Our faults and our stars:
some topics in comparative seismology



Benjamin A. Fernando

St Edmund Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary 2021

Acknowledgements

First of all, an enormous thank you to my family for putting up with my overly-excited ramblings about waves for the last five years. I hope that it has not been unbearable. Beyond that, the same to the new friends from Oxford and beyond who've made this project (mostly) so enjoyable.

This work would not have been possible without the contributions of my supervisors, Tarje Nissen-Meyer and Philipp Podsiadlowski. From the Seismology group, further thanks are also due to Kuangdai Leng, Karin Sigloch, Kasra Hosseini, Marta Pienkowska-Cote, Maria Tsekhmistrenko, Paula Koelemeijer, Thomas Garth, and Laura Ermert; as well as my fellow PhD students Claudia Haindl, Rebecca Colquhoun and Alice Turner. From the Astrophysics sub-department, the assistance and kindly ears of Ryosuke Hirai, Fabian Schneider, Matthew Clayton, Caroline Terquem, and Lucy Oswald are much appreciated.

More broadly from within the Oxford community, I'm very grateful to David Pyle, Victoria Forth, Lizzy Jeffers, Leslie Goldberg, and Keith Gull for their advice and help. From elsewhere, this work would not have been possible without the support of David Al-Attar, Eleonore Stutzmann, Lucia Gualteri, Ingrid Daubar, Shravan Hanasoge, Bill Chaplin, Martin van Driel, Warrick Ball, Stephen Justham, Selma de Mink, Hannah Alpert, Nick Teanby, Lars Hansen, Dave Wallis and Bruce Banerdt. The same goes for my fellow martian PhD students, Natalia Wójcicka, Geraldine Zenhäusern, Alex Stott, and Hannah Alpert.

This work was primarily funded by UK Research and Innovation through the Natural Environment Research Council's Doctoral Training Partnership at Oxford. Additional support for InSight related work was provided by the UK Space Agency's Aurora program. Supplementary funding was received from St Edmund Hall, the Royal Astronomical Society, the Worshipful Company of Scientific Instrument Makers, the High Altitude Observatory at Boulder, Colorado, and the Burdett-Coutts research fund at the University of Oxford.

Abstract

Seismology is one of the few geophysical probes which allows us to peer beneath the surface of a planet, moon, or star and map its structures and properties at a variety of scales and resolutions. This makes it a technique of enormous scientific power. Its application, however, suffers from problems of mathematical ill-posedness and poor conditioning. To alleviate some of these difficulties, we turn to computational seismic modelling.

Seismic modelling is not without its own challenges, of cost and complexity, which have limited its applications. In this thesis, material will be presented which advances the frontiers of seismic modelling, especially in the domain of coupled solid-fluid media. Application of these improvements offers the opportunity to better understand the seismic behaviour, and hence the interior structures, of the worlds in question.

In this thesis, we will cover advances in the seismic modelling of the Earth's oceans, Mars' atmosphere, the Sun, and other stars. The first five data chapters will cover the generation and analysis of synthetic data from the Earth, Mars (three times) and the Sun respectively. We will examine the influence that the oceans and their bathymetry have on surface wave propagation, followed by a presentation of results from the Oxford Seismology team's participation in the blind test for the NASA InSight mission to Mars. Next, we will present modelling which was used to inform InSight's attempts to detect the arrival of the Perseverance Rover in February 2021, followed by the results from this unique experiment.

The penultimate chapter covers the extension of similar techniques to modelling seismic wave propagation in the Sun, and finally we will explore whether asteroseismic techniques may be used to study the internal configuration of the λ Boötis class of stars.

Contents

1 Introduction **1**

1.1 Comparative seismology 1
1.2 A tour of the Solar System (and beyond) 3
1.2.1 The Earth 3
1.2.2 The Moon 6
1.2.3 Mars 7
1.2.4 Elsewhere in the Solar System 8
1.2.5 The Sun 9
1.2.6 Other stars 12
1.3 Inverse and forward methods 15
1.3.1 Inverse methods 15
1.3.2 Forward methods 16
1.4 Thesis particulars 17
1.4.1 Roadmap 17
1.4.2 Structure 17

2 The Oceans **21**

2.1 Preface 22
2.2 Summary 22
2.3 Introduction 23
2.4 Methodology 27
2.4.1 Theory 28
2.4.2 Implementation in AxiSEM3D 31
2.5 Validation 33
2.5.1 Nomenclature 33
2.5.2 Benchmark parameters 34
2.5.3 Benchmark results 35
2.6 Evaluation of ocean loading 38
2.7 Bathymetry in a global ocean context 43
2.7.1 Bathymetric implementation 43
2.7.2 The effects of bathymetry on waveforms 45
2.8 Propagation through the water column 52

2.9 The effects of the ocean on observed data . . . 53
2.10 Discussion . . . 56
2.10.1 Computational cost . . . 56
2.10.2 Cost scaling . . . 58
2.10.3 Potential cost savings . . . 59
2.10.4 Limitations . . . 60
2.10.5 Extensions in a global ocean model . . . 61
2.10.6 Development of a localised ocean formulation . . . 63
2.11 Summary . . . 64
2.12 Epilogue . . . 65
2.12.1 Microseismic noise: an introduction . . . 65
2.12.2 Noise in the Indian Ocean . . . 65
2.12.3 Results . . . 66

3 Mars, Part I **69**
3.1 Preface . . . 70
3.2 Summary . . . 70
3.3 Introduction . . . 71
3.3.1 Martian seismology . . . 71
3.3.2 The InSight spacecraft . . . 73
3.3.3 Terrestrial background . . . 74
3.3.4 Martian challenges . . . 75
3.3.5 The InSight Blind Test . . . 76
3.4 Pre-processing . . . 78
3.4.1 Data read-in . . . 78
3.4.2 Filtering . . . 79
3.4.3 Initial rotation . . . 80
3.5 Analysis . . . 81
3.5.1 Event identification . . . 81
3.5.2 Body waves . . . 82
3.5.3 Surface waves . . . 88
3.6 Results . . . 89
3.7 Conclusions . . . 90
3.8 Postface . . . 93
3.8.1 Performance . . . 93
3.8.2 Mars seismology in reality . . . 93

4 Mars, Part II **95**
4.1 Preface . . . 96
4.2 Summary . . . 97
4.3 Introduction . . . 98
4.3.1 Motivation . . . 98
4.3.2 Terrestrial and lunar context . . . 99
4.3.3 Extension of these methods to Mars . . . 100
4.3.4 The Mars 2020 EDL sequence . . . 101
4.3.5 Aims . . . 103
4.4 Methodology . . . 103
4.4.1 Acoustic signal: source and propagation . . . 105
4.4.2 Coupled seismoacoustic signal . . . 107
4.4.3 Elastodynamic seismic signal . . . 109
4.5 Results . . . 116
4.5.1 Acoustic signal . . . 116
4.5.2 Coupled seismoacoustic signal . . . 118
4.5.3 Elastodynamic seismic signal . . . 118
4.6 Discussion . . . 122
4.6.1 Noise conditions . . . 122
4.6.2 CMBD impact: detection probabilities . . . 125
4.6.3 Extensions to this work . . . 127
4.7 Conclusions . . . 130
4.8 Appendix - seismoacoustic simulations in AxiSEM3D . . . 132
4.8.1 Introduction . . . 132
4.8.2 Choice of boundary condition . . . 133
4.8.3 Theory . . . 134
4.8.4 Results . . . 135
4.8.5 Discussion . . . 135

5 Mars, Part III **139**
5.1 Preface . . . 139
5.2 Summary . . . 140
5.3 Introduction . . . 141
5.4 Methods . . . 141
5.5 Data . . . 143
5.6 Discussion . . . 145
5.6.1 Distance-amplitude scaling relationships . . . 145
5.6.2 Seismic efficiency and mantle attenuation . . . 147
5.7 Conclusions . . . 151

6 The Sun **153**
6.1 Preface . . . 153
6.2 Summary . . . 154
6.3 Introduction . . . 154
6.3.1 Motivation . . . 155
6.4 Theory . . . 156
6.4.1 Included effects . . . 156
6.4.2 Neglected effects . . . 158
6.4.3 Other approaches . . . 161
6.5 Implementation . . . 162
6.5.1 Outer boundary condition . . . 162
6.5.2 Seismic profile . . . 164
6.5.3 Gravity . . . 168
6.6 Results . . . 171
6.6.1 Mesh design and simulation cost . . . 171
6.6.2 Wave propagation . . . 173
6.7 Conclusions . . . 175

7 Other Stars **177**
7.1 Preface . . . 178
7.2 A quick primer on stellar... . . . 178
7.2.1 Life . . . 179
7.2.2 Death . . . 181
7.2.3 Classification . . . 182
7.2.4 Measurements . . . 183
7.2.5 Terminology . . . 183
7.3 Summary . . . 185
7.4 Introduction . . . 186
7.4.1 Population characteristics . . . 186
7.4.2 Potential formation channels . . . 186
7.4.3 The horizontal branch? . . . 187
7.4.4 Asteroseismology . . . 188
7.4.5 Tycho B . . . 188
7.5 Methodology . . . 189
7.5.1 Models . . . 190
7.5.2 Asteroseismology . . . 194
7.6 Results . . . 195
7.6.1 Mode frequencies . . . 195
7.6.2 Mode separations . . . 196

7.6.3 Temporal evolution 197
7.6.4 Sensitivity to degree of stripping 197
7.6.5 Rotation 198
7.7 Mode excitation 199
7.8 Applications 200
7.8.1 Ongoing observation campaigns 200
7.8.2 Tycho B 200
7.9 Conclusions 201

8 Conclusions **203**
8.1 Concluding introduction 203
8.2 Summary 204
8.2.1 Chapter 2: The oceans 204
8.2.2 Chapter 3: Mars, Part I 205
8.2.3 Chapter 4, Mars, Part II 205
8.2.4 Chapter 5, Mars, Part III 206
8.2.5 Chapter 6: The Sun 206
8.2.6 Chapter 7: Other stars 207
8.2.7 A few last notes... 208
8.3 Concluding conclusions/closing remarks 210

Works Cited **213**

Earthquakes travelling through the interior of the globe are like many messengers sent out to explore a new land. The messages are constantly coming and seismologists are fast learning to read them.

— Reginald Daly, *Our Mobile Earth (1926)*

1 Introduction

Contents

1.1 Comparative seismology 1
1.2 A tour of the Solar System (and beyond) 3
1.2.1 The Earth 3
1.2.2 The Moon 6
1.2.3 Mars 7
1.2.4 Elsewhere in the Solar System 8
1.2.5 The Sun 9
1.2.6 Other stars 12
1.3 Inverse and forward methods 15
1.3.1 Inverse methods 15
1.3.2 Forward methods 16
1.4 Thesis particulars 17
1.4.1 Roadmap 17
1.4.2 Structure 17

1.1 Comparative seismology

It has been more than three thousand years since Chinese scholars recorded the first occurrence of what we would today call an 'earthquake' (Lee et al. 1976). For the majority of the intervening time, they were regarded as signatures of divine intervention, rather than phenomena to be studied within the purview of what we today would call 'science'.

If we were to choose a particular event as having spurred the development of 'seismology' as a field, the Lisbon Earthquake of 1755 would be a strong contender - and indeed, is often cited as a catalyst for the European Enlightenment itself (Oldroyd et al. 2007; Fuchs 2009). The two centuries following the destruction of Lisbon cemented seismology's position as an invaluable tool for the study of the Earth's interior and dynamic processes, enabled by the development in the 1880s of the first modern-type seismometers (Ben-Menahem 1995).

Important subsequent findings included the discovery of the core-mantle and crust-mantle boundaries, the determination that the Earth's outer core is liquid, and that within it is a solid inner core. These advancements were accompanied by significant developments in tectonic and geodynamic studies, as well as substantial improvements in our understanding of earthquake and seismic processes themselves.

For all of this time, however, it remained unclear as to whether the Earth was alone in playing host to these sorts of geophysical processes - or put more simply, whether the other planets and moons of the Solar System experienced seismic activity as well. It was clear that the answer to this question would have profound implications for understanding the Earth's place in space, and how it and the other planets came to be the way that they are.

After many years of speculation, the first hints of an answer to this question were received in 1969 when the lead instrument of the Apollo Seismometer Network was deployed. The late 1970s and 80s saw such experiments extended to Mars and Venus, though with more limited success.

Perhaps surprisingly, the knowledge that stars (including our Sun) experience 'seismic' behaviour had been identified much earlier - though the relevant terms, 'helioseismology' and 'asteroseismology' were not initially used to describe these phenomena.

In the next section, I will briefly detail the types of seismic behaviour present on the worlds which are considered in this thesis. Whilst the exact research questions considered are different for Earth, Mars, and the Sun, the general principle remains the same: to use seismic waves as probe of interior structure and processes.

The aim is not only to understand the structure and properties of these worlds today, but also their evolutionary history.

Whilst the focus of this thesis is seismic modelling, simulations are of greatest use when there exist data to compare them to. Hence, this chapter will touch upon what sorts of seismic measurements are being made on different worlds today, and how they are recorded.

The discussion here is not intended to be exhaustive - each chapter of this thesis includes a paper-style introduction specific to the planet or star in question. Rather, it is intended to give a brief overview of how the different parts of what I have termed 'comparative seismology' fit together.

Before we proceed, however, a cautionary note. As a physicist (by background) who dabbles in both geophysics and astrophysics; and expects (or more realistically, hopes) that his work will be of interest to those from both communities, it is difficult to know what level of prior knowledge to assume. In general, I have worked on the basis that the reader has some passing familiarity with the structure and dynamics of worlds in question (for example, knowing that the Earth has a crust and the Sun a convective surface), but have tried to define any new seismic nomenclature (P-wave, p-mode, etc.) at its first usage. The sub-sections on the Earth and Sun, as the first two examples of terrestrial-type and astrophysical-type seismology respectively, are somewhat longer than subsequent entries in each category. Regardless, there are most likely places in this thesis where the level of detail is either excruciatingly painful or somewhat lacking, for those, I can only apologise.

1.2 A tour of the Solar System (and beyond)

1.2.1 The Earth

To many geoscientists, the basic seismology of the Earth is likely to be second nature, and hence this sub-section may be somewhat redundant. However, in drawing parallels with other worlds, it may prove useful to highlight particular aspects of terrestrial seismic wave propagation and nomenclature which we will encounter later, albeit in a different context.

1.2.1.1 Sources

On a world like the Earth, which is has both active tectonic and other surface processes, and is inhabited, the sources of seismic signals are numerous.

In the natural category, the most obvious examples are earthquakes, produced by the release of built-up tectonic stresses. These are the largest contributors to the total global seismicity budget by far, but are by no means the only one. The oceans (Longuet-Higgins 1950) and atmosphere (Kobayashi and Nishida 1998) also excite oscillations, which (thankfully, from a seismologist's perspective) are very different in presentation to those induced by earthquakes. More exotically, glacial and ice-sheet dynamics (Ekström et al. 2003), the impact or burn-up of meteorites (Latham et al. 1970b), or indeed animals (Mortimer et al. 2018; McDonald et al. 1995) may also be studied using the seismic oscillations they produce.

The other class of seismic sources - those which are artificial - are even more numerous. Those of most interest here are explosions (both nuclear: Mitchell (2002) and conventional: Latham et al. (1970a)), and the re-entry of spacecraft into the atmosphere (e.g. Hilton and Henderson (1974), de Groot-Hedlin et al. (2008), and ReVelle et al. (2005)), on the grounds that these generally constitute controlled events with known magnitudes and locations.

The waves which these sources produce are recorded by seismometers. On Earth, there are many such instruments, the most sensitive of which are permanently buried underground and well-insulated from both natural and artificial noise sources. There are many more instruments on temporary deployments, for example located for short periods of time on particular regions of the seafloor. The use of large-scale networks enables accurate inversion for structure and source properties when seismic waves are detected. Regardless of the nature of the source, the waves which they produce are generally classified as either 'body waves', which follow ballistic paths through the Earth's interior, and 'surface waves', which are confined to the crust and uppermost layers of the mantle.

1.2.1.2 Seismic phases

Body waves consist of two distinct families: P-waves, which are compressional, and S-waves, which involve shearing. Conversion between these phases at interfaces is possible, and in general S-waves cannot propagate in a fluid medium (this ceases to be necessarily true in environments where magnetohydrodynamic effects become important).

When propagating in the ocean or atmosphere, P-waves may be known as hydroacoustic or acoustic/infrasound waves respectively. If the fluid medium they propagate through is stratified, they may become trapped in waveguides formed by local minima in sound speed. In the oceanic case, these may propagate many thousands of kilometres without significant attenuation (Tolstoy and Ewing 1950). Diurnal and meteorological effects mean that atmospheric waveguides may be more transient, but are nonetheless regularly observed (Green and Nippress 2019).

Distinct surface waves of the kind seen on Earth are unique, and not observed on any of the other worlds considered in this thesis. On our planet they are split into two categories, a fast, transverse 'Love' wave and a slower, surface 'Rayleigh' wave. Rayleigh waves, as discussed in Chapters 2 and 3, are affected by the presence of an ocean or an atmosphere, whereas Love waves are not (at least to first order). Surface waves on Earth are particularly sensitive to the intrinsic attenuation and scattering properties of the crust, being confined within it for their entire propagation between source and receiver. For distinct, Earthlike surface waves to be observed on other worlds, a combination of an appropriate medium with low attenuation and scattering is needed. It appears that either such conditions are not found on any of the Solar System objects explored through seismology thus far; or that sources are too weak or too deep to excite observable surface waves at the recorded distances.

In addition to the distinct, coherent phases observed on Earth, a background of seismic 'noise' is ubiquitous. In the terrestrial frequency ranges of interest in most of this thesis (around 10 s), the dominant source of noise is the interaction of ocean waves with each other and with the seafloor in deep water. This 'microseismic noise' is generally treated as completely incoherent, meaning that the signal at

each point in space and time is random, and that sources at different points are entirely uncorrelated with each other. This is generally a reasonable assumption. The crosscorrelation of waveforms from two different receivers may then be used to extract a Green's function for propagation from one to the other, enabling the noise to be used as a seismic tool in and of itself (e.g. Stutzmann et al. (2009), Basini et al. (2013), and Grob et al. (2011)).

1.2.2 The Moon

The first in-situ seismic data from a world other than the Earth was recorded by the Seismometer Experiment carried onboard Apollo 11, which was deployed onto the lunar surface in July 1969 (Latham et al. 1969). The first instrument lasted only a few weeks, but four subsequent incarnations (from Apollo missions 12, 14, 15 and 16) and a gravimeter onboard Apollo 17 functioned until they were decommissioned in September 1977. Whilst future lunar seismology missions have been proposed (e.g. Yamada et al. (2011) and Neal et al. (2019)), none yet has funding for a full mission.

During the operation of the Apollo seismometer network thousands of 'moonquakes', which proved to be quite unlike terrestrial earthquakes in their nature (Nakamura et al. 1982) were recorded. Seismograms showed no clear phases, with arrivals being emergent rather than impulsive. Subsequent analysis of these features suggested that they were due to incredibly strong scattering in the lunar crust (Nunn et al. 2020). Whilst the Moon is not the topic of any particular chapter of this thesis, the use of controlled-impact seismology will prove relevant in later chapters, and hence let us touch on it briefly before moving on.

The monthly periodicities observed in most moonquakes suggested that the majority were tidally rather than tectonically excited. Some, however, appeared to be impact-related, with the observed seismic waves excited by objects striking the lunar surface. Four spent Saturn-IVB rocket stages were deliberately de-orbited during the operation of the seismometer network as part of these impact seismology investigations (Latham et al. 1970b). As their collisions with the surface had known masses, timings, and locations, the seismic waves they produced could be used to

invert for the average wavespeed along the propagation path, and thus used as a sort of 'calibration'. The final upper-stage de-orbit, that of Apollo 17 in 1972, marked the last time that a spacecraft on another world tried to detect the arrival of one of its colleagues until InSight's attempted detection of Perseverance's landing on Mars in 2021, some 48 years later. That experiment is discussed further in Chapter 4.

1.2.3 Mars

The first seismometers deployed onto a planet other than the Earth were carried onboard NASA's Viking 1 & 2 probes, launched in 1975 (Anderson et al. 1977). Both probes carried seismometers, but only Viking 2's deployed properly. Unfortunately, the position of the instrument on the lander meant that it proved to be more effective as a wind sensor than a seismometer (Nakamura and Anderson 1979). Subsequent reanalysis has suggested that at least one of the detected events may have been a genuine 'marsquake', rather than wind noise, though this inference is not conclusive (Lorenz et al. 2016).

A number of follow-on Mars or martian-moon seismology missions met with unfortunate ends: Mars 96 failed to leave Earth orbit, as did Fobos-Grunt, whilst Netlander was cancelled before launch. The NASA InSight (Interior Exploration using Seismic Investigations, Geodesy, and Heat Transport) mission, which was selected in 2012 and arrived at Mars in 2018 (Banerdt and Russell 2017) hence marked the first successful landing of a seismology package on Mars since the days of Viking. The InSight spacecraft and its instrument package are discussed further in Chapters 3, 4, and 5.

Efforts to better constrain the internal structure of Mars are very much ongoing (e.g. Panning et al. (2017) and Smrekar et al. (2019)). Previous investigations have used a variety of different techniques, including measurements from radio science experiments (Folkner et al. 1997; Sohl et al. 2005), estimation of tidal deformation levels (Yoder et al. 2003; Rivoldini et al. 2011); in many cases informed by studies of martian meteorites (e.g. Bertka and Fei (1998) and Combs et al. (2019)). It is generally agreed that Mars is fully differentiated. The core is, like

Earth's, iron-nickel and is thought to be at least partially molten. The mantle is likely still convecting, and is bounded by a crust which shows much greater lateral variations in thickness than does the Earth's (Watters et al. 2007). The surface is covered by a dusty-regolith material, the product of many billions of years of erosion. The atmosphere is much more rarified than Earth's, and is overwhelmingly made up of carbon dioxide, which has important implications for the coupling of seismic waves into the atmosphere (Bass and Chambers 2001).

It is probably fair to say that Mars' interior structure and seismic properties are similar enough to the Earth's that direct comparisons between the two are still meaningful, but different enough that observations are often confusing and challenging to interpret.

1.2.4 Elsewhere in the Solar System

Although they are not discussed specifically in this thesis, the many other planets, moons, and asteroids in the Solar System are expected to play host to seismic processes. For example, the Soviet Venera spacecraft carried seismometers to Venus in the 1980s, though they lasted less than a day before succumbing to the intense heat and pressure (Ksanfomaliti et al. 1982). Whilst they recorded signals which may be akin to microseismic processes on Earth, no clear 'venusquakes' were recorded. The giant planets, too, are observed to support global-scale, acoustic normal mode oscillations (Gaulme et al. 2011). The presence of more Earth-like seismic phenomena on their icy and rocky moons has also been predicted (e.g. Vance et al. (2018) and Stähler et al. (2018)).

In terms of smaller bodies, seismic activity is thought to play an important role in the resurfacing of asteroids (Richardson et al. 2004). Recently, this hypothesis was experimentally tested through the use of an artificial impact-type experiment carried out by the JAXA Hayabusa2 spacecraft on the asteroid 162173 Ryugu. I was involved in the modelling for this particular experiment, but for the sake of brevity will not discuss it further in this thesis. For further details, see Nishiyama et al. (2021).

The Rosetta mission's partially-successful Philae lander also carried a seismometer-like instrument to comet 67P/Churyumov-Gerasimenko (Kochan et al. 2000).

1.2.5 The Sun

One world thus far missing from our discussion of seismicity in the Solar System is the Sun. Visualising the seismic behaviour of a sphere of superheated plasma is challenging, but the best comparison is perhaps the global, normal mode oscillations of the Earth which are excited after large earthquakes.

There are, however, important differences, not least in the Sun's structure. These are in some cases so extreme that they necessitate fundamental changes to the ways in which we solve the wave equations which govern the propagation of seismic phases. The mathematical specifics are detailed in Chapter 5, but for way of context, let us first briefly recap some of the Sun's key attributes.

1.2.5.1 The structure of the Sun

The Sun's mass and radius are approximately 330,000 and 100 times those of the Earth, respectively. It is composed almost entirely of elemental hydrogen (73%) and helium (25%), and supported against gravitational collapse by thermonuclear fusion reactions in the core.

Energy transport in the Sun occurs through two dominant mechanisms. The inner 70% or so of the Sun's radius involves the transfer of energy through radiation, whilst the outer 30% involves transfer through convection. The interface between the two is called the 'tachocline', and is arguably the most important solar seismic structure (Spiegel and Zahn 1992). This juncture marks the boundary between internal solid body rotation and surface latitudinally-differential rotation; and is thought to play a crucial role in the generation of the Sun's magnetic field (Tobias 2005).

The Sun's atmosphere is complex and significantly more rarified than the Earth's across most of its depth. Magnetohydrodynamic effects are also important, complicating the choice of an outer boundary condition.

As discussed below, stellar structures generally support two different sets of global oscillation modes: those restored by pressure ('p-modes') and those restored by gravity ('g-modes'). Mixed modes acoustic-gravity modes are important in some stars, but are not relevant in the solar case. The reader may have noticed that the fact that the Sun's outer layers are convecting necessarily indicates that they are not stably stratified, and gravity-restored modes cannot propagate in such a medium. This means that any g-modes present in the radiative zone become evanescent in the convective zone, and would arrive at the surface with extremely small amplitudes. Their detection has been 'announced' many times over the last the last half-century, and remains controversial (Appourchaux et al. 2010). They are not anyhow simulatable using the techniques discussed in Chapter 5, so we will refrain from discussing them further; and instead wait for the next 'proof' of discovery to be announced.

1.2.5.2 Enter, helioseismology

Because the normal modes of the Sun are challenging to conceptualise, it is worth briefly detailing the events which led to their identification. This ought to serve as a primer into how such waves manifest themselves in observations.

During the 1960s, Dopplergrams (maps of the frequency shifts of particular spectral lines across the Sun's surface) indicated the presence of large-scale velocity fields in the solar atmosphere (specifically the photosphere; Evans and Michard (1962) and Leighton et al. (1962)). The patterns observed were consistent with the vertical oscillations of fluid parcels in the photosphere, with a characteristic period of around five minutes.

Follow-up work in the 1960s revealed that the 'five-minute oscillations' were in fact standing acoustic waves (p-modes) trapped in the Sun's outer layers (e.g. Ulrich (1970)); which were later found to be excited and damped in an entirely stochastic manner through turbulent processes in the Sun's convective zone. This phenomenon was dubbed 'helioseismology', and rapidly became a tool for probing the properties of the solar interior, in ways which were previously thought impossible.

Arguably the most important contribution ever made by helioseismologists concerned the solar neutrino problem (Bahcall and Ulrich 1988). Models of the nuclear fusion processes within the solar interior made a clear prediction of the number of electron neutrinos which would be produced in the Sun's core, and hence how many of these neutrinos would be observed at Earth.

However, the number observed was only a fraction of what was predicted, indicating that the rate of nucleosynthesis within the solar core was also much lower than expected - indeed, too low to keep the Sun in equilibrium against the force of gravity.

The simplest solution to this problem was that the models of the solar interior were wrong, and hence the rate of fusion required (and therefore the number of electron neutrinos produced) were in fact lower than predicted. However, by the mid-1990s, helioseismic measurements were able to place stringent constrains on the core sound-speed, indicating that the models of the solar interior were correct.

This led to the shocking and more fundamental conclusion that modifications to the Standard Model of particle physics were needed: neutrino oscillations between mass and flavour eigenstates during propagation explained the electron neutrino deficit; but required that the neutrino flavours themselves have differing, non-zero masses. This was an entirely fundamental result which would have been far harder to arrive at without helioseismology.

Today, the Sun is monitored seismically on a continuous basis using a mixture of ground-based telescope arrays (such as SONG;7 Grundahl et al. (2008)), and orbiting satellites such as the Solar and Heliospheric Observatory (SOHO, Liang et al. (2018)). The radial sound speed profile, averaged over latitude and longitude, is well known - the standard 'Model S' from the mid-1990s constrained the sound-speed to within 1% (Christensen-Dalsgaard et al. 1996). However, dynamic and three-dimensional effects are arguably more important on the Sun than the Earth, as fluid flows impose background velocity fields (e.g. Jørgensen et al. (2018)), solar cycles induce modulations in the helioseismic frequency spectrum spectrum (Lin

and Chou 2018), and the Sun's differential rotation makes latitudinal variations significant (e.g. Proxauf et al. (2020)).

The need for more localised measurements, which do not consider only global-scale free oscillation behaviours, has led to the development of 'time-distance helioseismology' Duvall et al. (1993). Time-distance helioseismology is more akin to traditional terrestrial seismology, in that it considers individual wave-packets on a local scale. The closest analogue on Earth is ambient noise interferometry (which was actually developed *from* time-distance helioseismology). Here, cross-correlations of fluctuations at different points on the solar surface are used to infer the properties of the background medium along the intermediary path. Such efforts are of interest in that local-scale phenomena can be studied in isolation, rather than through their effects on the global frequency spectrum. This is particularly of interest when studying localised phenomena like sunspots (Zhao et al. 2001), their associated magnetic fields (Zhao and Kosovichev 2006), and near-surface flows (Duvall and Gizon 2000).

The reader may well now be wondering how all these different parts of the Solar System fit together into one thesis, so it is worth remarking here that it all began with attempts to build a new model for use in time-distance helioseismology. In undertaking this work (the results of which are discussed in Chapter 6), it became clear that the methodological developments required to simulate acoustic wave propagation in a solar-type, fluid sphere were also applicable to other sorts of fluid (or indeed coupled solid-fluid) spheres, like the Earth's oceans (Chapter 2) and Mars' atmosphere (Chapter 4).

1.2.6 Other stars

Before moving on to discuss some of the general methodologies associated with seismic modelling, let us touch upon the seismic-like behaviour of objects far beyond the outer edge of our Solar System - other stars. Ironically, the first detections of pulsational behaviour in other stars were made long before the detection of seismic

phenomena on the Moon, Mars, or the Sun, though of course the behaviour was not understood physically or termed 'asteroseismology' until much later.

1.2.6.1 Mechanisms

In 1784, the star Delta Cephei was observed to be vary in magnitude over a period of five or so days, with a slow decrease in brightness preceeding a sharp increase (Goodricke 1786). At the time, it was entirely unclear whether this was an intrinsic phenomenon caused by changes in the star itself, or something extrinsic.

Although many further 'Cepheid' variables were identified, it was not until 1908 that a mathematical relationship explaining the oscillations in the stellar light curves was developed. Eddington (1918) recognised that the variations in brightness were linked to heat-driven pulsations in the stars' atmospheres, the heating and cooling involved therein leading to the star's entire surface expanding and contracting, and hence brightening and dimming. Although the mechanisms driving seismic behaviour in the Cepheid variables are entirely unlike those driving terrestrial, lunar, or even solar seismicity, they did set the scene for early studies in 'asteroseismology'.

Today, many different classes of asteroseismically-active star are known, and studies use both ground-based telescopes (Kjeldsen 2003) and space-based exoplanet-hunter observatories, such as Kepler (Gilliland et al. 2010) and TESS (Campante et al. 2016).

Some such stars' modes are excited by Cepheid-like 'Eddington' (or 'κ') valves, others because of solar-like convective turbulence. Other excitation sources include tidal forcings in multiple star systems, or convective blocking in stars with particularly steep gradients at the base of the convective zone (Aerts et al. 2010). Because the surfaces of other stars are not generally angularly resolvable, the use of whole-surface Dopplergrams as done for the Sun is impossible. Rather, changes in the star's overall brightness over time are tracked. Most stars oscillate in more than one mode, meaning that the changes in brightness over time are a result of the superposition of multiple different oscillation modes, some brightening, some dimming.

1.2.6.2 Observations

Long time-series are therefore required to decompose a light-curve into a series of distinct frequencies, which in theory act as a distinct 'fingerprint' for the class of star in question. However, the picture is somewhat complicated by the fact that the modes of oscillation displayed by a particular star depend on many other parameters. At a given position on the Herztprung-Russell diagram (or equivalently, for a particular temperature-luminosity combination), it is generally possible to predict in advance which excitation mechanisms will be active (if any), and which pulsations (pressure-restored, gravity-restored, or mixed) will be supported, if any (Aerts et al. 2010). As an example, stars with solar-like luminosities and surface temperatures (and hence solar-like masses and surface gravities) generally display convective-turbulence induced acoustic oscillations.

However, the oscillation spectrum may be altered by many variables, including departures from sphericity (i.e. a rapid rotation rate), the presence of a strong magnetic field, chemical peculiarity, or evolutionary status (and hence structure). Such variables may also change the lineshapes of individual modes as well as the global spectrum, for example through splitting singlet peaks. As such, inversions in asteroseismology are in some ways more complex than they are in the terrestrial case (where, for example, the age of the planet is not a particularly relevant variable). Nonetheless asteroseismology has proved to be a powerful investigative tool for studying a whole host of astrophysical problems, from white dwarfs (e.g. Winget and Kepler (2008)) through to blue supergiants (Bowman et al. 2019).

Chapter 7 of this thesis is concerned with the asteroseismic properties of the Lambda Boötis class of stars. Given that these are relatively niche objects with a nomenclature and evolutionary history quite unlike the other worlds discussed thus far, we will confine a description of them to that chapter. To avoid confusion or direct comparison to the solar results presented in Chapter 6, however, let us be clear that they are of an entirely different spectral class (late-B to early-F versus mid-G), and much closer to the Delta Scuti, Gamma Doradus, and Subdwarf B pulsators than to the Sun - that is, very much not-solar-like in their oscillations.

Our tour of the objects relevant to this thesis is now complete. It is worth, before we go any further, explaining what the need for forward methods (applied to a different setting in each of Chapters 2-7) is in general.

1.3 Inverse and forward methods

Whilst the forward and inverse methods we have already touched upon are intricately linked to each other and cannot really be decoupled in their application, they are often posed from different perspectives - as such, let us take a minute to outline how they differ, but also how they fit together.

1.3.1 Inverse methods

The primary goal of the types of seismic investigations described above is clearly to elucidate the internal structure, properties, and behaviour of the world in question; most often for the purpose of understanding how and why it came to be the way that it is. This use of seismic observations to infer the physical parameters of a system is known as the 'inverse problem' (see Tarantola (2006) for an excellent discussion).

The inverse problem requires a model within which we may contextualise and interpret observations: the seismic displacement from a quake, as measured at the Earth's surface, can only be used to study the nature of the core-mantle boundary if we have a preliminary understanding of both the Earth's structure and of how different structural parameters in the mantle affect the observed properties of the waveform.

Unfortunately, inversions are mathematically challenging in the best of cases, requiring extremely high observational precision if a unique and robust inversion is to be performed. Seismic observations, with their sparsity across most of the Earth's surface, significant noise levels, and generally non-repeating behaviour are considerably far from any such 'best case', This is true is true for all the cases described in Sec. 1.1; even the Sun with its near-continuous data coverage going back several decades.

In introducing the forward problem, the main topic of this thesis, it is worth considering why robust seismic inversions are so hard to undertake. Even in an idealised seismic inversion scenario (e.g. one with a simple structural model and excellent, noiseless data coverage), there are many free or poorly-constrained parameters which have an effect on the observations. These include the properties of the background medium (sound speed, shear-wave speed, density, and so on at every point in space - and in some cases, time as well), the properties of the source (depth, mechanism, magnitude, for example), and of the receiver-system (which may be changing in time, for example as the temperature or wind speed vary). Even if a unique and well-behaved solution exists, finding it through numerical optimisation can be extremely challenging and computationally intensive; with many local minima traps along the way.

Regardless, even with unlimited computational resources undertaking a robust inversion would still be a challenge. This is on account of the fact that inverse problems in seismology are generally not well-posed and well-conditioned, meaning that a solution need not exist or be unique, and may vary not vary continuously with variations in input parameters (the condition of non-existence of a solution is easier to conceptualise when the effects of noise are considered).

1.3.2 Forward methods

Whilst the challenges of undertaking inversions are fundamental, some traction can be gained against them through seeking solutions to the 'forward problem'. Here, rather than using observations of seismic waves at the surface of the Earth or Sun to make inferences about their internal structure and source properties, a prediction is made of the seismic waves which would be expected at a given point in space and time from a particular structural model and source. This can be achieved by finding solutions to equations of motion for a particular set of imposed conditions. Solutions can be compared to observations, and discrepancies are generally indicative of unrepresentative models of the source or structure, or

of not properly considering the relevant terms in the mathematical formulation of the problem (for example, erroneously neglecting gravity).

Forward and inverse methods clearly work in conjunction with each other, for example running an ensemble of simulations for a particular portion of the structural parameter space, and examining whether observed waveforms are matched better by one particular choice as compared to the others.

Of course, this relies on having observations which predictions may be compared against - in Chapters 3, 4, and 6, forward models were used purely for their predictive power, but are still very much useful. In both these cases the aim was to use forward modelling as a tool to design and justify an observational campaign for which no prior observations of this type had been made.

The remainder of this thesis will focus on forward modelling: theory, method development, and results; in varying proportions in the different chapters.

1.4 Thesis particulars

1.4.1 Roadmap

In theses focussed on modelling, it is traditional to give an overview in the introduction of the equations being solved and the computational tools used to achieve this. For reasons that will become apparent in subsequent chapters, this is not appropriate here - and indeed, would be enormously confusing. Whilst the basic principles behind finding solutions to the wave equation are the same regardless of whether we consider the Earth or a Lambda Boötis star, the mathematical and computational tools used are very different.

1.4.2 Structure

The results chapters of this thesis cover topics as follows:

- Chapter 2: modelling the propagation of seismic waves in the Earth's oceans
- Chapter 3: contributions to the InSight blind test for martian seismicity

- Chapter 4: modelling the seismic signatures from the arrival of the Perseverance Rover at Mars
- Chapter 5: results from InSight's attempts to detect Perseverance's arrival
- Chapter 6: developing new simulation methods for time-distance helioseismology
- Chapter 7: asteroseismic investigations of Lambda Boötis stars

In Chapters 2-6, we are concerned with modelling propagating seismic waves, which are by their nature transient and localised - or in other words, time-evolving, but within a static background model. In Chapter 7, we will explore free oscillations (i.e. those involving standing waves), whose structures are constant on oscillatory timescales, but crucially in a scenario where the background structural model is temporally evolving on a much longer timescale.

In Chapters 2-6 the key outputs will be considered in the time domain (i.e. examining the time-series of surface displacement at specific points in space), whereas in Chapter 7 outputs are considered in the frequency domain (i.e. identifying the frequencies of global oscillation modes, which change in time as the stars evolve). These differing approaches result in the theory, computational modelling, and results being presented very differently in each case.

Rather than attempting to tie together all the differences between these two scenarios in a single introduction, I will instead detail the mathematical and computational tools used in Chapters 2-6 at the start of Chapter 2 (with further expansion in later sections as needed), and those relevant to Chapter 7 in Chapter 7.

It may become obvious in reading on that this thesis is in an 'integrated' format, in that each section has already been, was hoped to have been, or will be published as an independent paper. Each chapter opens with a description of where in the publication process I am, and with a statement of contributions. Where supplemental work post-publication has been conducted, it is appended at the end of each chapter in an epilogue.

You may note in this thesis some attempts at 'humour', with varying degrees of success. It is mostly self-deprecating, and probably none too witty - so please interpret it as a way of coping with having to write a thesis during the absolute mayhem which was 2020-21.

And now, onward.

An ocean traveller has even more vividly the impression that the ocean is made of waves than that it is made of water.

— Sir Arthur Stanley Eddington

2 The Oceans

Contents

2.1 Preface 22
2.2 Summary 22
2.3 Introduction 23
2.4 Methodology 27
2.4.1 Theory 28
2.4.2 Implementation in AxiSEM3D 31
2.5 Validation 33
2.5.1 Nomenclature 33
2.5.2 Benchmark parameters 34
2.5.3 Benchmark results 35
2.6 Evaluation of ocean loading 38
2.7 Bathymetry in a global ocean context 43
2.7.1 Bathymetric implementation 43
2.7.2 The effects of bathymetry on waveforms 45
2.8 Propagation through the water column 52
2.9 The effects of the ocean on observed data 53
2.10 Discussion 56
2.10.1 Computational cost 56
2.10.2 Cost scaling 58
2.10.3 Potential cost savings 59
2.10.4 Limitations 60
2.10.5 Extensions in a global ocean model 61
2.10.6 Development of a localised ocean formulation 63
2.11 Summary 64
2.12 Epilogue 65
2.12.1 Microseismic noise: an introduction 65
2.12.2 Noise in the Indian Ocean 65
2.12.3 Results 66

2.1 Preface

This paper was published as Fernando et al. (2020a):

Fernando, B., Leng, K., & Nissen-Meyer, T. - **Oceanic high-frequency global seismic wave propagation with realistic bathymetry**, *Geophysical Journal International*, 2020.

It has the following DOI, at which address supplementary material may be found: *10.1093/gji/ggaa248*. Further animations are hosted at *https://bit.ly/ocean-animations*.

My contributions to this paper included: code modifications and testing, executing simulations, undertaking comparisons to data, and writing the manuscript. Assistance with code editing and debugging was received from Kuangdai Leng.

The acknowledgements to this publication include: David Al-Attar, Martin van Driel, Maria Tsekhmistrenko, Kasra Hosseini, Claudia Haindl and Thomas Garth for their technical help on this topic, as well as the manuscript's reviewers and editors: Wenbo Wu, Christian Boehm, and Sidao Ni.

2.2 Summary

- We present a new formulation of the AxiSEM3D code to simulate high-frequency seismic wave propagation in and under the oceans, with the aim of enabling more realistic seismic modelling to be undertaken.

- An evaluation of the ocean load formulation is shown, which we find breaks down at longer periods than was previously believed.

- Consideration of ocean bathymetry is also made, wherein we find that surface waveforms are significantly modified by seafloor topography. Observational data from the Pacific are better fitted by synthetics which include a global

ocean and topography compared to those which include topography on the solid surface alone.

- This methodology facilitates avenues for research into microseism generation studies, t-phase simulation, and other ocean seismology topics.

2.3 Introduction

Oceans cover more than 70% of the Earth's surface, and have complex and nuanced effects on the propagation of seismic waves through the planet. However, with the exception of isolated stations on remote islands, global seismometer distribution is enormously skewed toward being continent-based and so biased toward the northern hemisphere. Seismology in an oceanic context has thus been somewhat neglected as compared to its land-based counterpart, and the lack of comparably high-quality global data coverage from south of the Equator can be an impediment to seismic studies. However, in recent times the deployment of ocean-bottom (OBS) and floating seismometers, together with increases in the land-based instrument density along coastlines, has begun to alleviate this problem. As a consequence there are now many large datasets which are in need of interpretation. Their complexity is such that observation alone is unlikely to prove a route to thoroughly understanding them; instead, we must make use of the synergy between observation and forward modelling.

Seismologically speaking the effects of the oceans are twofold, and are particularly pronounced at high frequencies ($\leq$10 s, where sensitivity to crustal-scale features becomes signficant). Firstly, new sources of seismic signal and noise are introduced. These include infrasounds, such as those associated with underwater volcanoes (Bohnenstiehl et al. 2013), as well as the ocean microseism, which is dynamically generated in the water column (Longuet-Higgins (1950), Ardhuin et al. (2015)). Secondly, a new set of seismic phases which couple across the solid-fluid seafloor interface become supported. pWP (an upgoing p-wave reflected off the ocean surface before being re-transmitted into the solid earth), and t-phases which are

trapped in the SOFAR channel (Tolstoy and Ewing 1950), are examples of such. These are inherently high-frequency effects: water column reverberations like pWP have a characteristic period of a few seconds, whilst t-phases are observed at frequencies of a few Hertz and higher.

In the microseism case, these are particularly useful for mapping tomographic structures (e.g. Shapiro et al. (2005), Basini et al. (2013)), tracking storms (Davy et al. 2014) and making inference about climatic trends (e.g. Grob et al. (2011), Stutzmann et al. (2009)). Use of the pWP phase can substantially improve source localisation in Wadati-Benioff zones (Engdahl et al. 1998), whilst t-phases may be detected in-situ in the water column, far from land, and thus are useful for tomographic inversion in otherwise sparsely sampled areas (Blackman et al. 1995). Additionally, as t-phases experience a reduced degree of geometric spreading (cylindrical, $\sim \frac{1}{r}$, rather than spherical, $\sim \frac{1}{r^2}$), they may offer greater sensitivity to more distant or smaller seismic events than would be obtained from the corresponding P or S-waves at the same epicentral distance; furthermore, the lower and better-constrained sound speed in water can yield tighter constraints on source location than are possible from signals in the solid Earth alone (Dziak et al. 2004). Other potential uses of hydroacoustic phases include early-warning systems for tsunamis (Sasorova et al. (2005), Lay et al. (2019)), acoustic thermometry for remote monitoring of long-term changes in ocean temperature (Dushaw et al. 1999), detection of glacial and ice-calving events (Chapp et al. 2005), as well as monitoring of cetacean populations (e.g. McDonald et al. (1995), Dréo et al. (2019)) and for illicit nuclear tests in violation of the Comprehensive Nuclear Test Ban Treaty (Mitchell (2002), Okal (2008)).

However, there are numerous areas of basic physics which remain to be explored within the context of ocean seismology. These include how P-waves which are refracted to near-normal angles at the seafloor are converted to horizontally travelling t-waves in the SOFAR channel (de Groot-Hedlin and Orcutt 2001), how primary microseismic noise with a transverse component is generated (Nishida and Takagi 2016), and how best to isolate or account for water column reverberations, especially over areas rough bathymetry (Blackman et al. 1995). The latter of these can

significantly complicate understanding of earthquake dynamics (Yue et al. (2017), Qian et al. (2019)), especially of the largest earthquakes which occur at subduction zones with rough bathymetry (Lay and Rhode (2019),Wu et al. (2020)).

Such puzzles are unlikely to be resolved by observation alone, given that comprehensive seismic data from the oceans' depths are still sparse as compared to those from on land. Making use of the interplay between observation and modelling is thus one route to exploring such questions. Multiple tools for exploring the modelling side of these problems exist, including SPECFEM3D Cartesian (Peter et al. 2011), the hybrid Direct Solution-Spectral Element Method of Wu et al. (2018), and Salvus (Afanasiev et al. 2019).

However, modelling in an oceanic context is not without its own challenges, even at low frequencies. Forward modelling in seismology can be done through finding numerical solutions to the weak form of the equations of motion, which in solid media implicitly include the Neumann-type traction-free surface boundary condition (Igel 2016). In the case of a fluid surface layer, a Dirichlet boundary condition is needed, and in our method must be explicitly solved.

Furthermore, the low p-wave speeds in the oceans ($\approx$1450 ms^{-1}) yield acoustic waves with shorter wavelengths than in the underlying crust, thus necessitating the use of small elements. These in turn require shorter time steps to ensure sufficient temporal sampling in explicitly time-stepped methods, increasing the computational expense. The implementation of realistic seafloor bathymetry is also especially difficult when combined with the need to ensure that a hexahedral mesh remains conformal, with elements which are not overly deformed or skewed.

These challenges have led to a number of simplifications being made in the current norm of seismic modelling, such that the simulations become computationally tractable. These have involved either approximating the ocean as a mass loading in a global formulation (where we consider 'global' to mean that the effects of the Earth's sphericity are important), or introducing a realistic ocean layer but restricting the simulation to a local geometry or to fully axisymmetric domains with flat seafloors.

Komatitsch and Tromp (2002) adopt the first approach, wherein the bulk modulus at the free surface is modified to account for the weight of the water column without explicitly including the ocean in the mesh. This 'ocean loading' formulation (also known as the 'water column approximation') is valid only at long periods (where the wavelength of the seismic waves is significantly larger than the ocean depth), and does not reproduce hydroacoustic phases. It does however demonstrate that the ocean delays the arrival of the Rayleigh wave train and considerably changes its dispersion characteristic, whilst leaving the Love waves unchanged. Zhou et al. (2016) present an initial evaluation of this approximation for two specific phases, as discussed further in Sec. 2.6.

Conversely, Cristini and Komatitsch (2012), Bottero et al. (2016) and Mazoyer et al. (2013) use local-scale simulations in SPECFEM3D, which reproduce hydroacoustic phases but are restricted to axisymmetric '2.5D' formulations. Such simulations reproduce in-plane scattering induced by local-scale features like seamounts, but are not suitable for use at planetary scales or for simulating processes where out-of-plane scattering is thought to be important, such as Love wave generation at the seafloor. The need for global seismic modelling with realistic ocean layers and which supports their associated hydroacoustic phases at the high frequencies used in modern seismology is, therefore, clear.

In this paper, we present such an implementation. Our work is based on the spectral-element methodology AxiSEM (Nissen-Meyer et al. 2014a), which was expanded to include full 3D structures (Leng et al. 2016) with undulating discontinuities and ellipticity (Leng et al. 2019). A realistic ocean layer in AxiSEM3D is now introduced, with a pressure-free surface boundary condition. A benchmark against the code YSpec (Al-Attar and Woodhouse 2008) is shown for the flat seafloor case, which introduces new water-depth seismic phases, as well as convergence to the ocean loading formulation at long periods. We show that in a global context, substantial differences between the ocean loading and realistic ocean formulations become apparent below dominant periods of $\sim$20 s. We also evaluate the effects of bathymetry on seismic waveforms, with a particular focus on Rayleigh waves.

A clear limitation of this implementation is that it requires a consistent surface boundary condition along any line of azimuth; that is, an ocean must be either cover the planet entirely (the 'global' ocean) or lie in a ring along the surface (what we term a 'local' ocean, which is effectively donut-shaped). This paper deals with only the first of these cases, and though with careful choice of simulation geometry realistic simulations may still be performed, an arbitrary 'patched' ocean, with some solid surface areas and other fluid ones cannot be supported at present on a global scale in AxiSEM3D.

Thus, this method is in no way presented as a complete replacement for synthetics generated in a solid-surface implementation, but rather as a complimentary approach which is widely applicable given the majority of our planet's surface covered by water, and the resulting implication that an arbitrary source-receiver path at teleseismic distances is more likely than not to include an oceanic section. Such an approach is also justified on observational grounds - by undertaking a comparison to recorded data from remote island stations (which provide valuable global coverage in sparsely sampled areas and by their nature are strongly affected by the presence of an ocean), we give examples where a better match to synthetics occurs when a fluid ocean is included than when it is not.

Such a methodology should enable higher-frequency simulations to be carried out in an oceanic seismology context, with arbitrary structural complexity in the underlying solid Earth. In this paper we restrict further detailed discussion to examination of bathymetric effects and water-depth phases; however extensions of our methodology to other fluid seismology contexts (the atmosphere or localised but axisymmetric oceans) are also possible.

2.4 Methodology

This section describes the elastodynamic theory of wave propagation in an Earth model which has both a fluid surface ocean and an undulating seafloor, as well as the implementation of such a setup in AxiSEM3D. As compared to Leng et al.

(2016) and Leng et al. (2019), the main difference is in the stress-free (and hence pressure-free) boundary condition on the fluid ocean surface.

2.4.1 Theory

We first consider a spherical Earth model that consists of a solid domain Ω_S (with density ρ and elasticity tensor $\mathbf{C}$) and a fluid domain Ω_F (also with density ρ and bulk modulus κ). These two domains can be separated by several solid-fluid interfaces, collectively denoted Σ, such as the ocean floor, the core-mantle boundary and the inner core boundary.

In the solid domain, the weak formulation of the equations of motion (ignoring attenuation and long-period effects such as gravitation and rotation) may be written as

$$\begin{aligned} &\int_{\Omega_S} \rho \mathbf{w} \cdot \partial_t^2 \mathbf{u} \, d^3\mathbf{r} + \int_{\Omega_S} \nabla \mathbf{w} : \mathbf{C} : \nabla \mathbf{u} \, d^3\mathbf{r} \\ =& \int_{\Omega_S} \mathbf{w} \cdot \mathbf{f} \, d^3\mathbf{r} - \int_{\Sigma} \partial_t^2 \chi \, \hat{\mathbf{n}} \cdot \mathbf{w} d^2\mathbf{r}, \end{aligned} \tag{2.1}$$

where $\hat{\mathbf{n}}$ denotes the outward-pointing normal of a solid-fluid interface, $\mathbf{f}$ is a body force source, $\mathbf{u}$ is the displacement vector, and $\mathbf{w}$ is an arbitrary, vector-valued test function. In the fluid domain, we define a scalar potential χ such that $\mathbf{u} = \rho^{-1}\nabla\chi$ in Ω_F, which is a descriptor for the dynamic pressure through Ω_F as $\partial_t^2\chi = -P$.

$$\begin{aligned} &\int_{\Omega_F} \kappa^{-1} w \, \partial_t^2 \chi \, d^3\mathbf{r} + \int_{\Omega_F} \rho^{-1} \nabla w \cdot \mathbf{I} \cdot \nabla \chi \, d^3\mathbf{r} \\ =& \int_{\Sigma} \mathbf{u} \cdot w \hat{\mathbf{n}} \; d^2\mathbf{r}. \end{aligned} \tag{2.2}$$

Here, w is an arbitrary scalar-valued test function. The rank-two identity tensor $\mathbf{I}$, which appears to be somewhat redundant here, is included to enable more natural extension to an aspherical earth model with undulating boundaries in eq. (2.7). Note that Ω_S and Ω_F are coupled by the two surface integrals over Σ.

Next, we consider the stress-free boundary condition on the Earth's surface. Without a fluid ocean, this takes the form of a Neumann boundary condition where $\mathbf{C} : \nabla\mathbf{u} \cdot \hat{\mathbf{n}} = \mathbf{0}$, which is automatically satisfied by eq. (2.1). However, in the fluid ocean (where the stress-free boundary condition becomes a pressure-free one as the off-diagonal elements of the stress tensor vanish), it becomes a Dirichlet boundary

condition. In this case, $P(\mathbf{x},t) = -\partial_t^2\chi(\mathbf{x},t) = 0$, which is not automatically satisfied by eq. (2.2) and must be explicitly prescribed. From the requirement for a steady-state solution it follows that $\partial_t\chi$ and χ must also be identically zero at all times on the surface; these conditions are thus imposed at the boundary at each time step. Such boundary conditions have been used previously in a seismological context, for example by Peter et al. (2011) and Bottero et al. (2016).

Such a pressure-free boundary condition is perfectly reflecting in an analytical sense, which is justified physically on the grounds that the acoustic impedance of air is very much less than that of water so coupling from the ocean surface into the atmosphere can reasonably be neglected (though, if desired, it in can be accounted for in our formulation so long as the atmosphere is included in the mesh). At the seafloor boundary, the continuity of traction and normal velocity are implemented in the same way as at the outer core boundary.

When 3D bathymetry is incorporated, eqs. (2.1) and (2.2) remain valid, but Ω_S, Ω_F and Σ become aspherical. As AxiSEM3D requires an axisymmetric computational domain which is spherical at a global scale, these equations cannot then be solved directly. Leng et al. (2019) describes the implementation of the particle relabelling transformation of Al-Attar and Crawford (2016) to handle such interface undulations.

Given a reference spherical configuration $\tilde{\Omega}$, and a deformed configuration Ω which is homeomorphic to $\tilde{\Omega}$, the radial coordinate of a particle at $\mathbf{r}$ in $\tilde{\Omega}$ is shifted along $\hat{\mathbf{r}}$ by an amount $\tau(\mathbf{r})$ in $\tilde{\Omega}$ such that

$$\xi(\mathbf{r}) = \mathbf{r} + \tau(\mathbf{r})\hat{\mathbf{r}}. \tag{2.3}$$

Note that this transformation only shifts the boundary inward or outward along $\hat{\mathbf{r}}$, which is commensurate with the fact that the boundaries in question (e.g. seafloor bathymetry) have different radii at different positions, but do not need to be shifted laterally.

ξ, where $\xi : \tilde{\Omega} \to \Omega$ (and hence also $\xi : \tilde{\Sigma} \to \Sigma$), represents the undulation mapping, which is a kinematically permissible deformation with associated deformation gradient $\mathbf{F}$. In spherical coordinates, where $\mathbf{F} = \mathbf{F}(r, \theta, \phi)$, the deformation gradient can be expressed as:

$$\mathbf{F}\left(\tau, \mathbf{r}\right) = \left[\nabla\left(\tau\hat{\mathbf{r}}\right)\right]^{\mathrm{T}} = \begin{pmatrix} \hat{\mathbf{r}} & \hat{\theta} & \hat{\phi} \end{pmatrix} \begin{pmatrix} \partial_r \tau & \dfrac{\partial_\theta \tau}{r} & \dfrac{\partial_\phi \tau}{r \sin\theta} \\ 0 & \dfrac{\tau}{r} & 0 \\ 0 & 0 & \dfrac{\tau}{r} \end{pmatrix} \begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\theta} \\ \hat{\phi} \end{pmatrix}, \tag{2.4}$$

whilst the associated Jacobian for this transformation may be expressed as:

$$\mathbf{J}(\tau, \mathbf{r}) = \mathbf{I} + \mathbf{F}(\tau, \mathbf{r}). \tag{2.5}$$

The weak form of the undulating, three-dimensional model can thus be established in the reference configuration $\tilde{\Omega}$ in the solid (corresponding to eq. (2.1)) as

$$\begin{aligned} &\int_{\tilde{\Omega}_S} \tilde{\rho}\tilde{\mathbf{w}} \cdot \partial_t^2 \tilde{\mathbf{u}}\, d^3\mathbf{r} + \int_{\tilde{\Omega}_S} \nabla\tilde{\mathbf{w}} : \tilde{\mathbf{C}} : \nabla\tilde{\mathbf{u}}\, d^3\mathbf{r} \\ &= \int_{\tilde{\Omega}_S} \tilde{\mathbf{w}} \cdot \tilde{\mathbf{f}}\, d^3\mathbf{r} - \int_{\tilde{\Sigma}} \partial_t^2 \tilde{\chi}\tilde{\mathbf{n}} \cdot \tilde{\mathbf{w}} d^2\mathbf{r}, \end{aligned} \tag{2.6}$$

whilst in the fluid domain (corresponding to eq. (2.2)) the formulation becomes

$$\begin{aligned} &\int_{\tilde{\Omega}_F} \tilde{\kappa}^{-1}\tilde{w}\, \partial_t^2 \tilde{\chi}\, d^3\mathbf{r} + \int_{\tilde{\Omega}_F} \tilde{\rho}^{-1} \nabla\tilde{w} \cdot \tilde{\mathbf{I}} \cdot \nabla\tilde{\chi}\, d^3\mathbf{r} \\ &= \int_{\tilde{\Sigma}} \tilde{\mathbf{u}} \cdot \tilde{w}\tilde{\mathbf{n}} d^2\mathbf{r}, \end{aligned} \tag{2.7}$$

where the equivalent material parameters are given by

$$\tilde{\rho}(\mathbf{r}) = \rho(\xi(\mathbf{r}))|\mathbf{J}(\tau, \mathbf{r})|, \tag{2.8}$$

$$\tilde{\kappa}(\mathbf{r}) = \kappa(\xi(\mathbf{r}))|\mathbf{J}(\tau, \mathbf{r})|, \tag{2.9}$$

$$\tilde{\mathbf{C}}(\mathbf{r}) = \mathbf{J}^{-1}(\tau, \mathbf{r}) \cdot \mathbf{C}(\xi(\mathbf{r})) \cdot \mathbf{J}^{T}(\tau, \mathbf{r})|\mathbf{J}(\tau, \mathbf{r})|, \tag{2.10}$$

and the equivalent source by

$$\tilde{f}(\mathbf{r}) = f(\xi(\mathbf{r}))|\mathbf{J}(\tau, \mathbf{r})|, \tag{2.11}$$

and $\tilde{\mathbf{I}}$ and $\tilde{\mathbf{n}}$ respectively by

$$\tilde{\mathbf{I}} = \mathbf{J}^{-1}(\tau, \mathbf{r}) \cdot \mathbf{J}^{-T}(\tau, \mathbf{r}), \tag{2.12}$$

$$\tilde{\mathbf{n}}(\mathbf{r}) = \mathbf{J}^{-T}(\tau, \mathbf{r}) \cdot \hat{\mathbf{n}}(\xi(\mathbf{r}))|\mathbf{J}(\tau, \mathbf{r})|. \tag{2.13}$$

The solutions to eqs. (2.6) and (2.7), in $\mathbf{u}$ and χ respectively, are related to the solutions of eqs. (2.1) and (2.2) by

$$\tilde{\mathbf{u}}(\mathbf{r}, t) = \mathbf{u}(\xi(\mathbf{r}, t)), \tag{2.14}$$

and

$$\tilde{\chi}(\mathbf{r}, t) = \chi(\xi(\mathbf{r}, t)). \tag{2.15}$$

This formulation is sufficient to describe all variables needed to undertake the task of finding solutions to the elastodynamic equations of motion in an aspherical Earth which possesses undulating boundaries.

2.4.2 Implementation in AxiSEM3D

Eqs. (2.1) and (2.2) may be solved in a fully three-dimensional mesh, as is the case in SPECFEM (Komatitsch and Tromp 2002). AxiSEM3D, however, makes use of the axisymmetric formulation of Nissen-Meyer et al. (2007). The smoothness of most global tomographic models (which have significantly smaller gradients in seismic parameters in the two lateral directions than along their radius) means that the global seismic wavefield is, in general, also significantly smoother in the azimuthal direction than it is in a meridional plane. As such the use of a pseudospectral parameterisation in the azimuthal direction which exploits this wavefield smoothness and accounts for the periodicity of the solution over the interval $[0, 2\pi)$ can offer a substantial computational speed-up as compared to methods like SPECFEM3D Globe which rely on a full 'cubed-sphere' mesh. This speedup scales with increasing frequency - at 10 s, AxiSEM3D is approximately two orders of magnitude faster (Leng et al. 2019).

The choice of the highest term which must be included in the azimuthal Fourier expansion (N_u) depends on the complexity of the model. In a radially symmetric (1D) model, this representation becomes analytic for a second-order (quadropolar) moment tensor when $N_u = 2$. For more complex models with significant 3D structures, N_u can be increased as needed to capture azimuthal structures.

