine targets. This study bridges the gap by presenting an efficient 3d immersed boundary method whose flexible implementation can be easily ported to existing FD solvers. This opens up possibilities of applying full-waveform inversion to a whole new range of important targets.

Outlook The developed tools will be applied to refine the travel-tomography model directly beneath the Santorini caldera. Another natural near-term application is to the NZ3D experiment targeting the nature of slow-slip events at the Hikurangi subduction zone in New Zealand. The immersed-boundary code, due to its modular design, may be easily ported to other finite-different codes. With the ongoing developments in computing, in particular in stencil optimisation (Louboutin et al., 2019), finite-difference solvers are here to stay for longer than expected by the competing numerical methods. Developers of maturing finite-different codes may benefit from this study.

Final remarks

This interdisciplinary project was concerned with applying the established FWI code to a new, challenging dataset, extracting from it a geologically sound information, as well as with developing new capabilities of this code for future research. Apart from the substantial programming output, the thesis delivered a potentially interesting geological result, which contributed to making the Santorini volcanic field one of the best studied systems of this kind in the world. Where other tomographic studies of continental volcanoes had to stop, PROTEUS experiment has just started. Unlocking more secrets of the Santorini volcano field and generalising the conclusions to other volcanoes is under way.

Bibliography

Aki, K., Christoffersson, A., and Husebye, E. S. (1977). Determination of the three-dimensional seismic structure of the lithosphere. *Journal of Geophysical Research*, 82:277.

Aki, K. and Richards, P. G. (2009). *Quantitative Seismology*. University Science Books.

Alkhalifah, T. and Plessix, R.-É. (2014). A recipe for practical full-waveform inversion in anisotropic media: An analytical parameter resolution study. *Geophysics*, 79:R91–R101.

Andinisari, R., Konstantinou, K., and Ranjan, P. (2021). Seismicity along the Santorini-Amorgos zone and its relationship with active tectonics and fluid distribution. *Physics of the Earth and Planetary Interiors*, 312:106660.

Annen, C., Blundy, J. D., and Sparks, R. S. (2006). The genesis of intermediate and silicic magmas in deep crustal hot zones. *Journal of Petrology*, 47(3):505–539.

Arntsen, B. and Carcione, J. M. (2000). A new insight into the reciprocity principle. *Geophysics*, 65:1604 – 1612.

Asimow, P. D. and Ghiorso, M. S. (1998). Algorithmic modifications extending MELTS to calculate subsolidus phase relations. *American Mineralogist*, 83:1127–1132.

Baydin, A., Pearlmutter, B., Raduls, A., and Siskind, J. (2017). Automatic Differentiation in Machine Learning: a Survey. *Journal of Machine Learning*, 18:1–45.

Baysal, E., Kosloff, D. D., and Sherwood, J. W. (1983). Reverse time migration. *Geophysics*, 48:1514–1524.

Berkhout, A. J. (1984). Multidimensional linearized inversion and seismic migration. *Geophysics*, 49(11):1881–1895.

Berryman, J. G. (1980). Long-wavelength propagation in composite elastic media II. Ellipsoidal inclusions. *Journal of the Acoustical Society of America*, 68:1820–1831.

Best, M. G. (2002). *Igneous and Metamorphic Petrology*. Wiley-Blackwell, 2nd edition.

Beyer, R. P. and LeVeque, R. J. (1992). Analysis of a one-dimensional model for the immersed boundary method. *SIAM Journal on Numerical Analysis*, 29:332–364.

Bharadwaj, P., Mulder, W., and Drijkoningen, G. (2016). Full waveform inversion with an auxiliary bump functional. *Geophysical Journal International*, 206:1076–1092.

Bijwaard, H., Spakman, W., and Engdahl, E. R. (1998). Closing the gap between regional and global travel time tomography. *Journal of Geophysical Research*, 103.

Bingol, O. R. and Krishnamurthy, A. (2019). NURBS-Python: An open-source object-oriented NURBS modeling framework in Python. *SoftwareX*, 9:85–94.

Birch, F. (1947). Finite elastic strain of cubic crystals. *Phys. Rev.*, 71(11):809–824.

Bleibinhaus, F. (2014). Full-Waveform inversion of Controlled-Source seismic data. In *Encyclopedia of Earthquake Engineering*, pages 1–13.

Bleibinhaus, F. and Rondenay, S. (2009). Effects of surface scattering in full-waveform inversion. *Geophysics*, 74:WCC69–WCC77.

Boettcher, A. L. and Wyllie, P. J. (1968). Melting of granite with excess water to 30 kilobars pressure. *J. Geol.*, 76(2):235–244.

Bohlen, T., Daniel, K., and Jetschny, S. (2015). SOFI3D - seismic modeling with finite differences, 3D - acoustic and viscoelastic version. *Karlsruhe Institute of Technology*, pages 1–54.

Bohnhoff, M., Makris, J., Papanikolaou, D., and Stavrakakis, G. (2001). Crustal investigation of the Hellenic subduction zone using wide aperture seismic data. *Tectonophysics*, 343:239–262.

Bohnhoff, M., Rische, M., Meier, T., Becker, D., Stavrakakis, G., and Harjes, H.-P. (2006). Microseismic activity in the hellenic volcanic arc, greece, with emphasis on the seismotectonic setting of the Santorini–Amorgos zone. *Tectonophysics*, 423(1):17–33.

Brenguier, F., Shapiro, N. M., Campillo, M., Ferrazzini, V., Duputel, Z., Coutant, O., and Nercessian, A. (2008). Towards forecasting volcanic eruptions using seismic noise. *Nature Geoscience*, 1:126–130.

Brocher, T. M. (2005). Empirical relations between elastic wavespeeds and density in the Earth's crust. *Bulletin of the Seismological Society of America*, 95:2081–2092.

Brossier, R., Operto, S., and Virieux, J. (2009). Seismic imaging of complex onshore structures by 2D elastic frequency-domain full-waveform inversion. *Geophysics*, 74(6):WCC105–WCC118.

Bunks, C., Saleck, F. M., Zaleski, S., and G., C. (1995). Multiscale seismic waveform inversion. *Geophysics*.

Burgisser, A. and Bergantz, G. W. (2011). A rapid mechanism to remobilize and homogenize highly crystalline magma bodies. *Nature*, 471(7337):212–217.

Burridge, R., Lapwood, E. R., and Knopoff, L. (1964). First motions from seismic sources near a free surface. *Bulletin of the Seismological Society of America*, 54:1889–1913.

Cadoux, A., Scaillet, B., Druitt, T. H., and Deloule, E. (2014). Magma storage conditions of large Plinian eruptions of Santorini volcano (Greece). *Journal of Petrology*, 55:1129–1171.

Calderón Agudo, O., Vieira da Silva, N., Warner, M., Kalinicheva, T., and Morgan, J. (2018). Addressing viscous effects in acoustic full-waveform inversion. *Geophysics*, 83:R611–R628.

Cantner, K., Carey, S., and Nomikou, P. (2014). Integrated volcanologic and petrologic analysis of the 1650AD eruption of Kolumbo submarine volcano, Greece. *Journal of Volcanology and Geothermal Research*, 269:28–43.

Carey, S., Nomikou, P., Bell, K. C., Lilley, M., Lupton, J., Roman, C., Stathopoulou, E., Bejelou, K., and Ballard, R. (2013). CO2 degassing from hydrothermal vents at Kolumbo submarine volcano, Greece, and the accumulation of acidic crater water. *Geology*, 41:1035–1038.

Cashman, K. V., Sparks, R. S. J., and Blundy, J. D. (2017). Vertically extensive and unstable magmatic systems: A unified view of igneous processes. *Science*, 355.

Cashman, K. V., Stephen, R., and Sparks, J. (2013). How volcanoes work: A 25 year perspective. *Bulletin of the Geological Society of America*, 125:664–690.

Cerjan, C., Kosloff, D., Kosloff, R., and Reshef, M. (1985). A nonreflecting boundary condition for discrete acoustic and elastic wave equations. *Geophysics*, 50(4):705–708.

Chen, X. D., Xu, G., Yong, J. H., Wang, G., and Paul, J. C. (2009). Computing the minimum distance between a point and a clamped B-spline surface. *Graphical Models*, 71:107–112.

Chen, X. D., Yong, J. H., Wang, G., Paul, J. C., and Xu, G. (2008). Computing the minimum distance between a point and a NURBS curve. *CAD Computer Aided Design*, 40:1051–1054.

Christensen, N. I. (1979). Compressional wave velocities in rocks at high temperatures and pressures, critical thermal gradients, and crustal low-velocity zones. *Journal of Geophysical Research*, 84(B12):6849–6858.

Christensen, N. I. and Stanley, D. (2003). Seismic velocities and densities of rocks. *International Geophysics*, 81:1587–1594.

Claerbout, J. (1971). Toward a unified theory of reflector mapping. *Geophysics*, 36.

Claerbout, J. F. (1985). *Fundamentals of geophysical data processing: with applications to petroleum prospecting*. Blackwell Scientific Publications.

Cmiral, M., Fitz Gerald, J. D., Faul, U. H., and Green, D. H. (1998). A close look at dihedral angles and melt geometry in olivine-basalt aggregates: A TEM study. *Contributions to Mineralogy and Petrology*, 130:336–345.

Collier, J. and Sinha, M. (1990). Seismic images of a magma chamber beneath the lau basin back-arc spreading centre. *Nature*, 346(6285):646–648.

Cooper, K. M. and Kent, A. J. (2014). Rapid remobilization of magmatic crystals kept in cold storage. *Nature*, 506(7489):480–483.

Creager, K. C. and Dorman, L. M. (1982). Location of instruments on the seafloor by joint adjustment of instrument and ship positions. *Journal of Geophysical Research*, 87:8379–8388.

da Silva, N. and Warner, M. (2017). Global Visco-Acoustic Full Waveform Inversion. In *79th EAGE Conference & Exhibition*, pages 12–15.

da Silva, N. V., Yao, G., and Warner, M. (2019). Semiglobal viscoacoustic full-waveform inversion. *Geophysics*, 84:R271–R293.

Davy, R. G., Morgan, J. V., Minshull, T. A., Bayrakci, G., Bull, J. M., Klaeschen, D., Reston, T. J., Sawyer, D. S., Lymer, G., and Cresswell, D. (2018). Resolving the fine-scale velocity structure of continental hyperextension at the Deep Galicia Margin using full - waveform inversion. *Geophysical Journal International*, 212:244–263.

Debens, H. A. (2015). *Three-Dimensional Anisotropic Full-Waveform Inversion*. PhD thesis.

Detrick, R. S., Buhl, P., Vera, E., Mutter, J., Orcutt, J., Madsen, J., and Brocher, T. (1987). Multichannel seismic imaging of a crustal magma chamber along the East Pacific Rise. *Nature*, 326:35–41.

Diebold, J. B., Tolstoy, M., Doermann, L., Nooner, S. L., Webb, S. C., and Crone, T. J. (2010). R/V Marcus G. Langseth seismic source: Modeling and calibration. *Geochemistry, Geophysics, Geosystems*, 11:1–20.

Dimitriadis, I., Karagianni, E., Panagiotopoulos, D., Papazachos, C., Hatzidimitriou, P., Bohnhoff, M., Rische, M., and Meier, T. (2009). Seismicity and active tectonics at Coloumbo Reef (Aegean Sea, Greece): Monitoring an active volcano at Santorini Volcanic Center using a temporary seismic network. *Tectonophysics*, 465:136–149.

Dimitriadis, I., Papazachos, C., Panagiotopoulos, D., Hatzidimitriou, P., Bohnhoff, M., Rische, M., and Meier, T. (2010). P and S velocity structures of the Santorini-Coloumbo volcanic system (Aegean Sea, Greece) obtained by non-linear inversion of travel times and its tectonic implications. *Journal of Volcanology and Geothermal Research*, 195:13–30.

Druitt, T. H. (2014). New insights into the initiation and venting of the Bronze-Age eruption of Santorini (Greece), from component analysis. *Bulletin of Volcanology*, 76:1–21.

Druitt, T. H., Costa, F., Deloule, E., Dungan, M., and Scaillet, B. (2012). Decadal to monthly timescales of magma transfer and reservoir growth at a caldera volcano. *Nature*, 482:77–80.

Druitt, T. H., Edwards, L., Mellors, R. M., Pyle, D. M., Sparks, R. S. J., Lanphere, M., Davies, M., and Barreirio, B. (1999). Santorini volcano. *Geological Society Memoir*, 19.

Druitt, T. H., Mercier, M., Florentin, L., Deloule, E., Cluzel, N., Flaherty, T., Medard, E., and Cadoux, A. (2016). Magma storage and extraction associated with plinian and interplinian activity at Santorini Caldera (Greece). *Journal of Petrology*, 57:461–494.

Edmonds, M., Cashman, K. V., Holness, M., Jackson, M., and Edmonds, M. (2019). Architecture and dynamics of magma reservoirs Subject Areas : Author for correspondence :. *Philosophical Transactions of the Royal Society of London. A: Mathematical, Physical and Engineering Sciences*, 377.

England, P. C. and Katz, R. F. (2010). Melting above the anhydrous solidus controls the location of volcanic arcs. *Nature*, 467(7316):700–703.

Esser, E., Guasch, L., Herrmann, F. J., and Warner, M. (2016). Constrained waveform inversion for automatic salt flooding. *The Leading Edge*, 35:235–239.

ExxonMobil (2021). Full-waveform inversion. https://www.youtube.com/watch?v=VWF--1OnitM. Accessed: 2021-07-07.

Fabbro, G. N., Druitt, T. H., and Costa, F. (2017). Storage and eruption of silicic magma across the transition from dominantly effusive to caldera-forming states at an arc volcano (Santorini, Greece). *Journal of Petrology*, 58:2429–2464.

Fichtner, A., Kennett, B. L. N., Igel, H., and Bunge, H. P. (2008). Theoretical background for continental- and global-scale full-waveform inversion in the time-frequency domain. *Geophysical Journal International*, 175:665–685.

Fichtner, A., Saygin, E., Taymaz, T., Cupillard, P., Capdeville, Y., and Trampert, J. (2013). The deep structure of the North Anatolian Fault Zone. *Earth and Planetary Science Letters*, 373:109–117.

Fichtner, A. and Trampert, J. (2011). Resolution analysis in full waveform inversion. *Geophysical Journal International*, 187:1604–1624.

Fichtner, A., Zunino, A., and Gebraad, L. (2019). Hamiltonian Monte Carlo solution of tomographic inverse problems. *Geophysical Journal International*, 216:1344–1363.