Working in a cylindrical basis $(\hat{\mathbf{s}}, \hat{\phi}, \hat{\mathbf{z}} = \hat{\mathbf{g}}_\mathbf{1}, \hat{\mathbf{g}}_\mathbf{2}, \hat{\mathbf{g}}_\mathbf{3})$, the velocity $\mathbf{u}$ and potential χ may be expressed as:

$$\mathbf{u} = \hat{\mathbf{g}}_\mathbf{i}(\phi)\, u_i(s, \phi, z; t) = \sum_{|\alpha| \leq N_u} u_i^\alpha(s, z; t) \Psi^\alpha(\phi) \hat{\mathbf{g}}_\mathbf{i}(\phi) \tag{2.16}$$

and

$$\chi = \chi_i(s, \phi, z; t) = \sum_{|\beta| \leq N_u} \chi_i^\beta(s, z; t) \Psi^\beta(\phi), \tag{2.17}$$

respectively, where $\Psi^\alpha(\phi) = e^{\sqrt{-1}\alpha\phi}$ are the terms with order $\alpha \leq |N_u|$, and similarly for $\Psi^\beta(\phi)$.

As per eqs. (2.16) and (2.17), the decomposition of the solution into Fourier modes means that a two-dimensional mesh may be used (see Nissen-Meyer et al. (2007) and Leng et al. (2016)). Three-dimensional structural complexity may be added as required, subject to an increase in the Fourier order N_u, whilst topography may be implemented on structural discontinuities (e.g. the seafloor and Moho) through the particle relabelling technique described in Sec. 2.4.1.

Thus, full three-dimensional structures may be reliably accounted for in simulations involving AxiSEM3D. An example of an explicitly meshed surface fluid layer is shown in Fig. 2.1. Note that we use the terms 'bathymetry' and 'seafloor topography' interchangeably to describe the variation in ocean depth with position.

This enables inclusion of heterogeneities, anisotropy (van Driel and Nissen-Meyer 2014b) and attenuation (van Driel and Nissen-Meyer 2014a). In this paper, all simulations were performed without ellipticity, rotation or gravitation as we use only intermediate periods (up to 100 s), and with attenuation enabled in the solid earth only due to the comparatively high Q value of acoustic wave propagation in water. Simulations are performed on an anisotropic PREM mesh (6368 km radius without ocean, and 6371 km with ocean). The ocean is modelled as a 3 km thick homogeneous layer with $v_p = 1450$ ms^{-1} and $\rho = 1040$ kgm^{-3}, though this may be arbitrarily modified if required. Two elements per wavelength are used in all AxiSEM3D simulations.

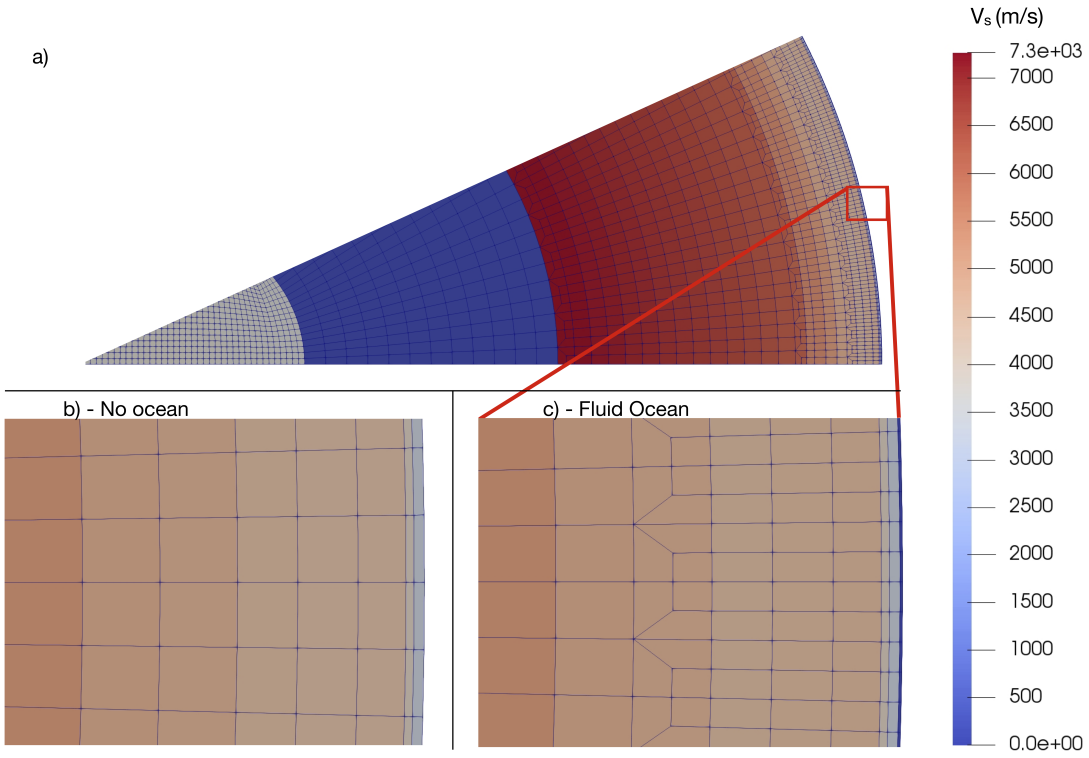

Figure 2.1: A two-dimensional mesh as used in AxiSEM3D, with the colour scale corresponding to the value of V_s. Panel a) shows the entire mesh in the ocean case, from surface to core, whilst panel b) shows a detail of the surface layers in the no ocean case and panel c) is the equivalent detail for the ocean containing mesh. The surface is leftmost in all cases. The differences between panels b) and c) are slight at this scale: the ocean mesh has a thin blue layer at its outer edge (where $V_s = 0$) and an extra mesh refinement layer midway through the displayed section to account for the velocity contrast at the seafloor, whilst the mesh without an ocean has neither. Note that in all regions of the mesh, the smallest non-zero velocity (i.e. V_s in the solid and V_p in the fluid) is the dominant determinant of the element size.

2.5 Validation

2.5.1 Nomenclature

For clarity, here we will briefly summarise the terms used to describe the different ocean configurations used in this paper. In Secs. 2.5 and 2.6, all models are spherically symmetric (i.e. use 1D seismic profiles only). 'No Ocean' refers to a PREM model of radius 6368 km with a solid surface, 'Ocean Load' is identical but accounts for the weight of a 3 km thick ocean layer without explicitly meshing it, whilst 'Fluid Ocean' indicates a PREM model with a global, explicitly meshed

3 km thick homogeneous ocean layer and outer radius 6371 km.

From Sec. 2.7 onward, non-spherically symmetric (fully 3D) models are used. The exact details are specified in Sec. 2.7, but in short, 'No Ocean' indicates a 6368 km radius model with undulation on the solid surface (surface topography), 'Fluid Ocean without Bathymetry' indicates a model with global ocean, radius 6371 km and no undulation on the seafloor (i.e. no bathymetry), and 'Fluid Ocean with Bathymetry' is otherwise identical except for the inclusion of seafloor undulation. As it is not the main topic of this paper, the ocean loading approximation is not considered beyond the end of Sec. 2.6.

2.5.2 Benchmark parameters

In order to verify the reliability of the AxiSEM3D simulations, we perform a benchmark against YSpec (Al-Attar and Woodhouse 2008). YSpec is a semi-analytical method which uses direct radial integration to compute full waveform synthetics, chosen specifically because it is not a spectral element method, and hence provides a robust benchmark. Unlike AxiSEM3D, YSpec does support calculations with full self-gravity, but these are not used here. For reproducibility, it should be noted that we edited the YSpec source code to remove an automatic ellipticity correction.

A radially symmetric model with an ocean layer is chosen as this is the most complex configuration for which a reference solution with a global ocean is available. It should be noted that we are not therefore undertaking a full, global benchmark with 3D structures (which in the fluid ocean case is not possible with open-source software at high frequency, and in the no ocean case was done by Leng et al. (2016)). Rather, we seek to confirm that the addition of a fluid ocean layer does not detract from the reliability of the synthetics.

The source parameters used are given in Table 2.1. In YSpec the maximum angular degree used is 22,000 and the maximum frequency is 600 mHz, whilst in AxiSEM3D the dominant simulation period is 2 seconds. Note we define the 'dominant period' as the minimum globally resolved period in the mesh, whilst

Table 2.1: Source Parameters. To enable exact reproducibility of our benchmark, we give the full parameters of the source used here.

Latitude	90°
Longitude	0 °
Depth	12 km (no ocean), 15 km (ocean)
Receiver Position	Surface (no ocean), Seafloor (ocean)
Source-Time Function	Heaviside (YSpec), Error (AxiSEM3D)
Moment Tensor	M_{rr} = 1 x 10^{26} dyne-cm

Komatitsch and Tromp (2002) define it as the corner frequency above which no energy is observed in the simulation. The stations in the benchmark are located at the surface in the no ocean case, and on the seafloor in the ocean case (i.e. at radius of 6368 km in both). In AxiSEM3D, the order of the Gauss-Lobato-Legendre quadrature is 4.

The source is represented through a near-instantaneous impulse, with a source time function which is a Heaviside in YSpec and a narrow error function (half-width 0.5 seconds) in AxiSEM3D. Although the two are extremely similar in shape, they are not identical - YSpec uses a true (instanteous-step) Heaviside function, whilst the error function in AxiSEM3D has some width. In order to retain numerical stability, the width of the error function cannot be decreased to zero. The remaining differences between these Heaviside and error source-time functions are accounted for by a small (∼1 second) temporal shift and convolution of both traces with a 2 second half-width Gaussian function, as the implementations of the source-time functions in the codes themselves cannot be made any closer to each other. The result from these simulations is that a Green's function is extracted in both cases, with the output Butterworth-Bandpass filtered (filter order 4) between 2 and 100 seconds.

2.5.3 Benchmark results

The modifications induced by the addition of the ocean layer are apparent in Fig. 2.2, whilst Fig. 2.3 shows the time-distance record section of the observed waveforms in the fluid ocean case. The Rayleigh wave train is substantially extended, lasting many times its original length, whilst the peak surface wave amplitude is decreased

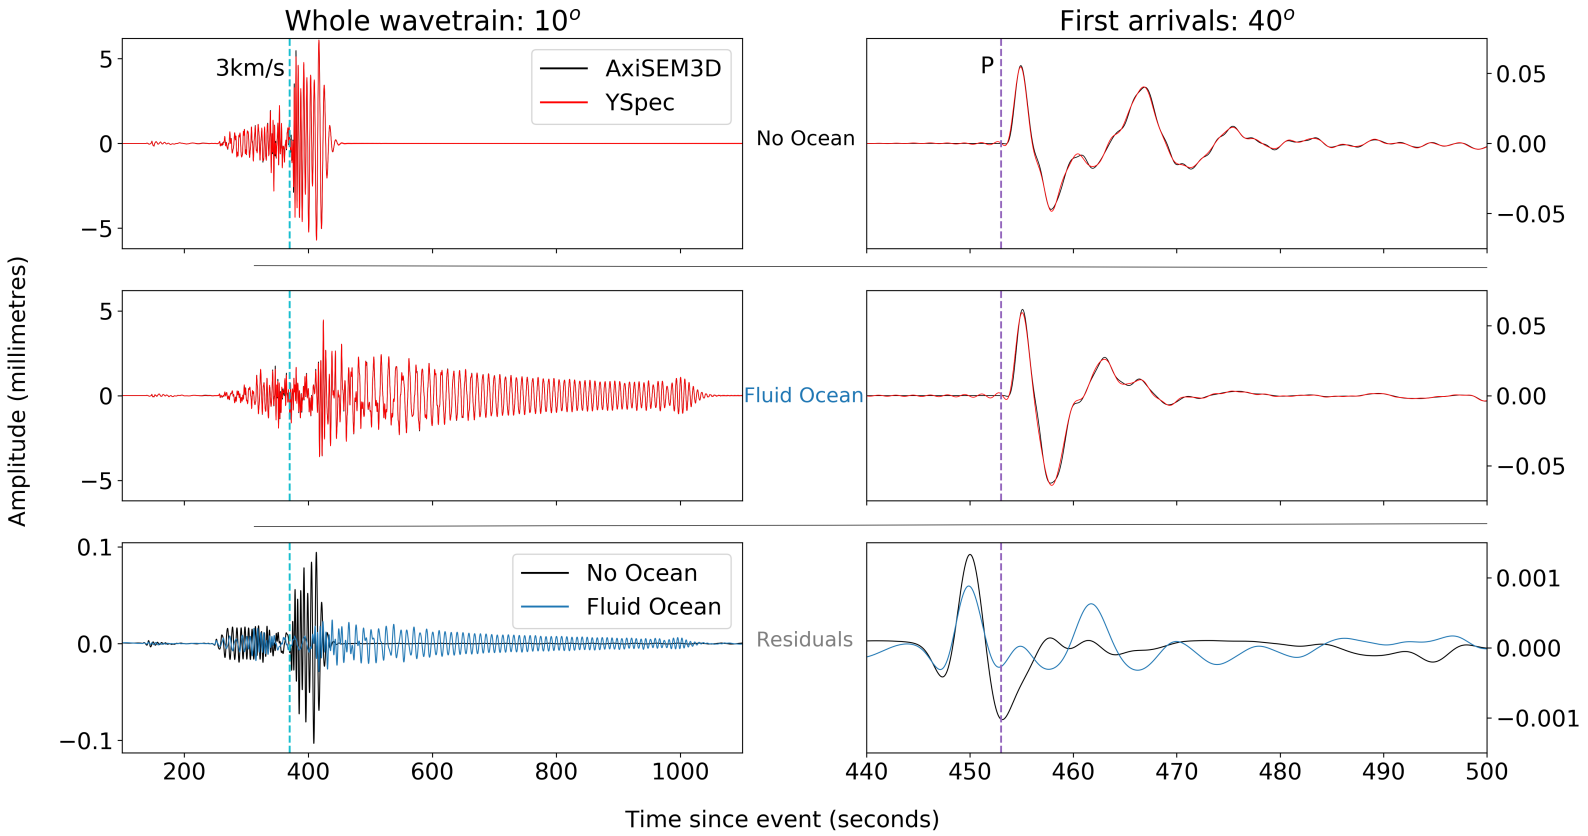

Figure 2.2: The modification to the waveforms from the addition of the ocean at 2 seconds resolution. The top row shows the seismograms in a PREM model with no ocean, whilst the bottom shows seismograms with a 3 km deep ocean layer in an otherwise identical setting. The vertical component is shown in all cases. The whole seismogram (including dispersive surface waves) is shown in the left column at 10° from the source; whilst the right column shows the modification to the P-wave train at 40° associated with the trapping of water depth phases in the ocean. For reference, the TauP arrival times of a 3 km/s surface wave and P phase are shown respectively in the left and right panels. The lowermost row shows the residuals ($u_{axisem} - u_{yspec}$), and it should be noted that the scales in the bottom two panels are different to those above.

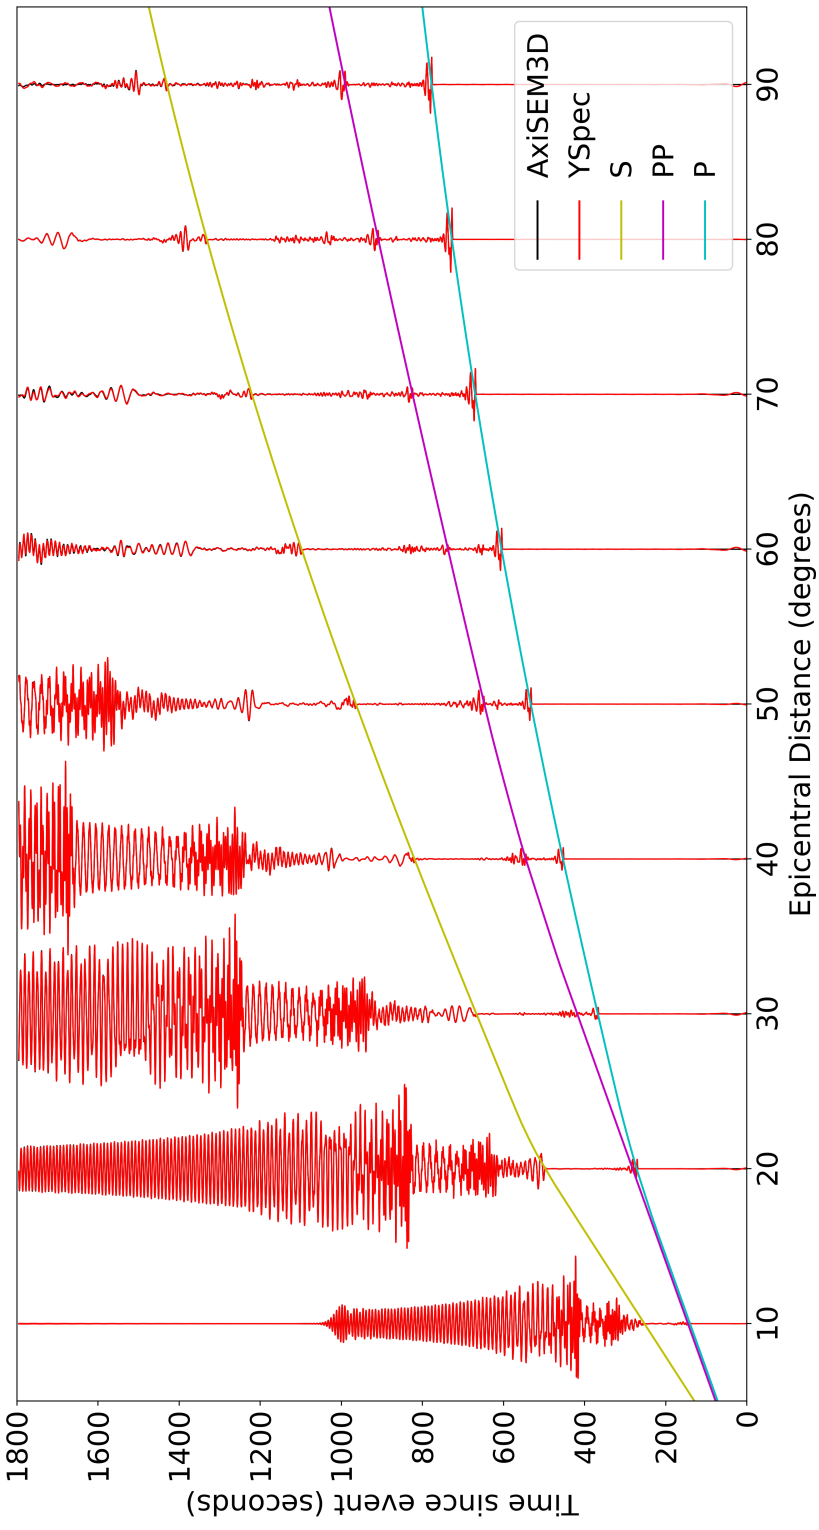

Figure 2.3: Benchmark results for the fluid ocean case between AxiSEM3D and YSpec, bandpass filtered between 2 and 100 seconds in the vertical component. The red trace shows the reference solution (YSpec), whilst the black trace shows the AxiSEM3D synthetics.

and the dispersion characteristic is modified. The changes to the body waveforms are more subtle, but so-called 'dog-leg multiples' (reverberations in the ocean column) become visible, and may be identified by their characteristic period of ∼5 seconds which corresponds to the vertical path length within the water column. These are most obvious between 460 and 470 s. Identical behaviour is seen in the radial component of the seismograms, whilst the transverse component remains unchanged in the presence of a homogenous ocean with uniform seafloor.

Strong agreement between AxiSEM3D and YSpec is demonstrated in both cases. To quantify this, we consider the time-frequency misfits between the two seismograms (after Kristeková et al. (2009)) using the ObsPy package (Beyreuther et al. 2010). The phase misfit increases slightly, from 0.02 to 0.03, when the ocean layer is added, whilst the envelope misfit remains unchanged at 0.01. According to the Kristekova ćlassification, both the envelope and phase misfits are 'excellent. Thus, we conclude that the implementation of the fluid surface layer in AxiSEM3D has been verified.

For completeness, we also demonstrate that the effects of the oceans become negligible at long seismic periods, justifying their lack of inclusion in low-frequency seismic simulations. Fig. 2.4 shows the convergence of the fully fluid ocean synthetics to those from the PREM case without ocean (and without ocean loading) at 50 seconds dominant period. Apart from the differences in the structural model, the simulation parameters are otherwise identical to those given in Table 2.1.

2.6 Evaluation of ocean loading

Having established the accuracy of AxiSEM3D simulations with a fluid ocean layer through means of a benchmark, we now consider the reliability of the ocean-loading formulation (Komatitsch and Tromp 2002). This is done by comparing AxiSEM3D seismograms in a PREM model overlain by a realistic ocean to those in PREM with an ocean load at the surface. The ocean loading formulation itself was validated in AxiSEM3D against SPECFEM by Leng et al. (2016). This is done for a variety of ocean depths and seismic periods at 20° from the source, which is distant enough

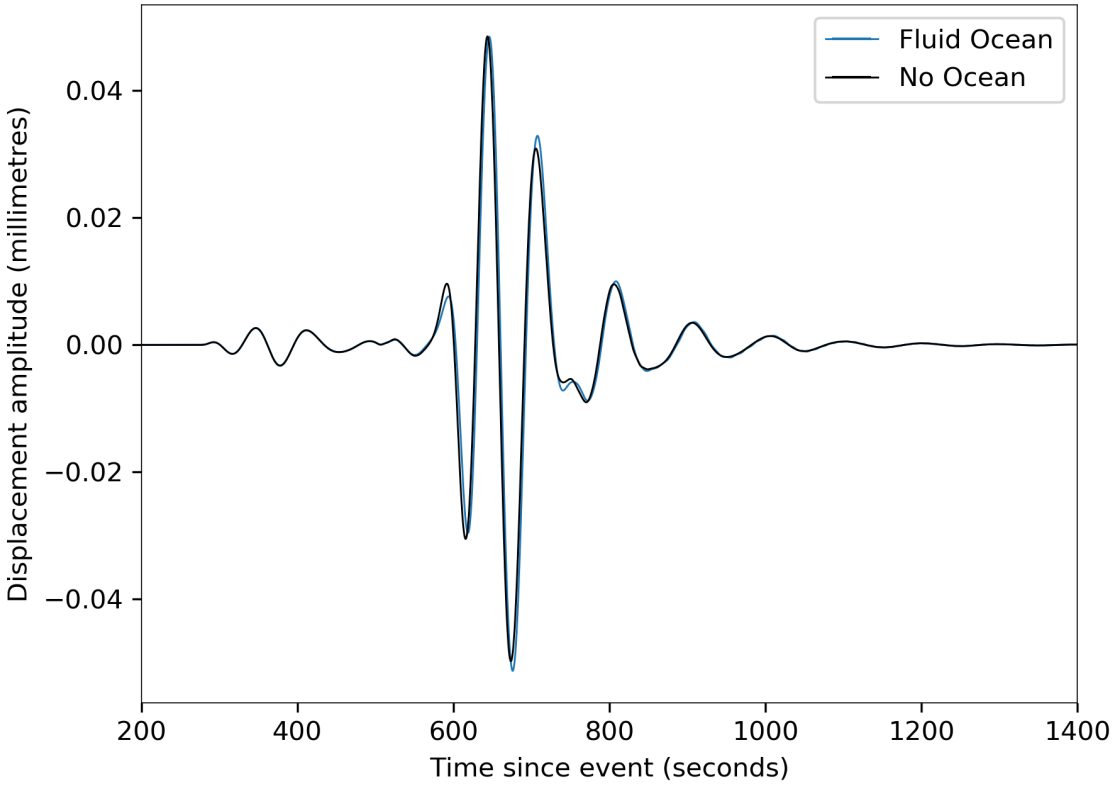

Figure 2.4: Comparison between the fluid ocean formulation and the PREM case without ocean, both simulated in AxiSEM3D, in the vertical component at 20° from the source. Synthetics are bandpass filtered between 50 and 100 seconds. No additional convolution with a gaussian function is needed in this case as both source-time functions are identical.

to capture the modified dispersion but close enough that the surface-wave train is well-recorded in the simulation interval.

Such an approach is similar in aim to that used by Zhou et al. (2016), but with two key differences. Firstly, we use a global, spherical Earth with realistic velocity and density structure, rather than a homogeneous half-space. Secondly, we also consider the effects of the ocean on the entire wavetrain, whilst their studies are restricted to the PP phase and Rayleigh waves.

Fig. 2.5 shows the evolution of the waveforms with increasing distance from the source for an ocean depth of 3 km in the vertical component, whilst Fig. 2.6 shows the reducing discrepancy between two formulations at increasing seismic periods. At greater than 30 seconds dominant period (lower right), no difference is resolvable

visually except in the final part of the surface wave coda. At less than 20 seconds, the extension to the Rayleigh wave coda is clear. By 10 seconds, the peak amplitudes are substantially over-predicted in the ocean loading formulation, whilst the coda length is much curtailed. At less than 7.5 seconds minimum period, differences in the body wave arrivals become apparent, whilst even the gross structures in the surface waveforms are missed, with only arrival times being predicted correctly. Animated versions of these comparisons, showing the transition in waveforms with decreasing period, are presented in the supplementary material.

Thus, we conclude that the ocean loading formulation works as expected above 20 seconds period for the entire wave train, but becomes unreliable at higher temporal resolutions. This result is consistent with the findings of Zhou et al. (2016), who identify modification to the Rayleigh wave dispersion characteristic and changes in the amplitude and arrival times of the body waves as being key features not predicted by the loading formulation.

For reference, neither formulation has any effect on the transverse component of the waveforms, confirming that in a radially symmetric model with an isotropic ocean, the ocean's effects are confined to the source-receiver plane. Of course, the ocean depth is not a uniform 3 km across the Earth's surface. For this reason, we also consider the envelope and phase misfits between the two implementations for a variety of ocean depths (Fig. 2.7).

A greater misfit is observed for the deeper oceans, as is expected where the ocean depth becomes more comparable to the seismic wavelength. A steep 'shoulder' is also clear in the intermediate period range, wherein the phase and envelope misfits both decrease rapidly with increasing dominant period, suggesting strong convergence to the ocean loading formulation at longer periods. In general, it appears that the degree of misfit is controlled by the ratio of minimum resolved seismic wavelength to ocean depth, though as the misfit calculation is non-linear and unreliable for large differences, no simple qualitative relationship for the transition is apparent. The degree of misfit also increases with increasing epicentral distance from the source, which we attribute to the increased length of the surface wave

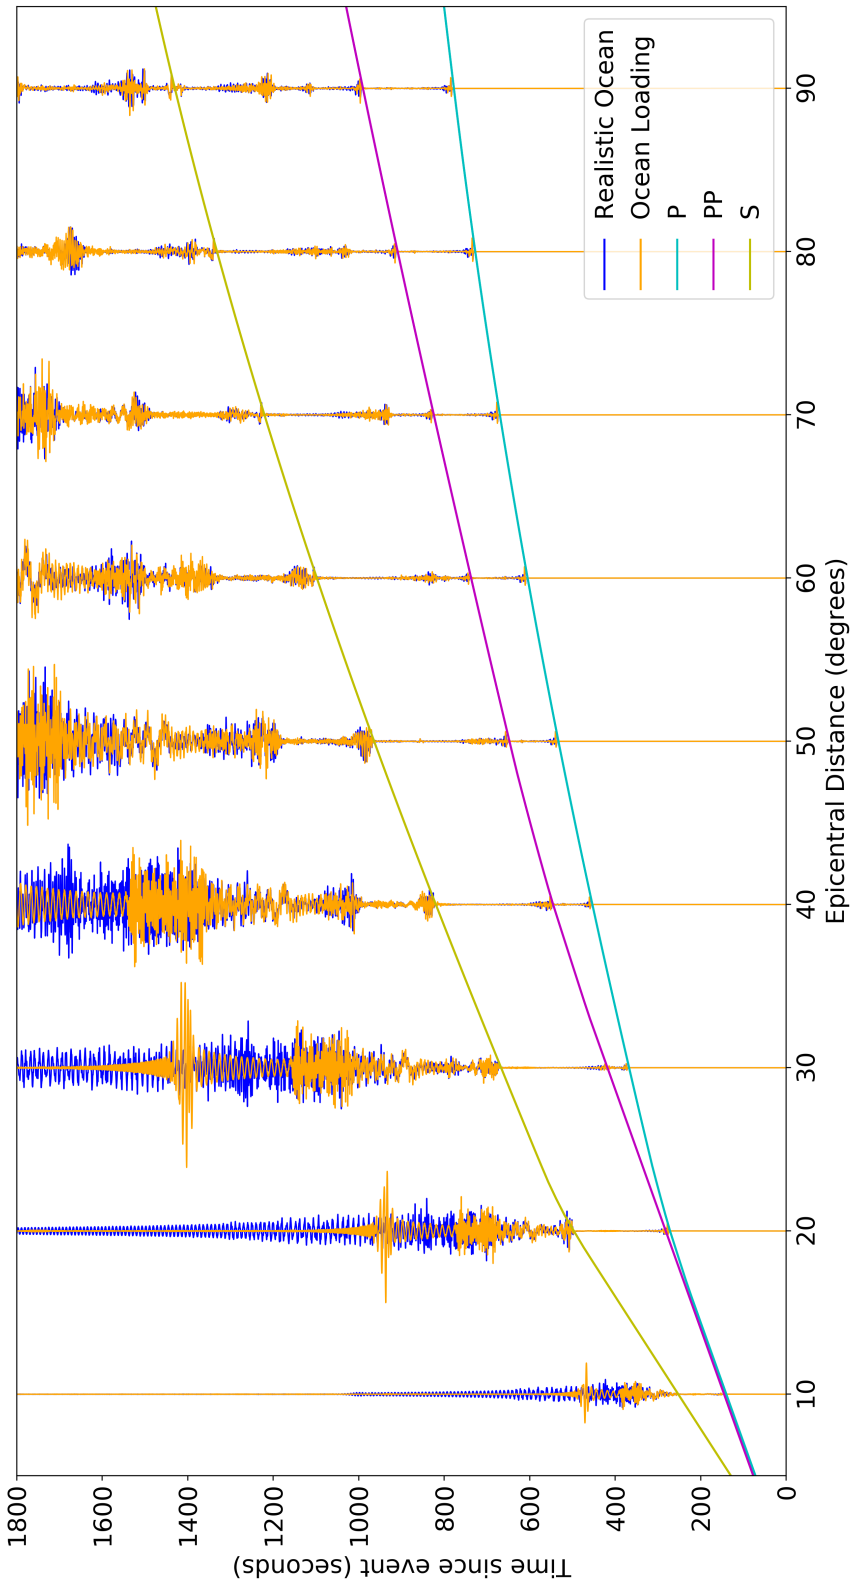

Figure 2.5: Comparison between the ocean loading and fluid ocean formulations for a 3 km deep ocean layer in the vertical component, both in AxiSEM3D. Seismograms are bandpass filtered between 2 and 100 seconds.

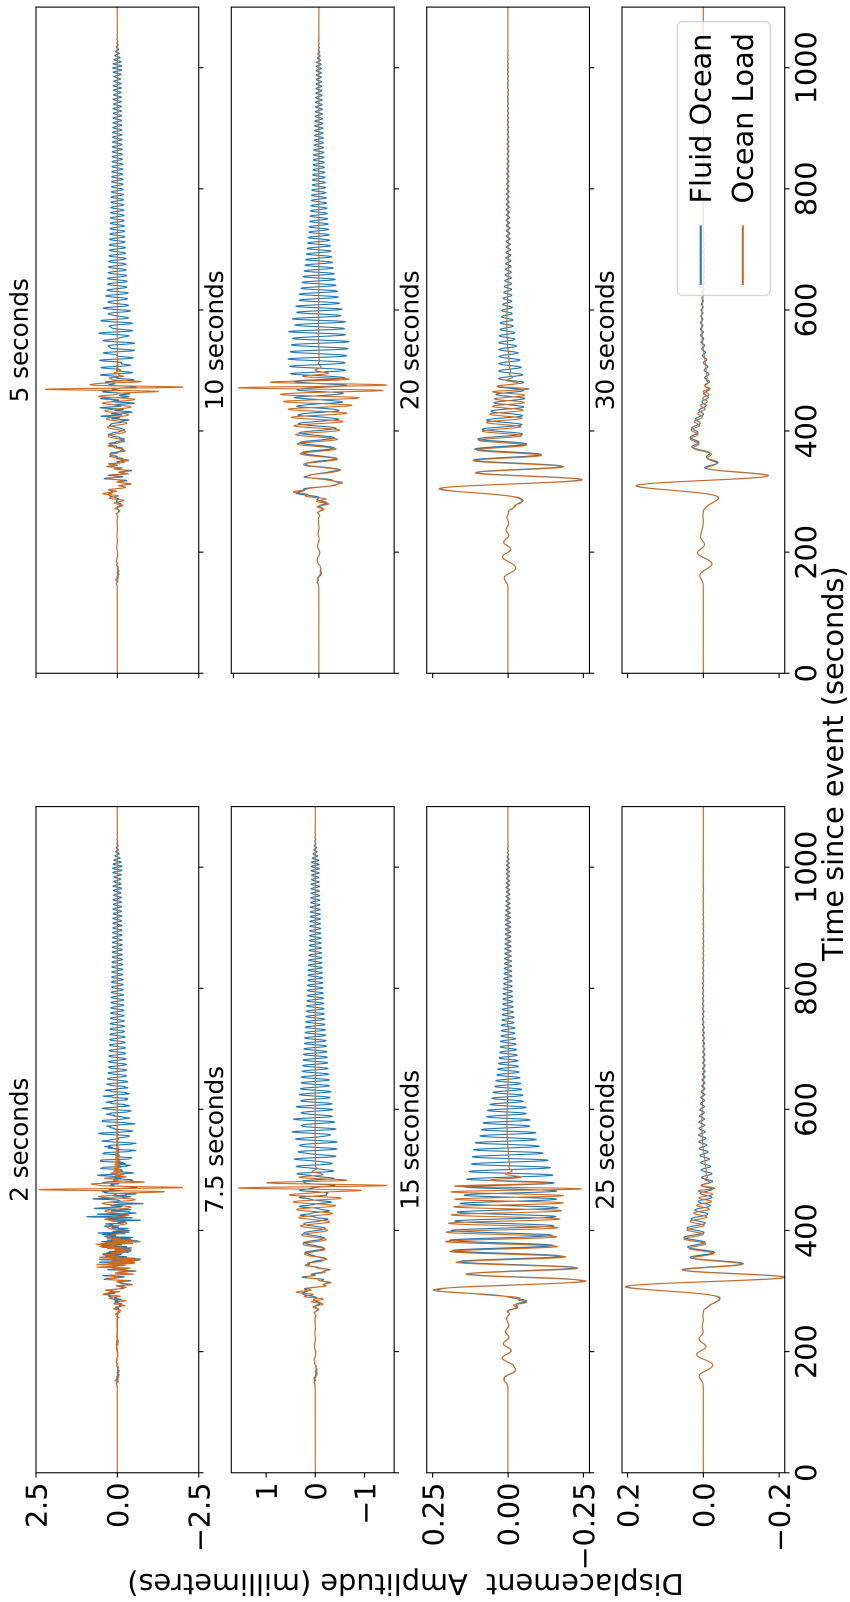

Figure 2.6: Comparison between the fluid ocean and ocean loading formulations for a 3 km depth ocean, in the vertical component, at a variety of different minimum periods 20° from the source. Note that the vertical scale is different for each row of figures.

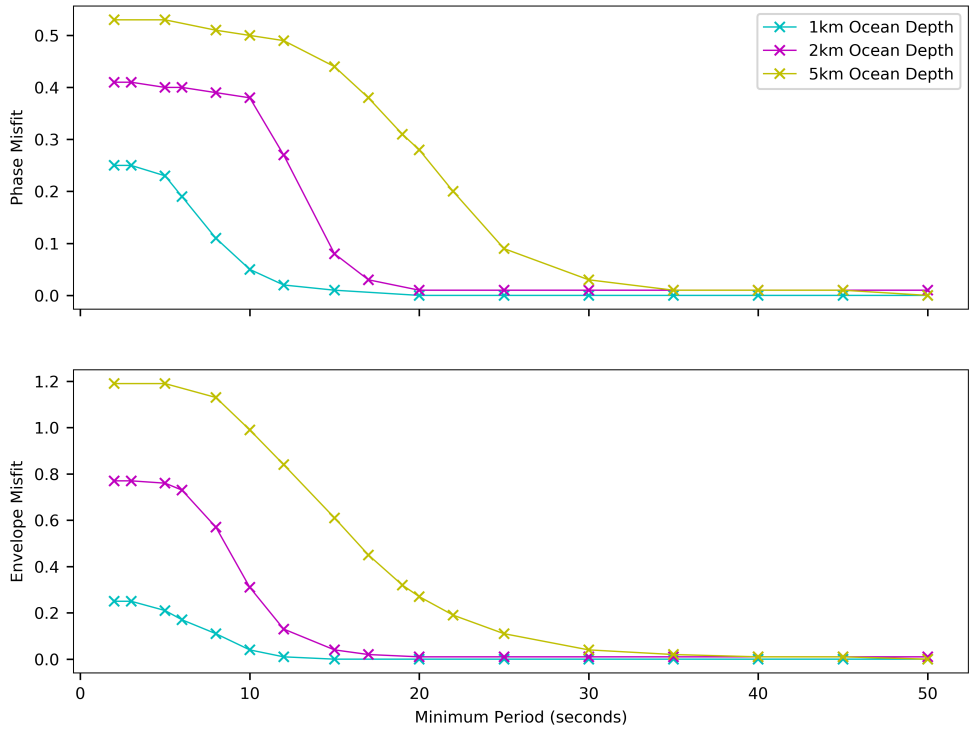

Figure 2.7: Phase misfits (top panel) and envelope misfits (bottom panel) between the ocean loading and realistic ocean implementations, with the vertical component sampled at 20° epicentral distance.

train at greater distances and the larger number of propagated wavelengths for in-ocean phases like pWP.

2.7 Bathymetry in a global ocean context

2.7.1 Bathymetric implementation

Having considered the effects of the addition of a constant-thickness fluid ocean layer on synthetic seismograms, we now consider the effects of undulating the seafloor solid-fluid boundary. In our current formulation, the ocean must cover the entire planet to accommodate homeomorphic boundary undulation, and ocean depths <500 metres (continents and their shelves) are linearly scaled to lie between 1500 and 500 m below sea level. This ensures that the time step does not become unfeasibly small in thin ocean layers overlying land areas. Naturally, this restricts

the domain of applicability to cover ocean basins and purely oceanic ray paths. The Earth's current continental configuration is such that many oft-used paths exist, for example at regional-to-teleseismic distances across the Pacific or Atlantic Oceans.

The ocean density and sound speed are constant throughout this model, but can be arbitrarily made to vary by latitude, longitude, and depth to reproduce geographical variations in seismic parameters caused by variations in water temperature, pressure, and salinity. Further work in this area will allow for global-scale meshes which support, for example, the propagation of high-frequency t-phases through the SOFAR channel. In these cases, the computational resource required may pose something of a challenge, and hence carefully defining the problem scope to minimise the grid size would be advised.

To illustrate the effects of bathymetry we consider a fluid ocean with and without bathymetry. Simulations are conducted at 5 seconds dominant period, apart from the animations which are at 10 seconds due to the enormous memory needed to produce them. All simulations use Crust 1.0 Moho (Laske et al. 2013), ETOPO1 bathymetry (sampled at 1° resolution, Amante and Eakins (2009)), and the SEMUCB-WM1 volumetric tomographic model of French and Romanowicz (2014) (sampled at 0.5° intervals and with $\delta v_p = 0.5\delta v_s$). The source used is the M_w 6.6 earthquake which occurred in New Britain, Papua New Guinea, on 2015 April 30. Arrival times for TauP (Crotwell et al. 1999) in PREM, i.e. without the fluid ocean, are also shown, with a small shift to account for the implementation of the source-time function in AxiSEM3D. To capture the complexity of the 3D models used, N_u is set to 1500 in the uppermost 100 km of the crust and mantle, and 200 below 400 km depth (depths are relative to 6371 km in all cases to ensure realism and consistency between the ocean and no ocean models). Between these depths linear interpolation of N_u is used. It should be noted that if smoother structural models are used, lower N_u may be used, decreasing computational cost significantly.

In undertaking particle relabelling on multiple boundaries (e.g. both the Moho and the seafloor), care must be taken to ensure that elements do not become too skewed (overly large bathymetric gradients being accommodated by too few

Table 2.2: Source Parameters

Date	2015 April 30
Time	10:45 UTC
Latitude	5.6°S
Longitude	151.7 °W
Depth	38.3 km (below mean sea level)
M_w	6.6
Half-duration	11.3 s

elements), or too small (the Moho and the seafloor moving toward each other such that the element shrinks and the time step becomes too small). With the coarse bathymetric sampling used here this did not pose an issue, but in simulations using stronger bathymetry this may be an issue. At higher frequencies, with more elements across which to spread a given gradient, this may be less of a concern.

Green's functions are generated for island stations in the Pacific Ocean, and convolved with the earthquake average source-time function from the SCARDEC database (Vallée and Douet 2016). Note that whilst the stations are located on land, we re-locate them to the seafloor in our global ocean formulation. Thus, any receiver-side conversions on the islands' submarine slopes are neglected; however these are not expected to be significant at these seismic periods given the small size of these islands (all have an above-water extent that is much less than 1° in any direction and hence they are not properly sampled anyhow in our implemented bathymetric model). Source parameters are given in Table 2.2 and receiver locations used in this study in Table 2.3, with seafloor-projected ray paths shown in Fig. 2.8.

2.7.2 The effects of bathymetry on waveforms

We begin by making a high-resolution comparison of the body waveforms in simulations with and without bathymetry. Such synthetics are presented in Fig. 2.9.

Comparison between the top and bottom panels reinforces the finding that the effects of bathymetry are less appreciable than those which come about from the addition of a fluid ocean layer at a resolution of 5 s, but nonetheless are noticeable. Examination of the bottom panel reveals that the coda associated with the ringing

Table 2.3: Receiver Locations

Name	Location	Latitude	Longitude	Distance
II.KWAJ	Kwajalein Atoll, Marshall Islands	8.8°N	167.6°E	21.6°
IU.WAKE	Wake Island, US Minor Outlying Islands	19.3°N	166.7°E	28.9°
IU.RAO	Kermadec Islands, New Zealand	29.4°S	177.9°W	37.3°
IU.JOHN	Johnston Atoll, US Minor Outlying Islands	16.7°N	169.5°W	44.6°

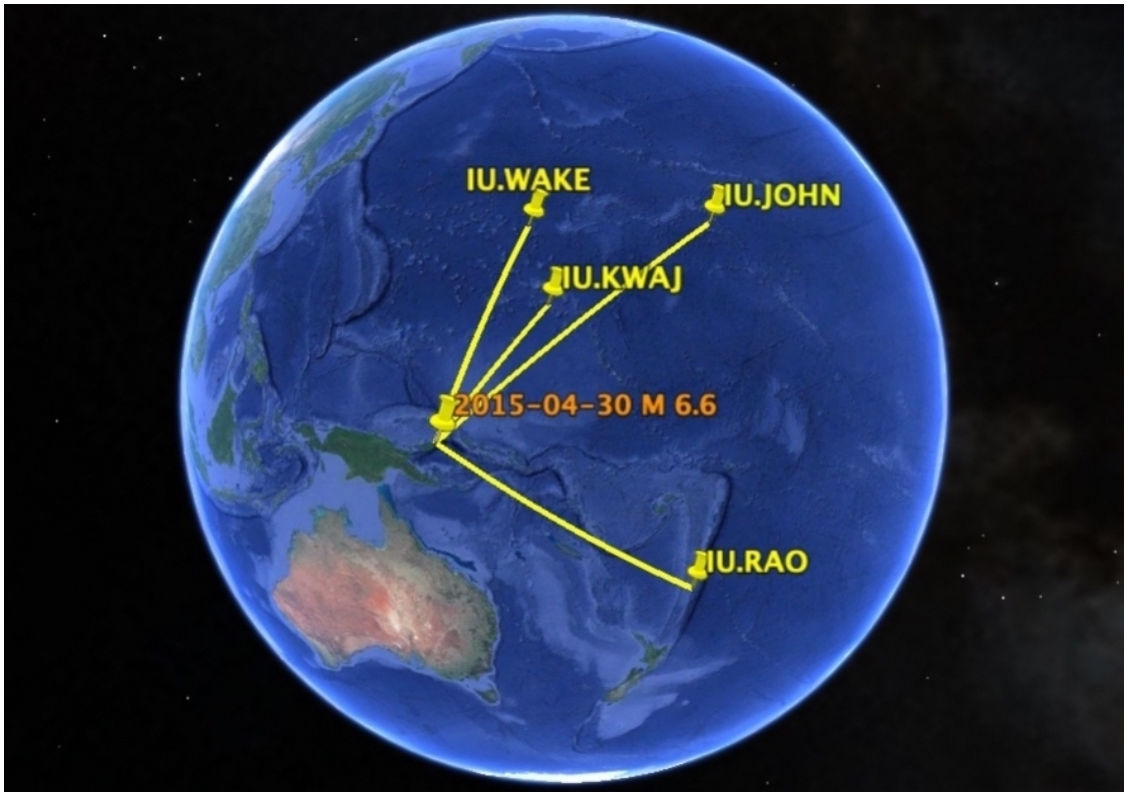

Figure 2.8: Source-receiver paths from the 2015 April 30 Papua New Guinea earthquake, plot generated in Google Earth Pro using Landsat/Copernicus/IBCAO images, SIO/NOAA/US Navy/NGA/GEBCO data)

of the body waves in the ocean column is modified by the roughening of the seafloor, and the modifications to PP (which bounces off the seafloor) are more substantial than those to P, which does not. The most substantial changes are in the surface waves, where the ratio of peak-to-mean amplitude is reduced, and a much more substantial coda is present. These effects are likely more pronounced at higher frequencies, as the shorter-wavelength surface waves are more sensitive to the small-scale bathymetric structures present along the seafloor.

Of course, such effects are strongly path-dependent. Fig. 2.10 illustrates this by considering the effects of the ocean and bathymetry on body waves at four separate stations. As expected, the most significant differences are observed at stations where the source-receiver path has strong bathymetric gradients (e.g. IU.RAO), whilst stations along reasonably flat source-receiver paths (e.g. IU.JOHN) do not.

We now present an animated comparison of the differences between the wavefields

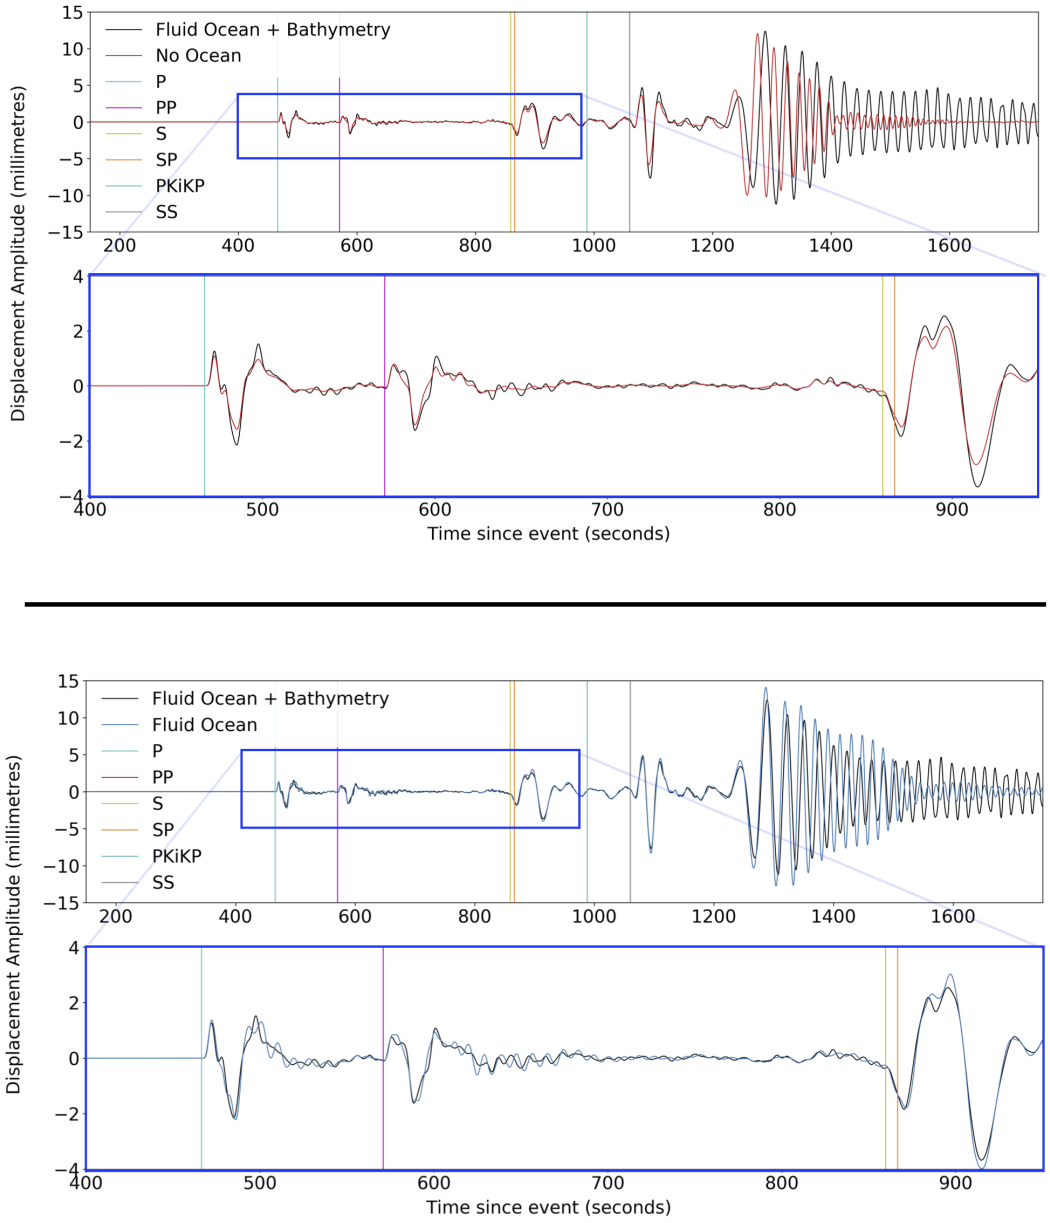

Figure 2.9: Synthetic seismograms for station IU.JOHN on Johnston Atoll, Hawaii. The upper panel compares the full, bathymetric ocean with the no ocean (PREM and Crust 1.0) case, whilst the lower panel shows the comparison between the bathymetry and no bathymetry cases. Green's functions are convolved with the SCARDEC source-time function and bandpass filtered at 5 s. The vertical component is shown and TauP times for the PREM case are overlaid.

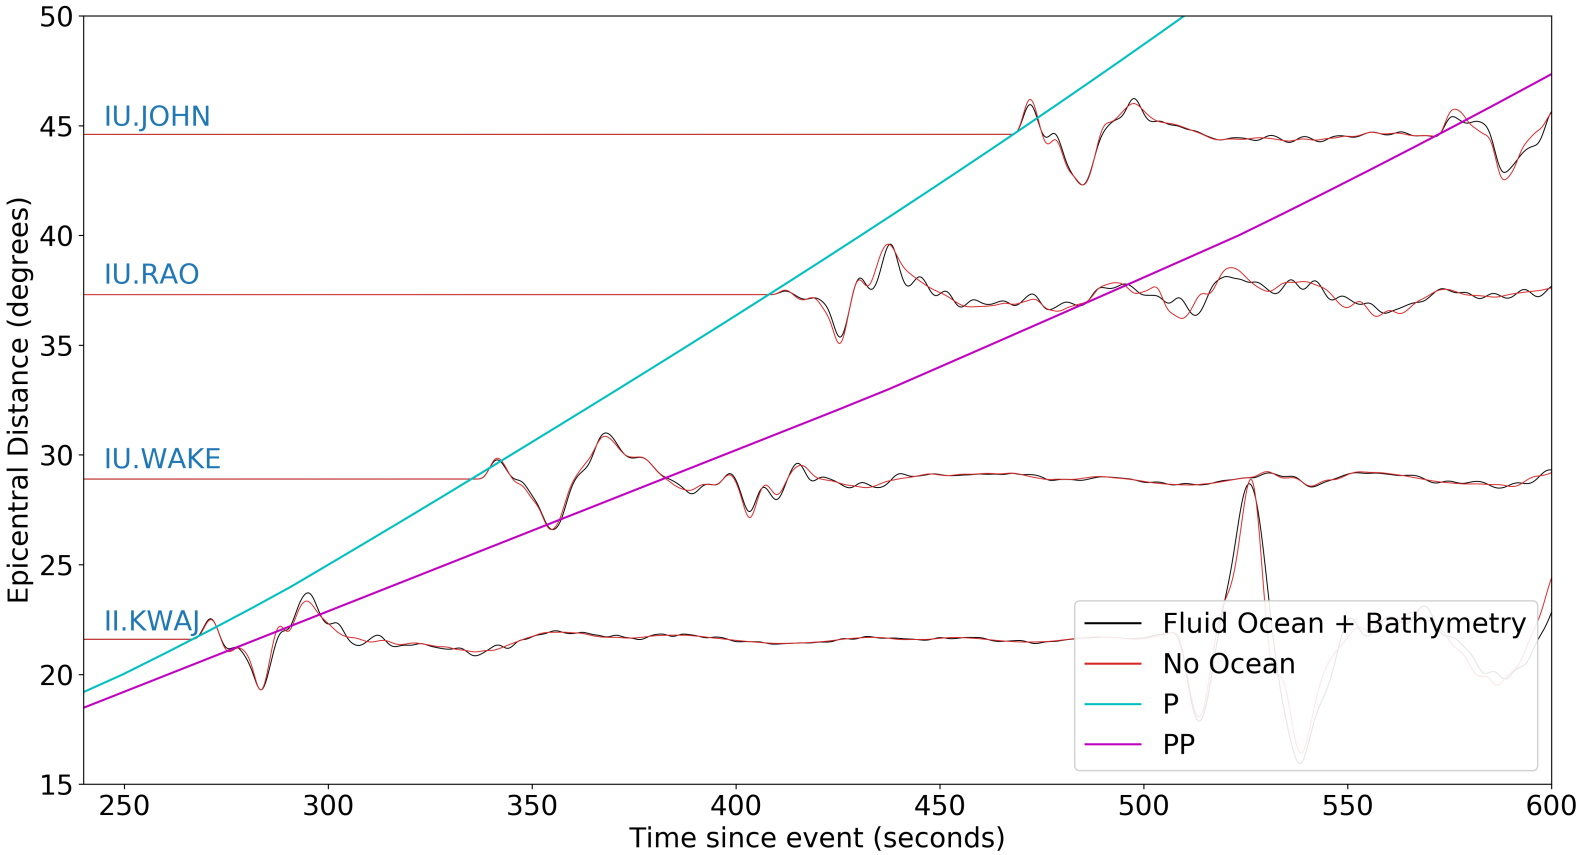

Figure 2.10: Detail of body wave arrivals in the vertical component, bandpass filtered between 5 and 100 seconds. TauP travel times for the phases P and PP in the PREM case are overlaid. Traces are normalised to the peak P-wave amplitude.

recorded at the seafloor in the case of the global fluid ocean with bathymetry and without bathymetry; with a focus on the surface waves. These visualisations are presented in Fig. 2.11. Note that in these simulations, conducted at 10 seconds dominant period, an 11 second half-duration gaussian source time function is used to ensure smooth interpolation and visualisation of the wavefield. The larger amplitude of the surface waves (leading to larger differences), and their greater sensitivity to near-surface conditions mean that they dominate the difference plots.

Video sequences were generated using ImageJ (Abràmoff et al. 2004). We include only four pertinent snapshots of the wavefield in this paper, however the full animations may be found on the 'Oxford Seismology' YouTube channel: https://www.youtube.com/c/SeismologyOxford.

In interpreting the differences in Fig. 2.11, it should be noted that whilst the simulations are consistent and accurate throughout, the wavefield observed in the sector clockwise from north-west to south-east is most realistic than that toward the south-west, where the scaling of the Australian continent to lie underwater occurs. The difference wavefield is not physically meaningful, but by tracking its 'propagation' the regions of bathymetry which significantly influence seismic wave propagation may be examined.

100 seconds after the event, two regions of significant difference are apparent. The northern portion corresponds to the passage of the surface waves over the island of New Britain and into the shallow Bismarck Sea, whist the southern portion emerges as the wavefronts pass over the steep topography of the New Britain Trench and into the deeper Bismarck Sea.

By 300 seconds, these differences have become more pronounced as elastically scattered remnants of the surface wave train appear to remain trapped in the bathymetric simulation for a substantial duration. The absence of these 'fingers' of remnant wavefield in the no bathymetry plots indicates that they are associated with the seafloor topography in these regions, rather than the Moho or crustal structure.

At 500 seconds these fingers remain pronounced, with the strongest appearing around the regions of rough bathymetry surrounding the Caroline Seamounts in

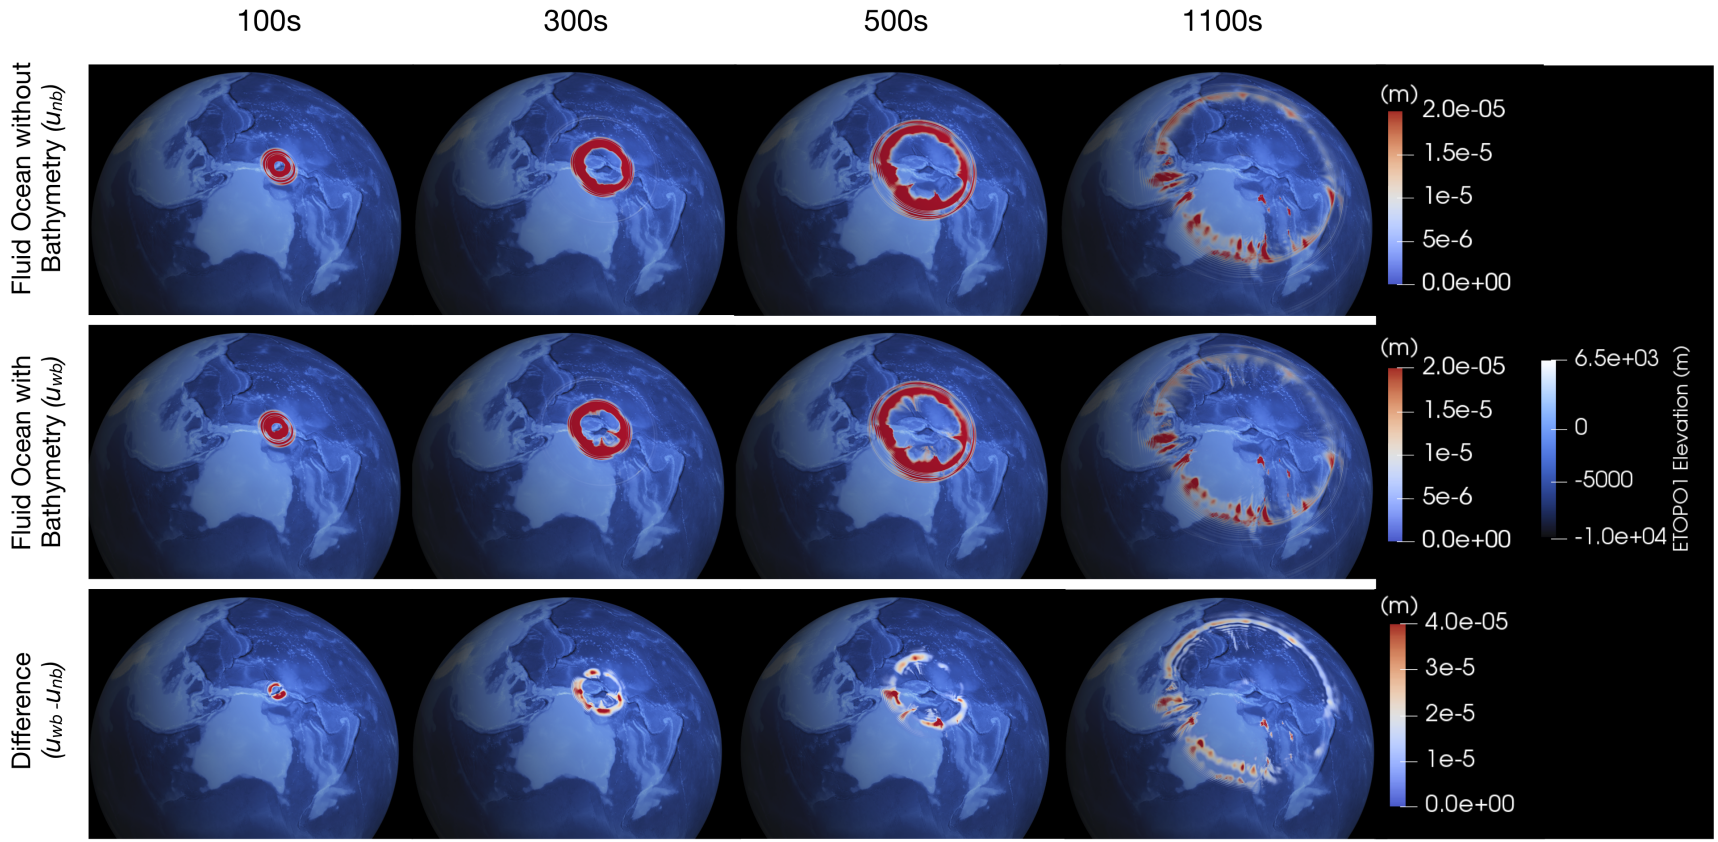

Figure 2.11: A visualisation of the simulated wavefield from the 2015 April 30 Papua New Guinea earthquake, as observed on the seafloor. The upper row of figures shows synthetics with a fluid ocean but without bathymetry, whilst the middle row is otherwise identical apart from the inclusion of bathymetry. The bathymetric colourbar applies to all three rows, and for ease of reference displays the physical elevation in ETOPO1 rather than the scaled elevation which we implement (for a detailed comparison to the scaling which we apply, see the supplementary material). The L2 norm of the two displacement vectors ($||u_{wb}||$ and $||u_{nb}||$ respectively) is shown in the top two rows, where red represents the largest displacements and white the smallest. The data range over which the colour is scaled in both is [0, 2x10^{-5} m] to ensure that the plots are not dominated by the surface wave amplitudes. For clarity, this means that displacements in the range [2x10^{-5} m,$\sim$4x10^{-5} m] all appear as the the deepest shade of red, where the latter value is the maximum amplitude observed in the simulations. The lowermost row is the L2 norm of the difference in displacements between the bathymetric and non-bathymetric simulations, i.e. $||u_{wb} - u_{nb}||$. In this row, the colour scale saturates at twice the maximum amplitude (4x10^{-5} m). An additional linear 'suppression' filter is applied to the wavefield opacities in the range [0, $\sim$1x10^{-5} m]. This has the effect of hiding small differences where $||u_{wb} - u_{nb}|| \sim 0$ (i.e. making them 'transparent') and ensures that the non-causal differences in numerical noise between the two simulations do not obscure the entire globe. North is upward in each screenshot, and the large pale area at lower centre is the Australian continent. New Zealand is visible at lower right, and the Indonesian Archipelago and Asian mainland at upper left. The Pacific Ocean, including numerous island chains referred to in discussions below occupies the right side of each snapshot.