Flinders, A. F. and Shen, Y. (2017). Seismic evidence for a possible deep crustal hot zone beneath southwest washington. *Sci. Rep.*, 7(1):7400.

Fouqué, F. (1879). *Santorin et ses éruptions*. G. Masson.

French, S., Lekic, V., and Romanowicz, B. (2013). Waveform tomography reveals channeled flow at the base of the oceanic asthenosphere. *Science*, 342.

French, S. W. and Romanowicz, B. (2015). Broad plumes rooted at the base of the Earth's mantle beneath major hotspots. *Nature*, 525:95–99.

Gao, K., Lin, Y., Huang, L., Alamos, L., Queen, J., City, P., Moore, J., City, S. L., and Berkeley, L. (2015a). Anisotropic elastic waveform inversion with modified total-variation regularization. *85th Annual International Meeting, SEG, Expanded Abstracts*, pages 5158–5163.

Gao, L., Brossier, R., Pajot, B., Tago, J., and Virieux, J. (2015b). An immersed free surface boundary treatment for seismic wave simulation. *Geophysics*, 80:1405–1409.

Gardner, G. H. F., Gardner, L. W., and Gregory, A. R. (1974). Formation velocity and density; the diagnostic basics for stratigraphic traps. *Geophysics*, 39.

Gelman, S. E., Gutiérrez, F. J., and Bachmann, O. (2013). On the longevity of large upper crustal silicic magma reservoirs. *Geology*, 41(7):759–762.

Górszczyk, A., Operto, S., and Malinowski, M. (2017). Toward a robust workflow for deep crustal imaging by FWI of OBS data: The eastern Nankai Trough revisited. *Journal of Geophysical Research: Solid Earth*, 122:4601–4630.

Gray, M., Bell, R. E., Morgan, J. V., Henrys, S., Barker, D. H. N., science parties Parties, t. I. E. ., and 375 (2019). Imaging the shallow subsurface structure of the North Hikurangi subduction zone, New Zealand, using 2D full-waveform inversion. *Journal of Geophysical Research: Solid Earth*, pages 9049–9074.

Guasch, L., Calderón Agudo, O., Tang, M.-X., Nachev, P., and Warner, M. (2020). Full-waveform inversion imaging of the human brain. *npj Digital Medicine*, 3:1–12.

Guasch, L., Štekl, I., Umpleby, A., and Warner, M. (2010). Elastic wavefield inversion in three dimensions. In *72nd EAGE Conference & Exhibition*, pages 1–5.

Guasch, L., Warner, M., and Ravaut, C. (2019). Adaptive waveform inversion: Practice. *Geophysics*, 84:R447–R461.

Gubbins, D. (2004). *Time Series Analysis and Inverse Theory for Geophysicists*. Cambridge University Press.

Gupta, I. N. and Kisslinger, C. (1966). Radiation of body waves from near-surface explosive sources. *Geophysics*, 31:1057–1063.

Hansen, P. C. (1992). Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Review*, 34:561–580.

Hashin, Z. and Shtrikman, S. (1963). A variational approach to the theory of the elastic behaviour of multiphase materials. *Journal of the Mechanics and Physics of Solids*, 11:127–140.

Heath, B. A. (2019). *Interactions between tectonism and magmatism at Santorini: insights from an active source seismic experiment*. PhD thesis.

Heath, B. A., Hooft, E. E., Toomey, D. R., Papazachos, C. B., Nomikou, P., Paulatto, M., Morgan, J. V., and Warner, M. R. (2019). Tectonism and Its Relation to Magmatism Around Santorini Volcano From Upper Crustal P Wave Velocity. *Journal of Geophysical Research: Solid Earth*, 124:10610–10629.

Hicks, G. J. (2002). Arbitrary source and receiver positioning in finite-difference schemes using Kaiser windowed sinc functions. *Geophysics*, 67:156–165.

Hildreth, W. and Moorbath, S. (1988). Crustal contributions to arc magmatism in the Andes of Central Chile. *Contributions to Mineralogy and Petrology*, 98(4):455–489.

Hobro, J. W. and Rickett, J. E. (2014). Explicit high-contrast surface refinement using full waveform inversion. *76th EAGE Conference and Exhibition 2014*.

Hobro, J. W. D. (1999). *Three-dimensional tomographic inversion of combined reflection and refraction seismic travel-time data*. PhD thesis.

Hooft, E. (2015a). PROTEUS - Plumbing Reservoirs Of The Earth Under Santorini. MGL1521 Cruise Report.

Hooft, E. (2015b). Report from Santorini: Marine Geophysics in Action.

Hooft, E. E., Nomikou, P., Toomey, D. R., Lampridou, D., Getz, C., Christopoulou, M.-E., O'Hara, D., Arnoux, G. M., Bodmer, M., Gray, M., Heath, B. A., and VanderBeek, B. P. (2017). Backarc tectonism, volcanism, and mass wasting shape seafloor morphology in the Santorini-Christiana-Amorgos region of the Hellenic Volcanic Arc. *Tectonophysics*, 712:396–414.

Hooft, E. E. E., Heath, B. A., Toomey, D. R., Paulatto, M., Papazachos, C. B., Nomikou, P., Morgan, J. V., and Warner, M. (2019). Seismic imaging of Santorini: Subsurface constraints on caldera collapse and present-day magma recharge. *Earth and Planetary Science Letters*, 514:48–61.

Hu, W. (2016). An improved immersed boundary finite-difference method for seismic wave propagation modeling with arbitrary surface topography. *Geophysics*, 81:T311–T322.

Hübscher, C., Hensch, M., Dahm, T., Dehghani, A., Dimitriadis, I., Hort, M., and Taymaz, T. (2006). Toward a risk assessment of central Aegean volcanoes. *Eos*, 87.

Hübscher, C., Ruhnau, M., and Nomikou, P. (2015). Volcano-tectonic evolution of the polygenetic Kolumbo submarine volcano/Santorini (Aegean Sea). *Journal of Volcanology and Geothermal Research*, 291:101–111.

Hudson, J. A. and Heritage, J. R. (1981). The use of the Born approximation in seismic scattering problems. *Geophysical Journal of the Royal Astronomical Society*, 66:221–240.

Huff, A. E., Nomikou, P., Thompson, L. A., Hooft, E. E. E., Walker, I. J., Huff, A. E., Nomikou, P., Thompson, L. A., and Hooft, E. E. E. (2021). Applying planetary mapping methods to submarine environments : onshore-offshore geomorphology of Christiana-Santorini-Kolumbo Volcanic Group , Greece. *Journal of Maps*, 00:1–11.

Huiskes, M. J., Plessix, R.-E., and Mulder, W. A. (2017). Acoustic VTI Full-waveform Inversion with 3-D Free-surface Topography. In *79th EAGE Conference & Exhibition*.

Igel, H. (2016). *Computational Seismology: A Practical Introduction*. Oxford University Press, 1st edition.

Jackson, M. D., Blundy, J., and Sparks, R. S. J. (2018). Chemical differentiation, cold storage and remobilization of magma in the Earth's crust. *Nature*.

Johnson, D. E. and Cohen, E. (1998). A Framework For Efficient Minimum Distance Computations. In *Proceedings of the 1998 IEEE International Conference on Robotics & Automation*.

Jolivet, L. and Brun, J.-P. (2010). Cenozoic geodynamic evolution of the Aegean. *International Journal of Earth Sciences*, 99:109–138.

Jolivet, L., Faccenna, C., Huet, B., Labrousse, L., and others (2013). Aegean tectonics: Strain localisation, slab tearing and trench retreat. *Tectonophysics*.

Kaiser, J. F. and Schafer, R. W. (1980). On the Use of the I0-Sinh Window for Spectrum Analysis. *IEEE TRANSACTIONS ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING*, ASSP-28(1):105–107.

Kak, A. C. and Slaney, M. (2001). *Principles of Computerized Tomographic Imaging*. Society of Industrial and Applied Mathematics.

Kalinicheva, T. I. (2019). *Full waveform inversion of narrow-azimuth towed-streamer seismic data*. PhD thesis.

Kamei, R., Pratt, R. G., and Tsuji, T. (2012). Waveform tomography imaging of a megasplay fault system in the seismogenic Nankai subduction zone. *Earth and Planetary Science Letters*, 317-318:343–353.

Karagianni, E. E., Papazachos, C. B., Panagiotopoulos, D. G., Suhadolc, P., Vuan, A., and Panza, G. F. (2005). Shear velocity structure in the Aegean area obtained by inversion of Rayleigh waves. *Geophysical Journal International*, 160:127–143.

Karato, S.-i. (1993). Importance of anelasticity in the interpretation of seismic tomography. 20(15):1623–1626.

Karátson, D., Gertisser, R., Telbisz, T., Vereb, V., Quidelleur, X., Druitt, T., Nomikou, P., and Kósik, S. (2018). Towards reconstruction of the lost Late Bronze Age intra-caldera island of Santorini, Greece. *Scientific Reports*, 8:1–8.

Karstens, J., Crutchley, G., Elger, J., Kühn, M., Schmid, F., Dalla Valle, G., Preine, J., and Nomikou, P. (2020). R/V Poseidon Cruise Report 538-THESEUS Tsunami hazard of explosive submarine eruptions. In *GEOMAR Helmholtz Centre for Ocean Research Kiel*.

Kent, G. M., Singh, S. C., Harding, A. J., Sinha, M. C., Orcutt, J. A., Barton, P. J., White, R. S., Bazin, S., Hobbs, R. W., Tong, C. H., and Pye, J. W. (2000). Evidence from three-dimensional seismic reflectivity images for enhanced melt supply beneath mid-ocean-ridge discontinuities. *Nature*, 406(6796):614–618.

Kilias, S. P., Nomikou, P., Papanikolaou, D., Polymenakou, P. N., Godelitsas, A., Argyraki, A., Carey, S., Gamaletsos, P., Mertzimekis, T. J., Stathopoulou, E., Goettlicher, J., Steininger, R., Betzelou, K., Livanos, I., Christakis, C., Bell, K. C., and Scoullos, M. (2013). New insights into hydrothermal vent processes in the unique shallow-submarine arc-volcano, Kolumbo (Santorini), Greece. *Scientific Reports*, 3:1–13.

Klaver, M., Carey, S., Nomikou, P., Smet, I., Godelitsas, A., and Vroon, P. (2016). A distinct source and differentiation history for Kolumbo submarine volcano, Santorini volcanic field, Aegean arc. *Geochemistry Geophysics Geosystems*, 17:3254–3273.

Ko, K. and Sakkalis, T. (2014). Orthogonal projection of points in CAD/CAM applications: an overview. *Journal of Computational Design and Engineering*, 1:116–127.

Komatitsch, D. and Tromp, J. (1999). Introduction to the spectral element method for three-dimensional seismic wave propagation. *Geophys. J. Int.*, 139:806–822.

Konstantinou, K. I. (2020). Magma chamber evolution during the 1650 AD Kolumbo eruption provides clues about past and future volcanic activity. *Scientific Reports*, 10:1–12.

Kreiss, H.-O. and Petersson, N. A. (2006). A Second Order Accurate Embedded Boundary Method for the Wave Equation with Dirichlet Data. *SIAM Journal on Scientific Computing*, 27:1141–1167.

Kristek, J., Moczo, P., and Archuleta, R. J. (2002). Efficient methods to simulate planar free surface in the 3D 4(th)-order staggered-grid finite-difference schemes. *Studia Geophysica Et Geodaetica*, 46:355–381.

Lailly, P. (1983). The seismic inverse problem as a sequence of before stack migrations. In Bednar, J. B., editor, *Conference on Inverse Scattering–Theory and Application*. SIAM.

Landau, L. D. and Lifshitz, E. M. (1970). *Theory of Elasticity*. Pergamon Press.

LD-EO (2021). Lamont-Doherty Earth Observatory webpage. https://www.ldeo.columbia.edu/research/office-of-marine-operations. Accessed: 2021-07-07.

Lees, J. M. (2007). Seismic tomography of magmatic systems. *Journal of Volcanology and Geothermal Research*, 167:37–56.

Levander, A. R. (1988). Fourth-order finite-difference P-SV seismograms. *Geophysics*, 53:1425–1436.

Lions, J. L. (1972). *Nonhomogeneous boundary value problems and applications*. Springer-Verlag, Berlin.

Lombard, B. and Piraux, J. (2004). Numerical treatment of two-dimensional interfaces for acoustic and elastic waves. *Journal of Computational Physics*, 195:90–116.

Lombard, B., Piraux, J., Gélis, C., and Virieux, J. (2008). Free and smooth boundaries in 2-D finite-difference schemes for transient elastic waves. *Geophysical Journal International*, 172:252–261.

Louboutin, M., Lange, M., Luporini, F., Kukreja, N., Witte, P. A., Herrmann, F. J., Velesko, P., and Gorman, G. J. (2019). Devito (v3.1.0): an embedded domain-specific language for finite differences and geophysical exploration. *Geoscientific Model Development*, 12:1165–1187.

Madariaga, R. (1976). Dynamics of an expanding circular fault. *Bulletin of the Seismological Society of America*, 66:639–666.

Magee, C., Stevenson, C. T. E., Ebmeier, S. K., Keir, D., Hammond, J. O. S., Gottsmann, J. H., Whaler, K. A., Schofield, N., Jackson, C. A.-L., Petronis, M. S., O'Driscoll, B., Morgan, J., Cruden, A., Vollger, S. A., Dering, G., Micklethwaite, S., and Jackson, M. D. (2018). Magma Plumbing Systems: A Geophysical Perspective. *Journal of Petrology*, 59:1217–1251.

Malinowski, M., Operto, S., and Ribodetti, A. (2011). High-resolution seismic attenuation imaging from wide-aperture onshore data by visco-acoustic frequency-domain full-waveform inversion. *Geophysical Journal International*, 186:1179–1204.

Martin, V. M., Holness, M. B., and Pyle, D. M. (2006). Textural analysis of magmatic enclaves from the Kameni Islands, Santorini, Greece. *Journal of Volcanology and Geothermal Research*, 154:89–102.

Mavko, G., Mukerji, T., and Dvorkin, J. (2009). *The Rock Physics Handbook*. Cambridge University Press.

McKenzie, D. (2020). The structure of the lithosphere and upper mantle beneath the Eastern Mediterranean and Middle East. *Mediterranean Geoscience Reviews*, 2(3):311–326.

McVey, B. G. (2019). *Magma accumulation beneath Santorini volcano from P-wave tomography*. PhD thesis.

McVey, B. G., Hooft, E. E., Heath, B. A., Toomey, D. R., Paulatto, M., Morgan, J. V., Nomikou, P., and Papazachos, C. B. (2020). Magma accumulation beneath Santorini volcano from P-wave tomography. *Geology*, 48.