Micronesia. Other substantial regions of difference appear in the Coral Sea off the eastern coast of Australia, which may be due to wavefield trapping in the 'valley' between the Australian coast and the shallower seas of the Lord Howe Rise; likewise in the Celebes Sea and to a lesser degree in the Mariana Trench region of the Philippine Sea.

By 1100 seconds, the largest differences are in the surface waves over the Australian continent, which are not realistically modelled and hence we do not consider their physical origin. Nonetheless, other substantial differences can be seen in regions with significant seafloor topography, including the northern part of the Tonga-Kermadec Ridge and the eastern Melanisian Basin.

Hence, it is clear from these animations that bathymetry has a significant effect upon the surface wave train, with particularly strong differences appearing in synthetics near regions of strong variation in seafloor topography, as would be expected. Whilst the effects on the body waves again appear to be relatively small in comparison, this suggests that analysis of observed surface waves should involve consideration of the roughness of the seafloor over the source-receiver propagation path.

2.8 Propagation through the water column

In AxiSEM3D, it is also possible to add receivers into the water column, which may be of interest in understanding the data recorded by MERMAIDS or at stations of the International Monitoring System for nuclear non-proliferation. To illustrate how waveforms within the water column vary as a function of depth, Fig. 2.12 shows the seismograms at regular depth intervals in the water column 'above' Wake Island. Of course, the seismometer on Wake Island is not actually underwater, but due to our relocation of the station to the seafloor, it has an ersatz water column above it.

The largest signals by amplitude in the vertical component occur at the surface, with noticeable changes in waveform shape descending through the water column. A detailed interpretation of these waveforms is beyond the scope of this paper, though the most significant body wave codas appear higher up in the water column.

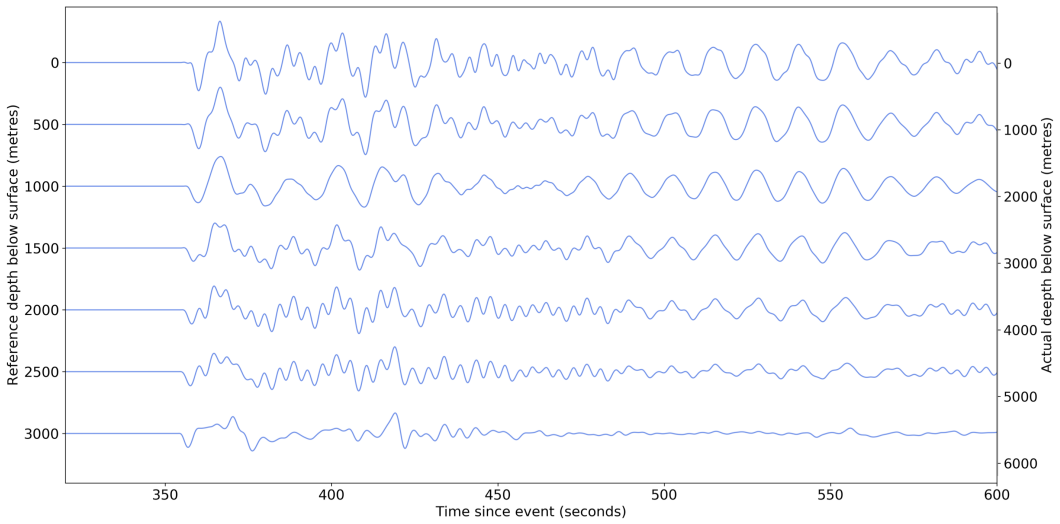

Figure 2.12: Seismograms in the water column above station IU.WAKE (Wake Island) at regularly spaced intervals. The vertical component, which is continuous across the seafloor interface, is shown.

The stations in the middling depths appear to have the smoothest waveforms, which may be due to complex reflections and scattering being more significant at the seafloor and surface.

2.9 The effects of the ocean on observed data

We now investigate whether the modifications to the waveforms induced by the addition of a fluid ocean and bathymetry are actually noticeable in observations. The comparison is made to both our new fluid ocean implementation with bathymetry, and also to the 'current standard', a model with a solid surface with topography upon it. Given the results presented earlier regarding the inapplicability of the ocean loading formulation at intermediate-to-high frequencies, we do not consider application of an ocean load to avoid conflating physical effects (the effects of the ocean's additional mass) with numerical errors (arising from the load's inaccuracy at the resolutions in question).

As detailed previously, Moho undulation and tomographic models are also included. As the 5 second dominant period used in the generation of synthetic

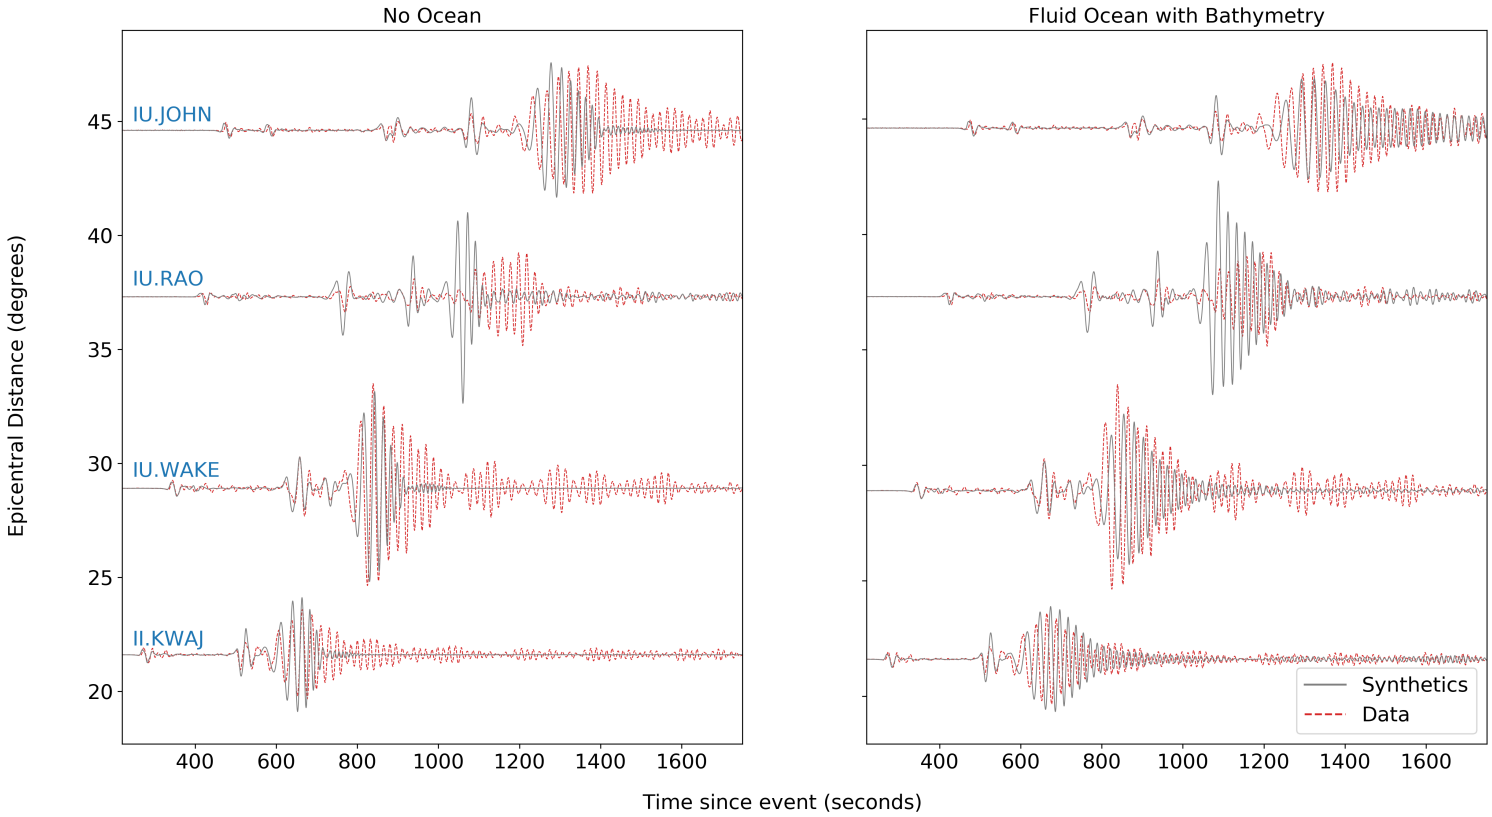

Figure 2.13: Comparisons between observed data (dashed lines) and synthetics (solid lines) for four island stations, with both data and synthetics log-Gabor filtered at 25 seconds and the synthetics convolved with the derived source-time function. Each trace is normalised to its peak p-wave amplitude, with the vertical component shown.

waveforms falls in the middle of the ocean microseism band (Longuet-Higgins 1950), we filter both data and synthetics using a log-Gabor filter with a central frequency of 25 seconds and focus on the surface waves.

Although this removes some of the higher frequency content of both synthetics and data, the prevalence of ocean noise at these island stations at shorter periods means that meaningful comparison would be extremely challenging. The choice of a log-Gabor filter over a bandpass filter is made as its response is more suited to removing the numerical artefacts above the maximum mesh-resolved frequency whilst also suppressing the intermediate-frequency ocean microseismic noise. It should be emphasised that the 25 second central frequency of the filter's instantaneous impulse response in this section is not directly comparable to the 25 second panel that in Fig. 2.6 - there it is the high-corner frequency of the bandpass filter above which little-to-no energy is transmitted, whilst here it is the frequency of maximal filter response for a relatively wide passband filter.

Fig. 2.13 shows the data overlaid with synthetics in the no ocean case on the left, and in the fluid ocean with bathymetry case on the right. There is a noticeable change in the qualitative fit between data and synthetics. Given that at 25 seconds resolution we expect the most significant changes to appear in the surface waves, we focus on these in this section.

It can be seen that the longer duration of the surface wave train, together with what appears to be a scattering-induced coda, are noticeable in the fluid ocean case. The length of the initial Rayleigh wave train is also better matched. The most substantial discrepancy between data and synthetics is at station IU.RAO in New Zealand's Kermadec Islands. This is likely attributable to the station's location on the edge of the Tonga Ridge, a region of significant tectonic and bathymetric complexity which is not accurately or reliably reproduced at the coarse sampling level of the existing crustal models used here.

Thus, from a qualitative inspection, we conclude that the contributions of a fluid ocean with realistic bathymetry account for more of the observed waveform features visible in data than are predicted from synthetics which include only

topography on the solid surface without a fluid ocean, at least for this particular set of source-receiver pairs. In other words we suggest that this implementation is a more reliable representation of the actual wave propagation physics.

Further work may focus on a more thorough characterisation of the goodness-of-fit difference between these traces. A proper exploration would require the computation of both frequency-domain and time-domain misfits, and ought to be undertaken for a range of passbands. This will be complicated by the fact that the noise is not white, but rather concentrated in the microseismic band between 5 and 10 s.

2.10 Discussion

2.10.1 Computational cost

The efficiency of AxiSEM3D's method is such that high-frequency simulations as described in this paper may be carried out using an intermediate-size cluster.

The 30 minute duration simulations at 10 seconds (as shown in the animations), with both crustal and tomographic models, require approximately 6000 core-hours on a Cray XC30 architecture. The addition of a 3 km fluid ocean with a flat seafloor increases the cost by approximately 90%, in part to the modest increase in the number of elements needed (from ∼124,300 to ∼142,500) to accommodate the ocean's volume, but mostly due to the 30% decrease in the minimum time step to accommodate its lower sound speed. The latter effect is more significant because AxiSEM3D uses the same time step across the mesh, and if the limiting global time step is set in the ocean, it acts as a constraint on the time step everywhere within the simulation volume.

At 5 seconds resolution (as in the seismogram plots), the corresponding cost for a simulation without an ocean layer is approximately 46000 core-hours, with a relatively smaller increase (65%) when a fluid layer is added. This illustrates the addition of the ocean layer is less of an additional computational burden at higher frequencies.

The increase in cost associated with the addition of bathymetry is strongly model-dependent, and evaluating it is somewhat more involved. At 10 s resolution

in the models discussed above, replacing the flat seafloor with an undulating one increases the cost by approximately 10%. In the 5 s case there is actually a cost decrease of approximately 6% associated with the addition of bathymetry. This may be somewhat counterintuitive, given that the particle relabelling transformation is computationally intensive and in-and-of itself always increase the computational cost, but can be understood by considering the deformation of elements at the location where the limiting time step is set.

In the 5 s case without bathymetry (but with undulation along the crust-mantle boundary), the minimum time step is set at the Moho at a position approximately 111° epicentral distance from the source. Because of the axisymmetric way in which we formulate the problem, a more exact geographic location cannot be identified. The addition of bathymetry deforms the elements between the Moho and seafloor again, in order to ensure that both interfaces are properly honoured.

In the examples considered here, the element in which the limiting time step was originally being set is stretched to a degree that it is no longer has the smallest time step globally. Instead, the new limiting time step is set at a location 106° from the source, though still at the Moho. This new element has a time step of 0.0236 s, ∼8% higher than the original value of 0.0218 s. Just as the decrease in minimum time step (of approximately 30%) with the addition of the ocean increased the overall simulation cost by a comparable value, this increase in dt decreases the simulation cost by a corresponding factor. The reason that the cost decrease is not exactly the same (6% rather than 8%) is due to the fact that the additional particle relabelling involves extra computation.

In the simulations presented here, this is clearly advantageous; but it should be noted that this effect may be somewhat esoteric and associated with our choice of 3D models and the exact resolution chosen (indeed, it occurs at 5 s but not 10 s). Hence it is not generally true that adding bathymetry to simulations will decrease the computational cost. However, at high frequencies with high resolution implementations of the Moho (which necessitate very small elements), the addition

of a second undulating boundary may occasionally yield such cost savings, though it should be noted that the demands on computer memory will be higher regardless.

2.10.2 Cost scaling

For a fixed structural model in AxiSEM3D it is possible to derive approximate cost scalings with frequency, which can be compared to other tools such as SPECFEM. In the scenarios considered in this paper, the seismic wavelength is substantially smaller than the characteristic structural scale associated with the most densely sampled feature (the bathymetry). As such, an approximate scaling relationship with frequency ω and azimuthal order N_u(derived by Leng et al. (2016) and Leng et al. (2019)) is given by $O(\omega^3 \cdot N_u)$.

This linear behaviour in N_u holds regardless of whether the increase is global (across the entirety of the planet's radius) or confined to only the upper crust and mantle, though a localised increase in N_u is of course cheaper than a global one. Thus, AxiSEM3D offers significant flexibility and can accommodate suitable 3D heterogeneities of arbitrary complexity with a favourable cost scaling, so long as the seismic wavelength is shorter than the structural scales. If the model becomes extremely complicated such that the seismic wavelength is comparable to the characteristic scale of 3D features, the scaling tends toward $O(\omega^4 \cdot N_u)$, more akin to that of full 3D methods such as SPECFEM.

For reference, Leng et al. (2016) conduct a full cost comparison against SPECFEM3D-GLOBE, and at 10 s with a full 3D tomographic model (similar to the scenario here but without the fluid ocean and at lower resolution), the speed-up is $\sim$100. The advantage increases with frequency but decreases with model complexity.

We also provide an estimate of the computational resources required for a global t-wave simulation, in reference to the scenarios discussed in Sec. 6.3 above - though it should be noted again that this estimate is strongly model dependent. Assuming that simulations were carried out at around 1 Hz (likely the lowest frequency suitable for t-phase modelling), the flexibility of AxiSEM3D is such that the N_u profile could be substantially reduced below the Moho, a regional scale

Table 2.4: In each scenario, four values are specified in the AxiSEM3D N_u profile. The 'Greater N_u value' refers to the constant value of N_u used between the surface and the depth given by 'Greater N_u depth bound', whilst the 'Lesser N_u value' refers to the constant value of N_u used between the 'Lesser N_u depth bound' and the planet's core. Linear interpolation between the Greater and Lesser N_u values is used between the two boundaries. As an example, the 'Low' profile uses $N_u = 1000$ in the top 40 km of radius, and $N_u = 40$ below 30 km depth, and a linear scaling in between.

Regime	'High Cost'	'Low Cost'
Mesh Colatitude	180°	90°
Greater N_u value	1500	1000
Lesser N_u value	200	40
Greater N_u depth bound (km)	100	30
Lesser N_u depth bound (km)	400	400

mesh used, and mantle tomographic models removed. Depending on the scales of propagation required, we estimate that this would have a computational cost of $10^5 - 10^6$ processor-hours. Increasing the frequency to 10Hz would increase this by a factor of approximately $\sim$100.

2.10.3 Potential cost savings

All simulations conducted thus far in this paper have made use of a 'full' global mesh and an azimuthal expansion of N_u which is more than high enough to capture both body and surface wave interactions with three-dimensional structures. Using the flexibility of AxiSEM3D, we now consider the effect of the cost-saving measures discussed in Sec. 2.10.2. Such modifications may include using a regional mesh (one which does not solve the wave equation on a complete sphere using a 'D' shaped mesh, but rather uses a wedge-shaped 'chunk' to simulate a smaller area), or localising the regions of high N_u such that phases of interest are accurately simulated whilst others are not.

As an example, Fig. 2.14 compares the generated synthetics for two different sets of simulation parameters at the station II.KWAJ, with the profiles given in Table 2.4. As before, all depths are relative to the ocean surface at radius 6371 km.

The low case shows very little discrepancy compared to the high one, indicating that the solution is well converged and the 3D structures present along this particular

source-receiver path are adequately sampled. The only apparent differences are in the deeply penetrating body waves, as may be expected from their non-negligible sensitivity at depths where the sampling of N_u has been significantly reduced; and in the surface wave coda. The differences in the coda are likely caused by the interacting reflections from the edge of the regional-scale mesh.

It should be noted that the degree of discrepancy between the different traces is likely to be strongly azimuth-dependant - a source-receiver path with smooth bathymetric and Moho undulation will require comparably lower N_u to achieve the same degree of convergence than one with sharp variations in the interfaces. This must be accounted for when optimising the N_u profile for a particular study; and of course regional meshes with a reflecting boundary (as here) are not suitable for studying surface wave codas due to contamination of the signal by reflections.

In this scenario substantial reductions of N_u are not, therefore, suitable for high-resolution studies of body wave surface multiples (such as PP or PPP), or generation of long-duration synthetics (where spurious reflections from the edge of the truncated mesh become problematic), but with careful optimisation of the N_u profile significant computational savings can be made if only specific phases are being investigated. For example, the 'Low' profile reduces the computational cost by 97% at 5 s. This enormous speedup is associated both with the reduction in N_u and the decrease in the size of the computational domain, the latter of which also increases the time step as the previous time step limiting element in the mesh (at 106° distance) is not included.

2.10.4 Limitations

As mentioned previously, the use of a global time step means that the addition of an ocean significantly increases the computational cost of simulations; the positive caveat being that for a given model the degree of temporal oversampling decreases with increasing frequency. As a result, higher-frequency simulations are more expensive than lower-frequency ones, and simulations with an ocean layer are more expensive than ones without; but the efficiency savings from moving to higher

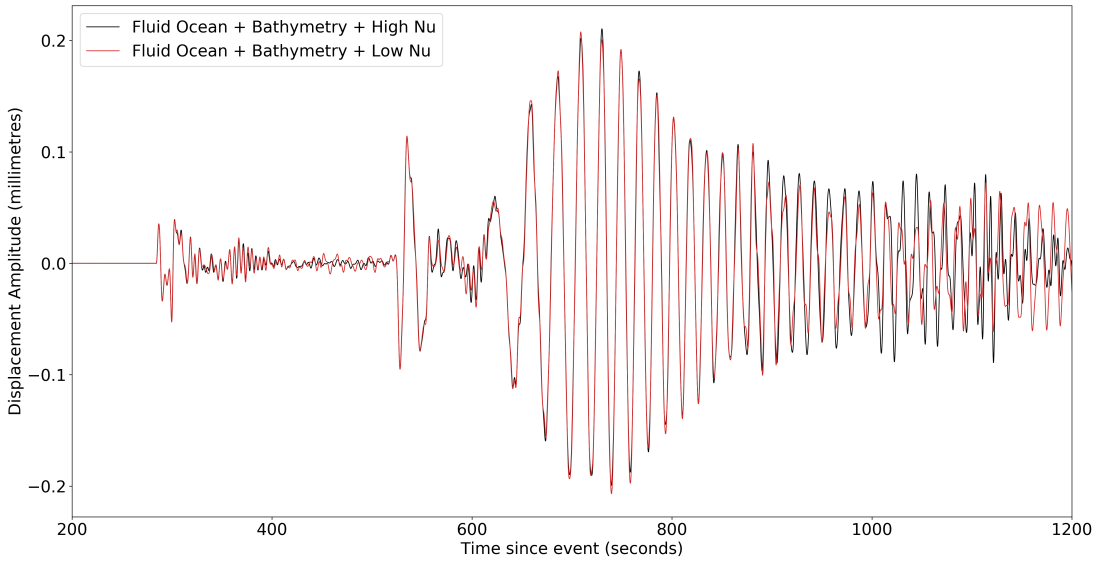

Figure 2.14: Synthetics for the station II.KWAJ in the Marshall Islands, showing the sensitivity to a reduction in the value of N_u. Note the largest discrepancies in the surface-multiple body waves, and in the surface wave coda.

frequencies in AxiSEM3D are greater in the case with an ocean than the case without. The addition of localised time stepping (e.g. Rietmann et al. (2015)) is one route through which the computational cost can be further decreased without sacrificing simulation accuracy.

The fact that the ocean must cover the entire planet is a clear current limitation of this method. The comparison to observations in Sec. 2.9 indicates that this is still an advance upon current numerical modelling methods, which can be justified physically on the grounds that the increase in the complexity of wave propagation physics captured with a fluid ocean, can, with careful use, outweigh the reductions in reliability associated with the relocation of stations on isolated islands to the seafloor. A more realistic representation of ocean-continent boundaries would be far more flexible in its application, and hence remains a long-term development goal.

2.10.5 Extensions in a global ocean model

In simulations constrained to use a global ocean model, the simplest increase in the realism (and hence of reliability) of the synthetics would be a reduction in the minimum ocean depth (here 500 m) over the continents. This would mean

that at a given seismic period, the sensitivity to the ersatz fluid layer would be reduced. Quantifying the exact degree to which this would change the synthetics at a given location is difficult, given that the effects of bathymetry are significant and highly spatially inhomogeneous, and that the addition of the fluid layer affects each component and seismic phase differently. Nonetheless, some insight may be gleaned from Fig. 2.12.

A rough estimate based on the convergence of the fluid surface layer case to the no-fluid surface layer case with increasing period (see supplementary material listed under DOI in preface) suggests that where the ocean depth is at least an order of magnitude less than the seismic wavelength, the effects of the ersatz fluid layer can be completely neglected. At 5 s resolution, this would equate to the ocean being no shallower than ∼200 m over continents, which would decrease the time step by factor ∼2.5 and hence increase the computational costs reported in Sec. 2.10.1 by roughly the same factor. This option may be especially attractive at higher frequencies, where as discussed in Sec. 2.10.4 the inclusion of the ocean layer is less of an 'inefficiency' - though care would need to be taken to account for the flattening of continental topography.

Nonetheless, if we restrict ourselves to purely oceanic source-receiver paths, various interesting phenomena not considered here can still be simulated on global scales. These could include scenarios either without or with a source in the fluid ocean. In the first case, an exploration of the effects of ocean bathymetry on studies of source properties using intermediate frequency (10-20 s) body waves would be pertinent, given that such effects are of particular concern when attempting to understand the dynamics of the largest earthquakes, which occur at subduction zones with rough bathymetry. Alternatively, a source in the fluid ocean could be used to study the generation of the primary ocean microseism by considering the propagation of acoustic waves from a near-surface pressure source down toward an area of rough bathymetry.

With the addition of a stratified or globally varying water column structure and the use of a regional-scale mesh with an optimised N_u profile, simulation of t-phase

propagation in the SOFAR channel is in theory achievable, if the computational resources needed to reach the necessary high frequencies (∼5 Hz) are available. Sources here could either be in the water column, for example representing illicit nuclear tests; or in the solid earth to study conversion of earthquake-generated P-waves. Care would need to be taken here to ensure that bathymetry along the propagation paths here is realistic, as t-phases are more sensitive to small-scale bathymetric variations than body waves, and hence our relocation of small islands to be underwater may mis-represent the degree of 'propagation blockage' which such waves experience. AxiSEM3D is ideally suited to such studies, as the value of N_u in the crust can be adjusted to ensure sufficient sampling of the bathymetry without adding unnecessary computational cost in the deeper mantle.

2.10.6 Development of a localised ocean formulation

The development of a localised, more realistic ocean formulation is also an avenue for potential future exploration. In the AxiSEM3D formulation it would be possible to create a 'localised' ocean which is axisymmetric, but which nonetheless honours a realistically sloping seafloor which meets the continental surface at its edges. This can be envisaged as a donut-shaped ocean 'trench' circling the entire planet about its equator, for example. This would require manual editing of the mesh, but would enable more realistic examination of the effects of the bathymetric profiles of continental shelves, if appropriate source-receiver paths are chosen. Such a setup might, for example, be of interest in understanding the generation of the transverse component of ocean microseismic noise.

Going one step further, development of a completely realistic 'patched' ocean would be extremely challenging due to the use of global basis functions in the AxiSEM3D Fourier expansion. These basis functions impose the constraint that the physical principles, such as the boundary conditions, must be the same on any line of azimuth. Hence, the arbitrary switching between fluid and solid areas which would be required for a patched ocean is not at present achievable. It ought to be possible to replace the global pseudospectral representation with a localised one, for example

one in which the basis functions are discontinuous wavelets or Slepian functions (e.g. Simons (2010)). Such a basis would not use globally constrained interpolating functions, and hence could model different boundary conditions at different azimuths.

2.11 Summary

We have presented a novel, open-source method for simulating global seismic wavefields in a model with a fluid ocean layer and 3D bathymetry. This method is embedded in the spectral-element method AxiSEM3D, introduced by Leng et al. (2016), which uses a pseudospectral representation in the azimuthal direction to significantly reduce computational cost. Our implementation is benchmarked for a radially symmetric model against the code YSpec to ensure its validity. An evaluation of the ocean-loading formulation previously used in global seismology is conducted, and we find that the approximation is not valid for the entire wave train unless the period reaches longer than 20 s for a 3 km ocean depth. Such findings concur with previous evaluations of this formulation which used a simplified testing setup (Zhou et al. 2016).

In a non-radially symmetric (i.e. fully 3D) model, we are able to investigate the effects of bathymetry through use of a global ocean formulation, and show that seafloor topography introduces substantial modifications to the surface waves. The modification to the body waves, whilst appreciable at 5 seconds dominant period, is much more slight. A suggestive comparison to data was also made, wherein we find that more of the gross features of the observed surface wave train are reproduced when a fluid, bathymetric ocean is added to the structural model. Whilst this study considers only specific oceanic propagation paths, the prevalence of oceanic source-receiver paths in modern seismic analysis and the importance of interpreting data from remote stations in otherwise poorly-sampled areas suggests this implementation can be a useful advance in computational seismic modelling.

2.12 Epilogue

Since the acceptance of this paper, I have been exploring the use of this AxiSEM3D formulation to study the generation of microseismic noise in the oceans, a possible extension which was discussed in Sec. 2.10.5. This work has been carried out in collaboration with Lucia Gualteri, Eleonore Stutzmann, and Lise Retailleau.

2.12.1 Microseismic noise: an introduction

The aim is to investigate the effect that bathymetry has on microseismic noise generation and propagation. Hurricanes and cyclones are used as exemplar sources, on account of their clear spatial and temporal localisation, which creates noticeable increases in the global ambient noise field recorded during periods where they are active.

Once the effect of bathymetry on noise generation and propagation is more thoroughly understood, it should be come easier to both: 1) remove noise as a contaminant from datasets of interest and 2) utilise it as a seismic tool, for investigations such as those described in Sec. 2.3.

Such investigations, including simulations, may be carried out on either a local or a global scale. The former offers information on how local bathymetry affects the signals recorded by individual seismometers, whilst the latter can be used to explore the generation of body waves from pressure perturbations in the water column more generally.

2.12.2 Noise in the Indian Ocean

One example of how this may be applied to better understand seismic noise is presented here for completeness. This particular simulation is designed to replicate the excitation of microseismic noise observed by the RHUM-RUM experiment's ocean-bottom seismometers in the Indian Ocean during Cyclone Dumile in January 2013 (Davy et al. 2014). Instruments at similar distances from the cyclone's centre (but along different azimuths) on the seafloor recorded very different microseismic noise signals, indicating that the source-receiver distance was not the sole determinant

of noise behaviour at any particular location. To investigate the effects that local bathymetry may have, a pressure perturbation in the water column immediately below the cyclone's centre was used as a source, and receivers were placed at the locations of RHUM-RUM OBS instruments. By running such simulations both with and without bathymetry, it is possible to isolate the contributions that it makes to the observed noise signals from, say, the effects of mantle heterogeneity.

2.12.3 Results

Data from an initial simulation is shown in Fig. 2.15.

It is clear that stations at comparable distances from the source (but in different directions) record very different surface wave waveforms.

The signals presented here are Green's function solutions for an impulsive pressure source in the water column, and as such the frequency content is not directly comparable to observations. Convolution with an appropriate source-time function (in this case, a random spectrum or the like) is needed to make a direct comparison, and such efforts are ongoing. Nonetheless, these trends are not erased by such a convolution, and demonstrate that bathymetry does indeed have an effect on the noise signals recorded by the RHUM-RUM instruments.

Further work on this topic is ongoing, but it is anticipated that modelling of this type (which may be directly compared to data and is not possible using existing methods) will represent a significant advance in the modelling of bathymetric noise generation and propagation.

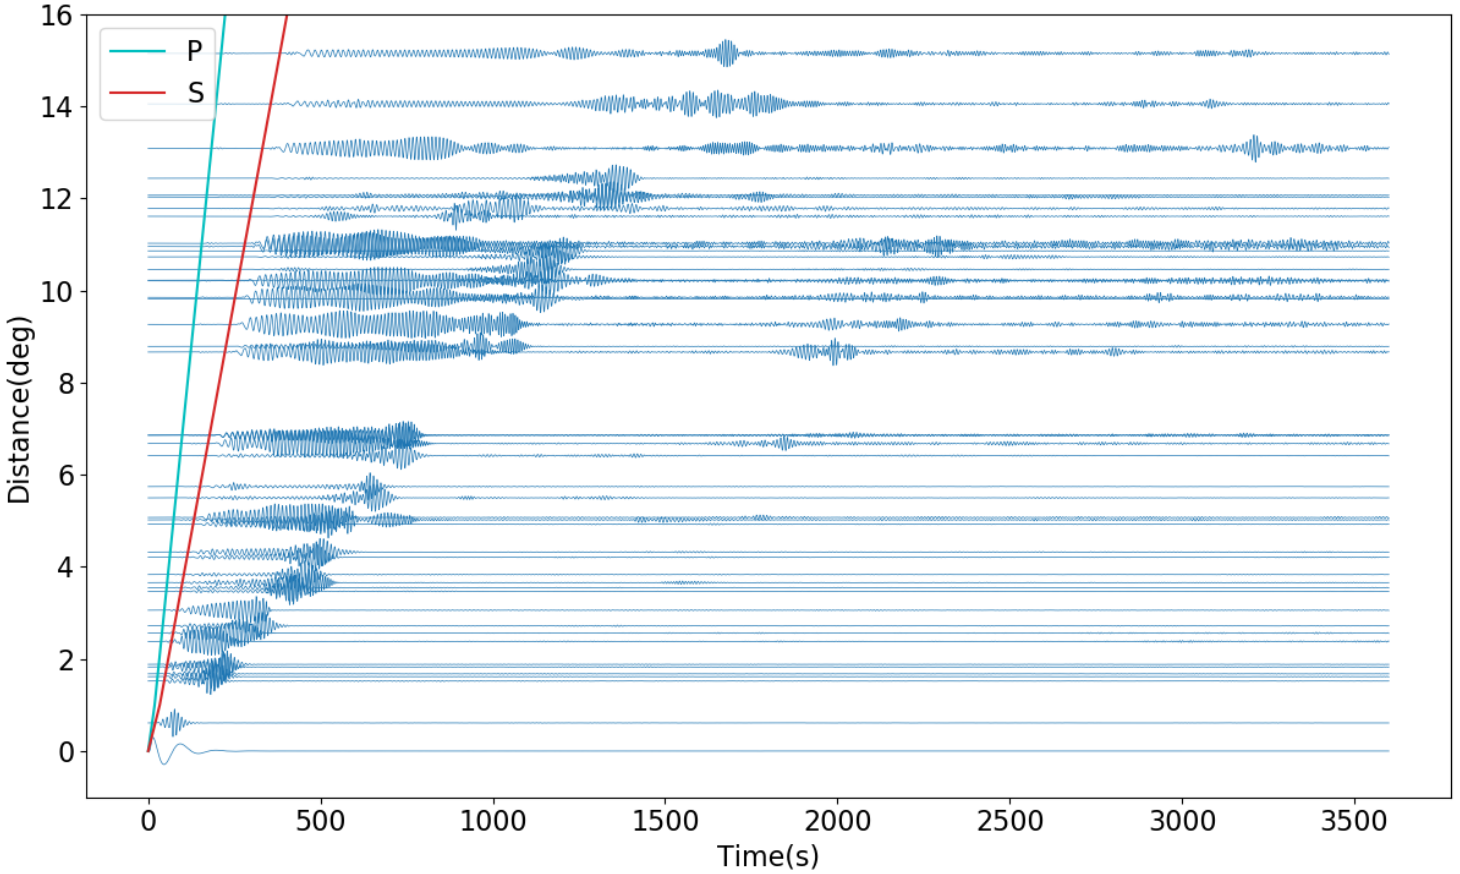

Figure 2.15: A one-hour record section showing the results of a simulation conducted at 10 s resolution. Traces are plotted as a function of station distance from the pressure perturbation source, with the P and S direct-arrival only travel time curves from TauP overlain. The vertical component of the seismograms is shown. Stations are realistically simulated as being on the ocean floor. Note that the waveforms shown are surface waves, with body waves being too low in amplitude to notice at these distances. This simulation includes a fluid ocean, bathymetry, and crustal structure.

The Viking Lander on Mars is a toy grown to large size, a metaphor of a dream; that dream of extending our will, our hand, our seeing eye to another world. It is not a machine, it is us.

— Ray Bradbury

3

Mars, Part I

Contents

3.1 Preface 70
3.2 Summary 70
3.3 Introduction 71
3.3.1 Martian seismology 71
3.3.2 The InSight spacecraft 73
3.3.3 Terrestrial background 74
3.3.4 Martian challenges 75
3.3.5 The InSight Blind Test 76
3.4 Pre-processing 78
3.4.1 Data read-in 78
3.4.2 Filtering 79
3.4.3 Initial rotation 80
3.5 Analysis 81
3.5.1 Event identification 81
3.5.2 Body waves 82
3.5.3 Surface waves 88
3.6 Results 89
3.7 Conclusions 90
3.8 Postface 93
3.8.1 Performance 93
3.8.2 Mars seismology in reality 93

3.1 Preface

This paper is primarily based upon a technical report submitted to the InSight mission's pre-landing Blind Test (Clinton et al. 2017) by the Oxford Seismology group, within which I was the project lead.

A summary of the Oxford group's efforts was published in as Fernando et al. (2018):

Fernando, B., Tsekhmistrenko, M., & Hosseini K. - **Training martian seismologists for InSight**, *Astronomy and Geophysics*, 2018.

It has the following DOI: *10.1093/astrogeo/aty230*.

The collected submissions to the blind test, and the analysis involved therein, were published in Seismological Research Letters as van Driel et al. (2019):

van Driel, M. et al (+ Fernando, B.) - Preparing for InSight: evaluation of the blind test for martian seismicity, Seismological Research Letters 90 (4), 2019.

That paper has the following DOI: *10.1785/0220180379*.

As this work was originally submitted to the group who had designed the blind test, it did not include background or an introduction, rather focussing on an exacting and technical description of the methods used. The introduction is largely taken from the summary of the Oxford group's efforts, as published in Fernando et al. (2018). This paper has not been updated to include post-landing data or analysis, though in some places a post-2018 reference is given for completeness and ease of reading where no alternative is available. My role within the Oxford team involved managing the project, identifying events, backazimuth and particle motion analysis, and writing the manuscript.

3.2 Summary

- In preparation for the InSight mission to Mars, a catalogue consisting of a year's worth of synthetic marsquakes has been released to help test inversion methods, build community participation, and develop new techniques for data analysis in the context of single-station planetary seismology.

- Catalogues were based on a single structural model chosen from a set of 14, though teams were not aware of which model was used from within the suite.
- Our effort uses differential travel times and particle motion for P and S-waves and dispersion curves fo surface waves to identify and locate 8 events within the dataset. A further 18 events are provisionally identified but cannot be located due to their low signal-to-noise ratio.
- For events where both body and surface waves are apparent, independent localisations are undertaken.
- We infer that the background model used in generation of the synthetic dataset is most likely either EH45TCold, EH45TColdCrust1b or DWTHot on the grounds that they fitted more than half the events which we analysed.
- This exercised reinforced the suspicion that event detection and identification would be especially challenging on Mars due to the single-station constraint.
- One particularly apparent area of challenge will be the discrimination between impact-generated seismic signals and those of a tectonic origin.

3.3 Introduction

3.3.1 Martian seismology

The last hundred years have seen enormous advancement in our knowledge of the Earth's internal structure. Today's geophysicists have at their disposal an enormous number of techniques with which we can probe the dynamics and composition of our planet's interior: gravimetry, geomagnetics, magnetotellurics and seismology, to name a few. However, when it comes to Mars, or indeed any of the other planets in the Solar System, we know comparatively very little about them.

This is largely because the most detailed geophysical investigations require high quality in-situ measurements. Mars missions are both expensive (NASA's latest attempt, InSight, will end up costing over US$ 800 million) and challenging

(historically, around 50% of the spacecraft launched to the Mars have failed). Orbital missions are certainly useful, but the lack of a global magnetic field limits the amount that can be learnt about the core from space, and whilst gravimetric surveys are available they are at far lower resolution than for the Earth. Thus, there is no substitute for surface experiments capable of recording precision seismic and geodetic data.

Attempts have previously been made to conduct seismic studies on Mars. In 1976, NASA's twin Viking spacecraft landed the first seismic packages on another planet. Unfortunately, they returned little useful data; on Viking 1 the instrument failed to deploy properly. Viking 2's seismometer did deploy, but because it was mounted on the spacecraft rather than properly coupled to the ground it mostly recorded the shaking of the lander in the Martian wind. This noise contaminated the entire dataset, enabling an upper bound to be estimated for the frequency of large 'marsquakes', but precluding detailed investigation of individual seismic events (Anderson et al. 1977).

Whilst the Apollo missions of the 1970s demonstrated the enormous science returns possible from the inclusion of seismometer packages on planetary missions (Nakamura et al. 1982), none of the missions which have successfully operated on Mars in the last 42 years have carried such an instrument.

This will change on November 26, when InSight will attempt to land on a broad open plane near the planet's equator known as Elysium Planitia (Banerdt and Russell 2017). Launched in May of this year from Vandenberg Air Force base atop an Atlas V rocket, it aims to shed light on topics which, at present, we know comparatively little about.

As an example, little is known about Mars' core (e.g. Arkani-Hamed and Olson (2010)). There is clearly no longer a strongly convecting metallic fluid region at depth, but it is unclear when this stopped being the case, and hence when Mars lost its magnetic field. The loss of the magnetic field would have had important consequences for the planet's atmospheric retention, and hence its past habitability (Fang et al. 2010). As a result, being able to constrain the

planet's evolutionary history from detailed measurements of its core properties would be invaluable to planetary scientists.

Additionally, many global-scale models of the planet's structure present a dichotomy in the expected temperature of the Mars' crust, between a geologically 'hot' one and a geologically 'cold' one, which in part reflects the discrepancy in crustal properties between Mars' northern and southern hemispheres (Watters et al. 2007). The temperature of the crust influences the planet's tectonic behaviour, for example in determining the depth at which the transition from rocks deforming under stress in a brittle manner to a ductile manner takes place, hence constraining where quakes may nucleate.

Other dilemmas include how long ago the planet lost its volcanism, why extinct volcanoes are so strongly clustered (Golabek et al. 2011), and what the degree of differentiation within the planet's interior is (Grotzinger et al. 2011). Terrestrial geophysicists are also very curious about such matters, as they offer both context and comparison to what is observed on Earth.

3.3.2 The InSight spacecraft

To investigate these questions, InSight carries onboard one of the most sophisticated science packages ever launched to the Red Planet.

The Heat Flow and Physical Properties Package (HP^3) aims to penetrate three to five metres into the Martian crust to measure the heat fluxes emanating from the planet's interior (Spohn et al. 2018).

The Seismic Experiment for Interior Structure (SEIS) contains instruments capable of measuring seismic waves at frequencies of up to 20 Hz during normal operation, and up to 100 Hz in the highest sampling rate mode (Lognonné et al. 2019). Initial commissioning data from SEIS, parts of which were switched on for testing during the cruise phase, are reported to look promising.

To ensure proper coupling to the ground and to avoid the problems experienced by Viking, a small robotic arm mounted on the spacecraft will lift SEIS onto the ground after landing, at which point it will be covered by a wind and thermal

shield (WTS). This is expected to dramatically reduce the amount of noise induced by dynamic shaking of the instrument in the wind and by thermal heating and cooling of the instrument (Murdoch et al. 2017).

Additional wind, pressure and temperature sensors are also carried to record data that will prove useful in the post-processing removal of some of the atmospheric contamination (Banfield et al. 2019). Other instruments being carried include two cameras and a radio experiment to track the planet's rotation, will prove useful in more accurately determining its moment of inertia, and hence its internal structure.

3.3.3 Terrestrial background

Even once seismic data has been gathered, its interpretation presents further challenges. On Earth, after a seismic event occurs data from receiving stations around the world can rapidly be combined to provide estimates of the source parameters like magnitude, focal mechanism and depth. This is done through analysis of the ground motions induced by the passage of different kinds of seismic waves past the receiver location.

Over time, the detection of waves emanating from many different seismic events allows for the development of a model of the Earth's interior. These models can include information about both the planet's gross structure, such as the presence of a liquid outer core; as well as information about local-scale heterogeneities such as magma plumes or sedimentary basins.

Although the signals from different events may look substantially different at first sight, their fundamental structure is normally similar. A three-component seismometer, that is, one capable of recording ground motion in the vertical direction as well as in the horizontal plane will first detect the arrival of body waves (initially compressional P-waves and later shear S-waves), followed by surface waves (Love waves, then by Rayleigh waves). These different waves are referred to as 'seismic phases'. These waves can be identified both by the relative timing of their arrivals and the components of the seismometer in which their signatures present themselves.

In the source-centred 'RTZ' coordinate system which is common in seismic analysis, the radial (r) direction is on a line pointing from the source to the receiver, the vertical (z) points outwards through the surface in the source-receiver plane, whilst the transverse component (t) points out of the source-receiver plane to complete the triad. P and Rayleigh waves oscillate almost exclusively in the r-z plane, whilst SV and Love waves are most prominent in the transverse component.

This picture is somewhat complicated by the reflection and diffraction of seismic phases off of various interfaces. For example, PP is a P-wave which has bounced once off the planet's surface before reaching the receiver, whilst PKP is a P-wave which has travelled through the outer core. Additional modifications to a signal can be introduced by seismic anisotropy, whereby rocks respond differently to waves incident from different directions - this can cause a fraction of the energy to 'leak' from one component to another. Nonetheless, by measuring the components of oscillation, relative arrival times amplitudes and phases of the different types of seismic wave it is possible to differentiate between them.

Further investigation can be conducted using synthetic modelling, whereby simulations are run for a variety of different source and structural parameters. A fitting can then be performed to find the closest match to the data. However, because the parameter space to be searched is so enormous and a unique inversion cannot necessarily be defined, this can be a computational challenge and on Earth is normally only done with some a priori knowledge of the likely input parameters.

3.3.4 Martian challenges

On Mars, however, the situation is very different and the inversion is much more challenging. InSight consists of only one receiving station, meaning that multi-station triangulation is impossible (e.g. Panning et al. 2015; Khan et al. 2016). Secondly, the a priori knowledge of the planet's structure is less well constrained than for the Earth, making the parameter space that must be searched when performing a fit much larger.

Additional difficulties are introduced by the operating environment: the dominant atmospheric and thermal noise on Mars is very different to the ocean-generated noise seismologists are experienced in dealing with on Earth. Additionally, the planet's atmosphere is much more rarefied than the Earth's, and as a result has a much lower heat capacity. This is expected to create extremely strong wind signals during the large-scale temperature changes around sunrise, which may overprint any seismic signals present.

The nature of seismic excitation and attenuation may also be quite different, and hence the analysis techniques developed for use on the Earth may require some adjustment.

In the case of the former, the thin atmosphere provides far less heating to incoming projectiles, enabling more meteorites to reach the surface than on Earth (Stevanović et al. 2017; Daubar et al. 2018). This means that impactor events will likely generate a larger proportion of the seismicity on Mars than on Earth, and will need to be analysed differently, with the important source quantities being impactor size, angle of entry and explosion altitude rather than depth and rupture magnitude.

In the case of the latter, a significant portion of the seismic attenuation on Earth comes about from water trapped along grain boundaries in the rock. The expected differences in the water content of the crust and mantle on Mars may change the planet's attenuation profile, and this would need to be investigated and accounted for to ensure inversion for the source parameters is robust.

3.3.5 The InSight Blind Test

Clearly there are a number of challenges involved in analysing InSight's data: for example given the constraints of having a single station, or of having potential impact signatures mixed in with tectonic events. Whilst some terrestrial seismology techniques (such as the minimisation of P-wave energy on the transverse component) have the potential to be directly applied to Mars, others (such as multi-station triangulation) this clearly cannot be done. For others, such as using P-S arrival

time differences to infer distances, it is unclear how good the spatial localisation of such a method might be, given the uncertainty in structural models.

For this reason, the InSight team decided to run a 'blind test', where members of the community would have the chance to develop methods for analysing synthetic data and to make inference about Mars' structure, in preparation for the return of genuine data next year (Clinton et al. 2017). This would also enable some of the uncertainties surrounding the usability of the methods themselves in an unknown structural setting to be quantified, by using a range of possible interior models.

Fourteen published and potential one-dimensional models of Mars' interior structure were chosen, and travel-time curves for the different seismic phases were provided to the participants for each possible model. From these models, one was chosen and used to generate an Earth year's worth of broadband synthetic Martian seismic data; however the participants were not told which model had been used.

Participants were told that the year's worth of data contained roughly 250 events with a variety of different source parameters, with randomly distributed locations. Magnitudes and event depths followed a skewed Gaussian profile, with peaks in magnitude of around 2.5 and in depth of approximately 35 km. The largest events were stated to be no larger than M_w 5 and the deepest no more than 100 km. Synthetic seismograms were generated using AxiSEM (Nissen-Meyer et al. 2014a).

A model of the dynamic response to the wind shaking the probe was used to create synthetic atmospheric noise, which was superimposed on the seismic data to simulate realistic data returns. The three-component seismometer data and wind pressure data were provided at a frequency of 2 Hz, whilst the wind speed, direction and temperature were given at 0.1 Hz and the magnetic field at 0.2 Hz.

The aim was for each of the participating teams to identify as many of these events within the data as they could, and where possible locate them on the planet and determine source parameters. The end result was expected to be correct identification of which of the 14 models was used in the data generation.

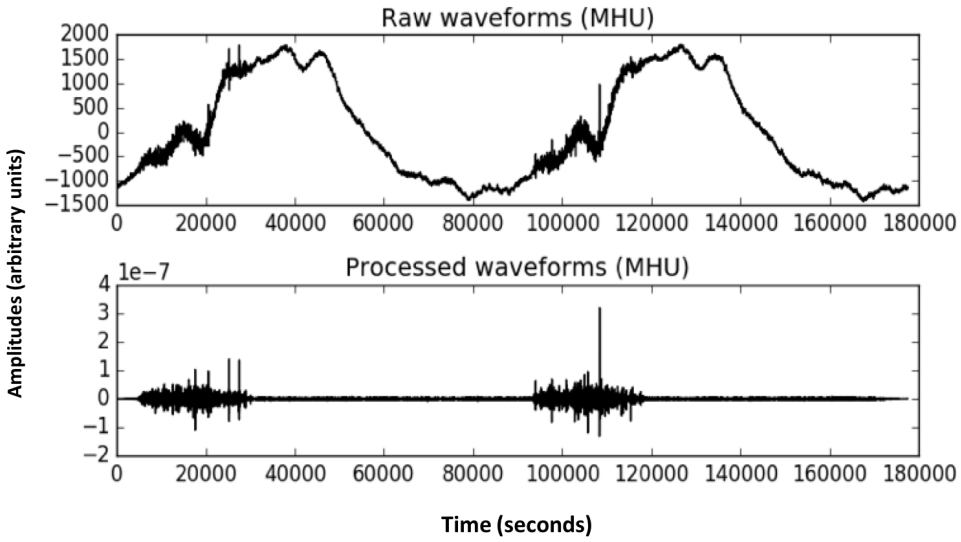

Figure 3.1: Raw waveforms (top) and processed waveforms (bottom) for the MHU component of the spacecraft seismometer for two days, Sols 052 and 053, in which no event is immediately obvious. The removal of the strong diurnal trend after processing is clear, though several hours of intense wind noise remain each day.

3.4 Pre-processing

We choose to analyse the data using a mixture of custom-written algorithms and those from the Python ObsPy seismic package. Henceforth, all references to specific days of data are given in Martian Solar Days ('Sols', equivalent to 24 hours and 39 minutes).

One component of the raw, unprocessed velocity data for two Sols (approximately 180,000 seconds total) are shown in Fig. 3.1 (upper panel). The coordinate system 'UVW' is spacecraft-centred and not geographically meaningful, but a number of phenomena can immediately be identified. The first is the clear diurnal trend caused by thermally induced expansion and contraction of the ground, with the observed background amplitude rising and falling over the course of a sol.

3.4.1 Data read-in

Removal of the instrument response was done on a sol-by-sol basis, using the *ObsPy remove_response* method. This step also computes the on-the-fly integration to

move from velocity to displacement representations. Overamplification is prevented using tapering in the frequency domain, with the four corner frequencies of the taper set at 0.008 Hz, 0.012 Hz, 1.0 Hz and 2.0 Hz. Similarly, removal of the daily trends presumably induced by atmospheric temperature variations used the *ObsPy* *detrend* and *taper* methods, with the taper fraction set at 5%. These data shall be henceforth referred to as 'pre-processed seismograms'.

3.4.2 Filtering

Clear periodicities in the noise are visible. The most significant is a period of around 6 hours at the start of each sol, which from the supplied air speed and pressure data is attributable to wind shaking the lander. Similar behaviour was observed by Anderson et al. (1977). We refer to this region as having 'wind noise', whilst the noise in the remainder of the day is considered 'quiet noise'. The former displays strongly time-varying amplitudes up to $5\text{x}10^{-6}$m in the pre-processed waveforms, whilst the latter is more constant and at least an order of magnitude lower in most cases. A plot of the noise spectrum is shown in Fig. 3.2.

Taking a Fourier transform of a pre-processed seismogram in which no event can be obviously identified shows two regions of pronounced noise in the spectrogram - between 0.008 and 0.08 Hz and at around 1 Hz. We suppose that the latter noise band is an aliasing response arising from the instrument response deconvolution. A number of different filters, including series of Gabor filters, were tested to try and optimise the event signal-to-noise ratio. A four-corner bandpass filter with a range of 0.08 Hz to 0.5 Hz was found to be optimum, and all subsequent data analysis was filtered with this bandpass. The amplitude of both the wind noise and the quiet noise are thus reduced significantly, though not entirely eliminated.

Following filtering, the average noise floor throughout each sol is much reduced. A large event on sol 257, being obvious above the noise floor across a broad frequency range, was used as an exemplar in the development of subsequent filtering and data processing methods. Successive plots show the seismograms from this event.

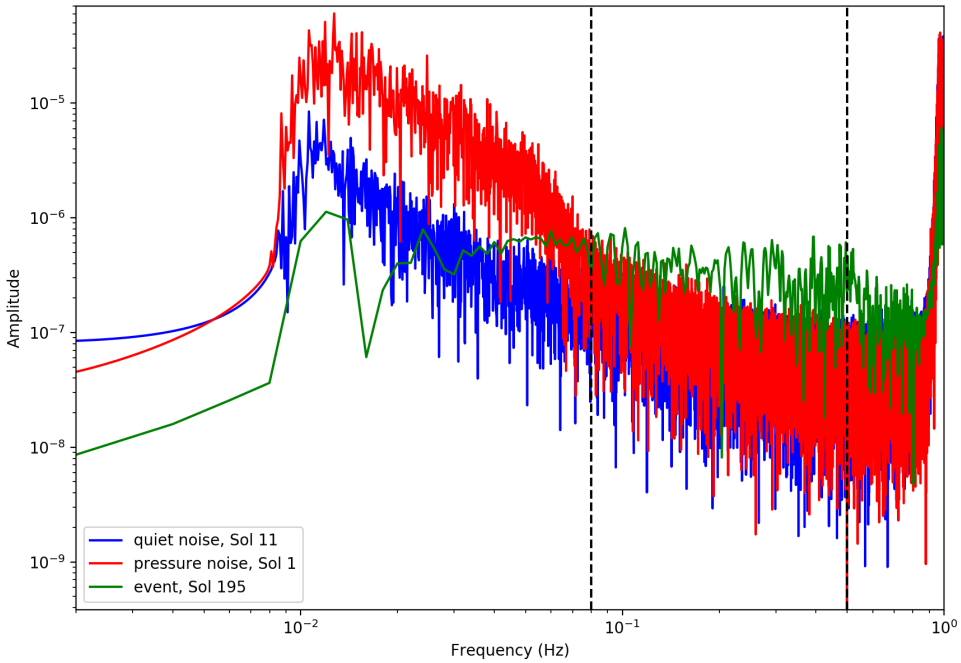

Figure 3.2: The quiet-noise (blue, from sol 11) and noise in the pressure spectrum (red, from sol 1) spectra are shown in comparison to the spectrum of a marsquake which occurred on sol 195. The dashed vertical lines indicate the passband limits of the filter used on the data.

3.4.3 Initial rotation

The next step performed was to convert from the provided data format (UVW) to easting-northing-vertical (ENZ). This represents a transformation from the coordinate system used by the spacecraft's seismometers to a geographically meaningful coordinate system. This was done using the rotation matrix shown below (similar to Wielandt (2009), though with a 30° offset from North for the U-channel seismometer), noting the positioning of the three seismometers specified in the metadata. For the exemplar event on sol 257, the results are shown in Fig. 3.3 for a window of approximately 400 s on either side of the time of peak amplitude.

$$\begin{pmatrix} N \\ E \\ Z \end{pmatrix} = \begin{pmatrix} sin(54.7°)cos(30°) & sin(54.7°)cos(150°) & 0° \\ sin(54.7°)sin(30°) & sin(54.7°)sin(30°) & -sin(54.7°) \\ cos(54.7°) & cos(54.7°) & cos(54.7°) \end{pmatrix} \begin{pmatrix} U \\ V \\ W \end{pmatrix} \quad (3.1)$$

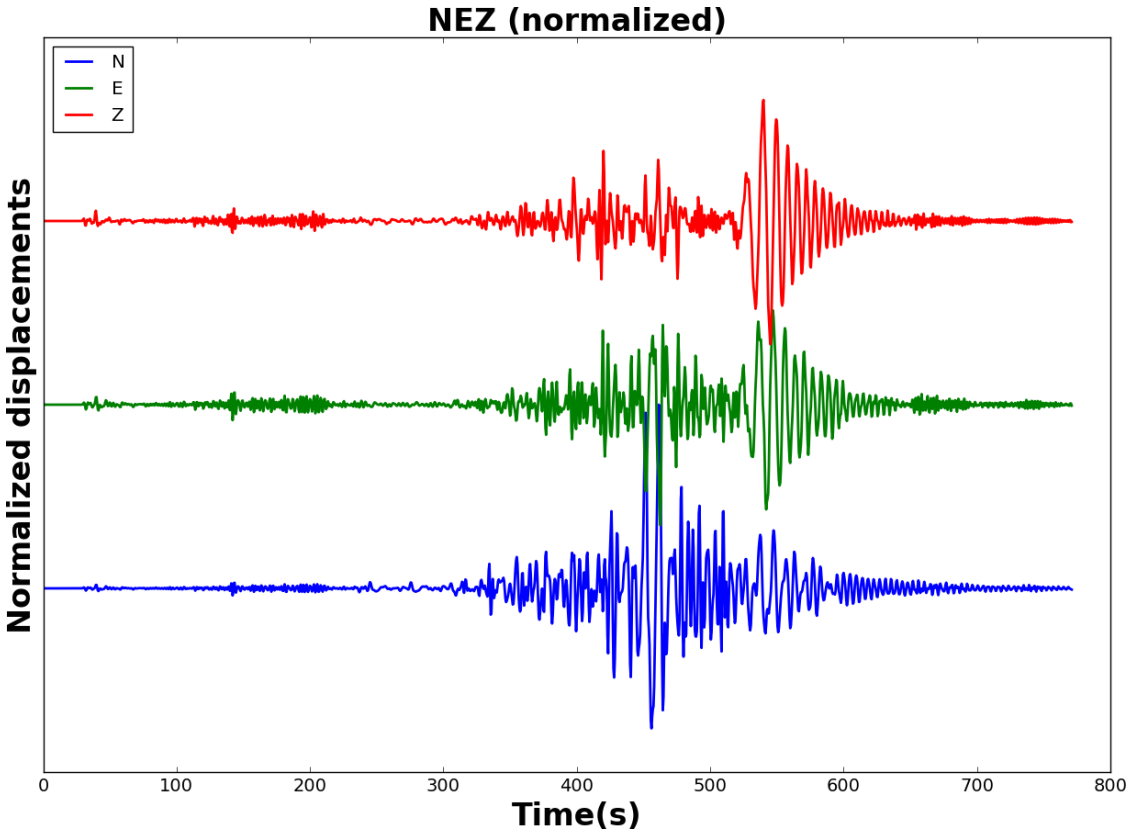

Figure 3.3: ENZ (easting-northing-vertical) components of the large event on sol 257. Detail of the large event on sol 257. The 'start' of the event is picked manually to be the first point at which the wave amplitude rises significantly above the background. The P-waves are clear approximately 50 seconds in, followed by the S-waves at approximately 250 s. The largest amplitudes are formed from the Love waves (first) and Rayleigh waves (second, with long coda).

3.5 Analysis

3.5.1 Event identification

Attempts were made to use an STA/LTA algorithm to pick out events in the filtered dataset. However, we were not able to remove the spurious peaks induced by the wind noise. Specifically, it appears that some sols contain a short duration (order 1000s) of high-amplitude noise with a distinctly different frequency content to the other windy sections. We were not able to develop a method for removing this within the time available, and hence the STA/LTA was unusable.

As a result, identification of events in the filtered dataset was only conducted visually. This was made somewhat easier by discarding signals in the velocity

channels which were strongly correlated with those in the pressure channel. However, this correlation was done by eye and hence significant room for improvement exists.

Alternative techniques which were tried and which failed include cross-correlating the waveforms from large, obvious events with the data as a whole (and hence attempting to find areas with similar properties elsewhere in the data), and stacking seismograms from multiple sols. In the case of the former, the strongest peaks in the cross-correlation again came at times where there was also a strong peak in the pressure channel (suggesting a wind-noise origin). In the case of the latter, the loss of temporal resolution meant that the data were not useful. We considered comparative stacking (measuring the residuals between the stacked seismograms from all sols and those from all but one sol) to try and resolve the days on which significant events occurred, as it was hoped that the average of the quiet noise would tend to zero if it were randomly generated. However, this proved to only be sensitive to the largest events, which we could already identify visually.

3.5.2 Body waves

For events where both P and S-waves were clearly visible, we were able to infer event azimuth, epicentral distance and the incidence angle of the arriving rays. Where possible, more than one method was used to ensure the accuracy of the inferences.

3.5.2.1 Backazimuth

We used three separate techniques to estimate the backazimuth of the source. These all rely on accurate manual selection of the times of P-wave first arrival.