Menke, W. (2012). *Geophysical Data Analysis: Discrete Inverse Theory*. Elsevier, 3rd edition.

Métivier, L. and Brossier, R. (2016). The SEISCOPE optimization toolbox: A large-scale nonlinear optimization library based on reverse communication. *Geophysics*, 81(2):F1–F15.

Métivier, L., Brossier, R., Mérigot, Q., Oudet, E., and Virieux, J. (2016). Measuring the misfit between seismograms using an optimal transport distance: Application to full waveform inversion. *Geophysical Journal International*, 205:345–377.

Métivier, L., Brossier, R., Virieux, J., and Operto, S. (2013). Full waveform inversion and the truncated newton method. *SIAM J. Sci. Comput.*, 35(2).

Miller, D., Oristaglio, M., and Beylkin, G. (1987). New Slant on Seismic Imaging: Migration and Integral Geometry. *Geophysics*, 52:943–964.

Mittal, R., Dong, H., Bozkurttas, M., Najjar, F. M., Vargas, A., and von Loebbecke, A. (2008). A versatile sharp interface immersed boundary method for incompressible flows with complex boundaries. *Journal of Computational Physics*, 227:4825–4852.

Mittal, R. and Iaccarino, G. (2005). Immersed Boundary Methods. *Annual Review of Fluid Mechanics*, 37:239–261.

Moczo, P., Kristek, J., and Gális, M. (2014). *The finite-difference modelling of earthquake motions: Waves and ruptures*. Cambridge University Press.

Montes, F. (2013). Research Vessel Marcus G. Langseth. http://www.shipspotting.com/gallery/photo.php?lid=1831846. Accessed: 2021-07-07.

Mora, P. (1987). Nonlinear elastic inversion of multioffset seismic data. *Geophysics*, 52:1211–1228.

Mora, P. (1988). Elastic wave-field inversion of reflection and transmission data. *Geophysics*, 53:750–759.

Mora, P. (1989). Inversion = migration + tomography. *Geophysics*, 54.

Morgan, J., Warner, M., Arnoux, G., Hooft, E., Toomey, D., VanderBeek, B., and Wilcock, W. (2016). Next-generation seismic experiments - II: Wide-angle, multi-azimuth, 3-D, full-waveform inversion of sparse field data. *Geophysical Journal International*, 204:1342–1363.

Morgan, J., Warner, M., Bell, R., Ashley, J., Barnes, D., Little, R., Roele, K., and Jones, C. (2013). Next-generation seismic experiments: Wide-angle, multi-azimuth, three-dimensional, full-waveform inversion. *Geophysical Journal International*, 195:1657–1678.

Morgan, J. V., Warner, M. R., Collins, G. S., Grieve, R. A., Christeson, G. L., Gulick, S. P., and Barton, P. J. (2011). Full waveform tomographic images of the peak ring at the Chicxulub impact crater. *Journal of Geophysical Research: Solid Earth*, 116:1–14.

Mulder, W. A. (2017). A simple finite-difference scheme for handling topography with the second-order wave equation. *Geophysics*, 82:1–46.

Mulder, W. A. and Huiskes, M. J. (2017). A simple finite-difference scheme for handling topography with the first-order wave equation. *Geophysical Journal International*, 210:482–499.

Nangoo, T. P. (2013). *Seismic Full-Waveform Inversion of 3D Field Data – From the Near Surface to the Reservoir*. PhD thesis.

Nash, S. G. (1985). Preconditioning of Truncated-Newton methods. *SIAM J. Sci. and Stat. Comput.*, 6(3):599–616.

Nemeth, T., Wu, C., and Schuster, G. T. (1999). Least-squares migration of incomplete reflection data. *Geophysics*, 64:208–221.

Newman, A. V., Stiros, S., Feng, L., Psimoulis, P., Moschas, F., Saltogianni, V., Jiang, Y., Papazachos, C., Panagiotopoulos, D., Karagianni, E., and Vamvakaris, D. (2012). Recent geodetic unrest at Santorini Caldera, Greece. *Geophysical Research Letters*, 39:1–5.

Nigbor, R. (1994). Six-degrees-of-freedom ground-motion measurement. *Bull. Seismol. Soc. Am.*, 84.

Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*.

Nomikou, P., Carey, S., Bell, K. L. C., Papanikolaou, D., Bejelou, K., Cantner, K., Sakellariou, D., and Perros, I. (2014a). Tsunami hazard risk of a future volcanic eruption of Kolumbo submarine volcano, NE of Santorini Caldera, Greece. *Natural Hazards*, 72:1375–1390.

Nomikou, P., Hübscher, C., and Carey, S. (2019). The Christiana-Santorini-Kolumbo Volcanic Field. *Elements*, 15:171–176.

Nomikou, P., Hübscher, C., Ruhnau, M., and Bejelou, K. (2016). Tectono-stratigraphic evolution through successive extensional events of the Anydros Basin, hosting Kolumbo volcanic field at the Aegean Sea, Greece. *Tectonophysics*, 671:202–217.

Nomikou, P., Papanikolaou, D., Alexandri, M., Sakellariou, D., and Rousakis, G. (2013). Submarine volcanoes along the Aegean volcanic arc. *Tectonophysics*, 597-598:123–146.

Nomikou, P., Parks, M. M., Papanikolaou, D., Pyle, D. M., Mather, T. A., Carey, S., Watts, A. B., Paulatto, M., Kalnins, M. L., Livanos, I., Bejelou, K., Simou, E., and Perros, I. (2014b). The emergence and growth of a submarine volcano: The Kameni islands, Santorini (Greece). *GeoResJ*, 1-2:8–18.

Nuber, A., Manukyan, E., and Maurer, H. (2016). Ground topography effects on near-surface elastic full waveform inversion. *Geophysical Journal International*, 207:67–71.

Operto, S., Gholami, Y., Prieux, V., Ribodetti, A., Brossier, R., Metivier, L., and Virieux, J. (2013). A guided tour of multiparameter full-waveform inversion with multicomponent data: From theory to practice. *The Leading Edge*, 32:1040–1054.

Operto, S., Miniussi, A., Brossier, R., Combe, L., Métivier, L., Monteiller, V., Ribodetti, A., and Virieux, J. (2015). Efficient 3-D frequency-domain mono-parameter full-waveform inversion of ocean-bottom cable data: application to Valhall in the visco-acoustic vertical transverse isotropic approximation. *Geophysical Journal International*, 202:1362–1391.

Operto, S., Virieux, J., Dessa, J. X., and Pascal, G. (2006). Crustal seismic imaging from multifold ocean bottom seismometer data by frequency domain full waveform tomography: Application to the eastern Nankai trough. *Journal of Geophysical Research: Solid Earth*, 111:1–33.

Opršal, I. and Zahradník, J. (1999). Elastic finite-difference method for irregular grids. *Geophysics*, 64:240.

Owusu, J. C., Podgornova, O., Charara, M., Leaney, S., Campbell, A., Ali, S., Borodin, I., Nutt, L., and Menkiti, H. (2016). Anisotropic elastic full-waveform inversion of walkaway vertical seismic profiling data from the Arabian Gulf. *Geophysical Prospecting*, 64:38–53.

Papazachos, B. C., Karakostas, V. G., Papazachos, C. B., and Scordilis, E. M. (2000). The geometry of the Wadati-Benioff zone and lithospheric kinematics in the Hellenic arc. *Tectonophysics*, 319:275–300.

Papazachos, C. (2017). Comments for source-time function generation on the basis of hydrophone-to-vertical comparisons for both water and body waves.

Parks, M. M., Biggs, J., England, P., Mather, T. a., Nomikou, P., Palamartchouk, K., Papanikolaou, X., Paradissis, D., Parsons, B., Pyle, D. M., Raptakis, C., and Zacharis, V. (2012). Evolution of Santorini Volcano dominated by episodic and rapid fluxes of melt from depth. *Nature Geoscience*, 5:749–754.

Parks, M. M., Moore, J. D. P., Papanikolaou, X., Biggs, J., Mather, T. A., Pyle, D. M., Raptakis, C., Paradissis, D., Hooper, A., Parsons, B., and Nomikou, P. (2015). From quiescence to unrest: 20 years of satellite geodetic measurements at Santorini volcano, Greece. *Journal of Geophysical Research : Solid Earth*, 120:1195–1209.

Paulatto, M. (2017). Bathymetry grid converted to local coordinates.

Paulatto, M., Laigle, M., Galve, A., Charvis, P., Sapin, M., Bayrakci, G., Evain, M., and Kopp, H. (2017). Dehydration of subducting slow-spread oceanic lithosphere in the Lesser Antilles. *Nature Communications*, 8:1–11.

Paulatto, M., Moorkamp, M., Hautmann, S., Hooft, E., Morgan, J. V., and Sparks, R. S. J. (2019). Vertically Extensive Magma Reservoir Revealed From Joint Inversion and Quantitative Interpretation of Seismic and Gravity Data. *Journal of Geophysical Research: Solid Earth*, 124:11,170–11,191.

Perrin, A., Goes, S., Prytulak, J., Rondenay, S., and Davies, D. R. (2018). Mantle wedge temperatures and their potential relation to volcanic arc location. *Earth and Planetary Science Letters*, 501:67–77.

Peskin, C. S. (1972). Flow patterns around heart valves: A numerical method. *Journal of Computational Physics*, 10:252–271.

Piper, D. J. W., Pe-Piper, G., Perissoratis, C., and Anastasakis, G. (2007). Distribution and chronology of submarine volcanic rocks around santorini and their relationship to faulting. *Geological Society, London, Special Publications*, 291(1):99–111.

Piper, D. J. W. and Perissoratis, C. (2003). Quaternary neotectonics of the south aegean arc. *Mar. Geol.*, 198(3):259–288.

Piraux, J. and Lombard, B. (2001). A New Interface Method for Hyperbolic Problems with Discontinuous Coefficients: One-Dimensional Acoustic Example. *Journal of Computational Physics*, 168:227–248.

Piromallo, C. and Morelli, A. (2003). P wave tomography of the mantle under the Alpine-Mediterranean area . *Journal of Geophysical Research: Solid Earth*, 108(B2):1–23.

Plessix, E. and Perkins, C. (2009). 3D full-waveform inversion with a frequency-domain iterative solver. In *71st EAGE Conference and Exhibition incorporating SPE EUROPEC 2009*. European Association of Geoscientists & Engineers.

Plessix, R. E. (2006). A review of the adjoint-state method for computing the gradient of a functional with geophysical applications. *Geophysical Journal International*, 167:495–503.

Plessix, R. É., Baeten, G., de Maag, J. W., ten Kroode, F., and Rujie, Z. (2012). Full waveform inversion and distance separated simultaneous sweeping: A study with a land seismic data set. *Geophysical Prospecting*, 60:733–747.

Płonka, A., Blom, N., and Fichtner, A. (2016). The imprint of crustal density heterogeneities on regional seismic wave propagation. *Solid Earth*, 7:1–28.

Pratt, R. G. (1989). *Wave equation methods in cross-hole seismic imaging*. PhD thesis.

Pratt, R. G. (1999). Seismic waveform inversion in the frequency domain, Part 1: Theory and verification in a physical scale model. *Geophysics*, 6411:888–901.

Pratt, R. G., Shin, C., and Hicks, G. J. (1998). Gauss-Newton and full Newton methods in frequency-space seismic waveform inversion. *Geophysical Journal International*, 133:341–362.

Pratt, R. G. and Shipp, R. M. (1999). Seismic waveform inversion in the frequency domain, part 2: Fault delineation in sediments using crosshole data. *Geophysics*, 64(3):902–914.

Pratt, R. G. and Worthington, M. H. (1990). Inverse theory applied to multi-source cross-hole tomography.. part 1: Acoustic wave-equation method1. *Geophys. Prospect.*, 38(3):287–310.

Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterline, W. T. (1990). *Numerical Recipes: The Art of Scientific Computing*. 4 edition.

Prieux, V., Brossier, R., Operto, S., and Virieux, J. (2013). Multiparameter full waveform inversion of multicomponent ocean-bottom-cable data from the Valhall field. part 1: Imaging compressional wave speed, density and attenuation. *Geophysical Journal International*, 194:1640–1664.

Ravaut, C., Operto, S., Improta, L., Virieux, J., Herrero, A., and Dell'Aversana, P. (2004). Multiscale imaging of complex structures from multifold wide-aperture seismic data by frequency-domain full-waveform tomography: Application to a thrust belt. *Geophysical Journal International*, 159:1032–1056.

Rawlinson, N. and Spakman, W. (2016). On the use of sensitivity tests in seismic tomography. *Geophysical Journal International*, 205:1221–1243.

Rayleigh, J. W. S. (1878). *The theory of sound, II.* Dover Publications, Inc. (reprinted 1945).

Ren, Y., Xu, X., Yang, S., Nie, L., and Chen, Y. (2020). A physics-based neural-network way to perform seismic full waveform inversion. *IEEE Access*, 4.

Rienstra, S. W. and Hirschberg, A. (2019). *An introduction to acoustics*. IWDE 92-06 lecture notes.

Riley, K. F., Hobson, M. P., and Bence, S. J. (2006). *Mathematical methods for physics and engineering*. Cambridge University Press.

Rizzo, A. L., Caracausi, A., Chavagnac, V., Nomikou, P., Polymenakou, P. N., Mandalakis, M., Kotoulas, G., Magoulas, A., Castillo, A., and Lampridou, D. (2016). Kolumbo submarine volcano (Greece): An active window into the Aegean subduction system. *Scientific Reports*, 6:1–9.

Robertsson, J. O. A. (1996). A numerical free-surface condition for elastic/viscoelastic finite-difference modeling in the presence of topography. *Geophysics*, 61:1921–1934.

Robertsson, J. O. A. and Holliger, K. (1997). Modeling of seismic wave propagation near the earth's surface. *Physics of the Earth and Planetary Interiors*, 104:193–211.

Robin, A., Singh, S. C., Noble, M., and Williamson, P. R. (2019). Pressure Versus Vertical Velocity Full-Waveform Inversion of Wide-Angle , Ocean-Bottom Node Data from the North Sea. In *81st EAGE Conference & Exhibition 2019*.

Rudnick, R. L. and Fountain, D. M. (1995). Nature and composition of the continental crust: A lower crustal perspective. *Rev. Geophys.*, 33(3):267.

Sambridge, M. and Mosegaard, K. (2002). Monte carlo methods in geophysical inverse problems. *Reviews of Geophysics*, 40.