1. **Minimisation of P-wave energy** - it is expected that P-waves from an event will oscillate primarily in the radial-vertical plane of a source-centred (RTZ) coordinate system, assuming no anisotropy. By performing grid-search rotation about all possible backazimuth angles in 1° increments, it is possible to find a pair of angles 180° apart along which the t-component energy of the first arrival is minimised (small deviations from zero in this channel are expected due to a variety of effects). Thus, the azimuth of the source is

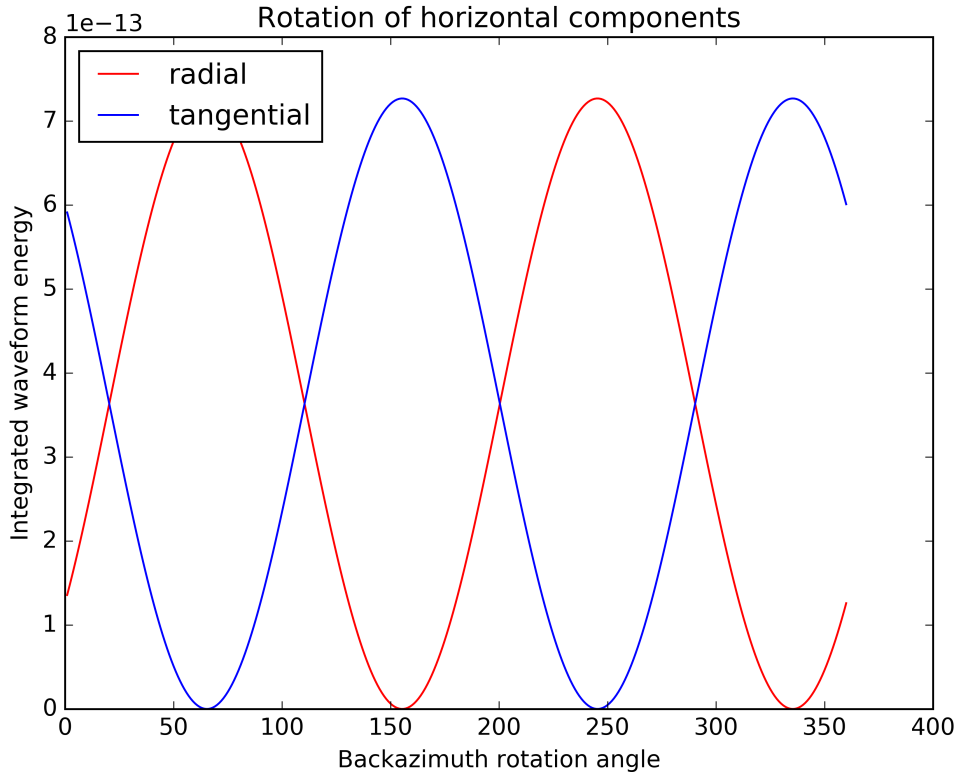

Figure 3.4: Output of the rotation grid-search, showing the energy integrated over the P-wave train. Two backazimuths corresponding to a minimum of energy in the t-channel and a maximum in the r-channel are clear at approximately 65° and 245° from north. This defines the line along which the azimuth of the source lies.

constrained to being along a line, though this method in itself does not tell us from which direction along that line the signal originates. An example of the grid-search output is shown in Fig. 3.4.

2. **ObsPy Particle Motion** - Using the build-in methods of ObsPy's particle motion analysis suite, it is possible to constrain the backazimuth again to lie along a line.

3. **Custom-developed Particle Motion** - We developed a routine in Python for analysing particle motion which allowed estimation of both the source azimuth and the associated error (the latter not being offered by ObsPy). A 3D point cloud representing the particle motion is plotted, and a three-dimensional orthogonal distance regression is fitted using a singular value

decomposition upon the cloud. By finding the eigenvectors of the covariance matrix of the point cloud, we are able to determine the dominant directions of motion. Two views of the point cloud, showing the dominant axes, are shown in Fig. 3.5.

For most events, these three methods are in good agreement. Identifications with larger errors correspond to events with smaller signal-to-noise ratios for the P-waves and hence less well-defined first arrival times, leading to poorer constraints on the duration over which to minimise the t-channel energy.

Using the given backazimuth, we perform a rotation from the ENZ coordinate system to a source-centered RTZ one. For the exemplar event on sol 257, the results are shown in 3.6 for a window of approximately 400 s on either side of the time of peak amplitude.

3.5.2.2 Angle of incidence

As with the backazimuth, two methods of finding the angle of incidence of the first arrival are tested. These agree very closely, allowing us to determine in which direction along the line of azimuth the source originated from (in effect, the P-wave rotation constrains the source to be somewhere along a great circle, this methods tells us which way along said great circle energy arrived from). The first method uses ObsPy's particle motion features, whilst the second is an extension of the point cloud method described above.

3.5.2.3 Epicentral distance

One method of estimating the epicentral distance of identified events uses the travel time delay between P and S-waves. For those events where both P and S phases are clearly visible (the former lying predominantly in the r-z plane and the latter in the t component), the differential travel time is measured. Unlike the methods described thus far, this process is far more sensitive to the structural model used in the generation of synthetics, as this controls the travel-time differences between P and S-waves.

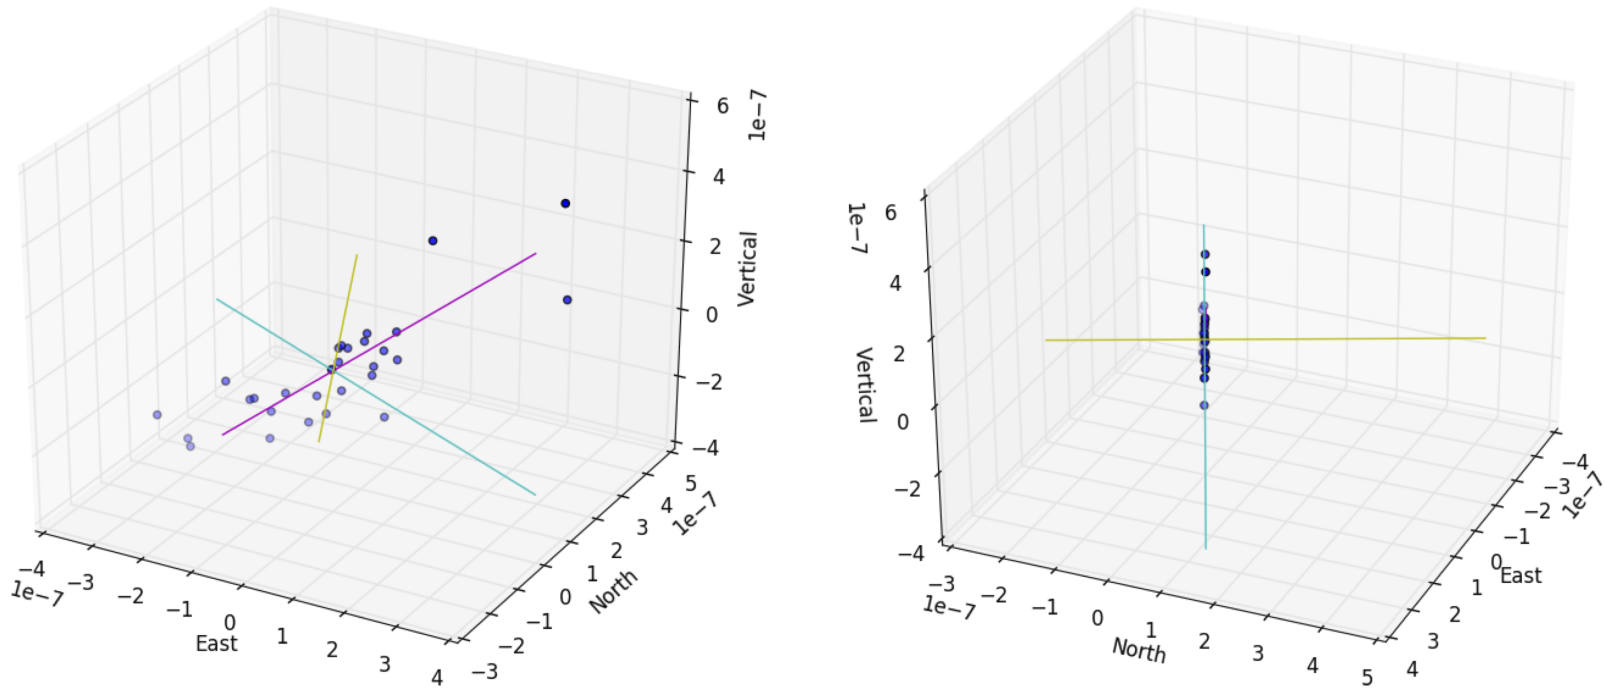

Figure 3.5: Two views of the point cloud for the sol 257 event, with the dominant axis shown in purple. The positions of the points in a plane is clearly visible in the right figure. Note that in all cases the axes plot displacement in arbitrary units.

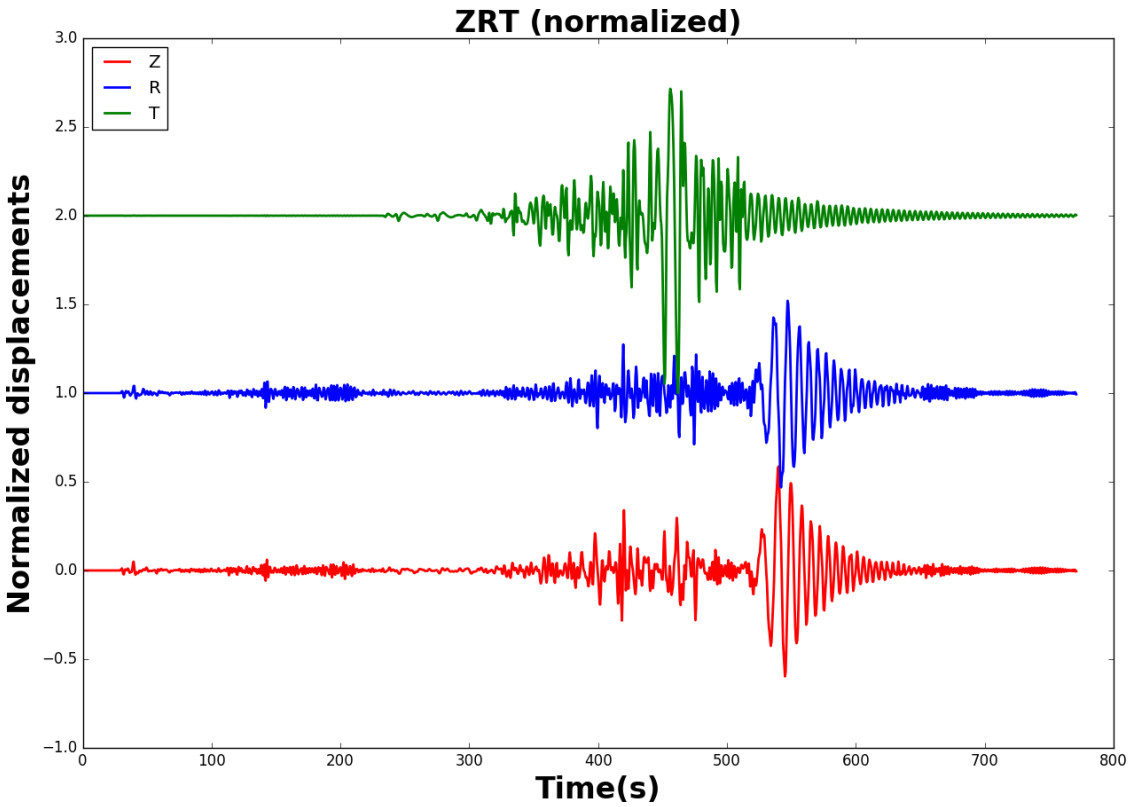

Figure 3.6: RTZ components of the large event on sol 257. The presence of P-waves (around 20 seconds) in only the r-z plane and S-waves predominantly in the t-channel is clear. Love waves, predominantly in the t-channel are also the first surface waves to arrive (400-500 s), followed by Rayleigh waves in the r-z plane (500+ s). We suspect that the structure of the coda on the Rayleigh waves is a known artefact of AxiSEM synthetic generation.

For each of the 14 background models, a catalogue of expected differential travel times is produced for 1° steps in epicentral distance and 1 km increments in depth using *TauP*. Those background models for which an expected travel time can be generated which matches the observed one with an error up to 15° in incidence angle and 15 s error in differential travel time are considered 'candidate models' (except for the sol 12 event where both tolerances were increased to 20° and 20 s due to a poor fit).

For some models, there is no solution for a particular event (indicating that for that model, the event would have to originate in either the P or S-wave shadow zones, for which there would necessarily be no phases of that type visible. Observations of phases within the predicted shadow zones would indicate that said model is unlikely

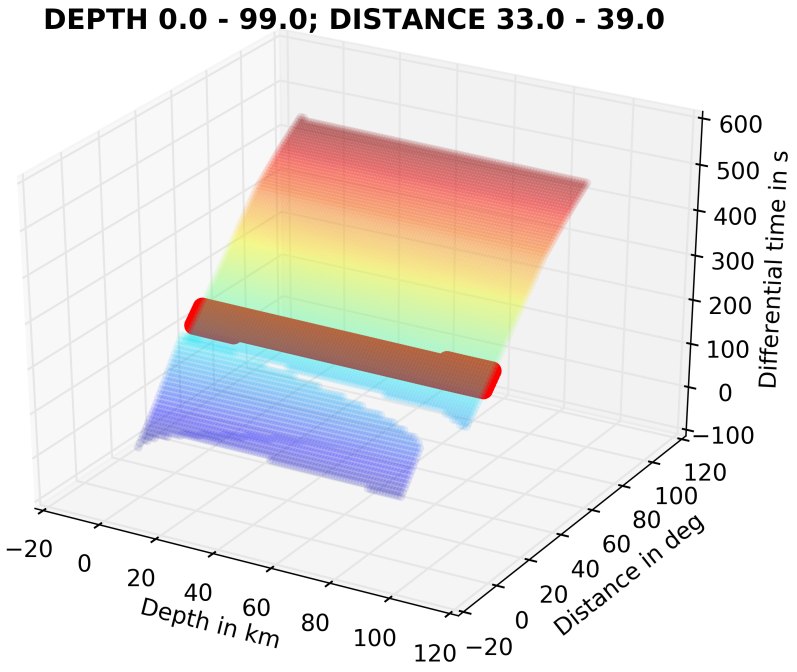

Figure 3.7: The travel-time space for different combinations of depth and distance in the EH45TCold structural model (one of 14 examples, which for this event provided a good fit). The gradation runs from red (longer travel times, more distant events) to blue (shorter, closer). The red shading represents the range of possible origins in parameter space for the sol 257 event based on its measured travel time difference. The white wedge represents a shadow zone from which no waves are expected. For this event and structural model, the observed travel time differential is compatible with any depth between 0 and 39 km, whilst the distance is constrained to be between 33° and 39°. Note the greater sensitivity of this method to distance than to depth - the lines of constant travel time are nearly parallel to the depth axis.

to have been the one used in the generation of synthetics. An example of the *TauP* output and constraints imposed by the event are shown in Fig. 3.7.

For those events in which surface waves are also clearly visible, the P and S arrivals need not be distinctly observable to estimate the epicentral distance. When both are visible, an independent verification is possible.

3.5.3 Surface waves

In many of the events, surface waves are clearly identifiable. In some cases, discrimination between Love waves (arrive first, t-component variation only) and Rayleigh Waves (arrive later, r-z plane oscillation) is possible. This allows us to use the dispersion of the wave trains to infer the epicentral distance, independently of the body-wave methods described above. Additionally, we are able to eliminate a number of potential background models by discounting those for which the dispersion is not consistent with that which is observed.

All analysis done thus far uses only the minor arc (i.e. waves which have travelled the 'short' way round the great circle, often denoted G1 and R1), though in some events the major arc also appears to be visible at much lower amplitude. Note that for consistency, even if visible the major arc is not included in our definition of the overall event duration. Where major arc can be positively identified, this would offer a possible check on the dispersion-based results for each event, however this is not yet implemented in our analysis.

Separate dispersion curves are created for each surface wave by filtering each pre-processed seismogram to achieve optimum signal-to-noise ratios. A tapering window is then used to isolate the event in the time domain. The power spectrum of each event is calculated in rolling 0.5 second windows to ascertain the point of maximum amplitude in each of 20 linearly-spaced frequency bands between 0.005 and 0.1 Hz. Frequency bands less than 0.02 Hz are excluded from the analysis, as they sample at less than the Nyquist frequency imposed by the use of the tapering window. The dispersion spectrograms for both wave types are shown in Fig. 3.8 for the exemplar event on sol 257.

Delay times between Love and Rayleigh waves are calculated in each remaining frequency band. Using the given Love and Rayleigh wave group velocities for the different frequencies, a linear regression is performed to evaluate the epicentral distance at which the time delay between the two waves best matches the expected delay. This is done for the group velocities given in each of the 14 models (except

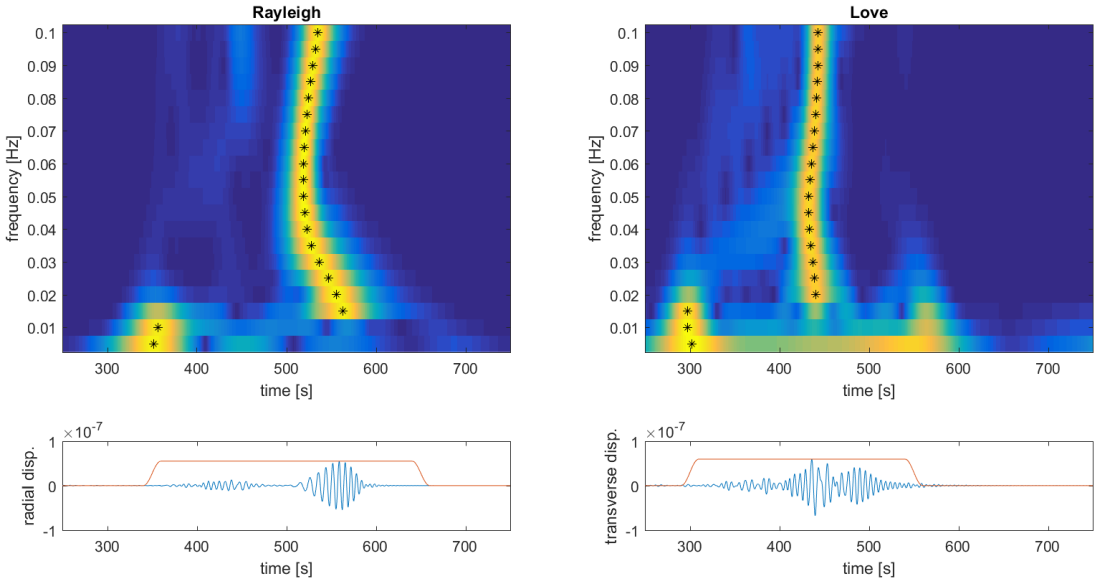

Figure 3.8: Dispersion spectra for the sol 257 event. The upper panels show the spectra, Rayleigh on the left and Love on the right. The lower panels show the waveforms in blue and the tapering window used in red. Note that the lowest frequency bands (<0.02 Hz) sample at less than the Nyquist frequency, so are not included in analysis.

'*Gudkova*' which has no Love wave data given). Those models for which there is no good fit for the observed travel time delay are discounted.

For the event on sol 257, a comparison between the observed and predicted surface wave dispersion is shown for one badly fitting model (Fig. 3.9, SANAK) and two well-fitting models (Fig. 3.10; EH45TCold and DWTHotCrust1b respectively).

3.6 Results

Using the procedures detailed above, we find 26 events in the Earth-year long worth of data. Of these, 8 are localisable in space (backazimuth and epicentral distance) using body waves, whilst for the remaining 18 body wave phases are not clearly identifiable to the degree needed. Three of the eight localisable events are also subjected to a surface wave dispersion analysis to determine the epicentral distance.

The best-fit models for the surface wave analysis are in good agreement with the best-fit models for the body wave analysis; though the two models are not necessarily the same. However as the best-fit cases in both are variants of 'EH45TCold' and 'DWTHot', we suggest that the discrepancies are due to small changes in the crustal

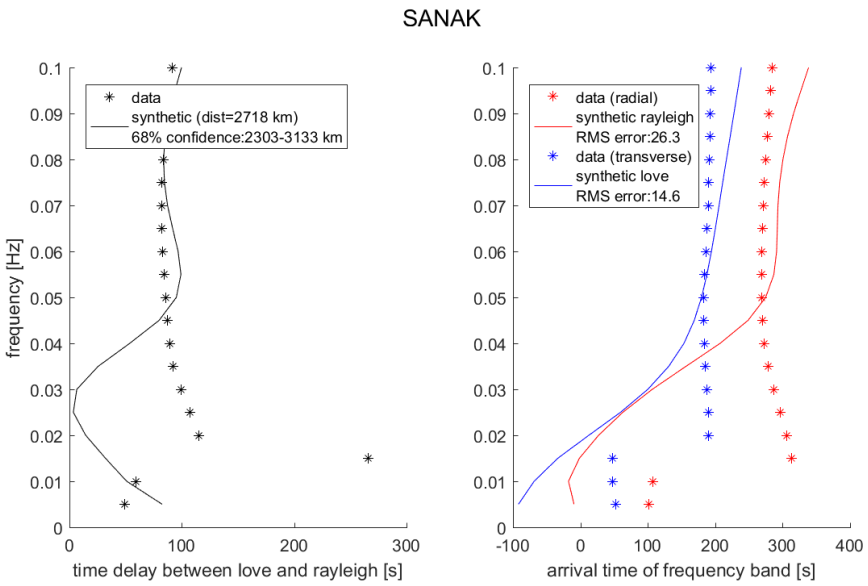

Figure 3.9: The observed dispersion curve for the sol 257 event and that predicted from the SANAK model. The left panel shows how well the travel-time delay (crosses) match the model predictions (solid line), in each frequency band. The right panel shows the relative arrival times of each frequency band, against the model predictions. It is clear that neither the travel-time differences between G1 and R1, nor the dispersion within each phase, matches the model predictions.

velocity profile near the surface, which change the refraction of the waves and hence can significantly alter the incidence angle, especially for shallow incoming waves expected from teleseismic events. This implies a shortcoming in the ray theory methods used to locate large events, and full waveform modelling may be more appropriate. The number of events that each model was successfully able to fit is shown in Fig. 3.11.

3.7 Conclusions

Three models fitted more than half the events analysed: EH45TColdCrust1, DWTHot, and EH45Cold. We consider that these are the three models most likely to be the 'correct' one, noting that two are based on very similar geological models.

We suspect that there may be further bias in our methods for events within a certain distance, however without a larger sample we suspect that any statistical analysis will lack meaningful significance. Given the limited time the authors were

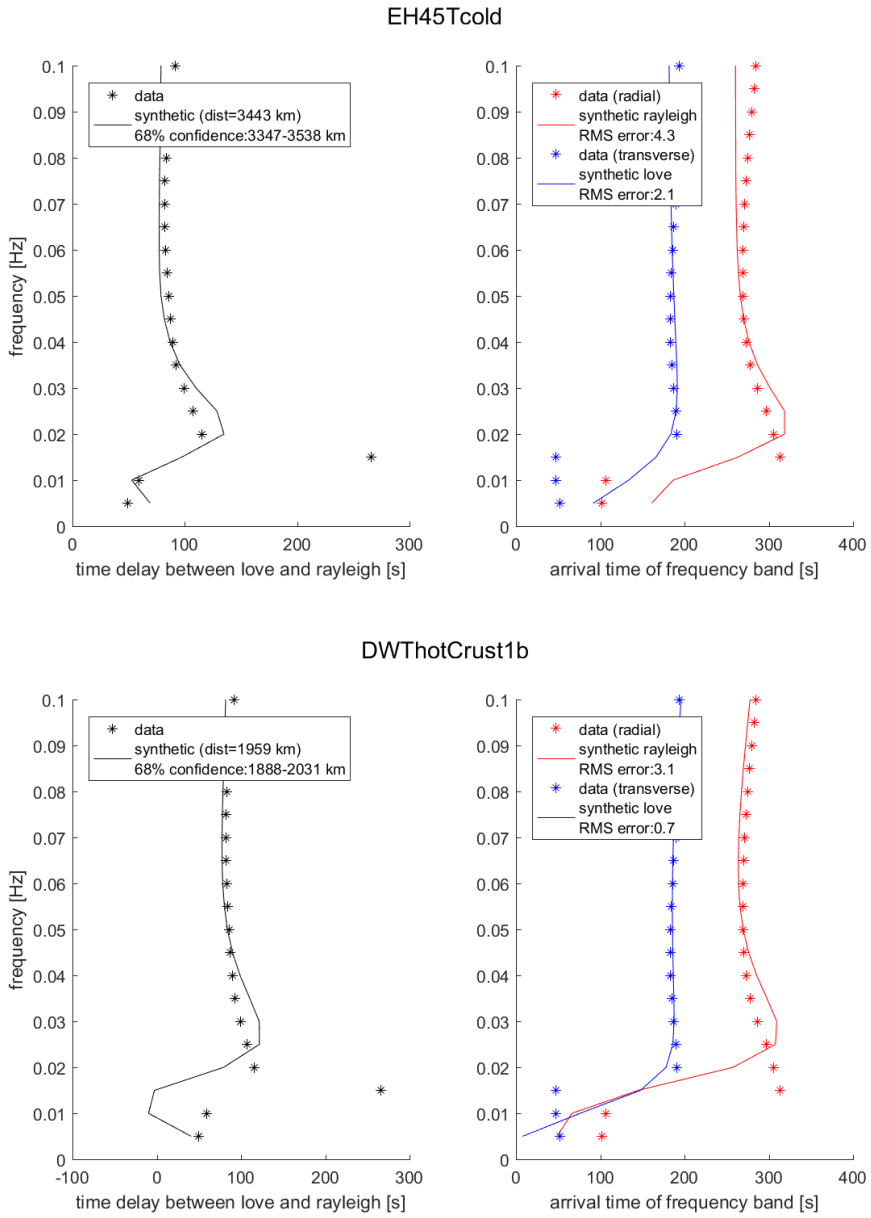

Figure 3.10: Observed and predicted dispersion curves for both Rayleigh and Love waves in two models (EH45TColdCrust1b at upper, DWTHotCrust1b at lower). Note that whilst both models give a good fit to data, the distances they predict based upon this dispersion analysis (1959 km and 3443 km respectively) are very different. This highlights the particular challenges of using only a single station in inversions - without more events, more stations, or independent constraints, this method offers no information on which is the 'true' distance (and hence, which model is correct). Data below 0.02 Hz is not included in analysis.

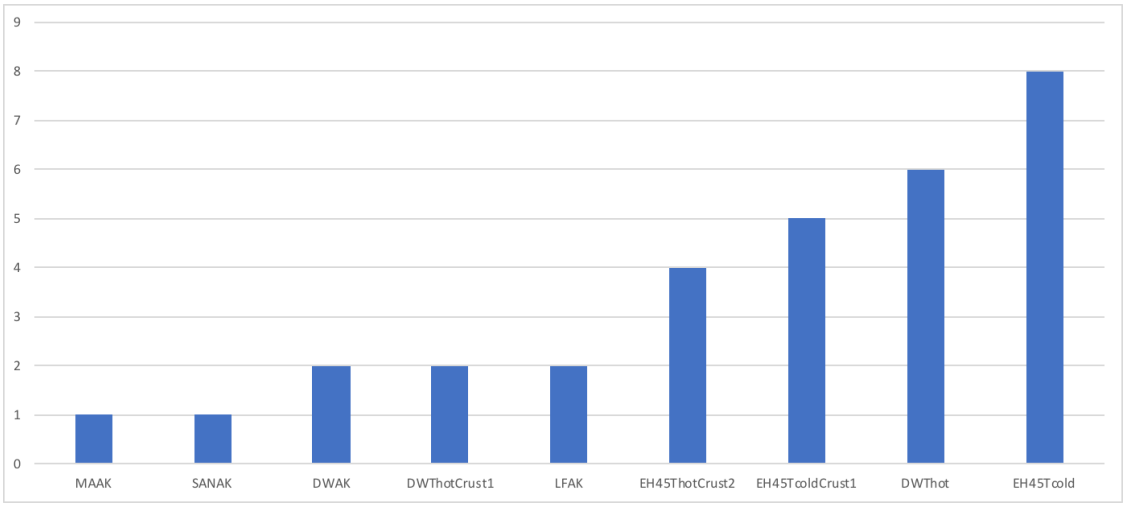

Figure 3.11: The number of events successfully fitted by each background model. For models not listed, no events were compatible with the given seismic structure to within the error bounds used. EH45TCold fitted all events analysed.

able to invest in the blind test alongside their PhD studies, we were unable to carry out all the methods of analysis that we hoped to. Nonetheless, we suggest a number of potential improvements:

- Generation of synthetic waveforms for each background model would prove useful in determining the expected waveforms for P and S-wave first arrivals. The lack of a current training set means that picking is a challenge, and may introduce significant error into results through mis-timings of first arrivals or even potential misidentification of phases, especially given the authors' lack of experience in phase picking.

- A multifrequency approach may prove beneficial for P and S-waves, whereby different frequency bands are analysed separately to determine epicentral distance and backazimuth.

- It should be possible to constrain event positions using surface wave travel times. This method is not robustly tested on Earth, but should be considered for development.

- Extending the surface wave analysis to seek major arc (G2/R2) or even multiple orbit (G3/R3 or greater) waveforms may be useful in further constraining

source distance, as the increased travel time will accentuate the effects of dispersion. However, detection may prove difficult for all but the very largest events.

- Event magnitude estimation should be possible given the observed peak waveform amplitudes and derived epicentral distances. This is likely very sensitive to the attenuation parameters, but these may be independently estimated.

- Measuring the ratio of P-wave to S-wave amplitudes should enable differentiation between impact (point-force) and tectonic (extended) sources, assuming that non-linear effects generating shear waves from impacts are small on Mars (as they are on Earth).

3.8 Postface

3.8.1 Performance

As per (van Driel et al. 2019), 84 teams registered for the blind test. The dataset itself consisted of 204 synthetic marsquakes of geological origin, as well as 36 impacts, though the majority of these were around or below the noise floor. The Oxford team identified 21 events, of which two were within the L1 (Level 1) mission requirements for quake location. Fig. 3.12, taken from van Driel et al. (2019), shows the performance of our effort. This ranked amongst the top six submitted catalogues for completeness.

The actual model used was EH45TColdCrust1b (Sanloup et al. 1999), which we identified as one of our three most likely candidates.

3.8.2 Mars seismology in reality

In terms of this thesis, one of the most accurate 'predictions' from the blind test was that distinguishing between impact-generated and tectonic-generated signals in InSight data was going to be extremely challenging. Indeed, as of mid-2021 no unambiguous impact signal has been detected in the data. One likely reason for

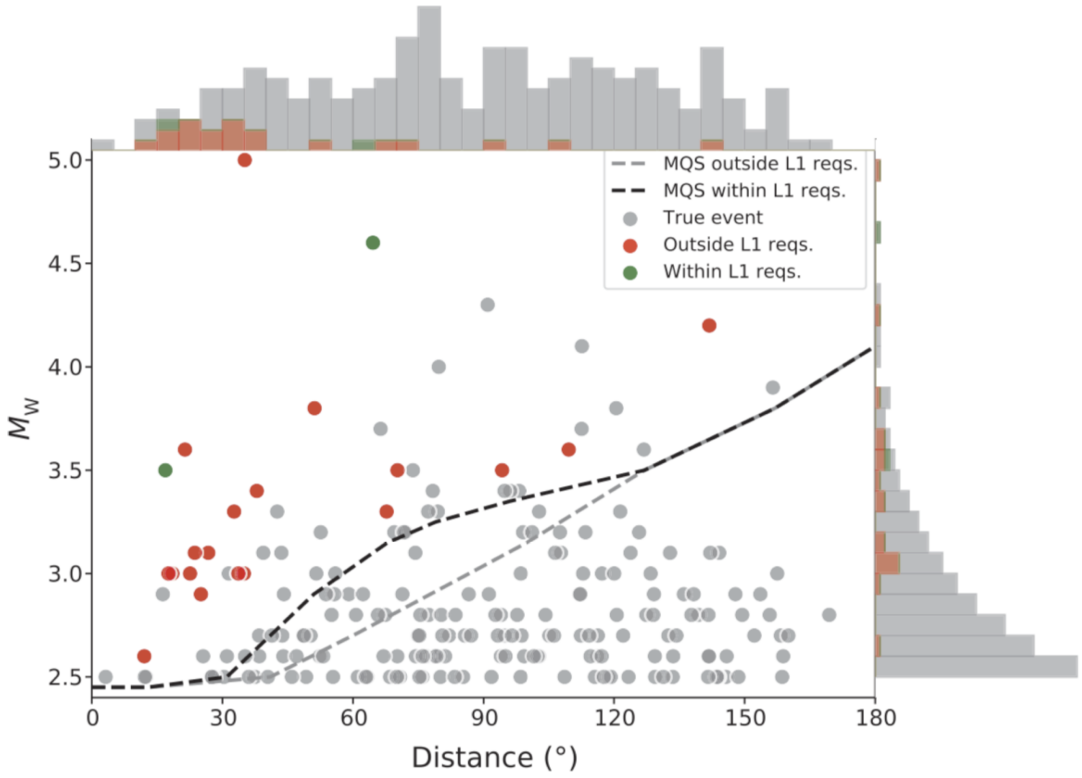

Figure 3.12: The results of our submission to the Blind Test catalogue. Correctly detected events are in red, correctly identified events are in green, and those missed are in gray. The dashed black and gray lines indicate the location and detection thresholds respectively for the Marsquake Service (MQS), who handle data returns from Mars. Right and top histograms show performance as a function of distance and magnitude.

this is that the impactor flux at the surface is likely to be lower than was previously thought (Daubar et al. 2020), but it remains unclear whether or not impact signals would be identifiable even if detected. This is very much an ongoing debate, and we will touch upon it later in the coming chapters.

The blind test also assumed that marsquakes would be similar in presentation to earthquakes - with many obviously distinguishable phases. In general this has not been the case - whether or not any surface waves exist in the data is another ongoing debate, and even body wave arrivals are harder to pick than on Earth. In this regard the blind test was not as illuminating as it would have been if martian seismology was more earthlike, but adding the necessary complexity (3D effects, thin low-velocity surface layers) to the structural models would probably have rendered it cumbersome, and the inversions too complex for such an open, community project.

Perseverance, the secret of all triumphs.

— Victor Hugo

4
Mars, Part II

Contents

4.1 Preface 96
4.2 Summary 97
4.3 Introduction 98
4.3.1 Motivation 98
4.3.2 Terrestrial and lunar context 99
4.3.3 Extension of these methods to Mars 100
4.3.4 The Mars 2020 EDL sequence 101
4.3.5 Aims 103
4.4 Methodology 103
4.4.1 Acoustic signal: source and propagation 105
4.4.2 Coupled seismoacoustic signal 107
4.4.3 Elastodynamic seismic signal 109
4.5 Results 116
4.5.1 Acoustic signal 116
4.5.2 Coupled seismoacoustic signal 118
4.5.3 Elastodynamic seismic signal 118
4.6 Discussion 122
4.6.1 Noise conditions 122
4.6.2 CMBD impact: detection probabilities 125
4.6.3 Extensions to this work 127
4.7 Conclusions 130
4.8 Appendix - seismoacoustic simulations in AxiSEM3D 132
4.8.1 Introduction 132
4.8.2 Choice of boundary condition 133
4.8.3 Theory 134
4.8.4 Results 135
4.8.5 Discussion 135

4.1 Preface

This paper was published as Fernando et al. (2021a):

Fernando, B. et al Wójcicka, N., Froment, M., Maguire, R., Stähler, S. C., Rolland, L., Collins, G. S., Karatekin, O., Larmat, C., Sansom, E. K., Karakostas, F., Leng, K., Nissen-Meyer, T., Kawamura, T., Giardini, D., Lognonné, P., Banerdt, W. B., and Daubar, I. J.

- Listening for the Landing: Detecting Perseverance's arrival at Mars with InSight, *Earth and Space Science*, 2021.

It has the following DOI: *http://doi.org/10.1029/2020EA001585* and further supplementary information may be found in a computer-readable format at *10.5281/zenodo.4291898*.

My role involved developing the plan for the paper and coordinating inputs from multiple authors. I also conducted our initial seismic and coupled seismoacoustic wave propagation simulations, which were not included in the final manuscript but are appended at the end of this chapter. I was also responsible for liaising with the Mars 2020 and InSight operations teams and devising the instrument operation sequence. Wave generation simulations were carried out by N.W., G.S.C., and M.F, whilst R.M., C.L., K.L., and I were responsible for wave propagation modelling. This paper represents the combination of many different modelling and scaling approaches.

The acknowledgements to this paper include Richard Otero, Erisa Stilley, and Ian Clark of the Jet Propulsion Laboratory on behalf of the Mars 2020 team, and Raphael Garcia. Further thanks is offered to Elizabeth Barrett of the InSight operations team.

The full funding sources for all research incorporated here, besides mine, include: the ETH Zürich ETH+ 'Planet Mars' program, the Center for Earth and Space Science of the Los Alamos National Laboratory (supported by the US Department of Energy), the Centre National D'Études Spatiales and the Agence National

de le Recherche (France), the Belgian Science Policy Office through the ESA PRODEX program, and the Australian Research Council. InSight is managed and operated by the Jet Propulsion Laboratory at the California Institute of Technology on behalf of NASA.

A follow-up paper presenting our results is presented in the next chapter, along with full data access details. Since this paper was published, onboard data from the Mars 2020 flight computer has become available which provides marginal refinement to the timings, locations, and distances discussed in this paper. For consistency, this paper has *not* been updated to include these; but they are discussed in the next chapter.

4.2 Summary

- The detection of seismic waves from a source of known magnitude and temporal/spatial origin would be of significant benefit to the InSight mission through acting as a 'calibration' signal for other marsquakes.

- The entry, descent, and landing of NASA's Mars 2020 Perseverance Rover may fulfil this need, and we evaluate whether any signal from Perseverance is likely to be detectable at InSight's position.

- We evaluate the detectability of acoustic, coupled seismoacoustic, and seismic signals using InSight's seismometer, wind, and pressure sensors.

- The elastodynamic signal produced by the impact of the spacecraft's cruise balance mass devices with the surface presents the strongest candidate for detection, with predicted P-wave ground velocities at InSight's position between 2.0×10^{-14} ms^{-1} to 1.3×10^{-10} ms^{-1}.

- Based on analysis of the ambient conditions from one martian year prior to landing, at the upper end of this range we expect the signal to be detectable above the noise floor 40% of the time.

- A positive detection would be of enormous value in constraining the interior seismic properties of Mars, whilst a negative one would enable us to place upper limits on the physics of impact processes.

4.3 Introduction

4.3.1 Motivation

NASA's Interior Exploration using Seismic Investigations, Geodesy and Heat Transport (InSight) mission landed on Mars' Elysium Planitia in November 2018, and since then has detected a number of 'marsquake' events which are thought to be tectonic in origin (Banerdt et al. 2020).

InSight faces a number of peculiar challenges associated with single-station seismology (Panning et al. 2015). Without independent constraints on source properties, robust seismic inversions are more challenging than they would be on Earth. Impact events (where meteoroids hit the planet's surface) offer an opportunity to overcome some of these challenges as they can be photographically constrained in location, size and timing from orbital images. In theory, this should allow a positive impact detection to be used as 'calibration' for other seismic measurements.

However, no impact events have yet been conclusively detected and identified by InSight, despite pre-landing expectations that impacts would make a significant contribution to martian seismicity (Daubar et al. 2018). A meteorite impact which formed a new 1.5 m impact crater only 37 km from InSight in 2019 was not detected (Daubar et al. 2020).

A number of possible reasons may explain the absence of impact detections thus far. These include uncertainties in the impactor flux entering Mars' atmosphere (Daubar et al. 2013) and in the seismic efficiency (the fraction of impactor kinetic energy converted into seismic energy) of ground impacts that form metre-scale craters (Wójcicka et al. 2020), as well as high ambient noise through much of the day, which makes detecting faint signals challenging. Should a seismic signal excited by an impact be detected, distinguishing it from tectonic events remains challenging

due to intense scattering in the shallow crust of Mars (see van Driel et al. (2019) or Daubar et al. (2020) for further discussion).

If a seismic signal recorded by InSight could be identified as impact-generated, conclusive attribution to a particular spatial and temporal location would require identification of a new crater on the surface. Sparse orbital imaging coverage of the martian surface at the required resolution, coupled with large error bounds on event distance and azimuth estimations (e.g. Giardini et al. (2020)), make this extremely challenging. This also offers no information about seismic or infrasonic signals induced by those impactors which either burn up or explode in the atmosphere as airburst events (Stevanović et al. 2017) and do not form new craters.

These challenges may be overcome by using as seismic sources the Entry, Descent, and Landing (EDL) sequences of objects with known entry ephemerides (meaning a priori calculated or independently constrained entry/re-entry timings and locations). The Mars 2020 mission, landing in February 2021, offers an opportunity for this possible measurement.

4.3.2 Terrestrial and lunar context

Spacecraft re-entering the atmosphere are comparatively common on Earth. These have trajectories which are often known prior to their arrival, meaning that seismic observation campaigns can be planned in advance. This has been done for a variety of spacecraft, including the Apollo command capsules (Hilton and Henderson 1974), the Space Shuttle (Qamar 1995; de Groot-Hedlin et al. 2008), Hayabusa-1 (Ishihara et al. 2012), Genesis (ReVelle et al. 2005) and Stardust (ReVelle and Edwards 2007). In these cases, seismic and infrasonic data were used to study the entry dynamics of the spacecraft in question, for example the mechanics of energy dissipation into the atmosphere.

Naturally occurring impact events on Earth may not have trajectories which are known in advance, but their flight paths may be independently reconstructed from photographic evidence (e.g. Devillepoix et al. (2020)) or the recovery of fragments. Examples include the Carancas impact which occured in Peru in 2007

(Le Pichon et al. 2008; Tancredi et al. 2009) and the Chelyabinsk airburst in Russia in 2013 (Borovička et al. 2013). In such studies, seismic and infrasonic measurements (Tauzin et al. 2013; de Groot-Hedlin and Hedlin 2014) are used to study both the entry dynamics and also the properties of the meteoroids themselves, for example radii, masses, and rates of ablation.

On Earth the density of seismic stations and frequency of tectonic events means that impacts are not needed for calibration purposes. However, the Apollo Seismic Experiment did use artificial impacts for calibration on the moon (Nakamura et al. 1982). In this case, the sources were the impacts of the spent upper stages of the Saturn V rockets or derelict Lunar Modules with the lunar surface, which were detected by a network of seismometers deployed by the Apollo astronauts. These events had a known time and location of impact, enabling exact identification of travel times and ray propagation paths for the resulting seismic waves to be made.

4.3.3 Extension of these methods to Mars

On Mars, spacecraft entering the atmosphere are rare. The presence of an atmosphere complicates modelling of impact processes as compared to the lunar case (Nunn et al. 2020), and the entirely different surface and atmospheric compositions mean terrestrial analogues are not directly applicable either (Lognonné et al. 2016). Specifically, the presence of a dry, weakly cohesive surface regolith layer on Mars is expected to reduce the seismic efficiency of impacts as compared to Earth (Wójcicka et al. 2020), whilst the high CO_2 concentration in the atmosphere attenuates high-frequency acoustic signals much more rapidly on Mars than on Earth (Williams 2001; Lognonné et al. 2016; Bass and Chambers 2001).

The landing of NASA's Mars 2020 Rover (Perseverance) on February 18, 2021 is the first time that an EDL event has occurred on Mars during the lifetime of the InSight lander. This paper informs the first ever attempted EDL detection on surface of another planet.

If detected, seismic signals from EDL events would enable us to place substantially better constraints on the seismic efficiency of small impacts on Mars (for

those parts of the EDL apparatus which strike the surface) and the generation of seismic waves by impacts.

An artificial impact also confers the advantage that the impactor mass, velocity, radius, and angle of flight with respect to the ground are all known to within a high degree of precision well in advance, and post-landing return of flight trajectory data and imaging of the resultant craters can provide further constraints (Bierhaus et al. 2013).

A positive detection would also be of substantial benefit to planetary geophysics more generally, enabling us to calibrate the source and structural properties derived from other marsquake events which do not have a priori known source parameters. A negative detection would also be useful, enabling us to place upper bounds on these signals' amplitudes and hence to better constrain the scaling relationships used to predict the amplitudes of seismic waves from impact events.

Finally, we hope that the workflow developed here to evaluate the seismic detectability of EDL signals will be of use in future planetary seismology missions.

4.3.4 The Mars 2020 EDL sequence

Perseverance's landing is targeted for approximately 15:15 Local True Solar Time (LTST) on February 18, 2021. This corresponds to 18:55 LTST at InSight (4.50°N, 135.62°E), or roughly 20:55 UTC on Earth. The centre of the 10 km by 10 km landing ellipse is within Jezero Crater at 18.44°N, 77.50°E (Grant et al. 2018). At atmospheric interface (125 km altitude), the spacecraft's entry mass is 3350 kg and the heat shield is 4.5 m in diameter. At this point the spacecraft's velocity is approximately 19,200 km/h, and it is accelerating.

This is a distance of 3452 km nearly due west from InSight. During descent the spacecraft trajectory is along an entry azimuth trajectory of approximately 100° (Figs. 4.1 and 4.3a), or pointing eastward (azimuth 105°) and directed almost exactly towards InSight.

Two portions of the EDL sequence are likely to produce strong seismic signals. The first is the period during which the spacecraft is generating a substantial Mach

shock as it decelerates in the atmosphere, and the second is the impact of the spacecraft's two Cruise Mass Balance Devices (CMBDs) on the surface.

The spacecraft will generate a sonic boom during descent, from the time at which the atmosphere is dense enough for substantial compression to occur (an altitudearound 100 km), until the spacecraft's speed becomes sub-sonic, just under 3 minutes prior to touchdown. This sonic boom will rapidly decay into a linear acoustic wave, with some of its energy striking the surface and undergoing seismoacoustic conversion into elastodynamic seismic waves, whilst some energy remains in the atmosphere and propagates as infrasonic pressure waves.

The second part of the EDL sequence which will generate a seismic signal is the impacts of the CMBDs on the ground. The CMBDs are dense, 77 kg unguided tungsten blocks which are jettisoned high in the EDL sequence (around 1,450 km altitude). Due to their high ballistic coefficients, they are expected to undergo very limited deceleration before impact. Based on data from the Mars Science Laboratory/Curiosity Rover's EDL in 2012 (Bierhaus et al. 2013) and simulations of the Mars 2020 EDL, the CMBD impacts are expected to occur at about 4000 m/s, less than 100 km from the spacecraft landing site, and at an impact angle of about 10° from horizontal.

In the case of Curiosity, the CMBDs formed several craters between 4 and 5 m in diameter, and the separation between CMBDs or their resulting fragments was no more than 1 km at impact, implying a difference in impact time of less than 1 second between them.

It should be noted that the CMBDs are not the only parts of the EDL hardware which will experience an uncontrolled impact. The heat shield, backshell and descent stage are also expected to reach the surface intact. However, in an optimal landing scenario these are expected to be at sub-sonic speeds (less than 100 ms^{-1} for masses of 440, 600, and 700 kg respectively). Six smaller 25 kg balance masses are also ejected much closer to the surface, and at considerably lower speeds. As such, no other component of the EDL hardware impacting the surface is expected to produce a seismic signal of comparable magnitude to the CMBD impact.

4.3.5 Aims

There is a clear scientific case for 'listening' for Perseverance's landing using InSight's instruments. Doing this requires comprehensive modelling of the propagation of such signals (both in the atmosphere and in the solid ground) from Perseverance's landing site to InSight, in order to estimate signal amplitudes and travel times. This paper presents this modelling work, which is being used to inform the configuration of InSight's instruments in advance of the landing.

4.4 Methodology

To assess their detectability at InSight, we consider three distinct types of signals generated by Perseverance's EDL. Each of these represents wave propagation in a different medium or combination thereof: atmosphere, coupled atmosphere-ground, and ground. Corresponding to the labels in Fig. 4.1, these are:

1. Acoustic signal: A linear, acoustic wave propagating in the atmosphere as an infrasonic (low frequency, <20 Hz) pressure wave, generated by the decay of the sonic boom produced during descent. The modelling methodology for this portion of the signal is presented in Sec. 4.4.1, and the results in Sec. 4.5.1.

2. Coupled seismoacoustic signal: A coupled air-to-ground wave, produced by the sonic boom, or its linear decay product, impinging upon the surface and creating elastodynamic body waves. On Earth, this would usually produce detectable surface waves too - however on Mars these are rapidly scattered away to non-detectable levels and hence are not depicted here. Methodology and results are in Secs. 4.4.2 and 4.5.2 respectively.

3. Elastodynamic signal: A'conventional' seismic wave travelling in the solid part of the planet, excited by the impact of the CMBDs. Methodology and results are in Secs. 4.4.3 and 4.5.3 respectively.

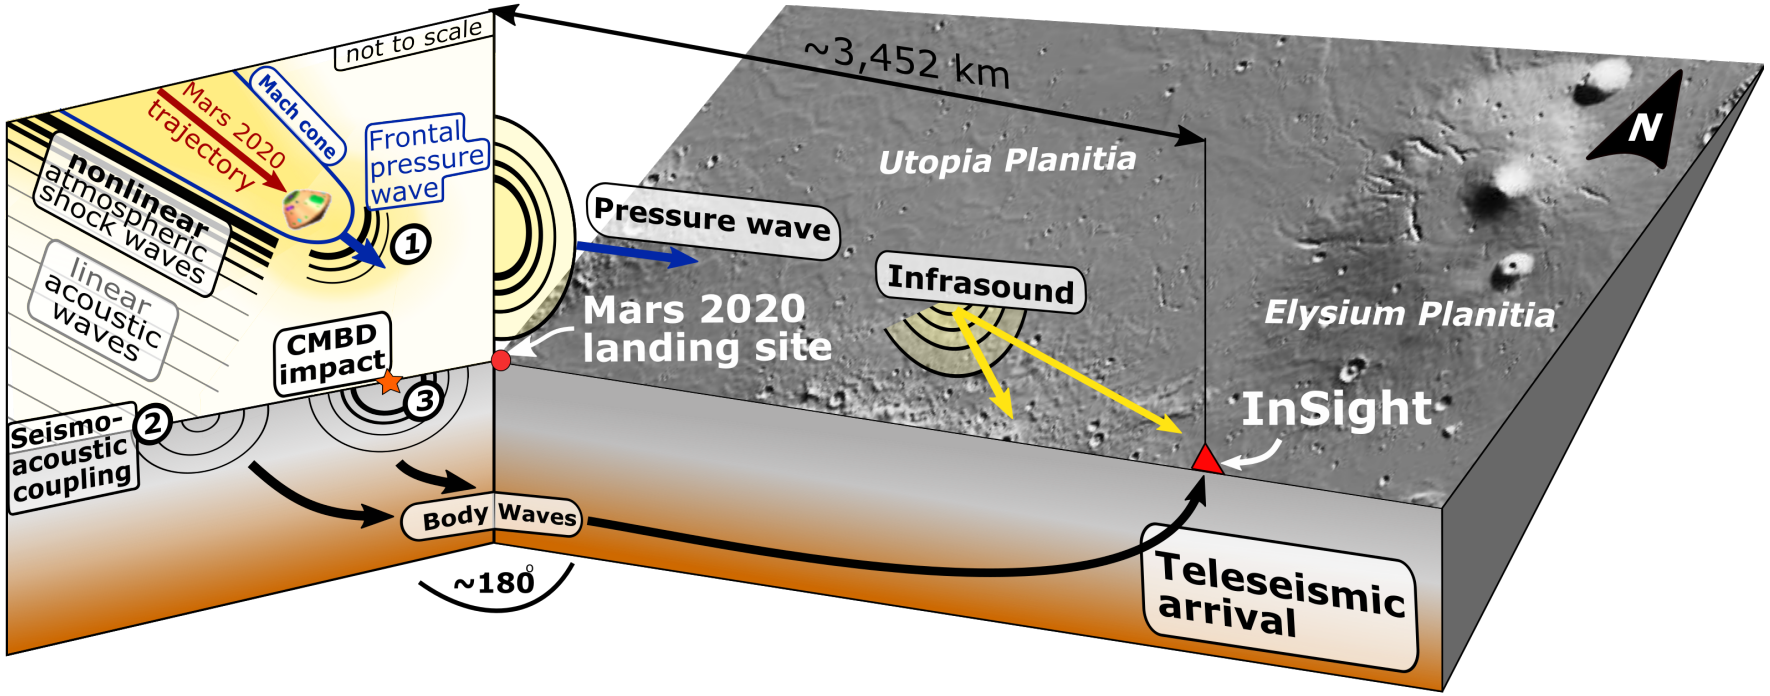

Figure 4.1: Schematic illustration of the seismic signals produced by the Mars 2020 EDL sequence (not to scale). Numbered features are: (1) the atmospheric acoustic signal, (2) the coupled seismoacoustic signal, and (3) the seismic signal propagating in the ground. The thickest airborne black lines represent non-linear shockwaves, decaying to weakly non-linear (thin black lines) and finally linear acoustic waves (thin gray lines). Surface waves, which on Mars do not appear to propagate at teleseismic distances, are not shown here. Black lines with single arrowheads represent body waves. The spacecraft's trajectory at entry is eastward along an azimuth of 100°, almost exactly pointing toward InSight, i.e. the two panels are angled toward each other at nearly 180°, but are shown as they are here to acknowledge remaining uncertainties in the exact entry trajectory which exist at the time of writing. Note that this figure shows all three **potential** sources of seismic signal, and is not intended to suggest that these all reach InSight at detectable amplitudes.

4.4.1 Acoustic signal: source and propagation

The shockwave produced by the hypersonic deceleration of the spacecraft will rapidly decay through viscous frictional processes into a linear acoustic wave. The resultant acoustic (pressure) waves will propagate in the atmosphere following paths determined by the atmospheric structure. The amplitude of any potential signal at the location of InSight is determined by the decay of the signal with increasing distance; due to attenuation in the atmosphere, transmission into the ground, and geometrical spreading.

4.4.1.1 Signal amplitudes

The spacecraft is treated as a cylindrical line source, which is justified on the grounds that the opening angle of the Mach cone is small at hypersonic velocities. Solving the weak shock equations (ReVelle 1976; Edwards 2009; Silber et al. 2015), with additional calculations based upon Varnier et al. (2018), enables us to estimate the energy dissipated into the atmosphere by the spacecraft's entry with increasing distance from its trajectory line.

As per the weak shock theory and sonic boom formulations of the above literature sources, the overpressure decreases with increasing distance from the source as $x^{-3/4}$ and the source wave period increases as $x^{1/4}$; where $x = \frac{r}{R_0}$ is the distance from the line source r, normalised by the blast wave relaxation radius (R_0). R_0 is the distance from the line source at which the overpressure approaches the ambient atmospheric pressure, and for a spherical source is approximately equal to the impactor diameter multiplied by its Mach number.

The calculations account for the gradual nature of the transition between a weakly non-linear and fully linear propagation regime; but do not include attenuation (this is discussed further in Sec 4.6).

As discussed further below, acoustic energy in the atmosphere may be trapped in waveguide layers, which enable low-attenuation long-distance propagation of atmospheric waves. The decay in amplitude with increasing distance from the source r for waves propagating within a waveguide is poorly constrained, with both

terrestrial and martian predictions falling into a range between r^{-1} and $r^{-1.5}$ (Martire et al. 2020; Ens et al. 2012). If acoustic waves are trapped within a waveguide, these scaling laws enable us to predict their amplitude far from the source.

4.4.1.2 Wave trajectories

The acoustic wave trajectories are modelled using the WASP (Windy Atmospheric Sonic Propagation) software (Dessa et al. 2005). The propagation medium is a stratified atmosphere parameterised using a 1D effective sound speed (Garcia et al. 2017). This effective sound speed accounts for the presence of directional waveguides in the martian atmosphere at certain times of day, caused by wind.Wind effects are therefore fully resolved within this model. Such waveguides can potentially enable long-distance propagation of an infrasonic signal in the direction of the wind. However, atmospheric waveguides are comparatively rarer than on Earth and exist only in the presence of winds, unlike on Earth where temperature inversions may create waveguides without wind (Garcia et al. (2017), Martire et al. (2020)).

The adiabatic sound speed and horizontal wind speed along the great circle propagation path from Mars 2020 to InSight are computed from the Mars Climate Database (Millour et al. 2015), at the time and location of Perseverance's landing).Fig. 4.2 shows the variation in effective sound speed toward and away from InSight, highlighting that the effects of the wind are highly directional.

The atmospheric dust content, which significantly influences martian wind and weather patterns through changes in opacity, is chosen as an average for the solar longitude $L_s = 5^\circ$ (northern spring) season, in which dust storms are rare (Montabone et al. 2015).

Weather perturbations may cause second-order changes in the atmospheric conditions (Banfield et al. 2020), but would not change the overall dynamics of acoustic wave propagation considered here. Regardless, the martian atmosphere in the equatorial regions in the northern spring is typically predictable in its meteorology (Spiga et al. 2018).

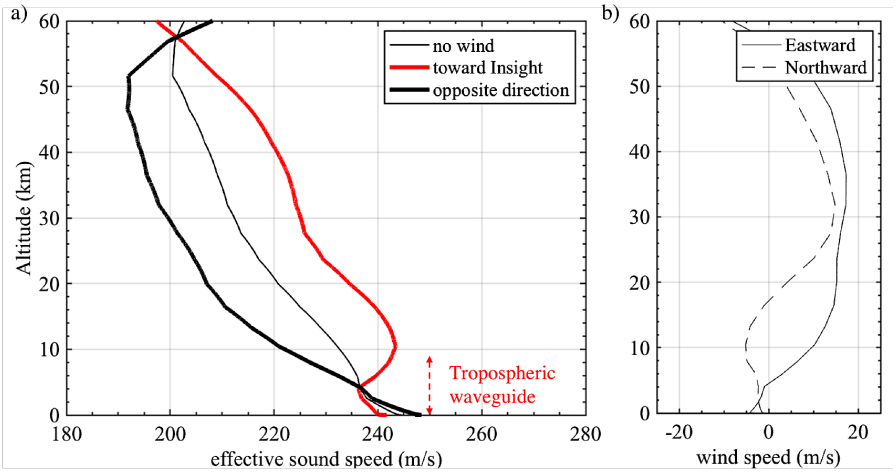

Figure 4.2: The atmospheric model used in simulations of acoustic wave propagation, plotted at the Perseverance landing site and representative of the atmosphere state along the great circle path toward InSight. The left panel shows the effective sound speed as a function of altitude to highlight that these effects are highly directional; whilst the right panel shows the horizontal wind strength. In the absence of wind the sound speed in the atmosphere monotonically decreases between the surface and ~ 50 km. This is not favourable for the long-range propagation of acoustic waves, but the eastward (zonal) wind modifies this to yield an 'effective' sound speed. This creates a tenuous tropospheric waveguide at the bottom of the atmosphere in the direction toward InSight (which is on an azimuth 104° from North from the landing site). The height of this waveguide is marked with a red dashed arrow. All parameters are derived from the Mars Climate Database (Millour et al. 2015).

Infrasonic signals, if at detectable levels, could be recorded directly by InSight's APSS (Auxiliary Payload Sensor Suite) instrument (Banfield et al. 2019); or indirectly by InSight's SEIS (Seismic Experiment for Interior Structure) instrument (Lognonné et al. 2019). The former records the actual atmospheric pressure perturbation, whilst the latter detects the compliance-induced displacement of the ground by atmospheric overpressure or underpressure.

4.4.2 Coupled seismoacoustic signal

The incidence of the linear acoustic waves from the atmosphere (the products of the decaying shockwave) upon the surface will excite elastodynamic (i.e. body and surface) waves in the solid ground. The dominant parameters in determining the amplitude of the elastodynamic waves in the solid ground are the air-to-ground coupling factor (which is a transmission coefficient), and the value of the

overpressure at the surface.

Using the method of Sorrells et al. (1971), we estimate the air-to-ground coupling factor by modelling the intersection of a planar acoustic wave with a regolith-like target material, with a density of 1270 kgm^{-3}, a P-wave velocity of 340 ms^{-1}, and S-wave velocity of 200 ms^{-1}. The effective sound speed is derived from the Mars Climate Database (see Fig. 4.2).

We obtain a value for the air-to-ground coupling factor of 4×10^{-6} $\mathrm{ms}^{-1}\mathrm{Pa}^{-1}$, which is of the same order of magnitude as on Earth. This value is also similar to values obtained by Garcia et al. (2017) and Martire et al. (2020).

The atmospheric overpressure at ground level is modelled as described in Sec. 4.4.1.1, and multiplied by the derived air-to-ground coupling factor to calculate the energy transmitted into the solid ground.

After the wave has coupled into the ground, its amplitude decay upon propagating through a 1D seismic Mars model is calculated using Instaseis (van Driel et al. 2015).

4.4.2.1 Body Waves

We focus on the prediction of P-wave amplitudes from seismoacoustic coupling, as these are expected to be the strongest of the direct-arrival body waves generated by atmospheric overpressure at the surface

Multiplying the value obtained for the overpressure at ground level by the air-to-ground coupling factor (Sec. 4.4.2) gives an upper bound for the velocity amplitude of the P-wave at the landing site.

The decay of this amplitude with distance to InSight's position can then be calculated using either waveform modelling or scaling laws (these are discussed below). The S-wave amplitude from the coupled seismoacoustic signal is expected to be much smaller, as the vertical incidence of the atmospheric acoustic wave (caused by the steeply descending trajectory of the spacecraft) produces much stronger pressure perturbations than shear perturbations in the solid ground.

The resulting body waves propagating in the solid ground could, if large enough in amplitude, be detected by SEIS.

4.4.2.2 Surface waves

Modelling of the excitation of atmospherically induced surface waves is discussed in detail by Lognonné et al. (2016) and Karakostas et al. (2018). However, the combination of a small transmission coefficient and strong seismic scattering in the portions of the crust where the surface waves propagate (Banerdt et al. 2020) means that the surface wave signal is extremely unlikely to be detected at InSight and we do not consider them further in this paper. If this procedure is applied to other planetary seismology settings where surface waves are expected, they should be considered as well as they may be greater in amplitude than the P-wave. Extending the use of Instaseis to achieve this is simple.

4.4.3 Elastodynamic seismic signal

The impact of the CMBDs at the Perseverance landing site will excite both surface and body waves. As was the case for the coupled seismocacoustic signal discussed in Sec 4.4.2, the surface wave phases are expected to be scattered away before they reach InSight. We therefore focus on the signals which we expect to have the largest amplitude, that is, the direct-arrival P-wave.

The entry trajectories of Perseverance and its CMBDs are obtained through aerodynamic simulations whose inputs (i.e. the initial trajectory state, vehicle mass, aerodynamic coefficients) are considered identical to those used by the Mars Science Laboratory Curiosity (Karlgaard et al. 2014), on account of the nearly identical EDL apparatus. The aerodynamic coefficients and vehicle mass are assumed to be constant along the trajectory. An ellipsoid gravity model is used, and atmospheric conditions are extracted from the Mars Climate Database (MCD) climatology scenario at the predicted landing location and time.