Sambridge, M., Rickwood, P., Rawlinson, N., and Sommacal, S. (2007). Automatic differentiation in geophysical inverse problems. *Geophysical Journal International*, 170:1–8.

Sears, T., Barton, P., and Singh, S. (2010). Elastic Full Waveform Inversion of Multi-Component OBC Seismic Data – Application to Alba Field, U. K. North Sea. *Geophysics*, 75.

Sears, T. J., Singh, S. C., and Barton, P. J. (2008). Elastic full waveform inversion of multi-component OBC seismic data. *Geophysical Prospecting*, 56:843–862.

Sedova, A., Messud, J., Prigent, H., Masclet, S., Royle, G., and Lambaré, G. (2019). Acoustic land full waveform inversion on a broadband land dataset: The impact of optimal transport. In *81st EAGE Conference and Exhibition*.

Sedova, A., Royle, G. T., Hermant, O., Retailleau, M., and Lambaré, G. (2017). High-resolution land full waveform inversion - A case study on a data set from the Sultanate of Oman. *79th EAGE Conference and Exhibition 2017*, pages 12–15.

Shah, N., Warner, M., Nangoo, T., Umpleby, A., Štekl, I., Morgan, J., and Guasch, L. (2012a). Quality assured full-waveform inversion : Ensuring starting model adequacy. In *SEG Technical Program Expanded Abstracts 2012*, pages 1–5.

Shah, N., Warner, M., Washbourne, J., Guasch, L., and Umpleby, A. (2012b). A phase-unwrapped solution for overcoming a poor starting model in full-wavefield inversion. In *74th EAGE Conference & Exhibition*, pages 1–5.

Shah, N. K., Warner, M., Guasch, L., Štekl, I., and Umpleby, A. P. (2010). A Strategy for Waveform Inversion without an Accurate Starting Model. In *72nd EAGE Conference & Exhibition*, pages 1–5.

Shankar, R. (1994). *Principles of Quantum Mechanics*. Plenum Press.

Shearer, P. M. (2009). *Introduction to Seismology*. Cambridge University Press.

Shepard, D. (1968). Two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of 1968 ACM National Conference*, pages 517–524.

Sheriff, R. E. and Geldart, L. P. (1995). *Exploration Seismology*. Cambridge University Press.

Shin, C., Yoon, K., Marfurt, K. J., Park, K., Yang, D., Lim, H. Y., Chung, S., and Shin, S. (2002). Efficient calculation of a partial-derivative wavefield using reciprocity for seismic imaging and inversion. *Geophysics*, 66:1856–1863.

Shipp, R. M. and Singh, S. C. (2002). Two-dimensional full wavefield inversion of wide-aperture marine seismic streamer data. *Geophysical Journal International*, 151:325–344.

Sigurdsson, H., Carey, S., Alexandri, M., Vougioukalakis, G., Croff, K., Roman, C., Sakellariou, D., Anagnostou, C., Rousakis, G., Ioakim, C., Goguo, A., Ballas, D., Misaridis, T., and Nomikou, P. (2006). Marine investigations of greece's santorini volcanic field. *Eos*, 87(34):337.

Silverton, A., Warner, M., Morgan, J., and Umpleby, A. (2016). Offset-variable density improves acoustic full-waveform inversion: a shallow marine case study. *Geophysical Prospecting*, 64:1201–1214.

Singh, S. C., Crawford, W. C., Carton, H., Seher, T., Combier, V., Cannat, M., Canales, J. P., Düsünür, D., Escartin, J., and Miranda, J. M. (2006). Discovery of a magma chamber and faults beneath a Mid-Atlantic Ridge hydrothermal field. *Nature*, 442:1029–1032.

SIO (2021). Scripps Institution of Oceanography "LC4x4-SP" ocean-bottom seismometer. https://obsic.whoi.edu/instrumentation/sio/short-period/. Accessed: 2021-07-07.

Sirgue, L., Barkved, O. I., Dellinger, J., Etgen, J., Albertin, U., and Kommedal, J. H. (2010). Full waveform inversion: The next leap forward in imaging at Valhall. *First Break*, 28:65–70.

Sirgue, L., Etgen, J. T., and Albertin, U. (2008). 3D frequency domain waveform inversion using time domain finite difference methods. *70th European Association of Geoscientists and Engineers Conference and Exhibition 2008: Leveraging Technology. Incorporating SPE EUROPEC 2008*, 3:1602–1606.

Sirgue, L., I. Barkved, O., P. Van Gestel, J., J. Askim, O., and H. Kommedal, J. (2009). 3D waveform inversion on valhall wide-azimuth OBC. In *71st EAGE Conference and Exhibition incorporating SPE EUROPEC 2009*. European Association of Geoscientists & Engineers.

Sirgue, L. and Pratt, R. G. (2004). Efficient waveform inversion and imaging: A strategy for selecting temporal frequencies. *Geophysics*, 69:231–248.

Slob, E., Wapenaar, K., and Treitel, S. (2020). Tutorial: unified 1D inversion of the acoustic reflection response. *Geophysical Prospecting*, 68:1425–1442.

Solano, J. M., Jackson, M. D., Sparks, R. S., Blundy, J. D., and Annen, C. (2012). Melt segregation in deep crustal hot zones: A mechanism for chemical differentiation, crustal assimilation and the formation of evolved magmas. *Journal of Petrology*, 53:1999–2026.

Spakman, W., Wortel, M. J. R., and Vlaar, N. J. (1988). The Hellenic subduction zone: a tomographic image and its geodynamic implications. *Geophysical Research Letters*, 15(1):60–63.

Sparks, R. S. J., Annen, C., Blundy, J. D., Cashman, K. V., Rust, A. C., and Jackson, M. D. (2019). Formation and dynamics of magma reservoirs. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 377:20180019.

Sparks, S. R., Sigurdsson, H., and Wilson, L. (1977). Magma mixing: A mechanism for triggering acid explosive eruptions. *Nature*, 267:315–318.

Stern, R. J. and Dickinson, W. R. (2010). The Gulf of Mexico is a Jurassic backarc basin. (6):739–754.

Stockwell, J. W. and Cohen, J. K. (2008). The New SU User's Manual.

Strang, G. (1986). *Introduction to applied mathematics*. Wellesley-Cambridge Press.

Symes, W. W., Terentyev, I. S., and Vdovina, T. W. (2008). Gridding Requirements for Accurate Finite Difference Simulation. In *SEG Las Vegas 2008 Annual Meeting*, pages 2077–2081.

Symes, W. W. and Vdovina, T. (2009). Interface error analysis for numerical wave propagation. *Computational Geosciences*, 13:363–371.

Tarantola, A. (1984). Inversion of seismic reflection data in the acoustic approximation. *Geophysics*, 49:1259–1266.

Tarantola, A. (2005). *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial and Applied Mathematics.

Tarantola, A. (2006). Popper, Bayes and the inverse problem. *Nature Physics*, 2:492–494.

Tarasewicz, J., White, R. S., Woods, A. W., Brandsdóttir, B., and Gudmundsson, M. T. (2012). Magma mobilization by downward-propagating decompression of the Eyjafjallajkull volcanic plumbing system. *Geophysical Research Letters*, 39(19):1–5.

Taylor, J. R. (1997). *An Introduction to Error Analysis: The Study of Uncertainties in Physical Measurements*. University Science Books, 2nd edition.

Tirel, C., Gueydan, F., Tiberi, C., and Brun, J. P. (2004). Aegean crustal thickness inferred from gravity inversion. Geodynamical implications. *Earth and Planetary Science Letters*, 228:267–280.

Tolstoy, M., Diebold, J., Doermann, L., Nooner, S., Webb, S. C., Bohnenstiehl, D. R., Crone, T. J., and Holmes, R. C. (2009). Broadband calibration of the R/V Marcus G. Langseth four-string seismic sources. *Geochemistry, Geophysics, Geosystems*, 10.

Toomey, D. R., Solomon, S. C., and Purdy, G. M. (1994). Tomographic imaging of the shallow crustal structure of the East Pacific Rise at 9°30'N. *Journal of Geophysical Research: Solid Earth*, 99:24135–24157.

Trefethen, L. N. (1982). *Wave Propagation and Stability for Finite Difference Schemes*. PhD thesis.

Tromp, J., Tape, C., and Liu, Q. (2005). Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels. *Geophysical Journal International*, 160:195–216.

Tsampouraki-Kraounaki, K. and Sakellariou, D. (2018). Seismic stratigraphy and geodynamic evolution of Christiana Basin, South Aegean Arc. *Marine Geology*, 399:135–147.

Ueki, K. and Iwamori, H. (2016). Density and seismic velocity of hydrous melts under crustal and upper mantle conditions. *Geochemistry Geophysics Geosystems*, 17:1312–1338.

Ulvrova, M., Paris, R., Nomikou, P., Kelfoun, K., Leibrandt, S., Tappin, D. R., and McCoy, F. W. (2016). Source of the tsunami generated by the 1650 AD eruption of Kolumbo submarine volcano (Aegean Sea, Greece). *Journal of Volcanology and Geothermal Research*, 321:125–139.

van Hinsbergen, D. J. J., Spakman, W., Meulenkamp, J., and Wortel, R. (2005). Nappe stacking resulting from subduction of oceanic and continental lithosphere below Greece. *Geology*, 33(4):325–328.

van Leeuwen, T. and Herrmann, F. J. (2013). Mitigating local minima in full-waveform inversion by expanding the search space. *Geophysical Journal International*, 195:661–667.

van Vossen, R., Robertsson, J. O. A., and Chapman, C. H. (2002). Finite-difference modeling of wave propagation in a fluid-solid configuration. *Geophysics*, 67:618–624.

Vasconcelos, I., Ravasi, M., and van der Neut, J. (2017). Subsurface-domain, interferometric objective functions for target-oriented waveform inversion. *GEOPHYSICS*, pages 1–22.

Vigh, D., Cheng, X., Jiao, K., Xu, Z., and Sun, D. (2018). Essential steps for successful full-waveform inversion using land data. *2018 SEG International Exposition and Annual Meeting, SEG 2018*, pages 1128–1132.

Virieux, J. (1984). SH-wave propagation in heterogeneous media: Velocity-stress finite-difference method. *Geophysics*, 49(11):1933–1942.

Virieux, J. (1986). P-SV wave propagation in heterogeneous media: Velocity-stress finite-difference method. *Geophysics*, 51.

Virieux, J., Asnaashari, A., Brossier, R., Métivier, L., Ribodetti, A., and Zhou, W. (2014). An introduction to full waveform inversion. In *Encyclopedia of Exploration Geophysics*.

Virieux, J., Etienne, V., Cruz-Atienza, V., Brossier, R., Chaljub, E., Coutant, O., Garambois, S., Mercerat, D., Prieux, V., Operto, S., Ribodetti, A., and Tago, J. (2012). Modelling Seismic Wave Propagation for Geophysical Imaging. In *Seismic Waves - Research and Analysis*.

Virieux, J. and Operto, S. (2009). An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74:WCC1–WCC26.

Wang, H., Burtz, O., Routh, P., Wang, D., Violet, J., Lu, R., and Lazaratos, S. (2021). Anisotropic 3D elastic full-wavefield inversion to directly estimate elastic properties and its role in interpretation. *Lead. Edge*, 40(4):277–286.

Wapenaar, C. P. A. and Berkhout, A. J. (1989). *Elastic wave field extrapolation. Redatuming of single- and multi-component data*. Elsevier.

Wapenaar, K., Brackenhoff, J., and Thorbecke, J. (2019). Green's theorem in seismic imaging across the scales. *Solid Earth*, 10:517–536.

Wapenaar, K. and Fokkema, J. (2004). Reciprocity theorems for diffusion, flow, and waves. *Journal of Applied Mechanics, Transactions ASME*, 71:145–150.

Warner, M. and Guasch, L. (2014). Adaptive waveform inversion: Theory. *Geophysics*, 81:1089–1093.

Warner, M., Morgan, J., Umpleby, A., Štekl, I., and Guasch, L. (2012). Which physics for full-wavefield seismic inversion? In *74th EAGE Conference & Exhibition*, pages 1–5.

Warner, M., Nangoo, T., Shah, N., Umpleby, A., and Morgan, J. (2013a). Full-waveform inversion of cycle-skipped seismic data by frequency down-shifting. *SEG Technical Program Expanded Abstracts 2013*, 176:903–907.

Warner, M., Ratcliffe, A., Nangoo, T., Morgan, J., Umpleby, A., Shah, N., Vinje, V., Štekl, I., Guasch, L., Win, C., Conroy, G., and Bertrand, A. (2013b). Anisotropic 3D full-waveform inversion. *Geophysics*, 78:R59–R80.

Warner, M., Štekl, I., and Umpleby, A. (2008). 3D wavefield tomography: synthetic and field data examples. In *SEG Las Vegas 2008 Annual Meeting*, pages 3330–3334.

Warner, M., Umpleby, A., Štekl, I., and Morgan, J. V. (2010). 3D full-wavefield tomography: imaging beneath heterogeneous overburden. In *72nd EAGE Conference & Exhibition*, pages 1–5.

WHOI (2021). Woods Hole Oceanographic Institution "D2" ocean-bottom seismometer. https://web.whoi.edu/obslab/instrumentation/short-period/. Accessed: 2021-07-07.

Woodhouse, J. H. and Dziewoński, A. M. (1984). Mapping the upper mantle: Three-dimensional modeling of earth structure by inversion of seismic waveforms. *Journal of Geophysical Research*, 89(B7):5953–5986.

Woodward, M. J. (1992). Wave-equation tomography. *Geophysics*, 57:15–26.

Wu, R.-S. and Toksöz, M. N. (1987). Diffraction tomography and multisource holography applied to seismic imaging. *Geophysics*, 52:11.

Yao, G., da Silva, N. V., Debens, H. A., and Wu, D. (2017). Accurate seabed modeling using finite difference methods. *Computers & Geosciences*, 22:469–484.

Zhang, W. and Chen, X. (2006). Traction image method for irregular free surface boundaries in finite difference seismic wave simulation. *Geophysical Journal International*, 167:337–353.

Zhao, S. (2010). A fourth order finite difference method for waveguides with curved perfectly conducting boundaries. *Computer Methods in Applied Mechanics and Engineering*, 199:2655–2662.

Zheglova, P. and Malcolm, A. (2020). Vector-acoustic full-waveform inversion: Taking advantage of wavefield separation and dealiasing. *Geophysics*, 85:R409–R423.

Zhou, C., Cai, W., Luo, Y., Schuster, G. T., and Hassanzadeh, S. (1995). Acoustic wave-equation traveltime and waveform inversion of crosshole seismic data. *Geophysics*, 60(3):765–773.