Two approaches are taken to estimate the peak P-wave amplitudes at InSight produced by the CMBD impacts, and hence to evaluate their detectability by SEIS. The first (Sec. 4.4.3) makes use of empirical amplitude scaling relationships to directly estimate the P-wave amplitude at InSight's position. The second makes use of wave propagation modelling, with the choice of source magnitude informed either

by scaling-based moment estimates (Sec. 4.4.3.1) or by shock physics simulations of the CMBD impact (Sec. 4.4.3.2). These approaches are complementary.

4.4.3.1 Method 1: Empirical amplitude scaling relationships

The first approach uses the empirical scaling relations of Teanby (2015) and Wójcicka et al. (2020) to estimate the peak P-wave amplitudes at InSight's location. The amplitudes of the S-wave are significantly harder to estimate (and are not predictable from the published scaling relationships discussed below), but are likely to be of the same order of magnitude as the P-waves.

The empirical scaling relationships are based on the measured P-wave amplitudes as a function of distance from artificial lunar (Latham et al. 1970b) and terrestrial missile impact experiments (Latham et al. 1970a), which follow a $r^{-1.6}$ relationship. The exponent of 1.6 is empirically derived using data from terrestrial explosions and differs from the exponent of 2 which would be expected from wavefronts spreading into an infinite space on account of the sound speed gradient and the presence of the Moho.

The approaches differ in how these relationships are rescaled to the CMBD impacts on Mars based on impactor properties.

The Teanby (2015) approach scales the empirically derived P-wave amplitude with the square root of the impactor's kinetic energy; whilst the Wójcicka et al. (2020) approaches use a scaling based on impactor momentum, either total or vertical. The estimated impact energy, total momentum and vertical-component momentum of the CMBD impact are 6×10^8 J, 3×10^5 Ns, and 5.2×10^4 Ns, respectively.

From these impact parameters the scaling approaches directly yield a predicted P-wave amplitude at InSight's position.

The application of lunar and terrestrial-derived scaling relationships to Mars is well-established (e.g. Daubar et al. (2020)). However, it should be noted that both these approaches involve considerable extrapolation in distance to reach the 3452 km separation from InSight. Extrapolation is required because comparable (i.e., controlled-source, and with the same momentum and energy) impact events

have not previously been seismically recorded on the Moon or Earth (or indeed Mars) at distances greater than 1200 km. Nonetheless, it should be noted that the lunar events are at comparable epicentral distances to this event, on account of the Moon's smaller radius.

4.4.3.2 Method 2: Wave propagation modelling using estimated moments

The second approach predicts the amplitudes of the elastodynamic waves recorded at InSight using wave propagation modelling. Because elastodynamic wave propagation is linear, the amplitude at InSight is directly proportional to the magnitude of the source, and calculations can be easily re-scaled for different estimates of source magnitude (which in these cases is a seismic moment) to yield a range of predicted amplitudes.

The seismic moment is thus the primary determinant of peak P-wave amplitude. Several approaches have been proposed to estimate the seismic moment of an impact, with an uncertainty that spans two orders of magnitude (Daubar et al. 2018). Here we derive two independent estimates of the seismic moment:

A) Scaling-based moment estimates Rearranging equations (5) and (6) of Teanby and Wookey (2011) provides an empirically-derived relationship between seismic moment (M) and impact kinetic energy (E), via $M = (k_s E/4.8 \times 10^{-9})^{0.81}$, where k_s is the seismic efficiency of the impact.

B) Impact physics hydrocode simulations To estimate the seismic moment of the CMBD impact in an independent way we use the iSALE2D (Amsden et al. 1980; Collins et al. 2004; Wünnemann et al. 2006) and HOSS (Hybrid Optimization Software Suite; Munjiza (2004), Lei et al. (2014), and Knight et al. (2020)) impact physics codes to simulate the impact and wave generation process on millisecond timescales.

Realistic simulations of highly oblique impacts such as the 10° from horizontal impact of the CMBDs are extremely challenging. Whilst HOSS is capable of such

simulations (iSALE2D is not), these are executable only with lower spatial resolution and over a shorter duration than simulations with vertical impactors.

Therefore, to provide the most robust prediction possible, we simulated the CMBD collision with the surface as a vertical impact using both iSALE2D and HOSS (at high resolutions and longer timescales), and as an oblique impact in HOSS (at a lower resolution and shorter timescales). These are labelled as scenario **(a)** and **(b)** below.

The trade-off between resolution and duration versus realism means that we cannot claim that one of these cases is 'better' or 'more accurate' than the other. Of the two, the vertical impact simulation is expected to provide an upper bound on the seismic moment as it maximises the coupling of the impactor's energy with the ground.

In oblique impacts such as this, the horizontal momentum contributes to the crater formation processes and the vertical component alone significantly underestimates the scalar seismic moment. Because of this, our 2D hydrocode simulations (i.e. those which use a vertical impactor) use the total momentum (3×10^5 Ns) as initial impact momentum.

The scalar seismic moment of the impact was calculated differently for the different simulation approaches. The scalar seismic moment calculated from the iSALE2D simulation results uses a combination of three methods that each provide a measure of either the scalar seismic moment or the diagonal components of the full seismic moment tensor (Wójcicka et al. 2020). The method used to determine the seismic moment from the HOSS simulation provides information about the full seismic moment tensor, including off-diagonal terms. Further details are provided in the Supplement.

Material models In the iSALE2D simulation, the balance mass was modelled using the Tillotson equation of state (Tillotson 1962) and the Johnson-Cook strength model (Johnson and Cook 1983) with parameters appropriate for tungsten.

To approximate the local geological conditions at Jezero Crater, the target was modelled as a porous basaltic regolith of bulk density $\rho = 1589$ kg/m^{-3} and sound speed $c_B = 857$ m/s, using the Tillotson equation of state combined with the ε-α porosity model (Wünnemann et al. 2006; Collins et al. 2011) and the Lundborg strength model (Lundborg 1968). Full material model parameters are available in supplementary material Table S1 or in user-ready format from Wójcicka and Froment (2020).

The HOSS model was configured to be as close to the iSALE2D initial conditions and material models as possible, enabling an accurate but independent method of verifying the derived seismic moment.

In HOSS, the CMBD was modelled using the same equation of state as detailed above for iSALE2D. The HOSS equation of state for the porous target material takes the form of a user-defined curve relating pressure and volumetric strain in a regime of elastic deformation at low stresses, followed by a regime of plasticity and pore-crushing at higher stresses. This model of martian regolith was recently validated based on laboratory hypervelocity impact experiments, conducted in a martian regolith proxy made of loose pumice sand (Richardson and Kedar 2013; Froment et al. 2020). In this work, parameters for porosity and sound speed were modified so that the material behaviour replicated, as far as possible, that used in the iSALE2D ε-α model.

iSALE2D modelling The shape of the CMBD in iSALE2D is approximated as a tungsten sphere of radius 9.6 cm and mass 75 kg. The mesh used in the simulations is cylindrically symmetric, approximately 30 m in radius. The impact-generated shockwave is tracked at high resolution until it decays to a purely linear elastodynamic wave.

To estimate the seismic moment in the vertical impact case with iSALE2D we follow the three approaches described by Wójcicka et al. (2020). The first approach is based on Müller (1973), which expresses seismic scalar moment, M_1, in terms of a hemispherical surface surrounding the impact that is moved by an

average residual displacement. The second approach is based on Walker (2003) and provides an estimate of the radial component of seismic moment, M_{rr}. The final approach was adapted from the Gudkova-Lognonné model (Lognonné et al. 2009) and returns the vertical seismic moment, M_{zz}, calculated from total momentum transferred to the target during impact. The arithmetic mean of the three seismic moment values was taken to produce a single representative value of the scalar moment, M_0, to be used in later calculations.

HOSS modelling HOSS uses the Lagrangian description, and is based on the Finite Discrete Element Method (Munjiza 2004; Lei et al. 2014)). This hybrid representation merges continuum solutions for the calculation of stresses as a function of deformation with the Discrete Element Method for the resolution of fracture, fragmentation, and contact interaction. Impact simulations are conducted in 3D, and unlike iSALE2D need not be cylindrically symmetric.

Two impact geometries are used to simulate a CMBD impact in HOSS:

(a) The first scenario assumes a vertical incidence and a 4000 m/s impact velocity. The target geometry is a 30° cylindrical sector with a height of 27 m. The minimum size of 3D elements is 1.2 cm.

(b) The second scenario (the 'oblique' case) accounts for the 10° from horizontal impact angle by modelling the target as a quarter of a sphere cut along the x-z and x-y planes, with radius 12 m. This scenario is conducted at lower spatial and temporal resolutions than scenario (a), with minimum element size of 1.8 cm – note that this is 50% larger than in scenario (b).

The approach used here to compute the seismic moment is different from that of iSALE2D and relies on the notion of *Stress Glut* developed by Backus and Mulcahy (Backus and Mulcahy 1976a; Backus and Mulcahy 1976b). This method was applied to planetary impacts in the work of Lognonné et al. (1994) and Gudkova et al. (2015). We derive a second rank seismic moment tensor with six independent components. Here, the effect of material shear strength is accounted for and contributes additional

diagonal and non-diagonal terms to the stress glut tensor. The expression of the stress glut, with opposite sign conventions to Lognonné et al. (1994), is the following:

$$\Pi_{ij}(t) = \Psi_{ij}(t) - S_{ij}(t) + (\rho v_i v_j)(t), \tag{4.1}$$

where Ψ_{ij} is the the modelled elastic Hooke stress deriving from impact-generated deformation, S_{ij} is the true stress in the material, and $\rho v_i v_j$ is the Reynolds momentum transport due to crater and ejecta formation. The expression of the time-varying moment tensor in the volume V of the impacted target is then (after Lognonné et al. (1994), eq. 16):

$$M_{ij}(t) = \int_V \Pi_{ij}(t)\ \mathrm{dV}. \tag{4.2}$$

From this tensor, a scalar seismic moment $M_0 = \frac{1}{\sqrt{2}}\sqrt{\sum_{ij} M_{ij}^2}$ is derived.

Wave propagation modelling Synthetic waveforms with an isotropic source are generated using Instaseis (van Driel et al. 2015) to retrieve pre-computed Green's function databases prepared for the InSight mission (Ceylan et al. 2017). These are accurate up to a frequency of 1 Hz, and are then rescaled using the scalar seismic moments derived for the CMBD impacts as detailed above.

In this paper, we consider the structural model EH45TcoldCrust1 with attenuation (Rivoldini et al. 2011), which has been used in previous benchmark modelling of impact signals on Mars (Daubar et al. 2018). While modelled waveform amplitudes vary slightly (less than 10%) between different structural models, the variations associated with different models are far lower than the uncertainty of the estimated seismic moment of the impact. Given the uncertainties in modelling the focal mechanism for a hypersonic impact (see Daubar et al. (2018) for more details), the use of an isotropic (explosive) source is a standard and justifiable assumption. If this methodology is applied to other contexts where a different source radiation pattern is desired, the extension to using a full second-rank moment tensor in Instaseis is simple.

4.5 Results

4.5.1 Acoustic signal

Fig. 4.3 presents the trajectories of the spacecraft and CMBD and acoustic ray-tracing simulations. The acoustic energy release at any point in time is dependent on both the velocity of the entry vehicle and the atmospheric density (and hence, the spacecraft altitude). The spacecraft dissipates the most energy into the atmosphere at the point of maximum aerodynamic deceleration, or approximately 30 km above the surface and 90 seconds after atmospheric entry interface.

Acoustic energy emitted at altitudes above 10 km reflects off the surface back into the atmosphere at too steep an angle to be refracted into the waveguide and propagate toward the lander. Therefore, the acoustic signal produced around the time at which Mars 2020 is undergoing maximum deceleration will not likely be detectable by InSight due to the geometry of the waveguide layer.

Below 10 km, acoustic energy from the decaying shock front may become trapped between the wind layers in the atmosphere and the surface, and hence propagate for long distances. However, the amount of acoustic energy emitted will decrease substantially as the entry vehicle's parachute deploys and it passes into the subsonic regime, around 140 s prior to landing and approximately 11 km above the surface. This signal will therefore, with high confidence, not be detectable. A more speculative discussion on this topic which includes a comparison to the APSS noise floor is included such that this methodology can be easily extended to other contexts in Sec. 4.6.3.3.

The impact of the CMBDs with the ground will generate an acoustic signal which will propagate up into the atmosphere. Due to the complexities of this signal's generation and propagation, it is not currently possible to meaningfully estimate its amplitude at InSight's position. Again, for a more speculative discussion, see Sec 4.6.3.3.

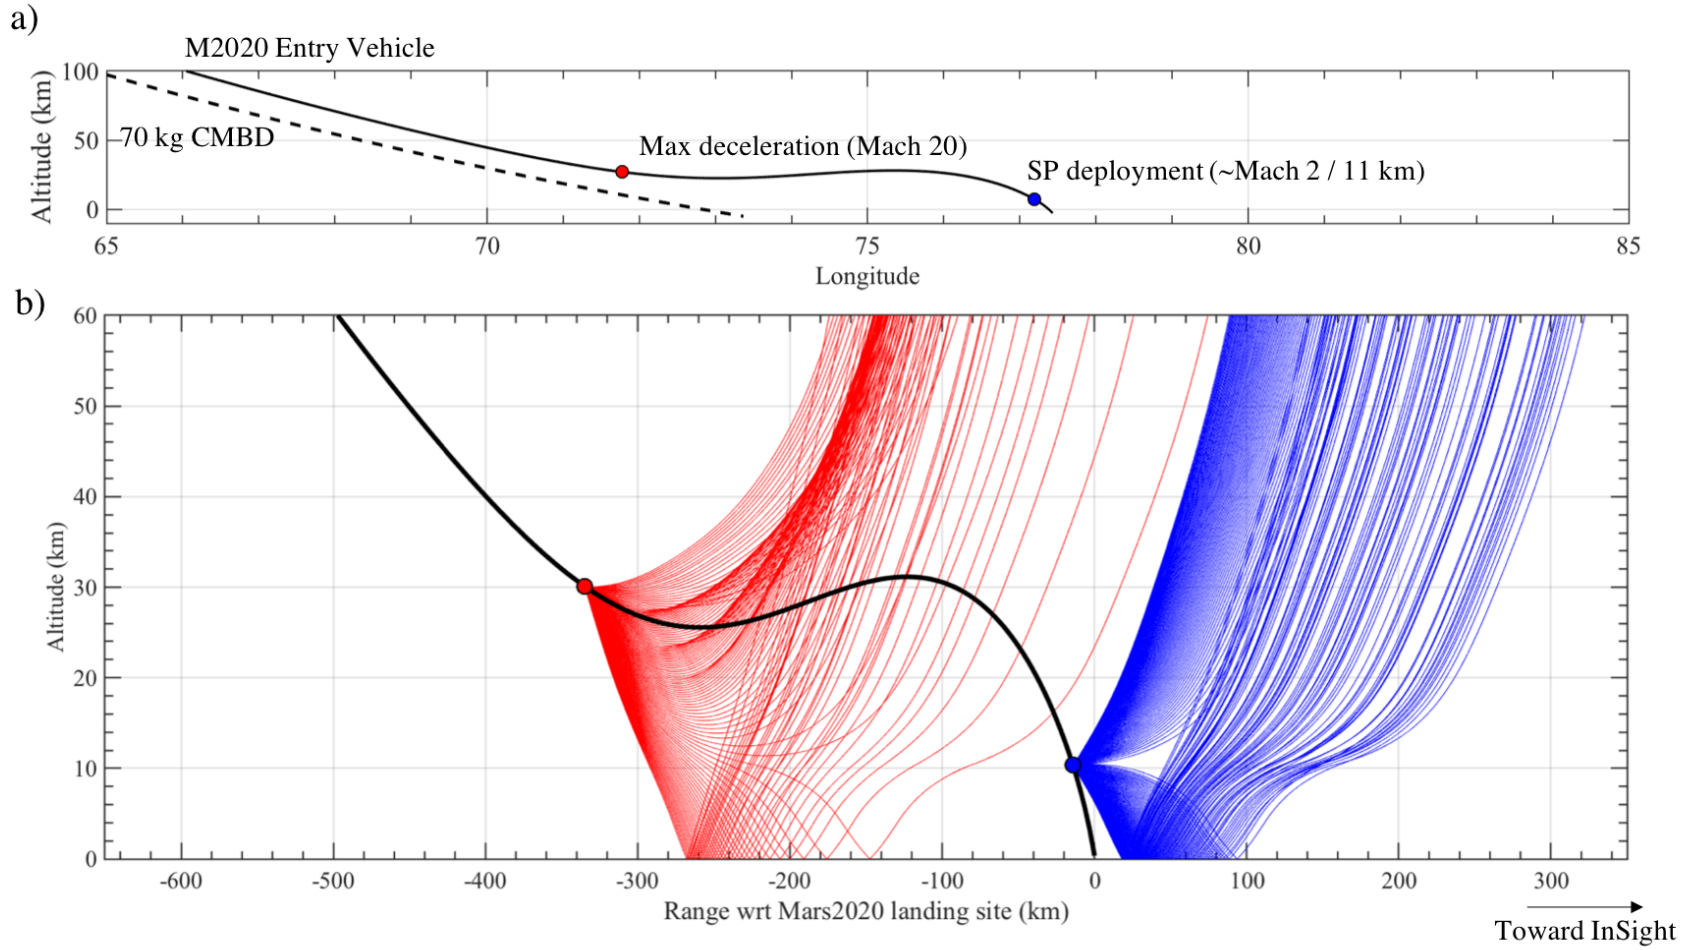

Figure 4.3: Panel a) shows the entry trajectories of the CMBDs and Mars 2020 entry vehicle (dashed and solid curves, respectively) CMBD separation occurs far off to the top left of the graphic (~1450 km altitude and ~3330 km downrange). The red dot marks the calculated point of maximum of deceleration (where the emission of acoustic energy into the atmosphere is highest),and the blue dot marks the estimated location of the Supersonic Parachute (SP) opening, after which the spacecraft rapidly becomes subsonic. Panel b) illustrates the infrasound propagation paths on Mars at the time of landing, in red for a source at 30 km height at the point of maximum deceleration, and in blue for an acoustic source at 11 km where the SP deployment occurs.

4.5.2 Coupled seismoacoustic signal

Acoustic ray tracing predicts a sonic boom swath (a 'carpet' in which waves reach the surface directly, i.e. without bouncing off of it first) of width no more than 100 km. We estimate a maximum surface overpressure in this region of 0.9 Pa with a fundamental frequency of 0.5 Hz, which is attributable to the portion of the sonic boom generated at 25 km height. At this position, the spacecraft is travelling fast enough to still generate a substantial shockwave (Mach 15).

Using our calculated air-to-ground coupling factor of 4×10^{-6} $\mathrm{ms^{-1}Pa^{-1}}$, the 0.9 Pa overpressure translates into a ground deformation velocity of 3.6×10^{-6} $\mathrm{ms^{-1}}$ at the landing site.

Modelling a seismic source of this magnitude using Instaseis suggests a maximum P-wave amplitude no larger than 2×10^{-11}m/s at InSight's location. The average noise spectrum is discussed below in Sec. 4.6.2, but this is substantially below the noise floor and hence will not be detectable.

4.5.3 Elastodynamic seismic signal

As per Sec. 4.4.3.2, the use of two independent methods (scaling laws and shock physics simulations) to estimate the amplitude of the elastodynamic seismic signal at InSight's position yields a spread of values for the seismic deformation velocity below the lander. From a detectability perspective, the highest of these values corresponds to the 'reasonable best case' scenario, as discussed in Sec. 4.6).

4.5.3.1 Method 1: Empirical scaling relationships

Fig. 4.4 presents estimates of the peak P-wave amplitude as a function of downrange distance for the CMBD impact, as compared to data from artificial lunar (Latham et al. 1970b) and terrestrial missile impact experiments (Latham et al. 1970a). These form the basis of our empirical scaling estimates.

The vertical offset between each scaling line occurs because of the different approaches used to scale the results of the experimental data to the Mars 2020 CMBD impact scenario (see Sec. 4.4.3.1 or Wójcicka et al. (2020) for more details):

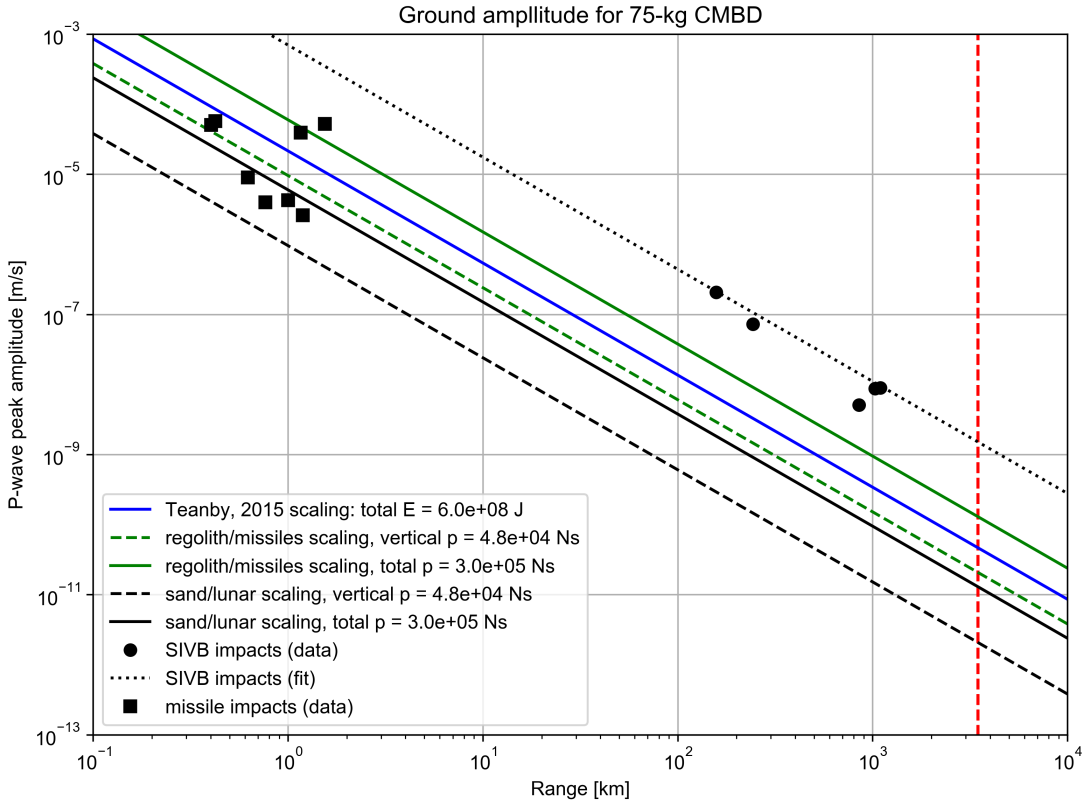

Figure 4.4: P-wave peak amplitude versus range estimates for the CMBD impact based on different scaling approaches discussed in Sec. 4.5.3.1. The solid blue line shows the estimate based on scaling P-wave peak amplitude by the square-root of the impact energy as described by Teanby (2015). The other lines show estimates based on scaling P-wave peak amplitude for the missile data (green lines) or lunar impact data (black lines), in each case with two lines corresponding to scaling by the total impactor momentum (solid lines) or vertical impactor momentum (dashed lines). The dotted black line is a fit to the lunar data, which as discussed below is distant in parameter space from the CMBD impacts and hence is significantly separated from the other scaling lines. The red vertical line marks the distance of InSight from the estimated CMBD impact point.

- The solid blue line shows scaling by the square root of impact energy.
- The green lines are based on terrestrial missile data, scaled by total (solid) and vertical momentum (dashed).
- The black lines are based on the lunar impact data scaled by total (solid) and vertical (dashed) momentum. The dotted black line is a fit to the data.

The Mars 2020 CMBD impactor momentum and kinetic energy of 3×10^5 Ns and 6×10^8 J respectively are of similar magnitude to the terrestrial missile

impacts (Latham et al. 1970a). Hence, the P-wave amplitude estimates based on extrapolation of these data (green and blue lines in Fig. 4.4) are comparable to the missile data (i.e. the scaling lines pass through the region of the datapoints).

The impact momentum, vertical impact momentum and kinetic energy of the lunar impacts (black lines), on the other hand, are approximately 120, 640 and 75 times larger than their corresponding values for the CMBD impact, respectively. As such, a sizable extrapolation in energy or momentum must be performed, in order to use the lunar data to make predictions of the peak P-wave amplitudes for the CMBD impact being considered here.

These differing approaches result in a large range in estimated P-wave peak amplitude at InSight when the trend line based on the lunar experimental data (dotted line) is re-scaled to the CMBD impact by momentum, vertical momentum or the square root of the kinetic energy (blue and black lines).

The lower and upper bounds on the P-wave amplitudes at InSight's position from the different methods are:

- 2.1×10^{-12} and 1.3×10^{-11} m/s from lunar-based impact momentum scaling
- 2.1×10^{-11} and 1.3×10^{-10} m/s from terrestrial-based missile scaling
- $5^{+10}_{-3.5} \times 10^{-11}$ ms^{-1} from the Teanby (2015) scaling

Note that in the first two cases the lower and upper bounds come from using the vertical and total momentum, respectively; whilst in the latter case the uncertainty is experimentally derived, non-symmetric, and hence differently presented.

The resulting overall range of peak P-wave velocities at the distance of InSight using these three methods is is between 2.1×10^{-12} and 1.3×10^{-10} ms^{-1} These results are plotted and compared to other derived values for the purposes of estimating detectability in Fig. 4.6.

4.5.3.2 Method 2: Wave propagation modelling with an estimated seismic moment

A) Scaling based moment estimate While there remains considerable uncertainty in the most appropriate value for the seismic efficiency of small impacts on Mars (Teanby and Wookey 2011; Daubar et al. 2018; Wójcicka et al. 2020), to derive a plausible upper bound on the seismic moment of the CMBD impact we adopt a value of $k_s = 5 \times 10^{-4}$ (Teanby 2015; Daubar et al. 2018), which yields a seismic moment $M = 1.3 \times 10^{11}$ Nm. This estimate has at least an order of magnitude uncertainty.

B) Impact physics hydrocode simulations In the case where the CMBD impact is approximated as a vertical (scenario **(a)** from Sec. 4.4.3.2), iSALE2D predicts a scalar seismic moment of $5.85 \pm 1.5 \times 10^8$ Nm whilst HOSS predicts a moment of 2.97×10^9 Nm. The factor-of-five discrepancy between these two values is likely due to differences in the way that the ejecta from the CMBD crater is modelled and in how the surface material is parameterised. As described in Sec. 4.4.3.2, each moment estimate must be computed using a different mathematical approach due to the simulation methods used, which will also introduce discrepancies

In the case of a highly oblique CMBD impact (scenario **(b)** from Sec. 4.4.3.2), the HOSS simulation results yield a scalar seismic moment of 0.92×10^9 Nm, which is within the range of estimates of the scalar moment of the vertical impact approximation. In this case, the scalar seismic moment presents a significant off-diagonal component of the moment tensor (shear in the vertical and along-trajectory directions), whereas the diagonal terms of the moment tensor dominate in the vertical impact case (Table 4.1). This suggests that the use of an isotropic moment tensor source approximation in our wave propagation modelling to represent a highly oblique impact source may introduce an additional uncertainty in P-wave amplitude that should be explored in further work but is beyond the scope of this paper.

The arithmetic mean of the estimates of scalar seismic moment suggests a moment of $\sim 1.5 \times 10^9$ Nm. While this estimate is more than two orders of magnitude less than the estimate of 1.3×10^{11} Nm based on the impact energy-moment scaling relationship of Teanby and Wookey (2011), as described in Sec.

4.4.3.2 A), it is consistent with other estimates of seismic moment (in both value and difference from other estimates) for impacts of similar momentum n terrestrial, lunar, and martian contexts (Gudkova et al. 2015; Daubar et al. 2018; Wójcicka et al. 2020). Possible reasons for this disparity are discussed in Sec. 4.6.2.

We therefore consider a predicted range for the seismic moment of 1.5×10^9-1.3×10^{11} Nm, which we are confident bounds the 'true' seismic moment. This can then be used to scale the results of our wave propagation modelling.

Using these limits on the source moment to linearly re-scale seismogram velocity amplitudes, as discussed in Sec 4.4.3.2, yields amplitudes in the range 2.0×10^{-14} ms^{-1} (corresponding to the lower bound of 1.5×10^9 Nm) and 2.0×10^{-12} ms^{-1} (corresponding to the upper bound predicted moment of 1.3×10^{11} Nm). Note that for this scenario, wave propagation modelling predicts that the peak S-wave amplitude will be lower than that of the P-wave.

These upper and lower values (v_u and v_l respectively) bound a predicted range of ground deformation velocities; note that these estimates are entirely independent of the scaling estimates presented in 4.5.3.1. Seismograms, showing these amplitudes as well as approximate arrival times, are shown in Fig. 4.5.

Possible reasons for the differences between the estimates produced the different methods are discussed below.

4.6 Discussion

4.6.1 Noise conditions

The upper range of the amplitude predictions of the elastodynamic seismic wave generated by the CMBD impact with the ground exceeds the noise floor for InSight's SEIS instruments at certain times of day. We now consider how likely this signal is to exceed a signal-to-noise ratio of 1.5 (a reasonable threshold for detection, based on InSight detections of tectonic events) at the predicted time of Perseverance's landing.

Given the highly repeatable meteorological patterns on Mars in the absence of a global dust storm, we estimate the likely noise levels at the time of Perseverance's landing (the local evening of February 18, 2021) using data averaged across twenty

Moment Component (HOSS results)	**Case (a), vertical**	**Case (b), oblique**
M_{xx} [Nm]	$(2.96 \pm 0.60) \times 10^9$	$(3.22 \pm 1.84) \times 10^8$
M_{yy} [Nm]	$(2.96 \pm 0.60) \times 10^9$	$(6.54 \pm 0.12) \times 10^8$
M_{zz} [Nm]	$(0.27 \pm 1.20) \times 10^9$	$(6.30 \pm 1.91) \times 10^8$
M_{xy} [Nm]	0	0
M_{yz} [Nm]	0	0
M_{xz} [Nm]	0	$(-6.21 \pm 0.1) \times 10^8$
Scalar Moment $\mathbf{M_0}$ [Nm]	$\mathbf{(2.97 \pm 0.30) \times 10^9}$	$\mathbf{(9.22 \pm 0.30) \times 10^8}$
Moment Component (iSALE2D results)	**Value**	
Radial seismic moment, M_{rr} [Nm]	4.2×10^8	
Vertical seismic moment, M_{zz} [Nm]	3.9×10^8	
Buried explosion moment, M_1 [Nm]	9.6×10^8	
Scalar Moment $\mathbf{M_0}$ [Nm]	$\mathbf{(5.85 \pm 1.5) \times 10^8}$	
Teanby and Wookey (2011) $\mathbf{M_0}$ [Nm]	$\mathbf{1.3 \times 10^{11}}$	

Table 4.1: Seismic moment of the CMBD impact obtained from the different hydrocode simulations (top two sections) and the (Teanby and Wookey 2011) method (bottom section). The top part of the table shows each moment tensor's components and the scalar seismic moment M_0 associated with impact scenarios **(a)** and **(b)** simulated with HOSS. Results from iSALE2D simulations of the CMBD, calculated using three methods described in Sec. 4.4.3.2 B, are shown in the middle part of the table.

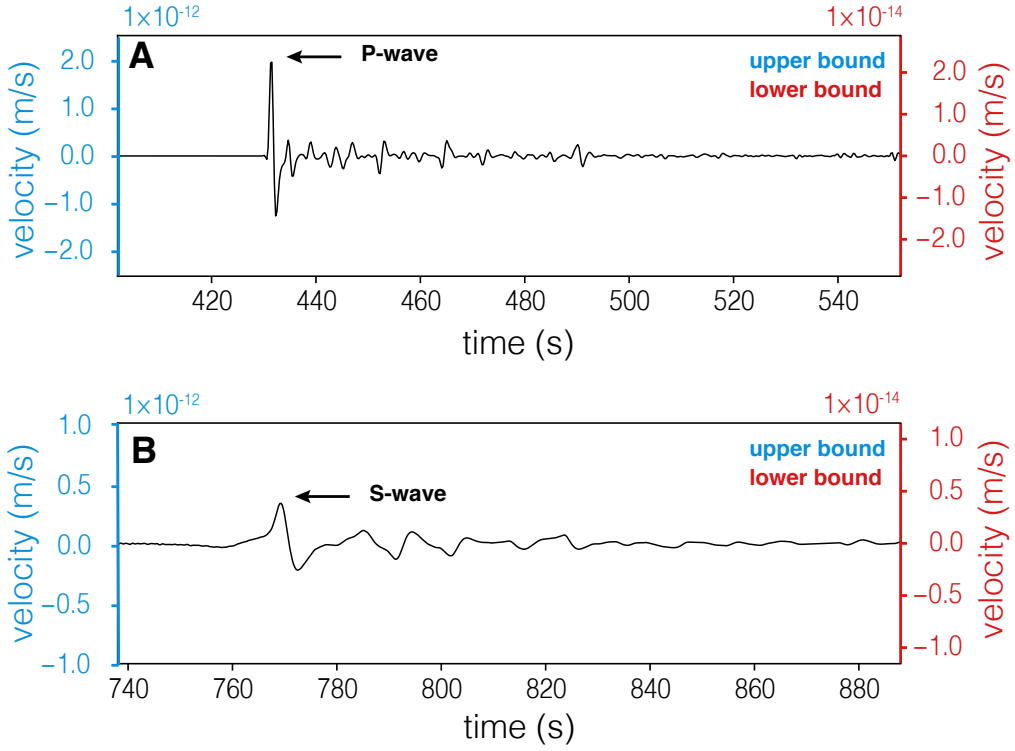

Figure 4.5: Vertical component Instaseis synthetics calculated for an isotropic moment tensor representation of the CMBD impact. Panels (A) and (B) show close ups of the P- and S-waves, respectively. The vertical scale shown in blue corresponds to velocities calculated assuming a scalar moment $M_0 = 1.3 \times 10^{11}$ Nm (the upper bound of moment estimates), and the vertical scale shown in red corresponds to velocities calculated assuming a scalar moment $M_0 = 1.5 \times 10^{9}$ Nm (the lower bound of moment estimates).

evenings from the same period the previous martian year (687±10 Earth days previously, UTC Earth dates 2019/04/01 to 2019/04/20).

In 2019, these spring evenings (18:30-20:00 LMST at InSight) on Mars were characterised by very low noise levels in the early evening post-sunset within the main seismic band used by the lander (0.2–0.9 Hz). To account for the temporal variability in the noise levels within this time, we consider the 'probability' of detection as being the fraction of time within the expected arrival window during which a signal of a given amplitude would be at least 1.5 times greater than the noise floor. For reference, we also plot the noise levels for the whole martian day (Sol) in Fig. 4.6; demonstrating that the noise is on average significantly lower during the evening.

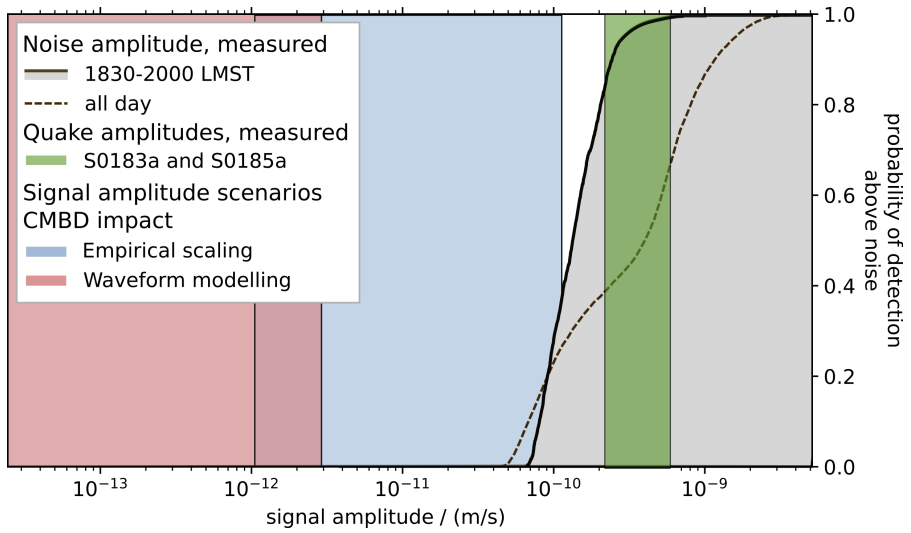

Figure 4.6: Detection probabilities for seismic signals of certain velocity amplitudes between 0.2 and 0.9 Hz. The solid black curve indicates the noise distribution considering the average signal amplitudes in only the early evening over 20 sols during the same martian season in 2019, whilst the dashed black curve is for the whole period of 20 sols. The shaded gray area indicates the regions in which signals are detectable. The blue and red bars mark the P-wave amplitude estimates of the 75 kg CMBD impact, using the empirical scaling and wave propagation modelling estimates, respectively, described earlier in this paper. Vertical lines bounding the different sectors correspond to the upper and lower bounds derived from these methods, for the blue and red sectors respectively (as an example, v_u and v_l are the vertical edges of the red sector). For comparison, the amplitudes of two previously detected tectonic marsquakes, S0183a and S0185a, located at comparable distances, are plotted in green.

4.6.2 CMBD impact: detection probabilities

The upper end of the peak amplitude estimates, derived from empirical impact scaling laws (Fig. 4.6), predicts an amplitude which exceeds the average early evening noise levels by a factor of 1.5 approximately 40% of the time. This implies that the elastodynamic signal propagating in the ground and induced by the CMBD impact may be detectable at InSight. However, the range of predicted peak ground velocities is substantial. This is not dissimilar to other amplitude predictions for martian impacts (Daubar et al. 2020). This wide range of predicted values is directly attributable to:

- Significant uncertainty in the efficiency of seismic wave generation of oblique impacts, especially in the relationship between impactor momentum and released seismic moment or between impact energy and seismic energy. This

is partially a consequence of no impacts having been seismically detected on Mars to date.

- A lack of prior examples of hypersonic impacts detected at distances greater than 1200 km on any body, making calibrating scaling relationships challenging. Different approaches to extrapolating these, coupled with differences in material properties between terrestrial soils, lunar regolith and the martian surface, yield estimates that differ by two orders of magnitude depending on the choices made.

- The frequency bands used in estimating scaling relationships are not identical to those used in waveform modelling and predicted noise levels. This is an unavoidable consequence of the frequency content of the available impact data, which are observed at ranges less than 1200 km, so have a somewhat higher frequency content at the receiver location than we expect for the CMBD impacts. For example, the lunar impacts have dominant frequencies of $\sim$2 Hz, whereas we expect the optimal detection band with the lowest noise is 0.2–0.9 Hz and waveform modelling is performed up to 1 Hz due to computational limitations.

As the range in estimated peak amplitudes stems from a fundamental lack of observed data in comparable contexts against which to check predictions and understanding of the relevant processes, the range of estimates described here cannot be constrained through further modelling; unless more observational data or more advanced modelling techniques become available. Rather, the uncertainties in our estimates reflect the general lack of knowledge of the excitation and propagation over large distances of impact-generated seismic waves.

Hence, even a single instance of impact detection from a source of known spatial and temporal localisation would therefore be of enormous value. It would offer the potential to better understand impact processes (especially seismic efficiency), enable us to make headway in understanding the sub-surface geology at the landing site

(through placing constraints on its seismic properties), as well as offering constraints on the attenuation and average propagation speed along the source-receiver path.

This strengthens the case for listening closely with InSight's instruments for the EDL sequence of Mars 2020. As the upper end of our certainly wide-ranging estimates suggests a reasonable probability of a signal being detected, a positive detection would be extremely useful in resolving the present uncertainty surrounding the propagation of the elastodynamic waves generated by impacts. The enormous advantage that this event holds in attempting to isolate its signal from the noise is that we know exactly the time and location at which it will be produced, and can reasonably estimate when these signals will reach InSight.

A non-detection would similarly enable us to further constrain the seismic detectability of impacts on Mars (in effect adding a datapoint on Fig. 4.4 at the level of the noise floor which represents an upper bound for the seismic signal amplitude), though admittedly by a smaller margin than a positive detection would. Because this measurement would quantify the maximum fraction of the CMBD's kinetic energy which was transmitted into seismic energy, it would in effect constrain the seismic efficiency of the impact. This quantity is particularly challenging to measure on account of its dependence on entirely local properties and geology, but is of significant interest as it effectively controls the degree of energy transmission into the solid ground during a planetary impact.

4.6.3 Extensions to this work

Having already discussed some of the limitations encountered in modelling of the CMBD impact signal which future work may seek to address (e.g. high-resolution simulations of oblique impacts and a better characterisation of the equivalent moment tensor), we briefly detail improvements to the modelling which may be made.

It is important to emphasise that we do not expect any of the effects not included in this paper to affect its conclusions, but these are discussed for completeness. They may well be more relevant to other applications of this methodology, for

example if an EDL event occurs with receivers in close proximity to the source, or applications to other planetary bodies.

4.6.3.1 Atmospheric attenuation

Both shock and linear acoustic waves experience an increased attenuation at height, impeding their long-range propagation. The 'classical' (terrestrial) acoustic attenuation due to viscosity, heat conduction, and diffusion is augmented on Mars by the strong molecular relaxation attenuation of CO_2 molecules (Williams 2001). This acts as a low-pass filter, limiting the range of infrasound frequencies which propagate in Mars' atmosphere to between 0.05Hz and the order of $\sim$1 Hz (Martire et al. 2020). The lower bound is the atmospheric cut-off frequency related to the pressure scale height (2.4 km), whilst the upper bound is related to the high (at least 2 db/100 km) molecular relaxation from CO_2. Further work is needed to more exactly constrain the attenuation dynamics in Mars' atmosphere (Petculescu and Lueptow 2007).

In this study we considered two *potential* classes of atmospheric signal: the long-distance propagation of acoustic waves in a waveguide (which we concluded do not reach InSight because of the geometry of the EDL and atmospheric structure), and the local-scale propagation of atmospheric-side portion of the coupled seismoacoustic signal in the region known as the sonic boom 'carpet'.

Because the latter is a local-scale effect, we neglected the effects of acoustic attenuation in the atmosphere (and as the predicted amplitude of the seismoacoustic signal at InSight is already below the noise floor, including it would not change our conclusions). However, in other applications of this methodology, it may become important to consider the attenuation of the waves in the atmosphere between the spacecraft and the ground.

4.6.3.2 Directionality

Our acoustic calculations are first-order and do not account for the direction of travel of the spacecraft (the fact that it is travelling almost exactly toward InSight upon arrival). The directionality is likely to have two effects: firstly, the amount of energy directed toward the lander may be reduced (as the majority of the energy

is radiated in a direction perpendicular to the vehicle's trajectory). Secondly, a small Doppler shift in the acoustic signal may be apparent. This will increase the frequency of the signal, and it will be more rapidly attenuated as a result.

Both of the consequences of the spacecraft's direction of travel being toward InSight are therefore to reduce the amplitude of any acoustic signal. As the signal is already well below the noise floor at the lander's location, neglecting these effects will not change the conclusions of this paper.

4.6.3.3 Low altitude acoustic sources

As discussed in Sec 4.5.2, there are two potential low-altitude (i.e. within the tenuous tropospheric waveguide) acoustic sources which occur as a result of the EDL sequence: the CMBD impact with the ground, and acoustic signal produced by the spacecraft once it is sub-sonic. The travel time to InSight for any such signal would be approximately 4-5 hours, though as discussed no detection is expected and this section is included to illustrate a methodology only.

Both of these will produce signals much weaker than the supersonic deceleration of the spacecraft, and in the case of the CMBD impacts quantitatively estimating the acoustic overpressure this will produce is challenging due to the multi-phase and non-linear nature of the problem.

For purely illustrative purposes, we consider an exemplar source of 1 Pa and 1 Hz frequency in the Mars 2020 landing region. This is a pressure perturbation $\frac{\Delta P}{P}$ of approximately 0.15%. This is substantially larger than most acoustic sources on Mars (Martire et al. 2020) and stronger and lower frequency than we expect either signal to be.

Even neglecting attenuation, the amplitude of this perturbation after propagating to the distance of InSight is no larger than 3x10^{-4} Pa (using a r^{-1} scaling) or 5x10^{-6} Pa (using a $r^{-1.5}$ scaling). For comparison, both values are far below the pressure noise floor of the APSS instrument ($\sim$2x10^{-3} Pa in the 0.1 to 1 Hz range) (Banfield et al. 2019; Martire et al. 2020).

4.6.3.4 Other relevant effects

For completeness, we also briefly detail other (lower-order) effects which may be relevant in applications of this methodology to other contexts.

The acoustic signal will be affected by surface topography and the temporal evolution of the atmosphere. If an acoustic signal were detected (which we do not expect to be the case here), more detailed consider of the non-linear infrasound propagation in the near-source region should be conducted if the signal is to be clearly associated with an EDL event. It is possible that other kinds of atmospheric waves (e.g. internal gravity waves) may also be excited by an EDL event, though again this is not relevant to the case discussed in this paper.

Modelling of the elastodynamic signal in the solid ground may also include consideration of the source mechanism (as discussed in Sec 4.5.3.2), and of three-dimensional heterogeneous effects including scattering and local geology.

4.7 Conclusions

We identified three possible types of seismoacoustic signals generated by the EDL sequence of the Perseverance rover: (1) the propagation of acoustic waves in the atmosphere formed by the decay of the Mach shock, (2) the seismoacoustic air-to-ground coupling of these waves inducing signals in the solid ground, and (3) the elastodynamic seismic waves propagating in the ground from the hypersonic impacts of the CMBDs.

(1) In the first case (atmospheric propagation), the stratification and wind structure in the atmosphere are such that the strongest signals produced will likely not be detectable at InSight, as they are reflected off the ground back up into the atmosphere. Signals produced in the lower 10 km of the atmosphere may be trapped and propagate for long distances, however the spacecraft will be subsonic by this point and will not be emitting substantial amounts of acoustic energy into the atmosphere. The Mach shock generated higher in the atmosphere will also have largely dissipated by the time it propagates down to this level. As such no

detectable signal is expected. The effects of attenuation and directionality were not included in this model, however both will serve only to reduce the amplitude of any signal at InSight.

(2) In the second case (air-to-ground transmission), the coupling is expected to be very weak. Combined with the substantial distance to InSight, we predict a maximum ground velocity amplitude at SEIS's position of 2×10^{-11} ms^{-1}. This is well below the noise floor at all times of day and hence is not predicted to be detectable.

(3) The generation of seismic waves by an impact comparable to the CMBD impact and the detectability of the seismic signal at large distance are not well understood. Using a combination of scaling relationships and wave generation/wave propagation methods, we estimate that the direct body wave arrivals from the impact may be detectable at InSight. In the realistic best-case (and assuming identical weather and noise spectra to the same period one martian year earlier), the requisite signal-to-noise ratio would be sufficient for a positive detection 40% of the time. It should be noted that our modelling was for only one of the two CMBD impacts. Based on data from the Mars Science Laboratory (Curiosity) landing in 2012, the two CMBDs will impact around 0.1 s and no more than 1 km apart. This separation is large enough that craters will not overlap spatially, and any interaction between the two signals will be in the linear propagation regime. As a result, the impact of two rather than one CMBD will increase the signal amplitude at InSight by no more than a factor of two, which is less than the uncertainty on the scaling estimates, as described above.

Such a P-wave signal would present itself as a sharp peak in the ground velocity recorded by InSight's SEIS instrument (Fig. 4.5) approximately 430 s after the impact of the CMBDs with the ground, just after 15:00 LMST (Perseverance time) or 20:30 LMST (InSight time). This is during the most seismically quiet part of the day at InSight (Banfield et al. 2020; Clinton et al. 2021). If detectable, the S-wave signal would be expected approximately 300 s later; and the travel-time difference would be of use in identifying the signal.

This is likely to be the only impact event with known source parameters during the lifetime of the InSight mission. The Chinese Tianwen-1 is also expected to land on Mars in the spring of 2021 (Wan et al. 2020), but due to a lack of published information on the EDL sequence and hardware, and the time and precise location of its landing, making predictions about the detectability of this signal is not possible; though we eagerly seek clarifying information.

As such, the case for listening for the Mars 2020 signal with InSight's instruments (SEIS and APSS) at the highest possible sampling rates is clear. Whilst this is the first time that such an event detection has been attempted on another planet, InSight's potential on this topic has already proved a source of inspiration in the popular media (*Away, Season 1, Episode 8* 2020).

Beyond Mars 2020 and Tianwen-1, this methodology may be extended to future missions including ExoMars (scheduled launch 2022) or Starship.

4.8 Appendix - seismoacoustic simulations in AxiSEM3D

4.8.1 Introduction

I undertook a number of preliminary simulations of seismoacoustic coupling to evaluate the sensitivity of the observed waveforms to the representation of the source mechanism prior to writing this paper. The rationale behind this was that if, at 5 s or so resolution, a seismoacoustic source in the atmosphere could be represented as a point force on the surface in the far-field, then the magnitude, geometry, and source-time function of the point force could be estimated independently using methods which are able to include non-linear (e.g. shock) effects.

This would give a more realistic description of the source's features than modelling the sonic boom as an impulsive point-source in the atmosphere with amplitudes estimated using inaccurate linear calculations alone. This would only be justified if the near-field propagation in the atmosphere itself made little difference to the waveforms in the far-field.

It was readily apparent after initial testing that the use of a point force on the surface as opposed to a pressure source in the atmosphere made little difference to

seismograms at the distance of InSight. Thus, the methods presented in Sec. 4.4.1.1, which include non-linear effects, proved to be a more effective way to calculate the surface overpressure than those detailed below. They also had the advantage that the effects of winds could be accounted for, which they could not be in AxiSEM3D.

As such, this method and plots were not included in the final version of the published paper. The methodology and a sample of the results are thus shown here for completeness as an Appendix should anyone wish to push the parameter space toward where atmospheric effects are more relevant - for example, on Venus or Titan.

In this section only, I will use a modified IASPI terminology to describe the phases being modelled: *i* will represent an upgoing wave in the atmosphere (i.e. from the source or ground toward space), and *I* a downgoing one. *iI* would hence be a wave which propagates upward, reflects off the 'top' of the atmosphere, and then hits the surface.

4.8.2 Choice of boundary condition

Seismoacoustic simulations for the atmosphere cannot in general use a pressure-free reflecting upper boundary condition, as was used for the ocean in Chapter 2, on account of the fact that the atmosphere has no definite upper edge from which all incident energy is reflected. Rather, an absorbing boundary condition is probably more appropriate, such that any energy in the atmosphere which propagates upward toward space is not reflected back downward. Of course, this does not preclude reflection within the atmosphere of energy from genuine density contrasts (e.g. in waveguides).

The implementation of an absorbing boundary condition is significantly more complex than that of a reflecting one (Clayton and Engquist 1977; Kosloff and Kosloff 1986), as energy must be effectively absorbed across a range of incidence angles and frequencies with a minimum amount of spurious reflection. The most basic absorbing boundaries function best for energy which is normally incident, and may be totally ineffective for grazing waves. For our purposes, such a simple basic boundary condition solved at the top of the atmosphere is sufficient as we

primarily seek to remove all waves which travel directly upward toward space from the source (*i*) and hence are at near-normal incidence. Removal of these waves' energy is sufficient for our purposes as they are most likely to produce spurious waveforms at regional or teleseismic distances.

Other spurious waves which this boundary condition will not remove do occur in simulations where they would not in reality. These include *iI* phases which propagate near-horizontally in the atmosphere before hitting the boundary (and hence are at the wrong angle to be absorbed) or surface multiples *IiI* (which hit the surface, reflect back up to the top of the atmosphere at non-normal incidence, and then back down to the surface). However, the significant amounts of time these phases spend in the low-velocity atmosphere means that when their energy finally strikes the ground and undergoes seismoacoustic coupling, it is delayed enough that it does not contribute to the amplitude of the P-wave, which is of most interest here (if one were studying Rayleigh wave codas created by seismoacoustic coupling, this may require more careful consideration).

4.8.3 Theory

Consider Eqn. 2.2 from Chapter 2, the weak form of the wave equation in the fluid domain Ω_F with solid-fluid boundary Σ, density ρ, and bulk modulus κ:

$$\begin{aligned} &\int_{\Omega_F} \kappa^{-1} w \, \partial_t^2 \chi \, d^3\mathbf{r} + \int_{\Omega_F} \rho^{-1} \nabla w \cdot \nabla \chi \, d^3\mathbf{r} \\ &= \int_{\Sigma} \mathbf{u} \cdot w \hat{\mathbf{n}} \;\; d^2\mathbf{r}. \end{aligned} \tag{4.3}$$

where $\mathbf{u}$ is the displacement vector, $\mathbf{w}$ is an arbitrary, test function, and we define a scalar potential χ such that $\mathbf{u} = \rho^{-1}\nabla\chi$ or equivalently $\partial_t^2 \chi = -P$ in Ω_F. Note that we restrict ourselves to a spherically symmetric model here for ease of reference, and hence drop the identity tensor $\mathbf{I}$ from Eqn. 2.2.

The addition of an absorbing boundary condition at interface Σ_{ABC} and pressure source f modifies this equation to:

$$\begin{aligned}&\int_{\Omega_F} \kappa^{-1} w\, \partial_t^2 \chi \, d^3\mathbf{r} + \int_{\Omega_F} \rho^{-1} \nabla w \cdot \nabla \chi \, d^3\mathbf{r} + \int_{\Sigma_{ABC}} (\rho c)^{-1} w \partial_t \chi \;\; d^2\mathbf{r} \\ =&\int_{\Sigma} \mathbf{u} \cdot w\hat{\mathbf{n}} \;\; d^2\mathbf{r} + \int_{\Omega_F} w f \, d^3\mathbf{r}.\end{aligned} \tag{4.4}$$

Which is identical to the formulation of Komatitsch et al. (2000), but represented in displacement rather than velocity. This equation is then solved in AxiSEM3D in the same way as described in Chapter 2.

4.8.4 Results

Initial testing using this boundary condition aimed to evaluate whether, at intermediate-to-high resolutions (around 5 s), the waveforms produced by seismoacoustic coupling were affected by how the source was represented (i.e. on the surface or in the air, with or without an actual atmosphere, as a point force or distributed pressure source).

As per Fig. 4.7, the seismograms at 10° distance (i.e. in the far field many seismic wavelengths away but still much closer than InSight) are very similar in each of the four cases. Small differences are apparent, for example in the peak-to-mean ratio of Rayleigh wave amplitudes in the 500-800 s range, but the broad character of the waveforms is the same. This is especially true for the body waves, which are the primary phases of interest on Mars due to the expected scattering of the surface waves. A minor difference in travel time would also be apparent (because in the atmospheric case, it takes around 50 s for waves produced at 10 km to reach the surface), though to ease comparison this was removed and seismograms in this Fig. 4.7 were aligned by P-wave arrival time.

4.8.5 Discussion

This experiment indicated that representing a seismoacoustic source in the atmosphere as a point force on the surface is reasonable at these distances and in this frequency range. It is also computationally cheaper as the low sound speeds in the atmosphere need not be explicitly meshed, with the further advantage that non-linear effects can be included in the atmospheric part of the simulations, providing a

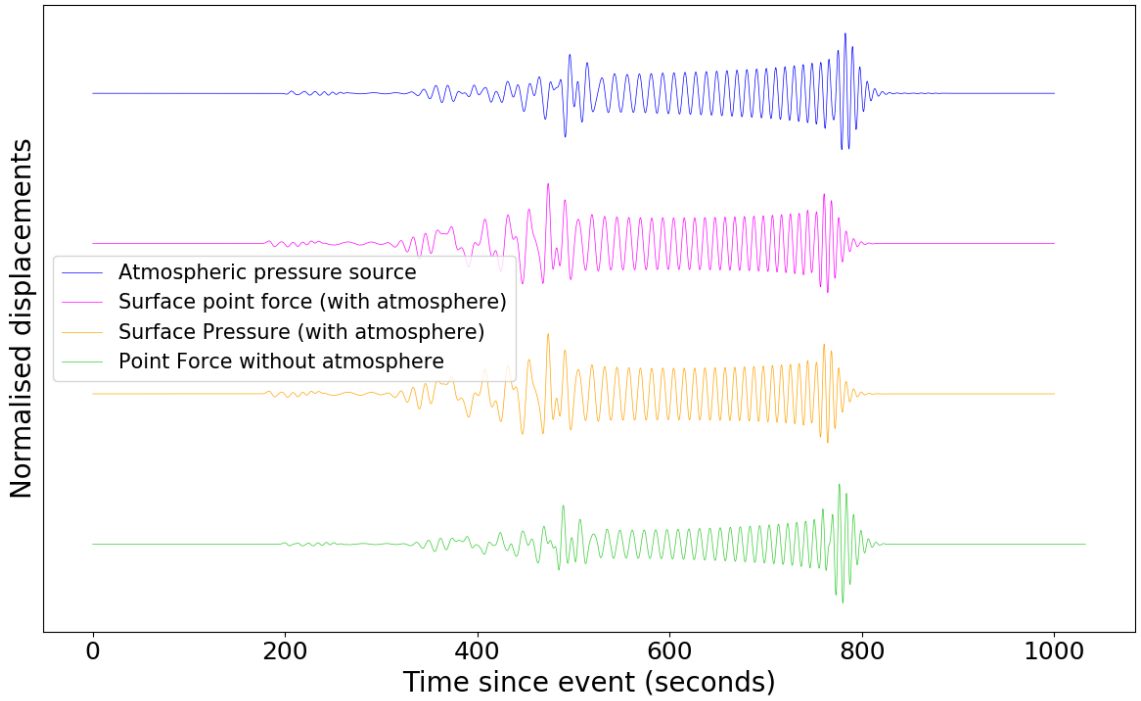

Figure 4.7: Seismograms for four different representations of a seismoacoustic source in the atmosphere, shown in the vertical component at 10° distance and bandpass filtered between 5 and 500 seconds. The blue, pink and yellow traces include an actual, global-scale fluid atmosphere; whilst the green trace does not. The blue trace shows the seismogram from an isotropic pressure source in the atmosphere, the pink trace uses a vertical point force at ground level instead, and the yellow trace uses an isotropic pressure source just below the surface. The green scenario is identical to the pink one, but without an atmosphere. Traces are normalised here by peak P-wave amplitude, as there is no direct and simple correspondence between the magnitude of a point force and of a moment tensor.

more accurate estimate of the overpressure at the surface which is used to scale the elastodynamic source amplitude. For these reasons, the methods used in Sec. 4.4.1.1 were used in seismoacoustic calculations rather than those discussed in this section.

The method described here retains some advantages which may be beneficial in other modelling scenarios not considered in this paper. Firstly, the full waveforms may be simulated in both the atmosphere and solid ground, meaning that frequency-dependent effects such as dispersion or waveguide trapping may be accurately modelled in the atmosphere as well. Ray-tracing methods cannot achieve this in the same way.

Secondly, in a case where the surface waves are of greatest interest, the small differences in amplitude and waveform shape caused by the presence of an atmosphere may also be examined, which cannot be done if a free-surface boundary condition with point force source is used. Scattering and the extreme bulk modulus contrast between the regolith and atmosphere mean that these effects are negligible on Mars, but on a planet with a much smaller bulk modulus contrast (e.g. Venus), the effects of the atmosphere may be more pronounced, akin to the effects on seismograms of the oceans on Earth.

I look back one more time. It's like a crater, a hole where something happened.

— Emma Donoghue

5 Mars, Part III

Contents

5.1 Preface 139
5.2 Summary 140
5.3 Introduction 141
5.4 Methods 141
5.5 Data 143
5.6 Discussion 145
5.6.1 Distance-amplitude scaling relationships 145
5.6.2 Seismic efficiency and mantle attenuation 147
5.7 Conclusions 151

5.1 Preface

This paper was published as Fernando et al. (2021b):

Fernando, B., Wójcicka, N., Maguire, R., Stähler, S. C., Stott, A. E., Ceylan, S., Charalambous, C., Clinton, J., Collins, G. S., Dahmen, N., Froment, M., Golombek, G., Horleston, A., Karatekin, O., Kawamura, T., Larmat, C., Nissen-Meyer, T., Patel, M. R., Plasman, M., Posiolova, L., Rolland, L., Spiga, A., Teanby, N. A., Zenhäusern, G., Giardini, D., Lognonné, P., Banerdt, W. B., and Daubar, I. J.

Seismic constraints from a Mars impact experiment using InSight and Perseverance, *Nature Astronomy*, 2021.

The full funding sources for all research incorporated here, besides mine, include: the ETH Zürich ETH+ 'Planet Mars' program, the Centre National D'Études Spatiales and the Agence National de le Recherche (France), and the Belgian Science Policy Office through the ESA PRODEX program. InSight is managed and operated by the Jet Propulsion Laboratory at the California Institute of Technology on behalf of NASA. The two spacecraft which took part in imaging of the CMBD impact site are the Mars Reconnaissance Orbiter (also operated by JPL) and the ExoMars Trace Gas Orbiter (operated by the European Space Agency in partnership with Roscosmos).

5.2 Summary

- We present results from the InSight mission's attempt to detect the arrival of the Perseverance Rover on Mars in February 2021.
- No seismic signal from the landing sequence is detected, though through thorough measurement of the noise levels during the landing window we are able to place an upper bound on the peak P-wave amplitude produced by the impact of the spacecraft's cruise balance mass devices with the surface of $1.1\text{x}10^{-10}$ m/s.
- This constrains the seismic efficiency of impacts in the landing region (the ratio of impactor kinetic energy to released seismic energy) to be less than 3% - the first time this quantity has ever been constrained in-situ on Mars.
- We are also able to exclude the total impactor momentum as being an appropriate determinant of peak P-wave amplitude for oblique impacts such as this.
- These conclusions may now be utilised in future studies of the dynamics of natural impact events.

5.3 Introduction

NASA's InSight (Interior Exploration using Seismic Investigations, Geodesy and Heat Transport) mission has operated a sophisticated suite of seismology and geophysics instruments on the surface of Mars since its arrival in 2018. On February 18, 2021, we attempted to detect the seismic and acoustic waves produced by the entry, descent, and landing (EDL) of the Perseverance Rover using the sensors onboard the InSight lander.

Similar observations have been made on Earth using data from both crewed (Hilton and Henderson 1974; Qamar 1995) and uncrewed (Ishihara et al. 2012; ReVelle and Edwards 2007) spacecraft, and on the Moon during the Apollo era (Latham et al. 1970b), but never before on Mars or another planet. This was the only seismic event to occur on Mars since InSight began operations which had an a priori known and independently constrained timing and location. It therefore had the potential to be used as a calibration for other marsquakes recorded by InSight.

Here we report that no signal from Perseverance's EDL is identifiable in the InSight data. Nonetheless, measurements made during the landing window enable us to **(1)** place constraints on the distance-amplitude relationships used to predict the amplitude of seismic waves produced by planetary impacts, and **(2)** to place in-situ constraints on martian impact seismic efficiency (the fraction of the impactor kinetic energy converted into seismic energy).

5.4 Methods

Pre-landing predictions based on modelling and on data from Mars Science Laboratory (Curiosity)'s EDL in 2012 suggested that the impact of the spacecraft's two Cruise Balance Mass Devices (CBMDs, also referred to as Cruise Mass Balance Devices, or CMBDs) might produce seismic waves of high enough amplitude to be detectable at InSight's position, approximately 3450 km to the east (Bierhaus et al. 2013; Fernando et al. 2021a).

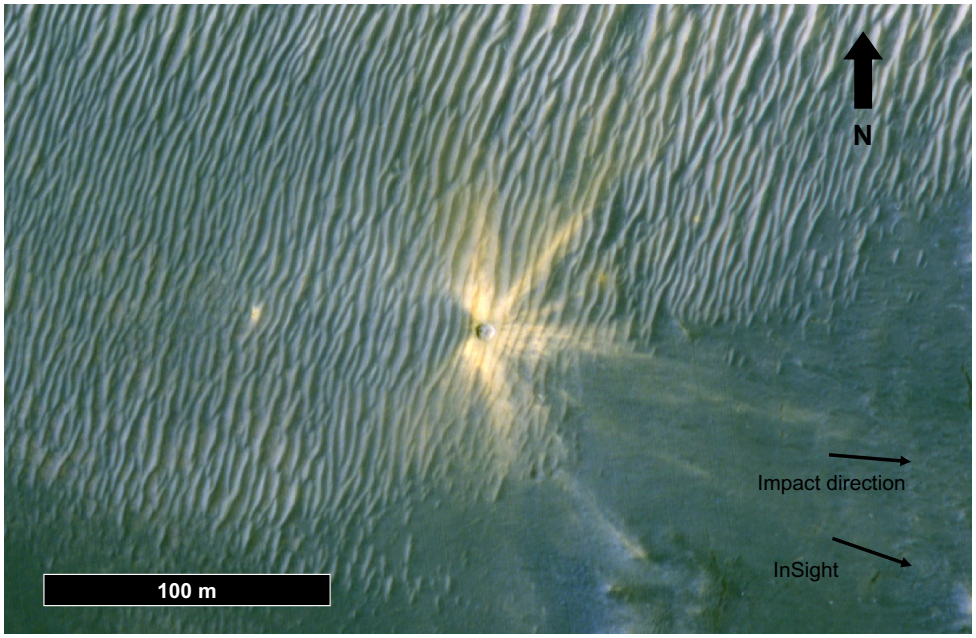

Figure 5.1: A high-resolution orbital image of one of the CBMD craters. This is the largest at ∼6 m in diameter, and is located at 18.956°N, 76.202°E. This image is a crop of the enhanced colour HiRISE image with observation ID ESP_069231_1990. North is up, and illumination is from the left. Arrows indicate the direction of impact as inferred from the asymmetric ejecta pattern and the direction towards the InSight lander. The abundant aeolian ripples that trend north-south demonstrate that the surface materials are poorly consolidated, fine grained sand harnessed by the wind. Image: NASA/JPL/University of Arizona.

The CBMDs are solid, 77.5 kg tungsten blocks used to adjust the spacecraft's lift-to-drag ratio during EDL. Data from the spacecraft's computer indicate that the commands were generated to fire the pyrotechnic releases holding them in place at 20:28:27 UTC, within 1 s of each other. At this time, the spacecraft was projected to be at an altitude of 1253 km and travelling at a planet-relative velocity of 4753 m/s.

NASA's trajectory reconstruction indicates the CBMDs hit the surface at 20:40:33 ±3 s UTC (around five minutes before the rover's touchdown), at a speed of approximately 3816 m/s and an oblique angle of around 10 degrees from the horizontal.

The impact craters from the CBMDs were imaged by the CTX (Malin et al. 2007) and HiRISE instruments (McEwen et al. 2007) onboard NASA's Mars Reconnaissance Orbiter on May 3, 2021. Images identified the craters at a position around 18.9°N 76.2°E (Fig 5.1).

The most promising candidate seismic phase for detection at InSight (4.5°N, 135.6°E) was expected to be a ballistic compressional (P) wave excited by the CBMD impact (Fig. 5 in (Fernando et al. 2021a)). The P-wave arrival time was predicted to be approximately 420 ±20 s post-impact, or 20:47 UTC (19:50 InSight Local Mars Solar Time, LMST, on Sol 793).

The two variables determining whether or not this signal would be identifiable in InSight data were the seismic noise during the arrival window (constrained post-landing from the seismometer recordings), and the amplitude of the impact-induced P-wave. The former was particularly low during the arrival window, as it occurred during the part of InSight's day when the atmospheric noise is lowest. Distance-amplitude scaling curves were used to make one prediction of the peak amplitude, (Teanby 2015; Wójcicka et al. 2020; Fernando et al. 2021a). These relationships come with large uncertainties as they are calibrated using only terrestrial and lunar data at closer range.

Data were recorded throughout the night (in UTC) of February 18-19 2021 on InSight's Very BroadBand seismometer (VBB) (Lognonné et al. 2019), and the pressure and wind sensors of the Auxiliary Payload Sensor Suite (APSS) (Banfield et al. 2019).

5.5 Data

Data from the arrival window are plotted in Fig. 5.2. The main signals observed are a large marsquake around 01:25 UTC, and irregularly spaced glitches which are artificial in origin. No other signal is observed which cannot be explained as noise excited by atmospheric phenomena, as corroborated using the wind speed and pressure measurements.

Therefore, we conclude that no signal associated with the CBMD impact is identifiable above the detection threshold (as defined in Fig. 5.2b), precluding a P-wave amplitude larger than $1.1\text{x}10^{-10}$ m/s.

As the marsquake detected at 01:25 UTC fits within the projected time window (4-5 hours post landing) for infrasound (low-frequency sound) waves to arrive from

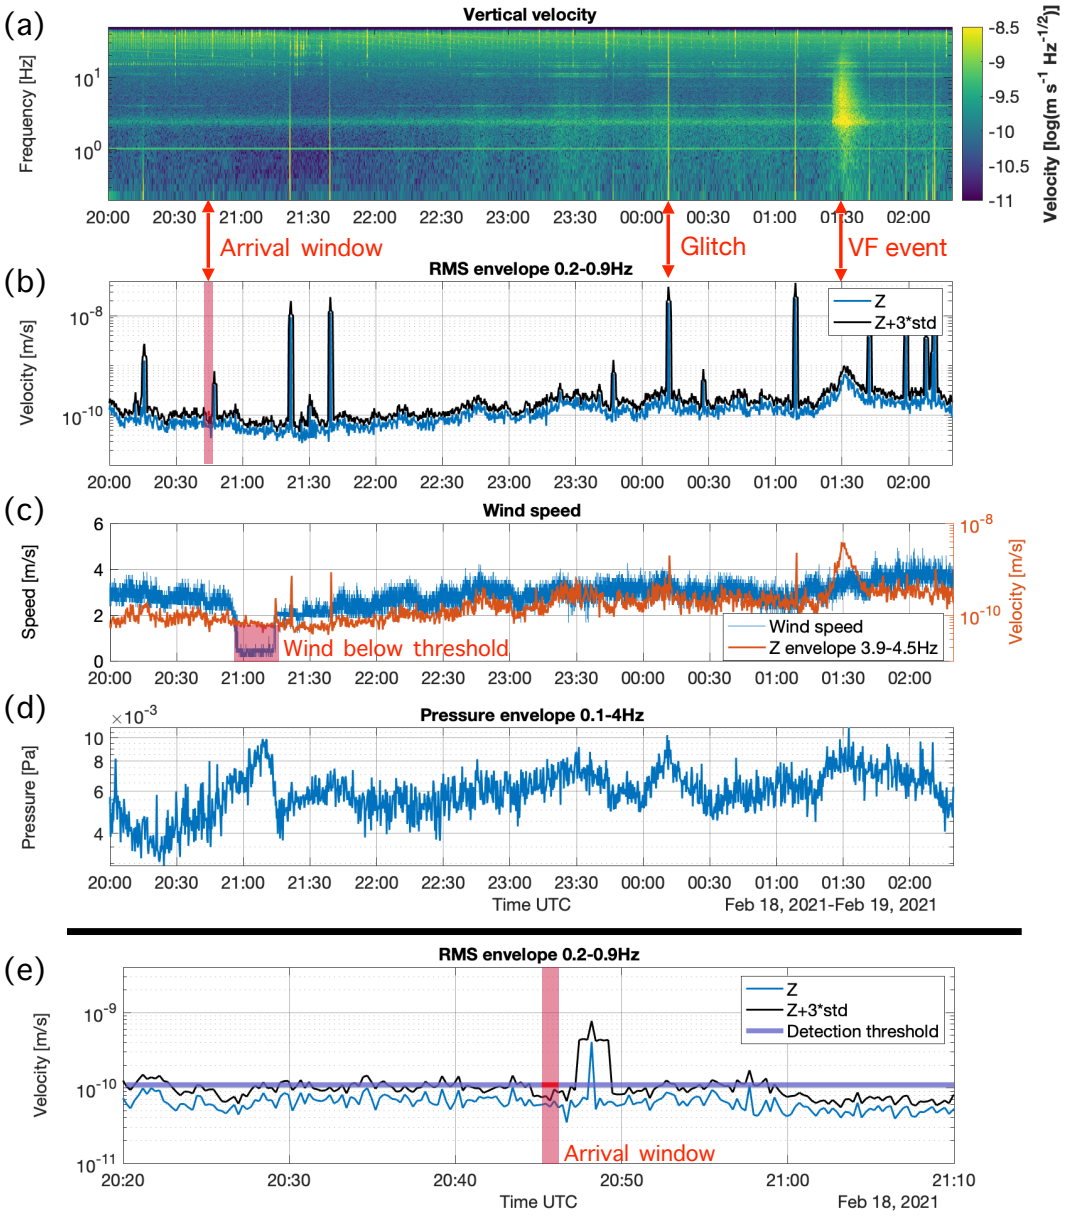

Figure 5.2: Data recorded by InSight during the landing window. **Panel (a)** shows a spectrogram of the vertical ground velocity recorded by the Very Broad-Band (VBB) seismometer, whilst **(b)** shows the RMS envelope of the vertical velocity in the 0.2–0.9 Hz frequency band most suited to isolating mantle-going phases (blue; labelled 'Z') as well as the RMS + 3 standard deviations (black). In both panels, glitches in the system are recorded as sharp vertical features in the spectrogram and peaks in the RMS envelope, with one exemplar highlighted. These glitches are easily distinguished from seismic signals (Scholz et al. 2020). The arrival window for the Cruise Balance Mass Device (CBMD) P-wave is highlighted in red, as is an unrelated marsquake observed around 01:45 UTC. **Panel (c)** shows the wind speed (blue) and the RMS envelope of the vertical seismometer velocity in the 3.9–4.5 Hz frequency band (red) - the latter, in this frequency range, contains a known oscillation mode of the spacecraft which is excited by the wind and can be used as a proxy for the wind speed (Ceylan et al. 2021). The absence of wind measurements around 21:00 UTC occurs where the wind speed drops below the instrument threshold. (Banfield et al. 2019). **Panel (d)** shows the RMS envelope of the atmospheric pressure in the 0.1–4 Hz band which is also correlated to wind speed (Charalambous et al. 2021). **Panels (e)** and **(f)** show a detail of the gray area in panels (a) and (b) respectively, where the largest amplitude event is a marked glitch just before 20:50 UTC, and our defined 'detection threshold' (an upper bound on the peak P-wave amplitude observed in the arrival window, defined as the RMS of Z + 3 standard deviations) is shown as a horizontal purple line. Note that the Local Mean Solar Time (LMST) values corresponding to the left and right edges of panels (a)-(d) are 19:07 (Sol 793) to 01:15 (Sol 794).

Perseverance's EDL, we will briefly outline why we believe it is unrelated and simply coincidental.

This assessment is based on the signal's amplitude, shape, and frequency content, which are entirely different to those of an infrasound wave. EDL-related infrasounds at these distances should have amplitudes below the noise floor (Fernando et al. 2021a); whilst the high frequencies observed in this event (up to ∼30 Hz) preclude an airborne propagation path as the high CO_2 concentration in the martian atmosphere would rapidly attenuate them (Bass and Chambers 2001). Rather, this event ('S0794a') is a standard 'Very High Frequency' (VF) marsquake (van Driel et al. 2021), in this case with an origin at a distance of ∼1,100 km, much closer than Perseverance. As of the end of April 2021, 40 of these events have been recorded by InSight, also often in the late evening (Clinton et al. 2021). These events are likely caused by tectonic or otherwise internal geological processes.

Having established that there is no signal from the EDL recorded in the data, we will now consider what this non-detection can be used to infer about impact processes on Mars.

5.6 Discussion

5.6.1 Distance-amplitude scaling relationships

There are a number of approaches to using distance-amplitude scaling relationships to predict peak P-wave amplitudes from impacts.

Depending on whether the impactor energy, total momentum, or vertical component of the momentum are used to scale the amplitude, and whether terrestrial missile-impact data or lunar spacecraft-impact data are the basis, predicted amplitudes vary by up to two orders of magnitude when extrapolated to distances of 3,500 km. Five such standard scaling relationships are shown in Fig. 5.3 - one based on impactor energy (Teanby 2015), and four based on impactor momentum (Wójcicka et al. 2020).

Of the latter, two are based on data from artificial lunar impacts, which occurred into almost cohesionless material (black curves, (Latham et al. 1970b)), and two on

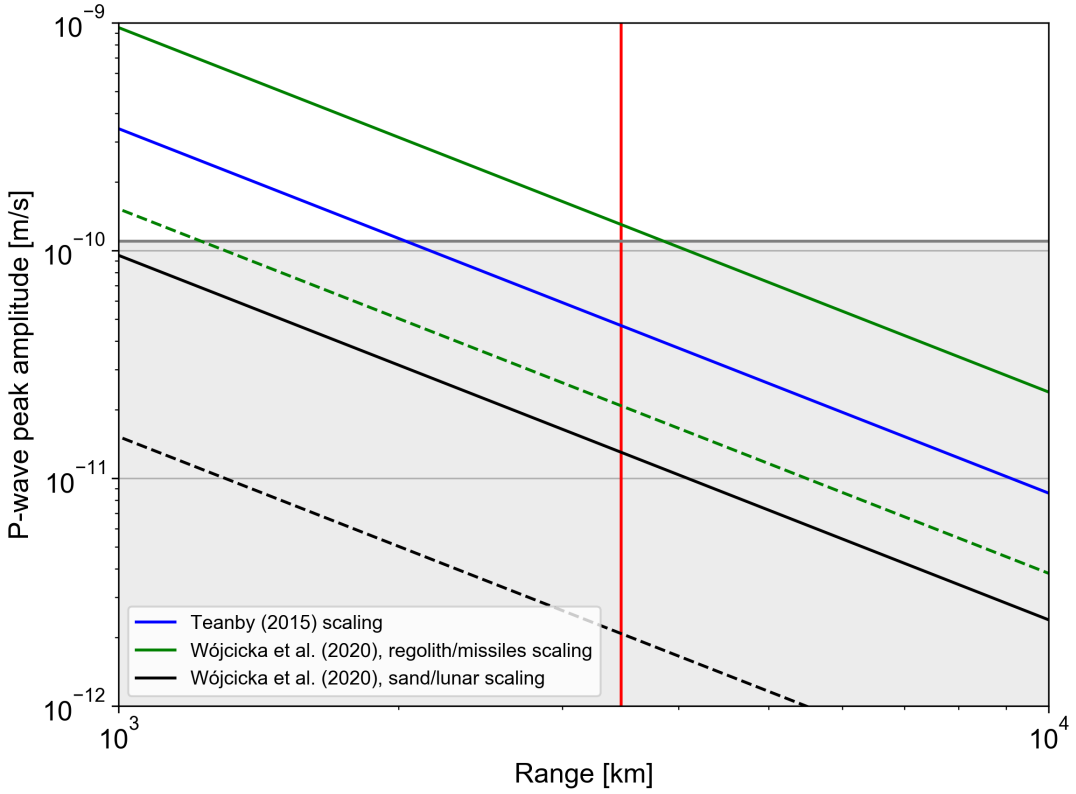

Figure 5.3: Predicted P-wave amplitudes for one Cruise Balance Mass Device (CBMD) impact calculated using different scaling relationships with distance from source. The (Teanby 2015) scaling (blue line) scales the amplitude with square root of impact energy ($E_k = 6 \times 10^8$ J). The (Wójcicka et al. 2020) scalings are based on either extrapolation of terrestrial missile impacts (green lines) or lunar artificial impact data (black lines), using either total impact momentum (solid lines, $p = 3 \times 10^5$ Ns) or the vertical component of impact momentum (dashed lines, $p_z = 4.8 \times 10^4$ Ns) to scale the peak P-wave amplitude. The red vertical line marks the distance between the Perseverance landing site and the InSight lander. The horizontal grey line indicates the detection threshold. Four of the five relationships are below the detection threshold (in the gray area), and one is above (white area).

terrestrial missile impacts into weakly cohesive regolith soils (green curves, (Latham et al. 1970a)). In each case, one curve uses the total impactor momentum as the determinant of peak P-wave amplitude (solid lines), whilst the other uses the vertical component of the impactor momentum (dashed lines).

The derived upper bound on the peak P-wave amplitude of 1.1x10^{-10} m/s is shown as a horizontal grey line in Fig. 5.3, and at the distance in question lies below the solid green scaling curve (total impactor momentum; terrestrial missile data). This means that if the solid green scaling was a robust predictor of amplitude,

a signal ought to have been detected.

This result indicates that the distance-amplitude scaling relationship based on terrestrial missile impacts into weakly cohesive soil, and the assumption that seismic wave amplitudes scale with total impactor momentum, are not appropriate in this case. There are three possible implications.

Firstly, for a highly oblique impacts such as this, this result suggests that the vertical component of the momentum may be a more appropriate quantity to scale by.

Alternatively, as the impact data that this model was based on was collected at much closer distances and on Earth, this may indicate a stronger attenuation of seismic wave amplitudes with distance on Mars. A third possibility is that artificial impacts on the lunar surface represent a better analogue for the seismic response of small impacts on Mars than those which occur into terrestrial soils.

5.6.2 Seismic efficiency and mantle attenuation

The upper bound on the peak P-wave amplitude may be used to place a joint constraint on the impact seismic efficiency (k_s, which is site and impact-specific) alongisde the average mantle attenuation (Q_μ) along the path from Perseverance's landing site to InSight.

Because k_s and Q_μ are entirely independent of each other (the former being related to generation of seismic waves on a local scale, and the latter to their propagation on a global scale), jointly constraining their values in this way is valid.

k_s is particularly poorly constrained, due to its high sensitivity to local conditions and the lack of relevant in-situ measurements (Daubar et al. 2018). All estimates for Mars thus far have therefore used modelling, simulation, or material analogues to estimate its value, and no in-situ measurements other than those in this paper exist.

Q_μ and Q_κ are the two quality factors used to describe viscoelastic (intrinsic/inelastic) attenuation properties within the solid part of the planet. The former is associated with shear properties, and the latter with bulk properties.

For typical martian mantle rheologies, $Q_\mu \gg Q_\kappa$, hence Q_κ has little influence on P-wave amplitudes. The P-wave attenuation Q_P can then be approximated as $Q_P = \frac{9}{4}Q_\mu$.

Tidal observations suggest that Q_μ does not vary strongly through the mantle (Bagheri et al. 2019), so we assume that our value of Q_μ is a reasonable descriptor of the average attenuation along the source-receiver path (see Sec. 5.4).

We do not treat the crust separately to the mantle, as the P-wave propagates almost entirely within the latter and previous studies show comparably negligible crustal attenuation (Lognonné et al. 2020; van Driel et al. 2021). In order to relate our observed upper bound on the P-wave amplitude to Q_μ and k_s, we first estimate the seismic moment associated with the CBMD impact using the following empirically-derived relationship from (Teanby and Wookey 2011):

$$M_0 = \frac{(k_s E_k)^{0.81}}{4.8 \times 10^{-9}} \tag{5.1}$$

where E_k is the impactor kinetic energy (in this case, 1.1 GJ). Note that the separation between the CBMDs in space and time at impact ($\sim$ 1 km and around 1 s) is large enough that the impact processes occur independently of each other. E_k is thus the kinetic energy of a single CBMD.

5.6.2.1 Accounting for geometric spreading

As the P-wave reduces in amplitude during its propagation through both attenuation and geometric spreading of the wavefront, making an estimate of Q_μ requires us to quantify the effects of the latter.

To do this, we undertake full-waveform seismic simulations in a purely elastic (non-attenuating) medium. Thus, the only energy loss in these simulations along the source-receiver path occurs due to the spreading of the wavefront. The effects of attenuation using an average Q_μ may then be applied as a post-simulation correction,

We use an interior velocity and density model called 'TAYAK' which combines geochemical data with geodetic constraints (Clinton et al. 2017; Khan et al. 2018). This was chosen on the basis of its good fit to other marsquake data (Brinkman

et al. 2021). The source is represented as an explosion just below the surface of magnitude M_0, as per (Daubar et al. 2018).

Simulations are conducted using the spectral element solver AxiSEM (Nissen-Meyer et al. 2014b) in combination with Instaseis (van Driel et al. 2015) at a dominant source period of 1 s. This allows us to account for amplitude decreases due to the geometric spreading of the wavefront alone.

Once the amplitude in the non-attenuating case, A_0, (where $A_0 = A_0(M_0)$) has been found, we account for the effects of attenuation, assuming an effective $Q_{\mathrm{P,eff}}$ that averages over the propagation path S:

$$Q_{\mathrm{P,eff}}^{-1} = \int_S Q_{\mathrm{P}}^{-1}(s)\mathrm{d}s. \tag{5.2}$$

The amplitude in a frequency band between f_1 and f_2 (here, 0.2 and 0.9 Hz) is estimated through application of the following equation:

$$A(Q_{\mathrm{P,eff}}, M_0) = A_0(M_0)\frac{\int_{f_1}^{f_2} \exp\left(-\frac{tf}{2Q_{\mathrm{P,eff}}}\right)\mathrm{d}f}{f_2 - f_1}. \tag{5.3}$$

where A is the amplitude at InSight. Note that the ratio $Q_{\mathrm{P}} = 9/4Q_\mu$ assumes a Poisson solid with a standard acoustic speed to shear speed ratio v_p/v_s of $\sqrt{3}$. $t = 420$ s is the predicted travel time of a P-wave and Q is assumed to be frequency-independent.

Combining Eqns 5.3, as derived above, and 5.1, allows us to determine which combinations of k_s and $Q_{\mathrm{P,eff}}$ (and hence k_s and Q_μ) produce a permissible A at InSight's position, i.e. one which is remains the detection threshold.

5.6.2.2 $k_s - Q_\mu$ constraints

Fig. 5.4 shows joint constraints which may be placed on k_s and Q_μ: as increasing the value of either parameter leads to larger predicted amplitudes, these are co-constrained such that the resultant P-wave amplitude does not exceed our detection threshold.

The quality factor has previously been determined to lie in the range 300-1000, with the lower bound derived from tidal observations (Bagheri et al. 2019) and

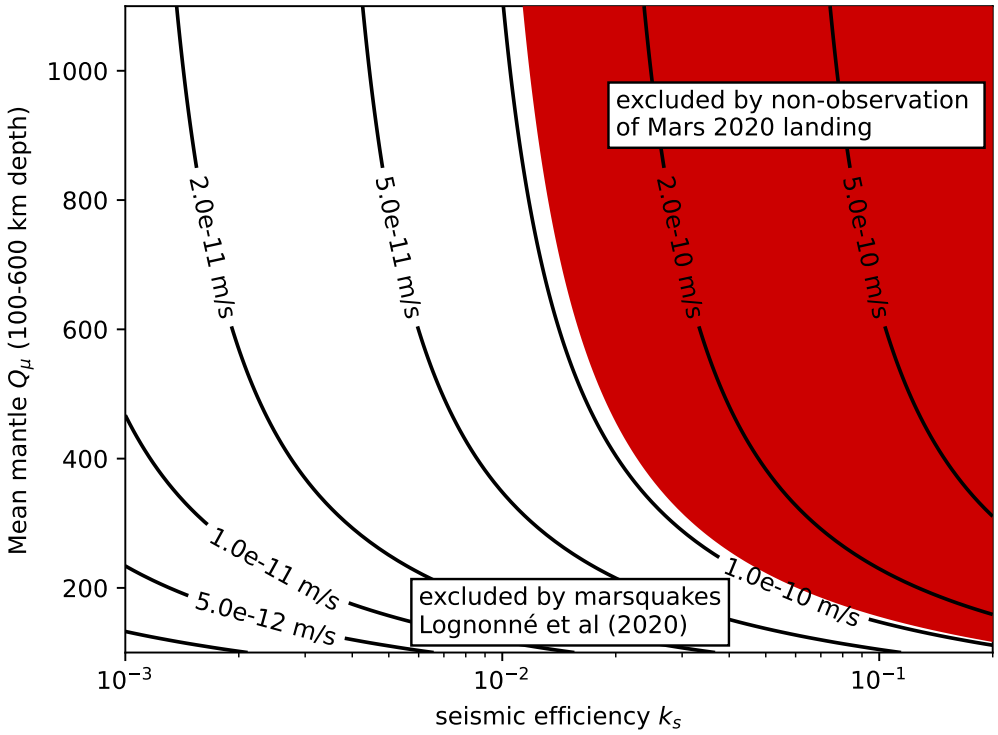

Figure 5.4: Constraints on seismic efficiency k_s, given mean mantle quality factor Q_μ, as derived from the non-observation of the CBMD (Cruise Balance Mass Device) impacts. Black curves are lines of constant amplitude, as would be recorded at InSight, and can be understood as either the line of valid k_s-Q_μ combinations if a P-wave were observed with said peak amplitude, or the upper bound on k_s-Q_μ combinations if the noise floor were of this amplitude. Thus, k_s-Q_μ combinations in the red zone are precluded as they would require P-wave amplitudes above the noise floor. The orange zones are incompatible with the spectra of previously recorded marsquakes and estimates of the quality factor based on tidal excitations. The green zone shows the range of combinations which would satisfy these observations (although smaller values have been proposed previously, we only plot $k_s \geq 10^{-3}$ as this is the range of interest from these results.) The rightmost point of this zone is 3% and marks the maximum possible value of k_s at this location on Mars.

the upper from observations of the spectra of marsquakes (Lognonné et al. 2020; Giardini et al. 2020). Thus, we constrain the maximum value of k_s to be 3%, corresponding to the rightmost corner of the green zone in Fig. 5.4.

This upper limit on k_s is compatible with experimental (terrestrial and lunar) and simulated (terrestrial, lunar, and martian) measurements, where k_s ranges from 10^{-6} (in porous sand/regolith) to 10^{-3} (in stronger, non-porous materials) (Latham et al. 1970b; Latham et al. 1970a; Richardson and Kedar 2013; Wójcicka et al. 2020;

Rajšić et al. 2021). Values up to 2-3% have been observed in underground nuclear explosions in stronger target materials (Patton and Walter 1993).

Fig. 5.1 shows one of the impact craters. The surrounding morphology is representative of this region of Mars, and consists of aeolian bedforms and unconsolidated/poorly consolidated regolith. The thermal inertia of the surface at the impact location is $\sim$250 $\mathrm{Jm^{-2}s^{-1/2}K^{-1}}$ (Putzig and Mellon 2007), indicative of a dominantly sandy surface. Although a quantitative relationship between k_s and target properties has yet to be derived, our upper bound on k_s is consistent with a surface material of this type.

The true seismic efficiency is likely to be lower than the maximum which we derived. Nonetheless, this does still demonstrate a practical method through which it may be constrained, and is the first time that such a constraint has been derived on Mars. From this, we may draw the robust conclusion that the conversion of kinetic energy into seismic energy on Mars is no more efficient than the most efficient such terrestrial coupling.

5.7 Conclusions

We used the non-detection of the seismic waves from the impact of Perseverance's CBMDs to show that the total impactor momentum is a poor predictor of amplitude at this distance and impact angle, if a terrestrial-based scaling is assumed. In this case, the lunar-based scaling or a relationship based on impactor energy may be more appropriate. This result could also indicate that effects of attenuation on impact-generated seismic waves are stronger than previously estimated.

We also used the non-detection to constrain the impact seismic efficiency to be less than 3%, which is compatible with geological analysis of the impact site, and commensurate with previous estimates which used modelling or proxies in place of in-situ data.

The methodology presented here provides a basis for using seismic (non) detection of artificial impacts to infer subsurface properties, and could in future be applied during geophysical missions to any of the Solar System's icy or rocky bodies.

Data availability

InSight APSS/TWINS/PS data are hosted by New Mexico State University (here). InSight SEIS data is available in the form of a seismic event catalogue and waveform data is hosted at the IPGP datacentre or from IRIS (here).

HiRISE data are available through the Planetary Data System (here).

At first sight it would seem that the deep interior of the Sun and stars is even less accessible to scientific investigation than any other region of the universe.

— Sir Arthur Stanley Eddington

6 The Sun

Contents

6.1 Preface 153
6.2 Summary 154
6.3 Introduction 154
6.3.1 Motivation 155
6.4 Theory 156
6.4.1 Included effects 156
6.4.2 Neglected effects 158
6.4.3 Other approaches 161
6.5 Implementation 162
6.5.1 Outer boundary condition 162
6.5.2 Seismic profile 164
6.5.3 Gravity 168
6.6 Results 171
6.6.1 Mesh design and simulation cost 171
6.6.2 Wave propagation 173
6.7 Conclusions 175

6.1 Preface

This material has not yet been published.

Assistance in editing the AxiSEM3D code to support wave propagation in the Sun was received from Kuangdai Leng, and support for mesh design from Martin van Driel.

As discussed below, consideration of hyperbolic formulations

of the gravitational potential were inspired by discussions with Ryosuke Hirai. I undertook the simulations, designed meshes, solved stability problems, and wrote this manuscript.

6.2 Summary

- The behaviour of the Sun affects the Earth, for example through space weather events and solar variability.
- The origin of these phenomena within the Sun lies deep below the surface, a realm in which helioseismology is ideally placed to make observations.
- Observations must be understood and contextualised within the framework of robust modelling, but the generation of synthetic seismograms on the Sun is significantly more challenging than it is on the Earth.
- Here, we present an adaptation of AxiSEM3D which enables pressure wave propagation to be simulated within the solar interior.
- This involves modifications to the source term and boundary condition; and consider we consider the code's performance as compared to existing methods.
- Additions required to account for the effects of stochastic excitation and gravitation are also discussed, though these are not yet implemented.

6.3 Introduction

As discussed in Chapter 1, the domain of helioseismology may be split into two parts: global and local heliosesimology. The former is concerned with analysis of normal-mode, free oscillations of the entire star; whilst the latter is more like the terrestrial seismic behaviour already discussed in this thesis, and indeed was the foundation for the set of terrestrial techniques known as ambient noise interferometry. We will focus on the latter in this Chapter, but the former is probably a more natural introduction to the subject, so let us begin there.

The observed normal mode oscillations of the Sun consist of acoustic waves trapped in resonant cavities. The edges of these cavities are bounded at their upper edge by the solar 'surface' (a term whose definition we will return to later), and at their lower edge by the increasing sound speed. The depth to which a particular acoustic mode penetrates is dependent on its frequency and its horizontal wavenumber, or equivalently its quantum numbers n (order) and l (angular degree). Three quantum numbers are sufficient to characterise any global mode: n, l, and m, the latter being the angular order. We will not worry too much about m in this chapter: in the absence of radial-symmetry breaking processes (rotation, fluid flows, magnetic fields - none of which we have in these models), the $2l + 1$ values of m for any given l correspond to eigenmodes degenerate in frequency.

The solar convective zone is dominated by p-modes, where pressure is the restoring force. The solar core is dominated by g-modes, where buoyancy is. Because of the unstably stratified nature of the convective zone, g-modes are evanescent in this region. This means that they reach the observable solar surface with extremely small amplitudes, and have yet to be conclusively detected (Appourchaux et al. 2010).

6.3.1 Motivation

Historically, scientific interest in helioseismology was motivated by a desire to understand the Sun's global internal structure: particularly to place constraints on its core properties, the generation of its internal magnetic field, and its variability. With the advent of better long-term monitoring and increased computational power, interest has also expanded to smaller-scale features and their variability - for example sunspots and sub-surface global-scale flows.

These observations are also of particular interest from a broader stellar physics perspective. Whilst the Sun is only one of many G-type stars in our galaxy, it is the only one whose physical properties can be measured nearly in-situ and at high-resolution. Further discoveries in solar physics through helioseismology therefore have clear potential to be more broadly of interest within the stellar seismology community.

6.4 Theory

It should already be clear that inverting for seismic properties on the Sun is going to require a somewhat different approach to that which we might use on the Earth. The same is, therefore, of course true for conducting seismic simulations as well - though in many cases the end goal - the development of sensitivity kernels which quantify the effects particular structures have on the wavefield - is the same.

It is worth briefly detailing the physical phenomena which make this the case: not all are accounted for in this model, but it is important to note which are and which are not incorporated.

6.4.1 Included effects

Let us begin by discussing effects which are included in these models (or rather, effects for which an *attempt* has been made to include them). We will deal with their actual implementation in Sec. 6.5, this section considers the relevant theoretical background.

6.4.1.1 Boundary condition

A change to the outer boundary condition used in finding solutions to the weak form of the equations of motion was effectively the topic of an entire chapter (Chapter 2), so it may come as little surprise to the reader that we begin this particular adventure with a discussion along the same lines. The trouble is that the question 'Which boundary condition should we use on the Sun?' has a rather complicated answer (Fournier et al. 2017).

The primary reason for this is that the boundary condition that is actually felt by waves within the Sun depends on the frequency of said waves. As noted previously, observed global normal-mode oscillations are the product of acoustic waves trapped in a resonant cavity between their lower turning depth and the surface. As such, the the surface boundary condition must be at least partially reflecting for waves of these periods (5 minutes/300 s) or longer.

It transpires that the upper boundary ceases to be reflecting at frequencies higher than ~ 5 mHz (periods of ~ 200 s), whereupon the waves cease to be trapped and instead propagate up into the chromosphere as travelling rather than standing waves (Jiménez et al. 2011). The exact radius at which waves are reflected from the 'surface', even at periods longer than this 'acoustic cutoff', also depends on their frequency (the relationship is complex - see Jiménez (2006)).

In this work, a reflecting outer boundary condition was used. For further details of the implementation used, see Sec 6.5.1.

Finally, we should note that measurements of displacements on the Sun do not necessarily correspond to 'surface' displacements. Dopplergram maps of the Sun necessarily use a single spectral line, and because the temperature structure in the solar photosphere is not constant, different spectral lines have their peak emissions at different depths. For example, two common chromospheric lines used in solar physics are hydrogen-alpha (656.3 nm) and calcium-K (393.4 nm), the latter of which is from cooler and deeper layer in the chromosphere. The Helioseismic and Magnetic Imager (HMI) onboard the SOHO spacecraft images the Fe (I) line at 617.3 nm (Nagashima et al. 2014), which is a photospheric line - and hence maps the motion within a different layer of the solar atmosphere to the other two.

None of this is too much of a problem, however - in AxiSEM3D the receivers can easily be placed below the surface, to ensure that the displacements measured are recorded at the desired depth.

6.4.1.2 Sources

The implementation of a pressure-domain source was discussed in the epilogue to Chapter 4, so we will only touch on it briefly here. In short, excitations in this formulation are represented as point-like pressure (isotropic) sources, i.e. corresponding to a diagonal moment tensor. This is a reasonable assumption: in the Sun, convective turbulence leads to stochastically generated pressure perturbations which excite (and damp) acoustic oscillations.

When written, AxiSEM3D was only able to handle single sources, due to the need to rotate the source onto the north pole to make the azimuthal Fourier expansion converge (see Chapter 2, Section 2 for full details). As the most recent release which supports distributed sources (Szenicer et al. 2020) has not been adapted to work for the Sun we will not discuss it here. In short, it would enable a more realistic representation of source geometries for the Sun, which in general do not include localised injections of seismic energy, unlike in terrestrial earthquakes.

6.4.1.3 Seismic profile

One might assume that the fully-fluid nature of the Sun would simplify the problem of seismic wave propagation, in that the lack of shear support precludes S-type waves from propagating and that the lack of hard interfaces prevents conversions between phases. Sadly, they would be wrong.

Although the mesher used in AxiSEM3D can easily be adapted to generate a computational grid for the solar sound speed profile, the gradients in seismic parameters involved cause all manner of problems when it comes to actually running simulations.

6.4.2 Neglected effects

Next, let us consider those effects which are not included in this version of AxiSEM3D. In each case I will expand upon the justification for neglecting each effect: though in some cases, this is based on what is computationally feasible within AxiSEM3D rather than some more physical explanation.

By the end of this sub-section the reader may be of the opinion that we have neglected more effects than we have included - but it is important to bear in mind that many of these are second-order, and are far more important in global than local helioseismology. Furthermore, the existing codes most comparable to AxiSEM3D are incapable of including all these effects either: for example as discussed further below, Gizon et al. (2017) do not include rotation, 3D effects, self-gravity (the perturbation to gravitational equilibrium induced by waves), or rotation either.

6.4.2.1 Gravity

The initial goal of this project was to incorporate the effects of gravitation into AxiSEM3D, which in theory would enable simulation of buoyancy-restored g-modes. In short, this turned out to be impossible in the solar case, and as such the simulations presented later in this chapter do not include any gravitational effects. Because this line of work involved some mathematical considerations, we will discuss further details and a potential implementation in Sec. 6.5.3.

6.4.2.2 Magnetic fields and solar cycles

Solar activity is modulated by a 11 year periodicity known as the 'solar cycle'. The Sun's behaviour at the solar maximum is very different to that at solar minimum, with changes in behaviour strongly modulated by the influence of the Sun's magnetic field. For example, the acoustic cutoff frequency is known to change throughout the solar cycle (Jiménez 2006), whilst the number of sunspot complexes (which have lower oscillation amplitudes, (Gizon et al. 2009)) also varies (Zhao et al. 2001). However, these are all second-order effects, which have a noticeable impact on solar behaviour, but do not change the bulk structure of the Sun. The models used in this experiment therefore represent averaged properties.

6.4.2.3 3D structures

The main aim of a project such as this would be to generate sensitivity kernels for structures in the Sun, many of which have some relationship to magnetic phenomena discussed above.

It is in theory possible to add in perturbations to the sound speed (effectively 3D heterogeneities, as described in Chapter 2), which could be used to simulate the changes to the seismic profile below features of interest. However, the capacity to do this in fluid regions has not yet been included in AxiSEM3D, though doing so is unlikely to pose any particularly technical challenges.

Once this has been done, a solar adjoint method may be used to generate sensitivity kernels as has been done for AxiSEM3D on the Earth (Szenicer et al. 2020).

6.4.2.4 Fluid flows

The circulation patterns within the Sun also have a noticeable impact upon the frequencies of p-modes, though again these are second order effects (Hanasoge et al. 2016; Mandal et al. 2018; Liang et al. 2018). This means that helioseismic measurements can be used to study the internal fluid dynamics of the Sun, and conversely to that accurate modelling should therefore also include the effects of these fluid flows.

Again, it is not possible to directly include the effects of velocity fields (i.e. advection) in AxiSEM3D, so such effects cannot be modelled here. We did speculate that introducing a seismic 'anisotropy' in the seismic profile might be one way to approximate such behaviour, where the effective acoustic wave velocity would cease to be isotropic, and would be greater in one direction and reduced opposite in another. However, this would require substantial modification to the pressure formulation used to represent the displacement in the fluid, and has the potential to introduce other numerical artefacts as well.

6.4.2.5 Attenuation

As noted previously, the mechanism responsible for both damping and excitation in the solar convective zone is the same: stochastic turbulence. However, AxiSEM3D does not currently include the capacity to attenuate seismic waves in fluid regions, meaning that no attenuation is included in these models. However, this is not particularly problematic from a stability perspective nor would it be particularly difficult to actually implement.

In reality, excitation processes on the Sun are continuous, but this is balanced out by continuous damping, which prevents the mode amplitudes from growing to infinity. In our models, however, we use a gaussian source-time function with half-width equal to the mesh resolution, so there is no threat of continuous energy input leading to an instability.

In a terrestrial model, neglecting attenuation would mean that that the absolute amplitudes of acoustic waves could not be calculated, but this is the case anyway

in our models of the Sun, as we do not include effects such as gravitation or use an appropriate boundary condition. The travel-times of waves should nonetheless be unaffected by neglecting attenuation effects.

6.4.2.6 Rotation

For a body undergoing slow rigid rotation (slow here implying a rotation rate much less than the breakup velocity at which rotational forces would tear the star apart), the effects of rotation are multiple. First, a departure from sphericity occurs. This flattening is easy to account for with AxiSEM3D. Secondly, the normal mode peaks are split by quantum number m. In a normal mode code it is possible to use a rough analytical relationship to constrain the degree of splitting, but this is not really possible in a time-domain code such as AxiSEM3D - however it is not too much of a problem, because rotation has less of an effect in small-scale, local helioseismic simulations. The combination of rotation and gravitation enables the propagation of new types of waves, for example Rossby waves, but we cannot consider these here because we have no gravitation in the models.

In actuality, the rotation of the Sun is much more complex and very much not rigid-body. The rotation velocity varies with both latitude and depth, leading to modifications to waveforms which are very much dependent on the source-receiver geometry. Because there is no way to simulate such effects in AxiSEM3D, we will not consider them here.

6.4.3 Other approaches

A number of other approaches to conducting local-scale helioseismic simulations exist. A common aim is to construct sensitivity kernels, which offer a quantitative explanation of how the wavefield around localised structures of interest is modified (Fournier et al. 2018). The development of sensitivity kernels for the Sun also forms the basis for imaging of the solar interior.

Early approaches used ray-theory based methods, due to their relative simplicity (Kosovichev 1996). More recent work used Born approximation methods, first in

cartesian and then spherical geometry (Böning et al. 2016). Such methods are based on normal mode summation, akin to similar work done on Earth.

The methods of Gizon et al. (2017), based on the original AxiSEM formulation (Nissen-Meyer et al. 2014b), are perhaps most closely related to AxiSEM3D, in that they use an axisymmetric background medium and neglect the effects of gravity. This method does restrict the user to a '2.5D' (axisymmetric) domain. Additionally, the solution is found in the frequency domain. This offers a number of advantages over the time-domain method employed by AxiSEM3D. Note however that frequency domain methods come with their own set of disadvantages - and the balance between these depends on the structure of the world being simulated. Whilst AxiSEM3D's speed and efficiency is an advantage, the overall complexities of time-domain simulations on the Sun probably tip the balance the other way as compared to Earth.

6.5 Implementation

For those areas where substantial modifications were made to the AxiSEM3D formulation beyond those presented in Chapters 2 and 4, I will provide a brief outline of the modifications which were attempted. These include modifications to the outer boundary condition, the seismic profile, and a discussion of gravitational effects. The implementation of the source is the same as that presented in Chapter 4, so I will spare the reader a repeat.

6.5.1 Outer boundary condition

There are two possible choices of boundary condition in the solar formulation of AxiSEM3D: reflecting and absorbing. Reflecting boundary conditions, as introduced in Chapter 2, are most suited in these settings to waves at periods longer than 200 s (i.e. frequencies below the acoustic cutoff frequency). Absorbing boundary conditions, as introduced in Chapter 4, are more suitable for waves with frequencies above the acoustic cutoff, which are not trapped and hence do not 'bounce' off of the surface.

6.5.1.1 Reflecting boundary conditions in the Sun

Eqn. 2.2 must be modified slightly to include sources in the fluid, and to discount solid fluid interfaces, Σ:

$$\begin{aligned} &\int_{\Omega_F} \kappa^{-1} w\, \partial_t^2 \chi \, d^3\mathbf{r} + \int_{\Omega_F} \rho^{-1} \nabla w \nabla \chi \, d^3\mathbf{r} \\ &= \int_{\Omega_F} w f \, d^3\mathbf{r} \end{aligned} \tag{6.1}$$

Where here we have a fully fluid domain Ω_F with bulk modulus κ. Again, $\mathbf{u}$ is the displacement vector, w is an arbitrary test function, and we define a scalar potential χ such that $\mathbf{u} = \rho^{-1}\nabla\chi$ or equivalently $\partial_t^2\chi = -P$ in Ω_F. Note that we restrict ourselves to a spherically symmetric model here for ease of reference, and hence drop the identity tensor $\mathbf{I}$ from Eqn. 2.2.

6.5.1.2 Absorbing boundary conditions in the Sun

The simplest absorbing boundary condition, as introduced in Eqn. 4.2, must also be modified now that we have no solid-fluid interfaces:

$$\begin{aligned} &\int_{\Omega_F} \kappa^{-1} w\, \partial_t^2 \chi \, d^3\mathbf{r} + \int_{\Omega_F} \rho^{-1} \nabla w \cdot \nabla \chi \, d^3\mathbf{r} + \int_{\Sigma_{ABC}} (\rho c)^{-1} w \partial_t \chi \;\; d^2\mathbf{r} \\ &= \int_{\Omega_F} w f \, d^3\mathbf{r}. \end{aligned} \tag{6.2}$$

6.5.1.3 Choice of boundary conditions

Because simulations were not conducted at periods shorter than 200 s, a reflecting boundary condition is used here. Ideally we would be able to replicate the transition from reflecting to absorbing boundary and conduct simulations at higher frequencies, with energy only reflected if incident on the surface at less than 5 mHz. However, use of a frequency-dependent boundary condition is not possible in the current formulation of AxiSEM3D, and would likely significantly increase the computational cost of such simulations. This does not pose too much of a problem in this scenario, as the absence of non-linear effects which would change wave frequencies means that seismic waves do not undergo frequency changes.

A more sophisticated boundary condition based on a perfectly-matched layer formulation (Berenger 1994) or a radiation condition (Gizon et al. 2017; Barucq et al. 2018) would no doubt be more accurate, but would involve considerably more complex modifications to the code structure. These would be substantially easier to implement in frequency-domain codes than time-domain codes like AxiSEM3D.

In this thesis, I have arbitrarily chosen the 'surface' to be at a radius of 695,700 km. This is the same value used by Christensen-Dalsgaard et al. (1996), but models here do not include the 500 km of atmosphere used in Model S, on the grounds that meshing it is simply too challenging and expensive. We will discuss that particular adventure in the next sub-section.

6.5.2 Seismic profile

As discussed in Chapter 2, AxiSEM3D uses the *Salvus Mesher Lite* package to generate its 2D computational meshes (Afanasiev et al. 2019). Previously, we showed that a mesh including and overlying ocean can be generated without too much difficulty, as the contrast in bulk modulus is relatively small. Adding in an atmosphere is a little more challenging as we saw in Chapter 4, due to the tendency for there to be a local maximum in the sound speed above the surface (at least on Mars), but slight smoothing of the velocity profile resolves this. On the Sun, however, the problem is far more serious, and requires much more careful simulation design.

6.5.2.1 Gradients

On the Sun, the density at the core is approximately 10 orders of magnitude higher than that at the surface. The variation in sound speed is smaller, at around two orders of magnitude - though still appreciable (Fig. 6.1). Whilst the absolute values of sound speed and density at the core and surface in isolation are not a problem from the point of view of generating the mesh, the gradient between the core and surface is.

The weak form of the wave equation, Eq. 6.2, may be written in matrix form as:

$$\mathbf{M}\nabla\ddot{\chi} + \mathbf{K}\nabla\chi = \mathbf{f} \tag{6.3}$$

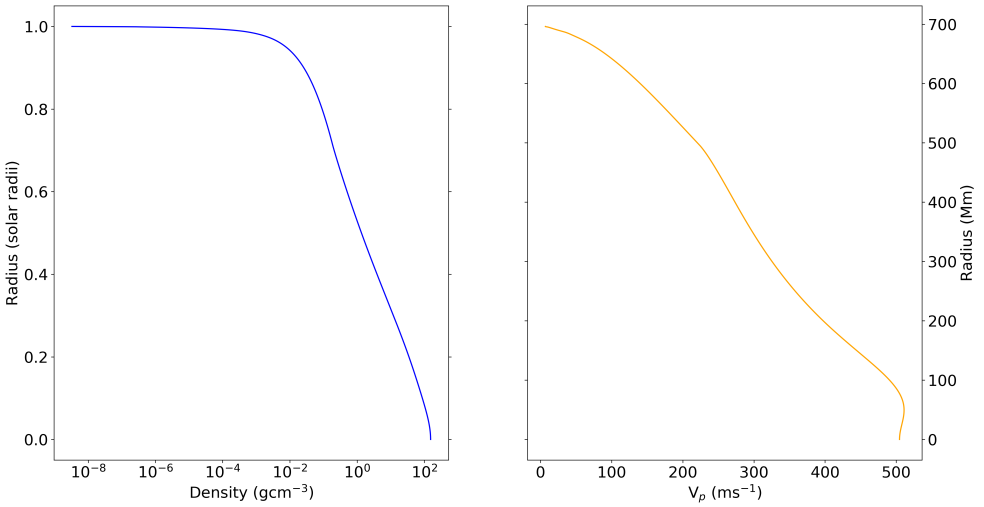

Figure 6.1: The solar seismic profile: density at left, sound speed at right. The left axis is in units of solar radii, the right axis is the equivalent in millions of metres. Note that in the left panel, the x axis is logarithmic. Features of particular note include the maximum in sound speed near the core corresponding to the second ionisation zone of helium, the change in curvature of the sound-speed profile around the tachocline (0.7 solar radii), and the enormous gradient in the density near the solar surface. Data are taken from 'Model S' of Christensen-Dalsgaard et al. (1996), and are derived from high-precision global helioseismic measurements.

Where **M** is the mass matrix and **K** is the stiffness matrix. It is in this form that we actually solve for the fluid potential χ.

In order to ensure that the mass matrix remains diagonal when assembled, AxiSEM3D uses Lagrange polynomials as interpolating functions. These have the form:

$$\phi_i(\xi) = \prod_{i \neq j}^{N+1} \frac{\xi - \xi_j}{\xi_i - \xi_j} \tag{6.4}$$

where ξ is our reference coordinate system on an elemental variable (see Eqn. 2.3 for more details of the mapping), and ξ_i and ξ_j are the interpolation points on our reference interval$[-1, 1]$, with $i, j = [-1, N + 1]$ for an N-th order interpolation. The choice of these polynomials ensures that the mass matrix is trivially invertible, saving a considerable amount of computational resource as inversion algorithms need not be employed at every timestep.

However, the strong radial gradients in density (and sound speed) involved across the Sun's entire radius also give rise to strong intra-element gradients within the mesh. These gradients must then be accommodated by the Lagrange polynomials - and Lagrange polynomials are not always the best behaved.

The reader may recall that interpolation often gives rise to 'overshoots', wherein the interpolating function takes on erroneously large values at particular points. Lagrange polynomials are no exception, and suffer from Runge's phenomenon - the occurrence of high-amplitude oscillations at the edges of an interpolation element (Fay and Webster 1996).

Initial simulations using a solar profile showed continued instability, which proved troublesome to diagnose. We eventually determined that the default polynomial order (4) was unable to accommodate the intra-element gradients, to the degree that the density was being interpolated as negative at some points. This led to the mass matrix acquiring negative entries, producing non-sensical displacements and hence instability.

Once diagnosed, it was possible to increase the interpolation order to 6, resolving the issue. However, as the choice of interpolating polynomials is tied to the use of a specific set of quadrature points in spectral-element type methods, this also necessitated increasing the quadrature order by a factor of 1.5.

An alternative method to deal with the inter-element density variations would be to remove the density terms from the mass matrix and evaluate them separately using a more appropriate quadrature scheme, however this is would require significant code editing and would be much more involved than increasing the quadrature order, even if it might produce cheaper simulations.

6.5.2.2 Lack of interfaces

The challenges of generating a computational mesh are always significant, and the difficulty of optimising any mesh are even greater. An optimal mesh must include the minimum number of solution points that can be achieved without

compromising resolution or stability, whilst also remaining conformal and honouring interfaces correctly.

To achieve this, seismic meshes of the type used in AxiSEM3D include refinement layers (see Fig. 2.1 in Chapter 2 for an illustration in the terrestrial case). These are conventionally placed along strong velocity contrasts, and are bounded above (shallower) by lower velocity layers with smaller elements, and below (deeper) by higher velocity layers with larger elements.

This enables seismic modelling codes to use the smallest number of elements required to satisfy a given Courant-Freidrichs-Levy condition (in this case, using a reasonable canonical value of 0.6). It is natural to place these refinement layers at the position of physical discontinuities, for example the seafloor or Moho. The *Salvus Mesher Lite* meshing software is designed to do this automatically, inserting refinements as needed in order to minimise the total number of elements within the mesh whilst maintaining stability.

On the Sun, the lack of hard interfaces and the strong gradients in seismic properties meant that the automated mesh design resulted in a grid which was so over-sampled as to make the number of elements (and hence the runtime and memory requirements of simulations) unnecessarily and unfeasibly large.

After some experimentation, the addition of 'fake' or ersatz discontinuities at specific radii was sufficient to bring the computational requirements down to something feasible on a supercomputer. In the *Salvus Mesher Lite* package, these consist of repeated lines in the seismic profile.

However, because these discontinues must be added manually, with visual inspection of the mesh used to determine potential placement locations, this procedure is far from optimised - but crucially, the resultant meshes are stable and suitable for use in simulations. One could no doubt come up with a more sophisticated way of deciding where to place refinement layers (as the mesher itself does for the Earth or Mars), but this is beyond the scope of this paper.

6.5.3 Gravity

Having adapted the seismic profile, source term, and outer boundary condition to the Sun, we turn our attention to consideration of the effects of gravitation. This section will detail some of the basic mathematics involved, before going on to explain why inclusion of gravitational effects in AxiSEM3D turned out to be unfeasible in the solar case - though would be entirely suitable in the lunar, martian, or terrestrial cases.

6.5.3.1 Elliptic formulation

The gravitational potential within the Sun can be considered in terms of the Newtonian gravitational potential ϕ through the Poisson equation for a mass density ρ and gravitational constant G:

$$\nabla^2\phi = 4\pi G\rho \tag{6.5}$$

This elliptic equation has no time dependence, meaning that in formulations using this definition of the gravitation, perturbations to the potential ϕ are felt everywhere within the body instantaneously. Resultantly, the simplest implementation of gravitational effects requires solving the Poisson equation everywhere within the mesh at every time step, which is enormously wasteful from a computational standpoint and resistant to parallelisation.

Most codes, such as SPECFEM, which include gravitational effects, make what is known as the Cowling approximation (Cowling 1941). Here, perturbations to the gravitational potential are neglected. In the stellar case, this is particularly justifiable when considering high-order p- and g-modes.

Physically, this involves discounting the mass redistribution which occurs as a result of the presence of high-amplitude seismic waves, which can appreciably deform the planet or star in question. This assumption is generally valid on Earth, expect for long-period normal mode oscillations. A similar approximation is generally made for the Sun, too.

We evaluated whether we might in AxiSEM3D be able to both 1) include self gravity and 2) overcome the computational cost induced by the non-causality of the

Poission equation through the use of a hyperbolic, rather than elliptic, formulation for the gravitational potential.

6.5.3.2 Hyperbolic formulation

A hyperbolic, causal formulation for the gravitational potential comes naturally from considering gravitation within the context of General Relativity, wherein changes in gravitational potential propagate at a finite speed (c). It is of course true that from a seismic perspective the difference between a change which propagates at lightspeed and one which is instantaneous is meaningless, in that the light-crossing timescale for the Earth is negligible even compared to the minimum reasonable timestep one might use in seismic simulations.

However, Hirai et al. (2016) present a formulation for the gravitational potential which is derived from the Einstein Field Equations but crucially which involves the propagation of gravity at a finite speed (c_g) which is less than c:

$$-\frac{1}{c_g^2}\frac{\partial^2\phi}{\partial t^2} + \nabla^2\phi = 4\pi G\rho \tag{6.6}$$

To understand this, let us assume a static background model in equilibrium with a single source. The only displacement which occurs is that which propagates away from the source location. If this source also involves mass redistribution (i.e. a large earthquake) the gravitational potential changes everywhere within the body almost instantaneously, but the seismic waves travel much more slowly. From our perspective, the perturbations to the potential are only important in the regions where seismic waves have already reached, because elsewhere there are no seismic waves to be affected by gravity.

Phrased another way, although the potential in the near-source region and at the antipode of the source change effectively instantaneously, we need only concern ourselves with the near source-region at first, as it takes some time for seismic waves to reach the antipode and hence changes in the gravitational field around said position are unimportant in the first few timesteps we run through.

A consequence of this is that it is reasonable to reduce the propagation speed of the perturbations to the potential from the speed of light in a vacuum, c, to a reduced velocity, c_g. This accounts for the fact that changes at distances to the gravitational field in regions which seismic waves have not yet reached are irrelevant from our seismic perspective.

The question then becomes how far c may be reduced without introducing artefacts into the solution - the driving factor being that the smaller the value of c_g, the larger the timestep may be and hence the more efficient simulations are.

It is probably reasonable to assume that so long as perturbations to the gravitational potential reach points on or within the body at least one seismic period before the seismic waves arrive, the behaviour will be a reasonable approximation to solutions to Eq. 6.5. This provides a lower limit for the value of c_g.

Furthermore, Eq. 6.6 is now hyperbolic, and hence much more straightforward to parallelise on high-performance computing architectures. Of course, this is not the only approach that can be taken to these problems: for example, we also considered the use of multigrid methods, which are naturally suited to solving differential equations where important behaviour is found on multiple different scales (e.g. see Li and Fleming (1997) for discussion of application of multigrid methods to coastal wave propagation). However, as we will see shortly, this approach would still cause issues relating to stability because of the unstable stratification of the convective zone.

6.5.3.3 Limitations

Regardless, after spending some time on the development of these methods, it became clear that they are not suited to use in a time-domain code of this type on the Sun. Specifically, the unstably stratified nature of the solar convective zone means that the effects of the gravitational potential are always to move fluid parcels upward in this region - i.e. simulated displacements would involve both an oscillatory part (due to the waves) but also an exponentially growing part (due to the motion of fluid parcels themselves).

In a linear, time-domain static code such as AxiSEM3D it is unclear how this would be dealt with, and would likely require substantial intervention at every timestep. In frequency domain codes, where the oscillation frequency can be split into a real and complex part trivially, adding in a damping term to prevent this instability is simple. For this reason, we halted attempts to introduce gravitational effects into the solar formulation of AxiSEM3D, though noting that for the case of rocky or stably stratified worlds, this would not be an issue. However, because gravitation is a long-period effect, and AxiSEM3D's performance advantages come about at higher frequencies, the case for integrating these effects in the non-solar formulation is weaker.

6.6 Results

Let us briefly cover the progress that was made using this code. The astute reader will note that the results section in this chapter is significantly shorter than in others, on account of this line of research being curtailed due to the technical challenges discussed above. Nonetheless, as a tribute to the many hours spent working on it, I have included it here.

6.6.1 Mesh design and simulation cost

The number of ersatz discontinuity layers required to stably mesh the Sun appears to depend on the seismic period required. Fig. 6.2 shows an exemplar mesh, at 10,000 s resolution.

In this example, two refinement layers are shown. However, for shorter periods, more are needed - experimentation during this part of the project used up to seven refinement layers, concentrated near the surface. If we wished to include the solar chromosphere as well, where the sound speed gradients are strongest, even more would likely be needed. On the other hand, in a local helioseismology setting, with simulations which do not extend all the way down to the core, fewer would likely be sufficient.

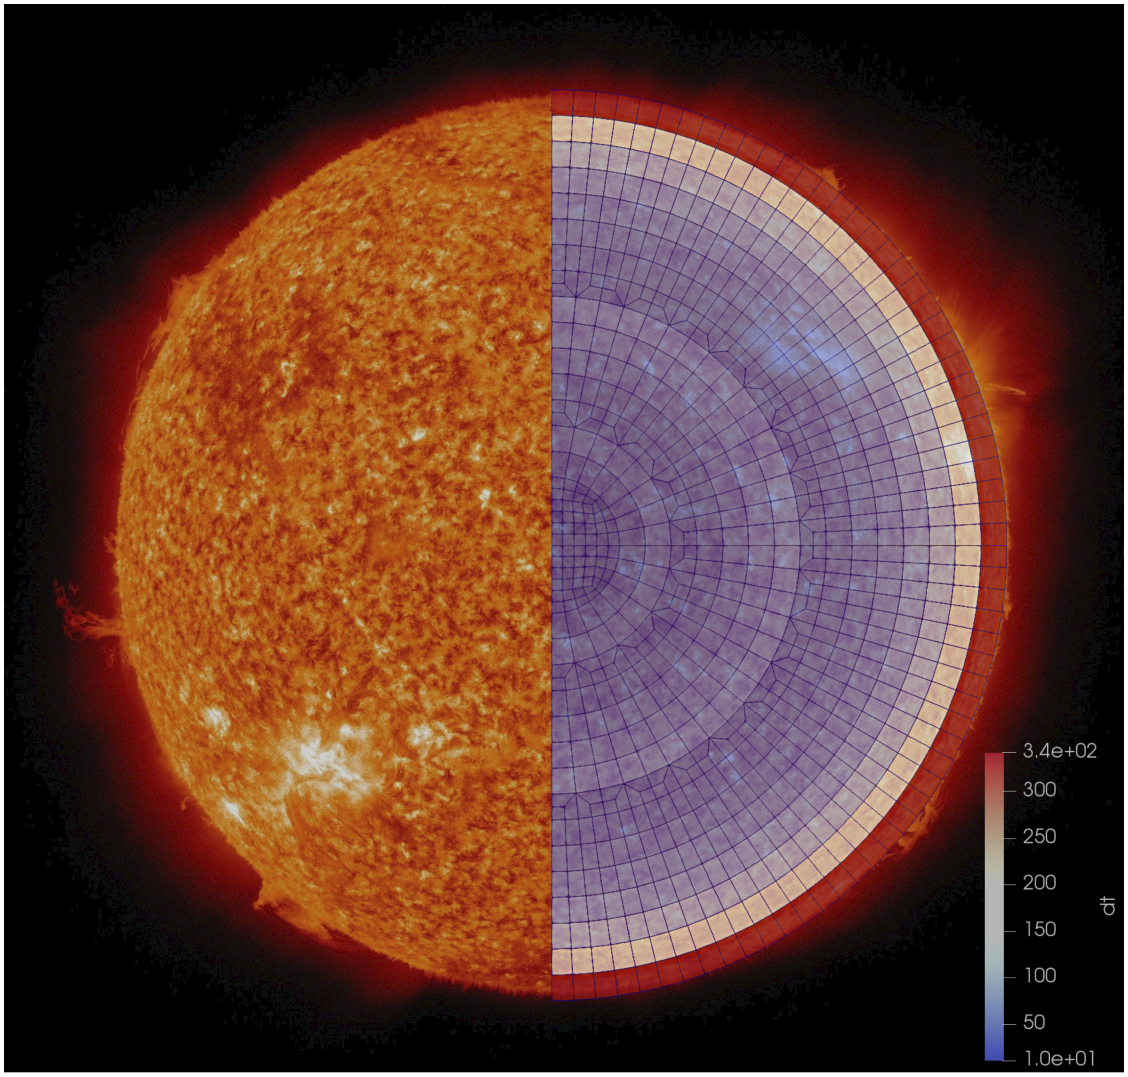

Figure 6.2: A 10000s resolution mesh for the Sun, from the surface to a radius of 695,700 km (i.e. ignoring the atmosphere). Two clear ersatz discontinuity layers are shown in this figure below which elements are larger and above which they are smaller; as is the transition from radially aligned elements to non-radially aligned ones in the core. Note that a resolution of 10,000 s is chosen in order to ensure that the near-surface elements remain visible. At this resolution, the seismic wavelength is on the order of 80,000 km, so the strong-near surface gradients in sound speed are not properly resolved and all elements are of comparable sizes. In a 200 s resolution mesh, where the seismic wavelength is closer to 1,600 km, the surface elements would be too small to be seen on a figure of this size. The time step colourbar for dt is in seconds.

At 200 s, as used below, the mesh consists of approximately 1.4 million elements. Note that this is well over ten times larger than for a whole-Earth bathymetric run at 10 s (of the type that we saw in Chapter 2), highlighting how fine global-scale meshes must be to conduct simulations of this type on the Sun. In the 200 s mesh, the minimum time step of 0.28 s is set in the core, at the point where the mesh transitions from radial to non-radial elements (the darkest shades of blue in Fig. 6.2).

In this case, a 200 s resolution run lasting 20,000 s comes at a cost of approximately 400 processor-hours. Although this is easily manageable on a supercomputer, the memory requirements are more challenging, given the enormous number of elements.

It should also be noted that in a radially symmetric model with a single pressure source, the order of the Fourier expansion in AxiSEM3D may be set to 0, i.e. using the monopole term only. If 3D structures are introduced, the marginal increase in cost will be larger than would be the case for a second-rank moment tensor on Earth, as the Fourier order would need to be increased from 0 to N_u rather than 2 to N_u.

6.6.2 Wave propagation

Fig. 6.3 shows the results of such a 200 s resolution simulation using a hard-coded Gaussian source-time function. The results are, to use a highly scientific term, 'interesting'.

On the one hand, the later arrival times with increasing distance follow the pattern which would be expected and indicate that the simulation is reasonably stable and propagating energy through the mesh in a causal manner. The receiver immediately above the source presents a waveform which appears reasonable given the choice of a gaussian source-time function, with two peaks in opposite directions consecutively - as would be expected from the passage of an acoustic wave past the receiver. Seismic phases are also observed - whilst these are not physically meaningful, they are expected given the simulation setup, and correspond to multiple bounces off of the surface before reaching the receiver.

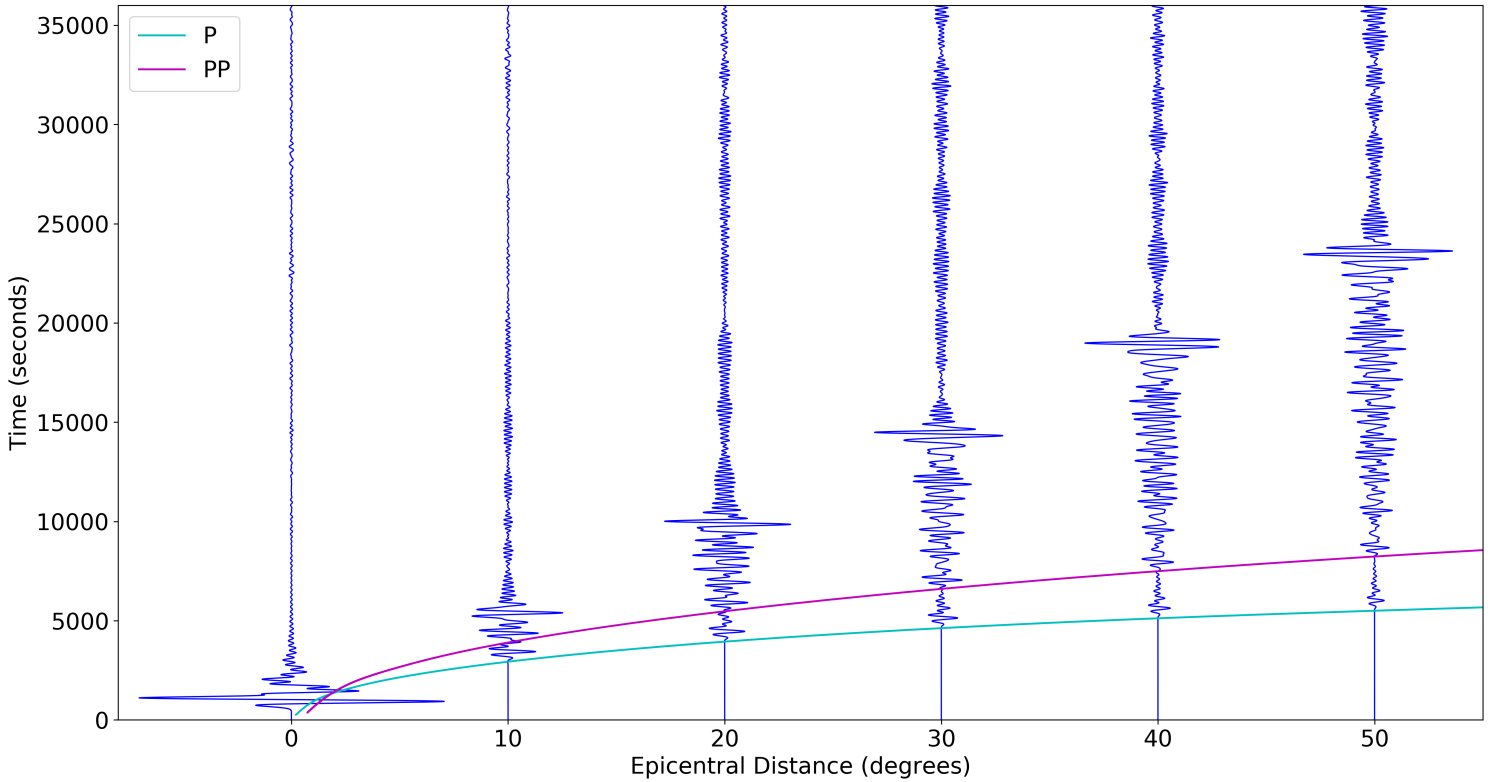

Figure 6.3: A 200 s resolution helioseismic run, bandpass filtered between 200 s and 1000 s. In this model the source was at a depth of 5,700 km and the receivers were located at the surface. Each vertical trace is linearly rescaled in amplitude with increasing distance from the source, to ensure that the plot is not dominated by the signal recorded at the receiver immediately above the source. Travel-time curves taken from TauP (using a Poisson-solid equivalent model, where $v_p = \sqrt{3} v_s$, as TauP requires a solid-domain source and surface). Whilst phases like P and PP are not observed on the Sun as discrete waves, acoustic waves experiencing a number of bounces off the surface between source and receiver are physically meaningful.

On the other hand, a terrestrial seismologist would identify surface-wave like phases as the largest amplitude in the seismograms. These features also show significant dispersion. Unlike the 'p-wave' (not be confused with p-mode) arrival, which decreases in amplitude with distance (due to spreading which is greater than $\frac{1}{r}$,) the surface-wave like phases maintain their amplitude when linearly rescaled by the epicentral distance, as would be expected. Whilst there is no hard velocity contrast in the Sun which is comparable to the Earth's crust, it is possible that these phases are P-waves in the limit of many reflections off of the surface between source and receiver.

Additionally, all traces appear to demonstrate significant noise - apparent at periods of several hundred seconds long after the peak in energy has passed. Again, it is unclear what the origin of this noise is - it does not appear to be completely random, but it is also not obvious what its computational origin is. It is present well above the dominant mesh period, suggesting that it is not simply a numerical artefact.

To actually use this method for modelling, further investigation would clearly be needed to either understand or 'tidy up' these simulations. However, in this case, where further development is not expected, we leave the reader with these mysteries, and move on to our final topic of discussion - asteroseismology.

6.7 Conclusions

Here, we have seen that AxiSEM3D is able to simulate acoustic wave propagation in the Sun. However, the realism with which these simulations are conducted is not ideal as the boundary condition is sub-optimal. However, this work does demonstrate that with careful mesh design, it is possible to undertake time-domain helioseismic simulations using AxiSEM3D. This is particularly difficult given that we have to deal with a ten order-of-magnitude variation in density between the core and the surface, which is a contrast approximately a billion times greater than that in the solid Earth.

Simulation results do, for the most part, look reasonable - though a thorough understanding and physical basis for each of the observed parts of the waveforms remains elusive.

If this line of work is pursued further, moving to a local or regional scale and introducing three-dimensional structure into the Sun (not currently supported in AxiSEM3D for fluid regions) would be an interesting extension.

At first sight it would seem that the deep interior of the Sun and stars is even less accessible to scientific investigation than any other region of the universe.

— Sir Arthur Stanley Eddington

7 Other Stars

Contents

7.1 Preface 178
7.2 A quick primer on stellar... 178
7.2.1 Life 179
7.2.2 Death 181
7.2.3 Classification 182
7.2.4 Measurements 183
7.2.5 Terminology 183
7.3 Summary 185
7.4 Introduction 186
7.4.1 Population characteristics 186
7.4.2 Potential formation channels 186
7.4.3 The horizontal branch? 187
7.4.4 Asteroseismology 188
7.4.5 Tycho B 188
7.5 Methodology 189
7.5.1 Models 190
7.5.2 Asteroseismology 194
7.6 Results 195
7.6.1 Mode frequencies 195
7.6.2 Mode separations 196
7.6.3 Temporal evolution 197
7.6.4 Sensitivity to degree of stripping 197
7.6.5 Rotation 198
7.7 Mode excitation 199
7.8 Applications 200
7.8.1 Ongoing observation campaigns 200
7.8.2 Tycho B 200
7.9 Conclusions 201

7.1 Preface

This material has not yet been published, but has been presented at two conferences: the Subdwarf OB Stars Meeting 9, in Hendaye, France; and the Dynamics of the Sun and Stars in Boulder, Colorado - both in 2019. The majority of the work contained herein was written as a manuscript which is planned to be submitted to Monthly Notices of the Royal Astronomical Society as:

Fernando, B., Justham, S., Antoci, V. and Podsiadlowski, P. - **Asteroseismic simulations of λ Boötis stars and potential application to Tycho B**.

The horizontal branch models used in this simulation are based upon those created by Stephen Justham. Suggestions regarding the content of this manuscript were also received from Selma de Mink and Victoria Antoci. I was responsible for running stellar evolution simulations, editing models, computing asteroseismic spectra, and interpreting them. Assistance in determining parallax distances to Tycho B using GAIA data was received from Payel Das, whilst evaluation of its potential as a potential target for TESS was aided by Ingrid Pelisoli.

This is still very much a work in progress - more remains to be done on the choice of evolutionary models, for example moving from a singular example to a more widely spaced grid in mass, especially for the horizontal branch.

7.2 A quick primer on stellar...

I hope the reader will not be insulted by my inclusion of a quick primer to the methodology discussed in this paper. Whilst Chapter 1 contained a short discussion of both helioseismology and asteroseismology, this was written in such a way as to be broadly applicable to both topics in the most general sense. Some of the specifics discussed in this paper, however, differ quite wildly from the methods used before, so I thought it might be worth outlining some of the background. This material is not included in the introduction to this paper as it would likely be entirely redundant

to the average reader from an astrophysics background, but might prove of use to a geophysicist. If not, please feel free to skip over it - if you have not already done so.

Of key use in this discussion will be the Hertzprung-Russell diagram, long beloved of astrophysicists. Fig. 7.2.3 shows one which illustrates the 'classical instability strip, where variables like the Cepheids that we have already mentioned reside.

7.2.1 Life

Stars generate their energy through the fusion of atomic nuclei, and the energy produced by this process supports them against gravitational collapse. Stars like our Sun are on the 'Main Sequence', meaning that they generate most of their energy from hydrogen fusion in their cores. What happens when the core hydrogen runs out very much depends on the mass of the star. Here, we're concerned mostly with stars around the mass of our Sun. These stars are reasonably well-behaved, by which I mean that they refrain from blowing themselves up when they run out of hydrogen, or doing other silly things like that.

When they have completed core hydrogen burning, Sun-like stars begin burning hydrogen in a shell around the core. Such stars are termed 'giants', and later in this thesis we will refer to this stage of their evolution as being 'on the giant branch'. The core itself consists at this point of helium, which in smaller stars becomes quantum degenerate matter. Eventually, the hydrogen in the shell begins to deplete and the reduction in radiation pressure threatens to take the star out of hydrostatic equilibrium. The star begins to contract, and the core heats up. Eventually, a 'helium flash' occurs, and helium burning begins, bringing the star back into hydrostatic equilibrium.

Stars which are burning helium in their cores are said to be on the 'horizontal branch'. Extreme examples of horizontal branch stars are the subwarfs, many of which likely form through binary interactions (i.e. in systems where more than one star is gravitationally bound together). These objects appear to have had most of their envelopes removed, and are small but hot.

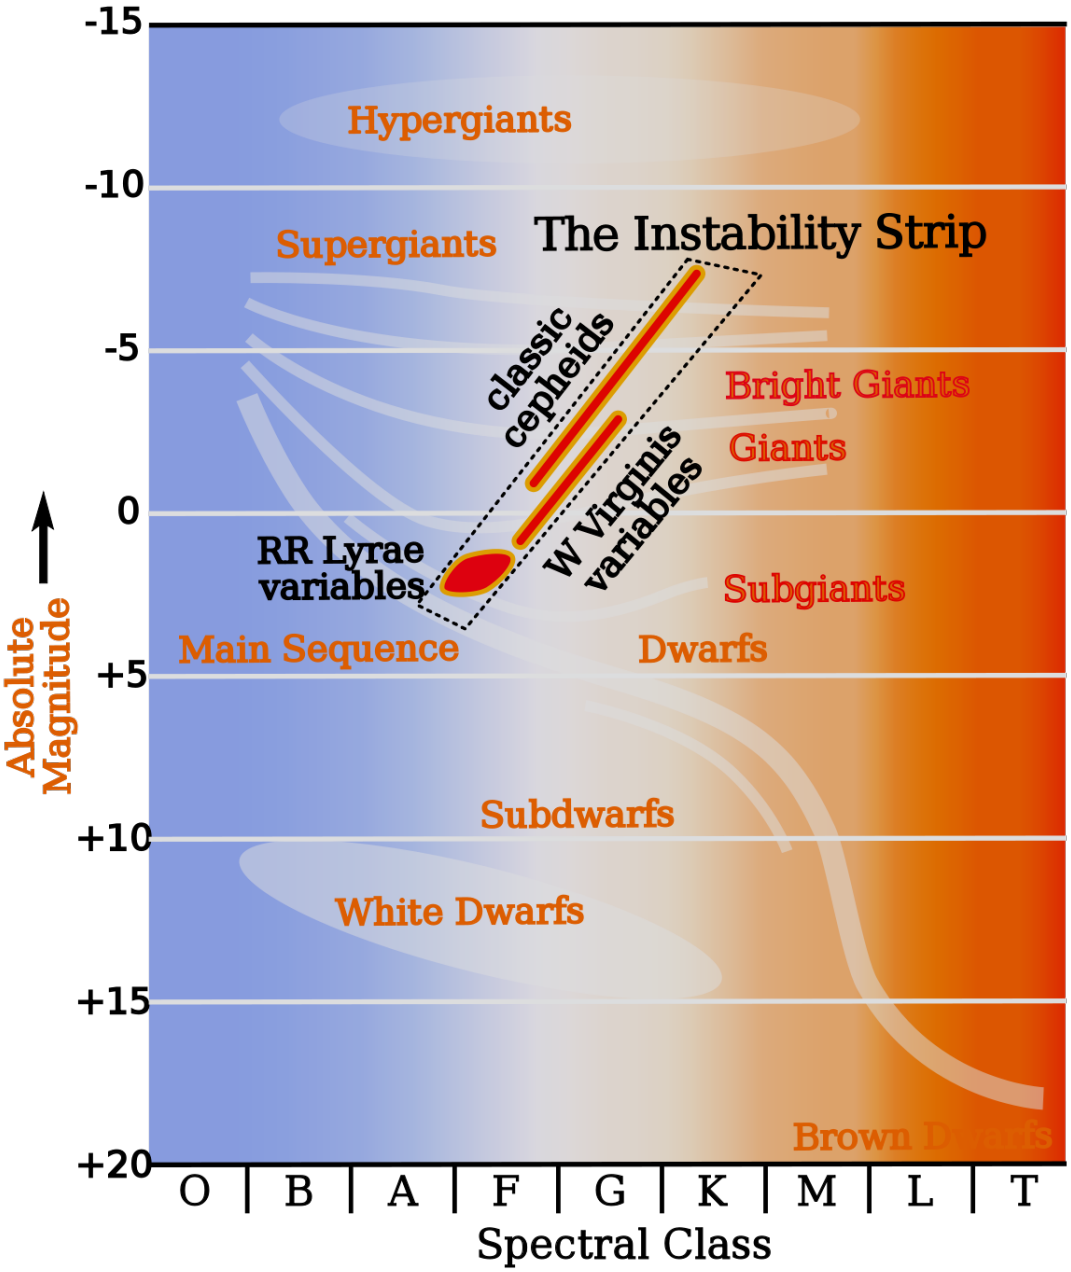

Figure 7.1: A Hertzprung-Russell diagram, plotting stars' temperatures (spectral class) against their luminosities (absolute magnitude). Note that hotter stars are toward the left and brighter stars are toward the top, and several natural families (white zones, referred to as Yerkes luminosity classes) are apparent. The vertical colour bands refer to the approximate colour the star would display if observed with the naked eye. The Sun is spectral type G and has an absolute magnitude of +4.8, putting it just around the 'D' of the 'Dwarfs'. Main sequence A-type stars are much brighter, at absolute magnitudes of around +2 (note that the magnitude scale is logarithmic). Horizontal branch stars are similar to subdwarfs, in that they are hot (spectral type A in our case), but dimmer than their main sequence counterparts. Diagram used under license CC BY-SA 3.0 from Wikimedia Commons.

7.2.2 Death

Whilst the stars that we deal with here are not prone to ending their lives in massive explosions (supernovae), the same is not necessarily true for their neighbours. The kinds of supernovae we are interested in here are 'Type Ia' events, involving white dwarf stars which exceed the Chandrasekhar Limit (about 1.4 times the mass of the Sun). They are distinct from more energetic Type II supernovae events, wherein giant stars run out of fuel and can no longer support themselves against gravitational collapse.

White dwarfs are stars which were not massive enough earlier in their lives to ignite the fusion of heavier elements ('metals') in their cores after helium was depleted. As a result, they shrug off their envelopes to leave behind Earth-sized remnants of electron-degenerate light metals, such as carbon and oxygen.

There are limits to how much mass may be supported through electron degeneracy pressure, however. If a white dwarf strays above this limit, it becomes unable to support itself against gravity and collapses. The resulting explosions are truly enormous - many billions of times brighter than the Sun, and are visible from the other side of the universe. Because all type Ia supernovae are of similar luminosity, they are used as 'standard candles' (or more properly, 'standardisable candles') to calibrate cosmic distance scales.

How white dwarfs end up going from below the Chandrasekhar Limit to above it remains a point of contention. Two competing theories exist - though it is entirely possible that both formation channels occur in practice.

The first is the 'single degenerate' scenario. Here, a white dwarf accretes matter gravitationally from a less evolved neighbour. The eventual explosion of the white dwarf may change the atmosphere of the 'donor', but is not expected to destroy it. The second mechanism involves the collision of two white dwarfs ('double degenerate') and likely destroys them both.

One way to show that the first channel occurs in practice is therefore to find one of these surviving donor stars, and to then demonstrate that it was in fact the donor rather than just an unlucky interloper. This is hard, however - only between two

and four (depending on who you believe) Type Ia supernovae have been observed in our galaxy since clear-record keeping began about two thousand years ago.

No donors have yet been identified - but in this paper we discuss whether asterosesimology might be able to shed some light on whether one particular star was a donor in Tycho's supernova, which was observed from all over the Earth in the constellation Cassiopea in the year 1572.

7.2.3 Classification

There are so many different families of stars that they make up a veritable zoological menagerie. A full description of all the different specimens found in our galaxy alone is well beyond the scope of this thesis, but a brief introduction might prove helpful.

Stars are generally classified according to their effective surface temperature and luminosity. For a given observed temperature and luminosity (and sometimes also surface gravitational field strength), it is possible to make a fair guess as to the internal composition of the star. Note that determining the mass of a star is hard - if it is in a binary some constrains on the mass ratio between the two stars can be derived from orbital dynamics, but for single stars in isolation, estimates of mass are usually derived from other quantities or modelling, rather than being directly measured.

Stars like our Sun are G-type, whereas the stars considered in this thesis are A-type (see Fig. for a comparison of luminosity and temperature). A-type main sequence stars are very-much the sensible, adult, members of the stellar zoo. They rotate fast, but do not have thick convective zones. Most (but not all) have weak magnetic fields, and hence weak stellar winds as well. This makes them reasonably easy to study remotely. Some A-type stars are found to be metal deficient in their outer layers, with elements heavier than sulphur very much depleted with respect to solar abundances. Such stars are termed 'λ Boötis stars', after their prototype - the 12th brightest star in the constellation of Boötes, the herdsman.

7.2.4 Measurements

There are many different ways of making measurements of stars - spectroscopically (measuring spectra), photometrically (measuring brightness or colour), and through asteroseismology. As we briefly discussed in the introduction, the most cutting-edge asteroseismic measurements are made today from satellites orbiting the Earth, as these may monitor many thousands of stars simultaneously, without the bother of an atmosphere, weather, or diurnal cycles. The most oft-used source of space-based asteroseismic measurements today is TESS, the Transiting Exoplanet Survey Satellite, launched in 2018.

7.2.5 Terminology

For the non-astrophysicist reader, a quick guide to terminology not detailed elsewhere may prove useful.

- **Cepheid**: the most famous class of variable star, which exhibits strong radial pulsations caused by time-dependent opacity variations in its outer layers.
- **Convective overshooting**: the continued motion of material driven by convection beyond the edge of the unstably stratified region, due to its momentum.
- **Core ignition**: the process through which helium in a star's core begins to fuse, following hydrogen exhaustion. Depending on the quantum state of the core helium, this may be either degenerate or non-degenerate.
- δ **Scuti star**: a class of young star which displays both radial and non-radial pulsations, which are opacity driven.
- **dex**: a power of 10, such that a change of +2 dex would be a change of 10^2.
- **Effective surface temperature**: the temperature that a body with the same luminosity as the star, but with an emmisivity of unity (i.e. a black body) would have.

- **Henyey track**: the evolutionary path followed by pre-main sequence stars roughly the mass of the Sun as they approach the main sequence.
- **Geometric cancellation**: the effect through which brightness changes on the surface of a star due to oscillations of high degree ($l > 3$) cancel out and become unresolvable from a distance.
- **Mass coordinate**: the fraction of a star's total mass contained interior to a particular point.
- **Mass fraction**: the fraction of a star's total mass made up by a particular chemical species.
- **Metal**: an element more massive than helium (i.e. lithium and anything beyond it on the periodic table!)
- **Node**: the point(s) along the star's radius at which the amplitude of a mode (a standing wave) is zero.
- **(Non-)adiabatic**: refers to whether or not oscillations exchange energy with the surrounding medium. Adiabatic oscillations, like adiabatic thermodynamic processes, are assumed not to exchange energy over time- and hence are significantly simpler to model.
- **Opacity**: the degree to which a material will absorb electromagnetic radiation. Increasing the opacity in a given region will increase its temperature, all other things being equal.
- **Pressure scale height**: the height range across which the hydrostatic pressure inside the star changes by a factor of e.
- **Stellar wind**: a tenuous, supersonic stream of charged particles ejected from a star's atmosphere. Mass loss through a stellar wind may be a significant factor in determining the evolutionary path of a star.

- **Surface gravity**: the acceleration experienced due to a star's mass, defined at an arbitrary 'surface'. The base-10 logarithm, log g, is often quoted for simplicity.

- **vsini**: the projection of a star's rotation velocity vector at the equator onto the line of sight to the observer.

7.3 Summary

- The Lambda Boötis class of stars are unusual objects which show chemical peculiarity in their outer layers. Exactly what drives their formation and evolution has been an open question in astrophysics for several decades.

- It is unclear at what stage in their evolution such stars are, with both Main Sequence (MS) and Horizontal Branch (HB) configurations suggested.

- Here, we evaluate whether asteroseimology may be used to distinguish between these possibilities.

- We find that although the mode spectra are sufficiently different to in principle enable discrimination through asteroseismology, a plausible excitation mechanism for either set of modes cannot be found.

- Thus, asteroseismic determinations of structure are likely to be extremely challenging, if not impossible.

- If excited modes could nonetheless be found in either type of star, the potential to use asteroseismology to study one particularly interesting λ Boötis star, Tycho B exists. This is considered on the grounds that it has been proposed as the potential donor-companion in the Type Ia supernova SN1572.

7.4 Introduction

7.4.1 Population characteristics

The λ Boötis stars are population I objects of spectral types late B to mid-F which show a characteristic chemical peculiarity in their outer layers. Deficiencies of up to -1.5 dex in elements from Sc to Fe, as well as in Mg, Si, Ca, Zn and Ba are observed in their spectra, with a population that is estimated to include no more than a few percent of the stars in the relevant spectral range (Paunzen 2004; Solano et al. 2001). Some canonical A-type stars, including Vega (α Lyrae) are found to display λ Boötis behaviour (Baschek and Slettebak 1988), providing strong motivation for further investigation into this phenomenon.

Beyond their metal-poor surface abundances, these stars appear unusual only in that they display broad hydrogen lines. Their rotation rates, although high, are compatible with those observed in other objects of a similar spectral type. The ages of the λ Boötis are also not well-constrained, with many of the models for their chemical peculiarity predicting specific stellar ages; though the absence of such stars in intermediate-age open clusters and their presence in OB1 associations suggests that they are relatively young objects (Gray and Corbally (1993), Gray and Corbally (2002)).

7.4.2 Potential formation channels

A number of hypotheses have been proposed which may explain the origin of the λ Boötis phenomenon, and of course multiple methods may operate simultaneously (as suggested by Murphy and Paunzen (2017)). These can be divided into two main categories: diffusion/mass loss and selective re-accretion. Michaud and Charland (1986) suggest that a combination of diffusion and mass loss can explain the observed underabundances if the stars are approaching the end of their main-sequence lifetimes. However, such a mechanism may be suppressed in the presence of even moderate rotation (Turcotte and Charbonneau 1993), and does not appear compatible with more recent observations suggesting an evolutionary position closer

to zero-age main sequence. Furthermore, the weak stellar winds from A-type stars do not favour a high-mass loss scenario.

Alternatively, Venn and Lambert (1990) demonstrate that a similar mechanism to that which depletes the gas in the interstellar medium of metals may produce the observed abundance pattern, wherein separation of refractory elements occurs in circumstellar material. Following this, the metal-poor gas is then accreted whilst the metal-rich dust grains remain in space surrounding the star. The presence of such a disc around the star would be compatible with a relatively young stellar system, in line with observational measurements; though the presence of a debris disc in a more evolved system is of course also possible.

Andrievsky (1997) propose that the λ Boötis stars are the end-results of mergers of W UMa-type contact binaries, in which the circumstellar disc is in fact the remnant of a circumbinary disc, with the metal-poor gas being re-accreted onto the merger product. However, as noted by Paunzen et al. (2003), not all λ Boötis stars show the infrared excesses which would be expected from a surrounding disc.

Jura (2015) propose accretion from the winds of hot Jupiters as an alternate mechanism which would not require a full circumstellar disc, though their remains significant observational uncertainty regarding the fraction of A-type stars which are host to such planets.

In summary, there is no universally accepted mechanism through which the λ Boötis phenomenon arises in main-sequence or pre-main sequence stars, though most theories suggest that the chemical peculiarity should be restricted to the stars' near-surface layers.

7.4.3 The horizontal branch?

Indeed, it is also possible that some stars of the λ Boötis type are not main sequence (MS) or near-MS objects at all, but rather are evolved horizontal branch (HB) stars in the core helium burning stage. Such stars would be subdwarf-like, helium-rich and with only a thin hydrogen envelope remaining, and may have formed as a consequence of binary mass transfer with the HB star as donor (Han and Podsiadlowski 2004).

Differentiating between these internal configurations on the base of brightness/distance and surface gravity measurements alone is difficult, particularly due to the objects' measured rapid rotation, chemical peculiarity and uncertainty in the degree of extinction associated with the possible presence of circumstellar dust shells.

7.4.4 Asteroseismology

Asteroseismic measurements offer the opportunity to overcome some of these challenges and to probe the λ Boötis stars' internal structures. Here, we present asteroseismic modelling of both HB and MS stars, highlighting the differences between the predicted seismic signatures which may be used as a diagnostic to differentiate between them. Similar techniques for differentiating between core and shell-helium burning subdwarf stars have been explored by Guo (2018).

Paunzen et al. (1998a) and Breger et al. (2006) suggest that a few λ Boötis stars exhibit δ Scuti-type behaviour. These consist of both radial and non-radial oscillations, as might be expected from their position within the classical instability strip. Bohlender et al. (1999) infer that non-radial pulsations are common in λ Boötis stars.

Paunzen et al. (1998b) and Casas et al. (2009) have attempted to use asteroseismology to date the stars, provisionally suggesting that the stars are relatively far advanced from the zero-age MS. More recent studies which take advantage of the increases in size of the λ Boötis catalogue in the last decade, are somewhat lacking.

7.4.5 Tycho B

A particular λ Boötis star of interest is Tycho B, which has an angular position close to the geometric centre of Tycho's supernova remnant (SN1572). SN1572 is one of very few (somewhere between 2 and 4) Type Ia supernova in the Milky Way observed in the last 2,000 years.

With a surface temperature of $\simeq$ 10,000 K, log g $\simeq$ 4.0 and a significant metal depletion of [Fe/H] $\simeq$ -1, Tycho B has been proposed a possible mass-donating companion to the white dwarf whose explosion was observed in the year 1572

(Kerzendorf et al. 2013). Recent data from Gaia DR2 suggest a distance for Tycho B of $2.03^{+0.19}_{-0.15}$kpc, which is compatible with the estimated distance to the remnant (Ruiz-Lapuente et al. 2018). It should of course be noted that determining the distance to the remnant itself, a diffuse feature, is extremely challenging.

Such investigations are of value as if Tycho B can be confirmed as a surviving companion to SN1572, this would place a first observational constraint on the proportion of Type Ia supernovae which occur through the single-degenerate channel (WD + MS/HB/Giant Branch companion) as compared to the proportion which occur through the double-degenerate (WD + WD) channel. Such a debate is ongoing in stellar physics and is of importance given the role of these explosions as 'standard candles' (Perlmutter et al. 1999; Riess et al. 1998) and sources heavy-element nucleosynthesis (e.g. Tsujimoto et al. (1995)).

Tycho B is not the only star to have been proposed as potential companion in the single-degenerate scenario. Ruiz-Lapuente et al. (2004) suggest instead Tycho G, though this hypothesis has been called into question more recently on account of Tycho G's slow rotation and position within the remnant (Kerzendorf et al. 2009). Ihara et al. (2007) propose Tycho E, though Kerzendorf et al. (2013) suggest that it is too far from the remnant to be a plausible candidate. Regardless, the quest for a potential companion to Tycho's SNR remains an exciting search.

7.5 Methodology

The stellar models in use were created using the `MESA` software package, version 10398 (Paxton et al. 2015); asteroseismic modelling was carried out using the `GYRE` software package, v5.2 (Townsend and Teitler 2013).

In this section, we describe the creation of the MS and HB models, and detail the procedure used to calculate the oscillation frequencies of the stars. In both cases, the outer 2% of the star by mass was stripped of all elements heavier than ^{32}S to replicate the λ Boötis phenomenon, reflecting that the chemical peculiarity is likely confined to the stars' outer layers.

We note that the chemical peculiarity we use likely extends much deeper than would be expected from the accretion or diffusion scenarios, and consider the sensitivity of our results to the degree of stripping in the Sec. 7.6.2. The abundances of elements lighter than ^{36}Ar were then renormalised to conserve the total mass of the star.

Step convective overshooting above the burning core only is used. The mixing coefficient for the overshoot is sampled 0.1 pressure scale heights deeper than the convective boundary, whilst the overshooting extends 0.1 pressure scale heights beyond the boundary.

7.5.1 Models

The MS model used had a mass of $2.3M_S$ and an initial solar metallicity of Z = 0.02. The stripping of the heavy metals from the stars' outer layer was conducted during the pre-main sequence stage, whilst the star was on the Henyey Track.

The HB model had a mass of $0.4M_S$, and was produced by stripping the envelope from a $2.6M_S$ star with initial Z = 0.02 and Y = 0.28. Type II opacities are used for better approximation of the stellar opacity following CO enhancement, with the initial base metallicity set to 0.02. The chemical peculiarity was introduced when the star reached its maximum radius on the giant branch, and helium ignition in this model occured in a non-degenerate core.

Following the stripping of the heavy, highly refractory metals from the stars' outer layers, both were allowed to evolve to T_{eff} = 10,000 K. At this point the MS model has a luminosity of $28L_S$ and log g = 4.3 whilst the HB model has a luminosity of $60L_S$ and log g = 3.2. As noted previously log g values for these types of stars can be hard to determine observationally, and indeed candidate members of the λ Bootis class have been proposed which span the entirety of this surface gravity range (e.g. Gray et al. (2017)). The difference between the log g and L modelled for the HB versus the MS star are not, therefore, at odds with observation.

The modifications to the composition are shown in Fig. 7.2, with their evolutionary tracks post-stripping shown in Fig. 7.3.

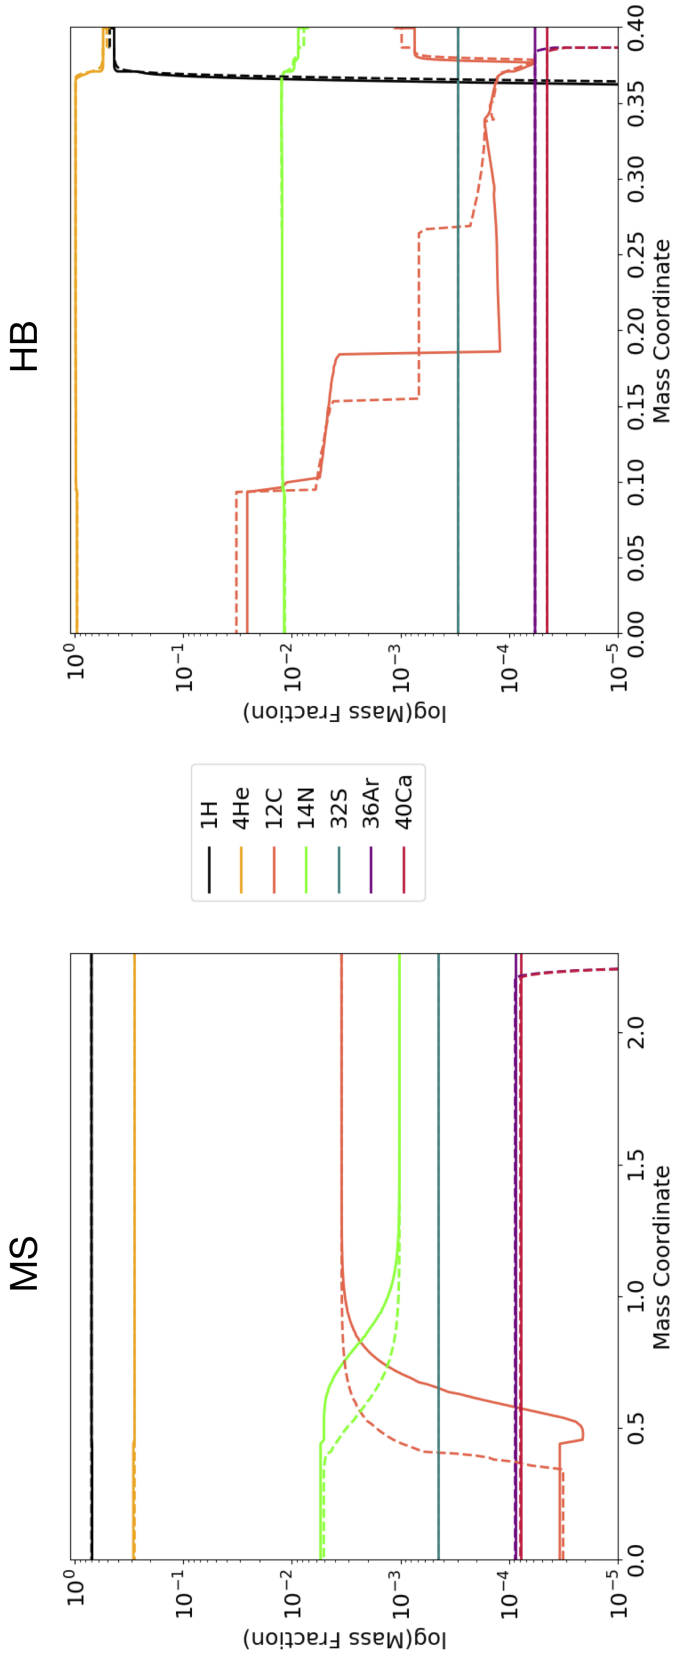

Figure 7.2: Abundance profiles for the MS star (left panel) and HB star (right panel). The solid lines show the composition of the star at the point where the asteroseismic measurements are made without stripping, the dashed lines show compositions at the same evolutionary position, but with stripping.

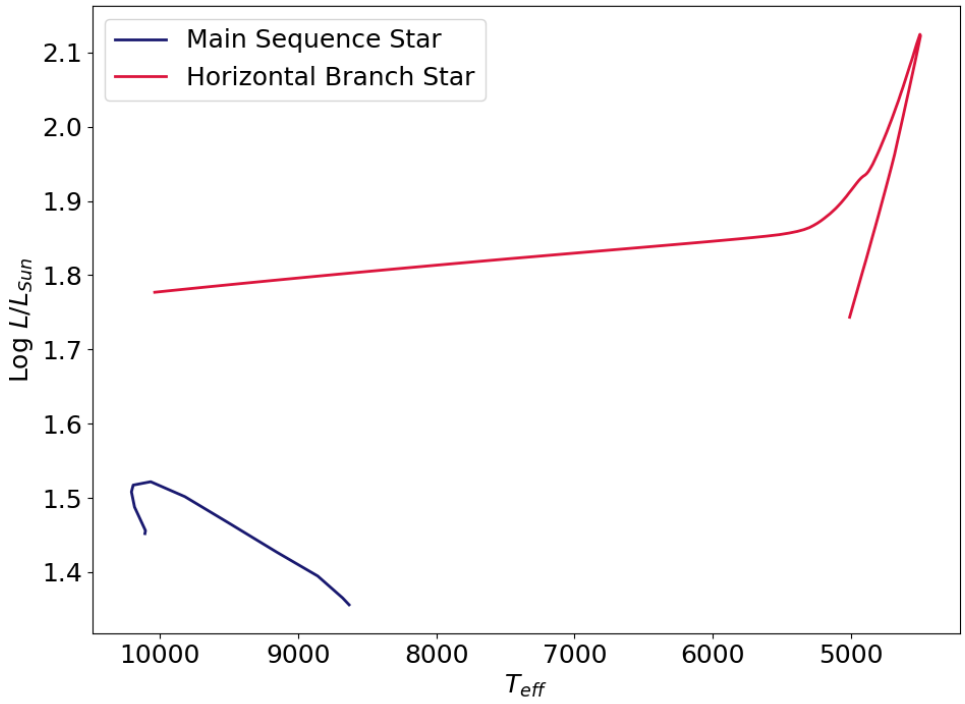

Figure 7.3: HR Evolutionary tracks for the MS and HB stars, both ending at an effective temperature of 10,000 K. The MS model shows the pre-main-sequence track as the fraction of the star's luminosity generated by hydrogen fusion in the core increases, whilst the HB model is plotted as it ascends the giant branch during hydrogen shell burning, through the tip of the giant branch, and onto the horizontal branch following non-degenerate helium core ignition.

The removal of metals from the outer layers of the HB star reduces its opacity, as can be seen in Fig. 7.4, thus increasing its luminosity by approximately 15%.

In the case of the MS star, the removal of the metals in isolation naturally reduces its opacity, but the renormalisaton of the remaining abundances increases the hydrogen mass fraction enough to compensate for this, slightly increasing the star's overall opacity and hence decreasing its luminosity by approximately 2%. The difference between the MS and HB cases come about because the MS case abundance renormalisation mostly adds hydrogen, whilst in the HB case it it mostly adds helium, with a smaller associated increase in opacity.

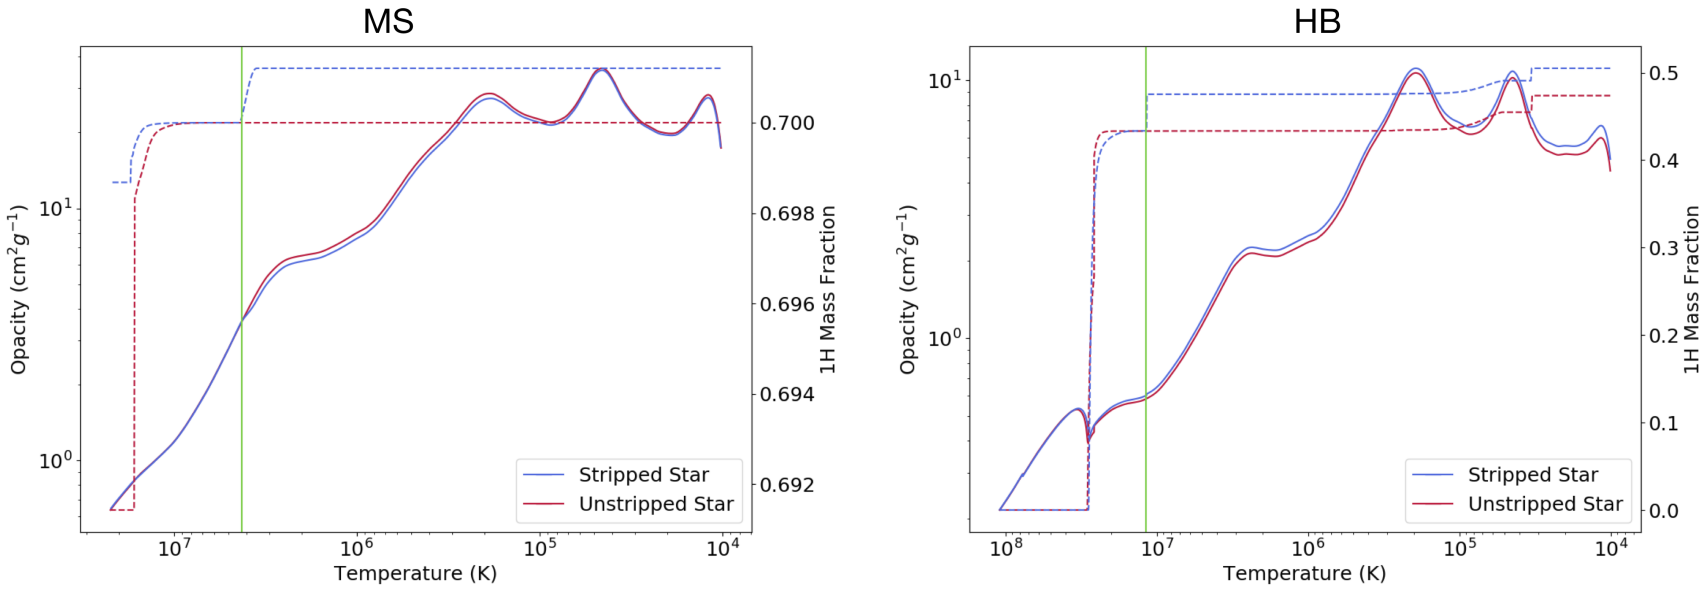

Figure 7.4: Opacities (left axis, solid lines) and hydrogen mass fractions (right axis, dashed lines) of the MS (left panel) and HB (right panel) stars as a function of temperature. The vertical green line indicates the lower limit of the metal stripping. The stars' surfaces, at a temperature of 10,000 K, are at extreme right. Note that the temperature gradient through the star is highly non-linear, hence 98% of the mass is contained to the left of the green line.

7.5.2 Asteroseismology

The adiabatic oscillation frequencies of the stars were calculated using `GYRE`. A vacuum outer boundary condition is set, with the oscillation frequencies of both radial ($l = 0$) and non-radial ($l \neq 0$) modes found. A second order Gauss-Legendre collocation scheme was chosen to solve the oscillation equations, on a linear grid for pressure (p)-modes and p-like mixed modes, and an inverse grid for gravity (g) and g-like mixed modes. 10 points per wavelength are used in the propagation regions, and 2 points per wavelength in the evanescent regions. Modes are computed for the range $l = 0$ to $l = 3$, as modes with $l > 3$ are expected to be generally unobservable due to geometric cancellation (Aerts et al. 2010).

Non-adiabatic oscillation frequencies were computed for the MS star, and found to be very close to the adiabatic frequencies. This is likely due to the star's structural profile being only super-adiabatic in the outermost regions. For the HB star, however, the `GYRE` software was unable to converge to the non-adiabatic frequencies for many of the modes. This is not uncommon for complex structural models, and is a result of the Magnus Multiple Scheme used in finding solutions to the oscillation equations failing to converge in the complex frequency plane.

For this reason, we restrict discussion to adiabatic calculations for both stars only; though noting that for those modes for which the non-adiabatic calculations could be performed, the imaginary parts of the eigenfrequencies were negative for both stars, indicating pulsational damping of these modes. We will discuss this in terms of possible excitation mechanisms later.

The radial order of the modes is defined as n_{pg} is defined as $n_{pg} = n_p - n_g$ where n_p corresponds to the number of acoustic nodes in the eigenmode and n_g the number of gravity nodes. Thus, purely modes in which pressure is the only restoring force (p-modes) have $n_g = 0$ and hence $n_{pg} > 0$, whilst those where gravity is the only restoring force (g-modes) have $n_p = 0$ and hence $n_{pg} < 0$. g-modes may only exist for $l > 0$, and are restricted to propagating in stably-stratified (i.e. non-convecting) regions. Thus, g-modes cannot propagate in the cores of either these HB or MS stars.

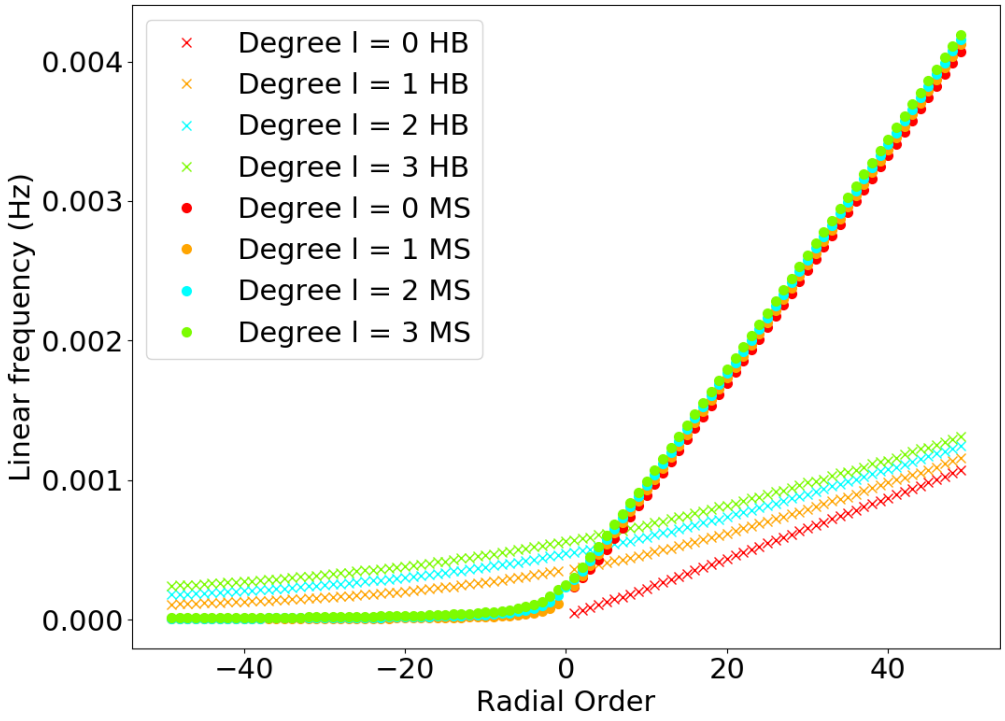

Figure 7.5: Modelled frequencies for the computed MS (circles) and HB (crosses) stars. The HB star g-like mixed modes are higher in frequency than the MS star pure g-modes, whilst the MS star pure p-modes are higher in frequency than the HB star's p-like modes.

For $l > 1$, a single intermediate mode, the fundamental (f) mode may also exist. Mixed modes, where the restoring force is pressure in some parts of the star and gravity in others, may also exist for $l > 0$. The sign of n_{pg} for these modes is determined by whether the mode has more nodes in pressure-restored regions ($n_{pg} > 0$) or gravity-restored regions ($n_{pg} < 0$).

7.6 Results

7.6.1 Mode frequencies

A comparison between the computed mode frequencies for both stars is shown in Fig. 7.5. Clear differences between the stars are apparent. The MS star shows a clear transition from p-modes at $n_{pg}>0$ which eventually become evenly spaced in frequency to g-modes at $n_{pg} < 0$ which become evenly spaced in period (e.g. for

p-modes see Fig. 7.6), with a single intermediate f-mode for $l > 1$. For the HB stars all modes with $l > 0$ are mixed. Thus, excluding the region around $n_{pg} \approx 5$, where the branches intersect, the computed mode frequencies are a potential method for discriminating between the two possible configurations, assuming of course that the relevant modes are excited (more on this later).

7.6.2 Mode separations

It can also be seen from Fig. 7.5 that at high positive radial orders the spacings (large separations) between successive p-modes in both the MS and HB stars approach a constant value, which is the characteristic asymptotic spacing $\Delta\nu_0$, given by $\Delta\nu_0 =< \nu_{n+1,l} - \nu_{n,l} >$ (Tassoul 1980). Plots of the successive large separations in terms of frequency are shown in Fig. 7.6. A similar characteristic separation in period can be found for the g-modes at high negative radial orders, however the value of the spacing is dependent on the mode degree.

The large spacings in frequency of the two stars are also substantially different, except in the region $n_{pg} \approx -5$, meaning that these too can be used as a differentiator between the MS and HB configurations.

It should be noted that the region where MS and HB stars show similar frequencies ($n_{pg} \approx 5$) does not overlap with the region where they show similar frequency spacings ($n_{pg} \approx$-5) meaning that in either case at least one of these should be a suitable diagnostic for differentiating between the structures if all the modes are observed. The convergence to $\Delta\nu_0$ with increasing n_{pg} is much faster in the case of the MS star than the HB star, and indeed the asymptotic separations are also substantially different. For the HB star $\Delta\nu_0$= 21.8 μHz, whilst for the MS star it is nearly four times larger, at $\Delta\nu_0$= 82.6 μHz.

We also note that the spectrum of the HB star appears to show significantly more avoided crossings, where the frequency spacing jumps substantially away from $\Delta\nu_0$, than in the case of the MS star. For a star of this type, this is expected to be a genuine asteroseismic signature, rather than a numerical artefact. This may also be a route to distinguishing between them on the basis of their oscillation signatures.

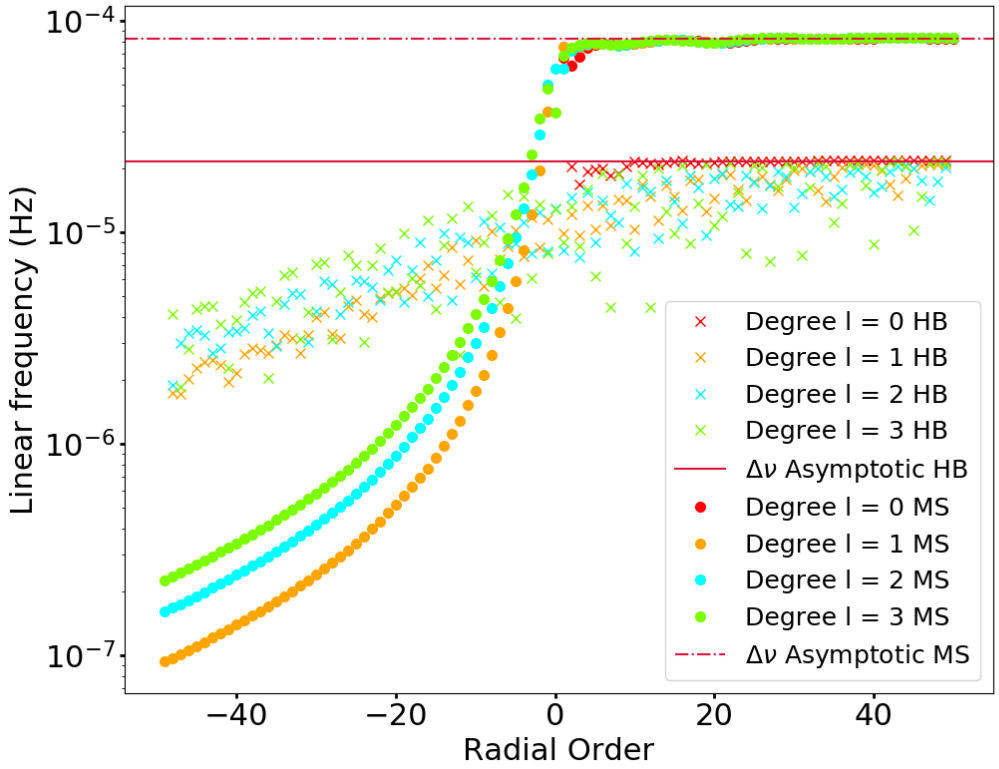

Figure 7.6: Mode separations in frequency. Note that the high-order p-modes are asymptotically spaced in frequency at large positive values of n_{pg}.

7.6.3 Temporal evolution

After Silvotti et al. 2007, we also consider whether the change in mode frequencies over time ($\frac{\mathrm{d}\Delta\nu_0}{\mathrm{dt}}$) an be used to differentiate between the MS and HB stars. For the MS star, we find the asymptotic separation $\Delta\nu_0$ slowly decreases at this point in its evolution, whereas for the HB star it is slowly increasing. Both these changes are smaller than 10^{-4}% per year, however if the measured pulsation frequencies are extremely stable over time (as might be expected for the HB star, given that it is the case for pulsating sdB stars), then they may be detectable.

7.6.4 Sensitivity to degree of stripping

In order to evaluate the sensitivity of our results to the degree of metal depletion in the star's outer layers (an *a priori* unknown), we also compute the asymptotic

separations for similar models which are not chemically peculiar. For the 'unstripped' (non-metal deficient) HB star, $\Delta\nu_0 = 82.2$ μHz, whilst for the unstripped MS star $\Delta\nu_0 = 19.8\mu$Hz. The decrease in asymptotic spacing is expected, with the removal of metals reducing the star's mean density and hence also reducing $\Delta\nu_0$. The fact that the change in $\Delta\nu_0$ is small compared to the difference between the MS and HB asymptotic frequencies indicates that uncertainty about the penetrating depth of the chemical peculiarity should not be an impediment to differentiating between a hydrogen-core burning and helium-core burning structure on the basis of measurements of $\Delta\nu_0$.

7.6.5 Rotation

Whilst measurements of $v\sin i$ for λ Boötis stars are relatively large (upwards of $\simeq 150\text{kms}^{-1}$), if their chemical peculiarity comes about through accretion from a metal-depleted disc, such rapid rotation is likely confined to the objects' outer layers, as stars in this mass range are not expected to be in solid-body rotation (Suárez et al. 2006).

Thus, if the stars' interiors are rotating at much less than the breakup velocity, because the rapid rotation is a near-surface phenomenon, the main effect of rotation in terms of asteroseismology is expected to be the splitting the mode peaks into multiplets. Here, the lth degree mode is split into $l(l+1)$ peaks, increasing in frequency from $-m$ to $+m$, with the $m = 0$ mode unchanged in frequency. Such splitting makes the pulsational spectrum of the star more dense, hence making observational mode identification even more challenging, but should not preclude differentiation between the HB and MS models so long as the main peaks within the spectrum are still identifiable. Indeed, it could allow for measurements of the stars' interior rotation profiles, which will be useful in constraining their evolutionary history.

On the other hand, if the entire star is rotating this fast (and hence at a substantial fraction of its breakup velocity), the deformation of the star from an equilibrium spherical structure is likely to be more significant than the frequency

splitting (Aerts et al. (2010), Suárez et al. (2006)). In this case, the mode structures become significantly more complex, and the frequencies of the m = 0 modes may also be modified. Such behaviour cannot currently be modelled with the `GYRE` software package, whilst additionally the influence of rotation-generated circulation flows may play a significant evolutionary role in this case. Further consideration of the effects of rotation is nonetheless required, though it is beyond the scope of this work.

7.7 Mode excitation

We have shown that there is a clear route to using asteroseismology to differentiate between possible structures of the λ Boötis stars, on the basis of 1) the computed mode frequencies, 2) mode separations and temporal evolution therein, and 3) the number of avoided crossings. However, this assumes that all the relevant modes in both stars are excited, as unexcited modes will not be detectable through observation.

For the MS star, the imaginary part of the complex mode eigenfrequency is negative, indicating that the modes are not pulsationally excited through an opacity-valve type mechanism. In the HB star, GYRE is unable to compute the complex part of the eigenfrequency, as this requires consideration of non-adiabatic effects and the solver cannot converge in this case (this is a common issue when dealing with complex stellar structures).

Of course, showing that the relevant modes would be excited in both stars is not necessary: if only one star is expected to show pulsation behaviour, and such pulsations are observed, this is an equally good indicator of the λ Boötis star's internal configuration. Given that it is not possible to determine from modelling whether or not modes are excited in the case of the HB star, we considered whether considering observations other stars of similar types, such as the δ Scuti variables, might provide some insight.

The δ Scuti stars have opacity valves regulated by the presence of ionised helium in their outer layers. However, such excitations are normally only found in stars with significantly lower effective temperatures, with a maximum of around 8,800 K (Uytterhoeven et al. 2011). As the relevant excitation mechanism is strongly

influenced by the thermal ionisation of helium, it is not expected to operate at higher temperatures than this within the range of log g of interest to us.

Similarly, γ Doradus type behaviour is not found in stars with temperatures approaching 10,000 K; and solar-type convective turbulence driven oscillations are not found in A-type stars.

However, recent research has found evidence of δ Scuti/γ Doradus pulsations at much higher temperatures, with an as-yet unexplained mechanism (Balona et al. 2016). Invoking an as-yet unexplained mechanism to justify that pulsations are excited in the HB star would of course be a weak argument - so in the absence of further evidence for excited modes from a modelling point of view, it would seem logical to consider the potential to turn to use observations in addition to simply modelling.

7.8 Applications

7.8.1 Ongoing observation campaigns

Campaigns to make asteroseismic measurements of a number of λ Boötis type stars are ongoing, for example Handler et al. (2020) recently published observations of pulsations in a close λ Boötis binary system. Endl (2018) are currently analysing TESS data as part of the search for planets around these stars, though the data produced will be of asteroseismic use as well.

As the fundamental periods of both the HB and MS stars considered are in excess of 1000s, and these modes are those which are most visible (as expected from δ Scuti like stars), then low-cadence TESS observations should be sufficient to resolve any oscillations λ Boötis stars .

7.8.2 Tycho B

We now consider application of these techniques to the previously mentioned λ Boötis star Tycho B.

The λ Boötis star Tycho B is located $2.03^{+0.19}_{-0.15}$ kpc away (Ruiz-Lapuente et al. 2018), near the geometric centre of the remnant of SN 1572. Its relatively rapid rotation, metal-deficiency and extinction from a potentially significant circumstellar disc mean that determination of its surface gravity and spectral type are difficult, with estimates ranging from B9 (Kerzendorf et al. 2013) to A8 (Ruiz-Lapuente et al. 2004).

This range of spectral types (and associated uncertainty in the degree of extinction) yields a spread in possible luminosities which covers the predicted luminosities of both our MS ($28\mathrm{L}_S$) and HB ($60\mathrm{L}_S$) models. Therefore, photometric measurements alone are rather insufficient for developing a comprehensive understanding of Tycho B's internal structure and past history, but asteroseismology offers us a promising route to doing so under the assumption that either one or both (HB or MS) of the configurations has at least some excited modes.

Unfortunately, Tycho B's proximity to the much brighter Tycho A means that it cannot be resolved on the TESS detectors. Its low apparent magnitude also means that extracting any pulsation signal from that amalgamated with the light curves of other stars from the relevant pixel would be extremely challenging. If a campaign were to be mounted to make asteroseismic measurements of Tycho B, a ground-based telescope would offer a more promising route to mode detection and identification.

7.9 Conclusions

We have computed oscillation frequencies using the `GYRE` package (Townsend and Teitler 2013), for two possible structures of the metal-depleted λ Boötis stars, using evolutionary models developed in `MESA` (Paxton et al. 2015). Asteroseismology offers a complementary technique through which some light may be shed on the internal structures of these stars.

We have shown that from 1) the calculated mode frequencies and 2) their separations and temporal evolution therein, and 3) the number of avoided crossings, it should be possible to differentiate between helium-core burning (horizontal branch) and hydrogen-core burning (main sequence stars) provided that at least one of these configurations has some excited modes.

Such measurements would hence be useful for constraining the age and properties of the λ Boötis stars, and may shed some light on the process through which their outer layers have become metal-depleted.

We have also considered the application of this work to the λ Boötis star Tycho B, which has been proposed as the surviving companion of the white dwarf which exploded in SN 1572. Asteroseismology of Tycho B would be able to offer significant insight into its evolutionary history.

If it can be confirmed as the surviving companion, this would be a substantial step forward in constraining the proportions of Type Ia supernovae which occur through the single-degenerate channel. Unfortunately, it is too dim and too close to another bright star to be observed by TESS, but may be resolvable through larger ground-based telescopes.

Men at some time are masters of their fates; the fault, dear Brutus, is not in our stars, but in ourselves

— Julius Caesar, Act 1, Scene II - William Shakespeare

8 Conclusions

Contents

8.1 Concluding introduction 203
8.2 Summary 204
8.2.1 Chapter 2: The oceans 204
8.2.2 Chapter 3: Mars, Part I 205
8.2.3 Chapter 4, Mars, Part II 205
8.2.4 Chapter 5, Mars, Part III 206
8.2.5 Chapter 6: The Sun 206
8.2.6 Chapter 7: Other stars 207
8.2.7 A few last notes... 208
8.3 Concluding conclusions/closing remarks 210

8.1 Concluding introduction

I'm not entirely sure what to say here - thanks for making it to the end? I hope that this has been at the very least a vaguely entertaining journey through some topics in comparative seismology. Given that each constituent paper has its own conclusion, I will attempt to keep summary and concluding remarks to a minimum.

8.2 Summary

Hopefully, this thesis has illustrated that developing new seismological models, especially those which go beyond the standard 'solid Earth' domain, can be enormously valuable. These extensions enable us to explore different kinds of seismic sources and processes (for example those coupled to the hydrosphere or atmosphere), but are not without their own challenges - for example, in the implementation of boundary conditions and increased computational costs.

8.2.1 Chapter 2: The oceans

In Chapter 2 I presented a modification to Oxford's in-house seismology modelling code AxiSEM3D, which supports the inclusion of an actual fluid ocean on a global scale in three dimensions. This was hitherto impossible using any open-source code. Using this new functionality, we saw that the effects of the oceans become noticeable in seismic waveforms at longer periods ($\sim$20 s) than was previously believed, and that the effects of bathymetry are quite noticeable on surface waveforms at periods of 10 s or less. By comparing observations to data from the April 2015 Papua New Guinea earthquake, we saw that these effects were not just model artefacts, but were borne out in data - or in other words, that synthetics with a fluid ocean and bathymetry included provided a better fit to data than synthetics without these features.

The routes to expanding upon this are numerous - the epilogue presented initial efforts to extend this formulation to simulate ocean microseisms excited by deep-water hurricanes, both for the purpose of better understanding their generation but also enhancing their use as a seismic probe and assisting in noise removal from other datasets of interest. Beyond this, at higher frequencies the potential exists to simulate t-phases, or to consider the effects of bathymetry on wave propagation from nuclear tests, which is very much an active topic of forensic seismology research in today's uncertain world.

8.2.2 Chapter 3: Mars, Part I

In Chapter 3, we reviewed the Oxford Seismology group's performance in the InSight Blind Test, which involved developing methods to detect, identify, and analyse signals of interest within one martian year's worth of synthetic data produced prior to InSight's landing in 2018. Many of the limitations of single-station seismology were apparent, but it was also clear that through clever use of surface wave dispersion analysis, differential travel-time measurements, and particle motion decomposition it was possible to constrain the properties of the largest marsquakes, though not with anything close to the ease or precision which would be possible on Earth. We also saw that the Oxford team's performance was rather good - of the 84 teams who registered, our catalogue was in the top 6 for completeness. Although we missed most of the events, so did most other teams!

The lessons learnt from the blind test and our participation in it informed development of analysis methods for actual seismic data from Mars - one particular topic where further work was clearly needed was in the identification of seismic signals from impacts, which just so happened to be the topic of the next chapter.

8.2.3 Chapter 4, Mars, Part II

The landing of Mars 2020 (the Perseverance Rover and Ingenuity Helicopter) at Jezero Crater in February 2021 offered the intriguing chance to use InSight to try and detect a seismic event of known magnitude, location, and timing. Such events are naturally very rare on Mars - spacecraft do not land all that often, and the source properties of marsquakes cannot be independently constrained without using seismology. If either the acoustic, seismoacoustic, or seismic signals from Mars 2020's entry, descent, and landing (EDL) were detected, they would have enabled us to calibrate all our other seismic measurements - as we would have a signal with known travel time and propagation path. If not detected, this would have enabled us to put the first experimental constraints on the seismic efficiency of impacts on Mars.

Either way, this experiment represented the first time ever that a spacecraft on the surface of another planet attempted to detect one of its colleagues' arrivals, and the first attempt to do this on a world other than the Earth since 1972.

Our pre-landing modelling predictions suggested that the acoustic and seismoacoustic signals were unlikely to be detectable above the noise floor, but in the best-case scenario (lowest environmental noise), there was a 40% chance that the seismic signal produced by the surface impact of the spacecraft's Cruise Mass Balance Devices (CMBDs) would be recorded by InSight. These estimates were used to inform the configuration of InSight's instruments for the landing period, data from which was returned in late February 2021.

8.2.4 Chapter 5, Mars, Part III

In this chapter, we explored the results from InSight's attempt to detect the seismic waves induced by the impact of Perseverance's CMBDs. Although no detection was recorded, the relatively low noise floor during the expected arrival window allows us to place two important constraints on the physics of impact processes.

Firstly, we were able to exclude a distance-magnitude scaling relationship which uses the total impactor momentum as the determinant of peak P-wave amplitude in the far field if we assume an earthlike target medium. Secondly, we were able to place a joint constraint on the impacts' seismic efficiency (the first such in-situ estimate ever made on Mars) and the mean mantle attenuation along the source-receiver path. We constrain the seismic efficiency to be less than 3%. This value is compatible with those derived from post-landing imaging of CMBD impact sites, and represents a novel method for estimating these quantities, which may be replicated in other planetary geophysics contexts.

8.2.5 Chapter 6: The Sun

The aim of this project was to extend the AxiSEM3D code which was developed for the Earth to simulate acoustic wave propagation in the Sun. However, this project proved to be more complex in its implementation than had been anticipated.

Firstly, the enormous gradients in sound speed between the solar surface and core proved a real challenge to accommodate using a seismic mesh, though it was not initially clear why this was the case. Solving this problem involved introducing ersatz discontinuity layers, removing some of the atmosphere from the model, and increasing the polynomial order of the solver, which in turn increased the computational cost of such simulations.

Beyond that, the choice between a reflecting and an absorbing boundary condition proved to be an unrepresentative dichotomy, as on the Sun the boundary condition's nature depends on the frequency range in question. The addition of self-gravity into our models was a topic of much discussion, and we appeared to have identified a potentially much cheaper way of doing this than the traditional approach, using a hyperbolic rather than elliptic formulation for the gravitational potential. However, it transpired that including a convectively unstable profile and the effects of gravitation without a mechanism to damp exponentially growing motion would have rendered our simulations unstable. In frequency domain codes this is relatively straightforward to deal with, but in AxiSEM3D there was no obvious route around this problem. Regardless, this is a potential addition that could stably be made to AxiSEM3D for rocky worlds.

Nonetheless, we were able to show that the AxiSEM3D formulation can be expanded to undertake simulations on a fully-fluid, high-gradient objects like the Sun. We were even able to generate seismograms which appeared to show sensible, causal behaviour. If someone were to move forward with this project, they might like to work on the inclusion of three-dimensional heterogeneities in fluid regions, and the extension of sensitivity kernels to quantify the effects of these on seismic wave propagation.

8.2.6 Chapter 7: Other stars

In Chapter 7, we took our first steps beyond the Solar System, and considered whether asteroseismology might be used to differentiate between two possible internal configurations for the λ Boötis class of stars. These stars display a chemical

peculiarity in their outer layers, and although they are canonically taken to be Main Sequence (MS) or near-MS objects, we also explored a potential Horizontal Branch (HB) configuration. Our interest was motivated by the λ Boötis object Tycho B, which has been proposed as a potential donor companion to the white dwarf which exploded in SN 1572. We hypothesised that asteroseismology might provide one route to differentiating between these two configurations, which has otherwise proved challenging.

We saw that the 1) computed mode frequencies, 2) their separations and evolution therein, and 3) the number of avoided crossings in the spectra of the MS and HB stars were very different - though these models will need some future refinement, for example to consider the effects of rapid rotation. However, we were unable to suggest a plausible excitation mechanism which would result in one or both of these stars showing observable pulsations, as they were too hot to be excited through the traditional opacity-valve type mechanisms for stars of their spectral class. However, all is not lost - pulsations have been observed in comparable stars of similar temperatures, though the mechanism(s) behind this remain entirely unexplained. If either one or both of these stars displays such an unexpected pulsation mechanism, then it might prove a route through which asteroseismic analysis could be used to differentiate between them. Sadly Tycho B is too faint and in too crowded a starfield to be studied by TESS, but the potential to search for pulsations using ground-based telescopes remains exciting.

8.2.7 A few last notes...

That's the end of the summary, but one more thing, before you go...

I'm told that once (if?) accepted, this thesis will be stored in the archive of the Bodleian Libraries at Oxford for all time, or something like that. As such, and given that it's meant to be a record of all the work I've undertaken over the last few years, I would like to mention a few other projects which I've worked on which we've not discussed so far. Most of these are not seismology related - and

I certainly don't expect the reader to go and look them up - but I would hope that you don't mind my adding them here.

8.2.7.1 Seismology-related papers

I've already briefly mentioned Nishiyama et al. (2021), but beg indulgence for a quick summary. In this experiment, a small explosive device was set off above the rubble-pile asteroid Ryugu by the Japanese Space Agency's Hayabusa2 spacecraft in 2019. By photographing the surface before and after the explosion, the hope was to constrain the seismic efficiency and propagation dynamics within the asteroid by evaluating where boulders had moved - which in the asteroid's weak gravity would be indicative of the seismic acceleration having exceeded the local gravitational field strength. As it happened, no seismic motions were observed anywhere outside of the crater rim. By conducting high-resolution ($\sim$kHz) simulations in AxiSEM3D, we were able to determine that this could only occur if the asteroid's bulk were substantially weaker than was previously expected, or if it was substantially more porous. Crucially, high attenuation was not found to be a plausible mechanism for preventing the propagation of seismic waves through the asteroid's interior to the required degree. My work involved helping to set up and run seismic simulations of wave propagation in the asteroid's interior.

In Daubar et al. (2020), the InSight impacts team reported on the formation of a new crater, around 1.5 m in diameter, around 40 km from InSight. We were able to constrain the crater to having formed within a period of 8 weeks from orbital images, however we were not able to unambiguously match it to any waveform observed in the seismic data. Indeed, modelling suggested that the highest-amplitude seismic waves induced by the impact were unlikely to have been above the seismometer noise floor more than 20% of the time. Because the timeframe within which we could constrain this impact to have occurred spanned weeks rather than seconds (as was the case in Chapter 5), the 'upper bound' on the peak P-wave amplitude was the highest noise level on the noisiest sol during this window. This is so

high an amplitude that constraints on seismic properties derived from this non-detection are not particularly enlightening. My role here also involved simulating the propagation of seismic waves from the crater.

8.2.7.2 Equality, diversity, and inclusion-related papers

These are undoubtedly the papers from my PhD of which I am most proud.

In Fernando et al. (2020b), a letter published in Nature, we detailed how COVID-19 related research funding received from the European Commission was vital to both UK-based and international efforts to combat the pandemic; and made the case for continued access to European funding programs for British and British-based researchers post-Brexit. This campaign was ultimately successful.

On a much less positive note, in Dowey et al. (2021) (published in Nature Geoscience), I helped undertake one of the first large-scale analyses of the ethnic makeup of the UK geoscience student population. We found that at every level in higher education, students from ethnic minorities are significantly under-represented as compared to their white peers. Furthermore, progress on this issue has been frustratingly slow, and in some places virtually non-existent. However, we were very proud to also present a detailed list suggestions to address some of the challenges raised.

Last but not least, those wishing to learn more about efforts along these lines I've helped to lead within my own department might wish to take a look at Fernando (2021) or Colquhoun and Fernando (2020), which focus on addressing the poor recruitment and retention within earth sciences specifically of those from minority ethnic backgrounds.

8.3 Concluding conclusions/closing remarks

All that remains to be done is to thank you for making it to the end! I hope you've found it interesting. More importantly, I hope that the not-so-witty attempts at humour managed to put a smile on your face at least once.

There is doubtless much more to be done on these topics, and many improvements to be made. For any shortcomings you've identified in this thesis, of which I am sure there are many, I can but leave you with a quote from *Julius Caeser*, which the man himself almost certainly never said or heard:

'Men at some time are masters of their fates: the fault, dear Brutus, is not in our stars, but in ourselves'.

Exit, PhD student, stage left.

Works Cited

Abràmoff, M. D., P. J. Magalhães, and S. J. Ram (2004). "Image processing with imageJ". In: *Biophotonics International* 11.7, pp. 36–42.

Aerts, C., J. Christensen-Dalsgaard, and D. W. Kurtz (2010). *Asteroseismology*. Springer Science & Business Media.

Afanasiev, M. et al. (2019). "Modular and flexible spectral-element waveform modelling in two and three dimensions". In: *Geophysical Journal International* 216.3, pp. 1675–1692.

Al-Attar, D. and O. Crawford (2016). "Particle relabelling transformations in elastodynamics". In: *Geophysical Journal International* 205.1, pp. 575–593.

Al-Attar, D. and J. H. Woodhouse (2008). "Calculation of seismic displacement fields in self-gravitating earth models-applications of minors vectors and symplectic structure". In: *Geophysical Journal International* 175.3, pp. 1176–1208.

Amante, C. and B.W. Eakins (2009). "ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis." In: *NOAA Technical Memorandum* NGDC-24.

Amsden, A.A., H.M. Ruppel, and C.W. Hirt (1980). *SALE: a simplified ALE computer program for fluid flow at all speeds*. Tech. rep. Los Alamos, NM (United States): Los Alamos National Laboratory (LANL).

Anderson, D. L. et al. (1977). "Seismology on Mars". In: *Journal of Geophysical Research (1896-1977)* 82.28, pp. 4524–4546.

Andrievsky, S. M. (1997). "On the possible origin of λ Boötis stars." In: *Astronomy & Astrophysics* 321, pp. 838–840.

Appourchaux, T. et al. (2010). "The quest for the solar g modes". In: *The Astronomy and Astrophysics Review* 18.1, pp. 197–277.

Ardhuin, F., L. Gualtieri, and E. Stutzmann (2015). "How ocean waves rock the Earth: Two mechanisms explain microseisms with periods 3 to 300s". In: *Geophysical Research Letters* 42.3, pp. 765–772.

Arkani-Hamed, J. and P. Olson (2010). "Giant impacts, core stratification, and failure of the Martian dynamo". In: *Journal of Geophysical Research: Planets* 115.E7.

Away, Season 1, Episode 8 (2020).

Backus, G. and M. Mulcahy (1976a). "Moment tensors and other phenomenological descriptions of seismic sources – I. Continuous displacements". In: *Geophysical Journal International* 46.2, pp. 341–361.

— (1976b). "Moment tensors and other phenomenological descriptions of seismic sources – II. Discontinuous displacements". In: *Geophysical Journal International* 47.2, pp. 301–329.

Bagheri, A. et al. (2019). "Tidal Response of Mars Constrained From Laboratory-Based Viscoelastic Dissipation Models and Geophysical Data". In: *Journal of Geophysical Research: Planets* 124.11, pp. 2703–2727.

Bahcall, J.N. and R.K. Ulrich (1988). "Solar models, neutrino experiments, and helioseismology". In: *Reviews of Modern Physics* 60.2, p. 297.

Balona, L.A. et al. (2016). "The hot γ Doradus and Maia stars". In: *Monthly Notices of the Royal Astronomical Society* 460.2, pp. 1318–1327.

Banerdt, W. B. and C. T. Russell (2017). "Editorial on: Topical Collection on InSight Mission to Mars". In: *Space Science Reviews* 211.1, pp. 1–3.

Banerdt, W. B. et al. (2020). "Initial results from the InSight mission on Mars". In: *Nature Geoscience* 13.

Banfield, D. et al. (2019). "InSight Auxiliary Payload Sensor Suite (APSS)". In: *Space Science Reviews* 215.

Banfield, D et al. (2020). "The atmosphere of Mars as observed by InSight". In: *Nature Geoscience* 13.3, pp. 190–198.

Barucq, H. et al. (2018). "Atmospheric radiation boundary conditions for the Helmholtz equation". In: *European Series in Applied and Industrial Mathematics: Mathematical Modelling and Numerical Analysis* 52.3, pp. 945–964.

Baschek, B. and A. Slettebak (1988). "The chemical composition of the Lambda Boötis stars". In: *Astronomy & Astrophysics* 207.

Basini, P. et al. (2013). "The influence of nonuniform ambient noise on crustal tomography in Europe". In: *Geochemistry, Geophysics, Geosystems* 14.5, pp. 1471–1492.

Bass, H. and J. Chambers (2001). "Absorption of sound in the Martian atmosphere". In: *The Journal of the Acoustical Society of America* 109.6, pp. 3069–3071.

Ben-Menahem, A. (1995). "A concise history of mainstream seismology: Origins, legacy, and perspectives". In: *Bulletin of the Seismological Society of America* 85.4, pp. 1202–1225.

Berenger, J.P. (1994). "A perfectly matched layer for the absorption of electromagnetic waves". In: *Journal of Computational Physics* 114.2, pp. 185–200.

Bertka, C. M. and Y. Fei (1998). "Density profile of an SNC model Martian interior and the moment-of-inertia factor of Mars". In: *Earth and Planetary Science Letters* 157.1-2, pp. 79–88.

Beyreuther, M. et al. (2010). "ObsPy: A python toolbox for seismology". In: *Seismological Research Letters* 81.3, pp. 530–533.

Bierhaus, E et al. (2013). "A fortuitous impact experiment at Mars". In: *Lunar and Planetary Science Conference, 2013*.

Blackman, D. K., J. A. Orcutt, and D. W. Forsyth (1995). "Recording Teleseismic Earthquakes Using Ocean-Bottom Seismographs at Mid-Ocean Ridges". In: *Bulletin of the Seismological Society of America* 85.6, pp. 1648–1664.

Bohlender, D. A., J.-F. Gonzalez, and J. M. Matthews (1999). "The incidence of nonradial pulsation in the lambda Boötis stars". In: *Astronomy & Astrophysics* 350, pp. 553–565.

Bohnenstiehl, D. R. et al. (2013). "Underwater acoustic records from the March 2009 eruption of Hunga Ha'apai-Hunga Tonga volcano in the Kingdom of Tonga". In: *Journal of Volcanology and Geothermal Research* 216.3, pp. 12–24.

Böning, V. G. A. et al. (2016). "Sensitivity kernels for flows in time–distance helioseismology: Extension to spherical geometry". In: *The Astrophysical Journal* 824.1, p. 49.

Borovička, J. et al. (2013). "The trajectory, structure and origin of the Chelyabinsk asteroidal impactor". In: *Nature* 503.7475, pp. 235–237.

Bottero, A. et al. (2016). "An axisymmetric time-domain spectral-element method for full-wave simulations: Application to ocean acoustics". In: *The Journal of the Acoustical Society of America* 140.5, pp. 3520–3530.

Bowman, D. M. et al. (2019). "Low-frequency gravity waves in blue supergiants revealed by high-precision space photometry". In: *Nature Astronomy* 3.8, pp. 760–765.

Breger, M. et al. (2006). "Pulsation of the Lambda Boötis star HD 210111". In: *Astronomy & Astrophysics* 455.2.

Brinkman, N. et al. (2021). "First focal mechanisms of marsquakes". In: *Journal of Geophysical Research: Planets*, e2020JE006546.

Campante, T.L. et al. (2016). "The asteroseismic potential of TESS: exoplanet-host stars". In: *The Astrophysical Journal* 830.2, p. 138.

Casas, R. et al. (2009). "Asteroseismological modelling of the multiperiodic λ Boötis star 28 Cygni". In: *The Astrophysical Journal* 697.

Ceylan, S. et al. (2017). "From Initial Models of Seismicity, Structure and Noise to Synthetic Seismograms for Mars". In: *Space Science Reviews*, pp. 1–16.

Ceylan, S. et al. (2021). "Companion guide to the marsquake catalog from InSight, Sols 0–478: Data content and non-seismic events". In: *Physics of the Earth and Planetary Interiors* 310, p. 106597.

Chapp, E., D. R. Bohnenstiehl, and M. Tolstoy (2005). "Sound-channel observations of ice-generated tremor in the Indian Ocean". In: *Geochemistry, Geophysics, Geosystems* 6.6.

Charalambous, C. et al. (2021). "A comodulation analysis of atmospheric energy injection into the ground motion at InSight, Mars". In: *Journal of Geophysical Research: Planets*, e2020JE006538.

Christensen-Dalsgaard, J. et al. (1996). "The current state of solar modeling". In: *Science* 272.5266, pp. 1286–1292.

Clayton, R. and B. Engquist (1977). "Absorbing boundary conditions for acoustic and elastic wave equations". In: *Bulletin of the Seismological Society of America* 67.6, pp. 1529–1540.

Clinton, J. F. et al. (2017). "Preparing for InSight: An Invitation to Participate in a Blind Test for Martian Seismicity". In: *Seismological Research Letters* 88.5, pp. 1290–1302.

Clinton, J. F. et al. (2021). "The Marsquake catalogue from InSight, sols 0–478". In: *Physics of the Earth and Planetary Interiors* 310.

Collins, G. S., H. J. Melosh, and B. A. Ivanov (2004). "Modeling damage and deformation in impact simulations". In: *Meteoritics & Planetary Science* 39.2, pp. 217–231.

Collins, G. S., H.J. Melosh, and K. Wünnemann (2011). "Improvements to the $\epsilon - \alpha$ porous compaction model for simulating impacts into high-porosity solar system objects". In: *International Journal of Impact Engineering* 38.6, pp. 434–439.

Colquhoun, R. and B. Fernando (2020). "An audit for action". In: *Astronomy and Geophysics* 61.5, pp. 5–40.

Combs, L. M. et al. (2019). "Petrology of the enriched poikilitic shergottite Northwest Africa 10169: Insight into the martian interior". In: *Geochimica et Cosmochimica Acta* 266, pp. 435–462.

Cowling, T. G. (1941). "The non-radial oscillations of polytropic stars". In: *Monthly Notices of the Royal Astronomical Society* 101, p. 367.

Cristini, P. and D. Komatitsch (2012). "Some illustrative examples of the use of a spectral-element method in ocean acoustics". In: *The Journal of the Acoustical Society of America* 131.3, pp. 229–235.

Crotwell, H. P., T. J. Owens, and J. Ritsema (1999). "The TauP Toolkit: Flexible Seismic Travel-time and Ray-path Utilities". In: *Seismological Research Letters* 70.2, pp. 154–160.

Daubar, I. J. et al. (2020). "A New Crater Near InSight: Implications for Seismic Impact Detectability on Mars". In: *Journal of Geophysical Research: Planets* 125.8.

Daubar, I.J. et al. (2013). "The current martian cratering rate". In: *Icarus* 225.1, pp. 506–516.

Daubar, I.J. et al. (2018). "Impact-Seismic Investigations of the InSight Mission". In: *Space Science Reviews* 214.

Davy, C. et al. (2014). "Tracking major storms from microseismic and hydroacoustic observations on the seafloor". In: *Geophysical Research Letters* 41.8825-8831.

De Groot-Hedlin, C. D. and J. A. Orcutt (2001). "Excitation of T -phases by seafloor scattering". In: *The Journal of the Acoustical Society of America* 109.5, pp. 1944–1954.

De Groot-Hedlin, C. D. et al. (2008). "Evaluation of infrasound signals from the shuttle Atlantis using a large seismic network". In: *The Journal of the Acoustical Society of America* 124.3, p. 1442.

de Groot-Hedlin, C. D. and M.A.H. Hedlin (2014). "Infrasound detection of the Chelyabinsk meteor at the USArray". In: *Earth and Planetary Science Letters* 402. Special issue on USArray science, pp. 337–345.

Dessa, J. X., J. Virieux, and S. Lambotte (2005). "Infrasound modeling in a spherical heterogeneous atmosphere". In: *Geophysical Research Letters* 32.12, pp. 1–5.

Devillepoix, H et al. (2020). "A global fireball observatory". In: *Planetary and Space Science* 191, p. 105036.

Dowey, N. et al. (2021). "A UK perspective on tackling the geoscience racial diversity crisis in the Global North". In: *Nature Geoscience*, pp. 1–4.

Dréo, R. et al. (2019). "Baleen whale distribution and seasonal occurrence revealed by an ocean bottom seismometer network in the Western Indian Ocean". In: *Deep-Sea Research Part II: Topical Studies in Oceanography* 161, pp. 132–144.

Dushaw, B. D. et al. (1999). "Multimegameter-range acoustic data obtained by bottom-mounted hydrophone arrays for measurement of ocean temperature". In: *IEEE Journal of Oceanic Engineering* 24.2, pp. 202–214.

Duvall, T. L. et al. (1993). "Time–distance helioseismology". In: *Nature* 362.6419, pp. 430–432.

Duvall, T.L. and L. Gizon (2000). "Time-distance helioseismology with f modes as a method for measurement of near-surface flows". In: *Helioseismic Diagnostics of Solar Convection and Activity*. Springer, pp. 177–191.

Dziak, R. P. et al. (2004). "P- and T-wave detection thresholds, Pn velocity estimate, and detection of lower mantle and core P-waves on ocean sound-channel hydrophones at the Mid-Atlantic Ridge". In: *Bulletin of the Seismological Society of America* 94.2, pp. 665–677.

Eddington, A. S. (1918). "The problem of the Cepheid variables". In: *Monthly Notices of the Royal Astronomical Society* 79, p. 2.

Edwards, W. N. (2009). "Meteor Generated Infrasound: Theory and Observation". In: *Infrasound Monitoring for Atmospheric Studies*. Ed. by A. Le Pichon, E. Blanc, and A. Hauchecorne. Springer Netherlands, pp. 361 –414.

Ekström, G., M. Nettles, and G. A. Abers (2003). "Glacial earthquakes". In: *Science* 302.5645, pp. 622–624.

Endl, M. (2018). *Planets Around Early-Type Stars Ii: A TESS Search For Transiting Planets Around Lambda Boötis Stars*.

Engdahl, E., R. D. Van Hilst, and R. Buland (1998). "Global teleseismic earthquake relocation with improved travel times and procedures for depth determination". In: *Bulletin of the Seismological Society of America* 88.3, pp. 722–743.

Ens, T.A. et al. (2012). "Infrasound production by bolides: A global statistical study". In: *Journal of Atmospheric and Solar-Terrestrial Physics* 80, pp. 208–229.

Evans, J.W. and R. Michard (1962). "Observational Study of Macroscopic Inhomogeneities in the Solar Atmosphere. III. Vertical Oscillatory Motions in the Solar Photosphere." In: *The Astrophysical Journal* 136, p. 493.

Fang, X. et al. (2010). "On the effect of the martian crustal magnetic field on atmospheric erosion". In: *Icarus* 206.1. Solar Wind Interactions with Mars, pp. 130–138.

Fay, T. H. and P. G. Webster (1996). "Lagrange interpolation and Runge's example". In: *International Journal of Mathematical Education in Science and Technology* 27.6, pp. 785–795.

Fernando, B. (2021). "Student-led diversity audits: A strategy for change". In: *EoS* 102.

Fernando, B., M. Tsekhmistrenko, and K. Hosseini (2018). "Training martian seismologists for InSight". In: *Astronomy and Geophysics* 59.5, pp. 5.17–5.21.

Fernando, B., K. Leng, and T. Nissen-Meyer (2020a). "Oceanic high-frequency global seismic wave propagation with realistic bathymetry". In: *Geophysical Journal International* 222.2, pp. 1178–1194.

Fernando, B. et al. (2021a). "Listening for the Landing: Detecting Perseverance's arrival at Mars with InSight". In: *Earth and Space Science* 8.4.

— (2021b). "Seismic constraints from a Mars impact experiment using InSight and Perseverance". In: *Nature Astronomy*.

Fernando, Benjamin et al. (2020b). "COVID-19 shows UK-EU collaborations are irreplaceable". In: *Nature* 586.7828.

Folkner, W.M. et al. (1997). "Interior structure and seasonal mass redistribution of Mars from radio tracking of Mars Pathfinder". In: *Science* 278.5344, pp. 1749–1752.

Fournier, D. et al. (2017). "Atmospheric-radiation boundary conditions for high-frequency waves in time-distance helioseismology". In: *Astronomy & Astrophysics* 608, A109.

Fournier, D. et al. (2018). "Sensitivity kernels for time-distance helioseismology-Efficient computation for spherically symmetric solar models". In: *Astronomy & Astrophysics* 616, A156.

French, S. W. and B. A. Romanowicz (2014). "Whole-mantle radially anisotropic shear velocity structure from spectral-element waveform tomography". In: *Geophysical Journal International* 199.3, pp. 1303–1327.

Froment, M. et al. (2020). "Lagrangian-based simulations of hypervelocity impact experiments on Mars regolith proxy". In: *Geophysical Research Letters* 47.13.

Fuchs, K. (2009). "The great earthquakes of Lisbon 1755 and Aceh 2004 shook the world. Seismologists' societal responsibility". In: *The 1755 Lisbon Earthquake: Revisited*. Springer, pp. 43–64.

Garcia, R. F. et al. (2017). "Finite-Difference Modeling of Acoustic and Gravity Wave Propagation in Mars Atmosphere: Application to Infrasounds Emitted by Meteor Impacts". In: *Space Science Reviews*.

Gaulme, P. et al. (2011). "Detection of Jovian seismic waves: a new probe of its interior structure". In: *Astronomy & Astrophysics* 531, A104.

Giardini, D. et al. (2020). "The seismicity of Mars". In: *Nature Geoscience* 13, pp. 205–212.

Gilliland, R. L. et al. (2010). "Kepler asteroseismology program: introduction and first results". In: *Publications of the Astronomical Society of the Pacific* 122.888, p. 131.

Gizon, L. et al. (2009). "Helioseismology of sunspots: a case study of NOAA region 9787". In: *Space Science Reviews* 144.1-4, pp. 249–273.

Gizon, L. et al. (2017). "Computational helioseismology in the frequency domain: acoustic waves in axisymmetric solar models with flows". In: *Astronomy & Astrophysics* 600, A35.

Golabek, G. J et al. (2011). "Origin of the martian dichotomy and Tharsis from a giant impact causing massive magmatism". In: *Icarus* 215.1, pp. 346–357.

Goodricke, J. (1786). "II. A series of observations on, and a discovery of, the period of the variation of the light of the star marked δ by Bayer, near the head of Cepheus. In a letter from John Goodricke, Esq. to Nevil Maskelyne". In: *Philosophical Transactions of the Royal Society of London* 76, pp. 48–61.

Grant, J. A. et al. (2018). "The science process for selecting the landing site for the 2020 Mars rover". In: *Planetary and Space Science* 164, pp. 106 –126.

Gray, R. O. and C. J. Corbally (2002). "A Spectroscopic Search for λ Boötis and Other Peculiar A-Type Stars in Intermediate-Age Open Clusters". In: *The Astrophysical Journal* 124.4.

Gray, R. O. et al. (2017). "The Discovery of λ Boötis Stars: The Southern Survey I". In: *The Astronomical Journal* 154.1.

Gray, R.O. and C.J. Corbally (1993). "A Search for λ Boötis Stars in OB Associations". In: *Colloquia of the International Astronomical Union* 138.

Green, D. N. and A. Nippress (2019). "Infrasound signal duration: the effects of propagation distance and waveguide structure". In: *Geophysical Journal International* 216.3, pp. 1974–1988.

Grob, M., A. Maggi, and E. Stutzmann (2011). "Observations of the seasonality of the Antarctic microseismic signal, and its association to sea ice variability". In: *Geophysical Research Letters* 38.11.

Grotzinger, J. et al. (2011). "Mars Sedimentary Geology: Key Concepts and Outstanding Questions". In: *Astrobiology* 11.1, pp. 77–87.

Grundahl, F. et al. (2008). "Stellar Oscillations Network Group–SONG". In: *Journal of Physics: Conference Series*. Vol. 118. 1. IOP Publishing, p. 012041.

Gudkova, T. et al. (2015). "Impact cutoff frequency–momentum scaling law inverted from Apollo seismic data". In: *Earth and Planetary Science Letters* 427, pp. 57–65.

Guo, J.J. (2018). "Distinguishing Core and Shell Helium-burning Subdwarf B Stars by Asteroseismology". In: *The Astrophysical Journal* 866.1, p. 58.

Han, Z. and Ph Podsiadlowski (2004). "The single-degenerate channel for the progenitors of Type Ia supernovae". In: *Monthly Notices of the Royal Astronomical Society* 350.4, pp. 1301–1309.

Hanasoge, S., L. Gizon, and K. R. Sreenivasan (2016). "Seismic sounding of convection in the Sun". In: *Annual Review of Fluid Mechanics* 48, pp. 191–217.

Handler, G. et al. (2020). "Tidally trapped pulsations in a close binary star system discovered by TESS". In: *Nature Astronomy* 4.7, pp. 684–689.

Hilton, D. A. and H. R. Henderson (1974). "Measurements of sonic boom overpressures from Apollo space vehicles". In: *The Journal of the Acoustical Society of America* 56.2, pp. 323–328.

Hirai, R. et al. (2016). "Hyperbolic self-gravity solver for large scale hydrodynamical simulations". In: *Physical Review D* 93.8, p. 083006.

Igel, H. (2016). *Computational Seismology*. Oxford University Press.

Ihara, Y. et al. (2007). "Searching for a companion star of Tycho's type Ia supernova with optical spectroscopic observations". In: *Proceedings of the Astronomical Society of Japan* 59.4, pp. 811–826.

Ishihara, Y. et al. (2012). "Infrasound/seismic observation of the Hayabusa reentry: Observations and preliminary results". In: *Earth, Planets and Space* 64.7, pp. 655–660.

Jiménez, A. (2006). "An estimation of the acoustic cutoff frequency of the Sun based on the properties of the low-degree pseudomodes". In: *The Astrophysical Journal* 646.2, p. 1398.

Jiménez, A., R.A. García, and P.L. Pallé (2011). "The acoustic cutoff frequency of the Sun and the solar magnetic activity cycle". In: *The Astrophysical Journal* 743.2, p. 99.

Johnson, G. R. and W. H. Cook (1983). "A constitutive model and data from metals subjected to large strains, high strain rates and high temperatures". In: *Proc. 7th Int. Symp. on Ballistics, The Hague, Netherlands*.

Jørgensen, S. et al. (2018). "Coupling 1D stellar evolution with 3D-hydrodynamical simulations on the fly–I. A new standard solar model". In: *Monthly Notices of the Royal Astronomical Society: Letters* 481.1, pp. L35–L39.

Jura, M. (2015). "Lambda Bpötis abudance patterns: accretion from orbiting sources". In: *The Astronomical Journal*.

Karakostas, F. et al. (2018). "Inversion of Meteor Rayleigh Waves on Earth and Modeling of Air Coupled Rayleigh Waves on Mars". In: *Space Science Reviews* 214.8, p. 127.

Karlgaard, C. D. et al. (2014). "Mars Science Laboratory Entry Atmospheric Data System Trajectory and Atmosphere Reconstruction". In: *Journal of Spacecraft and Rockets* 51.4, pp. 1029–1047.

Kerzendorf, W. E. et al. (2009). "Subaru high-resolution spectroscopy of star G in the Tycho supernova remnant". In: *The Astrophysical Journal* 701.2, p. 1665.

Kerzendorf, W. E. et al. (2013). "A high-resolution spectroscopic search for the remaining donor for Tycho's supernova". In: *The Astrophysical Journal* 774.2.

Khan, A. et al. (2016). "Single-station and single-event marsquake location and inversion for structure using synthetic Martian waveforms". In: *Physics of the Earth and Planetary Interiors* 258, pp. 28–42.

Khan, A. et al. (2018). "A geophysical perspective on the bulk composition of Mars". In: *Journal of Geophysical Research: Planets* 123.2, pp. 575–611.

Kjeldsen, H. (2003). "Space and ground based data for asteroseismology". In: *Asteroseismology Across the HR Diagram*. Springer, pp. 1–12.

Knight, E. E. et al. (2020). "HOSS: an implementation of the combined finite-discrete element method". In: *Computational Particle Mechanics*, pp. 1–23.

Kobayashi, N. and K. Nishida (1998). "Continuous excitation of planetary free oscillations by atmospheric disturbances". In: *Nature* 395.6700, pp. 357–360.

Kochan, H. et al. (2000). "CASSE—The ROSETTA Lander Comet Acoustic Surface Sounding Experiment—status of some aspects, the technical realisation and laboratory simulations". In: *Planetary and Space Science* 48.5, pp. 385–399.

Komatitsch, D. and J. Tromp (2002). "Spectral-element simulations of global seismic wave propagation - II. Three-dimensional models, oceans, rotation and self-gravitation". In: *Geophysical Journal International* 150.1, pp. 303–318.

Komatitsch, D., C. Barnes, and J. Tromp (2000). "Wave propagation near a fluid-solid interface: A spectral-element approach". In: *Geophysics* 65.2, pp. 623–631.

Kosloff, R. and D. Kosloff (1986). "Absorbing boundaries for wave propagation problems". In: *Journal of Computational Physics* 63.2, pp. 363–376.

Kosovichev, A. G. (1996). "Tomographic imaging of the Sun's interior". In: *The Astrophysical Journal Letters* 461.1, p. L55.

Kristeková, M., J. Kristek, and P. Moczo (2009). "Time-frequency misfit and goodness-of-fit criteria for quantitative comparison of time signals". In: *Geophysical Journal International* 178.2, pp. 813–825.

Ksanfomaliti, L.V. et al. (1982). "Microseisms at the Venera-13 and Venera-14 Landing Sites". In: *Soviet Astronomy Letters* 8, p. 241.

Laske, G. et al. (2013). "Update on CRUST1.0 - A 1-degree Global Model of Earth's Crust". In: *Geophys. Res. Abstracts*. EGU.

Latham, G, M Ewing, and G Sutton (1969). "The Apollo passive seismic experiment". In: *Science* 165.3890, pp. 241–250.

Latham, G., W. G. McDonald, and H. J. Moore (1970a). "Missile Impacts as Sources of Seismic Energy on the Moon". In: *Science* 168.3928, pp. 242–245.

Latham, G. et al. (1970b). "Seismic data from man-made impacts on the moon". In: *Science* 170.3958, pp. 620–626.

Lay, T. and A. Rhode (2019). "Evaluating the Updip Extent of Large Megathrust Ruptures Using Pcoda Levels". In: *Geophysical Research Letters* 46.10, pp. 5198–5206.

Lay, T., C. Liu, and H. Kanamori (2019). "Enhancing Tsunami Warning Using P Wave Coda". In: *Journal of Geophysical Research: Solid Earth* 124.10, pp. 10583–10609.

Le Pichon, A. et al. (2008). "Evidence for a meteoritic origin of the September 15, 2007, Carancas crater". In: *Meteoritics & Planetary Science* 43.11, pp. 1797–1809.

Lee, W., F. Wu, and C. Jacobsen (1976). "A catalog of historical earthquakes in China compiled from recent Chinese publications". In: *Bulletin of the Seismological Society of America* 66.6, pp. 2003–2016.

Lei, Z. et al. (2014). "A framework for grand scale parallelization of the combined finite discrete element method in 2D". In: *Computational Particle Mechanics* 1.3, pp. 307–319.

Leighton, R. B., R. W. Noyes, and G. W. Simon (1962). "Velocity Fields in the Solar Atmosphere. I. Preliminary Report." In: *The Astrophysical Journal* 135, p. 474.

Leng, K., T. Nissen-Meyer, and M. van Driel (2016). "Efficient global wave propagation adapted to 3-D structural complexity: A pseudospectral/spectral-element approach". In: *Geophysical Journal International* 207.3, pp. 1700–1721.

Leng, K. et al. (2019). "AxiSEM3D: broadband seismic wavefields in 3-D global Earth models with undulating discontinuities". In: *Geophysical Journal International* 217.3, pp. 2125–2146.

Li, B. and C. A. Fleming (1997). "A three dimensional multigrid model for fully nonlinear water waves". In: *Coastal Engineering* 30.3-4, pp. 235–258.

Liang, Z.-C. et al. (2018). "Solar meridional circulation from twenty-one years of SOHO/MDI and SDO/HMI observations-Helioseismic travel times and forward modeling in the ray approximation". In: *Astronomy & Astrophysics* 619, A99.

Lin, C.-H. and D.-Y. Chou (2018). "Solar-cycle variations of meridional flows in the solar convection zone using helioseismic methods". In: *The Astrophysical Journal* 860.1, p. 48.

Lognonné, P., B. Mosser, and F.A. Dahlen (1994). "Excitation of jovian seismic waves by the Shoemaker-Levy 9 cometary impact". In: *Icarus* 110.2, pp. 180–195.

Lognonné, P. et al. (2009). "Moon meteoritic seismic hum: Steady state prediction". In: *Journal of Geophysical Research* 114.E12, E12003.

Lognonné, P. et al. (2016). "Modeling of atmospheric-coupled Rayleigh waves on planets with atmosphere: From Earth observation to Mars and Venus perspectives". In: *The Journal of the Acoustical Society of America* 140.2, pp. 1447–1468.

Lognonné, P. et al. (2019). "SEIS: Insight's Seismic Experiment for Internal Structure of Mars". In: *Space Science Reviews* 215.1, p. 12.

Lognonné, P. et al. (2020). "Constraints on the shallow elastic and anelastic structure of Mars from InSight seismic data". In: *Nature Geoscience* in press.

Longuet-Higgins, M. S. (1950). "A Theory of the Origin of Microseisms". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 243.857, pp. 1–35.

Lorenz, R.D., Y. Nakamura, and J. Murphy (2016). "A bump in the night: wind statistics point to Viking 2 Sol 80 seismometer event as a real marsquake". In: *47th Lunar and Planetary Science Conference*, p. 1566.

Lundborg, N. (1968). "Strength of rock-like materials". In: *International Journal of Rock Mechanics and Mining Sciences and* 5.5, pp. 427–454.

Malin, M. et al. (2007). "Context camera investigation on board the Mars Reconnaissance Orbiter". In: *Journal of Geophysical Research: Planets* 112.

Mandal, K. et al. (2018). "Helioseismic inversion to infer the depth profile of solar meridional flow using spherical Born kernels". In: *The Astrophysical Journal* 863.1, p. 39.

Martire, L. et al. (2020). "Martian Infrasound: Numerical Modeling and Analysis of InSight's Data". In: *Journal of Geophysical Research: Planets* 125.6, pp. 1–34.

Mazoyer, C. et al. (2013). "T-wave generation and propagation: A comparison between data and spectral element modeling". In: *The Journal of the Acoustical Society of America* 134.4, pp. 3376–3385.

McDonald, M. A., J. A. Hildebrand, and S. C. Webb (1995). "Blue and fin whales observed on a seafloor array in the Northeast Pacific". In: *The Journal of the Acoustical Society of America* 98.2, pp. 712–721.

McEwen, A. S. et al. (2007). "Mars reconnaissance orbiter's high resolution imaging science experiment (HiRISE)". In: *Journal of Geophysical Research: Planets* 112.E5.

Michaud, G. and Y. Charland (1986). "Mass loss in A and F stars - The Lambda Boötis stars". In: *The Astrophysical Journal* 311, p. 326.

Millour, E. et al. (2015). "The Mars Climate Database (MCD version 5.2)". In: *European Planetary Science Congress 2015* 10.

Mitchell, B. J. (2002). "Monitoring the Comprehensive Nuclear-Test-Ban Treaty". In: *Pure and Applied Geophysics* 159.4, pp. 619–620.

Montabone, L. et al. (2015). "Eight-year climatology of dust optical depth on Mars". In: *Icarus* 251, pp. 65–95.

Mortimer, B. et al. (2018). "Classifying elephant behaviour through seismic vibrations". In: *Current Biology* 28.9, R547–R548.

Müller, G. (1973). "Seismic moment and long-period radiation of underground nuclear explosions". In: *Bulletin of the Seismological Society of America* 63.3, pp. 847–857.

Munjiza, A. (2004). *The combined finite-discrete element method*. Wiley.

Murdoch, N. et al. (2017). "Evaluating the wind-induced mechanical noise on the InSight seismometers". In: *Space Science Reviews* 211.1, pp. 429–455.

Murphy, S. J. and E. Paunzen (2017). "Gaia's view of the λ Boo star puzzle". In: *Monthly Notices of the Royal Astronomical Society* 466.1, pp. 546–555.

Nagashima, K. et al. (2014). "Interpreting the helioseismic and magnetic imager (HMI) multi-height velocity measurements". In: *Solar Physics* 289.9, pp. 3457–3481.

Nakamura, Y. and D. L. Anderson (1979). "Martian wind activity detected by a seismometer at Viking Lander 2 site". In: *Geophysical Research Letters* 6.6, pp. 499–502.

Nakamura, Y., G. V. Latham, and H. J. Dorman (1982). "Apollo Lunar Seismic Experiment—Final summary". In: *Journal of Geophysical Research: Solid Earth* 87.S01, A117–A123.

Neal, C. R. et al. (2019). "The lunar geophysical network mission". In: *AGU Fall Meeting Abstracts*. Vol. 2019, P33D–05.

Nishida, K. and R. Takagi (2016). "Teleseismic S wave microseisms". In: *Science* 353.6302, pp. 919–921.

Nishiyama, G et al. (2021). "Simulation of Seismic Wave Propagation on Asteroid Ryugu Induced by The Impact Experiment of The Hayabusa2 Mission: Limited Mass Transport by Low Yield Strength of Porous Regolith". In: *Journal of Geophysical Research: Planets* 126 (2).

Nissen-Meyer, T., A. Fournier, and F. A. Dahlen (2007). "A two-dimensional spectral-element method for computing spherical-earth seismograms - I. Moment-tensor source". In: *Geophysical Journal International* 168.3, pp. 1067–1092.

Nissen-Meyer, T. et al. (2014a). "AxiSEM: Broadband 3-D seismic wavefields in axisymmetric media". In: *Solid Earth* 5, pp. 425–445.

Nissen-Meyer, T. et al. (2014b). "AxiSEM: broadband 3-D seismic wavefields in axisymmetric media". In: *Solid Earth* 5.1, pp. 425–445.

Nunn, C. et al. (2020). "Lunar Seismology: A Data and Instrumentation Review". In: *Space Science Reviews* 216.

Okal, E. A. (2008). "The generation of T waves by earthquakes". In: *Advances in Geophysics* 49.07, pp. 1–65.

Oldroyd, D. et al. (2007). "The study of earthquakes in the hundred years following the Lisbon earthquake of 1755". In: *Earth Sciences History* 26.2, pp. 321–370.

Panning, M. P. et al. (2015). "Verifying single-station seismic approaches using Earth-based data: Preparation for data return from the InSight mission to Mars". In: *Icarus* 248.242, pp. 230–242.

Panning, M. P. et al. (2017). "Planned products of the Mars structure service for the InSight mission to Mars". In: *Space Science Reviews* 211.1, pp. 611–650.

Patton, H. J. and W. R. Walter (1993). "Regional moment:magnitude relations for earthquakes and explosions". In: *Geophysical Research Letters* 20.4, pp. 277–280.

Paunzen, E. et al. (1998a). "Pulsation in λ Boötis stars". In: *Astronomy & Astrophysics* 335.2.

Paunzen, E et al. (1998b). "Pulsation of the lambda Boötis stars HD 111786 and HD 142994". In: *Astronomy & Astrophysics* 330.

Paunzen, E. et al. (2003). "A study of λ Boötis type stars in the wavelength region beyond 7000 Å". In: *A & A* 404.2.

Paunzen, Ernst (2004). "The λ Boötis stars". In: *Proc IAU* IAUS224, pp. 443–450.

Paxton, B. et al. (2015). "Modules for Experiments in Stellar Astrophysics (MESA): Binaries, pulsations, and explosions". In: *The Astrophysical JournalS* 220.1.

Perlmutter, S. et al. (1999). "Measurements of Ω and Λ from 42 high-redshift supernovae". In: *The Astrophysical Journal* 517.2, p. 565.

Petculescu, A. and R.M. Lueptow (2007). "Atmospheric acoustics of Titan, Mars, Venus, and Earth". In: *Icarus* 186.2, pp. 413–419.

Peter, D. et al. (2011). "Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes". In: *Geophysical Journal International* 186.2, pp. 721–739.

Proxauf, B. et al. (2020). "Exploring the latitude and depth dependence of solar Rossby waves using ring-diagram analysis". In: *Astronomy & Astrophysics* 634, A44.

Putzig, N. and M. Mellon (2007). "Apparent thermal inertia and the surface heterogeneity of Mars". In: *Icarus* 191.

Qamar, A. (1995). "Space Shuttle and Meteroid–Tracking Supersonic Objects in the Atmosphere with Seismographs". In: *Seismological Research Letters* 66.5, pp. 6–12.

Qian, Y. et al. (2019). "Teleseismic Waveform Complexities Caused by Near Trench Structures and Their Impacts on Earthquake Source Study: Application to the 2015 Illapel Aftershocks (Central Chile)". In: *Journal of Geophysical Research: Solid Earth* 124.1, pp. 870–889.

Rajšić, A. et al. (2021). "Seismic Efficiency for Simple Crater Formation in the Martian Top Crust Analog". In: *Journal of Geophysical Research: Planets* 126.2, e2020JE006662.

ReVelle, D. O. (1976). "On Meteor-Generated Infrasound". In: *Journal of Geophysical Research* 81.7, pp. 1217–1230.

ReVelle, D. O. and W. N. Edwards (2007). "Stardust—An artificial, low-velocity "meteor" fall and recovery: 15 January 2006". In: *Meteoritics & Planetary Science* 42.2, pp. 271–299.

ReVelle, D. O., W. Edwards, and T. D. Sandoval (2005). "Genesis—An artificial, low velocity "meteor" fall and recovery: September 8, 2004". In: *Meteoritics & Planetary Science* 40.6, pp. 895–916.

Richardson, J. E., H. J. Melosh, and R. Greenberg (2004). "Impact-induced seismic activity on Asteroid 433 Eros: A surface modification process". In: *science* 306.5701, pp. 1526–1529.

Richardson, J.E. and S. Kedar (2013). "An experimental investigation of the seismic signal produced by hypervelocity impacts". In: *Lunar and Planetary Science Conference*. Vol. 44, p. 2863.

Riess, A. G. et al. (1998). "Observational evidence from supernovae for an accelerating universe and a cosmological constant". In: *The Astronomical Journal* 116.3, p. 1009.

Rietmann, M. et al. (2015). "Load-Balanced Local Time Stepping for Large-Scale Wave Propagation". In: *Proceedings - 2015 IEEE 29th International Parallel and Distributed Processing Symposium, 2015*.

Rivoldini, A. et al. (2011). "Geodesy constraints on the interior structure and composition of Mars". In: *Icarus* 213.2, pp. 451–472.

Ruiz-Lapuente, P. et al. (2004). "The binary progenitor of Tycho Brahe's 1572 supernova". In: *Nature* 431.7012, pp. 1069–1072.

Ruiz-Lapuente, P. et al. (2018). "Tycho's supernova: the view from Gaia". In: *The Astrophysical Journal* 135.2.

Sanloup, C., A. Jambon, and P. Gillet (1999). "A simple chondritic model of Mars". In: *Physics of the Earth and Planetary Interiors* 112.1-2, pp. 43–54.

Sasorova, E. V. et al. (2005). "Hydro-Acoustic Monitoring on the Kamchatka Shelf: A Possibility of Early Location of Oceanic Earthquake and Local Tsunami Warning". In: *Advances in Natural and Technological Hazards Research*, pp. 305–317.

Scholz, J.-R. et al. (2020). "Detection, analysis, and removal of glitches from InSight's seismic data from Mars". In: *Earth and Space Science* 7.11, e2020EA001317.

Shapiro, N. M. et al. (2005). "High-resolution surface-wave tomography from ambient seismic noise". In: *Science* 307.5715, pp. 1615–1618.

Silber, E. A., P. G. Brown, and Z. Krzeminski (2015). "Optical observations of meteors generating infrasound: Weak shock theory and validation". In: *Journal of Geophysical Research: Planets* 120.3, pp. 413–428.

Silvotti, R. et al. (2007). "A giant planet orbiting the 'extreme horizontal branch' star V 391 Pegasi". In: *Nature*.

Simons, F. J. (2010). "Slepian Functions and Their Use in Signal Estimation and Spectral Analysis". In: *Handbook of Geomathematics*. Springer.

Smrekar, S. E. et al. (2019). "Pre-mission InSights on the interior of Mars". In: *Space Science Reviews* 215.1, pp. 1–72.

Sohl, F., G. Schubert, and T. Spohn (2005). "Geophysical constraints on the composition and structure of the Martian interior". In: *Journal of Geophysical Research: Planets* 110.E12.

Solano, E. et al. (2001). "Physical parameters of lambda Boötis stars". In: *Astronomy & Astrophysics* 374.3, pp. 957–967.

Sorrells, G. G. et al. (1971). "Earth Motion Caused by Local Atmospheric-Pressure Changes". In: *Geophysical Journal of the Royal Astronomical Society* 26.1-4, 83–&.

Spiegel, E.A. and J.-P. Zahn (1992). "The solar tachocline". In: *Astronomy and Astrophysics* 265, pp. 106–114.

Spiga, A. et al. (2018). "Atmospheric Science with InSight". In: *Space Science Reviews* 214.7.

Spohn, T. et al. (2018). "The Heat Flow and Physical Properties Package (HP3) for the InSight Mission". In: *Space Science Reviews* 214.5, p. 96.

Stähler, S. C. et al. (2018). "Seismic wave propagation in icy ocean worlds". In: *Journal of Geophysical Research: Planets* 123.1, pp. 206–232.

Stevanović, J. et al. (2017). "Bolide Airbursts as a Seismic Source for the 2018 Mars InSight Mission". In: *Space Science Reviews* 211.1-4, pp. 525–545.

Stutzmann, E. et al. (2009). "Global climate imprint on seismic noise". In: *Geochemistry, Geophysics, Geosystems* 10.11.

Suárez, J. C., M. J. Goupil, and P. Morel (2006). "Effects of moderately fast shellular rotation on adiabatic oscillations". In: *Astronomy & Astrophysics* 449.2.

Szenicer, A., K. Leng, and T. Nissen-Meyer (2020). "A complexity-driven framework for waveform tomography with discrete adjoints". In: *Geophysical Journal International* 223.2, pp. 1247–1264.

Tancredi, G. et al. (2009). "A meteorite crater on Earth formed on September 15, 2007: The Carancas hypervelocity impact". In: *Meteoritics & Planetary Science Archives* 44.12, pp. 1967–1984.

Tarantola, A. (2006). "Popper, Bayes and the inverse problem". In: *Nature Physics* 2.8, pp. 492–494.

Tassoul, M. (1980). "Asymptotic approximations for stellar nonradial pulsations". In: *The Astrophysical Journal Supplement* 43.496.

Tauzin, B. et al. (2013). "Seismoacoustic coupling induced by the breakup of the 15 February 2013 Chelyabinsk meteor". In: *Geophysical Research Letters* 40.14, pp. 3522–3526.

Teanby, N. A. (2015). "Predicted detection rates of regional-scale meteorite impacts on Mars with the InSight short-period seismometer". In: *Icarus* 256, pp. 49–62.

Teanby, N. A. and J. Wookey (2011). "Seismic detection of meteorite impacts on Mars". In: *Physics of the Earth and Planetary Interiors* 186, pp. 70–80.

Tillotson, J. H. (1962). *Metallic Equations of State for Hypervelocity Impact*. General Atomic.

Tobias, S.M. (2005). "The solar tachocline: Formation, stability and its role in the solar dynamo". In: *Fluid dynamics and dynamos in astrophysics and geophysics* 1, p. 193.

Tolstoy, I. and M. Ewing (1950). "The T phase of shallow-focus earthquakes". In: *Bulletin of the Seismological Society of America* 40.1, pp. 25–51.

Townsend, R. H.D. and S. A. Teitler (2013). "GYRE: An open-Source stellar oscillation code based on a new magnus multiple shooting scheme". In: *Monthly Notices of the Royal Astronomical Society* 435.4, pp. 3406–3418.

Tsujimoto, T. et al. (1995). "Relative frequencies of Type Ia and Type II supernovae in the chemical evolution of the Galaxy, LMC and SMC". In: *Monthly Notices of the Royal Astronomical Society* 277.3, pp. 945–958.

Turcotte, S. and P. Charbonneau (1993). "Particle transport and the Lambda Boötis phenomenon. II - an accretion/diffusion model". In: *The Astrophysical Journal* 413.

Ulrich, R. K. (1970). "The five-minute oscillations on the solar surface". In: *The Astrophysical Journal* 162, p. 993.

Uytterhoeven, K. et al. (2011). "The Kepler characterization of the variability among A-and F-type stars-I. General overview". In: *Astronomy & Astrophysics* 534, A125.

Vallée, M. and V. Douet (2016). "A new database of source time functions (STFs) extracted from the SCARDEC method". In: *Physics of the Earth and Planetary Interiors* 257, pp. 149–147.

Van Driel, M. and T. Nissen-Meyer (2014a). "Optimized viscoelastic wave propagation for weakly dissipative media". In: *Geophysical Journal International* 1999.2, pp. 1078–1093.

— (2014b). "Seismic wave propagation in fully anisotropic axisymmetric media". In: *Geophysical Journal International* 199.2, pp. 880–893.

Van Driel, M. et al. (2015). "Instaseis: Instant global seismograms based on a broadband waveform database". In: *Solid Earth* 6, pp. 701–717.

Van Driel, M. et al. (2019). "Preparing for InSight: Evaluation of the blind test for martian seismicity". In: *Seismological Research Letters* 90.4.

Van Driel, M. et al. (2021). "High-Frequency Seismic Events on Mars Observed by InSight". In: *Journal of Geophysical Research: Planets* 126.2, e2020JE006670.

Vance, S. D. et al. (2018). "Vital signs: seismology of icy ocean worlds". In: *Astrobiology* 18.1, pp. 37–53.

Varnier, J., M. Le Pape, and F. Sourgen (2018). "Ballistic Wave from Projectiles and Vehicles of Simple Geometry". In: *AIAA Journal* 56.7, pp. 2725–2742.

Venn, K. A. and D. L. Lambert (1990). "The chemical composition of three Lambda Boötis stars". In: *The Astrophysical Journal* 363, p. 234.

Walker, J. D. (2003). "Loading sources for seismological investigation of asteroids and comets". In: *International Journal of Impact Engineering* 29.1-10, pp. 757–769.

Wan, W. X. et al. (2020). "China's first mission to Mars". In: *Nature Astronomy* 4.721.

Watters, T. R., P. J. McGovern, and R. P. Irwin III (2007). "Hemispheres Apart: The Crustal Dichotomy on Mars". In: *Annual Review of Earth and Planetary Sciences* 35.1, pp. 621–652.

Wielandt, E. (2009). "Seismic sensors and their calibration". In: *New Manual of Seismological Observatory Practice*. Deutsches GFZ, pp. 1–46.

Williams, J.P. (2001). "Acoustic environment of the Martian surface". In: *Journal of Geophysical Research: Planets* 106.E3, pp. 5033–5041.

Winget, D.E. and S.O. Kepler (2008). "Pulsating white dwarf stars and precision asteroseismology". In: *Annual Reviews in Astronomy and Astrophysics* 46, pp. 157–199.

Wójcicka, N. and M. Froment (2020). *Supplementary information for Fernando, B. et al, 2021.* Version v1.0.0.

Wójcicka, N. et al. (2020). "The seismic moment and seismic efficiency of small impacts on Mars". In: *Journal of Geophysical Research: Planets*.

Wu, W. et al. (2018). "An SEM-DSM three-dimensional hybrid method for modelling teleseismic waves with complicated source-side structures". In: *Geophysical Journal International* 215, pp. 133–154.

Wu, Z., T. Lay, and L. Ye (2020). "Shallow Megathrust Slip During Large Earthquakes That Have High P Coda Levels". In: *Journal of Geophysical Research: Solid Earth* 125.1.

Wünnemann, K., G. S. Collins, and H. J. Melosh (2006). "A strain-based porosity model for use in hydrocode simulations of impacts and implications for transient crater growth in porous targets". In: *Icarus* 180.2, pp. 514–527.

Yamada, R. et al. (2011). "Optimisation of seismic network design: application to a geophysical international lunar network". In: *Planetary and Space Science* 59.4, pp. 343–354.

Yoder, C.F. et al. (2003). "Fluid core size of Mars from detection of the solar tide". In: *Science* 300.5617, pp. 299–303.

Yue, H. et al. (2017). "Localized water reverberation phases and its impact on backprojection images". In: *Geophysical Research Letters* 44.19, pp. 9573–9580.

Zhao, J. and A. G. Kosovichev (2006). "Surface magnetism effects in time-distance helioseismology". In: *The Astrophysical Journal* 643.2, p. 1317.

Zhao, J., A. G. Kosovichev, and T. L. Duvall (2001). "Investigation of mass flows beneath a sunspot by time-distance helioseismology". In: *The Astrophysical Journal* 557.1, p. 384.

Zhou, Y. et al. (2016). "Accuracy of the water column approximation in numerically simulating propagation of teleseismic PP waves and Rayleigh waves". In: *Geophysical Journal International* 206.2, pp. 1315–1326.